Security Metrics and Security Investment Models

Rainer Böhme

International Computer Science Institute, Berkeley, California, USA
rainer.boehme@icsi.berkeley.edu

Abstract. Planning information security investment is somewhere between art and science. This paper reviews and compares existing scientific approaches and discusses the relation between security investment models and security metrics. To structure the exposition, the high-level security production function is decomposed into two steps: cost of security is mapped to a security level, which is then mapped to benefits. This allows to structure data sources and metrics, to rethink the notion of security productivity, and to distinguish sources of indeterminacy as measurement error and attacker behavior. It is further argued that recently proposed investment models, which try to capture more features specific to information security, should be used for all strategic security investment decisions beneath defining the overall security budget.

1 Introduction

According to recent estimates, global enterprises spent about US$ 13 billion on information security in 2009, and this figure is projected to grow by 14% in 2010 [1]. This amount is substantial even when broken down to the individual enterprise level. For instance, one in three surveyed firms in the US spends 5% or more of the total IT budget on information security [2]. In Japan, one in five firms spent 10% or more in 2007. However, the fraction of firms investing in security so intensively came down from one in three firms in 2006 [3]. This is not overly surprising as money allocated to security is not available for other purposes. So the key question in management of information security is if this money is being spent well. This question has attracted the attention of researchers from accounting, business, economics, computer science, and related disciplines.

This paper attempts to survey and systemize the literature, thereby extracting more mature facts as insights for practitioners and distinguishing them from untested hypotheses and open research questions for academic researchers interested in the field. In Section 3 we decompose the security investment process and discuss all key variables. Section 2 focuses on data sources and metrics for these variables. Section 4 gives an overview of recent directions in research deviating from the standard approach towards more domain-specific or empirically founded models. The paper concludes with a brief outlook.

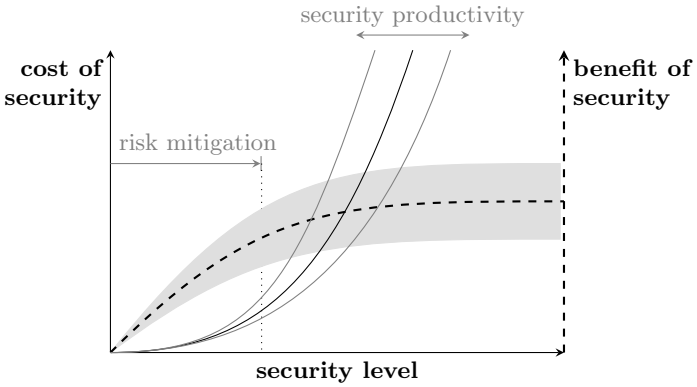

Fig. 1. Decomposition of the security production function into two steps

2 What to Measure

The key quantity in investment theory is the ratio of cost to benefit, or in terms of a production function, the amount of *output* per unit of *input*. The purpose of a *security investment model* is to describe this relation formally for the domain of information security. Every security investment model builds on *security metrics* which define the model's inputs, outputs, and parameters. If values are obtained from actual measurements, the model can predict whatever unknown variable it is solved for.

Undoubtedly the most famous security investment model has been proposed by Gordon and Loeb [4]. Standing in the tradition of the accounting literature, this model defines a *security breach probability function*, which maps the monetary value of security investment to a probability of incurring a defined loss. Under the assumption of a class of security breach probability functions, the authors derive a rule of thumb for setting the security investment as a fraction of the expected loss without security investment.[1] Several extensions of the Gordon–Loeb model criticize this conjecture [5], derive optimal investment rules for alternative forms of the security breach probability function [6], endogenize the probability of attack [7], or include timing decisions [8]. All variants have in common that security investment exhibits decreasing marginal returns: every additional dollar spent yields proportionally less benefit. This assumption can be justified intuitively [9] and it is also supported empirically on cross-sectional firm data [10].

From a measurement point of view, the high degree of abstraction of the Gordon–Loeb model and its variants can sometimes be inconvenient. This is so

[1] The precise conjecture states that for decreasing marginal returns, an upper bound for security investment is given by $1/e$ (or roughly 37%) of the expected loss without security investment [4].

because the *direct* mapping of inputs (monetary amounts of security investment) to outputs (probability of loss) neglects intermediate factors, notably the *security level*. In practice, intermediate factors are oftentimes better observable than the abstract parameters of the Gordon–Loeb model.

Therefore we use an alternative structure for our discussion of variables of interest. As depicted in Fig. 1, we decompose the *security production function* into two parts. First, the *cost of security* (in monetary terms) is mapped to the *security level* (solid lines in Fig. 1). Second, the *security level* stochastically determines the *benefits of security* (dashed lines and shaded area in Fig. 1). Indeterminacy is introduced to model attacker behavior. In the following we discuss each variable of interest and explain why this decomposition is useful.

2.1 Cost of Security

Cost of security seems to be the variable easiest to measure by summing up the expenses for the acquisition, deployment, and maintenance of security technology. Yet this reflects only the direct cost. Some security measures have non-negligible indirect cost, such as time lost due to forgotten credentials, the inconvenience of transferring data between security zones, or incompatibilities of security mechanisms slowing down essential processes. If security measures foster privacy or secrecy by enforcing confidentiality, some business decisions might have to be taken less informed and reach suboptimal outcomes compared to the fully informed case. This opportunity cost adds to the indirect cost of security.

It is sometimes useful to express the cost of security as a function of the economic activity in the core business: *fixed* costs are independent of the activity in the core business whereas *variable* costs grow proportionality to the activity. It is often sufficient to assume fixed cost of security. However, the cost of distributing security tokens to customers or indirect costs due to delayed business processes are clearly variable and should be modeled as such.

If the security investment model has a time horizon of multiple periods, one can distinguish the cost of security further by *onetime* and *recurring* (i.e., per-period) costs. While the acquisition and deployment of protection measures is naturally modeled as onetime cost, their maintenance and most indirect costs are recurring. In certain situations it is useful to consider *sunk* costs, which cannot be recovered when decommissioning protection measures [9]. Most security equipment (e.g., firewall devices) can be sold (at a discount) or repurposed (e.g., as routers), and staff transferred or fired [4]. But the expenses for training or for the distribution of security tokens to customers are irreversibly spent.

Whenever costs are distributed over several periods, effects of time-dependent discounting and non-linearities due to taxation can be considered [11]. This is common practice in general investment theory, but barely reflected in the specific literature on security investment so far. Given the pace of development and the short-term nature of most security investments, the errors introduced by ignoring these factors seem small compared to other sources of uncertainty and do not justify complicating the models excessively.

Whatever breakdown is used to account the cost of security, this variable should be considered as *deterministic* up to measurement noise. That is, a true value exists in theory, although it might not always be easy to measure it exactly.

2.2 Security Level

The security level is the variable in the model that summarizes the quality of protection. Like cost of security, it can be assumed to be embodied in a deterministic state, even though it is even more difficult to measure. The reason is that the quality of protection is not necessarily a scalar, but some discrete state which has to be mapped to (at least) an ordinal scale. Deterministic indicators include patch level, existence and configuration of intrusion detection systems, whether virus scanners are installed on end-user PCs, etc. [12]. Despite being often crude and noisy, these indicators convey some indication about the actual security level. This way, the various process models to evaluate security in organizations qualitatively (e.g., [13, 14]) can be connected with quantitative security investment models.

In addition, the security level can often we observed through stochastic indicators where—again—the indeterminacy reflects attacker behavior. Examples for this category are typical incident measures of intrusion detection systems and virus scanners, such as the actual false alarm and missed detection rates.

Observe that our decomposition of the security production function is useful if indicators of the security level are (partly) observable. Since in particular variables on the benefit side of security investment models are difficult to measure and error-prone, it can be of great help to include a supporting point by quantify the security level. This way, the first and second step of the security production function can be evaluated independently, checked for plausibility, and benchmarked against industry best practices.

A related remark concerns the notion of *security productivity*. While it is defined for both steps jointly in the Gordon–Loeb framework [4, 7]—in the absence of alternatives—we prefer to tie productivity more closely to the efficiency of the security technology and its ability to *mitigate* risk (as opposed to risk avoidance, transfer, and retention). As annotated in Fig. 1, security productivity is determined by the curvature of the function that maps the cost of security to the security level. It reflects the increase in security level per unit of security spending, possibly taking into account decreasing marginal returns.[2] Since the second function on the benefit side is much more specific to the individual organization (e.g., due to differences in the assets at risk), our definition of security productivity has advantages when comparing the efficiency of security spending between organizations.

[2] Intuitively, we expect that this characteristic applies to both mapping functions as depicted in Fig. 1. But this is not essential as long as the total effect prevails. There always exists a transformation of the security level so that only one function models the total effect of decreasing marginal returns.

2.3 Benefit of Security

The second step in the security production function involves the difficulty of mapping incidents to losses. More precisely, the security level is mapped to *prevented* incidents, which then can be translated to a benefit of security.[3]

Matsuura notes that fewer incidents can either be due to more attacks failing or due to fewer attacks. Most protection technology affects the first factor, but differences in security productivity could be used to balance investment along this dimension [7]. This is particularly relevant if the second factor (fewer attacks) is not specific to the organization, but affects others too (cf. Sect. 4.5).

As mentioned above, the benefit of security largely depends on the value of the assets at risk. This opens up the can of worms of valuating intangible information assets. For the sake of brevity, we spare a survey of this topic. Assume for now that the value of all assets affected by an incident is known. Then we can distinguish situations in which this value imposes an upper bound on the losses from situations where the losses can exceed the asset value. Examples for the latter include cases of liability or secondary costs to recover the asset [15]. We use the broader term of *recovery cost* to subsume all direct and indirect costs associated with a loss event.

By its very nature, losses and hence recovery costs are random variables that take positive values and oftentimes concentrate probability mass at zero (the case when no incident happens). These random variables can be summarized in scalars (e.g., by their moments), however not without losing information. We follow the convention in decision theory and express the expected benefits after a transformation by a *utility function*, which takes the *risk aversion* of the decision maker as parameter. If organizations are assumed to be risk neutral (this is justifiable for businesses), the utility function is the identity function.

It is needless to say that the random nature of losses complicates not only the ex-ante perspective of security investment ("What measures should we implement?"), but also ex-post evaluations ("Did we implement the right measures?") [16]. What appears right or wrong in one state of the world (i.e., realization of the random attack variable) is not necessarily the same in other states. This way or the other, a security manager's standing within an organization will always depend on a combination of skill and luck.

At least for the ex-ante perspective, very recent research points out that fuzzy logic might be the tool to deal with the large degree of uncertainty in security decision-making [17, 18]. However, it is too early to tell if these concepts are implementable in practice and whether they provide the right kind of signals that can be interpreted on a technical and managerial level alike.

3 How to Measure

With the three variable of interest defined, there remain open questions how to measure or estimate their values (Sect. 3.1) and how to calculate meaningful decision criteria (Sect. 3.2) for a specific investment decisions.

[3] Benefit is expressed in the same monetary unit as cost to calculate ratios.

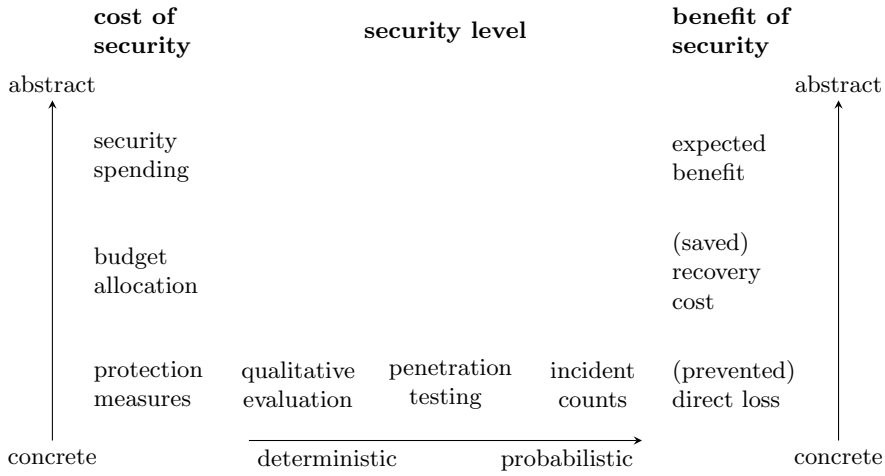

Fig. 2. Security investment indicators structured by level of abstraction; arrowheads point towards increasing difficulty in measurement

3.1 Data Sources

Data sources can broadly be divided into *internal* sources of the investing organization and *external* sources.

Figure 2 shows various security investment indicators from internal sources and their associated variable in the investment model. The indicators corresponding to cost and benefit of security are vertically ordered by their level of abstraction. Technical indicators of the security level, by their very nature, are concrete and specific to the technology in use [12]. Since the transition from in principle deterministic states to probabilistic quantities takes place at this level, it is convenient to organize these indicators along this dimension horizontally.

On the cost side, security spending means the total amount of the security budget of an organization. It is the indicator of interest to set the budget ("How much is enough?" [19]). For a given budget, the next more concrete level is to decide the security strategy ("Where to invest?"). This involves the allocation of budget to the typical alternatives in risk management (mitigation, avoidance, transfer, retention) and to types of security investment (proactive versus reactive, technical versus organizational, etc.). Even more concrete is the cost of individual protection measures. For many measures, this cost is easily observable (e.g., by the price tag). Measuring security costs on more abstract levels becomes increasingly difficult, as indirect costs emerging from certain measures and from the interaction between measures [9] have to be taken into account.

The hierarchy on the benefit side is roughly symmetric to the cost side. The only difference is that saved recovery cost and prevented direct loss are random

variables (or realizations in the ex-post perspective, if observable at all), whereas the expected benefits reflect an annualized[4] and risk-adjusted monetary value.

External data sources include *threat level indicators*, such as the number of active phishing sites, malware variants in circulation, breach disclosure figures, or the number of vulnerability alerts awaiting patches [20]. More and more of such indicators are collected and published on a regular basis by the security industry—mind potential biases [21]—, research teams, not-for-profit organizations, and official authorities. These indicators alone are certainly too unspecific for most organizations, but they can be helpful to update quantitative risk assessment models regularly and to adjust defenses tactically even if data from internal sources is only available at lower frequency or higher latency. By contrast, *market-based indicators* derived from price information in vulnerability markets have been proposed as alternatives to threat level indicators for their potential of being forward-looking [22, 16]. In prior work, we have identified bug challenges, bug auctions, exploit derivatives, and premiums charged by cyber-insurers as potential data sources. However, the most dominant type of vulnerability market in practice are vulnerability brokers, which emit the least signals to construct telling indicators [23].

3.2 Choice of Metrics

The main purpose of metrics is to compare between alternatives. While comparisons over time or across organizational units can be carried out with concrete technical indicators of the security level, comparisons between protection measures or budget allocation options require the underlying metrics to be on the same scale. This explains why the most regarded metrics in security investment are calculated as cost–benefit ratios on a higher level of abstraction.

Over the past decade, substantial work has been done in adapting principles and metrics of investment theory for security investment [19, 15, 16]. Most prominent is the notion of a *return on (security) investment* (ROSI/ROI). Among a handful of variants, we prefer the one normalized by the cost of security [24, 9],

$$\text{ROSI} = \frac{\text{benefit of security} - \text{cost of security}}{\text{cost of security}}. \tag{1}$$

Higher values of ROSI denote more efficient security investment. Note that the notion of *return* in ROSI is broad, as prevented losses do not constitute returns in a narrow sense.

Terminology feud aside, these metrics are also regarded with skepticism by practitioners who are familiar with the problems of statistical data collection for rare events. They see a main problem in obtaining annualized and risk-adjusted security benefit figures [12, 25]. Nevertheless, these metrics seem to remain as necessary compromise to justify security expenses within organizations.[5] It is

[4] or aggregated for any other fixed time horizon

[5] Another incontestable application of ROSI are result presentations for analytical models, e.g., [9].

common practice to make (or justify) budget decisions based on standard investment theory because it facilitates comparisons between investments in various domains. This has so often been noted that the largest annual survey among corporate information security managers in the US includes a specific question [2, Fig. 7]. According to that, ROSI is used by 44% of the responding organizations. The *net present value* (NPV) and the *internal rate of return*—two other standard investment indicators which allow for discounting, but share the same caveats—follow with 26% and 23%, respectively. Apparently security managers have little choice than adopting the terminology of corporate finance.

4 Recent Research Directions

Independent of the adoption of security metrics and investment models in practice, academia contributes to the formation and development of a security investment theory. This theory gets increasingly detached from its roots in accounting. Recent security investment models have been enriched with domain knowledge reflecting specific technical or environmental factors. While in the early days, security investment models were motivated with setting a security budget, newer models are devised to help setting a security strategy. The question has changed from "How much is enough?" [19] to "Where to invest?". In the following we will briefly review interesting recent developments.

4.1 Timing

Security investment inherently involves decision-making under uncertainty: will this threat realize or not? This uncertainty is reduced over time as incidents can be observed. An elegant way to model this is offered by real options theory, a branch of financial investment theory which accounts for deferred investment (unlike, for instance, the NPV metric). Gordon, Loeb and Lucyshyn [26] first adapted this line of thought to information security and proclaimed a "wait-and-see" tactic. Instead of over-investing into defenses that will never become relevant, it can be rational to wait until the first (non-catastrophic) incident happens, and then react. Herath and Herath [27] follow up and provide a comparison between ROSI-based security investment and the real options approach. Tatsumi and Goto [8] extend the Gordon–Loeb model [4] by a timing dimension.

Balancing proactive versus reactive security investment is also studied by Yue and Çakanyildirim [28] for the specific case of configuring an intrusion detection system (IDS), as well as in our "iterated weakest link" model [9]. This model combines several features specific to security investment—such as an attacker seeking to exploit the weakest link—in a repeated player-versus-nature game involving multiple threats over multiple rounds (unlike most real option models, which consider only two stages). The core idea is that the defender has some knowledge about the expected difficulty of pursuing several attack vectors, but remains uncertain about the true order. Accepting that some attacks may be successful enables more targeted security investment and thus reaches overall

better outcomes than blind over-investment. Thus in many cases, ROSI increases even after accounting for the losses of successful attacks.

4.2 Information Gathering

There are other ways to reduce the uncertainty in making security decisions than waiting for attacks. Sharing information with other defenders promises several benefits:[6]

1. *Early warning.* New attacks might not always hit all organizations at once. So the ones spared at the beginning do not need to wait until they get attacked, but can learn from their peers and upgrade just-in-time. On a technical level, this can be done by sharing IDS and anti-virus signatures.
2. *Noise reduction through aggregation.* Some types of incidents occur too rarely to estimate reliable probabilities of occurrence from internal observations only. By aggregating observations over many sites, even small probabilities can be determined more accurately.
3. *Forensic discovery of structure.* The nature of certain malicious activity online remains obscure to observers who see only a small fraction of the network. Sharing knowledge may give a 'bigger picture' and enable forensic investigations to find better defenses or prosecute perpetrators.

Gordon, Loeb and Lucyshyn [30] as well as Gal-Or and Ghose [31] proposed models to determine the optimal amount of information sharing between organizations. In their game-theoretic framework, security investment and information sharing turn out to be strategic complements.

Another way to gather information is to analyze precursors of attacks from internal sources via intrusion detection [32, 33] and prevention systems [28]. Since the deployment and maintenance of such systems constitutes an investment, it is quite natural to refine investment models to include this feature. A related feature are professionals services to test the resilience against attacks by exposing it to the latest attack techniques. Commissioning these so-called penetration tests can be seen as an investment in information acquisition. Hence it has its place in security investment models [34].

Note that the ROSI metric cannot be calculated separately for information gathering tasks because the acquired information can make planned security investments obsolete. These savings sometimes exceed the cost of information gathering, thus leading to a negative denominator in Eq. (1). As a rule of thumb, ROSI is a metric for the joint efficiency of the entire security investment strategy.

4.3 Information Security Outsourcing

Once the security budget is defined, it is rational to consider security as a service that is subject to a make-or-buy decision similar to most other operations,

[6] We list obvious benefits only. See for example [16, Table 1] for risks and [29] for ambivalent consequences of signaling information about the security level.

though with specific risks and benefits [35]. Outsourcing in general is best approached as a principal–agent problem where the provider is susceptible to moral hazard [36]. Ding et al. adapted this theory to the special case of information security outsourcing and mention the providers' long-term interest in a good reputation as limiting factor to moral hazard [37]. In a related analysis, the same team includes transaction costs in the investment model and warns that the decision to outsource security functions may bear hidden costs if principals find themselves locked into a relationship with their providers [38]. By contrast, Rowe [39] points to positive externalities of security outsourcing if multiple organizations share the same provider. These externalities arise from both economies of scale and improved information sharing. This is not only beneficial for the involved organizations but—depending on the model—also for others.

Schneier specifies that it is important to differentiate between outsourced functions [40]: penetration and vulnerability testing (see Sect. 4.2), security auditing, system monitoring, general system management, forensics, and consulting all involve different risks and incentive structures. This is partly reflected in the security investment model by Ceszar, Cavusoglu and Raghunathan [41], who analyze under which conditions it is optimal to outsource system management and system monitoring to a single or multiple independent providers.

4.4 Cyber-Risk Transfer

Aside from risk mitigation and risk avoidance, the financial risk of security incidents can be transferred to third parties, notably cyber-insurers. If the premium is lower than the difference between benefit and cost of security, this is a viable investment option. Note that if the insurance market is in equilibrium, this is only true if organizations are either risk averse or better informed about their specific risk than the insurer. However, the market for cyber-insurance seems underdeveloped in practice, presumably due to three obstacles characterizing cyber-risk: interdependent security, correlated risk, and information asymmetries [42].

If this situation changes in the future, insurers will most likely require that protection measures against all known threats are in place. Therefore cyber-insurance shall rather be seen as complementary to investing in protection measures or outsourced security operations, not as a substitute. To ensure that a defined security level is maintained, insurers might collaborate with security service providers and advise their clients to outsource security operations to them (see Fig. 3). Zhao, Xue and Whinston [43] study such a scenario and conclude that outsourcing (which they see as a substitute to cyber-insurance) is preferable to cyber-insurance. However, in this model security service providers assume full liability for potential losses. We are not aware of a single provider who offers this in practice. So effectively, this result should be interpreted as a combination of security outsourcing and cyber-insurance. Such a combination in fact promises better outcomes than cyber-insurance alone [42].

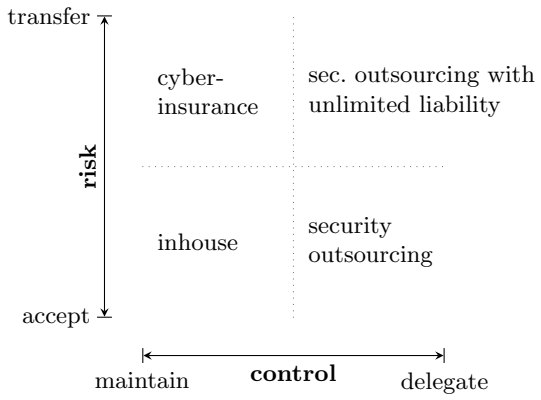

Fig. 3. Orthogonal relation of cyber-risk transfer and outsourcing of security operations

4.5 Private Versus Public Benefit

So far, this paper has taken the dominant perspective in investment theory: organizations seeking to maximize their private profit. A separate stream of related work has studied security investment as a problem of provisioning a public good. Varian [44] adapted Hirshleifer's [45] theory of public goods with different aggregation functions to the domain of information security. In independent work, Kunreuther and Heal [46] study security investment when it generates *positive externalities*, i.e., an organization's expected loss decreases not only with increasing own security level but also with increasing (aggregate) security level of other organizations connected in a network. Grossklags et al. [47] extend this work by distinguishing between two types of protection measures, one which generates positive externalities and one which does not. They describe the existence of equilibria in a game-theoretic setting as a function of the cost of both types of security investment. Cremonini and Nizovtsev [48] modify the setting by considering the case when security investment generates *negative externalities*. In general, if security investment creates positive externalities, profit-maximizing security investors try to free-ride and under-invest. The opposite is true if security investment creates negative externalities.

4.6 Empirical Underpinning

The academic literature on security investment suffers from a deficit in empirical validation with cross-sectional or longitudinal data[7], which can be explained by the difficulty of obtaining such data. The most regarded annual survey among US enterprises includes a number of relevant indicators, but its data quality if often criticized for ambiguous category definitions and low response rates indicating

[7] References to several case studies of single organizations can be found e. g. in [49].

potential coverage error [2]. Moreover, its results are not public since the 2008 edition, and the responses are not available in a disaggregated form.

The situations is better in Japan, where METI[8] data is available on a micro level. This data has been used to validate models of the Gordon–Loeb type [10]. Liu, Tanaka and Matsuura [49] also report evidence for the decomposed form of security investment models as advocated in this paper. They observe a broad indicator of security investment—including protection technology, organizational measures, and employee awareness raising—over several periods and find that consistency in security investment is a significant predictor for fewer incidents.

Eurostat has collected some indicators related to security in its annual ICT surveys of households and enterprises in Europe. However, the data is very fragmented and the indicators are not focussed on security investment [21]. A special survey module tailored to security is being administered in 2010. We are not aware of any literature testing security investment models with Eurostat data.

In [9], we present data from independent sources to support the basic assumptions in the iterated weakest link model. The model itself and its predictions, however, is not yet tested empirically.

5 Outlook

This paper has demonstrated that treating security investment as a science rather than an art is impeded by many factors, notably the difficulties of estimating probabilities for rare events and quantifying losses in monetary metrics. Some authors have suggested to abandon ROSI altogether. But what are the alternatives? No planning is not an option—it would be a miracle if about US$ 13 billion per year were spent effectively just by accident.

So the medium-term outlook is to refine measurements and models (in this order!). If ROSI and derived metrics are deemed unreliable, they should not be used for anything but negotiating a security budget. More specific models that link cost to security level and security level to benefit are better suited for setting the security strategy or deciding about individual protection measures. They might help to spend smarter and therefore less for the same effect.

As if managing information security investment in a scientific way was not already difficult enough, recent developments are likely to bring new challenges in the future. Ubiquitous network connectivity, novel architectures, and business models fostering massively distributed computing (aka cloud computing) are about to change the security landscape. On the cost side, this will make it more difficult to disentangle security investment from other expenses, e. g. for a redesign of the system architecture. Measures of the security level will become less reliable due to increasing interdependence between loosely connected and autonomous organizations. On the benefit side, detecting and measuring breaches in realtime will require sophisticated monitoring and forensics efforts (which themselves come at a cost). In addition, novel valuation methods will be

[8] the Japanese Ministry of Economy, Trade and Industry

needed to account for the value of (protected/breached/lost) information assets over time [50].

With the increasing dependence of organizations on information and information technology, the borderline between security investment and general risk management is about to blur. On the upside, this underlines the relevance of the subject. On the downside, it makes it even harder to keep an overview of the field and maintain a consistent terminology and conceptual framework.

Acknowledgements

Thanks are due to Kanta Matsuura and the organizers of IWSEC 2010 for the kind invitation to the conference. Kanta further gave helpful comments on an earlier draft and he contributed the security investment statistics for Japan. The paper also benefited from additional comments by Márk Félegyházi. The author gratefully acknowledges a postdoctoral fellowship by the German Academic Exchange Service (DAAD).

References

1. Canalys Enterprise Security Analysis: Global enterprise security market to grow 13.8% in 2010 (2010) `http://www.canalys.com/pr/2010/r2010072.html`.
2. Richardson, R.: CSI Computer Crime and Security Survey. Computer Security Institute (2008)
3. METI: Report on survey of actual condition of it usage in FY2009. `http://www.meti.go.jp/statistics/zyo/zyouhou/result-1.html` (June 2009)
4. Gordon, L.A., Loeb, M.P.: The economics of information security investment. ACM Transactions on Information and System Security **5**(4) (2002) 438–457
5. Willemson, J.: On the Gordon & Loeb model for information security investment. In: Workshop on the Economics of Information Security (WEIS), University of Cambridge, UK (2006)
6. Hausken, K.: Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability. Information Systems Frontiers **8**(5) (2006) 338–349
7. Matsuura, K.: Productivity space of information security in an extension of the Gordon–Loeb's investment model. In: Workshop on the Economics of Information Security (WEIS), Tuck School of Business, Dartmouth College, Hanover, NH (2008)
8. Tatsumi, K.i., Goto, M.: Optimal timing of information security investment: A real options approach. In: Workshop on the Economics of Information Security (WEIS), University College London, UK (2009)
9. 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)
10. Tanaka, H., Matsuura, K., Sudoh, O.: Vulnerability and information security investment: An empirical analysis of e-local government in Japan. Journal of Accounting and Public Policy **24** (2005) 37–59

11. Brocke, J., Grob, H., Buddendick, C., Strauch, G.: Return on security investments. Towards a methodological foundation of measurement systems. In: Proc. of AMCIS. (2007)
12. Jacquith, A.: Security Metrics: Replacing Fear, Uncertainty, and Doubt. Addison-Wesley (2007)
13. Alberts, C.J., Dorofee, A.J.: An introduction to the OCTAVE^{TM} method (2001) `http://www.cert.org/octave/methodintro.html`.
14. Bodin, L.D., Gordon, L.A., Loeb, M.P.: Evaluating information security investments using the analytic hierarchy process. Communications of the ACM **48**(2) (2005) 79–83
15. Su, X.: An overview of economic approaches to information security management. Technical Report TR-CTIT-06-30, University of Twente (2006)
16. 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
17. Sheen, J.: Fuzzy economic decision-models for information security investment. In: Proc. of IMCAS, Hangzhou, China (2010) 141–147
18. Schryen, G.: A fuzzy model for it security investments. In: Proc. of ISSE/GI-SICHERHEIT, Berlin, Germany (2010) to appear
19. Soo Hoo, K.J.: How much is enough? A risk-management approach to computer security. In: Workshop on Economics and Information Security (WEIS), University of California, Berkeley, CA (2002)
20. Geer, D.E., Conway, D.G.: Hard data is good to find. IEEE Security & Privacy **10**(2) (2009) 86–87
21. Anderson, R., Böhme, R., Clayton, R., Moore, T.: Security Economics and the Internal Market. Study commissioned by ENISA (2008)
22. Matsuura, K.: Security tokens and their derivatives. Technical report, Centre for Communications Systems Research (CCSR), University of Cambridge, UK (2001)
23. Böhme, R.: A comparison of market approaches to software vulnerability disclosure. In Müller, G., ed.: Emerging Trends in Information and Communication Security. LNCS 3995, Berlin Heidelberg, Springer-Verlag (2006) 298–311
24. Purser, S.A.: Improving the ROI of the security management process. Computers & Security **23** (2004) 542–546
25. Schneier, B.: Security ROI: Fact or fiction? CSO Magazine (September 2008)
26. Gordon, L.A., Loeb, M.P., Lucyshyn, W.: Information security expenditures and real options: A wait-and-see approach. Computer Security Journal **14**(2) (2003) 1–7
27. Herath, H.S.B., Herath, T.C.: Investments in information security: A real options perspective with Bayesian postaudit. Journal of Management Information Systems **25**(3) (2008) 337–375
28. Yue, W.T., Çakanyildirim, M.: Intrusion prevention in information systems: Reactive and proactive responses. Journal of Management Information Systems **24**(1) (2007) 329–353
29. Grossklags, J., Johnson, B.: Uncertainty in the weakest-link security game. In: Proceedings of the International Conference on Game Theory for Networks (GameNets 2009), Istanbul, Turkey, IEEE Press (2009) 673–682
30. Gordon, L.A., Loeb, M.P., Lucysshyn, W.: Sharing information on computer systems security: An economic analysis. Journal of Accounting and Public Policy **22**(6) (2003)
31. Gal-Or, E., Ghose, A.: The economic incentives for sharing security information. Information Systems Research **16**(2) (2005) 186–208

32. Lee, W., Fan, W., Miller, M., Stolfo, S.J., Zadok, E.: Toward cost-sensitive modeling for intrusion detection and response. Journal of Computer Security **10**(1–2) (2002) 5–22
33. 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
34. Böhme, R., Félegyházi, M.: Optimal information security investment with penetration testing. In: Decision and Game Theory for Security (GameSec), Berlin, Germany (2010) to appear
35. Allen, J., Gabbard, D., May, C.: Outsourcing managed Security Services. Carnegie Mellon Software Engineering Institute, Pittsburgh, PA (2003)
36. Jensen, M.C., Meckling, W.H.: Theory of the firm: Managerial behavior, agency costs and ownership structure. Journal of Financial Economics **3**(4) (1976) 305–360
37. Ding, W., Yurcik, W., Yin, X.: Outsourcing internet security: Economic analysis of incentives for managed security service providers. In Deng, X., Ye, Y., eds.: Prof. of WINE. LNCS 3828, Berlin Heidelberg, Springer-Verlag (2005) 947–958
38. Ding, W., Yurcik, W.: Outsourcing internet security: The effect of transaction costs o managed service providers. In: Prof. of Intl. Conf. on Telecomm. Systems. (2005) 947–958
39. Rowe, B.R.: Will outsourcing IT security lead to a higher social level of security? In: Workshop on the Economics of Information Security (WEIS), Carnegie Mellon University, Pittsburgh, PA (2007)
40. Schneier, B.: Why Outsource? Counterpane Inc. (2006)
41. Cezar, A., Cavusoglu, H., Raghunathan, S.: Outsourcing information security: Contracting issues and security implications. In: Workshop on the Economics of Information Security (WEIS), Harvard University, Cambridge, MA (2010)
42. Böhme, R., Schwartz, G.: Modeling cyber-insurance: Towards a unifying framework. In: Workshop on the Economics of Information Security (WEIS), Harvard University, Cambridge, MA (2010)
43. Zhao, X., Xue, L., Whinston, A.B.: Managing interdependent information security risks: A study of cyberinsurance, managed security service and risk pooling. In: Proc. of ICIS. (2009)
44. Varian, H.R.: System reliability and free riding. In: Workshop on the Economics of Information Security (WEIS), University of California, Berkeley (2002)
45. Hirshleifer, J.: From weakest-link to best-shot: The voluntary provision of public goods. Public Choice **41** (1983) 371–386
46. Kunreuther, H., Heal, G.: Interdependent security. Journal of Risk and Uncertainty **26**(2-3) (March-May 2003) 231–49
47. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: Proceeding of the International Conference on World Wide Web (WWW), Beijing, China, ACM Press (2008) 209–218
48. Cremonini, M., Nizovtsev, D.: Understanding and influencing attackers' decisions: Implications for security investment strategies. In: Workshop on the Economics of Information Security (WEIS), University of Cambridge, UK (2006)
49. Liu, W., Tanaka, H., Matsuura, K.: An empirical analysis of security investment in countermeasures based on an enterprise survey in Japan. In: Workshop on the Economics of Information Security (WEIS), University of Cambridge, UK (2006)
50. Berthold, S., Böhme, R.: Valuating privacy with option pricing theory. In: Workshop on the Economics of Information Security (WEIS), University College London, UK (2009)