IT SECURITY INVESTMENT AND GORDON-LOEB'S $1/e$ RULE

YULIY BARYSHNIKOV

ABSTRACT. Gordon-Loeb's model of investment into mitigation of Information Technology risks is simple and versatile, and thus attracted significant attention of both IT practitioners and economists. One of the claims of the original research was that the optimal investment level never exceeds $1/e$-th fraction of the value at risk, the result they verified for several of instances of their model. At the same time, subsequent works showed that the result is false in the full generality of GL model.

In this note we first put the GL model into a more general context, deriving their postulates from several verifiable axioms. Further, we prove that in this framework, the $1/e$ rule indeed holds in full generality, thus justifying the intuition of [2].

1. INTRODUCTION

Advances in information technology and its penetration of the business over the past few decades was one of the leading forces driving the productivity growth in the US and worldwide. This circumstance led to a surge of interest to economic aspects of information and communication technology.

In particular, the economic aspects of cybersecurity, addressing the specific risks and vulnerabilities of information technologies became increasingly scrutinized, especially given the relatively open, standards driven character of most communication and data processing protocols. We refer the reader to the survey [1] for comprehensive and up-to-date references.

We address here the one of the fundamental yet still poorly understood question of investment levels into IT security. How much a firm facing risks due to its IT vulnerabilities has to invest in mitigating these risks?

1.1. **Main results.**

The goal of this work is twofold. On one hand, I will try to go one step deeper into the understanding the risk mitigation process, postulating some axioms addressing the process as such. The properties of the residual vulnerability function S will then be derived from these axioms, not taken as primitive. Further I will analyze the optimal security investment for general problems satisfying the proposed axioms.[1] It will turn out that for such functions S, the $1/e$ rule of Gordon and Loeb holds.

1.1.1. *Basic assumptions.*

Here is a sketch of the basic setup of this study (details are found in the next section).

I posit that the mitigation process consists of a variety of independent actions (like installs of software patches), each of which reducing the loss probability insignificantly and similarly requiring small investment. A firm is free to choose a collection of the mitigating actions best addressing its demands, to maximize the total utility of such investment.

[1] We do not restrict attention to any particular functional family; indeed, the results are valid generally, and the set of functions S arising within the proposed framework is infinite-dimensional (indeed, an open convex subset of the Banach space $C([0,\infty))$).

To model mitigating actions, whose costs and effects each are individually negligible compared to the effect of the whole, we, following the standard paradigm, will introduce a measurable space, Ω, so that the elements of this space are *elementary actions*, and its subsets are the actions which can be deployed by a firm. We assume that the effects of the actions are independent and the costs are additive. Together these assumption would lead to integral representations of both costs and effects of the actions undertaking by a firm.

1.1.2. *Log-convexity and* $1/e$ *rule.* Under assumption of rationality (for a given budget, the agent always selects the actions within the budget minimizing the residual vulnerability), we prove that the residual vulnerability function S is not merely convex (thus justifying this part of Gordon-Loeb model), but *log-convex*. The main tool here is the Lyapunov convexity theorem.

And log-convexity of S turns out to be the key for the validity of $1/e$ rule: we prove that the optimizer z_* for (**??**) does not exceed R/e, as long as the function S is log-convex (the functional families of studied in [2] are log-convex). Thus the $1/e$ rule is valid for a very broad class of functions.

2. SETUP

We address the issue of rational investment into the security of information technology of a firm within the framework close to that of [2]. As the model is rather general and does not involve explicit modeling of IT risks or of ways of their mitigation, its scope does not restrict to IT investment only. However, we will postulate several properties that IT risks and actions mitigating them should possess; the properties of the model important to us will rely on these properties.

We formulate these properties, true to the fashion of economic literature, as *axioms A0-A3*.

One of the main differentiators of IT risks is the relative ease with which a given threat can be generated, and with which it can be deflected. Compared with the loss a single threat can, under favorable circumstances, inflict on an enterprise, the cost of either generating the threat or of its deflecting is vanishing. This can be seen, for example, in the fast changes of generations of computer viruses and in as efficient and rapid creation of anti-viral software.

A mitigating action against the plethora of IT vulnerabilities should therefore address many of them at once. Let us identify a mitigating tool with the set of risks it is neutralizing.

Accepting this leads one to the following set of axioms:

2.1. **Axioms.**

A0 We assume that *elementary protective actions* are elements of a measurable space of a separable measurable space $(\Omega, \mathcal{F})$, and that the protective actions form *measurable subsets* of Ω. [2] Thus, an admissible protective action is a measurable subset $A \in \mathcal{F}$.

We will assume further that to each (measurable) subset $A \subset \Omega$, we can associate

- the cost of protective measure A, $z(A)$, and
- the residual security risk, which we will denote as $s(A)$

[2]It will not restrict the generality a bit, to think of this measurable space as the unit interval equipped with the σ-algebra of Borel sets.

A1 We will assume that the costs of protective actions are additive: in other words, for disjoint actions A_1, A_2,

$$z(A_1 \amalg A_2) = z(A_1) + z(A_2).$$

Moreover, as, technically, we require z to be defined on all measurable subsets of Ω, we will assume that s is compatible with the σ-algebra structure, and therefore, z *is a positive measure on* Ω.

To formalize the idea that protective actions are infinitesimally small, we will assume that the measure s is *nonatomic*, i.e. that any measurable subset A of positive cost can be split into two disjoint subsets of lesser, yet positive costs (summing up, of course, to $s(A)$).

A2 Similarly, we will require that the residual security risks are *multiplicatively* independent, i.e. for disjoint A_1, A_2,

$$s(A_1 \amalg A_2) = s(A_1)s(A_2).$$

The rational behind this assumption should be clear: the probability for a security risk to escape two disjoint sets of independent protective actions is the product of the probabilities to penetrate each of the defenses.

The logarithm of s is additive, and as before we will require also the σ-additivity, so that

$$u := \log(s)$$

is a (non-positive) measure on Ω.

Involving the infinitesimal smallness of the individual, elementary protective actions, once again, we also postulate that u is nonatomic.

A3 Lastly, we will require that achieving perfect protection cannot be free, i.e. that the range of the *vector valued measure* (s, u),

$$\{(s(A), u(A)), A \in \mathcal{F}\}$$

does not contain $(0, -\infty)$.

Summarizing, our axioms [A0-A3] imply that the protective actions form a σ-algebra $\mathcal{F}$ of subsets of some measurable space of elementary protective actions Ω, equipped with two measures, non-negative, s, and non-positive, u, such that for a protective action $A \subset \Omega$, the cost is $z(A)$ and the residual risk is $s(A) = \exp(u(A))$. We will refer to these two measures as *cost measure* and *security breach* measure, respectively.

3. From measures to functions

Recall that the primitive in Gordon-Loeb's model was a function $S(z)$, called *security breach probability*, describing the residual probability of loss, given investment z. We will recover $S(z)$ assuming that given a budget z, the agent selects the most efficient protective action A.

More formally, let us define $S(z)$ as

$$S(z) = \inf_{A \in \mathcal{F}: z(A) \leq z} s(A) = \exp(\inf_{A \in \mathcal{F}: z(A) \leq z} u(A)).$$

In other words, the least residual risk one can hope to achieve under any protective action A whose cost does not exceed z.

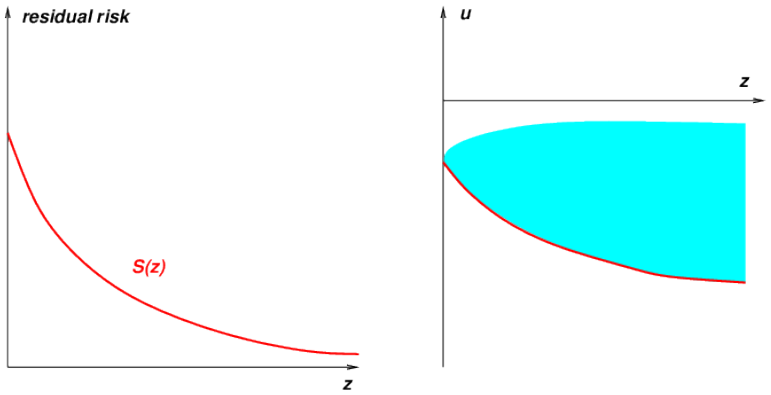

FIGURE 1. Residual risk probability (left) is the exponential of the lower envelope of the range of the vector-valued measure (right display)

3.1. **Main convexity result.**

Proposition 3.1. a) *The range of the (vector-valued) mapping*

$$A \mapsto (z(A), u(A))$$

is a convex closed subset $R \subset \mathbb{R}^2$ *(in fact, a proper subset of the forth quadrant* $\{z \geq 0, u \leq 0\}$*).*

b) *For any* z*, the value* $S(z)$ *is attained on a protective measure* $A \in \mathcal{F}$*;*

c) *The function*

$$v : z \mapsto \log(S(z))$$

is convex.

PROOF: The part a) of the Proposition are immediate corollaries of Lyapunov's convexity theorem, stating that the range of a vector-valued non-atomic measure is closed and convex. The properness of the range follows from our Axiom A3.

The part b) is then immediate by the continuity of exponential function.

The part c) is immediate from a) as well: the lower envelope of a convex set on the plane is a graph of a convex function. □

One immediate reformulation of the Proposition 3.1 is

Proposition 3.2. *The security breach probability* S *is a* log-convex *nonincreasing function of* z*, and therefore convex itself.*

PROOF: The first claim is a mere restatement of what we already know, while the second claim follows as exp is a convex monotonous function. □

The log-convexity of S is in fact fairly intuitive. Indeed, one can think of $S(z)$ as the best collection of protective measures one can acquire given the capital z. Obviously, one should invest into those tools whose (infinitesimal) return is highest. In our context, this return is the rate at which the residual breach probability reduces with incremental increase of the investment z. This rate is non-increasing if the investment is optimal, as the best protection is acquired first. That is equivalent to the log-convexity.

We also remark here, that while we have derived the convexity of S from our axioms A0-3, the differentiability of S (let alone twice differentiability) does not follow from these axioms. As we will see, it is not needed for our primary goal, the $1/e$ rule.

4. Optimal security investments and $1/e$ rule

Now we are ready to prove our main result, that the optimal security investment never exceeds the $1/e$ fraction of the total expected risk.

Theorem 4.1. *Let S be a non-increasing nonnegative log-convex function, and z_* is a solution to the optimization problem*

$$\min_{z\geq 0} \mathrm{LS}(z) + z.$$

Then

$$z_* \leq \mathrm{LS}(0)/e. \tag{1}$$

Proof: The definition of z_* means that

$$\mathrm{LS}(z_*) + z_* \leq \mathrm{LS}(z) + z, \tag{2}$$

for all $z \geq 0$. In other words, inequality (2) implies that the graph of the function $\mathrm{LS}(z)$ lies above the graph of the linear function

$$v(z) = \mathrm{LS}(z_*) + z_* - z \tag{3}$$

Let $e(z) = A\exp(-az)$ be the exponential function such that the values of e and v and of their derivatives coincide at z_*, i.e.

$$A\exp(-az_*) = \mathrm{LS}(z_*) \quad \text{and} \quad -Aa\exp(-az_*) = -1. \tag{4}$$

Taking the logarithms, we arrive at three functions

$$s_l = \log \mathrm{LS}, v_l = \log v, e_l = \log e,$$

such that

(1) $s_l(z_*) = e_l(z_*) = v_l(z_*)$;
(2) $s_l \geq v_l$ on $(0,\infty)$;
(3) *The linear function* e_l has the same derivative as v_l at z_*;
(4) the function s_l is convex.

It follows, that

$$s_l \geq e_l \quad \text{on} \quad (0,\infty).$$

Indeed, otherwise, for some $z' \geq 0$,

$$s_l(z') < e_l(z'), \tag{5}$$

and by convexity,

$$s_l(tz'+(1-t)z_*) < s_l(z_*)+t(s_l(z')-s_l(z_*)) < e_l(z_*)+t(s_l(z')-e_l(z_*)), 0 \leq t \leq 1. \tag{6}$$

On the other hand,

$$s_l(tz' + (1-t)z_*) \geq v_l(tz' + (1-t)z_*) = e_l(z_*) + t(e_l(z') - e_l(z_*)) + O(t^2) \tag{7}$$

(here we used the differentiability of v_l and the fact that v_l and e_l coincide with their derivatives at z_*). Combining (5),(6) and (7) we obtain a contradiction for small enough $t > 0$.

Hence,

$$\mathrm{LS}(z) \geq A\exp(-az), 0 < z < \infty.$$

In particular, $LS(0) \geq A$, and therefore, by (4),

$$LS(0)az_* \exp(-az_*) \geq z_*. \tag{8}$$

As $t \exp(-t)$ is bounded from above by $1/e$, our result follows. □

4.1. **Remark.** The proof of Theorem 4.1 would be even shorter if we assumed twice differentiability of S: in this case we could skip the introduction of the auxiliary function ν deploying just the infinitesimal optimality conditions.

4.2. **Example.** To somewhat juice up somewhat dry presentation, let us consider two functional families from [2], to check that they consist of log-convex functions:

First family is given by

$$S(z) = \frac{\nu}{(az+1)^b}, a, b > 0.$$

Hence

$$\log S = \log \nu - b \log(az+1),$$

which is a patently convex function.

For their second family,

$$S = \nu^{az+1},$$

logarithm is plainly linear, thus (non-strictly) convex as well.

5. Conclusion

One of the direction which would be very interesting to exploit deals with the properties of the residual loss probability function S in the situation where the elementary protective measures are not "sequential". In our model, essentially, the vulnerability, to succeed, has to pass through a sequence of "filters", each having its shot at eliminating the threat. In many situations, this model might be too limited: risks can have different nature, following, so to say, different routes. Understanding the properties of the range of (s, z) in this situation seems to be both interesting and challenging.

References

[1] Ross Anderson and Tyler Moore, The Economics of Information Security, *Science*, **314**, no. 5799 (2006), pp. 610 - 613.
See also www.cl.cam.ac.uk/~rja14/Papers/toulouse-summary.pdf

[2] Lawrence Gordon and Martin Loeb, The Economics of Information Security Investment, *ACM Trans. on Information and System Security*, **5**, no 4 (2002), pp. 438 - 457.

[3] Jan Willemson, On the Gordon&Loeb Model for Information Security Investment. *WEIS 2006*, http://weis2006.econinfosec.org/docs/12.pdf

[4] Kjell Hausken, Returns to Information Security Investment: The Effect of Alternative Information Security Breach Functions on Optimal Investment and Sensitivity to Vulnerability, *Information Systems Frontiers*, **5**, n.8, (2006).

Departments of Mathematics and ECE, UIUC, Urbana, IL
E-mail address: ymb@uiuc.edu