

# The Iterated Weakest Link

**S**ecurity breaches are in the news almost daily, each bigger and more costly than the last. Does this reflect flawed technology, policy, or simply ineptitude? Or, what if 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. Here, we describe a model that weighs the merits of such an approach.

## The Nature of Attacks

One key insight from information security economics literature<sup>1</sup> 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.<sup>2</sup> The most careless programmer in a software firm can introduce a critical vulnerability. Attackers have repeatedly exhibited a knack for identifying ways to bypass a system's security, even when the designer remains unaware of the particular weakness.

However, systems don't exist in a vacuum; rather, defenders respond to attacks by plugging known holes. And yet, as soon as defenders fix one flaw, attackers often identify and exploit another. So, a strong dynamic component is at play: attackers find the weakest link, defenders fix the problem, attackers find new holes that are then plugged, and so on, in a pattern that emerges repeatedly.

For instance, attackers construct networks of compromised machines (so-called botnets) to pester legitimate users by emitting spam, distributing malware, and hosting phishing Web sites. Attackers concentrate their efforts at the most irresponsible ISPs, moving on to others only after the ISP cleans up its act or shuts down.<sup>3,4</sup> Likewise, technical countermeasures to payment-card fraud have evolved over time, causing fraudsters to adopt new strategies as banks fix old weaknesses. When UK banks migrated from signatures to PIN verification for transactions, in-person retail fraud declined, whereas overseas ATM fraud and card-not-present fraud skyrocketed (see <http://cryptome.org/UK-Chip-PIN-07.pdf>).

## Defenders' Choice (under Uncertainty)

So, how can we grasp and model this dynamic interaction between attackers and defenders? Simply stated, a defender protects an asset of value against  $n$  possible threats. He or she can ward off each threat by investing in its corresponding defense. We represent the costs for each defensive countermeasure via an upper triangular matrix and model those costs as *independent*

(off-diagonal values zero), *complementary* (off-diagonal values negative), or *conflicting* (off-diagonal values positive). One nice property of arranging the cost matrix in this manner is that for positive off-diagonal elements, decreasing marginal utility of defenses becomes endogenous instead of appearing as an assumption, as in the Gordon-Loeb framework.<sup>5</sup>

Whether interdependent or not, we assume the defender's costs of implementing countermeasures are known. This is reasonable—security countermeasures such as firewalls and intrusion-detection systems come with a bill. By contrast, it's much harder for defenders to accurately predict the cost of different attacks in advance. Although they might possess some intuition about the relative difficulty of carrying out the  $n$  threats, such knowledge could very well be blurred.

To model this uncertainty, we order the threats 1, ...,  $n$  by increasing the attack's *expected cost*. By varying the level of uncertainty  $\sigma$  associated with different attacks' true costs, we can learn a great deal about why security investment often falls short of what technical experts desire. Figure 1a illustrates the role of uncertainty when ordering threats. Under uncertainty, expected and realized costs differ so that threat 4, not threat 3 as expected, is the weakest link if defenses 1 and 2 are in place.

To connect the model to a concrete example, consider the many threats to payment-card security Figure 1b illustrates. We might reasonably view face-to-face retail

**RAINER BÖHME**  
International  
Computer  
Science  
Institute,  
Berkeley

**TYLER MOORE**  
Harvard  
University

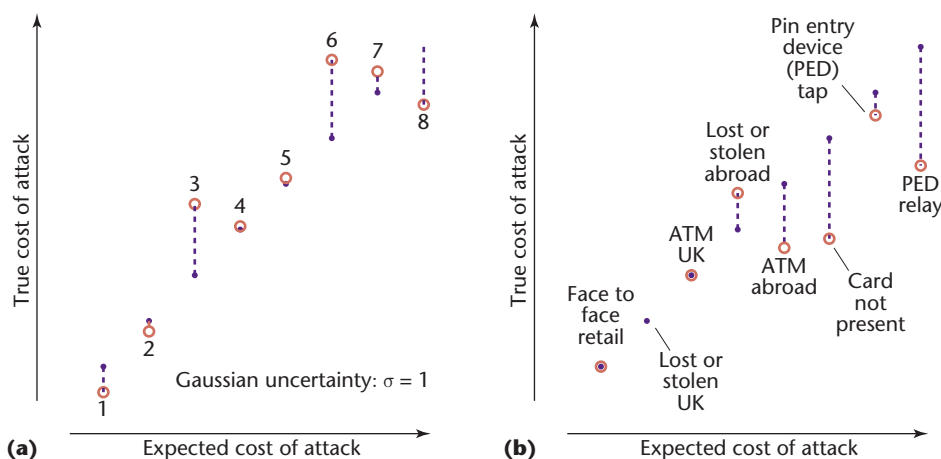


Figure 1. Uncertainty for attack costs (hypothetical values). (a) To model this uncertainty, we order the threats 1, ...,  $n$  by increasing the attack’s expected cost. By varying the level of uncertainty  $\sigma$  associated with different attacks’ true costs, we can learn a great deal about why security investment can fall short of what technical experts desire. Here, the Gaussian uncertainty of  $\sigma = 1$ . (b) We connected the model to a concrete example by looking at payment-card attack costs.

fraud (F2F) as the weakest link in the payment-card environment; its reduction following chip and PIN adoption supports this view. Similarly, the banks correctly anticipated that losses due to lost or stolen (L&S) credit cards inside the UK would drop once PINs were required for use. However, what the banks didn’t foresee was how swiftly the fraudsters could shift their tactics. Once UK payment processing was locked down by chip-based verification at retail stores and ATMs, fraud shifted to areas where protections weren’t in place. UK ATM cards continued to be cloned, but were now being cashed out at foreign ATMs lacking advanced protections. Likewise, the banks’ losses due to card-not-present fraud (CNP) rose much higher than forecast; unsurprisingly, many banks have subsequently begun deploying readers that verify PINs even for online transactions.

### Proactive vs. Reactive Security Investment

Our model is “run” in an iterated game: in each round, the defender decides which, if any, of the  $n$

threats to protect against. The attacker identifies and exploits the weakest link—that is, the threat least costly to him- or herself. Unlike the defender, the attacker is certain of each attack’s cost and doesn’t operate indiscriminately; rather, he or she attacks only when doing so is profitable.

We reach several interesting conclusions on examining the model. In the static case, in which defenders get only one chance to protect a system, increasing uncertainty about which link is weakest causes them to protect more assets, but only up to a point. When uncertainty is too high, defenders don’t know which asset to protect and so choose to protect none. If, instead, we allow for repeated defensive investments in the dynamic case, an uncertain defender will initially protect fewer assets and wait for the attacker to “identify” the weakest links, which the defender can fix in later rounds. So, it can be quite rational to under-invest in security until threats are realized.

Of course, security countermeasures can require significant capital investment from the out-

set. When we introduce sunk costs into our model,<sup>6</sup> we find that for moderate levels of uncertainty, such costs raise the proactive protection investment adopted in the dynamic case. We’ve translated our findings about optimal defensive strategies into accepted security indicators, such as return on security investment (ROSI), as Table 1 shows. For moderate levels of uncertainty ( $\sigma = 1$ ), moving from a static to dynamic defense strategy reduces security spending, leading to more observed attacks. However, gross returns increase, too. In fact, security spending is better targeted, overinvestment is reduced, and the security investment’s overall efficiency, as measured by the ROSI indicator, improves. Hence, we can draw an alternative interpretation to the omnipresent reports of security breaches in the media: rather than rashly framing them as engineering failures, we might also view breaches as unavoidable side-effects of smart defense strategies that balance the appropriate levels of proactive and reactive security investment.

Investment in countermeasures and, consequently, attack frequency depend fundamentally on the opportunity to defend reactively. When firms must do all security investment proactively, they might simply raise the white flag of surrender if they’re very uncertain ( $\sigma \geq 4$ ) about which threats are likely. Only a staged approach gives these investors an incentive to defend against the most aggressive threats. Given the chance to invest in later rounds, firms choose to protect the assets revealed to be weak, leading to a higher ROSI and reduced attack intensity. (Note that our model identifies the rational response to the private costs defenders face but ignores the public costs insecurity creates. So, whereas it might be narrowly better for some defenders to skimp on security initially, a public-policy

response could nonetheless be necessary to compensate for the negative externalities of insecurity such under-investment causes.)

Our proposed economic model explains why and under which conditions security under-investment can be rational, even against known threats for which defenses exist. Unlike other work that explains such under-investment with externalities or misaligned incentives, our model draws solely on uncertainty about where to invest in countermeasures. This result doesn't contradict or invalidate well-known explanations of market failure.<sup>1</sup> Rather, it complements the picture and highlights that market failure is a sufficient but not necessary cause for security under-investment. Real option frameworks that suggest a "wait-and-see" approach<sup>7</sup> also present the logic of initial security under-investment followed by reactive investment.

To conclude, we believe an iterated weakest-link model accurately captures the challenges of many information security threats today. Our findings suggest a need to reassess conclusions that condemn seemingly lax security practices found in the media. Our model can assist policy makers in reducing negative externalities as consequences (not causes) of insecurity by better predicting situations that hinder proactive investment. The model also helps identify influential factors—notably, uncertainty about attacks—so that firms and managers can derive incentive-based countermeasures. □

**Acknowledgments**

This essay was selected as winner of the inaugural Gordon Prize in Managing Cybersecurity Resources ([www.rhsmith.umd.edu/news/releases/2009/101409.aspx](http://www.rhsmith.umd.edu/news/releases/2009/101409.aspx)). For the full academic paper outlining the model

**Table 1. Security investment indicators.\***

Indicator	Level of uncertainty			
	$\sigma = 0$	$\sigma = 1$	$\sigma = 4$	$\sigma = 8$
<b>Static defense</b>				
Optimal number of defenses	11	12	0	0
Attack intensity (% rounds)	0.0	2.4	100.0	100.0
Return on security investment (ROSI; % security spending)	51.5	31.2	—	—
<b>Dynamic defense</b>				
Optimal number of proactive defenses	11	9	7	3
Attack intensity (% rounds)	0.0	6.1	15.7	32.7
ROSI (% security spending)	51.5	52.8	35.2	18.9

\*Asset value = US\$1 million, return on asset = 5%, loss given attack = \$25 000,  $n = 25$ , minimum expected cost of attack = \$15,000, gradient of attack cost = \$1,000.

presented here, see R. Böhme and T. Moore, "The Iterated Weakest Link: A Model of Adaptive Security Investment," *Workshop on the Economics of Information Security (WEIS)*, 2009; <http://weis09.infoseccon.net/files/152/paper152.pdf>.

**References**


1. R.J. Anderson and T. Moore, "The Economics of Information Security," *Science*, vol. 314, no. 5799, 2006, pp. 610–613.
2. H.R. Varian, "System Reliability and Free Riding," *Economics of Information Security*, L.J. Camp and S. Lewis, eds., Springer-Verlag, 2004, pp. 1–15.
3. T. Moore and R. Clayton, "Examining the Impact of Website Take-Down on Phishing," *Proc. Anti-Phishing Working Group eCrime Researchers Summit*, ACM Press, 2007, pp. 1–13; [www.cl.cam.ac.uk/~rnc1/ecrime07.pdf](http://www.cl.cam.ac.uk/~rnc1/ecrime07.pdf).
4. O. Day, B. Palmén, and R. Greenstadt, "Reinterpreting the Disclosure Debate for Web Infections," *Managing Information Risk and the Economics of Security*, M.E. Johnson, ed., Springer, 2008, pp. 179–197.
5. L.A. Gordon and M.P. Loeb, "The Economics of Information Security Investment," *ACM Trans. Information and System Security*, vol. 5, no. 4, 2002, pp. 438–457.
6. R. Böhme and T. Moore, "The

Iterated Weakest Link: A Model of Adaptive Security Investment," *Workshop on the Economics of Information Security (WEIS)*, 2009; <http://weis09.infoseccon.net/files/152/paper152.pdf>.

7. L.A. Gordon, M.P. Loeb, and W. Lucyshyn, "Information Security Expenditures and Real Options: A Wait-and-See Approach," *Computer Security J.*, vol. 14, no. 2, 2003, pp. 1–7.

*Rainer Böhme is a postdoctoral fellow at the International Computer Science Institute in Berkeley, CA. His research interests include economic and behavioral aspects of security and privacy, steganography and steganalysis, as well as multimedia forensics. Böhme has a PhD in computer science from Technische Universität Dresden in Germany. Contact him at [rainer.boehme@icsi.berkeley.edu](mailto:rainer.boehme@icsi.berkeley.edu).*

*Tyler Moore is a postdoctoral fellow at Harvard University's Center for Research on Computation and Society. His research interests include economics of information security, the study of electronic crime, and the development of policy for strengthening security. Moore has a PhD in computer science from the University of Cambridge. Contact him at [tmoore@seas.harvard.edu](mailto:tmoore@seas.harvard.edu).*

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.