NORTHFIELD | modeling financial markets worldwide

Optimization with Non-Divisible Position Sizes

Dan diBartolomeo

Webinar

November 2024

Introduction

A basic mathematical requirement needed to solve convex optimization problems is that the objective function must be continuous and differentiable with respect to portfolio weights.

In some real-world situations this mathematical requirement is not met.

If we are optimizing a real estate portfolio, it is generally impossible to buy or sell a portion of a building. in the same way we would choose to transact a desired number of shares of an equity.

Any transaction must be *"all or none"* which requires different methods to reach an optimal solution.

Another common example of this problem is in taxable portfolios wherein closing out an entire position will *usually* allow the transaction to escape the restrictions imposed by "wash sale" rules.

In this presentation we will discuss practical procedures that allow "all or none" type trades to be included in routine optimization problems.

A classic solution to this problem is to conduct the entire optimization using "integer programming" rather than "gradient" or "cutting plane" methods.

Integer programming solutions for rare but plausible problems can take hours or days to solve.

It's like doing a large jigsaw puzzle by "trial and error."

Defining An Optimal Portfolio

To be an optimal portfolio a vector of portfolio weights must generally meet what are called the Kuhn-Tucker conditions:

- The objective function is a polynomial that may include algebraic terms for return, risk, trading costs, taxes, and other portfolio specific preferences (e.g. a target income yield)

To meet the K/T requirements

- The first partial derivatives of the objective function with respect to the current weight of each portfolio asset must be equal unless constrained by a weight bound (i.e. weight min/max limits)
- The second partial derivatives for all portfolio weights not at a constraint boundary must all be negative.
- Together these requirements mean that *any further changes in portfolio weights will reduce* rather than improve the value of the objective function, so we know the current weight vector is optimal.

Solving the Portfolio Optimization Problem

There are two widely used methods for solving this kind of problem.

- Gradient (see Sharpe 1978)
- Cutting Plane (see Markowitz 1959, Van Hohenbalken 1975)

With very significant limitations on the complexity of the problem, some portfolio optimization problems can be solved in closed form.

- See Elton, Gruber, and Padberg (1978)
- EGP optimization with simple transaction costs was presented at our 2019 conference, Analytical Solutions of Optimal Portfolio Rebalancing

Integer programming methods are required to get *exact solutions* when certain types of constraints are involved in the problem.

- Maximum number of assets, minimum trade size, round lots, etc.
- However, since the inputs to a portfolio optimization *are estimates of the future* (not known values) getting exact solutions doesn't provide the expectation of meaningful improvement.

Making Optimization Intuitive in Two Steps

Sharpe (1978) suggests a method that should be intuitive to investors.

- Compute the first partial derivatives of the objective function with respect to each of the current portfolio weights.
- Sort the set of possible portfolio assets by their Marginal Utility (MU = the first partial derivative)
  - You want to sell the asset with the lowest MU (bottom of the list)
  - Use the proceeds of the sale of the worst asset to buy the highest MU (top of the list)
- The same asset may be involved in multiple “pairs” during the process.
- *By repeating this process over and over, all MU values will equalize, and the second partial derivatives will be negative for assets not bounded by a constraint.*

In practice, we need to add a second step to the process.

- Once we identify what we want to buy and sell for each pair, *we have to decide **how much** to buy or sell*?
- The optimal amount to trade can computed in closed form for most cases.
- *The correct amount to trade at each iteration is the lesser of the optimal amount and the maximum amount allowable under any constraints that are in force.*

The Problem with Non-Divisible Positions

If an asset within the optimization is indivisible (e.g. a building), it must be traded in an "all or none" fashion.

We are likely to have situations when the size of the indivisible position is larger than the optimal amount to trade within a given pair.

In some cases, the size of the indivisible position is larger than can be traded at all under some kind of hard constraint.

If the process cannot trade the *"best for worst"* due to constraints, we can skip to the next best pair (second best, second worst, or both) that can trade if other assets are available in the problem.

We can't use our "pairs" process in the usual way, because an indivisible asset cannot be broken up into pieces to participate in multiple pairs during the "climbing" of the gradient.

Illiquid Assets and the Fuzzy Efficient Frontier

One very practical way to solve the "indivisibility" problem is to recognize that inputs to portfolio optimizations are not known values, but rather are estimates of an uncertain future.

When assets are indivisible, illiquid, or both (e.g. most real estate), we have to recognize that optimal portfolio weights *need to be defined as a vector of ranges rather than scalar elements.*

- The weights still need to sum to 100%.
- We used bootstrap resampling to compute a "fuzzy" efficient frontier in Gold (1995)

Many illiquid assets have yet another aspect complicating use of optimization which is that the availability of any given asset to be bought may be sporadic.

- Belev and Gold (2015) won "Practitioner Paper of the Year" from the American Real Estate Society.
- This paper combined the illiquid nature of real estate (hard to get out), the fuzzy nature of optimal weights, and sporadic availability of assets to buy into a coherent, optimization procedure.
- Similar procedures can be used for other illiquid assets such as private equity, private credit.

A Common Problem with Taxable Portfolios

Subject to disagreement among tax attorneys is the issue of whether or not a "wash sale" can arise from the sale of an entire position.

- Many financial advisors take the position that closing out a position in a given security cannot cause any tax benefit associated with the event to be nullified as a violation of the wash sale rule.
- Our tax advisors have advised that while this view is true in the preponderance of cases, tax benefits associated a "close-out" can be construed as a wash sale violation.

If you take the position that a "close out" cannot cause a wash sale, there are obviously situations wherein economically material capital losses would be realized by selling off an entire position.

- *However, this is clearly an "sell all or sell none" kind of problem that is ill-suited for most optimization algorithms as it violates the requirement to be "continuous and differentiable."*

Defining the "Contra-Asset"

One way to make "all or none" problems more tractable is to define a suitable "contra-asset" that can always be the "other side" of the all or none trade.

- If we add (buy) or remove (sell) a portfolio position of known weight, the sum of the portfolio weights will no longer be 100% until we figure out what the "other side" of the trade might be.

After a "sell" our asset weights will now be less than 100%.

- *Plug the missing value with cash when defining return/risk in absolute terms (i.e. cash benchmark)*
- *Plug the missing value with an ETF for the benchmark when defining return/risk in index relative terms.*
- We could reweight (trade) the remaining portfolio positions to force them to once again sum to 100%

Defining the "contra asset" as cash for absolute return problems, and an ETF for the benchmark in relative return problems means that the **contra-asset effectively has no return and no risk when substituted into the portfolio for the sold asset.**

- *The math gets simplified a lot and you can be long or short the contra asset as required.*

Trading with a Defined Contra Asset

If we define the contra asset to be return, risk, and trading cost free, it will have no material contribution to the objective function.

We only need to calculate the change in portfolio utility that would arise from an "all or none" trade.

- This could be done with simple "before and after" calculations using the usual optimizer reports.
- If there are multiple indivisible assets involved, calculating "before and after" for each indivisible asset can be unnecessary work.

A more appealing alternative is to pre-process the "all or none" decision for each indivisible position before the optimization even starts.

- *If the determination of "all" or "none" is done in advance, the optimization constraints and other inputs can then be set to enforce the chosen decision when the optimization is run to address the other assets in the problem.*
- *There are a lot of required details all of which can be easily automated.*

Pre-Processing An "All or None" Trade for a Taxable Investor

- Assume the current price for security X is $95 and today is 8/30/2024.

- **Event Date Action**

8/15/2024 Buy 100 shares of X at $75 per share

6/25/2013 Bought 1000 of X at $125 per share

- The usual normal "sort" order of lots to be sold ignoring the wash sale rule would be to minimize capital gains.

6/25/2013 1000 shares at a loss of $30 per share or a $30,000 loss on the lot.

8/15/2024 100 shares at $75 at a gain of $20 per share or $2000

Since we've bought within the last 30 days, selling the any shares from the older lot would be a wash sale.

Pre-Processing An “All or None” Trade for a Taxable Investor

- However, if we believe that selling the entire position would NOT be a wash sale, we can for a possible “all or none” transaction.
  - *The investor can then claim a combined loss of $28,000, which is clearly economically attractive.*
  - The analytical problem is that irrespective of what we do for sorting of lots, proper trade hinges on an “all or none” decision, rather than considering the problem share by share or even lot by lot.
  - Let’s introduce a few more pieces of information that are likely to be relevant.

Portfolio value = $1 Million

Value of the X position = 1100 * 95 = 104,500 or 10.45%

Benchmark weight of X = 3%

Active weight = 7.45%

Alpha (X) = 2%

Trading cost (X) = $.10 per share

Amortization Constant = 100 (all costs are weighed against one year return and risk)

Portfolio tracking error = 5%, tracking variance = 25

RAP = 30

MV(initial, X) = 117 **(I just made this value up)**

Tax rate is 25%

Pre-Processing An "All or None" Trade for a Taxable Investor

- If we think of the problem in dollar profit and loss over one year, we get:
  - We have $104,500 invested in X that has an alpha of 2% per year.   If we keep the position, we add .02 * 104,500 to the value of the portfolio = $2090 over the year. If we sell, we lose that expected profit.
  - If we sell the position, we remove the 7.45% overweight contribution which contributes to tracking variance (if we assume the other side of the trade is an ETF for the benchmark)
  - If we sell the position, we create a new underweight of 3% which would contribute to the tracking variance to the tracking variance (if we assume the hypothetical other side of the trade is an ETF for the benchmark)
  - If we think of the change in the weight of X in isolation, we've reduced the active weight by 4.45% so there is a net decrease in tracking variance of **approximately** (.045 * 117 *.5) = 2.635 so the portfolio tracking variance goes from 25 to 22.365.
  - Since our RAP = 30, the portfolio level risk term goes down to 22.365 /30 = .7455%. In dollar terms that is an improvement of .007455 * $1 Million = $ 5553.98
  - If we sell whole position, we get the dollar benefit of the $28,000 tax loss = $28,0000 *. 25 = $7000 *(assuming that the investor has other gains that need to be offset)*.

Pre-Processing An "All or None" Trade for a Taxable Investor

- Let's add things up.

If we sell, we give up $2090 of alpha	-2090
If we sell, the portfolio penalty risk is smaller by	5553.98
If we sell, we have a trading cost of $110	-110
If we sell, we get a tax benefit of $7000	7000

- **Total = $10,353.98**

Since dollar value summation of the utility components is positive, we can conclude that we should do the "all" trade rather than the "none" trade for this specific case.

Forcing the “All” Trade in An Otherwise Regular Optimization

We want to force any subsequent optimization to include the “all” trade.

The relevant details are:

- Override the “basis date” for all recent lots (e.g. the 8/15/2024 lot) to any date that is more than a month ago but within the past year.
- Nothing changes for tax calculations, but the wash sale “impediment” disappears.
- *Doing this step alone is insufficient since there is a gain on the recent lot which is unlikely to be sold because of the tax cost.*
- Set the absolute Max weight of X to zero (or a MIN absolute weight of zero with very negative alpha like -100).

*For all other positions that are not designated as “**force a sell all,**” the optimization would operate as usual.*

Conclusions

There are a variety of circumstances where the decision to add or remove an asset from a portfolio must operate with a fixed quantity, rather than leaving the amount of the asset to be determined by the optimization procedure.

This situation often arises with indivisible assets (e.g. real estate) and there are other related aspects to optimization problems.

The indivisibility problem also arises in the situation of “close out” sales in taxable portfolios where there is an expectation that the wash sale rules would be violated by selling only part of the position.

We’ve illustrated how problems involving indivisible positions can be analyzed in an intuitive pre-optimization process by introducing the concept of a correctly defined contra-asset.