Some name company

Note on terminology

I’m using a mix of insurance and bond lingo. Afaik there’s no set terminology to use and we should end up using the one that makes buyers the most comfortable

What

We raise money from the financial market to allow companies to hedge their cyber security risk

Why

- Every company, regardless of their sophistication, has a security risk that cannot be eliminated

- Insurance policies don’t work because of lack of proper coverage, policy size and actuarial data

How

- We assess daily the Value at Risk of a company
- We raise from the market anywhere between 3 to 10 the VaR and keep it in a special purpose vehicle
- The company pays a coupon to the investors
- If an incident does happen, the investors lose the principal or a part of it depending on the incident cost

VaR

$$\mathrm{VaR}_\alpha(L) = \inf\{l \in \mathbb{R} : P(L > l) \leq 1 - \alpha\} = \inf\{l \in \mathbb{R} : F_L(l) \geq \alpha\}.$$

Our goal is to calculate a which is the confidence score in (0, 1) in R

In our case l is provided by the company and 'vetted' by us.

a is a score based on the infrastructure risk and the company risk profile

We raise from the market anywhere between 3*l and 10*l (in finance this is where banks are supposed to do montecarlo simulations or stress tests to see if they can handle that)

Note on VaR

VaR ignores fat tail risks, that's fine. To address those I think the only solution is to do a CAT-like bond based on the stock price of a company (or something like that)

Risk correlation

- To reduce the chances of investors losing the entire principal in case of an incident companies are split in tranches
- Investors buy in a tranche containing uncorrelated companies
- Tranches are rated like bonds (AAA to junk) based on the risk of each tranche (as the combination of the risk of the companies in each tranche)

Why tranches?

Mostly because of the central limit theorem..

“Towards infinity the mean of a sufficiently large set of independent random variables approximates a normal distribution” (given certain conditions)

Intuitively it means that if you get enough companies their average risk will approximate a normal distribution

Why tranches?

Practically it means that anything at 3-4sigma to the left can be considered AAA and anything at 3-4sigma to the right should be considered Junk

In other words, “bond” holders should feel pretty confident about the underlying risk they are buying

One last note on tranches

What explained in the previous two slides is by *no means* guaranteed to be true, but intuitively it should be.

Plus, we only need it to be approximately true anyway.. Essentially we need to make sure that investors will not lose their entire investment if one or two companies are indeed breached

Achtung

- The key here is: correlation. If we mess up and stuff is too correlated we'll create another mortgage subprime-like crisis.. Not cool

- Also this implies that for each company we should have N bonds not just 1

The tech details

- When a company is enrolled the CISO and the Board decide which part of their infrastructure they want to “insure” and its value
- We enumerate all points of entry to that part of the infrastructure
- We assign probability of compromise to each point of entry
- The probability of the incident is the probability of the least secure point of entry

Point of entry evaluation

- Analyze the software stack installed
- Spawn honeypots that replicate part of the infrastructure and observe their rate of compromise
- Use collected data on similar infrastructure to assess the risk
- Potentially run Capture the flag-style competitions for exotic infrastructures (a kaggle competition for security roughly)
- Rate the set of defenses added on the top of the insured infrastructure based on our internal scoring system

Motives is also important

- We gather data about compromised companies from the internet
- We split companies by sector and analyze connections between them
- We rate companies based on the industry sector, the connectedness to other companies and regulatory environment

Technology risks 1

- For highly customized environments you might not be able to replicate the infrastructure on the outside w/o leaking confidential data. In that case you have to do a classic pentest and then rate it (less reliable)
- If the data they care to protect is too spread out in the company it might be hard to isolate the entry points. That can be fixed by rating the bond as Junk?

Technology risks 2

- Especially at the beginning this thing might not scale very well. It will eventually once you have enough data points that you can assess the risk w/o doing any labor intensive work

“bond” triggers

- We select a number of verified partners that can attest if an incident happened
- When appropriate we also require the company to disclose the incident
- We use our own data and our lawyers to check the company claims

Fee structure

- We earn a % (1-4%) of the coupon companies pay to “bond” holders, plus a fixed (by company size) fee for the initial assessment
- We can license our risk model to insurances and reinsurances companies
- When the market grows we can also setup an option market for “bond” holders and we collect fees on the trades

Financial risk

- S&P + Deloitte take us over?
- SEC shuts us down?
- Companies want us to buy the ‘bond’ and then place them? (we don’t have enough capital for that)
- Companies bypass us and do this on their own with their IB?

Non-obvious side effects (positive)

- We can start publish our rating for security products and eliminate the crappy ones from the market
- The market will eventually set the price and the confidence in a company which should put pressure on them to improve their security
- Kaggle-like/Capture the flag competitions might do what bug bounties are doing for bugs = give an alternative option for people to monetize their offensive skills

Non-obvious side effects (negative)

- Companies might start to become reckless in the short term (we can probably fix that)
- Corporate espionage might raise because of our rating?
- If somebody creates a way to short this instrument, people might be incentivized to attack companies and make a profit on the market

Team

- 4 security guys
- 2 data analysts
- 2 financial analysts
- 2 quants/actuaries
- 1-2 lawyers (could be a contractor)
- 2 sales guys (could be a VAR on the tech side, could be the IB on the financial side)
- 2 devs

That's roughly $2.9m/year for salary + overhead

Connections needed

- Aon/AIG/insurance/reinsurance
- IBs/hedge fund(s)
- Incident response company
- Law firm

Open questions

- Will CFO/companies go for this?
- Will investors go for this?
- How to gain trust from both? Is it worth giving our model for free (at first) to insurances & reinsurances to build a brand?
- What are the legal requirements for the financial side?
- How to get a critical mass of companies to create the tranches?