NORTHFIELD | modeling financial markets worldwide

Advanced Optimization with Downside Risk

Dan diBartolomeo

Webinar
March 2023

Introduction

- In his seminal paper on portfolio theory Markowitz (*Journal of Finance*, 1952) wanted to use "semi-variance" as a measure of the downside risk.
  - The optimization math was intractable, so he settled for using variance as the risk measure.
  - It wasn't until Levy and Markowitz (*American Economic Review*, 1979) that the mean variance objective function became the common expression of investor utility.

- The idea that risk was not symmetric and that investors cared about downside risk was succinctly summarized in Roy (*Econometrica*, 1952)

- There have been many subsequent variations on the theme of downside risk including metrics like VaR, CVaR and Sortino Ratio (*Journal of Investing*, 1994)

Outline

- In today's presentation, we will consider three different forms of optimization problems that in some way recognize that investors will have different levels of concern for "upside" and "downside" risk.
- The first is the case where investors are comfortable with the standard mean-variance objective function, but where the expected distribution of asset returns is non-normal (e.g. presence of skew and/or kurtosis).
- The second are the set of cases where investors define the objective as "maximize return" subject to some dollar limit on expected loss (e.g. CVaR, expected shortfall).
- *Our new topic for today is the third case where the investor wants to minimize the likelihood of underperforming a fixed return target, such as the actuarial return assumption of a pension fund.*

Mean-Variance with Non-Normal Returns

- Since 2019, the Northfield Optimizer has had ability to incorporate higher moments of asset return distributions in risk calculations.
- The overall process is covered in diBartolomeo (“Performance Analysis When the Distribution of Returns is Non-Normal”, *Journal of Performance Measurement*, forthcoming 2023).
- The optimization problem is identical to the usual optimization set up except that “Odd Shaped Distribution (OSD)” is enabled.

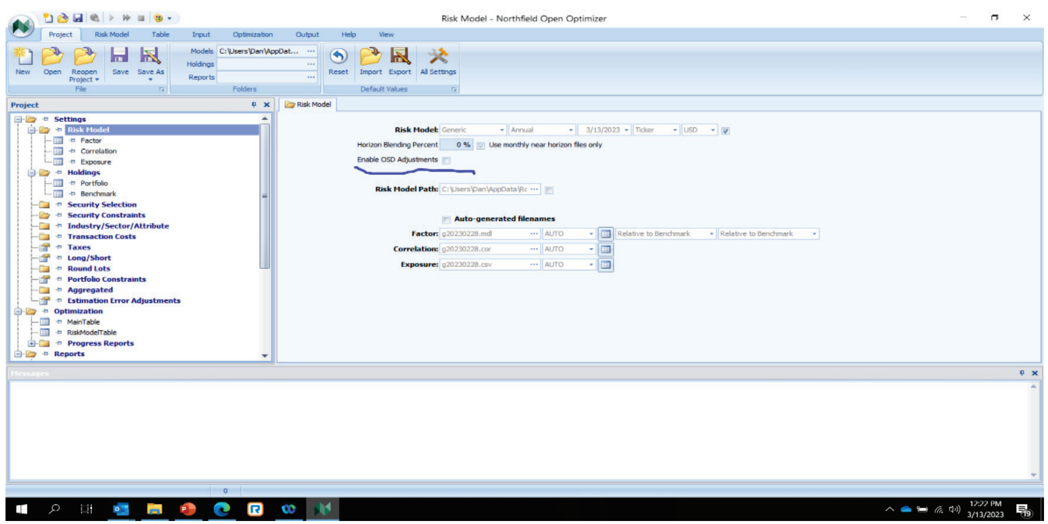

Parameterizing the Higher Moment Optimization

- For higher moments to matter, at least some assets in the problem must have non-normal distributions.
  - In our Everything, Everywhere model, the negative skew and kurtosis associated with fixed income default risk is included.
  - The default magnitude of the asset level higher moment values are scaled to 10 trading days to be consistent with UCITs rules but adjust to user horizon.

- *At the level of the portfolio, the amount of risk inclusive of higher moments is expressed as a "volatility equivalent"*
  - The magnitude of the change in risk by inclusion of higher moments is influenced by the "guaranteed survival time horizon".
  - A highly levered hedge fund can go broke in one day due to a margin call. A sovereign wealth fund may decline in value but will survive indefinitely.
  - Higher moments diversify away even faster than variances, so higher moments will matter less in more diverse portfolios.

Volatility Equivalence

- Volatility equivalence inclusive of higher moments (aka "tail risk") is computed on the fly during optimizations as the portfolio evolves from the initial to optimal state.
- Every common risk measure including variance (as in mean/variance), VaR, CVaR, expected shortfall, etc. *can be expressed as a simple scalar of the volatility equivalent.*
- For example, 95% VaR and CVaR for horizon H are:

VaR(95%, H) = portfolio value * 1.645 * volatility equivalent (H)

CVaR(95%, H) = portfolio value * 1.96 * volatility equivalent (H)

Computing Volatility Equivalence

- The full calculation of portfolio return properties including tail risk appears in Satchell and Hall (*Journal of Asset Management*, 2013).
- Given a user chosen confidence interval, a useful shortcut is the method of Cornish and Fisher (*Review of the International Statistical Institute*, 1938).

$$V_t = \sigma_t W_p / Z_p$$

$$W_p = Z_p + (S/6) * (Z_p^2 - 1) + Z_p * ((K-3)/24) * (Z_p^2 - 3) - Z_p * (S^2/36) * (2Z_p^2 - 5)$$

W_p = the Cornish Fisher "tail weight" parameter
σ_t is the standard deviation of the returns in the sample period
Vt = the "volatility equivalent" for the returns in the sample period

Volatility Equivalence Example

- Let's work a simple example:

  Mean = 8%
  Standard deviation = 10%
  Skew = -4
  Kurtosis = 5 (excess kurtosis = 2)

- The Cornish-Fisher tail weight parameter is -2.44 for the 5% lower tail

- At the 5% tail, the Z-score is -1.645

- The "volatility equivalent" is 10 * (-2.44/-1.645) = 14.83%

Nuances of VaR and CVaR Optimization

- When people talk about using CVaR, they often use it as equivalent to "expected tail loss" (ETL) or "expected shortfall" (ES).
  - Like VaR, CVaR is normally denominated in a dollar amount as in "I don't want to there to be more a P% (e.g. 2.5%) of a loss of $L dollars (e.g. $1 million) at the end of a period of T time."

- Are we measuring from the initial value of the portfolio or the expected ending value?
  - As in "My portfolio is starting at $10 million and I want to limit loss to $1 million, implying that I keep the value over $9 million *at the end of the period with probability (1-P)*.
  - My portfolio has an expected return of X%, so I expect my $10 million portfolio to be worth $10 * (1 + x/100) Million at the end of 1 year. *If I want to limit my loss L to $1 million, the ending value of my portfolio should be above (*$10 * (1 + x/100)-1) Million *with probability (1-P).*

Continuous Time Risk Limits

- We could define the problem in continuous time through the entire period:
  - My initial portfolio is $10 million and want to limit my potential loss to $1 million with P probability (i.e. a floor value of $9 million) throughout the entire period from today to 1 year from today.
  - This is not the same as saying $9 million at the end of the period, *since there is a non-zero probability that the portfolio value would go below the floor sometime during the period but go back up above $9 million before the end of the period.*
  - See Kritzman and Rich *(Financial Analyst Journal*, 2002*).*

- If the time horizon is very short all these definitions converge to mean the same thing.
  - For longer horizons they are different because over longer horizons the effect of the "drift" (positive expected return) is material.

Downside Risk with a Target Return

- Many investment situations are best expressed as “minimize my risk of getting a return less than T%”.

- *With this objective we care about higher returns only to the extent they represent a lower likelihood of underperforming the target return T*

- See Harlow (*Financial Analyst Journal*, 1991)

- Multiple objectives are possible, which all have a coincident solution under the volatility equivalent method.
  - LPM0, likelihood of underperforming the target
  - LPM1, expected value of underperformance
  - LPM2, squared expected value of underperformance.

Risk Relative to Target Returns

- Even if asset returns are normally distributed, the asset return distributions will have skew relative to the target return.
- Assets with expected returns higher than the target will have a lower "volatility equivalent" and those with expected returns lower than the target will have a higher volatility equivalent.
- *The expected magnitude of the skew relative to target is calculated similarly to Pearson's coefficient of skew.*

Skew = 3 * (Mean – Target) /standard deviation

- We can now calculate the volatility equivalent inclusive of the induced skew.

Optimizing for Target Returns

- Once we know the ratio of the volatility equivalent of an asset (with higher moments) to the standard deviation (assuming normal), we can manipulate the optimization inputs to find the minimum downside risk portfolio.
  - Set Risk Acceptance close to zero (minimum variance problem)
  - Multiply each asset weight by its respective ratio of volatility equivalent/standard deviation.
  - The weights will now sum to something other than 100%
  - Adjust the cash position so that the weights sum to 100%

- Optimize
  - Constraining the weight in cash back to the pre-adjustment value is optional
  - The resultant portfolio will be the minimum downside risk portfolio available

Conclusions

- The concept of making decisions about portfolio composition based on "downside" risk has been around for a very long time.

- The Northfield Optimizer allows for a number of different optimization objectives that allow investor risk to be structured around various representations of economic loss.

- Today, we discussed three different structures for downside risk problems
  - Mean/variance equivalence with asset level higher moments
  - VaR / CVaR optimizations
  - Minimizing lower partial moments relative to a return target.

- Structuring such problems in the Optimizer is minimal extra effort for users.