Gordon-Loeb model and the mystery of $1/e$

yuliy **baryshnikov**
Mathematics and ECE
UIUC

October 13, 2011

ROI on IT security

Large and small, for-profit or not, modern organizations are struggling with managing their information technology (IT) assets, trying to make them secure at reasonable cost.

The problem is to maintain right balance between over-investing into IT security (which often negatively influences productivity of co-workers) and under-investing (with contingent potential catastrophic losses, if the IT infrastructure is compromised).

One should remember that the costs of the investment into security are not just direct ones, but the inherent externalities: how many minutes every day your Windows machine goes through all the security checks as it boots?

One the other hand, the externalities of underinvesting are even more pronounced: botnets do not do harm to the computer they infect, only to others...

ROI on IT security

Large and small, for-profit or not, modern organizations are struggling with managing their information technology (IT) assets, trying to make them secure at reasonable cost.

The problem is to maintain right balance between over-investing into IT security (which often negatively influences productivity of the workers) and under-investing (with contingent potential catastrophic losses, if the IT infrastructure is compromised).

One should remember that the costs of the investment into security are not just direct ones, but the inherent externalities: how many minutes every day your windows machine goes through all the security checks as it boots?

One the other hand, the externalities of underinvesting are even more pronounced: botnets do not do harm to the computer they infect, only to others...

ROI on IT security

Large and small, for-profit or not, modern organizations are struggling with managing their information technology (IT) assets, trying to make them secure at reasonable cost.

The problem is to maintain right balance between over-investing into IT security (which often negatively influences productivity of the workers) and under-investing (with contingent potential catastrophic losses, if the IT infrastructure is compromised).

One should remember that the costs of the investment into security are not just direct ones, but the inherent externalities: how many minutes every day your Windows machine goes through all the security checks as it boots?

One the other hand, the externalities of underinvesting are even more pronounced: botnets do not do harm to the computers they infect, only to others...

ROI on IT security

Large and small, for-profit or not, modern organizations are struggling with managing their information technology (IT) assets, trying to make them secure at reasonable cost.

The problem is to maintain right balance between over-investing into IT security (which often negatively influences productivity of the workers) and under-investing (with contingent potential catastrophic losses, if the IT infrastructure is compromised).

One should remember that the costs of the investment into security are not just direct ones, but the inherent externalities: how many minutes every day your Windows machine goes through all the security checks as it boots?

One the other hand, the externalities of underinvesting are even more pronounced: botnets do not do harm to the computer they infect, only to others...

Return on Investment in security

ROI in security in notoriously hard to quantify: one derives profit from the events that are not actualy happening. A random example from the literature:

> *Question: what is ROI in fire extinguishers? Answer: $3 for each $1 invested, according to some studies.*

This talk is addressing a specific model of business efficiency of investment in IT security: the Gordon-Loeb model, and some of its critiques and apologies.

Return on Investment in security

ROI in security in notoriously hard to quantify: one derives profit from the events that are not actualy happening. A random example from the literature:

> *Question: what is ROI in fire extinguishers? Answer: $3 for each $1 invested, according to some studies.*

This talk is addressing a specific model of business efficiency of investment in IT security: the Gordon-Loeb model, and some of its critiques and apologies.

Return on Investment in security

ROI in security in notoriously hard to quantify: one derives profit from the events that are not actualy happening. A random example from the literature:

> *Question: what is ROI in fire extinguishers? Answer: $3 for each $1 invested, according to some studies.*

This talk is addressing a specific model of business efficiency of investment in IT security: the Gordon-Loeb model, and some of its critiques and apologies.

Gordon-Loeb model

In 2002 Gordon and Loeb introduced a simple model for return on IT security, which proved very influential (*Google scholar* returned 433 citations today (October 11, 2011). Perhaps a right combination of simplicity and versatility was the root to its popularity.

It stipulated that a firm, facing potential loss L from a cyber-security risk, can invest a certain amount z to mitigate the risk, that is to reduce the probability of loss.

The defining primitive of the model is the *residual vulnerability* $S(z)$, that is the probability of a loss to happen, given the investment level z.

Parenthetically, they actually considered some extra parameters, like the *vulnerability* of the software, etc, but these parameters do not play any role in the overall analysis.

Gordon-Loeb model

In 2002 Gordon and Loeb introduced a simple model for return on IT security, which proved very influential (*Google scholar* returned 433 citations today (October 11, 2011). Perhaps a right combination of simplicity and versatility was the root to its popularity.

They stipulated that a firm, facing potential loss L from a cyber-security risk, can invest a certain amount z to mitigate the risk, that is to reduce the probability of loss.

The defining primitive of the model is the *residual vulnerability* $S(z)$, that is the probability of the loss to happen, given the investment level z.

Parenthetically, they actually considered some extra parameters, like the *vulnerability* of the software, etc, but these parameters do not play any role in the overall analysis.

Gordon-Loeb model

In 2002 Gordon and Loeb introduced a simple model for return on IT security, which proved very influential (*Google scholar* returned 433 citations today (October 11, 2011). Perhaps a right combination of simplicity and versatility was the root to its popularity.

They stipulated that a firm, facing potential loss L from a cyber-security risk, can invest a certain amount z to mitigate the risk, that is to reduce the probability of loss.

The defining primitive of the model is the *residual vulnerability* $S(z)$, that is the probability of the loss to happen, given the investment level z.

Parenthetically, they actually considered some extra parameters, like the *vulnerability* of the software, etc, but these parameters do not play any role in the overall analysis.

Gordon-Loeb model

In 2002 Gordon and Loeb introduced a simple model for return on IT security, which proved very influential (*Google scholar* returned 433 citations today (October 11, 2011). Perhaps a right combination of simplicity and versatility was the root to its popularity.

They stipulated that a firm, facing potential loss L from a cyber-security risk, can invest a certain amount z to mitigate the risk, that is to reduce the probability of loss.

The defining primitive of the model is the *residual vulnerability* $S(z)$, that is the probability of the loss to happen, given the investment level z.

Parenthetically, they actually considered some extra parameters, like the *vulnerability* of the software, etc, but these parameters do not play any role in the overall analysis.

Gordon-Loeb model, cont'd

Henceforth, our basic assumption is that given investment z, the expected loss is $LS(z)$ for some function S.

As only *risk neutral* firms are considered, we absorb the probability of the loss at $z = 0$ into the factor L, thus normalizing $S(0) = 1$.

Risk neutrality implies that the optimal security investment should be the value z_* minimizing the total expected loss and the costs of mitigation, solving

$$\min_{z>0}(LS(z) + z). \tag{1}$$

Gordon-Loeb model, cont'd

Henceforth, our basic assumption is that given investment z, the expected loss is $LS(z)$ for some function S.

As only *risk neutral* firms are considered, we absorb the probability of the loss at $z = 0$ into the factor L, thus normalizing $S(0) = 1$.

Risk neutrality implies that the optimal security investment should be the value z_* minimizing the total expected loss and the costs of mitigation, solving

$$\min_{z>0} (LS(z) + z). \tag{1}$$

Gordon-Loeb model, cont'd

Henceforth, our basic assumption is that given investment z, the expected loss is $LS(z)$ for some function S.

As only *risk neutral* firms are considered, we absorb the probability of the loss at $z = 0$ into the factor L, thus normalizing $S(0) = 1$.

Risk neutrality implies that the optimal security investment should be the value z_* minimizing the total expected loss and the costs of mitigation, solving

$$\min_{z>0}(LS(z) + z). \tag{1}$$

Gordon-Loeb model, cont'd

Thus we are facing a minimization problem:

$$\min_{z>0}(LS(z)+z). \tag{2}$$

Following customary economic intuition Gordon and Loeb postulated that the function S is

- differentiable;
- non-increasing;
- converges to 0 as $z \to \infty$, and
- is convex.

Gordon-Loeb model, cont'd

Further, Gordon and Loeb investigated two natural parametric families of functions S satisfying these requirements and found, remarkably, that

> *the optimal investment* z_* does not exceed $1/e$-th fraction of the total value at risk:
>
> $$z_* \leq L/e.$$

The families G&L restricted their attention to were

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

Gordon-Loeb model, cont'd

Further, Gordon and Loeb investigated two natural parametric families of functions S satisfying these requirements and found, remarkably, that

> *the optimal investment* z_* does not exceed $1/e$-th fraction of the total value at risk:
>
> $$z_* \le L/e.$$

The families G&L restricted their attention to were

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

refutations

And what about other functions?

It is easy to construct functions S, satisfying all of the properties above, yet such that the optimal investment levels z_* are *arbitrarily* close to the total value at risk L, see **Willemson'06, Hausken'06** (obviously, a risk neutral agent will not spend more on risk mitigation than the expected risk itself).

Here is the idea:

refutations

And what about other functions?

It is easy to construct functions S, satisfying all of the properties above, yet such that the optimal investment levels z_* are *arbitrarily close* to the total value at risk L, see **Willemson'06, Hausken'06** (obviously, a risk neutral agent will not spend more on risk mitigation than the expected risk itself).

Here is the idea:

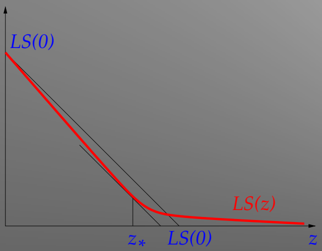

back to basics

My contribution to the problem is twofold:

- I introduce an axiomatic framework allowing to recover G& L's setup from first principles, and
- I show that the resulting functions S do in fact imply the $1/e$ rule.

Let's start.

the model

Where these functions S are coming from? Let's look at the roots.

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.

We introduce several axioms formalizing these notions.

the model

Where these functions S are coming from? Let's look at the roots.

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.

We introduce several axioms formalizing these notions.

the model

Where these functions S are coming from? Let's look at the roots.

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.

We introduce several axioms formalizing these notions.

axioms

A0 We assume that *elementary protective actions* are elements of a separable measurable space $(\Omega, \mathcal{F})$, and that the *protective actions* are tantamount to *measurable subsets* of Ω.
To each (measurable) subset $A \subset \Omega$, we can associate

- the cost $z(A)$ of protective measure A, and
- the residual security risk $s(A)$.

Informally, this axiom expresses the *smallness* of individual protective actions.

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

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

i.e. z is a positive *non-atomic* measure on Ω.

axioms

A0 We assume that *elementary protective actions* are elements of a separable measurable space $(\Omega, \mathcal{F})$, and that the *protective actions* are tantamount to *measurable subsets* of Ω.
To each (measurable) subset $A \subset \Omega$, we can associate

- the cost $z(A)$ of protective measure A, and
- the residual security risk $s(A)$.

Informally, this axiom expresses the *smallness* of individual protective actions.

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

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

i.e. z is a positive *non-atomic* measure on Ω.

axioms, cont'd

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

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

so that $u := \log(s)$ is a (non-positive), non-atomic measure on Ω.

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

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

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

axioms, cont'd

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

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

so that $u := \log(s)$ is a (non-positive), non-atomic measure on Ω.

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

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

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

from measures to functions

The *range*

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

of the measures z, u encodes the potential possibilities of investment into IT security for the firm.

We can now recover $S(z)$ assuming that this is the best residual risk a protective action A feasible under the budget z can achieve.

from measures to functions, cont'd

Formally, let us define

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

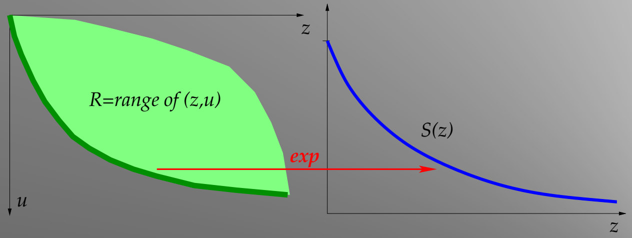

Lyapunov convexity

Proposition

Under Axioms A0-A3,

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.*

Lyapunov convexity, cont'd

In other words, S is well defined, non-increasing, and *log-convex* (hence, convex).

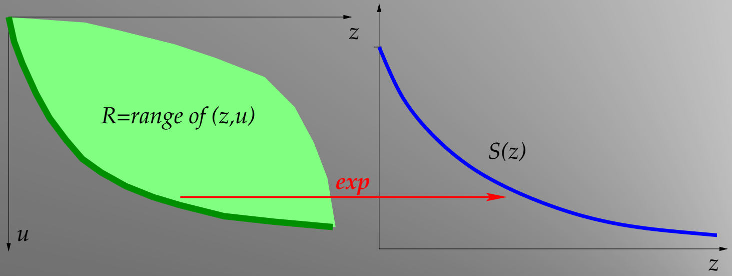

$1/e$ rule vindicated

Theorem

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

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

Then

$$z_* \leq L/e. \tag{3}$$

Log convexity, unlike mere convexity, implies the $1/e$ rule.

$1/e$ rule vindicated

Theorem

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

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

Then

$$z_* \leq L/e. \tag{3}$$

Log convexity, unlike mere convexity, implies the $1/e$ rule.

quick proof

Denote by $f(z) := LS(z)$, and set z_* to be a point where $f(z) + z$ attains its minimum on $[0, \infty)$. Then f lies above the linear function $l(z) := f(z_*) + (z_* - z)$ (and touches it at z_*), and, by log-convexity, also lies above some exponential function $a\exp(-qz)$ that is tangent to l at z_*.

Now [illegible] has

$$aq\exp(-qz_*) = 1 \quad \text{and} \quad f(z) \geq a\exp(-qz),$$

in particular $f(0) \geq a$.

Now

$$\frac{z_*}{f(0)} \leq \frac{z_*}{a} = \frac{qz_*}{\exp(qz_*)} \leq 1/e.$$

quick proof

Denote by $f(z) := LS(z)$, and set z_* to be a point where $f(z) + z$ attains its minimum on $[0, \infty)$. Then f lies above the linear function $l(z) := f(z_*) + (z_* - z)$ (and touches it at z_*), and, by log-convexity, also lies above some exponential function $a\exp(-qz)$ that is tangent to l at z_*.

Now one has

$$aq\exp(-qz_*) = 1 \quad \text{and} \quad f(z) \geq a\exp(-qz),$$

in particular $f(0) \geq a$.

Now

$$\frac{z_*}{f(0)} \leq \frac{z_*}{a} = \frac{qz_*}{\exp(qz_*)} \leq 1/e.$$

quick proof

Denote by $f(z) := LS(z)$, and set z_* to be a point where $f(z) + z$ attains its minimum on $[0, \infty)$. Then f lies above the linear function $l(z) := f(z_*) + (z_* - z)$ (and touches it at z_*), and, by log-convexity, also lies above some exponential function $a \exp(-qz)$ that is tangent to l at z_*.

Now one has

$$aq \exp(-qz_*) = 1 \quad \text{and} \quad f(z) \geq a \exp(-qz),$$

in particular $f(0) \geq a$.

Now

$$\frac{z_*}{f(0)} \leq \frac{z_*}{a} = \frac{qz_*}{\exp(qz_*)} \leq 1/e.$$

back to Gordon-Loeb

One can easily check that both of the families, considered in the G& L's paper are log-convex:

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

Somewhat remarkable, they guessed the correct result!

back to Gordon-Loeb

One can easily check that both of the families, considered in the G& L's paper are log-convex:

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

Somewhat remarkable, they guessed the correct result!

beyond Gordon-Loeb

Gordon-Loeb model - or rather our reinterpretation - can be thought of as the optimization of the sequential chain: to succeed, the attacker has to go through several independent filters, each of which reduces its probability of success.

An obvious and interesting generalization is to move to general topologies: where are the optimal placements of filters there? are there any convexity properties? thus far, unknown...

beyond Gordon-Loeb

Gordon-Loeb model - or rather our reinterpretation - can be thought of as the optimization of the sequential chain: to succeed, the attacker has to go through several independent filters, each of which reduces its probability of success.

An obvious and interesting generalization is to move to general topologies: what are the optimal placements of filters there? are there any convexity properties? thus far, unknown...

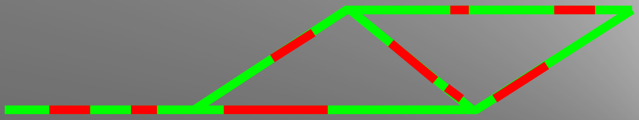

The End