

Optimal Information Security Investment with Penetration Testing

Rainer Böhme and Márk Félegyházi

International Computer Science Institute, Berkeley, California
{rainer.boehme|mark}@icsi.berkeley.edu

Abstract. Penetration testing, the deliberate search for potential vulnerabilities in a system by using attack techniques, is a relevant tool of information security practitioners. This paper adds penetration testing to the realm of information security investment. Penetration testing is modeled as an information gathering option to reduce uncertainty in a discrete time, finite horizon, player-versus-nature, weakest-link security game. We prove that once started, it is optimal to continue penetration testing until a secure state is reached. Further analysis using a new metric for the return on penetration testing suggests that penetration testing almost always increases the per-dollar efficiency of security investment.

1 Introduction

Information security investment decisions have recently attracted the attention of researchers from computer science, economics, management science, and related disciplines. The emerging topic of the economics of information security aims at formalizing these decisions, but there is still a gap between the formal models and experiences in practice [1]. In particular, information gathering options of defenders of computer systems differ from other scenarios. Penetration testing (short: pentesting), the focus of our paper, is an example of proactive information gathering options specific to computer systems. Penetration testing is widely used in practice, but its effects have not been reflected in the information security investment literature.

In this paper, we build a model on a simplified version of the *iterated weakest link* (IWL) model of dynamic security investment [2, 3] which emphasizes the role of uncertainty in security decision making. The original IWL model explains why a defender facing uncertainty about which threats are most likely to realize might defer security investment and learn from observed attacks where the investment is most needed. The benefits of more targeted investment may outweigh the losses suffered through non-catastrophic attacks, thereby increasing the *return on security investment* (ROSI). We extend the IWL model by an option to commission pentests as a means to reduce uncertainty. Indeed, waiting for actual attacks need not be the only way of gathering information to guide security investment. Uncertainty can also be reduced by observing pre-cursors of

attacks or near misses [4], information sharing [5, 6], or investment in information gathering. Penetration testing can be seen as information gathering prior to investing into protection against so-identified threats.

Penetration testing is also referred to as “ethical hacking” because the commissioned penetration testers investigate the target system from an attacker’s point of view, reporting weaknesses rather than exploiting them. The aim of this work is to study the added benefits and costs of penetration testing to the entire system defense. The similarity between pentesting and attacks leads to the intuition that information revealed by pentests should be modeled in exactly the same way as information revealed by attacks. Yet there exist differences on the cost side: pentests cause calculable up-front costs, whereas costs associated with successful attacks are typically more volatile, much higher, and borne ex post. For all other modeling decisions, we stay close to the original IWL model, and we refer the reader to [2] for a more detailed discussion of its features.

This paper makes the following contributions:

- it provides a first attempt to study information gathering options by pentesting in the framework of the economics of security investments;
- it contains a proof that in this model, pentesting should be done consistently once started;
- it defines a metric for *return on penetration testing* (ROPT);
- and it demonstrates that pentesting not only increases total profit for the defender, but also increases (most of the cases) the per dollar efficiency of security investments.

The remainder of this paper is organized as follows. After recalling the context of related work in Section 2, we describe in Section 3 our approach to include penetration testing as information gathering step into an established model of security investment. Section 4 presents solutions of the model. Section 5 defines the ROPT metric and demonstrates how the model can be applied in investment decision making. The final Section 6 concludes with discussion and outlook.

2 Related Work

Information security investment have been studied from the economics perspective. Gordon and Loeb [7] formulate a basic economic model. They argue that taking both the risk profiles of vulnerabilities and the cost to protect them into account, the best investment strategy for a defender is to protect the mid-range of vulnerabilities.

Intrusion detection systems (IDS) build a solid line of defense against most outside attackers, but the systems are notoriously difficult to configure. Cavusoglu et al. [8] study the value of intrusion detection systems and argue that the main benefit of IDSs is not the increased detection rate, but the deterrence of the system and the increased availability of information for forensics. Using their analytical model, they found that an IDS is only valueable if the detection

rate is high enough. The authors show that the threshold for an IDS to be valuable is determined by the attacker’s benefit. The attacker’s benefit is difficult to assess in practice [9], that makes the model difficult to apply in practice. In a subsequent paper, Ogut et al. [10] study intrusion detection policies using a decision-theoretic framework. They describe a scenario where defenders wait and gather more information about potentially malicious users instead of acting on IDS signals immediately. They propose an optimal waiting strategy as well as a myopic heuristic that relies on less parameters; hence it is easier to apply in practice.

Penetration testing is an important method to assess the vulnerability of a computer system before it is deployed. Geer and Harthorne [11] argue that penetration testing requires special skills because attacks are unknowable and hence innumerable in advance. They informally discuss the value of penetration testing and connect it to the formulation of the return on security investments (ROSI). The authors advocate the evaluation of penetration test results in the light of a risk assessment. Arkin et al. [12] provide an insight into software penetration testing practices. They mostly argue for better integration of testing during the development cycle of software systems. They agree with [11] that penetration test results should not be considered as a final checklist, but rather as a sample from the potential problems. They emphasize that decision makers often stop penetration testing after an initial round, because “having found the issues” gives them a false sense of security.

A leading survey of industry participants ([13], Fig. 20) reveals that the majority of responding firms performs penetration testing in practice. Yet, we are unaware of prior theoretical work that formalizes penetration testing as a specific tool available to the information security manager.

3 Model

Our model extends [2]. The *defender* operates a system that represents an asset of value a yielding a return r per period. The defender protects this system against a dispersed set of attackers. We do not distinguish between different attackers, rather we consider the group of attackers as a *single attacker* entity with enhanced capabilities. There exist n possible components of the system that are threatened by an attack.¹ Each threat can be prevented by investing into the protection of the specific component and we assume that a protection is always effective. The defender orders the threats according to their *expected* cost $\bar{x}_1 \leq \bar{x}_i \leq \bar{x}_n$, but the *true* costs to attack x_i is hidden from the defender. This reflects the opinion of many practitioners who remain skeptical about the quantifiability of attack probabilities but reckon it is possible to order threats by severity. There are no restrictions about the source of prior beliefs about this order. It can result from individual judgement, semi-formal aggregation of expert opinions, or formal calculations of threat prioritization using system models[14].

¹ Alternatively, the notion of ‘components’ can be substituted by ‘attack vectors’.

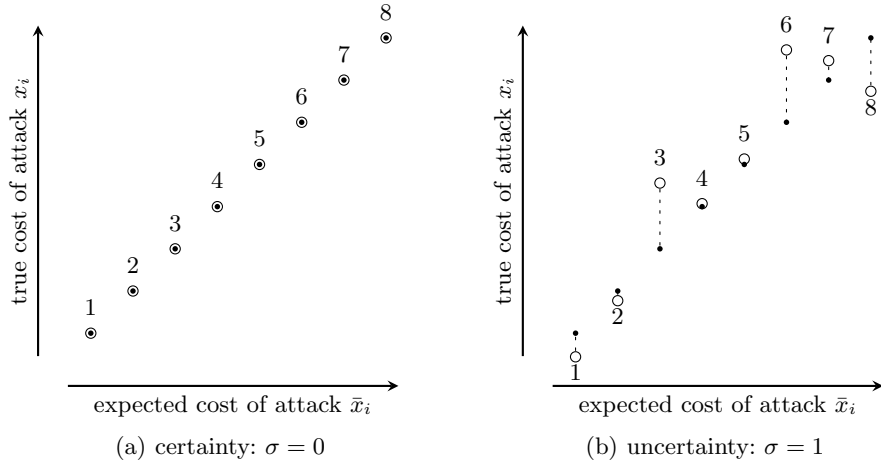


Fig. 1. The defender forms expectations about the order of attack costs for different threats (here: $i = 1, \dots, 8$) but remains ex ante uncertain about the true costs

We model the true costs to attack as

$$x_i = \sup(0, \bar{x}_i + \chi_i) \quad \text{with} \quad \chi_i \sim \mathcal{N}/(\Delta x)^2 \quad \text{and} \quad \bar{x}_i = \bar{x}_1 + (i - 1) \cdot \Delta x, \quad (1)$$

where \mathcal{N} is a mean-free Gaussian random source with standard deviation $\sigma \geq 0$. This parameter controls the degree of uncertainty and it is key to analyze the usefulness of penetration testing as uncertainty varies. Figure 1 visualizes the influence of σ in introducing noise, i.e., adding unknown offsets between actual and expected costs. Higher values of parameter σ indicate that the defender's order of threats differs more from the true order of costs to attack.

The order is relevant because in each round, the attacker exploits the *weakest link*. That is, she attacks the unprotected component with the least *true cost*² and loots a fraction of z from asset a . The attacker is opportunistic in that she attacks only if the benefits (i.e., the defender's losses) exceed the cost of attack. Hence, a secure state can be reached when all vulnerable components that can be attacked below the reservation cost of attack are protected.

We model the interaction between the defender and the attacker as a dynamic, discrete time, finite horizon, player-versus-nature game. In the initial round ($t = 0$), the defender chooses her *defense configuration* to protect against the k most possible threats $1, \dots, k$. A unit cost of 1 is incurred per protection and round. As the realizations of (x_1, \dots, x_n) are unknown, any other initial configuration would lead to inferior outcomes on average.

In the following reactive rounds ($t = 1, \dots, t_{\max}$), four steps are iterated:

² The intuition is that our attacker model represents the ensemble of individual attackers that are likely to discover the weakest link.

1. The defender chooses whether or not to commission a pentest at cost $c > 0$.
2. If a pentest has been executed, it succeeds with probability $p > 0$ and reveals the next weakest link i with true attack cost $i = \arg \min_j x_j$ over all unprotected links j . The defender protects i . This increases her defense cost by 1 in the current and all subsequent rounds.³
3. An attack occurs if at least one $x_i \leq z \cdot a$. If so, the defender learns which link i was the weakest and incurs a loss of $z \cdot a$. Otherwise the defender learns that the system has reached a secure state.
4. The defender chooses whether to upgrade the defense configuration and protect against the threat revealed in the last attack. This increases her defense cost by 1 for all subsequent rounds.

Observe that steps 3 and 4 exactly correspond to the original model in [2], steps 1 and 2 are new to introduce pentests as means of information gathering. For simplicity, we do not consider sunk costs or interdependent defenses here, i.e., $\lambda = 0$ and $\rho = 0$ in the notation of [2]. Like the original model, the defender is risk neutral.

4 Analysis

Although the model is simple, its solution is not trivial. Figure 2 depicts an excerpt of the defender's optimization problem in extensive form. Observe the pairwise alternation of moves by player and nature and the repetition of steps 1 to 4 in each round.

The defender starts at node S and chooses the initial level of defense k . Nodes annotated with T are terminal nodes:

T_0 : this singular case corresponds to knowingly indefensible situations, i.e., if $r \geq z$ and $a \cdot (r - z) < E[\frac{1}{n} \sum_i x_i]$. In this case, the defender refrains from investing in security and rather accepts the losses due to attacks in each round. The value of this node is

$$T_0 = t_{\max} \cdot a \cdot (r - z). \quad (2)$$

T_1 : this case corresponds to the arrival at a secure state. The defender's goal is to reach a node of type T_1 as soon as possible. The deterministic value of these nodes is a function of t , k , and the number of successful pentests $|\mathcal{M}^+|$,

$$T_1(t, k, |\mathcal{M}^+|) = (t_{\max} - t + 1) \cdot (r \cdot a - t - k - |\mathcal{M}^+| + 1). \quad (3)$$

T_2 : this is the case when the system is found indefensible only after revelation of realizations of nature. In an indefensible situation, the defender would always

³ As defense costs are constant for each threat, the decision to defend upon revelation is cogent. Otherwise it is always better not to commission the pentest in step 1.

prefer T_0 over any T_2 (where costs for ineffective defenses are unrecoverably sunk). The deterministic value of these nodes is a function of t ,

$$T_2(t) = (t_{\max} - t + 1) \cdot a \cdot (r - z). \quad (4)$$

Since the influence of nodes T_2 is negligible for the parameter settings used throughout this paper, we do not consider them in the analysis for brevity.

The dashed branches leading to an asterisk node are decisions not to pentest even though at least one pentest has been commissioned in an earlier round. Theorem 1 states that these paths are strictly dominated by the alternative choice and can indeed be eliminated to simplify the extensive form representation.

Theorem 1 *Once the defender starts pentesting, she will keep doing it until a secure state is reached.*

Proof. We use Lemma 1 proven in the appendix. It gives us the following expression for the total profit of a defender as a function of the initial defense k , the set of rounds in which pentests are commissioned $\mathcal{M} = \{m_1, \dots\}$, and a fixed number of unprotected components K :

$$G = \sum_{t=1}^K (a(r-z) - (k+t-1)) + \sum_{i=1}^{|\mathcal{M}|} (az - c - (K - i - m_i + 2)) + \sum_{t=K+1}^{t_{\max}} (ar - (k+K)). \quad (5)$$

The contribution of each pentest i to the total profit depends on the round m_i when the pentest is commissioned. The second sum of Eq. (5) shows that the marginal benefit of penetration testing increases with the number of rounds in the game. Thus, if a defender decides to commission pentests in round t , then she will keep doing it in each round $u > t$ until all weak links are discovered. \square

Pentests are successful with probability p . For $p < 1$, \mathcal{M} can be partitioned ex post into two disjoint subsets $\mathcal{M} = \mathcal{M}^+ \cup \mathcal{M}^-$ of rounds with successful, respectively unsuccessful penetration tests. Since the pentests are independent events, we can simply multiply their contribution with their respective probability. Hence the expected value of (5) becomes:

$$E[G] = \sum_{t=1}^K (a(r-z) - (k+t-1)) + p \sum_{i=1}^{|\mathcal{M}|} (az - c - (K - i - m_i + 2)) + \sum_{t=K+1}^{t_{\max}} (ar - (k+K)). \quad (6)$$

An intuitive way to analyze the composition of the revenue is a graphical representation, as depicted in Fig. 3. The figures show a schematic representation with infinitesimally small rounds. The costs are proportional to the shaded areas defined by the asset value a , the loss due to attacks z , the total number of weak links K , the number of proactive defenses k , the cost of a pentest c and the probability of a successful pentest p .

The first figure shows the case when no pentests are commissioned and the weak links are discovered one-by-one until all K are protected. During this time,




Fig. 3. Comparison of costs in scenario without (left) and with pentesting ($p = 1$, center) for infinitesimally many rounds; costs are proportional to areas; the success probability p defines the slope of the gradient towards reaching the secure level K (center versus right); note that K is a random variable unknown to the decision maker ex ante; pentesting is worthwhile if the expected value of the light area (savings) exceeds the expected value of the dark area (direct and indirect cost of pentesting)

the attacker loots $(K - k) \cdot a \cdot z$ profit from the asset. While protecting the asset, the defender spends the proactive protection cost $(K - k) \cdot k$ and the reactive protection cost $\frac{K-k}{2}$. In round $K - k$, all weak links are protected and the defender maintains the defense cost K for all subsequent rounds. That prevents the attacker from looting the asset.

The second figure shows⁴ that pentesting with $p = 1$ introduces two additional costs: the cost of pentests c and the cost of the resulting protection; both costs are shown as dark areas in Fig. 3. Pentesting has a benefit (the light grey areas in Fig. 3) of discovering weak links earlier than without pentests and this reduces the looting cost from the attacker. Having $p = 1$ doubles the speed of discovering weak links (the slope of protection costs is two) and halves the total looting cost. The defender chooses to perform pentests as long as the benefit due to prevented attacks is higher than the pentesting costs. The third figure shows a case when pentests are less efficient and hence their effect to reduce cost due to attacks decreases. Nonetheless, the defender has to pay the cost of pentesting for each try.

From Theorem 1, we know that the defender performs pentests from m_1 until the attacks stop. Then, we can write the expected number of pentest as:

$$E[|\mathcal{M}|] = \left\lceil \frac{K + 1 - m_1}{1 + p} \right\rceil \quad (7)$$

⁴ We show the most likely case where the pentesting starts from the first round, but the figures can easily be adapted to the case when pentesting starts at a later round.

The number of pentest depends on whether the last weakest link is protected following a pentest or an attack. Let ϵ be an indicator variable showing if the last weak link is fixed after a pentest or an attack. If the last weak link is discovered by a pentest, then $\epsilon = 1$, otherwise $\epsilon = 0$.

From (7) and using ϵ , we can derive the optimal number of pentests and the optimal time to start pentesting. The derivation is in the appendix. Substituting the optimal number of pentests into (6), we obtain an expression for the expected total profit with optimal number of pentests for a fixed K . The direct application of this expression, however, is impeded by the fact that the overall profit is largely determined by a discontinuous boundary condition and the randomness of K .

While k and m_1 are choice variables, K is a discrete random variable with known distribution but a priori unknown realization. Similar to [2], the optimal defense strategy can also be found by numerically summing up the expected total profit over the domain of K and finding the maximum of a grid search for the tuple of choice variables (k, m_1) . A result of a numerical maximization with a selected set of parameters is shown in Fig. 4, indicating the optimal strategies with and without access to pentests. For our example set of parameters, the total profit of the defender is optimal if she starts pentesting in the first round. We also observe that the initial investment in defenses k is lower in the case of pentesting as opposed to the case where no pentests are commissioned. The reason is that some resources spent on proactive protection are now reallocated to a more efficient discovery of weak links using pentests.

5 Return on Penetration Testing (ROPT)

Several definitions exist to measure the return on security investment (ROSI) [15]. We follow the approach in [2] and choose an indicator normalized by the average security investment per period [16]. Without pentesting we have,

$$\text{ROSI}_{\text{NPT}} = \frac{\text{ALE}_0 - \text{ALE}_{\text{NPT}} - \text{avg. security investment}}{\text{avg. security investment}}, \quad (8)$$

where ALE is the *annual* (i.e., per period) *loss expectation* for two cases:

- ALE_0 : a baseline case where no security investment is made,
- ALE_{NPT} : *with* security investment but *without* pentesting.

Higher values of ROSI_{NPT} denote more efficient security investment. A natural extension to penetration testing is to define:

- ALE_{PT} : loss expectation *with* security investment *and* pentesting.

However, there is no straightforward way to measure the specific return on penetration testing by plugging both ALE_{NPT} and ALE_{PT} in the numerator of Eq. (8). The reason is that the fraction of security investment related to penetration testing is difficult to identify since it consists of direct costs c and indirect costs from defenses set up earlier than without pentesting (cf. Fig. 3). Yet another source of *indirect benefits* is not visible in Fig. 3. The possibility to do




Fig. 4. Expected total return surface as a function of choice variables (k, m_1) for the following parameters: asset value $a = 1000$, return $r = 5\%$, loss given attack $z = 2.5\%$, profile of expected attack costs $(x_1, \Delta x) = (15, 1)$, uncertainty $\sigma = 4$, pentest cost $c = 0.5$, pentest success $p = 100\%$, $n = t_{\max} = 25$; $m_1 = 0$ means no pentest at all

pentests can lead to a lower optimal initial defense k (cf. Fig. 4). If these benefits match or exceed the direct and indirect costs of pentesting, then the defender can face situations where she invests equal or less and still achieves higher security than without pentesting. In this case, funds are shifted from investment in protective measures towards spending on information gathering. This seemingly odd result once again demonstrates the special role of pentesting and the need to appropriately reflect it in security investment models. For such special cases, the normalizing term based on the simple difference becomes zero or negative. This would lead to undefined values for ROSI.

To fully characterize the effects of pentesting, we propose the following metric called *return on penetration testing (ROPT)*:

$$\text{ROPT} = \text{ROSI}_{\text{PT}} - \text{ROSI}_{\text{NPT}}, \quad (9)$$

where ROSI is calculated according to Eq. (8) with ALE_{PT} and ALE_{NPT} , respectively. Consistent with the interpretation of ROSI as the dollar amount of prevented losses per dollar of security spending, ROPT can be understood as the dollar amount of *additionally prevented* losses per dollar if security investment is optimized *with* penetration testing.

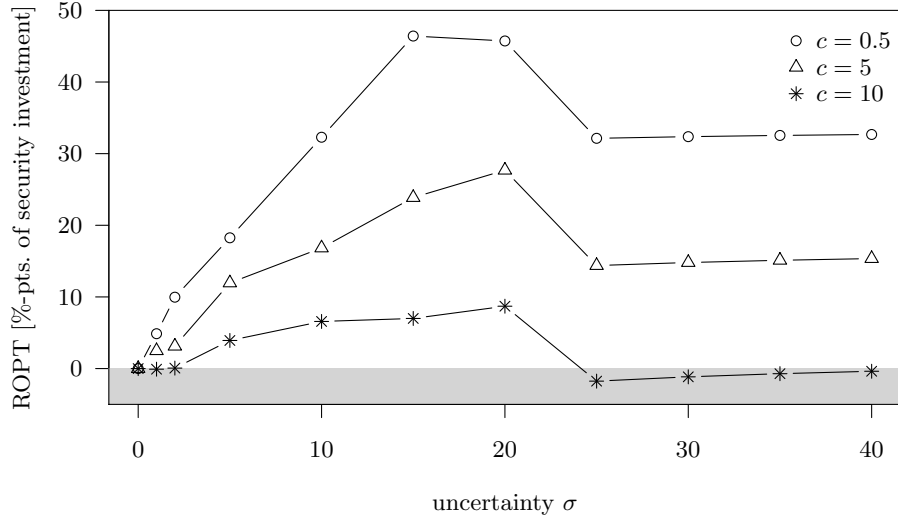


Fig. 5. Profile of return on penetration testing (ROPT) as uncertainty increases for varying cost of pentesting ($c = 0.5$ is half of the cost per protection measure and round)

Fig. 5 shows numerical values of ROPT as a function of uncertainty for the example set of parameters used in Fig. 4 and the pentest cost c . Observe for each curve that ROPT first increases with σ , then decreases until it approaches a constant value. The first increase in the ROPT values can be explained by the benefit of pentesting to gather additional information. A defender who commissions pentests invests less into proactive defenses and more into pentests and this benefit increases as uncertainty about true attack costs increases. ROPT starts to decrease when proactive defenses with pentest reach zero (e.g., $k_{PT} = 0$) meaning that a defender using pentests relies exclusively on reactive defenses, whereas a defender who does not commission pentests still invests in proactive defenses ($k_{NPT} > 0$). ROPT becomes constant when uncertainty is so high that both defense strategies avoid the proactive defense period ($k_{PT} = k_{NPT} = 0$).

The ROPT value is mostly positive (can be as high as 50%), meaning that pentesting brings a significant per dollar efficiency to security investments. However, if the pentesting cost c becomes relatively high ($c = 10$ is one order of magnitude higher than the defense cost of a weak link) then ROPT might turn negative, indicating that pentesting is a more costly security investment alternative than no pentesting. We emphasize that even in this case, the total profit of the defender increases with pentesting until the defense cost including pentesting reaches the looting cost. Thus we conclude that pentesting is a beneficial defense option for a wide range of parameters.

6 Discussion and Conclusion

In this paper, we leveraged the iterated weakest link model of [2] to propose a framework that accounts for penetration testing, an important information gathering option when making security investment decisions. To the best of our knowledge, this is the first paper that explicitly models penetration testing and shows its potentially catalyzing effect on the efficiency of security spending. We are also the first to propose ROPT, a metric to account for the efficiency of penetration testing.

Our model formalizes much of the informal discussion in other papers about security investments and pentesting. Ogut et al. [10] study intrusion detection policies and propose a model for optimal waiting time to act on intrusion signals. Our model recovers the same mentality by allowing the defender to invest less in proactive security and fix the weak links reactively after an attack occurs. Our model also formalizes the arguments of Geer and Harthorne [11] who emphasize that the results of penetration testing should be considered in the light of risk assessment rather than perceived as a security todo list. Arkin et al. [12] iterate on this view by stating that security decision makers should follow-up on the insights uncovered by pentesting. We proved that in the IWL framework, commissioning pentests is the best strategy for the defender until all weak links with feasible attack costs are protected.

We conjecture that our model captures the basic mechanisms in security investments with pentesting. Nonetheless, as any formal model, it has its shortcomings. We model the security investment process as a finite-horizon game between the defender and the nature player. A natural extension of this paper to consider the attackers as rational players. We acknowledge this future direction, but point to the fact that the profit functions of the attackers are relatively difficult to model [9]. There is some initial work to understand the profits of the attackers in real life [17], but we are lacking of deeper understanding to properly model attackers in security games. Our model assumes that the set of weak links does not change within the finite horizon of the game. There are two improvements one can consider regarding this assumption. First, the game is typically a dynamic game that can be considered as an infinite game with discounting. Second, the dynamics of changing weak links are worth exploring as well. Yet another direction involves further refinement of the model to capture even more specific details of security investment, such as the difference between black-box and white-box testing. This choice defines the distribution of information and should be modeled to affect the heuristic potential of pentesting to discover weak links similar to a real attacker.

Our paper provides a theoretical framework for penetration testing. While this exposition focused on the defender’s decision, we note that this kind of model can also be solved for the cost of penetration testing to inform providers of pentest services and guide their price setting. One major question is how this model and its conclusions fit to real data from industry sources. Obtaining such a confirmation is a potential future work. Finally, we will extend the IWL framework considering other options for uncertainty reduction beyond penetration testing.

Acknowledgements

The first author received a Postdoctoral Fellowship by the German Academic Exchange Service (DAAD) to support his visit at the International Computer Science Institute.

References

1. Su, X.: An overview of economic approaches to information security management. Technical Report TR-CTIT-06-30, University of Twente (2006)
2. Böhme, R., Moore, T.W.: The iterated weakest link: A model of adaptive security investment. In: Workshop on the Economics of Information Security (WEIS), University College London, UK (2009)
3. Böhme, R., Moore, T.W.: The iterated weakest link. *IEEE Security & Privacy* **8**(1) (2010) 53–55
4. Panjwani, S., Tan, S., Jarrin, K.M., Cukier, M.: An experimental evaluation to determine if port scans are precursors to an attack. In: Proc. of Int'l Conf. on Dependable Systems and Networks (DSN 2005), Yokkohama, Japan (2005)
5. Gordon, L.A., Loeb, M.P., Lucyszyn, W.: Sharing information on computer systems security: An economic analysis. *Journal of Accounting and Public Policy* **22**(6) (2003)
6. Gal-Or, E., Ghose, A.: The economic incentives for sharing security information. *Information Systems Research* **16**(2) (2005) 186–208
7. Gordon, L.A., Loeb, M.P.: The economics of information security investment. *ACM Transactions on Information and System Security* **5**(4) (2002) 438–457
8. Cavusoglu, H., Mishra, B., Raghunathan, S.: The value of intrusion detection systems in information technology security architecture. *Information Systems Research* **16**(1) (2005) 28–46
9. Barth, A., Rubinstein, B., Sundararajan, M., Mitchell, J., Song, D., P.L., B.: A learning-based approach to reactive security. In Radu, S., ed.: *Financial Cryptography and Data Security (Proc. of FC 2010)*. LNCS 6052, Berlin Heidelberg, Springer-Verlag (2010) 192–206
10. Ogut, H., Cavusoglu, H., Raghunathan, S.: Intrusion detection policies for it security breaches. *INFORMS Journal on Computing* **20**(1) (2008) 112–123
11. Geer, D., Harthorne, J.: Penetration testing: A duet. In: Proc. of the 18th Annual Computer Security Applications Conference (ACSAC), Las Vegas, NV, USA (2002)
12. Arkin, B., Stender, S., McGraw, G.: Software penetration testing. *IEEE Security & Privacy* **3**(1) (2005) 84–87
13. Richardson, R.: *CSI Computer Crime and Security Survey*. Computer Security Institute (2007)
14. Miura-Ko, R.A., Bambos, N.: SecureRank: A risk-based vulnerability management scheme for computing infrastructures. In: *IEEE International Conference on Communications (Proc. of ICC)*. (2007) 1455–1460
15. Böhme, R., Nowey, T.: Economic security metrics. In Eusgeld, I., Freiling, F. C., Reussner, R., eds.: *Dependability Metrics*. LNCS 4909, Berlin Heidelberg, Springer-Verlag (2008) 176–187
16. Purser, S.A.: Improving the ROI of the security management process. *Computers & Security* **23** (2004) 542–546

17. Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G., Paxson, V., Savage, S.: Spamalytics: An empirical analysis of spam marketing conversion. In: Conference on Computer and Communications Security (Proc. of ACM CCS), Alexandria, Virginia (2008) 3–14

A Appendix

A.1 Lemma: Independence of Pentests

Let us write the total profit for the defender as

$$F = \sum_{t=1}^{t_{\max}} f(k, t). \quad (10)$$

In (10), $f(k, t)$ is the profit per round and can be written as follows:

$$f(k, t) = a(r - q \cdot z) - c_t, \quad (11)$$

where q is an indicator variable that takes the value of 1 if an attack is successful and 0 otherwise; and c_t is the cost in round t .

The defender chooses the initial defense k and fixes one defense at the time as discovered by the attacker. Let us now assume that without penetration testing, the number of rounds with successful attacks is K (where $k + K$ threats are warded off). In this case, we can write the total profit for the defender F as

$$F = \sum_{t=1}^K f(k, t) + \sum_{t=K+1}^{t_{\max}} f(k, t) = \sum_{t=1}^K (a(r - z) - (k + t - 1)) + \sum_{t=K+1}^{t_{\max}} (ar - (k + K)). \quad (12)$$

Let G be the total profit for the defender when commissioning pentests. We can show that penetration test independently contribute to the total profit of the defender.

Lemma 1 *For every fixed proactive defense k , the contributions to the expected total profit from individual reactive defenses to attacks and individual penetration tests are independent and additive.*

Proof. Let $\mathcal{M} = \{m_1, m_2 \dots, m_{|\mathcal{M}|}\}$ be the ordered set of rounds where the defender commissions a pentest. Assuming $p = 1$ for now, we obtain

$$G = \sum_{t=1}^{K-|\mathcal{M}|} (a(r - z) - c_t) + \sum_{t=K-|\mathcal{M}|+1}^K (ar - c_t) + \sum_{t=K+1}^{t_{\max}} (ar - c_t),$$

where c_t is the cost at round t .

Let us now separate the saved losses due to pentesting in the second sum,

$$\begin{aligned}
G &= \sum_{t=1}^{K-|\mathcal{M}|} (a(r-z) - c_t) + \sum_{t=K-|\mathcal{M}|+1}^K (a(r-z) - c_t) + \sum_{t=K-|\mathcal{M}|+1}^K (az) + \sum_{t=K+1}^{t_{\max}} (ar - c_t) \\
&= \sum_{t=1}^K (a(r-z)) + \sum_{t=K-|\mathcal{M}|+1}^K (az) - \sum_{t=1}^{K-|\mathcal{M}|} (c_t) - \sum_{t=K-|\mathcal{M}|+1}^K (c_t) + \sum_{t=K+1}^{t_{\max}} (ar - c_t). \quad (13)
\end{aligned}$$

We now develop the costs c_t for each period of the game. In the period of attacks, each pentest contributes one more to the total number of protected threats. In addition, each pentest costs c to perform. After the attacks stop, all links are protected and the defense cost remains $k + K$ for the rest of the game:

$$\begin{aligned}
G &= \sum_{t=1}^K (a(r-z)) + |\mathcal{M}| \cdot (az) - |\mathcal{M}| \cdot c - \\
&\quad - \sum_{t=1}^{m_1-1} (k+t-1) - \sum_{i=1}^{|\mathcal{M}|-1} \sum_{t=m_i}^{m_{i+1}-1} (k+t+i-1) - \sum_{t=m_{|\mathcal{M}|}}^{K-|\mathcal{M}|} (k+t+|\mathcal{M}|-1) - \\
&\quad - \sum_{t=K-|\mathcal{M}|+1}^K (k+K) + \sum_{t=K+1}^{t_{\max}} (ar - (k+K)).
\end{aligned}$$

Note that if the last weak link is fixed by a pentest, then $m_{|\mathcal{M}|} = K - |\mathcal{M}| + 1$ and the 5th sum does not exist.

Now splitting the last but one sum results in

$$\begin{aligned}
G &= \sum_{t=1}^K (a(r-z)) + |\mathcal{M}| \cdot (az) - |\mathcal{M}| \cdot c - \\
&\quad - \sum_{t=1}^{m_1-1} (k+t-1) - \sum_{i=1}^{|\mathcal{M}|-1} \sum_{t=m_i}^{m_{i+1}-1} (k+t+i-1) - \sum_{t=m_{|\mathcal{M}|}}^{K-|\mathcal{M}|} (k+t+|\mathcal{M}|-1) - \\
&\quad - \sum_{t=K-|\mathcal{M}|+1}^K (k+t-1) - \sum_{t=K-|\mathcal{M}|+1}^K (K-t+1) + \sum_{t=K+1}^{t_{\max}} (ar - (k+K)).
\end{aligned}$$

This algebraic manipulation allows us to separate the contribution of attacks and pentest to the total profit,

$$\begin{aligned}
G &= \sum_{t=1}^K (a(r-z) - (k+t-1)) + |\mathcal{M}| \cdot (az - c) - \\
&\quad - \sum_{i=1}^{|\mathcal{M}|-1} \sum_{t=m_i}^{m_{i+1}-1} (i) - \sum_{t=m_{|\mathcal{M}|}}^{K-|\mathcal{M}|} (|\mathcal{M}|) - \\
&\quad - \sum_{t=K-|\mathcal{M}|+1}^K (K-t+1) + \sum_{t=K+1}^{t_{\max}} (ar - (k+K)).
\end{aligned}$$

Instead of writing the costs of pentesting per round, we rewrite them as a sum of costs per pentest,

$$G = \sum_{t=1}^K (a(r-z) - (k+t-1)) + |\mathcal{M}| \cdot (az-c) - \sum_{i=1}^{|\mathcal{M}|} \sum_{t=m_i}^{K-i+1} (1) + \sum_{t=K+1}^{t_{\max}} (ar - (k+K)).$$

Finally, we can write the expression for the total profit as

$$G = \sum_{t=1}^K (a(r-z) - (k+t-1)) + \sum_{i=1}^{|\mathcal{M}|} (az-c - (K-i-m_i+2)) + \sum_{t=K+1}^{t_{\max}} (ar - (k+K)). \quad (14)$$

The first sum is the contribution of attacks to the profit, the second sum shows the individual contributions of pentests and the last sum is the profit after the original attacks would have stopped without pentests.

A.2 Optimal Number of Pentests

Now we show a detailed derivation for the optimal number of pentests. Pentests are successful with probability p . Let $E[g]$ be the contribution of pentesting to the expected total profit (i.e., the second sum in (6)) and let us have a closer look at it. Clearly, penetration testing has to contribute a positive profit to be worth performing, i.e.,

$$E[g] = p \sum_{i=1}^{|\mathcal{M}|} (az-c - (K-i-m_i+2)) > 0. \quad (15)$$

Now we use Theorem 1 and replace m_i by $m_1 + i - 1$,

$$\begin{aligned} E[g] &= p \sum_{i=1}^{|\mathcal{M}|} (az-c - (K-m_1-2i+3)) \\ &= p|\mathcal{M}| \cdot (az-c - K + m_1 - 3) + 2 \cdot \sum_{i=1}^{|\mathcal{M}|} (i) \\ &= p|\mathcal{M}| \cdot (az-c - K + m_1 - 3) + |\mathcal{M}| \cdot (|\mathcal{M}| + 1) \\ &= p|\mathcal{M}| \cdot (az-c - K + m_1 + |\mathcal{M}| - 2). \end{aligned} \quad (16)$$

From (7) and using ϵ , we have:

$$K - m_1 = (1+p)|\mathcal{M}| - 1 - \epsilon. \quad (17)$$

Hence, we can rewrite (16) as

$$\begin{aligned} E[g] &= p|\mathcal{M}| \cdot (az-c - (1+p)|\mathcal{M}| + 1 + \epsilon + |\mathcal{M}| - 2) \\ &= p|\mathcal{M}| \cdot (az-c - p|\mathcal{M}| - 1 + \epsilon). \end{aligned} \quad (18)$$

The series of pentests is worth performing if the expected profit $E[g]$ is positive, meaning that

$$E[g] = p|\mathcal{M}| \cdot (az - c - p|\mathcal{M}| - 1 + \epsilon) > 0. \quad (19)$$

Since $0 < m_1 \leq K$, we have $|\mathcal{M}| > 0$ from (7) and we can write the condition for pentesting:

$$\begin{aligned} az - c - p|\mathcal{M}| - 1 + \epsilon &> 0 \\ az - c - p \left\lceil \frac{K + 1 - m_1}{1 + p} \right\rceil - 1 + \epsilon &> 0. \end{aligned}$$

We can obtain the optimal number of pentests $|\mathcal{M}|^*$ as the value that maximizes (18),

$$|\mathcal{M}|^* = \max_{|\mathcal{M}|} p|\mathcal{M}| \cdot (az - c - p|\mathcal{M}| - 1 + \epsilon) = -p^2|\mathcal{M}|^2 + p|\mathcal{M}| \cdot (az - c - 1 + \epsilon).$$

Derivation gives us the maximum value as follows:

$$|\mathcal{M}|^* = \frac{az - c - 1 + \epsilon}{2p}, \quad (20)$$

where $0 \leq |\mathcal{M}|^* \leq \left\lceil \frac{K + \epsilon}{1 + p} \right\rceil$. Note that the expression in (20) returns a real number that is optimal only asymptotically. The decision criterion can be discretized rounding off to the nearest integer or applying a randomized strategy.

A.3 Optimal Time to Start Pentesting

The substitution of (20) into (17) also gives us the optimal time to start pentesting m_1^* ,

$$m_1^* = K - \frac{1 + p}{2p}(az - c) + \frac{1 + 3p - \epsilon(1 - p)}{2p}, \quad (21)$$

where $0 < m_1^* \leq \left\lceil \frac{K + \epsilon}{1 + p} \right\rceil + \epsilon$ holds.

A.4 Expected Total Profit of Pentesting

Substituting the optimal number of pentests into (6), we obtain an expression for the expected total profit with optimal number of pentests,

$$\begin{aligned}
E[G] &= \sum_{t=1}^K (a(r-z) - (k+t-1)) + p \sum_{i=1}^{|\mathcal{M}|} (az - c - (K-i-m_i+2)) + \\
&\quad + \sum_{t=K+1}^{t_{\max}} (ar - (k+K)) \\
&= K(a(r-z) - (k-1)) - \frac{K(K+1)}{2} + \\
&\quad + p|\mathcal{M}|^*(az - c - p|\mathcal{M}|^* - 1 + \epsilon) + (t_{\max} - K)(ar - (k+K)) \\
&= K(a(r-z) - (k-1)) - \frac{K(K+1)}{2} + E[g] + (t_{\max} - K)(ar - (k+K)),
\end{aligned} \tag{22}$$

where $E[g]$ takes the values depending on the conditions in (20) as

$$E[g] = \begin{cases} 0, & \text{if } |\mathcal{M}|^* = 0; \\ \left(\frac{az-c-1+\epsilon}{2}\right)^2, & \text{if } 0 < |\mathcal{M}|^* < \left\lceil \frac{K+\epsilon}{1+p} \right\rceil; \\ \frac{p}{1+p} \left\lceil \frac{K+\epsilon}{1+p} \right\rceil \left(az - c - K - 1 + \frac{p}{1+p} \left\lceil \frac{K+\epsilon}{1+p} \right\rceil \right), & \text{if } |\mathcal{M}|^* = \left\lceil \frac{K+\epsilon}{1+p} \right\rceil. \end{cases} \tag{23}$$