Quantitative Security

Colorado State University

Yashwant K Malaiya

CS 559

L6: Insurance and Probability

CSU Cybersecurity Center
Computer Science Dep

About this Course

CS 559 is a research-oriented course.

- 200-level classes: little student content
- 400-level: 5% student presentations/discussions
- 530: 10-15% student presentations/discussions
- 559: 25-40% student presentations/discussions

Quick Project Presentations

- Presentations Sept 15, 17 coming week
  - Schedule will be on Canvas discussions
- 12-15 min presentations for each person, 12 slides
  - Need time for discussions
  - When two persons have the same topic, they must collaborate, to minimize overlap. The presentations will be scheduled one after the other.
  - Post slides/abstract 24 hours in advance Canvas Discussions in Quick Research Abstract and presentation forum
  - Everyone should preview upcoming presentations
  - Online section: Submit abstract, presentation and a link to your video by Tuesday 5 PM.
    - Slides: Video: same limits
- On-campus: 5-10 minutes discussion for each presentation
- Everyone must view all presentations, comment/questions on at least half of the presentations
- Provide feedback using the Peer Review form due 9/19 Sat. Also do a detailed review of two assigned reports.

Peer Review form

- Will be available soon in Canvas. Due Sept 20, Sat.
- Evaluate each category using 10-point scale.
  - 10: top 25%, 9: next 25%, 8: next 25%, 7: bottom 25%.
  - 25% is about 7 students.
  - Overall score should be average of the other 4 scores.

Risk Transfer

- Risk can be transferred by buying insurance.
- Insurance companies spread the risk to the pool of their insured customers.
  - The society (a government) can assume some of the risk in exchange for collecting taxes.
- The insurance premiums and payouts depend on the risk (and operational costs and profits)
- Some insurance companies transfer some of their risk by buying reinsurance from reinsurance companies.

How does insurance work?

- Illness is unpredictable and need for treatment is uncertain, so individuals cannot predict their future health care expenditure
- Based on historical information, insurance providers can predict the probability of a risk for a large group of insured people and estimate the average cost of the risk
- True for security breach risks except ..

How does insurance work?

- Insurance takes all the risks in a group and puts them into a pool

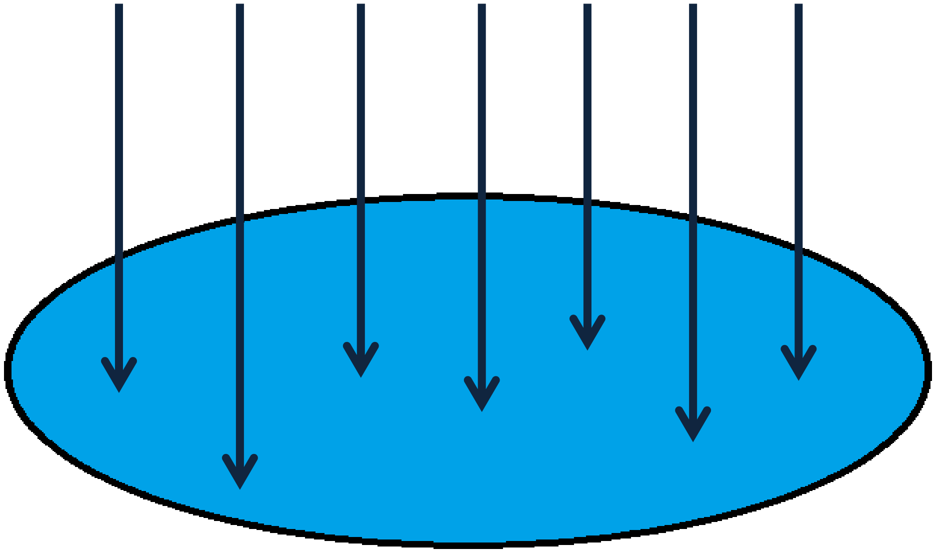

- Not all insured persons claim their benefits at the same time
- Contributions paid by all insured members are used to compensate for the financial consequences of the few persons who experience the risk

Colorado State University

Calculating insurance premiums

In Insurance that is *actuarially fair*, the premium is equal to expected claims (ignores overheads/company profit).

Example 1: Kate was born with a rare disease and has a 40% chance of relapse in a year. If relapse occurs, she has to visit a doctor for consultation once every 3 months. Each consultation costs 30$. Calculate the fair insurance premium.

Premium = probability of illness in a year x average no. of utilization of services per year x unit cost of each utilization

= 0.4 x 4 x 30$

= 48$

40% = 40/100 = 0.4

Actuaries analyze the financial costs of risk and uncertainty.

Colorado State University

Calculating insurance premiums

Example 2 (cont.): Further, in the consultation the doctor may either prescribe medication or recommend that she take laboratory tests. Probability of prescription is 80% while that of tests is 20%. Cost of medication is 50$ and cost of tests is 150$. Calculate the fair health insurance premium.

Premium = probability of illness in a year x average no. of utilization of services per year x unit cost of each utilization

= 0.4 x 4 x [30$ + (0.8 x 50$) + (0.2 x 150$)]

= 0.4 x 4 x [30$ + 40$ + 30$]

= 0.4 x 4 x 100$ = 160$

Colorado State University

Loss Ratio

- The loss **ratio is** the ratio of incurred losses and loss adjustment expenses to premiums earned.

$$Loss\ Ratio = \frac{incurred\ losses + loss\ adjustment\ expenses}{premiums\ earned}$$

- *Example:* A company paid $133.6 million in losses and $14 million in loss-adjustment expenses. Premiums earned were $205 million.

$$Loss\ Ratio = \frac{133.6 + 14}{205} = 0.72$$

- The loss ratio is often in the 65 percent to 75 percent range. Thus the insured pool with get back about 65 to75% of the premiums paid. Affordable Care Act (ObamaCare)requires it to be minimum 80% (85% for large groups).

Principles of risk Management and insurance, Rejda, McNamara, 2017

Asymmetric information

Adverse selection

- The insured person conceals information that places them in a high-risk bracket
  - Thus, average risk of the insured group increases and premium rises
  - May cause low-risk people to leave, high-risk people are left in the group
  - E.g. Bob must undergo a surgery within the next 10 months and joins a health insurance scheme, knowing that the operation will be covered when the waiting period is over. (*preexisting condition*)

Moral hazard

  - The insured person takes risks and is careless about their safety, knowing that they are protected from financial losses
    - E.g. Jenny has a good insurance policy and goes to see her general practitioner, allergist, gynaecologist and dermatologist at least once every month

Asymmetric information

Ways to minimize adverse selection and moral hazard

- establish mandatory insurance, so that high-risk and low-risk people are members of the risk pool
- exclude predictable events such as planned surgeries from the benefit package
- implement co-payments and limitations, e.g. maximum number of days of hospitalization, health expenditure reimbursement up to a maximum level
- establish long waiting periods
- put in place control mechanisms like pre-authorization of high-cost planned surgeries

Why buy insurance?

Statistically, insurance buyer is not likely to come ahead by buying insurance.

- However, a person or a company has a limited capacity of bearing a loss. Insurance will protect against catastrophic losses. Such losses are spread to the population of the insured in form of premiums.
- The Utility function gives the value of money to a person. A person may be willing to pay more than *actuarially fair* premium. The maximum willingness can be computed using the Utility function.

Utility functions and Risk aversion

- How much insurance should one buy?
- Depends on
  - Risk aversion
  - Utility function
- Common utility functions
  - Diminishing marginal utility of wealth
  - Common forms
    - $u(w) = \ln(w)$
    - $u(w) = r(1- e^{-w/r})$ where $r>0$.
    - $u(w) = w^{1/2}$ (used in a simple example next)

Utility In Insurance

Utility in insurance quantifies the benefit a buyer gains by transferring cyber risk to an insurer, thereby reducing the variability (or uncertainty) of financial losses and increasing the firm's overall expected satisfaction or economic value.

- Risk-averse organizations prefer **more predictable outcomes**, even if the expected value is slightly lower. Cyber insurance provides that.
- Utility function for a risk-averse organization with respect to payment can be of the form (in the example next)

$$u(w) = 1 - w^{0.5}$$

- The goal is to **maximize expected utility**, not just expected profit.

Example: Football Player 1

Player's utility function is $u(w) = 1\text{-}w^{0.5}$

where w is the earning from the game.

- Player earns \$1 million if healthy.
- There is a 10% probability the player is injured and earns \$0

E(u) = Prob(healthy).u(if healthy)+(1- Prob(healthy)).u(if injured)

$= 0.9 \times (1{,}000{,}000)^{0.5} + 0.1 \times (0^{0.5}) = 900$

Actuarially fair insurance premium

=prob(loss) x size of loss = 0.1 x 1,000,000 = \$100,000

If the player buys insurance ..2

Earnings with insurance (premium paid is $ 100,000)

- Healthy: \$1 mill. -100,000 = \$900,000
- Injured: \$0 - 100,000 + \$1 mill. (from insurance)

  = \$900,000 same! (as expected)

But consider utility

E(utility with insurance)

$$= 0.9 \times (900{,}000^{0.5}) + 0.1 \times (900{,}000^{0.5})$$

$$= 948.7$$

which is better than E(utility without insurance) of 900.

Max willingness to pay 3

- The max the player will be willing to pay for insurance with be such that utility is not less than that without insurance.
- Utility without insurance is 900 units.

  $900 = w^{0.5}$

  $w^* = 810{,}00$
- Max willingness to pay

  $\$1 \text{ mill} – w^*$

  $= \$1 \text{ mill} – 810{,}00 = \$190{,}000$
- It can also be seen graphically using a plot of u(w)

Cyber coverage may be obtained for ..1

- **Network Security Liability**: liability to a third party as a result of a failure of your network security to protect against destruction, deletion, or corruption of a third party's electronic data, denial of service attacks against internet sites or computers; or transmission of viruses to third party computers and systems
- **Privacy Liability**: liability to a third party as a result of the disclosure of confidential information collected or handled by you or under your care, custody or control. Includes coverage for your vicarious liability where a vendor loses information you had entrusted to them in the normal course of your business.
- **Regulatory Investigation Defense**: coverage for legal expenses associated with representation in connection with a regulatory investigation, including indemnification of fines & penalties where insurable.
- **Event Response and Crisis Management Expenses**: expenses incurred in responding to a data breach event, including retaining forensic investigator, crisis management firm and law firm. Includes expenses to comply with privacy regulations, such as communication to impacted individuals and appropriate remedial offerings like credit monitoring or identity theft insurance.

Cyber coverage may be obtained for.. 2

- **Cyber Extortion**: ransom &/or investigative expenses associated with a threat directed at you that would cause an otherwise covered event or loss
- **Network Business Interruption**: reimbursement of your loss of income and / or extra expense resulting from an interruption or suspension of computer systems due to a failure of technology. Includes coverage for dependent business interruption.
- **Data asset protection**: recovery of costs and expenses you incur to restore, recreate, or recollect your data & other intangible assets (i.e., software applications) that are corrupted or destroyed by a computer attack.

The Cyber Insurance Market

Source: NetDiligance, marketsandmarkets.com

Market capacity:

- Over 50 markets selling or participating in cyber insurance

Appetite & Approach: Different for each insurer and varies by:

- Size: revenue, record count, transaction volume
- Industry: Healthcare, Retail, Finance, Higher Ed, etc.
- Jurisdiction: USA, Canada, Europe, Asia, etc.

Principal Markets:

- For larger risks: AIG, Beazley, Zurich, Chubb, Safehold (representing certain Lloyd's Syndicates)
- For SME, key markets: capacity is plentiful--One Beacon, Philadelphia, etc.

Market Size:

- projected to grow from USD 11.9 billion in 2022 to USD 29.2 billion by 2027

The Cyber Insurance Market

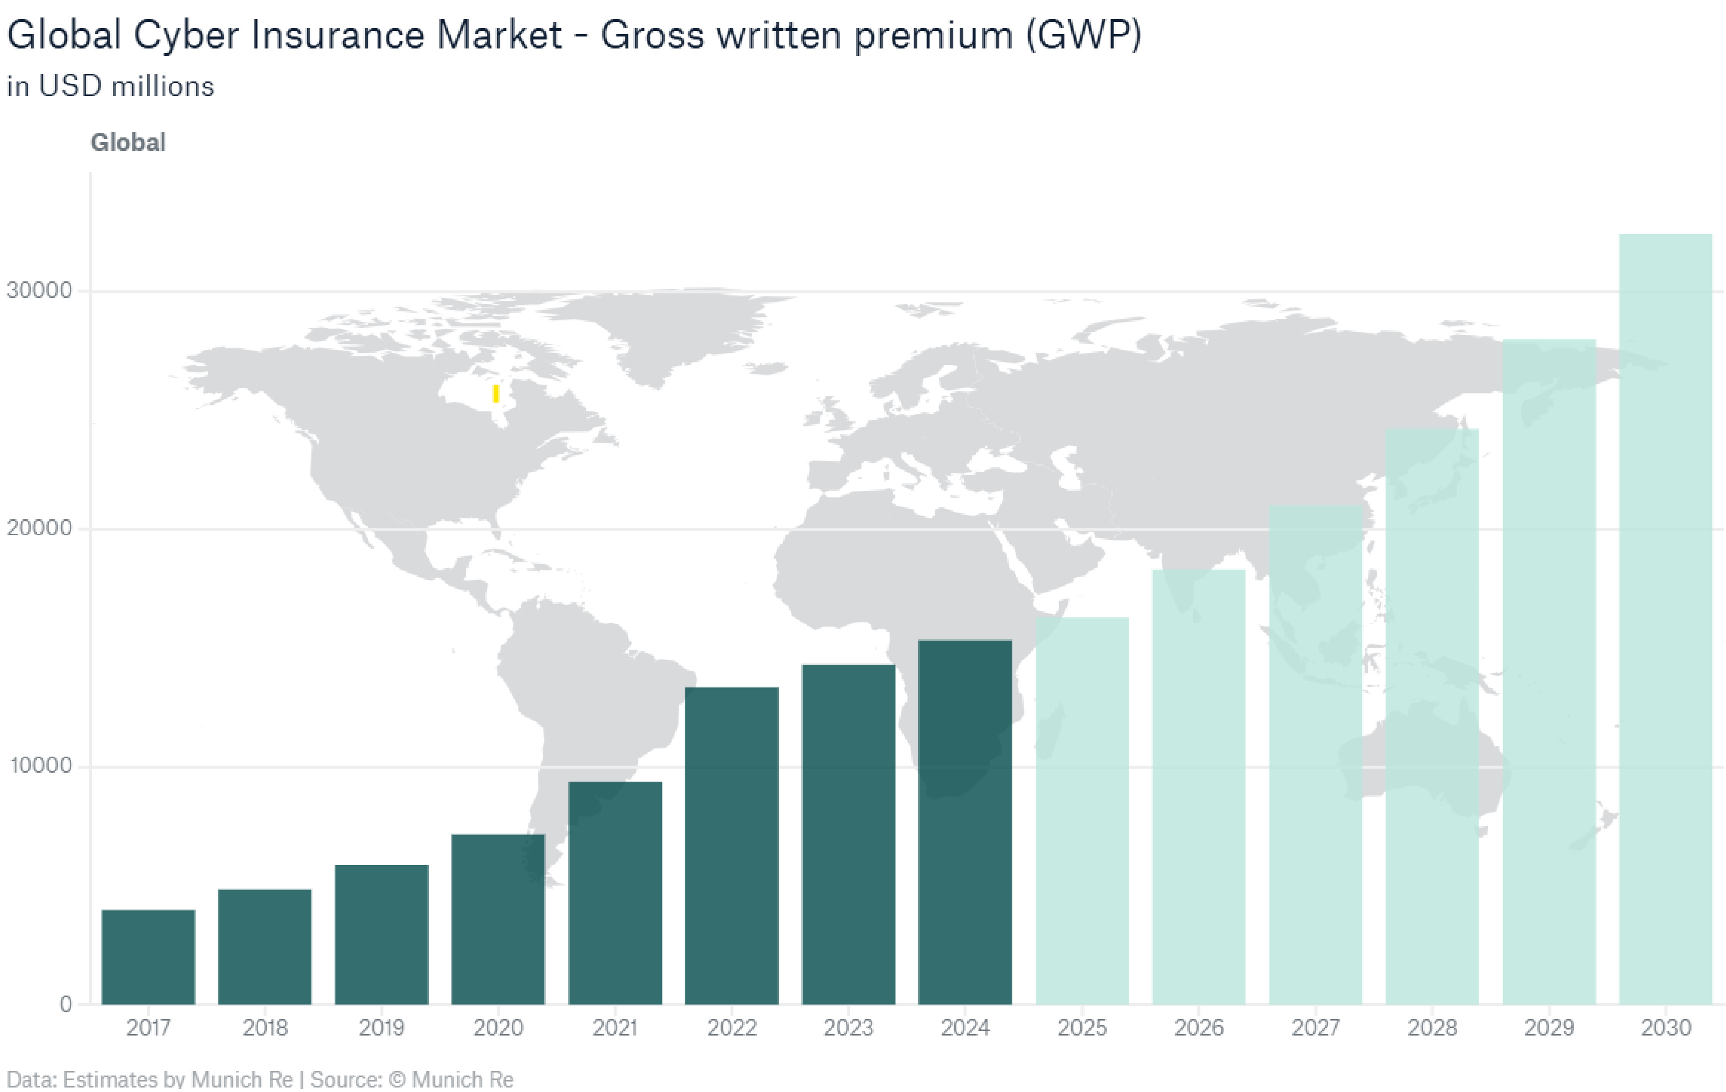

Colorado State University

Change in Cyber Insurance Premiums, 2014-2025

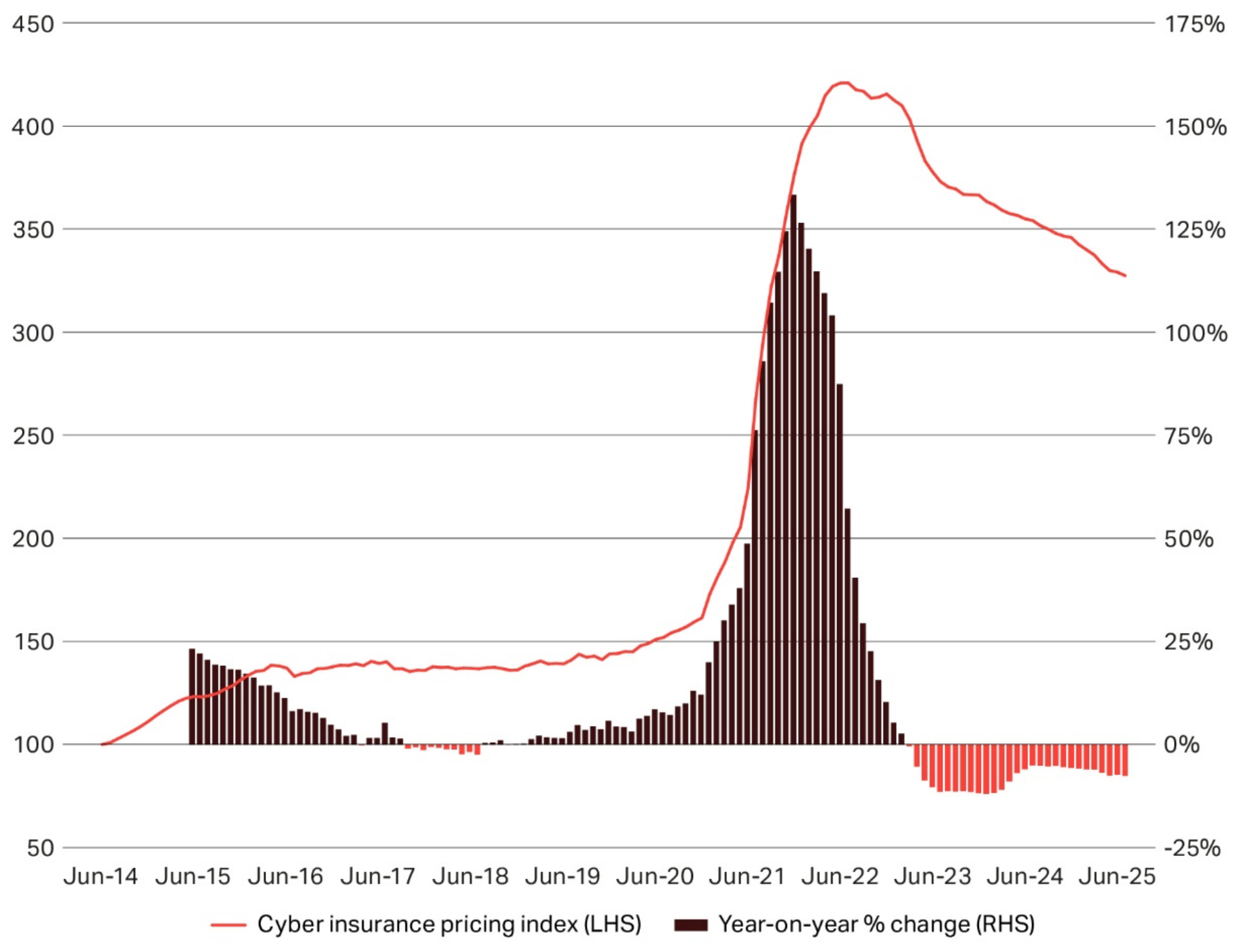

String competition and improved cyber hygiene has reduced costs since 2022, [according to Howden.]()

Colorado State University

Average Cost of Cyber Insurance

Quotes and rate filings from major insurance companies in Connecticut.

Cyber Liability Limit	Deductible	Average Annual Insurance Premium
$1,000,000	$10,000	$1,588
$500,000	$5,000	$1,146
$250,000	$2,500	$739

The higher the number of sensitive records or financial transactions stored, the higher your company's insurance premiums will be.

Most frequent cyber insurance claims
Hacking claims: when a hacker breaks into your company's computer network and steals data,
Ransomware attacks: your company's data or critical software is threatened unless you pay a ransom.
Phishing attacks induce your employees to disclose login credentials to hackers
Employee negligence claims can arise from something as simple as an employee losing a laptop that contains sensitive customer or employee data.

https://advisorsmith.com/data/average-cost-of-cyber-insurance/

Cost of Cyber Insurance

- Data from from BlackFire Cyber Insurance, April, 2021
- I have plotted the Coverage Limits and Annual Premiums for several industries.
- Observations?

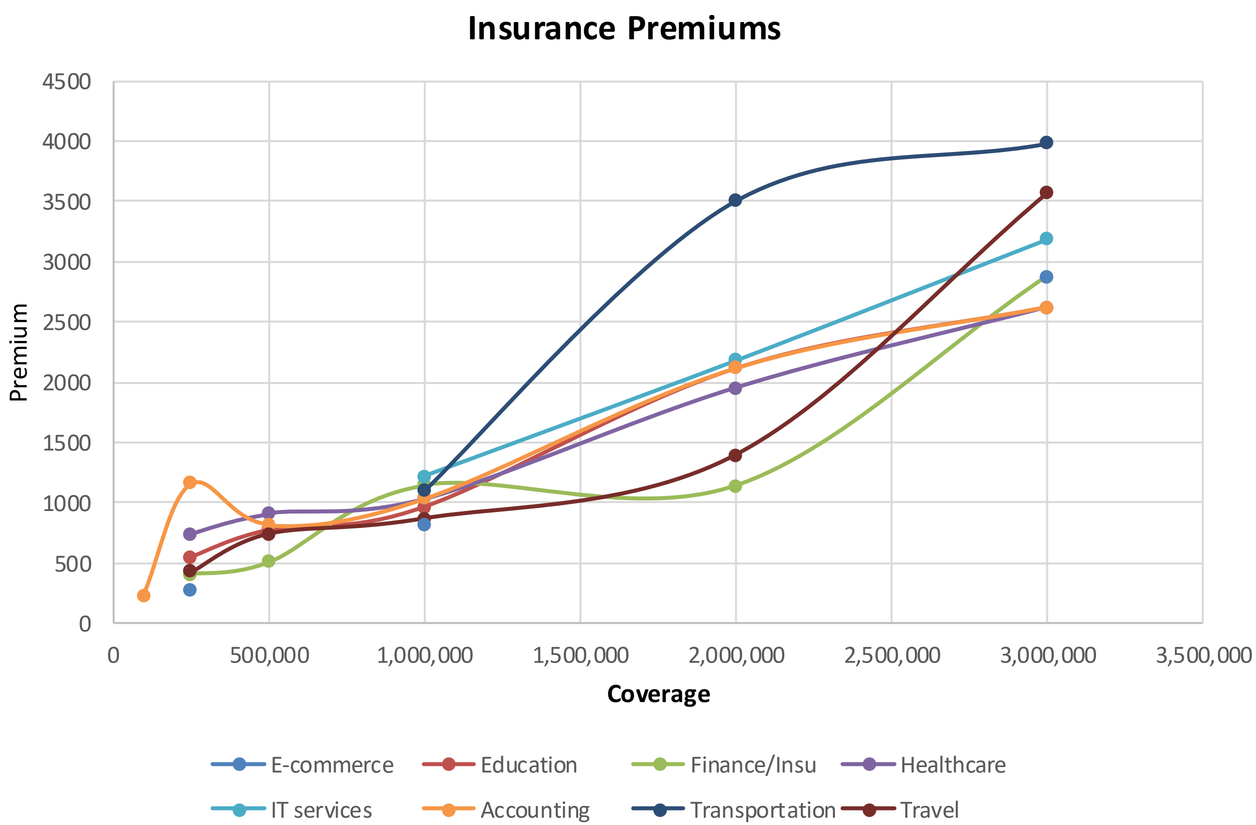

Colorado State University

Loss ratio for U.S. standalone cyber policies – 2016 to 2025

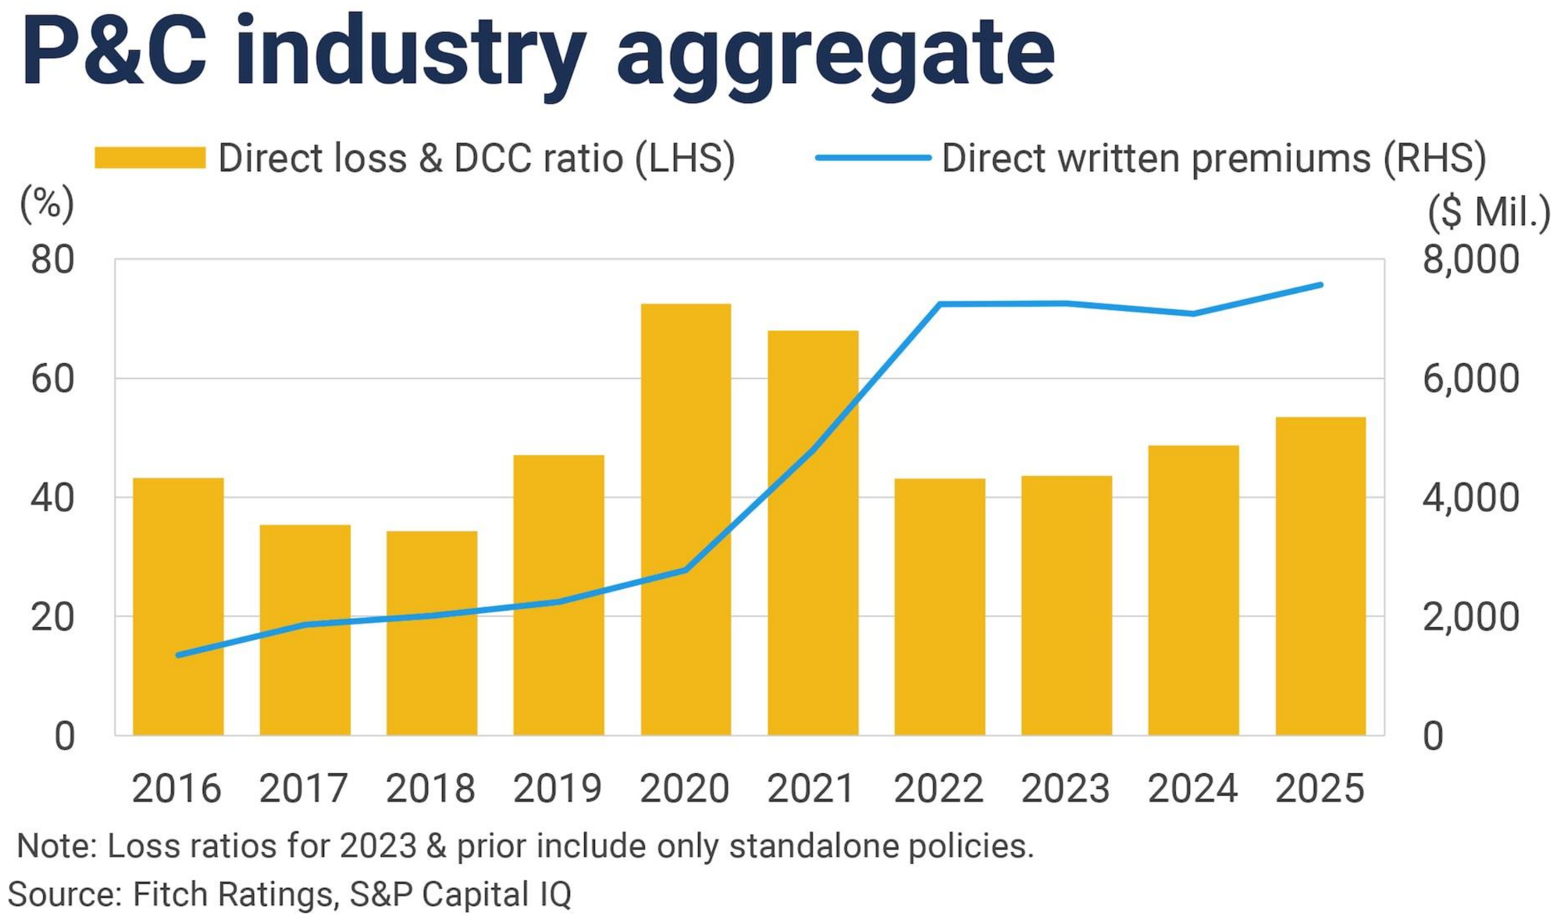

https://www.theinsurer.com/cyber-risk/news/fitch-us-cyber-insurance-market-swung-to-7-written-premium-growth-in-2025-2026-06-03/

Colorado State University

Quantitative Security

Colorado State University

Yashwant K Malaiya

CS 559

Probabilistic Perspective

CSU Cybersecurity Center

Computer Science Dep

Perspectives: Deterministic vs Probabilistic

You are a security person in your organization. What is your security perspective?

- You have done everything required. Thus, there is noting to worry about. **Deterministic view.**
- There is really no way to predict what will happen tomorrow. **Fatalist view.**
- Assess your risks probabilistically. Allocate your budget to reduce the risk as much as possible. **Probabilistic view.**

Discuss presentations

- Next two lectures
- Tuesday: on-campus
- Thurs: online presentations
- Slides must be posted by Monday 5 PM
- Peer review form will be available soon

Probabilistic Methods

- Much of this may be a review of probability and statistics you have taken elsewhere.
- We cannot predict exactly when something will fail, but we can calculate the probability of a failure, and what can be done to reduce that.
- This is similar to what insurance industry does: they may not know when a person will die, but they can compute life-expectancy of someone who is say, 45 years old, and maintains an ideal weight.

Probabilistic Methods: Overview

- We can have concrete numbers even in presence of uncertainty.

Topics:

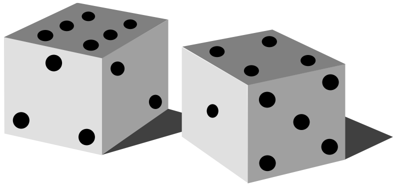

- Probability
  - Disjoint events
  - Statistical dependence
- Random variables and distributions
  - Discrete distributions: Binomial, Poisson
  - Continuous distributions: Gaussian, Lognormal, Beta, Exponential
- Stochastic processes
  - Markov process
  - Poisson process

Basics

- Probability of an event A

$$P\{A\} = \frac{n}{N}$$

if A occurs n times among N equally likely outcomes.

- Probability is a number between 0 and 1.
- Ex: Roll of a die

$$P\{odd\} = \frac{3}{6} = 0.5$$

- If more information is available, probability of the same event changes. If we know die is *loaded*, perhaps

$$P\{odd\} = 0.6$$ is possible.

Basics Concepts

- Prob. Of union of two events:
  - $P\{A \bigcup B\} = P\{A\} + P\{B\} - P\{A \bigcap B\}$

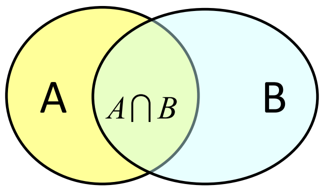

- Ex: Roll of a die

$$P\{outcome\ even \bigcup outcome \leq 3\}$$

$$= P\{even\} + P\{\leq 3\} - P\{even \bigcap \leq 3\}$$

$$= \frac{3}{6} + \frac{3}{6} - \frac{1}{6} \overset{\cup}{=} \frac{5}{6}$$

- If A and B are disjoint, i.e. if $A \bigcap B = \varphi$ (i.e. empty set),

$$P\{A \bigcup B\} = P\{A\} + P\{B\}$$

- $P\{\overline{A}\} = 1 - P\{A\}$

Conditional Probability

- Conditional probability

$$P\{A \mid B\} = \frac{P\{A \cap B\}}{P\{B\}} for P\{B\} > 0$$

P{A|B} is the probability of A, given we know B has happened.

- If A and B are independent, P{A|B}= P{A}. Then

$$P\{A \cap B\} = P\{A\}P\{B\}$$

- Example: A toss of a coin is independent of the outcome of the previous toss.

Conditional Probability

- If A can be divided into disjoint A_i, i=1,..,n, then

$$P\{B\} = \sum_i P\{B \mid A_i\} P\{A_i\}.$$

- Example: A chip is made by two factories A and B. One percent of chips from A and 0.5% from B are found defective. A produces 90% of the chips. What is the probability a randomly encountered chip will be defective?
- P{a chip is defective} = (1/100)x0.9 + (0.5/100)x0.1

  =0.0095 i.e., 0.95%

Bayes' Rule 1763

- Conditional probability

$$P\{A \mid B\} = \frac{P\{A \cap B\}}{P\{B\}} for P\{B\} > 0$$

P{A|B} is the probability of A, given we know B has happened.

- Bayes' Rule

$$P\{A \mid B\} = \frac{P\{B \mid A\} P\{A\}}{P\{B\}} for P\{B\} > 0$$

$$P\{A|B\} = \frac{P\{B|A\}P\{A\}}{P\{B|A\}P\{A\}+P\{B|\sim A\}P\{\sim A\}}$$

Bayes' Rule Example

P{A|B} is the probability of A, given we know B has happened.

- Bayes' Rule

$$P\{A \mid B\} = \frac{P\{B \mid A\} P\{A\}}{P\{B\}} \; for \, P\{B\} > 0$$

- **Example:** A drug test produces 99% true positive and 98% true negative results. 0.5% are drug users. If a person tests positive, what is the probability he is a drug user?

$$P\{DU \mid P\} = \frac{P\{P \mid DU\} P\{DU\}}{P\{P \mid DU\} P\{DU\} + P\{P \mid nDU) P\{nDU\}}$$

$$= \quad ?$$

Bayes' Rule

Example: A drug test produces 99% true positive and 98% true negative results. 0.5% are drug users. If a person tests positive, what is the probability he is a drug user?

$$P\{DU|P\} = \frac{P\{P|DU\}P\{DU\}}{P\{P|DU\}P\{DU\}+P\{P|nDU\}P\{nDU\}}$$

$$= \frac{0.99\times0.005}{0.99\times0.005+0.02\times0.995} = 0.1991$$

= 19.91%%

Person is a drug user and test result is positive: True Positive
P{P|DU} = 0.99 (99%)

Person is not drug user and test result is negative: True Negative
P{N|nDU} = 0.98 (98%)

P{P|nDU} = 1-0.98 = 0.02

Frequentist statistics:

- probability as the limit of the relative frequency of an event after many trials

Baysian statistics:

- based on the Bayesian interpretation of probability - expresses a degree of belief in an event. The degree of belief may be based on the results of previous experiments, or past personal experiences.

Bayes' Rule: Posterior Probability

- Implications of Bayes' rule:

$$P\{A \mid B\} = \frac{P\{B \mid A\} P\{A\}}{P\{B\}} \; for \; P\{B\} > 0$$

- P{A} represents prior probability, when we did not know about B.
- P{A|B} represents posterior probability, after we know B.
- Probability changes if we know more!

- Average age of student = 23 years
- Average age of student| a graduate student = 33 years

40

Bayes' Rule: OJ Simpson trial Example

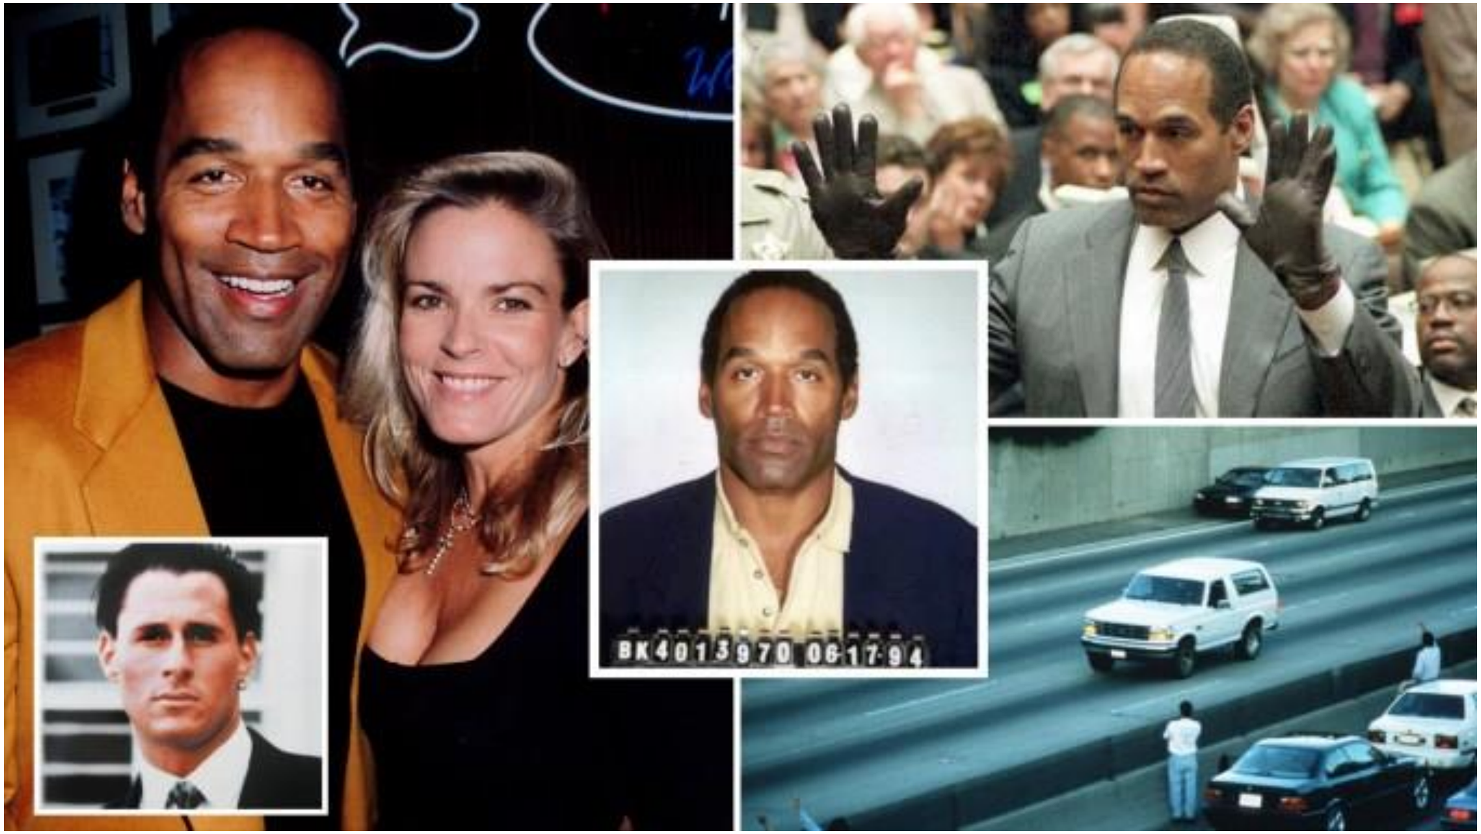

Bayes' Rule: OJ Simpson trial Example

OJ Simpson Trial: There was a prior belief of guilt. There was a blood match. What is the updated belief.

- Given Information on Blood Test (T+/T-)
  - True Positive Rate (Sensitivity): P(T+ | Guilty)=1
  - True Negative Rate (Specificity): P(T- | Innocent)=.9957 $\Rightarrow$ P(T+| Inn)=.0043
- Suppose you have a prior belief of guilt: P(G)=p*
- What is "posterior" probability of guilt after seeing evidence that blood matches: P(G | T+)?

$$P(T+) = P(T^+G) + P(T^+I) = P(G)P(T^+ \mid G) + P(I)P(T^+ \mid I) =$$

$$= p^*(1) + (1-p^*)(.0043)$$

$$P(G \mid T^+) = \frac{P(T^+G)}{P(T^+)} = \frac{P(G)P(T^+ \mid G)}{P(T^+)} = \frac{p^*(1)}{p^*(1)+(1-p^*)(.0043)} = \frac{p^*}{.9957p^* + .0043}$$

B.Forst (1996). "Evidence, Probabilities and Legal Standards for Determination of Guilt: Beyond the OJ Trial", pp. 22-28

Bayes' Rule: Example

$$\text{Prior Probability of Guilt}: \qquad P(G) = .10 \Rightarrow$$

$$P(G \mid T^{+}) = \frac{.10(1)}{.10(1) + .90(.0043)} = \frac{.10}{.10387} = .9627$$

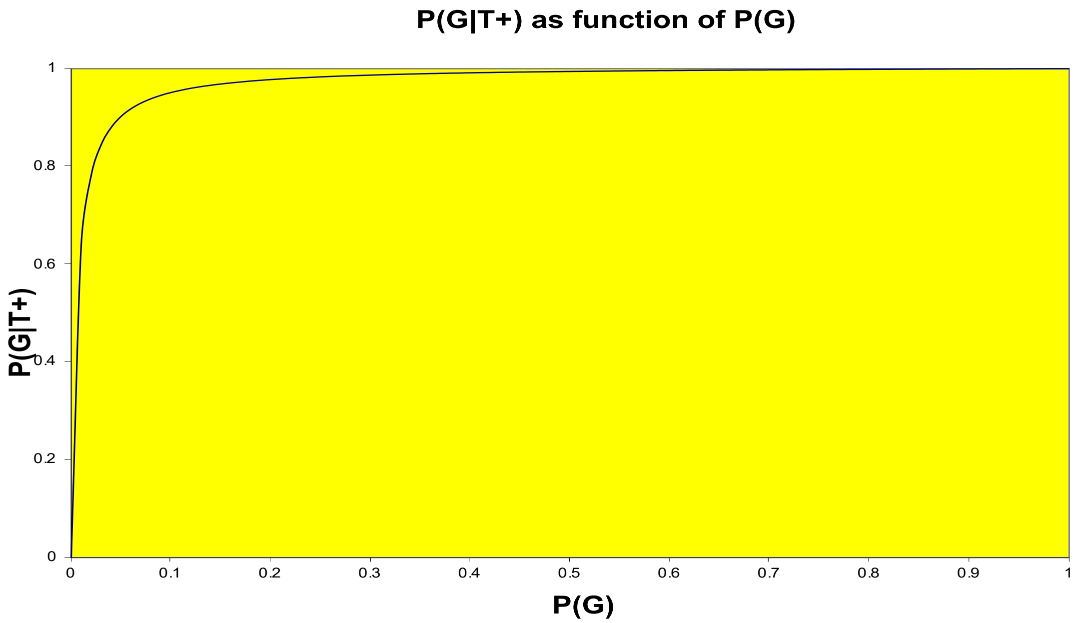

Even if the prior probability of guilt is low, positive test outcome makes it almost certain.

Binary Classification Problem

Binary classification problem

- **There are no perfect tests for diseases.**
- **There is no intrusion detection algorithm that is perfect.**

	Disease +	Disease -
Test +ve	TP	FP
Test –ve	FN	TN

We may want to know

- If the person has the disease, what is the prob test is positive?
- If the person does not have the disease, what is the prob test is indeed negative?
- If the result is positive, what is the prob it is true?
- If an intrusion detection algorithm claims it has detected an intrusion, what is the probability there is actually no intrusion?

Collect available past data to estimate these.

Confusion Matrix

	Disease +	Disease -
Test +ve	TP	FP
Test –ve	FN	TN

- Sensitivity = TP/(TP+FN)   also TPR true pos rate
  - If the person has the disease, what is the prob test is positive?
- Specificity = TN/(FP+TN)   also TNR true neg rate
  - If the person does not have the disease, what is the prob test is indeed negative?
- Precision = TP/(TP+FP) PPV positive predictive value
  - If the result is positive, what is the prob it is true?
  - Ex: TP= 580, FP = 21, FN = 308, TN = 105
    Precision = 580/(580+21) = 0.965
- Several other measures used. For example
  - F1 score = 2TP/(2TP + FP + FN)

RT-PCR s	Disease +	Disease -
Test +ve	580	21
Test –ve	308	105

Example: Intrusion Detection

If an ID scheme is more sensitive, it will increase false positive rates.

- Ex Car alarm

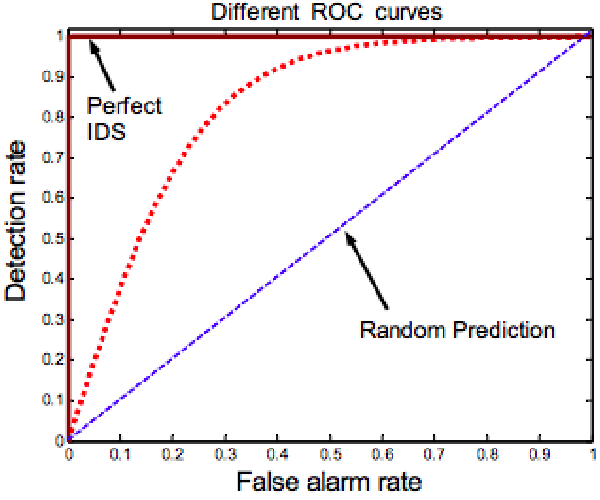

Figure 2-5. ROC Curves for different intrusion detection techniques

- Detection rate (True Positive rate, sensitivity) vs False Positive Rate. We want detection rates to be higher for the same false positive rate.
- Area under the ROC curve is a good measure of the ID scheme.

"ROC" **receiver operating characteristic** we don't care why.

Intrusion Detection A Survey, Lazarevic, Kumar, Srivastava, 2008