

The Iterated Weakest Link

A Model of Adaptive Security Investment

Rainer Böhme^{a)} and Tyler Moore^{b)}

^{a)} Technische Universität Dresden, Inst. of Systems Architecture, Germany;

`rainer.boehme@tu-dresden.de`

^{b)} Harvard University, Center for Research on Computation and Society, USA;

`tmoore@seas.harvard.edu`

Abstract. We devise a model for security investment that reflects dynamic interaction between a defender, who faces uncertainty, and an attacker, who repeatedly targets the weakest link. Using the model, we derive and compare optimal security investment over multiple periods, exploring the delicate balance between proactive and reactive security investment. We show how the best strategy depends on the defender's knowledge about prospective attacks and the sunk costs incurred when upgrading defenses reactively. Our model explains why security underinvestment is sometimes rational even when effective defenses are available and can be deployed independently of other parties' choices. Finally, we connect the model to real-world security problems by examining two case studies where empirical data is available: computers compromised for use in online crime and payment card security.

Keywords: Economics of Information Security, Security Investment, Weakest Link, Uncertainty, Dynamic Model

1 Introduction

We hear about security breaches in the news almost daily, each bigger and more costly than the last. Does this reflect flawed technology, policy, or simply ineptitude? What if, instead, allowing some attacks to succeed is entirely rational? Rather than over-invest proactively, companies could wait to observe which attacks work and use this knowledge to better allocate security spending. In this paper, we describe a model that weighs the merits of such an approach.

One key insight from the economics of information security literature [1] is that attackers bent on undermining a system's security operate *strategically*. Moreover, information systems are often structured so that a system's overall security depends on its *weakest link* [2]. The most careless programmer in a software firm can introduce a critical vulnerability. The Internet's global, distributed architecture leads to security being dominated by the weakest link – attackers compromise machines hosted at Internet service providers (ISPs) with lax enforcement policies or located in uncooperative countries. Attackers have repeatedly exhibited a knack for identifying the easiest way to bypass a system's

security, even when the system’s designer remains unaware of the particular weakness.

However, systems do not exist in a vacuum; rather, defenders respond to attacks by plugging known holes. And yet, as soon as one weakest link is fixed, another weak point is often identified and exploited. Therefore, a strong dynamic component is at play: attackers find the weakest link, defenders fix the problem, attackers find new holes which are then plugged, and so on. We see this pattern emerge repeatedly. For instance, attackers construct networks of compromised machines (so-called botnets) to pester legitimate users by emitting spam, distributing malware and hosting phishing websites. Attackers concentrate their efforts at the most irresponsible ISPs, moving on to others only after the ISP cleans up its act or is shut down. Likewise, technical countermeasures to payment card fraud have evolved over time, causing fraudsters to adopt new strategies as old weaknesses are fixed. For example, when UK banks migrated to PIN verification of transactions rather than signatures, in-person retail fraud declined while overseas ATM fraud and card-not-present fraud skyrocketed.

In this paper, we devise a model that reflects this dynamic interaction between attackers and defenders. Our model captures the *iterative* aspect of attack and defense; we exclusively study the case where security depends on the *weakest link*; consequently, we model the *iterated weakest link*. Key characteristics of the model include:

1. Defensive countermeasures can be represented as interdependent; thus, the oft-cited diminishing marginal return of information security investment (e. g., [3]) becomes endogenous in our model.
2. Defender uncertainty about which components are weakest is captured.
3. A repeated game setting reflects the iterative process of attacking and defending successive weakest links.

We reach several interesting conclusions upon examining the model. By comparing the static case (a single round of attack and defense) to the dynamic one (multiple rounds), we find different defender strategies may prevail. When the defender only gets one chance to protect a system, increasing uncertainty about which link is weakest causes the defender to protect more assets, but only up to a point. When uncertainty is too high, the defender does not know which asset to protect and so chooses to protect none. If instead we allow for repeated defensive investments, an uncertain defender will initially protect fewer assets and wait for the attacker to ‘identify’ the weakest links to be fixed in later rounds. Hence, it can be quite rational to *under-invest* in security until threats are realized. Unlike in other theories, this type of under-investment is not driven by the interrelation with other market participants and resulting incentive systems. Of course, security countermeasures may require significant capital investment from the outset. When we introduce sunk costs to our model, we find that for moderate levels of uncertainty, higher sunk costs can raise the proactive protection investment.

We then translate these findings about optimal defensive strategies into accepted security indicators such as annual loss expectancy (ALE) and return on security investment (ROSI). Return on investment drops as uncertainty about

known attacks rises, even as the defender grows increasingly reactive when countering realized threats.

The remainder of this paper is organized as follows: Section 2 specifies the model, which is solved later on in Section 3. We apply the model to real problems in two case studies on current topics of information security, namely online crime (Sect. 4.1) and payment card security (Sect. 4.2). Section 5 puts our contribution into perspective with prior art, and Section 6 concludes the paper.

Readers who want to glance over the paper without getting into the formal details may find the model summary in Section 2.5 along with the interpretation of example results in Section 3.3 most useful. Also the case studies (Sect. 4) and the general conclusion (Sect. 6) do not require deep understanding of the model.

2 Model

Imagine a simplistic world, in which a *defender* protects an *asset* of value a against a dispersed set of possibly heterogeneous *attackers*. There exist n possible *threats*, which can be regarded as distinct attack vectors against a single system.¹ Each threat can be warded off by investing in its corresponding *defense* (or control). In other words, we assume a one-to-one mapping between threats and defenses, and defenses are always effective.

2.1 Defender's options

The model is ‘run’ in an iterated game with discrete time t . In each round, the defender makes a security investment decision to define his *configuration* of defenses \mathbf{d}_t and extracts a net return $r \cdot a$ from his asset before the attacker penetrates the system and, if successful, loots a fraction z of the asset. Gross returns are consumed or distributed so that the asset value does not accumulate over time.

Let elements d_i of the binary column vector $\mathbf{d} \in \{0, 1\}^n$ indicate whether a defense against the i -th threat is implemented ($d_i = 1$) or not ($d_i = 0$), and let $k = \sum_i^n d_i$ be the number of defenses in place.

The cost of defense c_t in round t can be calculated from an $n \times n$ upper triangular cost matrix \mathbf{C} to reflect possible interdependent defenses,

$$c_t = \mathbf{d}_t \mathbf{C} \mathbf{d}_t . \tag{1}$$

Diagonal elements $C_{i,i}$ hold the cost to implement a defense against the i -th threat, and off-diagonal elements $C_{i,j}, j > i$ indicate the extra costs if the i -th defense is implemented together with the j -th defense. If all off-diagonal elements are zero, then the defenses are independent:²

¹ An alternative interpretation is that threats represent distinct targets in a distributed system that together form asset a . This interpretation comes close to the notion of *weakest target* as opposed to weakest link games in [4].

² We do not consider higher-order interdependence such as extra costs if three or more defenses are involved.

$$\mathbf{C} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix} \qquad \mathbf{C} = \begin{bmatrix} 1 & 0 & 1/2 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1/2 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

- (a) independent defenses (b) conflicting defenses: 1 and 3
complementary defenses: 3 and 4

Why go to the trouble of modeling interdependent defenses? First, organization science suggests that complexity adds super-linear administrative costs, for instance to pay a manager who directs several employees where each is an expert for a particular control. Conflicting defenses also emerge from incompatible systems, e.g., running two virus scanners on the same machine to increase coverage slows down the machine and causes errors. Lastly, human behavior fundamentally limits the composability of defenses: a password policy that requires both special characters (defense 1) and frequent password changes (defense 2) encourages people to stick their password on their monitor.

A nice property of the cost matrix is that for positive off-diagonal elements, decreasing marginal utility of defenses has become endogenous to our model. This compares favorably to, say, the Gordon–Loeb framework, in which this property appears as an assumption (A3 of [3, p. 443]).

Sunk costs

So far we assume that the defender can upgrade his configuration \mathbf{d}_t at any time. This may be necessary to adjust it to new information on the threat level or to changing risk appetite. For example, a start-up company is exposed to so many risks that is might tolerate a moderate level of information security risk. As the venture grows and develops a brand name, its risk aversion will increase to reflect the higher damage caused by potential reputation losses. Conversely, market competition (alongside herd behavior and short-sighted incentives to management) may drive firms to decrease their risk aversion for the sake of higher expected profits.

As is argued in [3], security disinvestment is often possible to a certain degree, since personnel can be fired or equipment sold (though sometimes at high transaction costs). However, for many businesses, updating \mathbf{d}_t may be very expensive and the costs are ‘sunk’ in the sense that money is spent irrevocably. For instance, the vast majority of the cost to incorporate new security features into bank notes or payment cards are borne when the changes are first made. Producing and distributing tailored tokens or devices to a large and dispersed community is expensive, and unlikely to be repeated often. As sunk costs considerably affect the strategy of security investment in an iterated setting, we include them in our model as follows:

$$s_t = \begin{cases} 0 & \text{if } t = 1 \text{ or } \mathbf{d}_t = \mathbf{d}_{t-1} \\ \lambda \cdot a & \text{else.} \end{cases} \quad (2)$$

Hence, parameter $\lambda \geq 0$ controls the amount of sunk costs.

2.2 Defender’s knowledge

While the defender may possess some intuition about the relative difficulty of carrying out the n threats, such knowledge may very well be blurred. To model this uncertainty, we order the threats $1, \dots, n$ by *increasing expected cost of attack* (the expectation is taken from the point of view of the defender).³ This constitutes our notion of an *attack profile*. We define a simple functional form for the expected attack costs \bar{x}_i of the i -th threat as follows,

$$\bar{x}_i = \bar{x}_1 + (i - 1) \cdot \Delta x \quad , \text{ with } \Delta x > 0 . \quad (3)$$

The unknown true cost x_i , however, is modeled as a Gaussian random variable with mean \bar{x}_i and standard deviation $\sigma/\Delta x$ (censored to values $x_i \geq 0$),

$$x_i = \sup(0, \chi_i) \quad \text{with} \quad \chi_i \sim \mathcal{N}(\bar{x}_i, \sigma/\Delta x) . \quad (4)$$

Note that the realizations χ_i are drawn only once and keep constant over time. We can vary the level of uncertainty by adjusting parameter σ . Modeling uncertainty in this way is crucial to the model, since it captures the difficulty defenders face in anticipating which of the undefended threats turns out to be the weakest link exploited by the attacker. Figure 1 illustrates the role of uncertainty when ordering threats. The left graph plots a perfect matching between expected and realized costs of attack under certainty; under uncertainty, by contrast, the right graph shows that threat 4, not threat 3 as expected, is the weakest link if defenses 1 and 2 are in place.

The defender’s knowledge about the attack profile may increase over time when observed attacks reveal which threat is the weakest link with respect to a given configuration \mathbf{d}_t .

2.3 Attacker’s options and knowledge

Our attacker model is very simple: the attacker identifies and exploits the weakest link, i. e., the threat least costly to the attacker. We do not require the attacker to make a trade-off between cost and potential gains by assuming the same utility is received for exploiting all threats. If the attacker succeeds, regardless of how, he profits $z \cdot a$, which is added to the defender’s cost. Unlike the defender, the attacker is certain of the cost for each realization of x_i . The attacker does not operate indiscriminately; rather, he only attacks when it is profitable to do so, that is, if the term $\max_i(z \cdot a - x_i)$ is not negative.

³ Ordering the inputs this way is also convenient because it allows us to ignore the threats deemed too unlikely when deciding what to defend. This way, we can think of the n threats as those anticipated by the defender.




Fig. 1. Attack profile: Modeling uncertainty about true attack costs

2.4 Simplifying assumptions

We acknowledge that the attacker assumptions are quite strong and probably unrealistic in some cases. For targeted attackers in particular, searching for the weakest link could be very expensive (though the chance that one of many opportunistic attackers [5] finds the hole accidentally is quite high.) Moreover, the attacker’s gain from successful attacks may only represent a small fraction of the defender’s losses. But this is not crucial here since we do not consider attackers’ profits; we simply need an asset-dependent quantity to relate to the cost x_i . A level shift of x_i could easily account for the negative balance (destruction of wealth) after successful attacks. This assumption is also common in the literature [6], and justified there by attackers who mainly care about inflicting harm on the defender. In sum, we consider it prudent to err in the direction of a more powerful attacker than an unrealistically weak one.

To make the analysis more tractable, we restrict the space of possible cost matrices \mathbf{C} by fixing all diagonal elements $C_{i,i}$ to a unit cost of value one and all off-diagonal elements $C_{i,j}$, $j > i$ to the same constant $\varrho \geq 0$. Hence, we can now derive from Eq. (1) a simple expression for the cost per round to maintain k defenses,

$$c_t = \frac{\varrho}{2}k^2 + \left(1 - \frac{\varrho}{2}\right)k. \quad (5)$$

Observe that $c_t \sim \mathcal{O}(k^2)$ if $\varrho > 0$ and $c_t \sim \mathcal{O}(k)$ if $\varrho = 0$. Of course, given better data or assumptions on the cost of defenses, these assumptions can be relaxed and \mathbf{C} populated with other values.

We further assume a *risk neutral* defender, that is, one who maximizes intertemporal utility as a linear function of (expected) revenues and costs. We defer

consideration of other attitudes towards risk to future work in order to keep the core model lean.

2.5 Model summary and example parameters

Our model can be applied to many scenarios. For now, we consider an imaginary online music store to demonstrate its parameters in action (we describe other, more compelling scenarios in Sect. 4). The music store’s asset is the library of audio files and the associated property rights. Let the total value of the asset be one million dollars. To avoid large numbers, we count in thousands of dollars, so asset value $a = 1\,000$. Every year, indexed by t , the net return from online sales amounts to $r = 5\%$ of the asset value.

Table 1. Overview of model parameters

Parameter	Symbol	Example values
Business model		
asset value	a	1 000
time in years	t	$1, \dots, t_{\max}$
return	r	5.0 %
Attacker		
number of threats	n	25
loss given attack	z	2.5 %
expected minimum attack cost	\bar{x}_1	15
gradient of attack cost	Δx	1
level of uncertainty	σ	$0, \dots, 16$
Defense		
cost of each individual defense	1	1 (unit definition)
defense interdependence	ϱ	0.1
sunk costs	λ	0

Surrounded by malicious competitors and professional cyber-criminals, the store is exposed to various threats, say, $n = 25$. In case of a successful attack, the store loses $z = 2.5\%$ of the asset value, or \$25 000, due to forgone sales and incurred damage. It is known (from industry sources) that, *on average* over all similar businesses, the cost of a successful attack is at least 15 000 dollars ($\bar{x}_1 = 15$). This could be the price for a standard exploit on the vulnerability black market [7, 8]. Attack costs also rise with novelty and difficulty of implementation, *on average* by 1 000 dollars for each level of sophistication. Hence, the gradient of the attack cost is $\Delta x = 1$. However, these are only average values and it is unknown how much the actual costs are to attack the particular business. This uncertainty can be expressed by the deviation σ of the attack costs around the

average. To put this parameter in perspective, $\sigma = 0$ means that all attacks can be predicted perfectly, $\sigma = 1$ means 96 % of the attacks can be predicted correctly,⁴ and for $\sigma = 16$, predicting attacks is only 65 % successful. This is still somewhat better than random guessing.

The store owner is not entirely helpless: he may spend on security measures to defend against each of the $n = 25$ threats. Each defense costs 1 000 dollars per round (this is a constant that defines the unit, so no symbol is required). On top of that, defense interdependence with parameter $\varrho = 0.1$ accounts for additional cost for every pair of installed defenses. As a consequence, every investment in security must be carefully considered, since countermeasures added later cost more due to increased complexity and interoperability challenges. Finally, parameter λ controls (optional) sunk costs. It describes the cost to update the defense configuration as a fraction of the asset value a . In the example of a music store, sunk costs are negligible (i. e., $\lambda = 0$) since security upgrades do usually not render existing investment useless. We include sunk costs mainly to model scenarios where upgrading defenses requires replacing existing infrastructure, such as exchanging all customers' payment cards or security tokens.

Table 1 summarizes the parameters of our model and their values used in the subsequent examples. The following section will show how defenders can optimize their security spending in this model. Our main interest is on the reaction of the optimal strategy to changes in the level of uncertainty σ ; we also elaborate on the role of sunk costs λ .

3 Analytical results

We distinguish between *static* analysis, where the defender chooses the configuration in the first round ($t = 1$) and keeps it constant further on, and *dynamic* analysis, where the configuration is updated in every round. In case of certainty ($\sigma = 0$), a static solution exists that is at least as good as any dynamic strategy. Sunk costs only occur in the dynamic setting.

3.1 Static solution

We consider a defender with defense vector \mathbf{d} where the k lowest cost threats are defended, leaving the remaining $n - k$ threats unprotected: $d_1, \dots, d_k = 1$ and $d_i = 0$, $k < i \leq n$. Let $f(k)$ be the return function for this defender,

$$f(k) = a(r - zq) - c_t \tag{6}$$

$$= a(r - zq) - k - \frac{\varrho}{2} (k^2 - k) \ , \tag{7}$$

⁴ all other parameters as specified in Table 1

where q is an indicator variable that takes value one if the attacker is successful and zero otherwise. In the case of **certainty**, q is fixed for all rounds, so that two outcomes are possible:

$$f_{\text{att}} = a(r - z) \quad (\text{attack}) \quad \text{and} \quad (8)$$

$$f_{\text{att}}(k) = ar - k - \frac{\theta}{2} (k^2 - k) \quad (\text{no attack}). \quad (9)$$

The cost term disappears in Eq. (8) because if we accept attacks anyway, then there is no need to spend money to make it more difficult for the attacker (i. e., $k = 0$). Otherwise, the minimal number of defenses k^* to defeat all attacks can be calculated by rearranging Eq. (3) and comparing it to the attacker's gain za ,

$$k^* = \left\lfloor \frac{za - \bar{x}_1 + \Delta x}{\Delta x} \right\rfloor. \quad (10)$$

Typically, the best strategy for the defender is to follow either Eq. (8) or (9), whatever is higher. When Eq. (9) is negative and Eq. (8) is positive, the defender finds it profitable to operate without investing in any security measures. When both equations are negative, the cost of attacks and the cost of defending them are both so high that the defender's business becomes unviable.

In the case of **uncertainty** ($\sigma > 0$), q becomes a random variable and a risk neutral defender maximizes $E(f(k))$, where expectations are taken over the realizations of q . For a static solution, it is still rational to use configurations in which lower-order threats are defended first. So we can calculate $E(q)$ as a function of k as the probability that the cost of at least one of the unprotected threats falls below the threshold za ,

$$E(q(k)) = 1 - \prod_{i=k+1}^n (1 - \Phi(z \cdot a; \bar{x}_1 + (i-1) \cdot \Delta x, \sigma/\Delta x)) \quad , \quad k < n \quad (11)$$

where $\Phi(x; \mu, \sigma)$ is the cumulative distribution function of the Gaussian normal distribution. The product term is the probability that none of the $n - k$ undefended threats materializes, so one minus this term is the probability of at least one successful attack. Inserting into Eq. (7), we obtain

$$k^* = \arg \max_k E(f(k)) = \arg \max_k a(r - z E(q(k))) - k - \frac{\theta}{2} (k^2 - k) \quad . \quad (12)$$

Figure 2 depicts return estimates of the static solution for typical parameters as a function of k . Each curve corresponds to different levels of uncertainty; the utility-maximizing selection of k^* for each uncertainty level is highlighted in yellow. In the case of certainty ($\sigma = 0$), the optimal security investment amounts to $k = 11$ defenses (Eq. (10)). Increasing uncertainty affects the optimal strategy in a non-linear manner: while some noise in the predicted attack costs lets the defender take more proactive counter-measures ($\sigma = 1, 2$), the cost of this over-investment start to outweigh the protected revenues when uncertainty is too high ($\sigma = 4, 8$). Eventually, it becomes rational for more uninformed




Fig. 2. Static solutions for the following set of parameters: asset value $a = 1\,000$, return $r = 5\%$, loss given attack $z = 2.5\%$, min. expected cost of attack $\bar{x}_1 = 15$, gradient of attack cost $\Delta x = 1$, defense interdependence $\varrho = 0.1$, and $n = 25$. The optimal number of defenses is highlighted in yellow for each curve.

defenders to refrain from *any* defense and instead accept moderate losses in each round. In other words, knowing your enemy is a prerequisite for taking tailored countermeasures, and tailored defense is the only option if comprehensive defense (i. e., $k = n$) is prohibitively expensive.

3.2 Dynamic (iterated) solution

When $\sigma > 0$ (uncertainty), the static solution is often sub-optimal because observed attacks reveal additional information about the true cost of attack which should have been used to re-configure the defenses adaptively. Under the dynamic case, uncertainty is repeatedly eliminated for targets revealed to be the weakest link. Consequently, this reduction in uncertainty for later rounds can lead to a better outcome for the defender than in the static case. In general, the inter-temporal outcome is the sum of all returns over a period of time $t = (1, \dots, t_{\max})$,

$$y = \sum_{t=1}^{t_{\max}} f(\mathbf{d}_t, \mathbf{d}_{t-1}) = \sum_{t=1}^{t_{\max}} (a(r - zq_t) - c_t - s_t) , \quad (13)$$

and can be maximized by adjusting \mathbf{d}_t adaptively. At each point in time t , the defender learns about the materialized threat (if any) of period $t - 1$. In Eq. (13), c_t , s_t and q_t , are functions of the sequence of configurations \mathbf{d}_t and the random

realization of \mathbf{x} . Due to this stochastic component and its subsequent path dependencies through adaptation, the analysis must be based on expectations over all possible realizations. We will first discuss the simpler case without sunk costs (i. e., $\lambda = 0$) before moving on to the general case.

Without sunk costs, the defender has to choose a *proactive defense* \mathbf{d}_1 against the k_1 most probable threats, so that $d_{i,1}, \dots, d_{k_1,1} = 1$ and $d_{i,1} = 0$, $k_1 < i \leq n$. Unless this configuration is already secure, the defender would iteratively add another defense in each round as the attacker targets new weak links. The order of these *reactive defenses* is determined by the order of x_i , and thus by the realization of the random vector \mathbf{x} . This way, one weakest link is fixed after another. It follows from the path dependencies that once a secure configuration is found, the defender would not touch \mathbf{d} anymore: in this model, it is always better to start with a lower proactive defense k_1 than tinkering with existing defenses after a couple of rounds. (This is so because the uncertainty about x_1, \dots, x_{k_1} is not reduced. And the direct and indirect ($\varrho > 0$) cost of the k_1 -th defense has to be born for all intermediate rounds.)

Let $t_{\text{att}} \leq t_{\text{max}}$ be the number of rounds with successful attacks, then $\mathbf{d}_{t_{\text{att}}+1} = \mathbf{d}_t \forall t > t_{\text{att}}$ is the final secure configuration that is eventually reached if $t_{\text{att}} < t_{\text{max}}$. So Eq. (13) without sunk costs can be rewritten as

$$y = \sum_{t=1}^{t_{\text{att}}} (a(r-z) - c_t) + \sum_{t=t_{\text{att}}+1}^{t_{\text{max}}} (ar - c_t) \quad (14)$$

$$= t_{\text{att}} \cdot a(r-z) - \sum_{t=1}^{t_{\text{att}}} c_t + (t_{\text{max}} - t_{\text{att}})(ar - c_{t_{\text{max}}}), \quad (15)$$

after plugging in Eq. (5) we obtain:

$$= t_{\text{att}} \cdot a(r-z) + (t_{\text{max}} - t_{\text{att}})(ar - c_{t_{\text{max}}}) - \frac{\varrho}{2} \sum_{k=k_1}^{k_1+t_{\text{att}}} k^2 - \left(1 - \frac{\varrho}{2}\right) \sum_{k=k_1}^{k_1+t_{\text{att}}} k. \quad (16)$$

Eq. (15) holds because c_t is constant when \mathbf{d}_t is constant. A closed-form expression for Eq. (16) can be obtained by replacing the remaining sums by first and second order arithmetic progression formulas. This expression gives the output y as a function of the proactive defense k_1 and the number of rounds under attack t_{att} . The latter parameter in fact depends on the random realization of \mathbf{x} . So the expectation of the term in Eq. (16) can be calculated by the sum of all possible values for $0 \leq t_{\text{att}} \leq t_{\text{max}}$, weighted by the respective probability of occurrence $\mathcal{P}(t_{\text{att}})$.

To calculate $\mathcal{P}(t_{\text{att}})$, we interpret the condition for a successful attack $x_i \leq z \cdot a$ as Bernoulli random variable with probability of attack

$$p_i = \Phi(z \cdot a; \bar{x}_1 + (i-1) \cdot \Delta x, \sigma/\Delta x). \quad (17)$$

Hence, \mathbf{p} is the parameter vector of non-homogeneous independent Bernoulli variables. For reasonable parameter choices, the sum of these random vectors can




Fig. 3. Dynamic solutions for the same set of parameters as in Fig. 2: asset value $a = 1000$, return $r = 5\%$, loss given attack $z = 2.5\%$, min. expected cost of attack $\bar{x}_1 = 15$, gradient of attack cost $\Delta x = 1$, defense interdependence $\varrho = 0.1$, and $n = t_{\max} = 25$. The optimal number of proactive defenses is highlighted in yellow for each curve.

be approximated by the Gaussian distribution with location $\mu = \sum_{i=1}^n p_i$ and variance $\varsigma = \sum_{i=1}^n p_i(1 - p_i)$. This completes the analysis. Next, we compute the optimal choices for the proactive defense k_1^* (that maximize the expected intertemporal outcome) with numerical methods by searching for the supremum in the integer range $k_1 = 0, \dots, n$.

Figure 3 shows return estimates *per round* of the dynamic solution for the same parameters as in Figure 2, now as a function of the *proactive defense* k_1 . (As in Figure 2, the optimal proactive defense for each value of σ is highlighted in yellow.) Obviously, in the case of certainty ($\sigma = 0$), the optimal proactive defense $k_1^* = k^* = 11$ equals the result of the static analysis. If there is no uncertainty, a defender does not gain any information from observed attacks and can define the optimal configuration entirely proactively. However, as soon as uncertainty arises, the optimal strategies in the static and dynamic case diverge. While some uncertainty ($\sigma > 0$, but small) makes a static defender more cautious (to reduce the probability of being vulnerable) and leads him to spend more, a dynamic investor knows that he will be able to adjust later and can thus take on more risk in the first place. The higher the uncertainty, the more valuable is information gleaned from observed attacks, and the lower the optimal share of proactive defense becomes. Since it is not economically viable for static defenders to compensate for uncertainty with ever higher security spending, there exists a threshold above which any security investment cannot defeat attacks. In this

case, the defender ‘surrenders’ and accepts the losses in each round. Observe that this state is already reached for $\sigma = 4$ in Fig. 2, whereas much higher uncertainty is needed to refrain from all proactive investment in the dynamic case ($\sigma = 16$ in Fig. 3). So we gain the insight that the sheer possibility of adjusting defenses reactively can stimulate proactive security investment if the uncertainty is very high.

To consider **sunk costs**, Eq. (13) is rearranged to the form of Eq. (14) as follows:


$$y = \sum_{t=1}^{t_{\text{att}}} (a(r - z) - c_t - s_t) + \sum_{t=t_{\text{att}}+1}^{t_{\text{max}}} (ar - c_t) . \quad (18)$$

The second sum does not include s_t because \mathbf{d}_t remains constant for $t > t_{\text{att}}$. Rearranging and inserting Eq. (2) yields


$$y = \lambda \cdot a + t_{\text{att}} \cdot a(r - z - \lambda) - \sum_{t=1}^{t_{\text{att}}} c_t + (t_{\text{max}} - t_{\text{att}})(ar - c_{t_{\text{max}}}) . \quad (19)$$

Now we proceed as with Eq. (16) in the case without sunk costs. It is possible to stick with the same strategy of iterated investment because the elements of the random vector \mathbf{x} are independent and observing one attack does not shed any light on what the next weakest link will be. So even with sunk costs, it would be irrational to invest in bundles of reactive defenses if one had not done so right at the beginning. Nevertheless, the existence of non-negligible sunk costs substantially changes the effort spent on proactive measures. For a finite t_{max} , however, it is rational for the defender to reflect on the strategy after having observed an attack in t_1 : for very high λ , he might prefer not to defend at all and accept losses in the remaining $t_{\text{max}} - 1$ rounds. This is like a reduction to Eq. (8) of the static case. So effectively, introducing sunk costs in our model can be seen as an increase in z for t_{att} rounds (if the iterated patching starts) combined with the option to fall back to the static case (without the increase in z).

In Figure 4, we fix two settings for the level uncertainty, moderate ($\sigma = 4$) and high ($\sigma = 7$), and plot expected returns per period as a function of the proactive defense k_1 for various settings of λ . Observe in Fig. 4 (a) that higher sunk costs trigger greater precaution and force the defender to invest more in proactive security. This jibes with common sense. However, if the uncertainty is high, then there exists a threshold for the sunk costs above which a defender ‘surrenders’ and prefers to cope with the attacks, as visible in Fig. 4 (a) for values of $\lambda > 1\%$ of the asset. The similarity to the static case can also be seen by comparing the similar shape of the curves. Note that these results should be interpreted with caution as they depend on a finite horizon t_{max} , which might not apply to all business cases.



(a) moderate uncertainty: $\sigma = 4$



(b) high uncertainty: $\sigma = 7$

Fig. 4. Dynamic solutions with sunk costs for moderate (top) and high (bottom) uncertainty. Same parameters as in Fig. 2: asset value $a = 1000$, return $r = 5\%$, loss given attack $z = 2.5\%$, min. expected cost of attack $\bar{x}_1 = 15$, gradient of attack cost $\Delta x = 1$, defense interdependence $\varrho = 0.1$, and $n = t_{\max} = 25$. λ controls the height of sunk costs (as fraction of a), and σ the degree of uncertainty. The optimal number of proactive defenses is highlighted in yellow for each curve.

3.3 Quantitative results

As we have analyzed the optimal strategies of the defender in the previous sections, we can now use the model to calculate some key indicators that relate to practical experience. Table 2 compares typical security indicators derived from the model for different degrees of uncertainty (in columns) and different defense strategies (row sections – static / dynamic / dynamic with sunk costs). The list of indicators includes:

- **Optimal defense** is a memo item that describes the optimal defense strategy $k_{(1)}^*$. All other indicators are computed on the assumption that a rational defender follows this strategy.
- **Attack intensity** measures the probability of a successful attack per round, averaged over all t_{\max} rounds.
- **Average gross return** is the objective function of the optimization problem, scaled to a percentage of the asset value a . The average is computed over all t_{\max} periods and expectations are taken over all realizations of the stochastic component \mathbf{x} , weighted by their respective probability of occurrence. Higher values are better.
- **Average security spending** is the average effort spent on defenses, comprising direct, indirect and sunk costs (if applicable). Averages and expectations are computed as above. Lower values are not necessarily better; what matters is the efficiency of the security spending, i. e., how security spending compares against prevented losses. This depends on how targeted the defense configuration is to the realization of the attack profile.
- **Annual loss expectation (ALE)** measures expected foregone profits per period [9].⁵ ALE can be computed for two scenarios: ALE_0 is the (hypothetical) ALE if no defenses were deployed, and ALE_1 is the ALE of the respective defense strategy. Expectations are taken over realizations of the random vector \mathbf{x} .
- **Return on information security investment (ROSI)** is a summary measure to quantify the effectiveness of security investment. We follow the approach in [10] and normalize the indicators by the security spending, i. e.,

$$ROSI = \frac{ALE_0 - ALE_1 - \text{avg. security spending}}{\text{avg. security spending}}. \quad (20)$$

Higher values denote more efficient security investment. The indicator is not defined for cases where no defenses are optimal.

Unsurprisingly, the three modes of adaptivity concur in the case of certainty (leftmost column in Table 2). However, uncertainty creates some remarkable results, which we now discuss.

For moderate levels of uncertainty ($\sigma = 1$), replacing a static with a dynamic defense strategy primarily reduces security spending, which leads to more observed attacks. However, gross returns increase, too. So in fact, security spending

⁵ We stick to the term ‘annual’ to be consistent with the literature and assume that time periods in the model correspond to financial years.

Table 2. Key security investment indicators derived from the model

Indicator	Level of uncertainty			
	$\sigma = 0$	$\sigma = 1$	$\sigma = 4$	$\sigma = 8$
Static defense				
optimal defense k^*	11	12	0	0
attack intensity (% rounds)	0.0	2.4	100.0	100.0
avg. gross return (% asset)	3.4	3.1	2.5	2.5
avg. security spending (% asset)	1.6	1.9	0.0	0.0
ALE ₁	0.0	0.6	25.0	25.0
ROSI (% security spending)	51.5	31.2	—	—
Dynamic defense w/o sunk costs				
optimal proactive defense k_1^*	11	9	7	3
attack intensity (% rounds)	0.0	6.1	15.7	32.7
avg. gross return (% asset)	3.4	3.2	3.0	2.7
avg. security spending (% asset)	1.6	1.5	1.6	1.4
ALE ₁	0.0	1.5	3.9	8.2
ROSI (% security spending)	51.5	52.8	35.2	18.9
Dynamic defense w/ sunk costs				
optimal proactive defense k_1^*	11	10	9	0
attack intensity (% rounds)	0.0	2.9	9.8	100.0
avg. gross return (% asset)	3.4	3.2	2.8	2.5
avg. security spending (% asset)	1.6	1.6	1.9	0.0
ALE ₁	0.0	0.7	2.5	25.0
ROSI (% security spending)	51.5	50.6	15.7	—
Memo items: no defense				
avg. gross return (% asset)	2.5	2.5	2.5	2.5
ALE ₀	25.0	25.0	25.0	25.0

Parameters: asset value $a = 1000$, return $r = 5\%$, loss given attack $z = 2.5\%$, min. expected cost of attack $\bar{x}_1 = 15$, gradient of attack cost $\Delta x = 1$, defense interdependence $\varrho = 0.1$, sunk costs $\lambda = 2.5\%$, $n = 25$, $t_{\max} = 25$. All figures rounded to 1 digit.

is better targeted, over-investment is reduced, and the overall efficiency of security investment, as measured by the ROSI indicator, improves. From this result, we can draw an alternative interpretation to the omnipresent reports of security breaches in the media: rather than rashly framing them as engineering failures, one can also view them as unavoidable side-effects of smart defense strategies that balance the amount of proactive and reactive security investment.

Regarding attack intensity, a converse effect is observable when uncertainty is higher ($\sigma \geq 4$). Here, the attack intensity drops substantially from (almost) 100 % in the static case to a very low rate in the dynamic case. At the same time, average returns increase. The reason for this is that high uncertainty lets some defenders refrain from deploying any defenses (and thereby absorbing losses). Only a staged approach gives these investors an incentive to defend against the most aggressive threats. Note that the model identifies the rational response to the *private* costs faced by defenders. We ignore the *public* costs created by insecurity. While not addressed by this paper, negative externalities of poor security practices [2, 11, 12] are widely considered to be a significant public problem. Hence, while it may be narrowly better for some defenders to skimp on security when they are unsure whether they will be targeted, a public policy response may be necessary to compensate for the negative externalities of insecurity caused by such under-investment.

The presence of sunk costs muddies this picture somewhat. Sunk costs raise incentives towards more proactive security investment. This leads to higher proactive defenses and thereby lower attack intensity. However, sunk costs also eat up margins, so gross returns tend to decline. Even more importantly, the overall efficiency of the security investment drops, and the differential grows as uncertainty rises. So moderate uncertainty combined with expensive infrastructure upgrades can stimulate over-investment – but only to a certain point. Eventually, high uncertainty can demotivate any security investment, similar to what happens in the static case. In the real world, we would expect to see relatively more weight on proactive security when sunk costs are high (e. g., payment cards) and uncertainty is low or moderate, compared to businesses with negligible sunk costs (e. g., web sites). This prediction could be tested against cross-sectional empirical data in future work. In any case, the above-mentioned excuse of breaches as smart defense strategy does not apply to defenders with high sunk costs. The policy dimension here is to reduce the uncertainty when it is so high that it discourages security investment. This can be done, for instance, by encouraging information sharing [13, 11] so that the cost of reducing the uncertainty can be shared between several defenders.

Note that all these conclusions remain fairly robust for other parameters than the example set in Table 1. For brevity, we refrain from discussing the influence of all other parameters. Instead, we demonstrate the validity of our model by applying it to two case studies where real data is available.

4 Empirical evidence for iterated weakest links

The observation of imperfect security and gradual improvement is at the core of the endeavor to study information security with means and methods borrowed from economics [14, 1]. In the following, we will point the reader to past ‘iterated arms races’ in the realm of information security. And we argue that such dynamics, against the backdrop of our model, may in fact be rational and should not simply be framed as ‘failures’. We have selected two cases, one with seemingly negligible sunk costs (Sect. 4.1) and another one where sunk costs should not be ignored (Sect. 4.2).

4.1 Case 1: Phishing and online crime

It can be argued that the security of computers comprising the Internet follow an iterated weakest link model. Fundamental to the Internet’s design is the notion that any connected computer can communicate with every other computer. Attackers have relentlessly exploited this structure, constructing networks of compromised machines (so-called botnets) to pester legitimate users by emitting spam, distributing malware and hosting phishing websites. Internet service providers (ISPs) are usually well-placed to detect infection, because evidence of a user’s infection flows over an ISP’s communication systems. Moreover, large companies that act as Internet service providers have technical staff who can detect and clean up infected machines, while domestic users and small businesses are mostly helpless to even recognize when they are compromised, let alone to take appropriate remedial action. However, for every responsible ISP that keeps a clean network, there are countless others who find it more cost-effective to ignore the problem and let infections flourish.

Moreover, there is substantial evidence that attackers concentrate their efforts at the most irresponsible ISPs, moving on to others once the ISP cleans up its act or is shut down. For instance, until 2007, a lot of the world’s malware was hosted by the Russian Business Network (RBN), which refused to comply with requests from international law enforcement [15]. After Russian Business Network suddenly went offline in late 2007, malware distribution shifted to other ISPs. In November 2008, a journalist from the Washington Post persuaded upstream bandwidth providers to shut off their connection to San Francisco-based McColo [16], which led to a temporary fall of almost 70% in the volume of spam worldwide. Apparently, many botnet herders had been using McColo to host their control machines. Spam volume has since recovered as the spammers found new safe havens. And recently EstDomains, which served as the primary domain name registrar for malicious websites hosted by Russian Business Network [17], became the first domain registrar to have its accreditation terminated by the Internet Corporation for Assigned Names and Numbers [18], the not-for-profit corporation that plays a leading role in coordinating the Internet’s naming system. Unfortunately, other irresponsible registrars remain.

Attackers can often move on to exploit the next weakest link in the Internet far faster than defenders can knock them out. Moore and Clayton [19] showed

how one leading group of online fraudsters, the rock-phish gang, operates. It registers many malicious web domain names to carry out attacks. Periodically, the gang picks a firm for registering Internet domain names that it has never used before and registers hundreds of web domains, using false names and stolen credit cards. These domains do not resemble normal bank names, so a registrar may not identify foul play, and the first domains to be used for false purposes are not removed quickly.




Fig. 5. Take-down speed for phishing attacks targeting different domains.

Figure 5 presents scatter plots of phishing site lifetime based on the date reported from data of [19]. Both .hk (Hong Kong) domains (left) and .cn (China) domains (right) lasted much longer in their first month of use than in later months. The gang targeted Hong Kong domains first in March 2007, followed by Chinese domains in May 2007 after the Hong Kong authorities wised up. Other researchers have found that websites hosting most malware move from one registrar to the next [20].

4.2 Case 2: Payment card security

Financial transactions have long been subject to fraud, and banks and payment processors have recognized that fraud has to be managed rather than eliminated. In 2000, the UK banks decided to adopt the EMV card payment system designed by Europay, Mastercard and Visa, popularly known as ‘Chip and PIN’ [22]. The Chip and PIN roll-out provides a compelling example of taking security decisions to tackle the weakest link, only to find that the attackers react by exploiting new holes previously ignored. A May 2007 internal report prepared by APACS, the UK payment association, compared the expected benefits from EMV adoption to fraud incidence over time. While the report was intended to be kept confidential, it was accidentally leaked and is now hosted on the Cryptome website [21]. We use the data gleaned from this report to demonstrate how the bank’s fight against card fraud has essentially been a sequence of patches to plug weak links.




Fig. 6. Financial losses due to fraud over time. On top, fraud rates as a fraction of overall turnover from 1972–2006. On bottom, ten-year fraud projections made in 2000 comparing whether Chip and PIN is adopted (blue and red lines), plus actual fraud rates (green line) (Source: [21]).

Figure 6 (top) plots the annual percentage of transaction turnover found to be fraudulent from 1972 to 2006. Notably, there are significant swings in the rate. According to the report, “the industry has responded to crises in card fraud by making significant investment in new preventative measures,” including “terminalisation in the early 80s and, later, a major increase in levels of authorisation during the mid 90s.” Such behavior is consistent with our model of iterating through successive weakest links.

Around 2000, a fear arose that another uptick in the fraud rate was imminent. The blue diamond line in Figure 6 (bottom) plots the forecast losses due to fraud based on maintaining the status quo as of 2000. To stem this expected rise in fraud, UK banks decided to adopt the EMV standard (Chip and PIN) over the course of the next few years. The estimated cost of implementing EMV across the UK was just over £1 billion, to be shared between the banks and retailers. We can interpret this investment to roll out EMV as *sunk costs*. The red square line in Figure 6 (bottom) plots the forecast losses following EMV adoption. The uptick to 2003 is due to the phased roll-out of the technology.

In fact, the actual losses look rather different (green triangle lines in Figure 6 (bottom)). The anticipated explosion of fraud in the early 2000s did not materialize; strikingly, the full adoption of Chip and PIN (completed in 2005) did not significantly alter the fraud totals. Instead, a slow but steady increase was observed, both before and after EMV adoption.

Digging a little deeper, we can see that the switch to EMV did significantly alter attacker strategy, even if it did not significantly shift the aggregate fraud rates. Instead, as some types of fraud diminished, other techniques quickly filled the gap. For instance, Chip and PIN has caused a dramatic drop in face-to-face retail fraud, since forging a signature is easier than guessing the right PIN. In 2004, face-to-face fraud totaled £219 million; in 2006, after Chip and PIN became mandatory, it fell to £72 million (see the table in the top of Figure 7). However, card-not-present (CNP) fraud (where only the card number is used, as in online transactions) has sharply risen since the introduction of Chip and PIN, from £150 million in 2004 to £216 million in 2006, and it is projected to grow sharply thereafter. This is not surprising, since PINs are not verified for CNP transactions.

Another example of attackers moving to weaker links can be found by looking at where ATM fraud takes place. With the adoption of Chip and PIN, ATMs in the UK stopped verifying PINs using magnetic stripes, requiring the PIN be verified using the more secure chip. However, many ATMs outside the UK still allow the less secure verification via magnetic strip. As UK ATMs switched over during 2006, fraudsters immediately adapted to cashing out via less secure international ATMs. Figure 7 (bottom) shows how this transformation happened. In December 2005, when most UK ATMs still allowed magnetic stripe PIN verification, £6 million out of the total £6.5 million stolen from ATMs took place within the UK. By the end of the year, after the switch to more secure verification, the total monthly fraud increased to over £9 million, with nearly 75% of the losses

Fraud type	2004	2006
Face-to-face retail	£219m	£72m
Card-not-present (UK)	£100m	£138m
Card-not-present (Int'l)	£50m	£78m




Fig. 7. Shifting attacker strategy following EMV adoption: the table shows how card-not-present fraud has risen as face-to-face retail fraud has dropped; the bar graph shows how ATM fraud has moved overseas to ATMs which haven't made the switch to chip verification, but still rely on less secure magnetic stripes (Source: [21]).

occurring overseas. This is compelling evidence of how quickly attackers can shift to find the next-weakest link once the biggest hole has been plugged.

The EMV standard actually provides several options for implementation that vary in the cost and security achieved. While the UK EMV implementation was certainly not inexpensive, its designers did opt for many of the lower-cost implementations offered by the standard. For instance, PIN verification is normally done offline, between the card's chip and the PIN entry device (PED). This enables a relay attack whereby a small, legitimate transaction authorized by a compromised terminal can be used to approve a simultaneous, much larger fraudulent transaction elsewhere [23]. Furthermore, the communications within the PED between the card's chip and the PED's processor are not encrypted. Combined with some shoddy tamper-proofed devices, this enables PINs to be read using a paper clip tapped inside the PED [24]. While entirely practical, there is no evidence that these attacks are being used by attackers at present. Instead, they seem to be focusing on the more obvious weak links of card-not-present transactions and magnetic stripe verification at foreign ATMs. If (and




Fig. 8. Example of uncertainty for attack costs (hypothetical values)

when) these holes are plugged⁶, one might expect these attacks to be deployed, which would trigger defensive countermeasures by the banks.

4.3 Attack profiles for case studies

Figure 8 plots the expected and realized costs of different attacks for the case studies just presented. Figure 8(a) plots a sequence of top-level domains that potentially targeted by the rock-phish gang. First, many Chinese domains were registered by the gang, followed by domains from Hong Kong (as shown in Figure 5). The next top-level domain targeted – Austria – came as a surprise. Austria has an extensive Internet security community, and its law enforcement has long cooperated with cyber investigations from other countries. Nonetheless, the rock-phish gang registered many `.at` domains, and they were not removed for a very long time. It turns out that the Austrian domain registrar `nic.at` balked at removing the illicit domains without more extensive evidence documenting their illegality. The domains were only removed following a public row between the e-mail-blacklist operator Spamhaus and `nic.at` [26]. Finally, after `.at` domains began getting removed promptly, the Turkish top-level domain `.tk` was targeted.

Meanwhile, many top-level domains have not been targeted by the gang. This includes domains such as Paraguay (`.py`), which we might have expected to be

⁶ Unfortunately, card readers used to secure online banking (and potentially all card-not-present e-commerce) have recently been demonstrated to be insecure [25].

an easy target for attackers⁷, as well as domains such as `.se`, `.edu`, and `.gov` that are harder to register.

The many threats to payment card security, along with their expected and actual costs, are shown in Figure 8 (b). Face-to-face retail fraud might reasonably be seen as the weakest link in the payment card environment; the reduction in face-to-face retail fraud following the adoption of Chip and PIN supports this view. Similarly, the banks correctly anticipated that losses due to credit cards lost or stolen inside the UK (L&S in the figure) would drop once PINs were required for use. By contrast, fraud rates due to cards stolen outside the UK were lower than the banks expected. Card fraud at UK ATMs is a bit harder than retail fraud since a PIN is required. One area where the banks' expectations were not met is with ATM fraud on UK cards outside the UK. It turns out that fraudsters can easily clone stolen UK cards and use them in foreign ATMs; hence the true cost of attack is lower than expected.

The banks' losses due to card-not-present fraud were much higher than forecast; unsurprisingly, many banks have now decided to deploy readers that verify PINs in order to carry out online banking. Meanwhile, uncertainty remains regarding the attacks on PIN entry devices uncovered in [23] and [24], though it is apparent that the costs are higher than for card-not-present fraud and foreign ATM fraud.

5 Related work

The relevance of developing quantitative models of information security has been widely recognized in accounting [3, 27], information security [9, 28], and dependable computing [29–32]. Of this and subsequent literature, we briefly describe several works with similarities to our model.⁸

Gordon et al.'s “wait-and-see approach” [27] comes closest to our model. The authors explain deferred security spending in the framework of real option theory and conclude that “it is economically rational to initially invest a portion of the information security budget and defer remaining investments until security breaches actually occur.” Similar to our argument, they name uncertainty about the attack surface as main reason why they “expect organizations to use security breaches as a critical determinant of their actual [...] expenditures on information security.” In this sense, our model can be regarded as a substantial

⁷ We do not intend to single out Paraguay; many smaller countries have not been targeted by the rock-phish gang.

⁸ For the sake of brevity, we do not review in detail the extensive literature explaining security underinvestment by interactions with other market players (e.g., by under-supply in a lemon market [33], moral hazard in risk sharing contracts [34], market failure in residual risk transfer [35], defense effectiveness conditional to other players' actions [2], or general externalities [12]). A trade-off between proactive and reactive measures is also studied in the literature on optimal configuration of response in intrusion detection and prevention systems [36]. We deem this related, but not very relevant, due to the narrow scope of adjusting a technical decision threshold.

extension of the two-period case study they discuss. Gordon et al.'s work differs from this paper in that they treat security spending as an allocation problem of a fixed security budget rather than seeing it as capital investment to be put into relation to the assets at risk. Also, the weakest link attacker model is unique to our work. Both Gordon et al. and us back their claims with empirical evidence. While they draw on cross-sectional survey data collected from senior information officers, we cite time series of actual attacks in Section 4 to emphasize the iterated dimension.

Only recently, Herath and Herath [37] have transferred Gordon et al.'s [27] real option approach (ROA) to a case study of e-mail spam filtering. They compare net present values (NPV), which imply now-or-never investments, to ROA. Their results concur with our comparison between the static and dynamic solution: for their selected parameters, NPV suggests no security spending whereas ROA indicates positive returns for deferred partial security investment.

Real options are in fact another tool to model uncertainty, learning, and sequential investment, so it is just plausible (and reassuring) that models based on them come to similar conclusions. Both ROA and our iterated weakest link (IWL) have specific merits. While ROA is more general and builds on mainstream financial mathematics, IWL is more specific to information security and avoids the complexity of nested option contracts and abstract parameters, such as project volatility, which are difficult to interpret and even harder to estimate in practice. Both approaches are compatible so that a CISO might use IWL to make his decisions and express it in terms of ROA to convince his CFO.

Bier et al. [38] use a general game-theoretic setting to study strategic interaction between a single attacker and a defender who optimizes the allocation of defenses to multiple targets (two in their formal model). Unlike in our model, only a single period is modeled; the model is designed to capture terrorism threats. Uncertainty of the defender is a common element in both models, in which the defender has to cope with uncertainty about an assumed hidden preference of the attacker to target a particular target. If our notion of asset is interpreted as a distributed system in which each target potentially gives the attacker access to the entire system (similar to our case study in Sect. 4.1), then targets can be interpreted as threats and (hidden) attack costs as preferences. A topic that has no counterpart in our work is Bier et al.'s reference to signaling theory to study whether the defense configuration should be kept secret or made public.

Another single-period game theoretic model is proposed in Grossklags et al.'s [4] work on strategic security investment. Unlike our work, the authors assume a truly strategic attacker. Therefore they extend Varian's [2] security games, one of which implements the weakest link principle, by the possibility for the defender to use insurance. Although risk transfer via cyber-insurance is not considered as an investment option in our model, accepting losses due to security under-investment can be interpreted as self-insurance.

Finally, the model by Cremonini and Nizovtsev [39] explains security investment decisions with uncertainty of the attacker. They too draw on signaling theory and conclude that a defense is good enough if, and only if, it looks strong

enough to divert the attackers' attention towards other, seemingly weaker targets. In this sense, their work can be regarded as a kind of mirror image to our work, which deals with the uncertainty on the defender's side.

6 Conclusions and outlook

In this paper, we have presented an economic model that explains why and under which conditions under-investment in security can be rational, even against known threats for which defenses exist. Under-investment might reasonably occur when a) reactive investment is possible, b) uncertainty exists about the attacker's relative capability to exploit different threats, c) successful attacks are not catastrophic, and d) the sunk cost to upgrade the defense configuration is relatively small.

Unlike in other work explaining security under-investment with market failures, the ingredients to our model solely draw on the relation between defender and attacker, and do not involve actions of other market participants. Our result does not contradict or invalidate the well-known explanations by market failure. It rather complements the picture and highlights that market failure is a sufficient, but not a necessary cause for security under-investment.

Our findings suggest a need to reassess conclusions which condemn seemingly lax security practices observable in practice. While it has long been known it may be rational for some ISPs to overlook security, our model can assist policy makers in reducing negative externalities as consequences (not causes) of insecurity by better predicting situations where proactive investment is hindered. The model helps identify influential factors, notably uncertainty about attacks and sunk costs of defense, so that incentive-based countermeasures can be derived.

We believe an iterated weakest link model accurately captures the challenges facing many information security threats today. We discussed two timely examples in the paper. First, miscreants committing high-volume online crime exploit the Internet's global, distributed architecture by compromising insecure machines and defrauding uninformed registrars, wherever they are located. Second, the security of payment cards has evolved over time by reacting to fraudster's actions. Empirical data on losses shows how the introduction of Chip and PIN in the UK has simply diverted attacker strategy to other methods, notably card-not-present and foreign ATM fraud.

As with any stylized model, ideas for extensions are abundant and implementation should be guided by the principle of parsimony, especially in the absence of reliable data to validate model assumptions. Nevertheless, a handful of extensions appears reasonable to make the model more applicable. One is to let the cost of attacks decrease over time to reflect learning and the creation of automated tools. Another direction could be to allow more ways for reducing uncertainty. This can be done either by including information gathering as an option for security investment (think of penetration testing and its associated cost), or by relaxing the independence assumption between the stochastic realizations of the attack profile. The latter would allow the defender to infer

knowledge about the probability of other threats from observed attacks. In the presence of sunk costs, this opens new strategies of bundled investment in new defenses to better amortize the sunk cost. Lastly, it may be of – at least academic – interest to consider other attack strategies. For example, attackers might not choose the weakest link to confuse the defender and trigger misallocations of security spending. Anecdotal evidence for such behavior can be found in the actions of some spammers, who send waves of messages with no other apparent purpose than to wear out self-learning spam filters.

One key future challenge is to find a case study where enough information is available to calibrate the model to empirical data. More generally, it may be worthwhile to explore further the question of when to deal with a problem. When is it better to move first and take proactive measures, and when is it better to defer and respond to other’s actions?

Acknowledgements

The first author has been supported in part by the EU Network of Excellence “Future of Identity in the Information Society” (FIDIS, <http://www.fidis.net>).

References

1. Anderson, R.J., Moore, T.: The economics of information security. *Science* **314** (2006) 610–613
2. Varian, H.R.: System reliability and free riding. In Camp, L.J., Lewis, S., eds.: *Economics of Information Security*, Springer Verlag (2004) 1–15
3. Gordon, L.A., Loeb, M.P.: The economics of information security investment. *ACM Trans. on Information and System Security* **5** (2002) 438–457
4. Grossklags, J., Christin, N., John, C.: Secure or insure? A game-theoretic analysis of information security games. In: *Proc. of the 17th International World Wide Web Conference (WWW2008)*, Beijing, China (2008) 209–218
5. Collins, M., Gates, C., Kataria, G.: A model for opportunistic network exploits: The case of P2P worms. In: *Proc. of Workshop on the Economics of Information Security (WEIS)*, Cambridge, UK (2006) <http://weis2006.econinfosec.org/docs/30.pdf>.
6. Major, J.A.: Advanced techniques for modelling terrorism risk. *Journal of Risk Finance* **4** (2002) 15–24
7. Böhme, R.: Vulnerability markets – What is the economic value of a zero-day exploit? In: *22C3: Private Investigations*, Berlin, Germany (2005) http://www.inf.tu-dresden.de/~rb21/publications/Boehme2005_22C3_VulnerabilityMarkets.pdf.
8. Miller, C.: The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. *Workshop on the Economics of Information Security (WEIS)*, Carnegie Mellon University, Pittsburgh, PA (2007) weis2007.econinfosec.org/papers/29.pdf.
9. Soo Hoo, K.J.: How much is enough? A risk-management approach to computer security. In: *Workshop on Economics and Information Security (WEIS)*, Berkeley, CA (2002) <http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/>.

10. Purser, S.A.: Improving the ROI of the security management process. *Computers & Security* **23** (2004) 542–546
11. Anderson, R., Böhme, R., Clayton, R., Moore, T.: Security economics and the internal market. Study commissioned by ENISA (2008)
12. van Eeten, M.J.G., Bauer, J.M.: *Economics of Malware: Security Decisions, Incentives and Externalities*. STI Working Paper 2008/1, Information and Communication Technologies. OECD (2008)
13. Gal-Or, E., Ghose, A.: The economic incentives for sharing security information. *Information Systems Research* **16** (2005) 186–208
14. Anderson, R.J.: Why cryptosystems fail. *Communications of the ACM* **33** (1994) 32–40
15. The Economist: A walk on the dark side (30 August 2007) http://www.economist.com/displayStory.cfm?story_id=9723768.
16. Krebs, B.: Major source of online scams and spams knocked offline. *Washington Post* (2008) http://voices.washingtonpost.com/securityfix/2008/11/major_source_of_online_scams_a.html.
17. Krebs, B.: EstDomains: A sordid history and a storied CEO. *Washington Post* (2008) http://voices.washingtonpost.com/securityfix/2008/09/estdomains_a_sordid_history_an.html.
18. ICANN: Termination of registrar EstDomains to go ahead (12 November 2008) <http://www.icann.org/en/announcementsannouncement-12nov08-en.htm>.
19. Moore, T., Clayton, R.: Examining the impact of website take-down on phishing. In: *Proc. of the Anti-Phishing Working Group eCrime Researchers Summit*. (2007) 1–13 <http://www.cl.cam.ac.uk/~rnc1/ecrime07.pdf>.
20. Day, O., Palmen, B., Greenstadt, R.: Reinterpreting the disclosure debate for web infections. In Johnson, M.E., ed.: *Managing Information Risk and the Economics of Security*, New York, Springer (2008) 179–197
21. APACS: 2007 UK Chip and PIN report (2007) <http://cryptome.org/UK-Chip-PIN-07.pdf>.
22. EMVCo, L.: EMV 4.1 (2004) <http://www.emvco.com>.
23. Drimer, S., Murdoch, S.J.: Keep your enemies close: Distance bounding against smartcard relay attacks. In: *USENIX Security Symposium*. (2007)
24. Drimer, S., Murdoch, S.J., Anderson, R.: Thinking inside the box: System-level failures of tamper proofing. In: *IEEE Symposium on Security & Privacy*. (2008) 281–295
25. Drimer, S., Murdoch, S.J., Anderson, R.: Optimised to fail: Card readers for online banking. In: *13th International Conference on Financial Cryptography and Data Security (Financial Crypto)*. (2009)
26. Spamhaus: Report on the criminal ‘rock phish’ domains registered at nic.at (press release) (2007) <http://www.spamhaus.org/organization/statement.lasso?ref=7>.
27. Gordon, L.A., Loeb, M.P., Lucyshyn, W.: Information security expenditures and real options: A wait-and-see approach. *Computer Security Journal* **14** (2003) 1–7
28. Schechter, S.E.: Toward econometric models of the security risk from remote attacks. *IEEE Security and Privacy* **3** (2005) 40–44
29. Littlewood et al., B.: Towards operational measures of computer security. *Journal of Computer Security* **2** (1994) 211–229
30. Jonsson, E., Andersson, M.: On the quantitative assessment of behavioural security. In: *Information Security and Privacy (Proc. of ACISP’96)*. Volume LNCS 1172., Berlin, Springer-Verlag (1996) 228–241

31. Jonsson, E., Olovsson, T.: A quantitative model of the security intrusion process based on attacker behaviour. *IEEE Trans. on Software Engineering* **23** (1997) 235–245
32. Avižienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. on Dependable and Secure Computing* **1** (2004) 11–33
33. Anderson, R.J.: Why information security is hard – An economic perspective (2001) <http://www.cl.cam.ac.uk/~rja14/econsec.html>.
34. Kunreuther, H., Heal, G.: Interdependent security. *Journal of Risk and Uncertainty* **26** (2003) 231–249
35. Böhme, R., Kataria, G.: Models and measures for correlation in cyber-insurance. *Workshop on the Economics of Information Security (WEIS)*, Robinson College, University of Cambridge, UK (2006) <http://weis2006.econinfosec.org/docs/16.pdf>.
36. Yue, W.T., Çakanyildirim, M.: Intrusion prevention in information systems: Reactive and proactive responses. *Journal of Management Information Systems* **24** (2007) 329–353
37. Herath, H.S.B., Herath, T.C.: Investments in information security: A real options perspective with Bayesian postaudit. *Journal of Management Information Systems* **25** (2008) 337–375
38. Bier, V., Oliveros, S., Samuelson, L.: Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory* **9** (2007) 563–587
39. Cremonini, M., Nizovtsev, D.: Understanding and influencing attackers' decisions: Implications for security investment strategies. *Workshop on the Economics of Information Security (WEIS)*, Robinson College, University of Cambridge, UK (2006) <http://weis2006.econinfosec.org/docs/3.pdf>.