Risk and Portfolio Construction

Interview preparation summary · HF track

IN THIS SUMMARY

- Volatility and risk-adjusted return
- Diversification and tail risk
- Sizing and portfolio construction

INT.RP.1.a · Volatility and risk-adjusted return

Compute a standard deviation and scale volatility between daily, weekly, monthly and annual horizons.

Volatility is the standard deviation of returns, usually quoted per year. For a sample of returns, square the deviations from the mean, divide the sum by the number of observations minus one for a sample, and take the square root. To change horizon, variance grows with time when returns are independent, so volatility scales with the square root of time: annual volatility is monthly volatility times the square root of 12, or daily volatility times the square root of 252 on a common trading day convention (conventions vary, with 250 and 260 also used). Going the other way, divide by the square root. The rule fails when returns are autocorrelated: smoothed or illiquid assets show positive autocorrelation, so scaling a monthly figure understates true annual risk. Volatility targeting sets exposure as target volatility divided by the asset's volatility, so exposure falls when volatility rises. Interview point: state the day count and say that square root scaling assumes independent returns.

WORKED EXAMPLE: MONTHLY RETURNS TO ANNUAL VOLATILITY AND A TARGET

An interviewer gives five monthly returns for a strategy: 3%, minus 1%, 5%, 1% and minus 3%. Treat them as a sample from a longer history, assume returns are independent, and use 12 months in a year. What is the strategy's annual volatility, and what exposure gives a 10% annual volatility target?

Answer: The monthly standard deviation is 3.16%, the annual volatility is about 10.95%, and an exposure of about 91% of capital hits a 10% volatility target.

> **Interview tip:** Say the day or month count you are using and that square root scaling assumes independent returns, before the interviewer asks. The interviewer listens for the sample divisor of observations minus one and for the square root of time rather than a straight multiple.

INT.RP.1.b · Volatility and risk-adjusted return

Calculate a Sharpe ratio, compare strategies with it and explain its limits.

The Sharpe ratio is the average return in excess of the risk-free rate divided by the standard deviation of returns, both over the same period. It measures reward per unit of total volatility, which makes strategies of different risk comparable. A monthly ratio is annualised by multiplying by the square root of 12, not by 12. Leverage scales return and volatility together, so it leaves the ratio unchanged when borrowing is at the risk-free rate, and a lower volatility strategy can be levered up to match a riskier one. The limits matter in interviews: standard deviation treats upside and downside alike, ignores skewness and fat tails, and is flattered by smoothed or illiquid returns and by strategies that sell options, which earn small steady gains until a large loss arrives. A short track record gives a very noisy estimate. Sortino ratios and drawdown measures are common complements. Interview point: always add that a high Sharpe does not rule out tail risk.

WORKED EXAMPLE: COMPARING TWO FUNDS ON SHARPE RATIO

An interviewer says: Fund A earned a mean monthly return 0.5% above the risk-free rate, with a monthly standard deviation of 2%. Fund B returned 7% in a year with annual volatility of 7%. The annual risk-free rate is 2%, and returns are independent. Which fund has the better Sharpe ratio, and what would you add to your answer?

Answer: Fund A has an annualised Sharpe ratio of about 0.87 against 0.71 for Fund B, so it looks better per unit of volatility, subject to checks on tails and smoothed returns.

> **Interview tip:** Annualise by the square root of 12, never by 12, and then volunteer the limits of the ratio. The interviewer listens for the sentence that a high Sharpe ratio does not rule out tail risk.

INT.RP.2.a · Diversification and tail risk

Compute two-asset portfolio volatility from weights, volatilities and correlation, and explain the limits of diversification.

For two assets, portfolio variance equals the squared weight times variance for each, plus twice the product of the weights, the correlation and the two volatilities. At a correlation of one, portfolio volatility is just the weighted average, so there is no benefit. Lower correlation cuts volatility, and a correlation of minus one allows a perfect hedge, with zero risk when each weight times its volatility matches. With many equal volatility assets that share an average pairwise correlation, volatility falls toward a floor equal to the volatility times the square root of the correlation, so diversification cannot remove common risk. Correlations are unstable and tend to rise in crises, which is when diversification is most needed. Correlation measures linear co-movement and says nothing about tail dependence. Interview point: show the formula, then say that the correlation you use in calm times may understate the correlation you will face in stress.

WORKED EXAMPLE: TWO-ASSET VOLATILITY AT THREE CORRELATIONS

An interviewer says: asset A has annual volatility 10% and asset B has 20%. A portfolio holds them 50/50. What is the portfolio volatility if the correlation is 0, 0.5 and 1? What does the pattern tell you about diversification, and what would you say about a crisis?

Answer: Portfolio volatility is 11.2% at zero correlation, 13.2% at 0.5 and 15.0% at one, so diversification shrinks as correlation rises and disappears at one.

> **Interview tip:** Show the formula, give the correlation of one case as the no-benefit benchmark, then say the calm correlation may understate the stress correlation. The interviewer listens for that last sentence.

INT.RP.2.b · Diversification and tail risk

Describe drawdown, tail risk and value at risk, compute simple versions and state what they miss.

Drawdown is the fall from a previous peak to a later trough, and maximum drawdown is the worst such fall. Recovery is asymmetric: after a loss of 20% you need a gain of 25% to get back, since 1 divided by 0.8 is 1.25. Value at risk states a loss not expected to be exceeded at a stated confidence over a stated horizon. A parametric version multiplies the portfolio value by the volatility and a normal quantile, such as 1.645 for 95% and 2.326 for 99%, and a longer horizon scales by the square root of time. A historical version reads the loss at the relevant percentile of past outcomes. Value at risk says nothing about how bad losses are beyond the threshold, and the normal assumption understates fat tails. Expected shortfall averages the losses beyond the threshold and is more informative, and stress tests apply specific shocks. Interview point: explain what value at risk does not tell you before you are asked.

WORKED EXAMPLE: MAXIMUM DRAWDOWN AND PARAMETRIC VALUE AT RISK

An interviewer gives a fund's value at six month-ends: 100, 120, 108, 90, 99 and 105. Separately, a portfolio of USD 80 million has daily volatility of 1.2%. Using a normal distribution with zero mean, a z value of 2.326 for 99% and independent days, find the fund's maximum drawdown, the gain needed from the trough to regain the peak, and the one-day and ten-day 99% value at risk.

Answer: The maximum drawdown is 25%, a 33.3% gain is needed from the trough, and the 99% value at risk is USD 2.23 million for one day and USD 7.06 million for ten days.

> **Interview tip:** Give the number, then say straight away what value at risk does not tell you. The interviewer listens for you volunteering that it is blind beyond the threshold before being asked.

INT.RP.3.a · Sizing and portfolio construction

Size positions using risk budgets and explain the intuition and the dangers of the Kelly criterion.

Position size should reflect conviction and risk, not just a round number. A common risk based method sets the size so that hitting the stop loss costs a fixed share of capital: a 0.5% loss budget with a 10% stop gives a position of 5% of capital, and a wider stop gives a smaller position. The Kelly criterion picks the bet fraction that maximises the long-run growth rate of capital. For a bet that wins b times the stake with probability p and loses the stake with probability q, the fraction is p minus the quantity q divided by b, which is zero when there is no edge. Betting more than Kelly adds risk and lowers growth, and at twice Kelly the growth rate is about zero. Because estimates of edge are noisy, practitioners use a fraction of Kelly, often half, which gives up a quarter of the growth for much less volatility. Interview point: say Kelly is a ceiling on sizing, not a target.

WORKED EXAMPLE: SIZING AN IDEA WITH A LOSS BUDGET AND KELLY

An interviewer says: you will lose at most 0.6% of capital if a stop 8% below entry is hit. Separately, a repeatable bet pays even money and wins with probability 55%. How big is the position under the loss budget, what does the Kelly formula say about the bet, and what would you actually do? Use the expected log growth rate, with natural logs, to compare full and half Kelly.

Answer: The loss budget gives a position of 7.5% of capital, full Kelly is 10%, and half Kelly at 5% keeps about 75% of the long-run growth rate with much less risk.

> **Interview tip:** Give the risk based size first, then say Kelly is a ceiling and that you would bet a fraction of it because the edge is estimated. The interviewer listens for the point that betting twice Kelly has about zero growth.

INT.RP.3.b · Sizing and portfolio construction

Construct and monitor a long-short book using gross and net exposure, beta, concentration and liquidity.

In a long-short book, gross exposure is longs plus shorts as a share of capital and measures total market footprint and leverage, while net exposure is longs minus shorts and measures directional dollar risk. Net exposure can mislead: beta-adjusted net multiplies each side by its beta, so a book with 40% net dollar exposure can have much lower net beta if the shorts are higher beta. Market neutral usually refers to beta close to zero, not necessarily dollar neutral. A short position has asymmetric risk, because a rising price enlarges the position and the loss is unbounded, so shorts are sized smaller and watched for squeezes. Concentration limits cap single positions and sectors. Liquidity is measured in days to exit, position size divided by a fraction of average daily volume. Interview point: describe the exposures you would report each day, and say that practices on limits vary between firms.

WORKED EXAMPLE: EXPOSURES AND BETA OF A LONG-SHORT BOOK

An interviewer says: a fund has capital of USD 200 million, longs of USD 260 million with an average beta of 1.1, and shorts of USD 140 million with an average beta of 1.4. Report gross exposure, net dollar exposure and net beta exposure as shares of capital, say what hedge makes the book beta neutral using an index future with beta 1, and say how long a 6 million share position takes to exit if average daily volume is 3 million shares and the fund trades at most 20% of volume.

Answer: Gross exposure is 200%, net dollar exposure 60% and net beta exposure 45%, selling USD 90 million of index futures makes the book beta neutral, and the 6 million share position takes 10 days to exit.

> **Interview tip:** List the exposures you would report each day: gross, net, beta-adjusted net, concentration and days to exit. The interviewer listens for the point that market neutral means beta near zero, not necessarily equal dollars long and short.

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