

Optimization 101

Dan diBartolomeo
Webinar (from Boston)
October 22, 2013



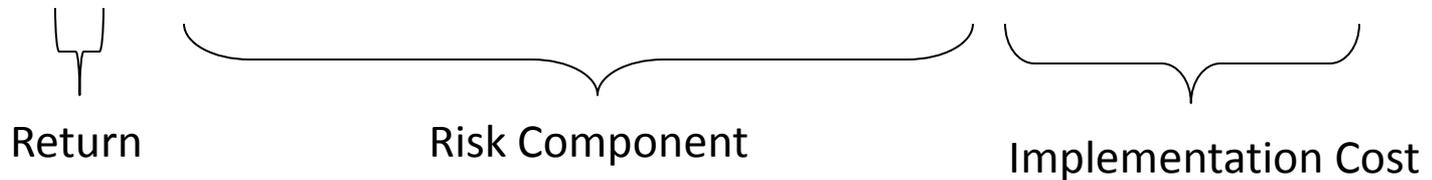
Outline of Today's Presentation

- The Mean-Variance Objective Function
- Optimization Methods, Strengths and Weaknesses
- Estimation Error and Input Scaling
- Constraints, Hierarchy and Costs
- Trading Costs and Turnover Limits
- Taxable Optimization
- Best Practices, Common Errors and Troubleshooting

Optimization Objective Function

- Investor objective is to maximize risk- adjusted returns net of costs:

$$U = \alpha - ((\sigma_s^2 / RAP_s) - (\sigma_u^2 / RAP_u)) - ((C + T) * A)$$



- Most commercial portfolio optimizers uses the objective function described in Levy and Markowitz (1979)
- Portfolio return variance is the proper measure of risk because the difference between the arithmetic average rate of return and the geometric average rate of return is proportional to the variance (see Messmore, 1995)

Objective Function Explained

- $U = \alpha - (\sigma_s^2 / RAP_s) - (\sigma_u^2 / RAP_u) - ((C + T) * A)$
 - α = the “certainty equivalent” expected portfolio return
 - σ_s^2 = portfolio variance risk due to common factors (correlation across securities)
 - σ_u^2 = portfolio variance risk due to stock specific risks
 - **RAP = risk tolerance**
 - C = transaction costs for the optimization
 - T = capital gain taxes for the optimization
 - **A = amortization constant**

Note that this is a convex function and the various scalars transform all terms into certainty equivalent expected return units

Quadratic Optimization Methods

- There are several computational methods that can be used to solve quadratic optimization problems
- Cutting Planes
 - Represent the solution as the vertex of a set of hyperplanes
 - Guarantees the solution in a finite number of steps and get faster as you add more constraints
 - Non-monotonic approach to optimality so often hard to interpret individual trades
- Integer Programming
 - Most precision on cardinality constraints such as position rounding and maximum number of assets
 - Downside is speed. Full integer methods could take thousands of years to solve certain common portfolio problems so heuristic short cuts are always needed

Quadratic Optimization Methods

- Conic Optimization

- Not really an optimization method at all.
- It's just a way to represent certain aspects of optimization as the equations that describe intersection of a plane and a cone
- If you change the angle between the plane and the axis of the cone you get describe different optimization objectives
- You still solve your problem with one of the other methods
- <http://www.northinfo.com/documents/299.pdf>.

- Interior-Point Gradient Optimization

- Monotonically increasing objective. Most similar to the way a human would operate by hand so most intuitive
- Find the minimum turnover feasible portfolio
- Do “pairs trading” from there climbing the gradient
- Disadvantage: Slows down when lots of linear constraints are employed
- Northfield uses this approach since modern computers are very fast

Estimation and Scaling

- Objective function is in units of “certainty equivalence”
 - Optimizers are stupid. When you tell them something they believe you
 - Epistemology: What do we really know, what do we think we know and what do we merely assume?
 - Alpha, risk and costs are often estimated over different horizons
- Optimization inputs must be adjusted to reflect certainty
 - Expected returns are the most uncertain
 - Risk models are pretty good but not perfect (you know the sign)
 - Trading costs models are pretty good but not perfect (you know the sign)
- Our optimizer offers several functions to help with input scaling
 - Black-Litterman, Jorion for return magnitude
 - Ledoit Wolf for risk magnitude
 - “Horizon Blending” to adjust risk forecast horizon to your typical holding period

Constraint Basics

- Any optimization can be run with a variety of three classes of constraints
- Linear combinations of security weights
 - Position limits
 - Industry and sector limits
 - Factor Exposure or Penalty File Variables
- Cardinality Constraints
 - Maximum number of assets
 - Minimum trade sizes
 - Rounding to round lots
- Stopping Criteria
 - Turnover limits
 - Maximum realized capital gains

Constraint Feasibility

- If there are conflicts among the constraints, the portfolio optimization is called “infeasible”
 - Our system will try to find a reasonable solution to infeasible problems but this does not mean that you can always get an intuitive result
- Consider the portfolio construction:
 - Max number of assets equal to 20
 - Maximum weight in a position is 4%
 - Sometimes conflicts are very subtle and hard to notice
 - We always produce warning messages: Pay attention

An Optimization in Steps

- First, the optimizer attempts to find a portfolio that meets all linear constraints, if one exists
 - Next it will find the feasible portfolio that has the least turnover from the initial portfolio. This is the real starting point of the process
- Once a feasible portfolio is found, the main optimization occurs
 - Our goal is to adjust the security weights to maximize the objective function, given your input settings
- Violations of cardinality constraints are checked when maximum objective is achieved
 - Solutions to most cardinality constraints are approximations that hold well only when dealing in small quantities. *Use only when actually needed.*
 - If there are any violations of cardinality constraints that can be fixed, the optimizer will adjust the security weights to fix them
- You're done when marginal utility of the best and worst positions are equal to within a selected tolerance

Constraint Priority

- If there are conflicts among the constraints, the Optimizer will try to come up with a reasonable compromise
 - Linear constraints have priority over cardinality constraints
 - Within linear constraints, position sizes are higher priority than grouping constraints, that are in turn higher than factor/penalty constraints
 - There is no specific priority among cardinality constraints, so the optimizer tries to find the best compromise based on the objective function inputs
 - Cardinality constraints can be violated they are economically irrational (i.e. they drive the portfolio far away from the unconstrained optimum)
- Constraints on stopping criteria can cause premature termination and unintuitive results
 - Imagine a concentrated portfolio where sector constraints require 65% two way turnover to be feasible. If you set a turnover constraint below 65% you won't get the fully correct answer. It can't be done.
 - Same concentrated portfolio but where almost all the positions have embedded capital gains. If you set a "Max Realized Capital Gain constraint" equal to something small (say zero), its not possible and the process will stop prematurely

Measuring the Cost of Constraints

- The cost of constraints in optimization has traditionally been expressed as what is called a Lagrangian multiplier.
 - Let's assume you had one or more hard constraints on a portfolio (e.g. $18 < P/E < 20$). One way to do this would be:

Use a quadratic penalty with a goal of $((\min + \max)/2)$

Start the penalty scale at zero

If the attribute is within the bounds, QUIT

If the attribute is outside the bounds, increase the scale and go back to the previous step

- When the process ends, the value of the scale is the Lagrangian multiplier for that constraint

More on Constraint Costs

- If you have multiple hard constraints they can conflict with each other so the Lagrangian multipliers don't really tell the whole story.
- If you have "stopping criteria" on the optimization process (e.g. turnover constraints), those won't be in the same units you have more interpretation problems.
- Scherer and Xu (2007) gives a translation of the traditional Lagrangian multiplier method into alpha units but this only works for hard weight constraints (e.g. long-only, max weight, etc.).
- There is a simple and flexible way to interpret the cost of all constraints. Our Optimizer already gives you the implied alphas for the securities.
 - You can just calculate the Spearman rank correlation of the implied alphas and the user's input alphas.
 - **If there are no binding constraints, there correlation should be close to one. If the correlation is close to zero, the constraints are so tight that the optimization isn't actually doing anything at all.**
 - Consider the Axa Rosenberg case

Trading Costs and Turnover

- The most widely made error by optimization users is to assume zero trading costs and use a turnover constraint instead
 - If trading costs are constant an approximate closed form solution to optimal turnover exists. See Grinold and Stuckelman (1993)
- Real world institutional trading creates market impact
 - Northfield provides a sophisticated market impact model free of charge that covers all 65,000 stocks in our global universe.
 - Use them unless you have something better
- MPT, CAPM and commercial optimizers assume the problem is for a single period
 - Real world is multi-period to which Markowitz and VanDijk (2003) propose a game theory solution that is impractical for large problems
 - Our optimizer contains a tractable approximation for amortizing trading costs more efficiently, <http://www.northinfo.com/Documents/500.pdf>

Special Aspects of Taxable Optimization

- Inclusion of tax lot level capital gain taxes makes the optimization objective function “linear piecewise” with respect to the lots
 - This makes calculating partial derivatives of the objective function much more complicated
 - Always use the NLTC setting “On” when doing problems with large concentrated positions. This breaks the optimization into lots of “small bites” that are less likely to be impacted by the piecewise nature of the problem
 - As previously noted, the use “max capital gain” stopping criteria can induce infeasible solutions
 - The incremental taxes between short term and long term gains have to be amortized separately. This is a built-in function
- The system allows users two sets of values for cost basis and acquisition date, one actual and one set of “user controlled overrides”
 - The override inputs can be used to tailor outcomes for special situations like “tax gain” harvesting that are otherwise non-convex.

Best Practices / Common Errors

- Understand what constraints are really appropriate
 - Constraints should be in place to make sure the optimizer doesn't do anything crazy if there is an error in one of your inputs. If all your inputs are correctly scaled to reflect your beliefs no constraint should be binding
 - Always try minimal constraints first to see what you get. If the result is not intuitive, there is something wrong and forcing the solution with constraints is a very poor response
 - The Optimizer produces reasonable outcomes from its internal hierarchy but resolving infeasible solutions is time consuming. Using lots of tight constraints is defeating the purpose anyway.
 - Don't overuse cardinality constraints. If you only exactly 50 stocks in your portfolio or \$500,000 minimum trades you need to be able to articulate an economic reason that is not better described with a continuous function.

Best Practices / Common Errors

- Users often want to do Sharpe Ratio or Information Ratio maximizing optimizations
 - This can be done in the Northfield optimizer by including an ETF for the benchmark to held long or short.
 - The rest of the positions are the ratio maximizing portfolio. Just rescale the weights to sum to 100%
 - Not economically rational in the absence of leverage or the presence of non-zero trading costs.
- Long/short portfolios require some thought
 - In a long only portfolio you change risk level through diversification
 - In a long/short portfolio you can change risk level by diversification or by changing leverage levels. The risk tolerances implicit in these two aspects of portfolio construction must make sense together.
 - Try a long/short without leverage constraints and see what you get

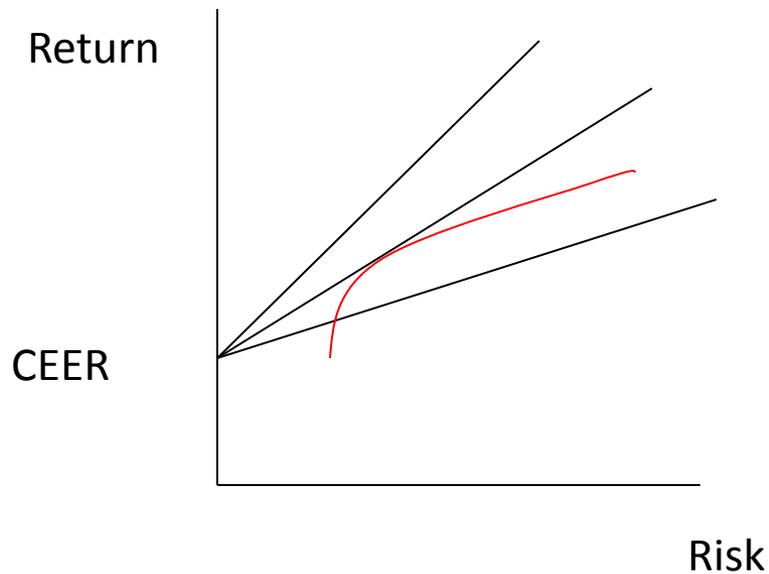
Best Practices / Common Errors

- Be careful about the specification of cross-listings, ADRs and GDRs
 - Always use composite assets to represent all but one of the positions
 - If you go long 100 X on the NYSE and short 100 X in London, the factor model will see those as two different positions, rather than a zero net position if you're not careful.
 - The Optimizer will automatically recognize the composite asset representation if present even if the data record exists.
- Augmenting the objective function to include both tracking error and absolute risk at the same time is easy
 - Absolute risk is just tracking error around a zero volatility asset (i.e. cash)
 - By definition, the risk free asset has no covariance with anything so the problem collapses to the regular problem with some cash mixed into the benchmark (proportional to relative risk aversion)

Best Practices / Common Errors

- In a computer no two numbers are ever equal. They are just equal to a specified number of decimal places.
 - Abide by the concept of significant digits. Avoid false precision.
 - Never set $\text{MIN} = X = \text{MAX}$ for any hard linear constraint. You must always leave a finite increment (can be very, very small) between the MIN and MAX of linear constraints so $\text{MIN} < X < \text{MAX}$
- RAP values (risk tolerance coefficients) have economic meaning
 - A good rule of thumb is 6 times your desired volatility level. See Wilcox (2003) on the Discretionary Wealth Hypothesis.
 - You can have different values for factor and specific risk but you have to be rational. You can't say you are 100 times more tolerant of one kind of risk or the other without a pretty good reason.
- Be sensitive to "basis" formation. The basis is the set of securities that will contain all members of the optimal portfolio (usually plus extras)
 - Set "Max Iterations" to no more than $3N$ where N is the maximum possible number of assets in the optimal portfolio

CEER & RAP



Slope = $1/RAP$

$$CEER_1 = CEER_0 + Risk/RAP$$

Where $CEER_0$ is risk free return