

# Quantitative Analyst interview questions

careertips.com.au: common questions and what they're really testing

## What to expect

Quantitative analyst interviews are usually a mixture of technical probing, coding or modelling exercises, and a check that you can explain your work to people who are not quants. Employers want evidence that you understand the finance behind the maths, not just the maths, and that you can be trusted with models that feed live risk and trading decisions.

- **Technical and modelling:** Questions on derivatives pricing, stochastic processes, statistics and numerical methods. Expect to derive or explain something on a whiteboard rather than recite a definition.
- **Programming and data:** Questions on Python or R code structure, SQL, version control and how you handle messy or large market datasets. Some employers set a short take-home or timed coding task.
- **Process:** Questions about how a model gets from research to production: validation, sign-off, monitoring and what happens when it drifts.
- **Behavioural:** Questions about working with traders, risk managers and engineers, disagreements over assumptions, and models that did not perform as expected.
- **Scenario and judgement under pressure:** Questions where you have to reason through a live problem, often with incomplete information and someone waiting on the answer.
- **Stakeholder communication:** Questions on explaining model limitations to a risk committee, auditor or regulator without overstating what the model can do.

Most processes start with a recruiter or hiring manager screen covering your background and motivation, followed by an online maths or coding assessment. From there come two or three technical rounds, often with a quant, a quant developer and sometimes a trader or risk manager, usually including a model or case discussion. A final round with a desk head or senior risk lead tends to focus on judgement, communication and how you would fit the team. Panel interviews are common, and you may be asked to walk through a past model or piece of research in detail.

## 1. Walk me through how you would price a European swaption, and which assumptions you would test first.

### Why they ask

This checks whether you can connect the maths to an actual instrument and whether you know which assumptions carry the most risk.

### How to structure your answer

A structured technical walk-through: instrument and payoff, choice of model, calibration, then the assumptions you would stress and why. Keep it in logical order and flag where you would stop and check rather than push ahead.

### Example answer

*"For a European swaption I would start from the underlying swap: fix the notional, tenor, fixed rate, day count and payment frequency, since those conventions drive the price as much as the model does. For a vanilla European swaption the market convention is to quote a Black implied volatility off a swaption cube, so my first pass would be Black's model, with the forward swap rate and annuity pulled off a properly built discount curve. The assumptions I would test first are the shape of the volatility surface, because a single implied vol ignores skew and smile, and the curve construction, particularly the interpolation between pillars and how I handle the long end where quotes are thin. I would also check the fixing and settlement conventions against the confirm, since a mismatch there is*

*the classic source of a small persistent P&L; break. Where the trade is longer dated or the surface is steep, I would move to a short rate or LMM framework and calibrate to the cube, then compare both results before deciding what goes to the desk.”*

## **2. How do you decide a model is ready to move out of research and into production?**

### **Why they ask**

Employers want to know you understand that a model is not finished when the numbers look right, and that production means other people are relying on it.

### **How to structure your answer**

A process walk-through: state your own bar for evidence, then the validation and sign-off steps, then monitoring. Be concrete about what you would refuse to sign off on.

### **Example answer**

*“I want three things before I would push a model into production. First, evidence that it behaves sensibly out of sample and under stressed market conditions, not just on the data I fitted it to, which usually means backtesting across a few different regimes and comparing against a simpler benchmark. Second, a written record of the assumptions and their known limitations, so that whoever inherits the model knows where it is fragile. Third, tests and version control so a change can be reviewed and rolled back, because a model in production is software as well as maths. I would also agree monitoring thresholds up front with the risk team, so we catch drift early rather than discovering it in a P&L; break. If I could not explain why the model works, or if it only worked with parameters I had tuned hard to the backtest, I would hold it back.”*

## **3. Tell me about a time a model you built did not perform as expected once it was in use.**

### **Why they ask**

This is the standard behavioural question for a quant role, and it tests honesty, diagnostic discipline and how you work with the people who were affected.

### **How to structure your answer**

STAR: situation, task, action, result. Spend most of your time on the action, meaning how you diagnosed the problem and what you changed, and be direct about the outcome.

### **Example answer**

*“I had built a statistical model to flag mispricings in a set of fixed income instruments, and in the first weeks live it was generating far more signals than the backtest suggested. The task was to work out whether the model was genuinely finding something or whether I had a data problem. I started by comparing the live inputs against the same dates in my backtest data and found that a vendor had changed how it handled stale quotes, so my model was flagging what was really just a missing print. I added a staleness filter and a data quality check that ran before the signal logic, and reworked the backtest to use the same inputs the live system actually received. The signal count came back to something the traders trusted, and the data check later caught a second feed issue before it reached the desk. The lasting lesson was that I now treat the data pipeline as part of the model, not something upstream of it.”*

## **4. A trader calls you an hour before the market opens insisting your pricing model is wrong because their P&L; does not tie out, and they want it fixed now. What do you do?**

### **Why they ask**

This is the everyday pressure of the role, and interviewers want to see whether you hold your ground on evidence without becoming obstructive.

### How to structure your answer

A judgement-under-pressure structure: acknowledge and gather facts, form a preliminary view, communicate clearly, then commit to a next step and a time. Show that you separate the urgent question from the correct answer.

#### Example answer

*"I would take the call seriously, because a P&L break usually means something real, but I would not change a production model on the basis of one day's break. First I would ask for the specific positions and the size of the discrepancy, and check whether it is a pricing difference or a booking and convention issue, since those look identical from the trader's screen. I would then pull yesterday's official marks and rerun the model on the affected trades, and if the model output matches the previous day and the market has moved, that points to the market data feed rather than the model. Either way I would give the trader a clear answer within the hour on what I have found so far and when they will have a proper one. If it turned out to be a genuine model issue, I would escalate to the model owner and risk before touching anything, because a change at that hour needs a second pair of eyes."*

## 5. How would you structure a Python codebase for a pricing library so that other quants and developers can maintain it?

### Why they ask

Quant analysts write code that other people have to validate, extend and audit, so employers probe your engineering habits, not just your maths.

### How to structure your answer

A practical walk-through with concrete choices: project layout, testing, version control, dependencies and documentation. Mention trade-offs rather than presenting one right answer.

#### Example answer

*"I would start by separating the core pricing logic from the data access and the reporting layer, so the maths can be tested without a live market data connection. That usually means a small number of modules with clear responsibilities, pure functions where possible, and configuration for curve and volatility inputs rather than hard coded values. I would write unit tests around known analytic results, for example checking a European option price against Black Scholes, and regression tests against a frozen set of market data so a refactor shows up immediately if it changes output. Everything would go through Git with review before merge, and dependencies pinned in a lock file so a laptop and the production environment give the same answer. For documentation, I would keep the model assumptions next to the code and maintain a short README explaining how to run and calibrate the library, because the person who needs it most is usually the one inheriting it."*

## 6. How would you explain the limitations of a model to a risk committee or an auditor who is not a quant?

### Why they ask

Quant analysts sit between the maths and the governance around it, and this question tests whether you can be honest about uncertainty without losing the room.

### How to structure your answer

A communication structure: state what the model does, then what it assumes, then where it fails, using plain language and a concrete example. Finish with what you are doing about it.

#### Example answer

*"I would start with what the model is actually for, in one sentence, so the committee knows what decision it supports. Then I would set out the assumptions in plain terms, for instance that the model assumes we can trade at quoted prices and that correlations between markets hold at the levels we observed over a particular window. I would give a concrete illustration of where that breaks, such as the way correlations move toward one during a liquidity event, which is exactly when we most need*

*the risk numbers to be reliable. I would be clear about what the model is not telling them, rather than leaving them to assume it covers everything. Finally I would set out the mitigations already in place, like stress testing alongside the model, monitoring thresholds agreed with the risk team, and a documented review date. In my experience committees respond better to a well explained limitation than to a confident answer that later turns out to have been oversold."*