Portfolio risk management using VaR

Susan Thomas

April 7, 2011

Goals

- Approaches and terminology in risk management.
- Systematic and specific VaR
- Marginal and Incremental VaR
- Benchmark VaR
- Conditional VaR – Extreme Tail Loss, Extreme Shortfall
- Methods of portfolio VaR: normal linear, historical simulation, Monte Carlo.

Where VaR fits into risk management?

- Risk management needs a single measure of risk, which can be
  - *disaggregated* – so that a central risk manager can hand down risk limits to different department risk managers. This is critical for *risk budgeting*.
  - *aggregated* – aggregation is more complicated. It needs to be done at the level of different departments, as well as across different risk factors.

Recap: Systematic and specific VaR

- Total risk = systematic + specific
- Systemic risk: portfolio risk that can be mapped to risk factors.
  For example, equity portfolios driven by market risk, interest rate risk.
- Specific risk: risks that sit outside of the portfolio risk.
  For example, commodity futures portfolios driven by changes in warehousing costs.
- Model risk is one type of specific risk.
  For example, in equity portfolios, risk cannot be fully explained by the factor models.
  - What if the model captures *too little* of the portfolio variance.
  - Moreover, residuals may have large conditional volatility.
- Solution: track VaR of model residuals.

Recap: Evaluating systematic VaR

- Focus on portfolio returns, not risk factors.
- Simple linear normal distribution setting:
  - VaR on historical portfolio returns. (Works for equity/commodities.)
  - Simulate portfolio returns using Monte Carlo. (Needs to be done for bond portfolios.)
- When a factor model captures most of the portfolio volatility, total risk can be measured as:
  - Systematic risk when it dominates total risk.
  - Sum of systematic and specific risk (making the assumption that these are uncorrelated).
- Regulatory requirements: An extra weight needs to be applied to systematic risk to obtain total risk.

Stand-alone VaR

- Systematic portfolio risk can be broken into *stand-alone* components driven by core risk factors.
- Aim: disaggregate risk into factors associated with specific asset classes – equity, interet rate, foreign exchange, commodity VaR.
- Example: the volatility of a portfolio of commodity futures contracts is driven by the volatility of underlying spot prices. But there is also interest rate risk because of cost of carry.
- Stand-alone VaR of a factor sets sensitivity of all other risk factors to zero.
  Measures the risk of a specific asset in isolation – capital that can be used to compare across trading activities.
- This assumes that trading activities are risk-managed separately. They should not be penalised or rewarded for overall diversification.

Stand-alone vs. total VaR

- Stand-alone VaR do not add up to total VaR (unless correlation = 1).
- Stand-alone capital is inappropriate for risk budgeting. Trading desks are within their trading limits but overall business could be in breach of limits.
- Solution: use conditional VaR (expected tail loss). CVaR is sub-additive.

Marginal VaR

- Marginal VaR provides the manager with *relative risk factor contributions* to systematic risk of a diversified portfolio.
- Marginal VaR is explicitly additive.
- For instance, assume that:VaR $= f(\theta)$.
- The first partial derivative across various factors is the vector $g(\theta) = (f_1(\theta), \ldots, f_k(\theta))'$, where

$$f_i(\theta) = \frac{\partial f(\theta)}{\partial \theta_i} \qquad \forall i = 1, \ldots, k$$

- Then, the first order Taylor approximation to VaR is

$$f(\theta) = \theta' g(\theta) = \sum_{i=1}^{k} \theta_i f_i(\theta)$$

where $\theta_i f_i(\theta)$ is called the marginal component VaR for the i^{th} factor.

Incremental VaR

- With linear portfolios, and normality, the marginal VaR based approximation is exact.
- When portfolios have non-linear payoffs, or when the returns are simulated using Monte Carlo (as with bonds), the marginal is not exact.
- It remains useful for testing the effect of partial hedges on a factor VaR limit: by examining the change in VaR for a small change in θ.
- Here, we use incremental VaR, where

$$f(\theta_1) - f(\theta_0) = (\theta_1 - \theta_0)' g(\theta_0)$$

Benchmark VaR

VaR to benchmark risk of portfolios

- Fund managers need to benchmark both returns *and* risk.
- Typically, this means identifying a *benchmark* – often, ambiguous.
- When returns are benchmarked, we consider the *active/net return* – difference between portfolio and benchmark return.
- Then, *benchmark VaR* is the α-quantile of the *active return* distribution.

Example #1: Benchmark VaR

- Question: What is the 1% *benchmark VaR* over one-year for 10 million invested in a fund with an expected active return equal to the risk free interest rate and a tracking error of 3% annualised?
  1. Expected active return == risk free rate. Then, net return on the portfolio = 0.
  2. Tracking error is $\sigma_{\text{net returns}}$.
- Answer:

$$\text{VaR}_{\text{one-year, 0.01}} = 2.3264 * 0.03 = 6.98\%$$

- This means 1% benchmark VaR of 697,904 INR over one year on the 10 million portfolio.

Example #1: Benchmark VaR

- Interpretation: Losses *relative to the benchmark* will not exceed 697,904 with 99% confidence over the next one year.
- Tracking error only tracks what is the risk of outperforming the benchmark. Does not count the expected net return.
- Benchmark VaR takes expected return into account. If the benchmark loses more on net, then we will commensurately lose more.
- Empirical observation: Expected net return has a linear effect on benchmark VaR.
  - When a portfolio is expected to outperform a benchmark, then risk drops.
  - When a portfolio is expected to underperform the benchmark, the risk (as measured by the benchmark VaR) increases.

Conditional VaR

Conditional VaR

- Standard VaR tells nothing of how much we can lose if VaR is breached.
- Conditional VaR tells us the average level of loss.
- Two Conditional VaR measures:
  - *Expected tail loss*, (ETL), at α level of confidence on a portfolio worth V :

    $$-E(X|X < -\mathrm{VaR}_\alpha)V$$

    Here, X denotes the portfolio return over the appropriate time interval, and *VaR* is the Value at Risk.
  - *Expected shortfall*, (ES), at α level of confidence:

    $$-E(\tilde{X}|\tilde{X} < -\mathrm{BVaR}_\alpha)V$$

    where $\tilde{X}$ denotes the *net return* on the portfolio worth V and *BVaR* is the benchmark VaR at α level of confidence.

Differences between VaR, BVaR, ETL, ES

- Context: Consider the returns on a portfolio and its benchmark for $N = 1000$ days. Calculate both absolute losses and losses relative to the benchmark.
- Then for a risk assessment at 99% confidence:
  - the 1% VaR is the 10th largest absolute loss;
  - the 1% ETL is the average of the 10 largest absolute losses;
  - the 1% BVaR is the 10th largest relative loss;
  - the 1% ES is the average of the 10 largest relative losses.
- Note: this is effectively an empirical approach to VaR estimation using the historical simulation approach.

Implementing portfolio VaR

How to calculate portfolio VaR?

Three categories of approaches:

1. Assume multivate-normally distributed risk factors that the portfolio is linear in.
   Model: *normal linear* VaR
2. Use historical data for as long as possible to "estimate" VaR with little assumption about the distribution of risk factors.
   Model: *Historical simulation* VaR
3. Assume some distribution for risk factors (simplest case: multivariate normal) to simulate the returns of the portfolio.
   Model: *Monte Carlo* VaR

Normal linear VaR

- Only applicable to a portfolio with returns that are linear in the risk factors/component security returns.
  Options do not fit.
- When VaR is measured over a short-horizon, portfolio excess returns is set to zero, and Portfolio VaR becomes a function of

  $$w'\Sigma w$$

  where Σ is the covariance matrix.
- The final covariance matrix can be a combination of several covariance matrices. (As when we face heteroskedasticity of returns.)
- Because portfolio VaR is determined by the covariance matrix of the risk factor/security returns.
  Sometimes called the *covariance VaR* model.

Historical simulation of VaR

- Assume that all possible variations have occured in the past data – then use the past data to estimate the VaR.
- Terminology clash: the previous normal linear VaR model also uses historical returns to estimate Σ.
- Sometimes called *non-parametric*.
  Problem: Sometimes, the approach would be to simulate historical scenarios on contemporaneous changes in risk factors to “simulate” possible portfolio values.
  Eg., bond portfolio returns.
- *Obvious:* Bad idea for short time series of returns.
- This gets exacerbated if different risk factors have different frequencies and/or different data spans.
  Eg., interest rates and equity market returns.

Monte Carlo simulation of VaR

- Assume a distribution – any distribution! – for risk factor/security returns.
- Monte Carlo simulation involves:
  - Generating a time series for reach risk factor/security.
  - Generating the associated time series of portfolio returns.
  - Picking out the VaR from the distribution of these generated returns.
- Can make the size of the time series very large.
- Can use *any* distribution.
- In reality, the performance of historical returns dictates what is the distribution used.
- Use backtesting of predicted VaR (from a selected distribution) against realised returns to drive what distribution to use.

Summary: Advantages/Disadvantages

Advantages

- Normal VaR is analytically tractable.
- Historical VaR makes little assumption about the distribution of returns.
- Monte Carlo VaR can accomodate *any* distribution.

Disadvantages

- Normal VaR is restrictive: the only real give is in using mixtures of distributions.
- Historical VaR assumes the past contains all possible occurances of returns.
- Monte Carlo VaR is computationally intensive and its very flexibility leaves open considerable simulation errors.