LEBANESE MATHEMATICS DAY 2026

The Mathematics of Risk

Actuarial Science, AI, and the Future of Prediction

Presented by Dr. Re-Mi Hage

- Today we study risk as something measurable, not mysterious.
- We will move from simple probability to AI-based risk prediction.
- Every concept will be linked to real life examples.

In one word, what comes to mind when you hear the word “Risk”?

OPENING

Why Risk Is Everywhere

- Students face academic, career, financial, and health risks.
- Companies face operational, cyber, market, and reputational risks.
- Governments face public health, disaster, and economic risks.
- Mathematics gives a common language for comparing very different risks.

One language, many domains

Health

Finance

Cyber

Insurance

RISK IS INTERDISCIPLINARY

Which risk do you think is hardest to predict?

Mathematics as the Language of Uncertainty

- Probability describes whether events occur.
- Statistics learns patterns from observed data.
- Optimization searches for the best decision.
- Modeling translates reality into mathematical structure.

Core tools

Probability
Statistics
Optimization
Modeling

MATHEMATICS IN ACTION

Match one tool to one real-life decision.

OPENING

Workshop Roadmap

- Part 1 — Basic risk maths: probability, expected value, distributions.
- Part 2 — Actuarial science: risk pooling, expected loss, and pricing.
- Part 3 — Risk management: mitigation, transfer, and monitoring.
- Part 4 — Data science: from raw data to a model you can test.
- Part 5 — Artificial intelligence: how algorithms learn to predict risk.
- Part 6 — Professional judgement: what AI cannot replace.
- Part 7 — Applied workshop: five ways to estimate claim severity.
- Part 8 — Closing: key takeaways, careers, and discussion.

From idea to decision

Risk

Model

Prediction

Action

EIGHT CONNECTED PARTS

PART 1

BASIC RISK MATH

Probability, expected value, and the shape of losses.

What Is Risk?

- Risk is uncertainty about an outcome that matters.
- A risk event is something that may or may not happen.
- A loss is the negative consequence if the event happens.
- A model helps us estimate and compare risks.

Risk event

May happen

Has impact

Can be modeled

START WITH DEFINITION

What is the risk event in car insurance?

Expected Loss = Probability × Impact

- Probability answers: how likely is the event?
- Impact answers: how serious is the loss?
- Expected Loss = Probability × Loss Severity, i.e. $E[X] = \Sigma\ x\ p(x)$.
- It is a long-run average — no single policyholder pays exactly E[X].
- Running example: one motor policy, 20% claim chance, $1,500 average loss.

$$E[X] = P \times L$$

Expected Loss

P(claim) = 0.20

Average loss = $1,500

$E[X] = 0.20 \times \$1{,}500 = \$300$

$0.8 \times \$0 + 0.2 \times \$1{,}500 = \$300$

CORE FORMULA · PRICING FOUNDATION

3% chance of a $20,000 loss. What is the expected loss?

Distribution of Losses

- A distribution shows all possible losses and their probabilities.
- Most claims may be small, while a few are very large.
- Insurance data often has skewness and heavy tails.
- Choosing a distribution is a modeling assumption.

Loss distribution

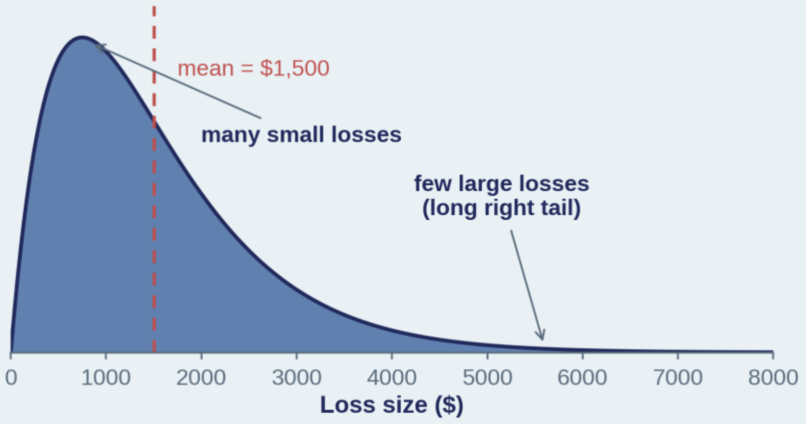

SHAPE MATTERS

Why are insurance losses rarely normally distributed?

Mini-Question: Which Is Riskier?

- Risk A: 50% chance of losing $100.
- Risk B: 1% chance of losing $10,000.
- Both have expected loss of $50? Check carefully.
- Riskier depends on average, variability, and ability to absorb loss.

A/B

Compare risks

A: frequent small loss

B: rare severe loss

Same? different?

RISK APPETITE

Vote: choose A or B, then justify your choice.

PART 2

ACTUARIAL SCIENCE

Turning uncertain individual risk into priceable cost.

What Is Actuarial Science?

- Actuarial science applies mathematics to evaluate risk.
- It combines probability, statistics, finance, economics, and data science.
- Actuaries estimate future losses and design fair, sustainable prices.
- They also assess solvency, reserves, and risk management strategy.

Actuary toolkit

Probability
Statistics
Finance
Modeling

MATHEMATICS OF INSURANCE

What is the difference between an actuary and a data scientist?

Insurance as Risk Pooling

- Many policyholders pay a premium into a pool.
- Only some policyholders will experience a claim.
- The pool pays claims using collected premiums.
- The law of large numbers makes the average loss more predictable.

Pooling

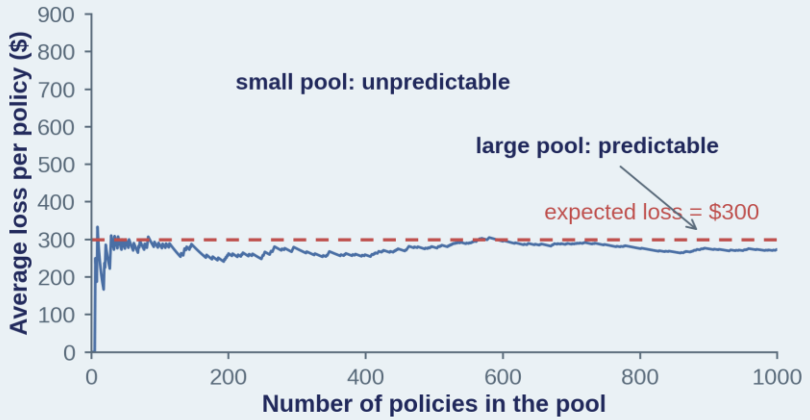

LAW OF LARGE NUMBERS

Expected Loss Formula

- Expected Loss = E[Frequency] × E[Severity].
- Occurrence models estimate the probability of at least one claim.
- Frequency models estimate the number of claims.
- Severity models estimate the cost if a claim occurs.

$\lambda \times \mu$

Frequency × Severity

How many?

How large?

Total expected cost

ACTUARIAL CORE

Hands-on: 0.20 expected claims × $1,500 severity = ?

Premium Calculation: Pure Premium and Loading

- Pure premium covers the expected claim cost.
- Loading covers expenses, uncertainty, profit, and capital cost.
- Final premium must be fair, competitive, and sustainable.
- Too low: insolvency risk. Too high: customers leave.

Premium

Same policy: expected cost 300→*price*375

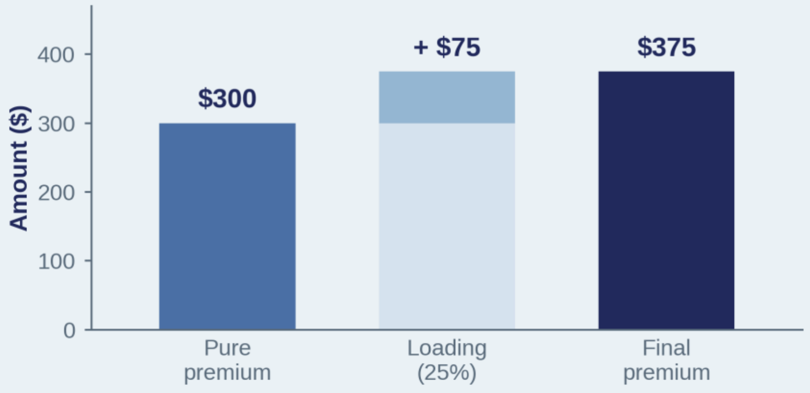

PRICE RISK CAREFULLY

Example: Motor Claim Prediction

- Target: will a driver report a claim this year?
- Features: age band, car type, region, past claims, mileage.
- Model output: estimated probability of claim.
- Decision: pricing, prevention, fraud monitoring, or underwriting review.

Motor risk profile

Driver data

Vehicle data

Claim probability

EXAMPLE WORKFLOW

List three variables you would include and explain why.

PART 3

RISK MANAGEMENT

Identify, measure, predict, and mitigate.

What Is Risk Management?

- Risk management is the process of dealing with uncertainty.
- It begins before the loss occurs, not after.
- It combines measurement, prediction, prevention, transfer, and monitoring.
- Good risk management turns mathematics into action.

Risk cycle

Identify

Measure

Mitigate

Monitor

CONTINUOUS PROCESS

Give one example of preventing a risk instead of paying for it later.

Risk Mitigation

- Risk mitigation is any action taken to reduce how often a loss occurs, how much it costs when it does, or both.
- Mitigation reduces probability, severity, or both.
- Examples: training, safety systems, deductibles, limits, diversification.
- Some mitigation is behavioral; some is technical or financial.
- A good model helps target mitigation where it matters most.

Mitigation levers

Reduce P

Reduce loss

Transfer risk

ACTIONABLE RISK

Why do insurance contracts use deductibles?

Risk Transfer Through Insurance

- Insurance transfers financial consequences from policyholder to insurer.
- The insurer pools many risks and charges premiums.
- Reinsurance transfers part of insurer risk to another insurer.

Transfer

Policyholder

Insurer

Reinsurer

SHARING LOSSES

PART 4

DATA SCIENCE

From raw data to models that forecast losses.

From Raw Data to Risk Prediction

- Data science begins with observations and ends with decisions.
- Raw data must be cleaned, checked, and organized.
- Models learn relationships between predictors and outcomes.
- Prediction must be interpreted carefully before action.

Data pipeline

Collect
Clean
Model
Decide

RISK ANALYTICS

What makes insurance data messy?

Features and Target Variable

- Features are the information used by the model.
- The target variable is the outcome to predict.
- For classification, target may be claim/no claim.
- For regression, target may be loss amount.

X→Y

Feature → target
X: inputs
Y: outcome
Model learns link

SUPERVISED LEARNING

Identify X and Y in “predict claim severity from policy data.”

Training and Testing a Model

- Training data is used to fit the model.
- Testing data is held aside to evaluate performance.
- Good performance on training data alone is not enough.
- The goal is generalization to future unseen cases.

Split data

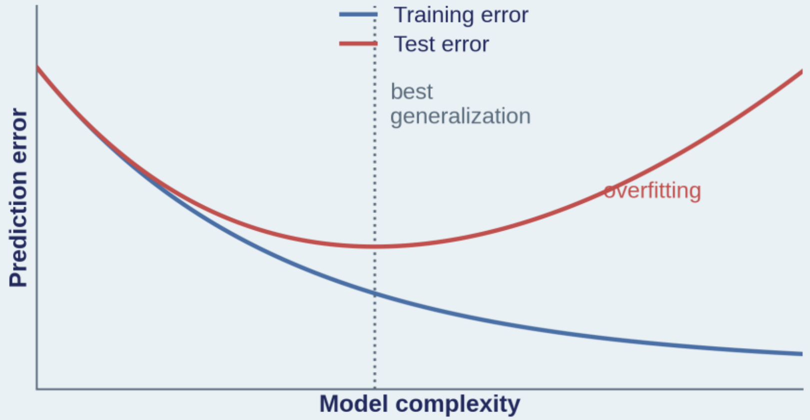

AVOID OVERFITTING

Why not train on 100% of the data?

PART 5

ARTIFICIAL INTELLIGENCE

How algorithms learn to predict risk.

What Is Artificial Intelligence?

- AI refers to systems that perform tasks requiring human-like intelligence.
- In risk, AI is often used for prediction, classification, anomaly detection, and decision support.
- AI is not magic; it is mathematics, data, and computation.
- Its usefulness depends on data quality and responsible use.

AI in risk

Predict

Detect

Explain

Support

NOT MAGIC

Where have you seen AI used in daily life?

What Is Machine Learning?

- Machine learning lets models learn patterns from examples.
- Instead of writing every rule, we fit a function from data.
- The model improves by reducing prediction error.
- Supervised learning needs known outcomes for training.

f(x)

Learn a function

Input X

Output Y

Fit f

PATTERN LEARNING

Neural Networks

- A neural network combines many simple weighted calculations.
- Inputs pass through layers to produce a prediction.
- Weights are adjusted during training.
- Neural networks can capture nonlinear patterns but need careful validation.

Network structure

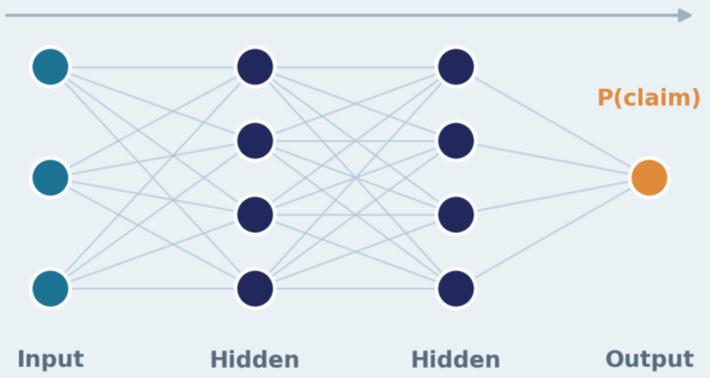

FLEXIBLE MODEL

Why might a neural network overfit small data?

Explainable AI: Why Did the Model Decide This?

- Risk decisions affect people and organizations.
- Users need to understand important drivers of predictions.
- Explainability helps detect errors, bias, and unrealistic patterns.
- A black-box prediction is not enough for responsible risk management.

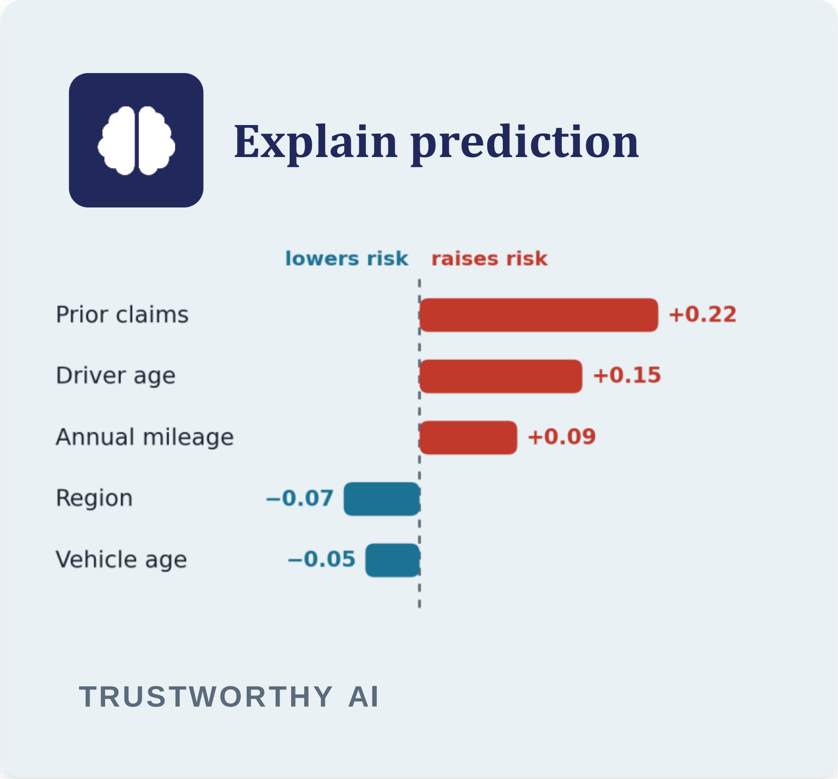

Why is explanation important in insurance pricing?

Worked Example: Who Will Claim Next Year?

- Policyholder H-2031 — age 27, no chronic condition, 0 visits last year.
- Policyholder H-2032 — age 58, diabetic, 11 visits, 2 claims in past 2 years.
- Policyholder H-2033 — age 41, no chronic condition, 1 small claim last year.
- Same three rows — who claims next year? Let the model score them.

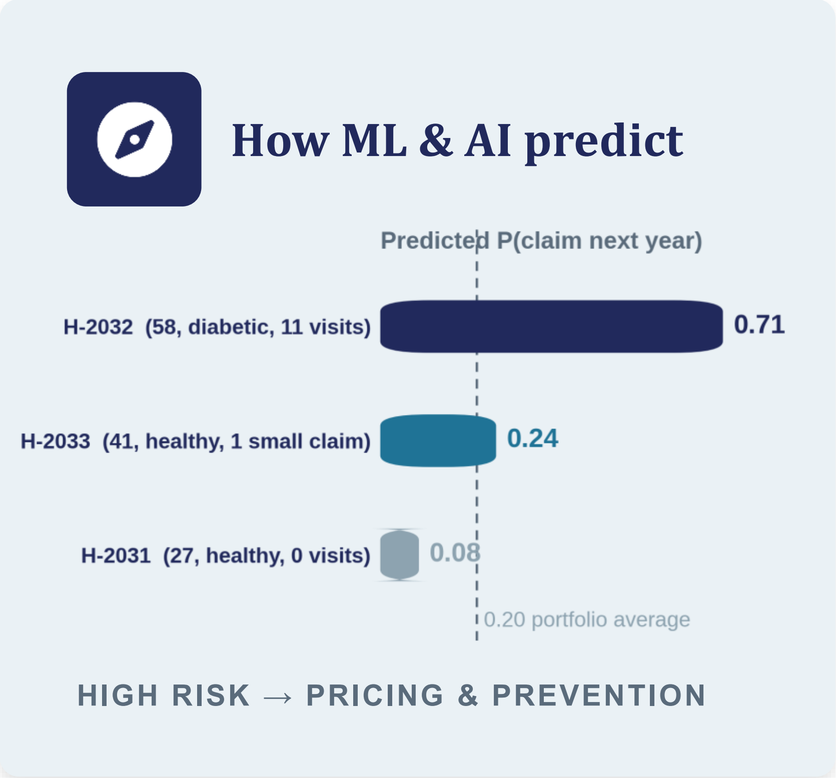

Raw data tells the story — the model only scores what it is given.

Example: AI for Claim Prediction

- Claim prediction is a classification problem: claim or no claim next year.
- Features may include age, region, chronic conditions, past utilization, and plan type.
- Model output is a claim probability for each policyholder.
- Predictions feed pricing, reserving, and prevention — with human oversight.

Prediction workflow

Policy data

Score P(claim)

Price & prevent

Actuary reviews

HUMAN + AI

Should AI set the price alone, or should an actuary review it?

Comparing Models: Who Predicts Claims Best?

- Same task, same data, five methods: will this policyholder claim next year?
- AUC measures ranking quality: 1.00 is perfect, 0.50 is a coin flip.
- Logistic regression, random forest, and a neural network all land near 0.75; boosting edges ahead at 0.76.
- The data sets the ceiling — better features beat fancier algorithms.

ROC-AUC by model

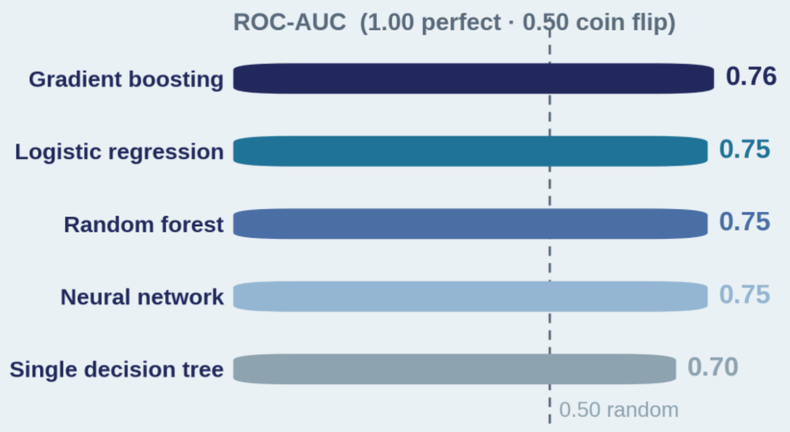

DATA SETS THE CEILING

Why don’t the fancier methods win by more?

Why Gradient Boosting Was Selected

- Two ways to be wrong: a false negative (FN) misses a real claimant; a false positive (FP) flags someone who never claims.
- The missed claimant is the costly error — that risk goes unpriced; a false alarm mainly costs review effort.
- At the same cutoff, boosting missed the fewest claimants (1,450 FN) without inflating false alarms (2,600 FP).
- Model selection weighs the errors that matter — not AUC alone.

FN vs FP by model

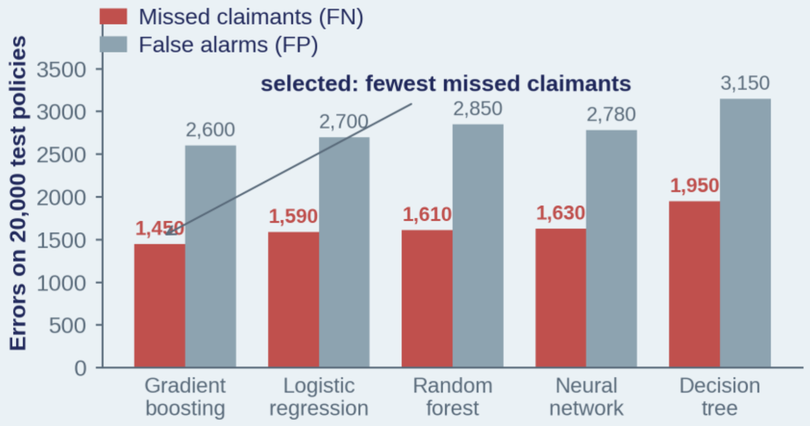

CHOOSE THE ERROR YOU CAN AFFORD

Which error is worse for the insurer — and which for the customer?

A High AUC Can Still Mislead

- If only 1 in 5 policyholders claims, predicting “no claim” for all is 80% accurate — and finds none.
- AUC ranks policyholders well across all thresholds — but the decision is made at one cutoff.
- What counts is recall (claimants found) and precision (flags that are real) — and they trade off.
- High AUC with low recall, or high recall with low precision — always read both, not one.

Read three numbers

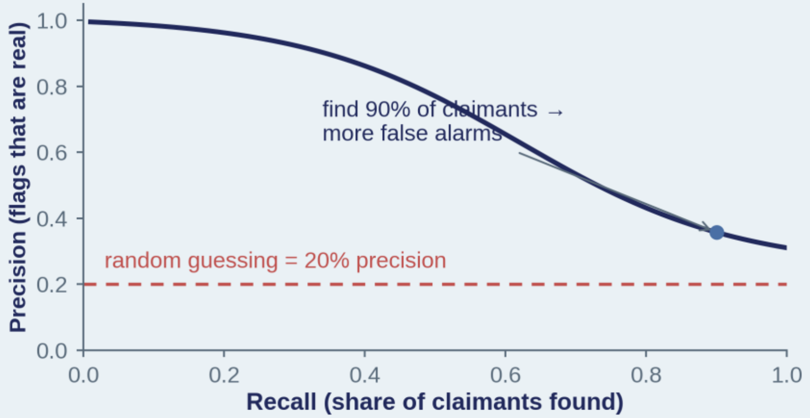

ONE METRIC HIDES; THREE REVEAL

A model is 80% accurate but finds no claimants — is it any good? Tie back to the 20% claim rate.

PART 6

PROFESSIONAL JUDGEMENT

The human layer that AI cannot replace.

AI Automates Calculation, Not Judgement

- AI produces outputs — predictions, scores, drafts — at scale and speed.
- But context, interpretation, ethics, and communication stay human.
- A model optimises a metric; only a person can own the decision.
- The more autonomous AI becomes, the more critical judgement becomes.

Humans remain responsible for

Context

Interpretation

Ethics

Decisions

AI PRODUCES; PEOPLE ANSWER FOR IT

Would you sign a report you cannot explain?

Why AI Cannot Replace the Actuary

- AI optimises an objective; it does not choose the right objective.
- It cannot set assumptions or judge whether data deserve trust.
- It cannot be held accountable to a regulator, board, or policyholder.
- It cannot be challenged — or explain itself — to real stakeholders.

What stays human

Accountability

Assumptions

Fairness

Trust

AI CHANGES HOW ACTUARIES WORK — NOT WHY THEY MATTER

The actuary's competitive advantage will increasingly be judgement.

Estimating a Claim Reserve: ML / AI

- Machine-learning estimate — input is structured columns: age, region, payments to date.
- A GLM or gradient-boosting model fits frequency × severity.
- Output: a point reserve, e.g. $12,400, with explainable drivers.
- Stable, auditable, regulator-ready — but blind to anything off-table.

AI / GenAI estimate

Reads notes & reports

Extracts new features

Updates as claim evolves

Powerful but opaque

SAME CLAIM, RICHER STORY — IF THE DATA IS JUDGED FIT

The actuary decides which estimate to trust, and why.

Data Judgement: Reading the Story in the Data

- Raw data is noise until it is cleaned, checked, and questioned.
- Good features turn columns into a story: who claims, when, and why.
- AI consumes data; the actuary judges whether it deserves to be consumed.
- More data does not automatically mean better decisions.

Key data questions

Relevant & reliable?

Complete & representative?

Hidden biases?

Fit for this purpose?

GARBAGE IN, GOVERNED OUT

What Should Never Be Fully Automated

- Data selection — is this the right data, and is it fair?
- Model choice — does the model suit the problem and the regulation?
- Assumption setting — which expert judgements drive the result?
- Interpretation & communication — what it means, and who owns it.

The common foundation

Technical expertise
Ethical principles
Professional standards
Public accountability

ALL FIVE REST ON PROFESSIONAL JUDGEMENT

What AI Changes — and What It Doesn't

- AI changes the work: more data, more automation, faster decisions.
- AI does not change accountability, professional judgement, or trust.
- Thrive in the AI era by governing it, challenging it, and translating it.
- The future actuary is not the best modeller — but the best judge.

Tomorrow morning

Learn AI & prompting

Use AI weekly

Challenge outputs

Stay accountable

AI CHANGES HOW ACTUARIES WORK,
NOT WHY THEY MATTER

Trust is not generated by AI — it is earned through judgement.

PART 7

APPLIED WORKSHOP

Estimating claim severity: probability, GLM, machine learning, deep learning.

Simulated portfolio of 8,000 claims

One question, five ways to answer it

1. **The one-size-fits-all guess**
   STEP 1 · PROBABILITY
   Fit one distribution to every claim and predict its mean — the same number for everyone. Gives quantiles and tail risk, but no covariates.

2. **The points system**
   STEP 2 · LINEAR REGRESSION
   Start from a base price, then add or subtract for each fact about the driver. Simple — but the error shape is wrong.

3. **The percentage system**
   STEP 3 · GLM — GAMMA, LOG LINK
   Same idea, but multiply: sports car × 1.83, city driver × 1.36. Multiplicative relativities — the industry pricing standard.

4. **The 20-questions machine**
   STEP 4 · MACHINE LEARNING
   Trees ask their own questions and discover non-linearities and interactions nobody specified.

5. **The tiny artificial brain**
   STEP 5 · DEEP LEARNING
   A neural network learns its own representation of risk — the technology behind modern AI, in miniature.

Each rung relaxes one assumption of the rung before it — and all five are scored the same way.

Insurance is a guessing game — played with your money

1 You pay a little

Every driver pays a premium — a fixed amount per year — whether they crash or not.

2 Some people crash

The company pays the repair bill. Some bills are $500. A few are $50,000.

3 The price must be fair

To set your premium, the company must guess, in advance, how expensive your crash would be.

The whole workshop in one sentence: we will teach the model to look at a driver — their age, their car, where they live — and guess the size of their repair bill **before it happens.**

THE DATA

A portfolio of 8,000 settled claims

One row per crash: the final bill, plus six facts about the driver and the car.

Variable	Range	What it means
DriverAge	18–85 years	U-shaped: the young and the senior crash expensively
VehicleCat	Sedan · SUV · Sport · Utility	Sport carries the largest surcharge
VehiclePower	50–260 kW	Small city car → powerful sports car; log-linear surcharge
Region	Urban · Suburban · Rural	Urban repairs are the dearest
BonusMalus	50–200	Driving record: clean history → many past accidents
VehicleAge	0–25 years	Older cars are cheaper to repair

Severity ~ Gamma, mean μ multiplicative in the factors — plus a quadratic age curve and a young × Sport interaction that no linear model can see.

The typical bill

Mean severity	$3,700
Median	$2,740
Largest claim	$50,000+
Train / test	6,000 / 2,000

MEAN ≫ MEDIAN — A HEAVY TAIL

One golden rule: 2,000 bills are hidden away. Every method studies the other 6,000, then guesses the hidden ones.

Severity is right-skewed — and that drives everything

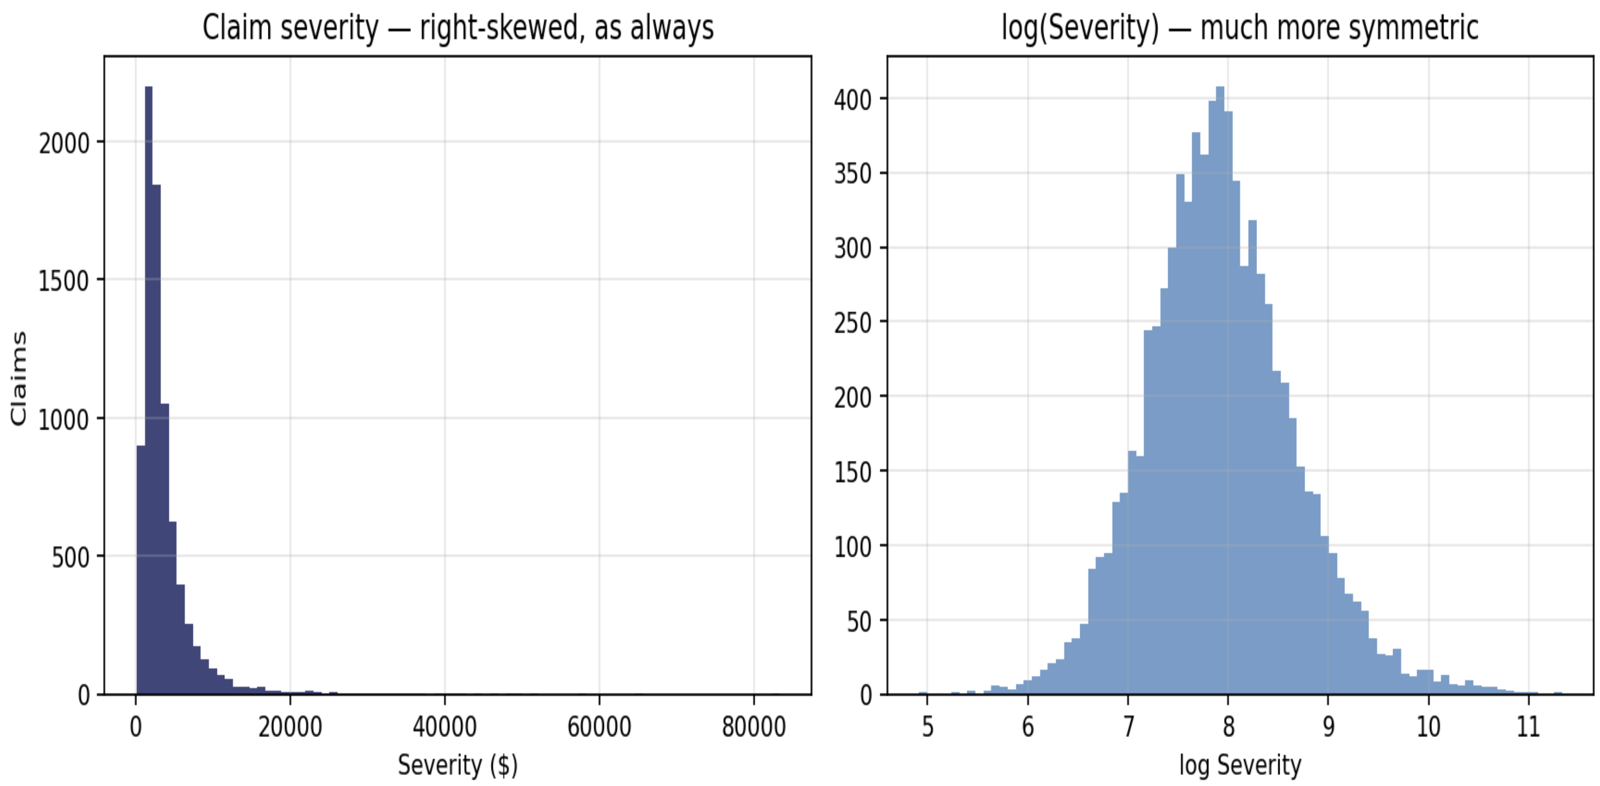

Heavy right tail. Mean ($3.7k) sits far above the median ($2.7k); the top claims dominate total cost.

Log helps. On the log scale the distribution is nearly symmetric — a hint that effects are multiplicative.

Consequence: methods assuming symmetric errors start on the wrong foot.

Fit a distribution, predict its mean

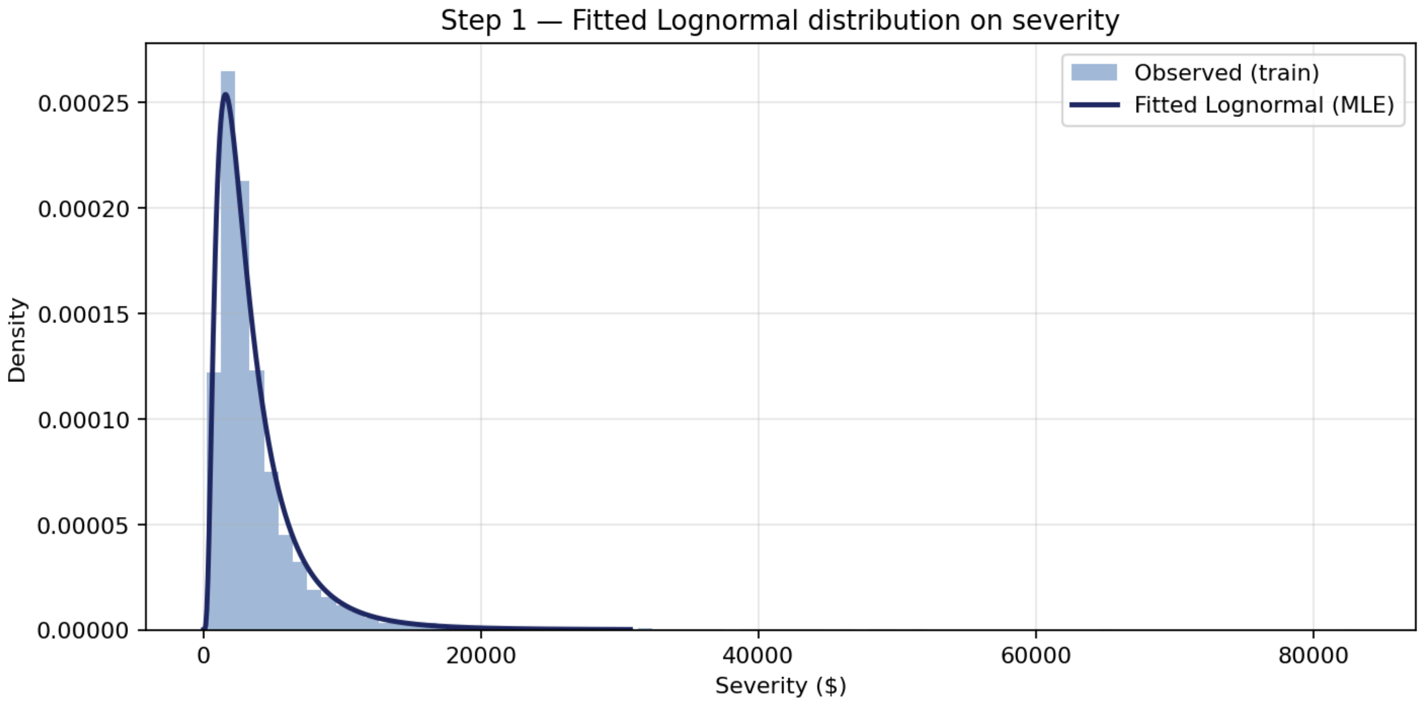

Distribution	AIC	KS
Lognormal ✓	**108,719**	**0.032**
Gamma	109,982	0.093
Weibull	110,543	0.096
Pareto	118,777	0.381

Maximum likelihood on the training set; select by AIC. Every claim gets the same prediction: the fitted mean.

R^2 = 0 by construction — yet this rung prices the portfolio, sets tail capital and reinsurance layers.

Answers "how much in total, and how bad can it get?" — not "which claim will be expensive?"

The points system — risk factors, added up

Instead of one number for everyone, build the guess like a restaurant bill:

Base price	$3,000
Sports car	+ $900
Lives in the city	+ $400
Older car (cheaper parts)	– $300
Your guess	**$4,000**

$$Y = \beta_0 + \beta_1 \cdot \text{Age} + \beta_2 \cdot \text{Power} + \dots + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2)$$

- Symmetric, constant-variance errors — but severity is skewed and its variance grows with the mean.
- Additive dollar effects — a Sport surcharge should scale with the claim, not add a flat amount.

Out-of-sample

Flat mean (Step 1)	R^2 0.000
OLS, raw dollars	R^2 0.226
OLS on log Y	R^2 0.268
Average miss	$1,882

RMSE $3,871 → $3,407 → $3,312; the miss falls $320 per bill.

COVARIATES BUY A REAL JUMP

The classic patch is to regress log Y — but exp(prediction) targets the median, so Duan's smearing correction is needed.

Keep the linear predictor, fix the distribution

$$Y \sim Gamma(\mu, \varphi), \quad \log \mu = \beta^T x$$

Gamma errors: right-skewed, variance $\propto \mu^2$ — constant coefficient of variation, exactly what claims data show.

Log link: e^β is a multiplicative relativity — a rating table a regulator can read. And predictions target the mean on the dollar scale: no retransformation bias, ever.

Result: $R^2 = 0.245$, RMSE $3,363 — at par with log-OLS, but every coefficient is defensible.

Factor	Relativity e^β
Sport (vs SUV)	**× 1.83**
Urban (vs Rural)	**× 1.36**
Suburban (vs Rural)	× 1.10
Utility (vs SUV)	× 0.82
Sedan (vs SUV)	× 0.78

Its limit: non-linearities and interactions must be specified by hand — and our data secretly holds both.

Let the computer play twenty questions

The idea. Stop writing rules. Give the computer the 6,000 study bills and let it invent its own questions:

"Is the driver under 26?"
→ yes: "Is it a sports car?"
→ **yes: expect a big bill**
→ no: "Do they live in the city?"
→ … and so on, hundreds of times

Each chain of questions is called a decision tree. Modern methods grow hundreds of trees and average their answers (a "random forest") or let each new tree fix the previous ones' mistakes ("gradient boosting").

The magic: nobody tells the computer which questions matter. It discovers, on its own, that young drivers in sports cars are a special — and expensive — combination.

THE SCORE: **$1,442 average miss** — a huge jump. Roughly $450 closer per bill than the percentage system. The next slide shows why.

The 20-questions machine finds what nobody told it

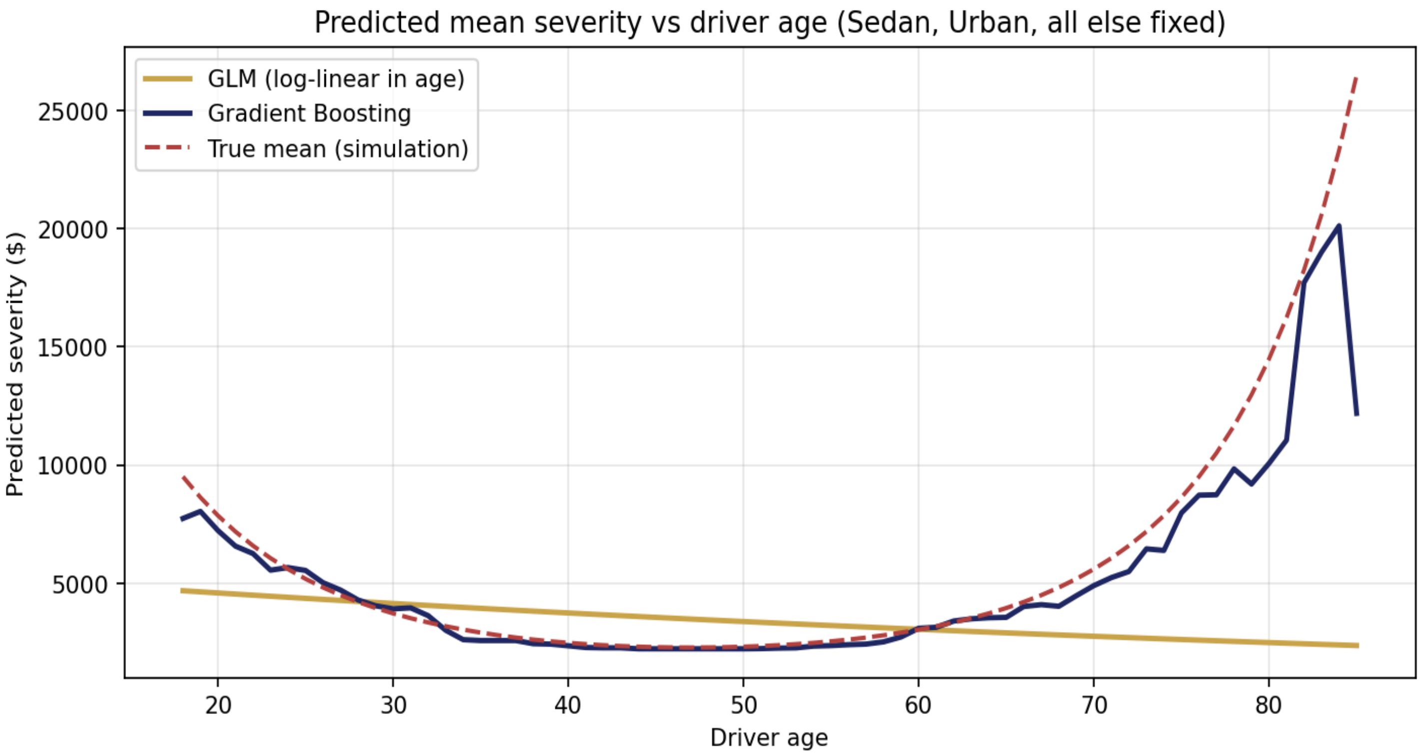

Driver age 18–85 against the expected repair bill: the truth (red dashes), the GLM (gold), the trees (navy).

What the trees found

Random Forest	R^2 0.618
Gradient Boosting	R^2 0.575
RMSE (best)	$2,394
Average miss	$1,442

The GLM is log-linear in age — it cannot bend twice. The trees traced the U with no hints.

ACCURACY WITHOUT A RATE CARD

The cost: no one-page rating table. Insight returns through feature importances and partial-dependence plots.

A tiny brain that builds its own risk factors

- A neural network is a stack of very simple calculators.
- Each layer turns the raw facts — age, car, region — into better clues.
- Nobody designs those clues; the network invents them while practising.
- Here: a small 64–32–16 network on log-severity, stopped early.

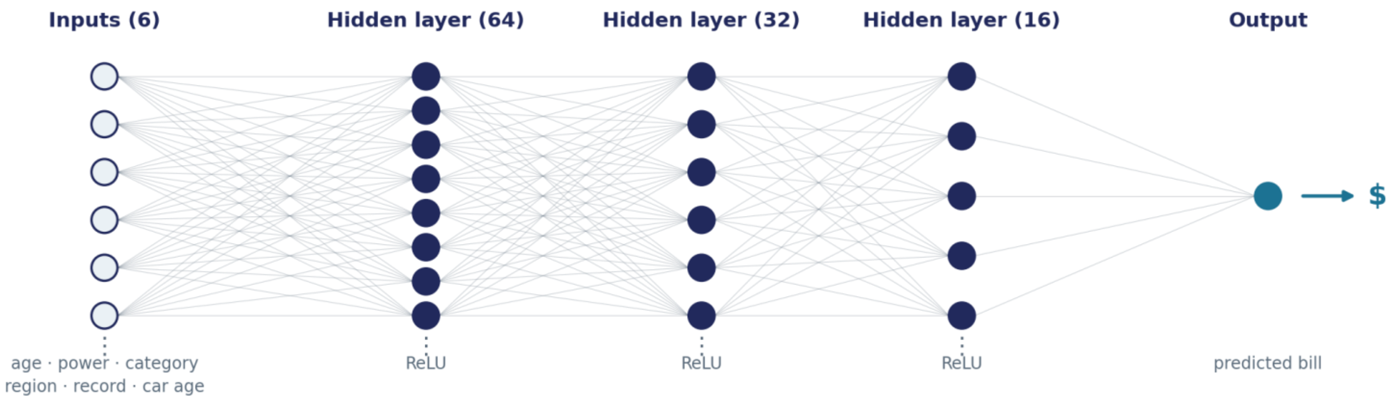

Every arrow carries a weight; training nudges all of them until the guesses stop improving.

Best of the ladder

R^2	0.666
RMSE	$2,236
Average miss	$1,412

Best of the day — $30 a bill better than the 20-questions machine.

ON SMALL TABLES, THE TREES TIE

Deep learning's edge appears at scale — telematics, images, written claim notes; not on a small table of numbers.

The ladder, measured on the same 2,000 claims

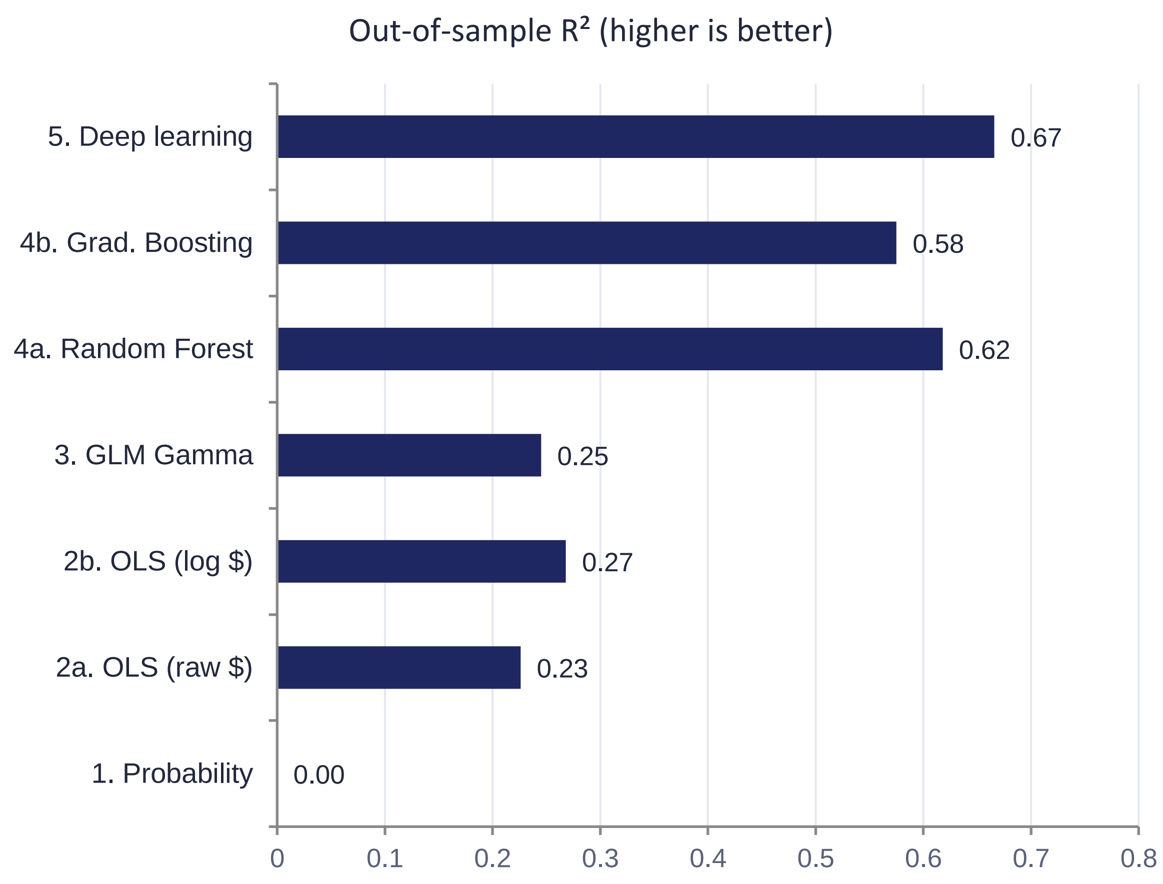

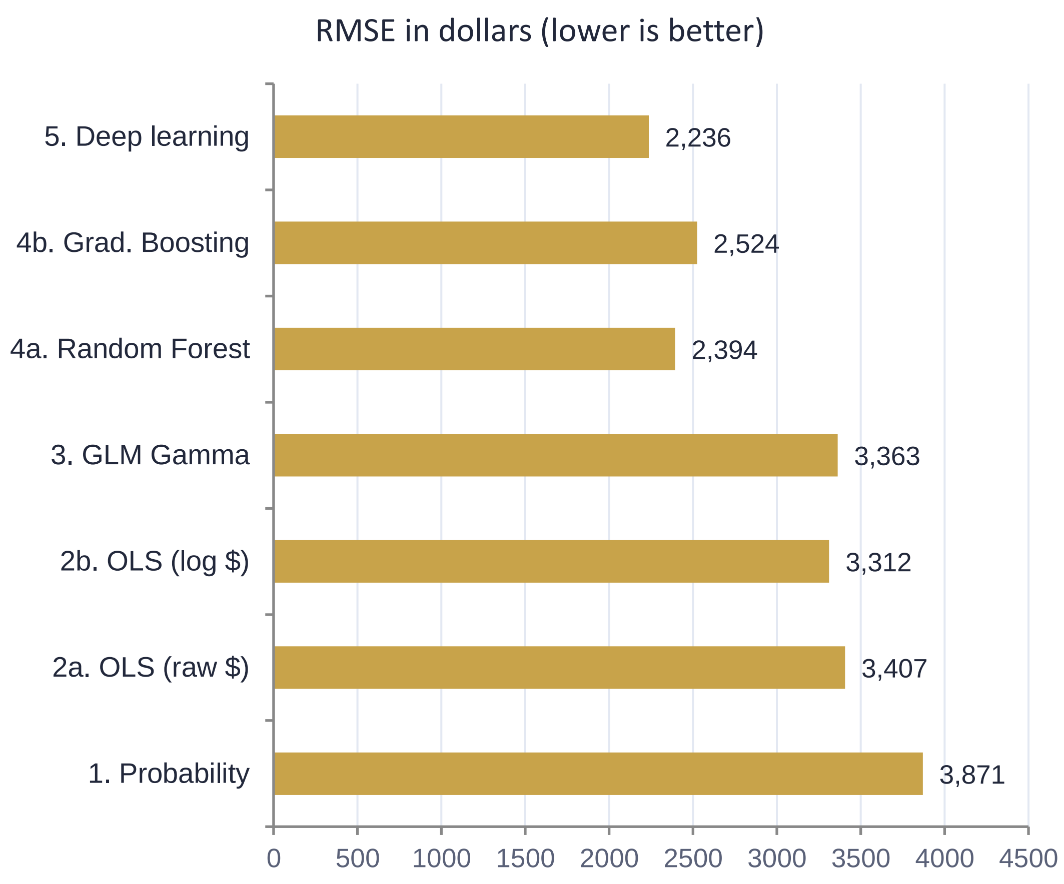

From flat mean to deep learning: RMSE falls 42% — $3,871 → $2,236 per claim.

What each rung assumes, gives — and costs

Rung	Core assumption	Best when…	Watch out for…
1. Probability	One distribution fits all claims	The one-size-fits-all guess: totals, tail risk, reinsurance layers	Blind to individual risk (R^2 0.00)
2. Linear regression	Additive effects, normal errors	The points system: a quick, simple first covariate model	Adds when life multiplies (R^2 0.27)
3. GLM — Gamma, log	Multiplicative effects, Gamma errors	The percentage system: prices a regulator can read	Only the patterns you name (R^2 0.25)
4. Machine learning	Nothing on the functional form	The 20-questions machine: hidden interactions matter	Needs XAI tools to explain (R^2 0.62)
5. Deep learning	None; the layers learn features	The artificial brain: huge or messy data — telematics, text	Data-hungry and opaque (R^2 0.67)

The skill is matching the rung to the question, the data and the audience — not knowing the fanciest method.

PART 8

CLOSING

Skills, careers, and the ideas that tie it together.

Key Takeaways

- Risk can be measured using probability, severity, and distributions.
- Actuarial science turns risk measurement into pricing, reserving, and solvency decisions.
- AI and data science improve prediction but require validation and explanation.
- Optimization and finance help choose better decisions under constraints.

Four messages

Measure

Predict

Explain

Decide

MATHEMATICS CREATES IMPACT

Final reflection: which application surprised you most?

INTERACTIVE

Final Discussion and Questions

- What risk would you like to model in your own project?
- What data would you need to predict it?
- Which mathematical tool would you start with?
- How would you explain the result to a non-mathematician?

YOUR TURN

Ask

Challenge

Connect

Create

Thank you