Quantitative Skills, Probability and Mental Maths

Interview preparation summary · IN track

IN THIS SUMMARY

- Probability and expected value
- Numbers under pressure

INT.NQ.1.a · Probability and expected value

Calculate an expected loss from a small table of outcomes and use it to reason about price, deductibles and limits.

Expected value is the probability-weighted average of the possible outcomes. For an insurer it is usually called the expected loss: multiply each possible claim amount by its probability and add the results. It is not the most likely outcome, because most policies have no claim at all, and no single policy will ever cost exactly the expected loss. It matters because across a large group of similar policies the average cost settles near it, so it is the starting point for a price. The premium is then the expected loss divided by the target loss ratio, plus whatever the pricing method adds for risk. A deductible removes the first slice of every claim and a limit removes everything above a cap, so each lowers the expected loss, and the effect is calculated claim size by claim size, never by subtracting the deductible from the average. Interviewers ask this to check you can do the arithmetic calmly and explain it in one sentence. Common mistakes are ignoring the no-claim outcome, multiplying instead of dividing by the loss ratio, and forgetting that small claims below a deductible disappear completely.

WORKED EXAMPLE: EXPECTED LOSS FOR A SHOP FIRE POLICY AT HARLOW MUTUAL

Harlow Mutual, an invented insurer, covers a small shop's stock against fire. In a year the shop has no loss with probability 94%, a partial loss of 5,000 with probability 4%, a larger loss of 20,000 with probability 1.5% and a total loss of 100,000 with probability 0.5%. Find the expected loss, the premium if the target loss ratio is 65%, and the effect of a deductible of 1,000 on every loss. Convention: one policy-year, ignore tax, expenses and investment income.

Answer: The expected loss is 1,000 and the premium at a 65% target loss ratio is 1,538.46. A 1,000 deductible cuts the expected loss to 940 and the premium to 1,446.15.

> **Interview tip:** Say it as: weight each outcome by its probability, add them up, then divide by the target loss ratio to get a price. Mention that the rare large loss drives much of the answer, and that a deductible is worked out loss by loss.

INT.NQ.1.b · Probability and expected value

Use independence, conditional probability and Bayes' rule, preferably with natural frequencies, to interpret a flag or a signal.

Two events are independent if knowing one has happened does not change the chance of the other, and only then can you multiply their probabilities to get the chance of both. Conditional probability, written P(A given B), is the chance of A among cases where B has happened. It is not the same as P(B given A), and mixing them up is the commonest error in interviews. Bayes' rule links the two and, in practice, is best done with natural frequencies: imagine 10,000 cases, count how many fall in each branch, then divide. The key lesson is the base rate. A fraud model that catches most fraud but also flags a few honest customers will, when fraud is rare, produce flags that are mostly false alarms. Insurers care because the same logic applies to claims triage, underwriting referrals and telematics scores. A good answer states the base rate first, builds the table of counts, gives the percentage, and says what decision follows, usually referral for a second look rather than automatic action.

WORKED EXAMPLE: A FRAUD FLAG AT CALDER MOTOR CLAIMS

Calder Motor Claims, an invented claims team, believes 2% of claims are fraudulent. Its screening model flags 80% of fraudulent claims and also flags 5% of genuine claims. A claim has just been flagged. What is the chance it is fraudulent, and what changes if fraud is twice as common? Convention: take 10,000 claims and treat the model's rates as exact.

Answer: A flagged claim has a 24.6% chance of being fraudulent at a 2% base rate, and 40.0% if fraud is 4%. The flag should trigger review, not automatic rejection.

> **Interview tip:** Always start with the base rate and use counts out of 10,000 rather than formulas. Finish by saying what the business should do with the number.

INT.NQ.1.c · Probability and expected value

Explain why pooling many independent risks makes results more predictable, quantify it with the square-root rule, and say when it fails.

Insurance works because of the law of large numbers: as the number of similar, independent policies grows, the average claim cost per policy settles close to its expected value. The total claims bill does get larger in dollars, but relative to its expected size it gets smaller. For independent policies the standard deviation of the total grows with the square root of the number of policies, while the expected total grows in direct proportion, so the coefficient of variation (standard deviation divided by mean) falls with the square root. In words: the number of claims on a large book behaves roughly like a Poisson count, whose standard deviation is about the square root of its mean, so 400 expected claims typically wobbles by about 20. The argument fails when risks are not independent. A storm, a cyber event or a pandemic hits many policies at once, so adding more policies in the same place does little. Interviewers want the square-root rule, an example, and the caveat about correlation and large single risks.

WORKED EXAMPLE: WHY A BIGGER POOL IS SAFER AT WEXMOOR HOME

Wexmoor Home, an invented insurer, writes identical home policies. Each independently has a 2% chance of a claim in the year and a claim costs exactly 10,000. Compare a pool of 100 policies with a pool of 10,000 policies. Convention: independent policies, use the binomial standard deviation, ignore expenses.

Answer: The coefficient of variation falls from 70% at 100 policies to 7% at 10,000 policies, a factor of the square root of 100, which is 10. This relies on independence, which a common event can break.

> **Interview tip:** Say the rule in one sentence (total risk grows with the square root of the pool, the average grows in proportion) and then volunteer the caveat about correlated events before you are asked.

INT.NQ.2.a · Numbers under pressure

Work with percentage changes, percentage points and successive changes accurately, including splitting a premium change into price and volume.

Under pressure, most percentage errors come from three habits. First, confusing percentage points with percent: a loss ratio moving from 60% to 66% is up 6 points, which is a 10% relative rise. Second, adding changes that should be multiplied: if the average price rises 6% and the number of policies falls 4%, total premium changes by 1.06 x 0.96 - 1, which is +1.76%, not +2%. Third, reversing a change on the wrong base: after a 20% fall you need a 25% rise, not 20%, to get back. Insurers talk in these terms constantly, splitting premium growth into rate (price for the same risk) and volume (more or fewer policies), and comparing claims inflation with rate changes. A good answer says the base, converts each change to a multiplier, multiplies, converts back, and quotes the result with its unit. Always say which comparison you mean, such as year on year or against plan, and whether you are quoting points or percent.

WORKED EXAMPLE: RATE, VOLUME AND CLAIMS INFLATION AT ASHGROVE DIRECT

Ashgrove Direct, an invented motor insurer, had premium of 50m and a loss ratio of 65% last year. This year the average price per policy rises 6% and the number of policies falls 4%. Claims cost per policy rises 5%. Find the new premium, the new loss ratio and what price rise would be needed to recover the lost volume's premium. Convention: all else equal, loss ratio is claims divided by premium.

Answer: New premium is 50.88m (+1.76%), the loss ratio is 64.39% (down 0.61 points), and regaining 4% of lost volume needs a 4.17% increase.

> **Interview tip:** Convert every change to a multiplier, multiply, and say points or percent out loud. Spotting that volume cancels out of the loss ratio shows you understand the structure.

INT.NQ.2.b · Numbers under pressure

Calculate loss, expense and combined ratios, find a break-even loss ratio, and do the arithmetic quickly in your head.

The loss ratio is claims divided by earned premium, the expense ratio is operating costs including commission divided by premium, and the combined ratio is their sum. As a common rule of thumb, a combined ratio below 100% means an underwriting profit and above 100% a loss, though practice varies by whether claims-handling costs and reserve changes are included. The break-even loss ratio is simply 100% minus the expense ratio. Because insurers also earn investment income on premiums held before claims are paid, a combined ratio slightly above 100% can still be profitable overall, which is why you should say what the benchmark ignores. Under pressure, use round-number tricks: 10% of a figure is a decimal shift, 20% is

double that, 5% is half of 10%, and one point of loss ratio on 80m of premium is 0.8m. Interviewers check you can build the three ratios from claims, expenses and premium, name the break-even, and keep the denominator straight, using earned premium for the loss ratio and the same base for expenses.

WORKED EXAMPLE: TARN INSURANCE, A COMMERCIAL PROPERTY BOOK

Tarn Insurance, an invented insurer, has earned premium of 80m, claims incurred of 52.8m, commission of 22% of premium and other expenses of 11% of premium. Investment income on its float is 4% of premium. Find the three ratios, the underwriting result, the break-even loss ratio with and without investment income, and the value of one point of loss ratio. Convention: pre-tax, investment income is stated as a percentage of earned premium.

Answer: Loss ratio 66%, expense ratio 33% and combined ratio 99%, giving an underwriting profit of 0.8m. Break-even loss ratio is 67% (71% including investment income) and each point is worth 0.8m.

> **Interview tip:** Give the three ratios first, then the break-even, then a sentence on what the benchmark leaves out, such as investment income and reserve movements. Use round-number tricks aloud so the interviewer hears the method.

INT.NQ.2.c · Numbers under pressure

Read a small table or chart aloud, use a weighted average rather than a simple one, and make and sense-check a quick estimate.

A numerical interview often hands you a small table and asks what you see. A strong approach has four moves: say what the table shows, including units and period; give the headline pattern; point out one exception or inconsistency; and say what you would ask or do next. Averages need care. When lines of business differ in size, the portfolio loss ratio is the premium-weighted average, equal to total claims divided by total premium, not the simple average of the line ratios. For estimation, round each input to one or two significant figures, multiply, and check the answer against a second route or against what is plausible, for example whether a loss ratio above 100% would be believable for a healthy book. State the rounding and give a range. Also check the denominator: loss ratios use earned premium, not written premium, which matters for a fast-growing book. Interviewers are watching for calm structure, correct weighting and honesty about precision, not for decimal places.

WORKED EXAMPLE: READING PENHALLOW GROUP'S LINES ALOUD

Penhallow Group, an invented insurer, reports earned premium and loss ratios by line: motor 120m at 72%, home 80m at 58%, and liability 50m at 81%. Describe the table aloud, find the group loss ratio, and sense-check it. Convention: premium is earned premium; round claims to one decimal place.

Answer: The group loss ratio is 69.32% (173.3m on 250m), about 1 point below the misleading simple average of 70.33%. The quick estimate of 68.8% confirms it.

> **Interview tip:** Open with what the table is, then the pattern, then one exception, then your next question. Use total claims divided by total premium whenever you combine ratios.