PTBA - Risk Selection In Cyber Insurance Underwriting

Ari Chatterjee, ACAS & Dr Raveem Ismail

ABSTRACT

Cyber is an emerging line of insurance, which has demonstrated tremendous growth potential over the next decade. Since it is also an anthropogenic peril, with evolving threat landscape and coverages, it is naturally challenging to underwrite. Here, we propose a new and simple measure, the PTBA (Propensity To Be Attacked). Its key advantages are that it is simple to calculate, and is driven by the interplay between attacker motivation and cybersecurity defence. It produces a single number as an output, and is therefore an ideal *risk score*, a familiar concept in the insurance world (e.g., the terrorism class), and pivotal to quick and practical relative risk appraisal required for underwriting decisions.

Keywords: Cyber, Insurance, Reinsurance, Underwriting, Pricing, Risk, Modelling, Catastrophe.

Table of Contents

1. RESEARCH CONTEXT & OBJECTIVE

Given the paucity of literature on cyber risk insurance, pricing and underwriting, this paper aims to outline a method to underwrite and select risks in a dynamic cyber threat landscape. The traditional methodology of risk classification fails to capture the dynamic nature of the threat landscape and very often, data collected by insurers is insufficient for constructing sophisticated risk classes. We believe the proposed will assist underwriters and actuaries in profitably underwriting cyber insurance.

2. DEFINING PROPENSITY TO BE ATTACKED (PTBA)

The expected income to an attacker from a cyber-attack is the value of each record hacked, plus any other value that might be derived from the target. I.e., if **I** is the *expected income*, then:

$$I = N_{PII}C_{PII} + N_{PHI}C_{PHI} + O ,$$

where:

- **N** is the number of records the attacker expects to exfiltrate from a target firm.
- **C** is the expected price per record.
- **O** is *other gains expected by attacker from target* (includes ransom, IP, possible trading insights, possibility of gaining access to larger targets, recognition, etc.).
- **PII** is Personally Identifiable Information (as defined by NIST, e.g., name, date of birth, credit card information, email address, etc.).
- **PHI** is Protected Health Information (as defined by HIPAA, e.g., names, medical records, biometric details, etc.). The estimated relative value of PHI to PII is 50:1 (World Privacy Forum[1]).

Attackers also have (daily) costs in order to achieve their income - "profits" are the difference between costs and potential income:

$$P = I – Kt ,$$

i.e.,

$$P = N_{PII}C_{PII} + N_{PHI}C_{PHI} + O – Kt ,$$

where:

- **P** is expected profit for attacker from target.
- **K** is the *daily cost of executing a cyber-attack* (including reconnaissance, infrastructure, outsourcing, cost of hiring insiders, paying for credentials, cost of zero-day vulnerabilities, consequences of getting caught, etc.).
- **t** is time required to breach the target.

Then, for the attacker, the aim is to maximise the profit function **P** across all targets:

$$Max(P) = Max(N_{PII}C_{PII} + N_{PHI}C_{PHI} + O – Kt) .$$

Therefore, for the attacker to ascertain target desirability simply means sorting targets in descending order by **P**. I.e., it will be preferable to attack firms with a higher **P** (profit function) first.

Each attacker will have their own, potentially unique, list in which a firm appears at a certain percentile rank **R**. NB:

Absolute Rank = No 1 ---> Percentile Rank = 0%

Absolute Rank = No [Last] ---> Percentile Rank = 100%.

Since any one firm will be a target for multiple attackers, with various different value of **R** in each attacker's desirability list, *the sum of **R**, across all considered attackers **n**, for any one target firm, is a measure of the overall susceptibility of the target to attackers.* We therefore define PTBA, the Propensity To Be Attacked, as:

$$PTBA = (\sum_n R) / n .$$

The higher this metric, the higher is the likelihood to be attacked (elevated risk).

3. SOME OBSERVATIONS ON KEY PARAMETERS

C (expected price per record):

- Given the sensitive nature and value of healthcare information, it is no surprise that $\mathbf{C_{PHI} > C_{PII}}$[1].

O (other gains expected by attacker from target):

- Is highly correlated to the target's industry. E.g., investment banking, hedge funds[2], law firms, accounting firms, etc., all have (motivating) gains, other than data exfiltration, for an attacker. E.g., ready funds to transfer, etc.
- May be high for smaller vendors working for larger corporations: attackers can leverage such a relationship to attack the larger organisation. The Target breach was via a HVAC vendor[3].
- For hacktivists, terrorists and nation states, **O** is non-monetary. As with terrorism, their aim is often to maximise propaganda-of-the-deed than monetary profit (**P**).
- The ransom demanded from ransomware victims is an example of **O**. Generally, the ransom is designed in a way that the victim is better off paying quickly without waiting long, thus ensuring that the cost of suffering (cost to recreate data + cost of unavailability of systems) is below the ransom amount. Globally, about 40% of victims pay[4].

K (daily cost of executing a cyber-attack):

- Depends on the type of attacker. A sophisticated and well-resourced attacker capable of absorbing larger expense would generally have a better chance against larger targets. For less sophisticated adversaries, a different victim set or different attack type (with less technical complexity e.g. Ransomware, DDoS) may maximize profits[5].

- Increases significantly[6] if there are legal or financial consequences the attacker faces for its action and can be a powerful deterrent to attack.
- To minimize **K** an attacker may try to re-use the same attack components on similar firms e.g., industry peers, or companies using similar technology. For example, Target and Home Depot hacks included variants of BlackPOS malware, the Sony hack used Destroyer which had code level similarities with Shamoon, used to attack Saudi Aramco[7].
- Attackers are opportunistic. They go after easiest targets first, not wasting time where quick results are not yielded. Attackers tend to quit when their target firm exhibits strong security[12].
- The time to deter the majority of attacks is less than two days. The longer an organisation can keep the attacker from executing, the more likely the attacker will move to the next target (a parallel from the terrorism space is *target substitution*). Higher IT maturity may therefore deter attackers from pursuit of the target firm[12].
- For calculating **K**:
    - 69% of the attackers are motivated by money. On average, attackers receive $28,744 annually for every 704 hours spent on attacks[12]. This is dissimilar to terrorism, where ideology and propaganda-of-the-deed are key.
    - Attacker technology and availability is improving, enabling more attacks. Technically proficient attackers spend an average of $1,367 for specialized tools to execute attacks[12].

  If we ignore the (possibly eventual) cost of extradition or legal costs to the attacker, we can calculate an average daily cost, the aforementioned **K**:

$$K = (\$28{,}744 + \$1{,}367) / 704 \text{ hours} = \$42.8 \text{ per hour.}$$

t (time required to breach target):

- Depends on both the maturity of IT security employed by the victim and the sophistication of the attacker.

PTBA (Propensity To Be Attacked):

- Annual revenues are not an exact indicator for PTBA, since the target could be in business of managing third party data (e.g., payroll processors, accountants) which could

be of a different value to what its own revenues might imply.

- A bank might have a very high desirability and a large payoff. What prevents it having a high PTBA, is that its stringent countermeasures attenuate its profit function, hence it is by no means guaranteed that a bank would be first in the percentile ranking of profit function.
- PTBA is dimensionless: regardless of how long the list of targets held/considered by each attacker, and regardless of how complete the attacker spectrum characterisation, it is a normalised score between zero and one.

4. CALCULATING EXAMPLE PTBAS

In principle, highly granular data on each individual attacker could be ascertained via the dark web and/or sinkholes. However, since representative (let alone exhaustive) compilation of these is not currently possible in practise, using *attacker groups*, a broader and more practical classification, covers all types of attacker.

Using VCDB data, we can calculate the PTBA across a range of industry classes for a spectrum of attacker types, across a two-year period (2015-2016):

Sector	PTBA			
	Crime	Hacktivist	Nation State	Malicious Insider
Accommodation	0.789	0.526	0	0.631
Administrative	0.315	0	0	0.473
Agriculture	0	0	0	0
Construction	0	0	0	0.157
Educational	0.684	0.526	0	0.842
Entertainment	0.315	0	0	0.263
Finance	0.894	0.842	0	0.894
Healthcare	1	0.842	0	1
Information	0.631	0.947	0.736	0.684
Manufacturing	0	0	0	0.526
Mining	0.315	0	0	0
Other Services	0.578	0.789	0.736	0.578
Professional	0.684	0.736	0.947	0.789
Public Sector	0.947	1	0.947	0.947
Real Estate	0	0	0	0.263
Retail	0.842	0.526	0.736	0.736
Trade	0.315	0.526	0	0.368
Transportation	0	0	0	0.421
Utilities	0.315	0	0.736	0.157

Since PTBA can be calculated for any granularity, we can also combine all attacker types to more simply look at how the threat landscape changes over time (2015-2016 to 2016-2017):

Sector	PTBA	
	2015-16	2016-17
Accommodation	0.4865	0.44425
Administrative	0.197	0.111
Agriculture	0	0
Construction	0.03925	0.097
Educational	0.513	0.5135
Entertainment	0.1445	0.49975
Finance	0.6575	0.666
Healthcare	0.7105	0.6665
Information	0.7495	0.72175
Manufacturing	0.1315	0.13875
Mining	0.07875	0.0555
Other Services	0.67025	0.6385
Professional	0.789	0.5275
Public Sector	0.96025	0.87475
Real Estate	0.06575	0
Retail	0.71	0.722
Trade	0.30225	0.208
Transportation	0.10525	0.111
Utilities	0.302	0.2775

Hence, we infer that the public sector is the most hazardous industry class, while agriculture is the least, borne out empirically, *and* according to the PTBA measure which objectively quantifies such risk.

5. CONCLUSION

We have shown that using readily available historical data, or forecasts for future events[15], that it is possible to calculate a single-number risk score: the PTBA (Propensity To Be Attacked). This metric takes into account both attacker motivations and cost, and defender cyber countermeasures. It varies correctly across time, industry, and attacker type. It is flexible and dimensionless: regardless of how long the list of targets held/considered by each attacker, and regardless of how complete the attacker spectrum characterisation, it is a normalised score between zero and one, making it ideal for underwriting both single risks and portfolios, for insurance and reinsurance.

REFERENCES

1. *2017 Cost Of Data Breach Study.* https://www.ibm.com/security/data-breach/
2. *Cyber Attackers Turn Their Focus On Hedge Funds.* http://usblogs.pwc.com/assetmanagement/cyber-attackers-turn-their-focus-on-hedge-funds/
3. *Target Hackers Broke In Via HVAC Company.* https://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company/
4. *Malwarebytes: Understanding The Depth Of The Global Ransomware Problem.* https://go.malwarebytes.com/OstermanRansomwareSurvey.html

5. *Ransomware As A Service.* https://documents.trendmicro.com/assets/resources/ransomware-as-a-service.pdf
6. *Three Chinese Hackers Fined $9 Million For Stealing Trade Secrets.* http://thehackernews.com/2017/05/chinese-hacker-trade-secrets.html
7. *Recycle, Reuse, Reharm: How Hackers Use Variants Of Known Malware To Victimize Companies And What Paypal Is Doing To Eradicate That Capability.* https://www.paypal-engineering.com/2015/11/19/recycle-reuse-reharm-how-hackers-use-variants-of-known-malware-to-victimize-companies-and-what-paypal-is-doing-to-eradicate-that-capability/
8. *Black Market Medical Record Prices Drop To Under $10, Criminals Switch To Ransomware.* http://www.csoonline.com/article/3152787/data-breach/black-market-medical-record-prices-drop-to-under-10-criminals-switch-to-ransomware.html
9. Chopitea, Thomas. 2012. *Threat Modelling Of Hacktivist Groups – Organization, Chain Of Command, And Attack Methods.* http://publications.lib.chalmers.se/records/fulltext/173222/173222.pdf
10. *Crowdstrike, Art Of Attribution* https://www.rsaconference.com/writable/presentations/file_upload/anf-t07b-the-art-of-attribution-identifying-and-pursuing-your-cyber-adversaries_final.pdf
11. *NIST Cybersecurity Framework* https://csrc.nist.gov/Projects/Program-Review-for-Information-Security-Assistance/Security-Maturity-Levels
12. *The Real Cost Of Attacks* https://media.paloaltonetworks.com/lp/ponemon/report.html
13. Holt, Thomas J and Smirnova, Olga. 2014. *Examining The Structure, Organization, And Processes Of The International Market For Stolen Data.*
14. VCDB Github, https://github.com/vz-risk/VCDB
15. Ismail, Raveem & Werner, Christoph. 2017. *Structured Expert Judgement for (Re)insurance: Forecasting Political Violence frequency.* Journal Of Terrorism & Cyber Insurance. Vol 1 No 1. https://1drv.ms/b/s!AjWDpOlDwNZGgbgsgU8u0E6BYYJ5-w

APPENDIX: WALKTHROUGH FOR REPLICATING CALCULATIONS

The PTBA formulation is agnostic to data used: here, we have used the VCDB database for its virtues of being open source and having good coverage (7,300 events at the time of writing).

The reproducible steps and assumptions are:

1. Group by attacker class (Malicious Insider, Nation State, Criminak, Hacktivist):
   - **Malicious Insider**: "actor.internal" = "TRUE". Accidental data releases do occur, but without capturing true motives, malicious intent can be assumed for conservatism.
   - **Nation State:** "actor.external.variety.Nation-state" or "actor.external.variety.State-affiliated" = "TRUE".
   - **Crime:** "actor.external.variety.Organized crime" = "TRUE".
   - **Hacktivist:** "actor.external.variety.Activist" = "TRUE".
2. Calculate the number of attacks by each actor category by year and industry class (simple pivot). We used victim.industry.name as the industry class, and timeline.incident.year as the year.
3. In the absence of further data, we assume that the number of events (for a particular attacker type and industry) is the manifestation of ranking in target desirability.
4. Calculate PTBA for a particular industry for a given year and an attacker category (our first table): $PTBA_{industry} = PercentileRank_{industry}$, where, in the absence of further data, the number of attackers is assumed the same for each attacker category.
5. For PTBAs without breaking out attacker types (our second table), we simply sum PTBAs across attacker type (each row in our first table) and divide by 4, exhibiting the additive utility of the risk score.

It should be noted that the PTBA calculation is a framework for dealing with a heterogonous spectrum of data quality. The VCDB data are extremely basic, but we have shown how they would plug into the PTBA calculation – comprehensive data on unknowable or challenging to acquire information (such as desirability, number of attackers, etc.) would go straight into the PTBA calculation and improve it, but even in its absence, a useful metric can be produced.