Catastrophe Modelling, Climate and Emerging Risks

Interview preparation summary · IN track

IN THIS SUMMARY

- Catastrophe
- Emerging risks

INT.NE.1.a · Catastrophe

Describe how a catastrophe model turns hazard, exposure and vulnerability into an insured loss, and explain why the quality of exposure data matters.

A catastrophe model estimates the insured loss from natural events such as windstorm, flood, earthquake and wildfire. It has linked parts. The hazard part describes how often events of different strength occur and how intense they are at each location, for example wind speed or flood depth. The vulnerability part says how much damage that intensity does to a building of a given type and age, as a percentage of its value. The financial part then applies the policy terms, deductibles and limits, to give the insured loss. The exposure data is what the insurer supplies for every risk: location, value, construction, occupancy and terms. A model is only as good as that data. A poor address places a building in the wrong flood zone, an understated sum insured understates the loss in proportion, and an unknown construction type is given an average damage curve. Interviewers ask because candidates often treat output as fact. A good answer says it is an estimate with uncertainty and that data quality should be tested. The usual mistake is to blame the model when the exposure data is the weak link.

WORKED EXAMPLE: ONE WINDSTORM EVENT OVER THREE SITES AT HARROWGATE PROPERTY

Harrowgate Property, an invented insurer, insures three sites in one coastal windstorm zone. For one modelled event the damage ratios (damage as a share of value) are 6% for masonry and 14% for timber frame, and the model gives an unknown-construction site the average of 9%. Site A is recorded at 40m, masonry. Site B is 25m, timber frame. Site C is 15m, construction unknown. Each site has a 1m deductible. Find the modelled insured loss. Then suppose Site A's true rebuild value is 48m and Site C is in fact timber frame, and find the correct loss and the size of the error. Figures are in dollars, for one event, ignoring tax.

Answer: The modelled loss is 4.25m and the corrected loss is 5.48m, so poor values and an unknown construction type understated the loss by 1.23m, about 22.4%.

> **Interview tip:** Name the three parts (hazard, vulnerability, financial), then say that the exposure data feeds all three. Give one concrete data error and its direction, and say you would test the data before trusting the answer.

INT.NE.1.b · Catastrophe

Explain return period, average annual loss and exceedance probability, and calculate them from a small event loss table.

A catastrophe model produces a list of possible events, each with a loss and an annual rate. From it come the two numbers interviewers expect you to know. Average annual loss, or AAL, is the long-run average yearly cost: multiply each event's annual rate by its loss and add them up. It is the pure cost of catastrophe cover before any loading for expenses, uncertainty or capital. The exceedance probability is the chance that the loss in a year reaches or goes above a given amount. The return period is one divided by that probability, so a 1-in-100 year loss has about a 1% chance of being exceeded in any year. It does not mean it happens once every 100 years, and it can happen in consecutive years. An occurrence exceedance curve looks at the largest single event in the year, while an aggregate curve adds all events in the year, so it sits higher. The usual mistake is to use AAL alone. Two portfolios can have the same AAL and very different tails, and capital and reinsurance are sized from the tail.

WORKED EXAMPLE: AN EVENT LOSS TABLE FOR FENWICK RE

Fenwick Re, an invented reinsurer, has a five-event table for a windstorm portfolio. Event rates per year and losses are: 0.10 at 5m, 0.04 at 20m, 0.02 at 45m, 0.01 at 80m and 0.004 at 150m. Treat events as independent. Find the AAL, the probability that a year has a loss of at least 45m, and the matching return period. Then say what the 80m figure means. Money is in dollars, for one year.

Answer: The AAL is 3.60m. A loss of at least 45m has a probability of about 3.3% in a year, a return period of about 30 years. A loss of at least 80m is about a 1-in-72 year loss.

Interview tip: Say 'one in thirty means about a three per cent chance in any year', never 'once every thirty years'. Give the AAL as the average cost and the exceedance figure as the tail, and say they answer different questions.

INT.NE.1.c · Catastrophe

Distinguish per-occurrence from aggregate terms, and show how accumulation of exposure in one zone drives the loss from a single event.

Accumulation means many insured risks that can be hit by the same event, for example every building in one flood plain, one windstorm footprint or one earthquake zone. Insurers track accumulation by adding up the sums insured in each zone, often with a scenario damage ratio, and set a budget for the loss they are willing to take from one event. This matters because catastrophe loss is driven by where the exposure sits, not by how many policies there are. Two books with the same total sums insured can have very different single-event losses if one is concentrated. Terms also matter. A per-occurrence deductible or limit applies to each event separately. An aggregate deductible or limit applies to the total of all events in a policy year, so it caps the worst year. Interviewers ask about this to see whether you understand why catastrophe business is managed by capacity and not only by price. The usual mistake is to look at premium and loss ratio without asking how much sits in the same place.

WORKED EXAMPLE: ACCUMULATION AND AGGREGATE TERMS AT BRACKENFIELD INSURANCE

Brackenfield Insurance, an invented insurer, has 1,000m of property sums insured, of which 600m is in one river valley and 400m is spread elsewhere. A flood event in the valley has a 4% damage ratio. Compare the loss with a version of the book spread equally over five zones of 200m, where the same event hits one zone at the same damage ratio. Then take a policy with a 2m per-occurrence deductible and a 10m annual aggregate limit, hit in one year by events of 3m, 7m, 6m and 4m, and find the payout. Ignore tax.

Answer: Concentration makes the single-event loss 24.0m instead of 8.0m, and the policy pays 10m in the bad year because the aggregate limit applies after the per-occurrence deductible.

Interview tip: Define accumulation in one sentence, give a zone example with numbers, then separate per-occurrence terms (each event) from aggregate terms (the whole year).

INT.NE.2.a · Emerging risks

Distinguish physical from transition climate risk, and show how a change in event frequency and severity changes expected catastrophe loss and the price needed.

Climate change reaches an insurer in two ways. Physical risk is the change in the weather itself: more frequent or more severe floods, storms, heatwaves and wildfires, plus sea level rise, which raise claims on property, crops and other lines. Transition risk comes from the move to a lower-carbon economy: policy changes, new technology and shifts in demand can cut the value of an insurer's investments in carbon-heavy companies and create new liabilities, such as claims against companies accused of causing harm. For physical risk the key idea is that historical losses may understate the future when the climate is changing, so models are adjusted for trends in frequency and severity, and exposure growth, such as more building on flood plains, is added on top. Because most property contracts last a year, insurers can reprice and change terms at each renewal, but affordability and fairness limit how far they can go. The usual mistake is to treat climate as a distant issue. It already affects this year's price.

WORKED EXAMPLE: A FLOOD BOOK AT THORNFIELD UNDERWRITING

Thornfield Underwriting, an invented insurer, writes flood cover. Its model gives 0.2 damaging events a year with an average loss of 12m. A reviewer proposes a climate adjustment of 10% more events and 8% higher average severity. The catastrophe premium is 4.0m. Find the AAL before and after, the effect on the loss ratio, and the premium needed to keep the loss ratio unchanged. Ignore expenses and tax; figures are for one year.

Answer: AAL rises from 2.40m to 2.85m, up 18.8%, and the premium would need to rise from 4.00m to 4.75m to keep the same loss ratio.

Interview tip: Separate physical from transition risk first, then show that frequency and severity changes multiply. Say the adjustments are assumptions with a range, not forecasts, and that annual contracts let the insurer reprice.

INT.NE.2.b · Emerging risks

Explain why cyber risk accumulates across many insureds, describe war and state exclusions as an idea, and size a systemic event against premium.

Cyber insurance covers losses from attacks and failures of computer systems: business interruption, data breach response, extortion and liability. Unlike a flood, a cyber event has no geography. A fault in widely used software, or an outage at a cloud provider, can hit thousands of insureds at once, so losses that look independent can arrive together. This is called systemic risk or aggregation. Insurers manage it by tracking which insureds depend on the same providers, by limiting dependent business interruption with sublimits and waiting periods, and by testing scenarios against capital. War and state exclusions are an idea meant to remove losses from hostile acts by or on behalf of states, which could be large and correlated. In practice the wording and the question of who is behind an attack are hard to settle, and practice varies. Interviewers ask because cyber is growing fast and has little history. The usual mistake is to judge it on claim frequency alone and ignore the scenario in which many insureds claim together.

WORKED EXAMPLE: A CLOUD OUTAGE AT LARCHMONT CYBER

Larchmont Cyber, an invented insurer, has 500 policies with an average annual premium of 25,000. In a scenario, a cloud provider fails for 3 days and 35% of the insureds depend on it. Each affected insured loses 40,000 a day of income, and the policies have a 12-hour waiting period (half a day). Find the scenario loss and compare it with premium. Then find the loss if dependent business interruption were sublimited to 80,000 per insured. Ignore other claims, expenses and tax.

Answer: The outage costs 17.5m, which is 140% of the 12.5m premium, and the 80,000 sublimit brings it to 14.0m, or 112%.

Interview tip: Contrast cyber with natural catastrophe: the 'zone' is a shared technology, not a place. Give a scenario in numbers, compare it with premium and capital, and name the controls: sublimits, waiting periods and dependency tracking.

INT.NE.2.c · Emerging risks

Apply the tests of an insurable risk to pandemic and liability trends, and show how social inflation erodes the margin of a liability book.

A risk is usually insurable when several conditions hold. The loss is accidental and outside the insured's control. It is definite and measurable. There are many similar exposures, so losses can be pooled. They are not all hit at once. The probability can be estimated well enough to price. And the premium is affordable. Pandemic business interruption is the standard example of a risk that fails the pooling test: one event hits most policyholders together, so insurers typically exclude it or cap it tightly. Liability trends test the pricing condition. Social inflation is the rise in liability claim costs faster than general inflation, driven by more litigation, higher awards, third-party funding of claims and changing attitudes. Liability claims take years to settle, so a book priced on a 4% claims trend can drift badly if the true trend is 9%, and the problem shows first in reserves and later in the loss ratio. Interviewers ask this to see whether you can judge when to decline, cap or reprice. The usual mistake is to treat a short, good run of results as proof that the price is right.

WORKED EXAMPLE: SOCIAL INFLATION IN THE LIABILITY BOOK AT MARLOWE CASUALTY

Marlowe Casualty, an invented insurer, wrote 100m of liability premium priced at a 65% loss ratio, assuming claims severity would rise by 4% a year. Expenses are 30% of premium. Severity actually rises 9% a year for three years. Find how far claims costs have drifted from the price after three years, the loss ratio and the combined ratio, ignoring changes in volume, investment income and tax.

Answer: After three years claims cost 15.1% more than priced, the loss ratio is 74.8% and the combined ratio is 104.8%, so a planned underwriting profit becomes a loss.

Interview tip: List the insurability tests, name the one each example fails, and keep the numbers simple: a small gap in annual trend compounds. Offer options: reprice, cap limits, change terms or decline.

Summary of the Charta interview-preparation topic. Companies and figures in the examples are invented. Draft content, not yet expert reviewed.