Portfolio optimisation

J.E. Beasley

CARISMA, ***C***entre for the ***A***nalysis of ***R***isk and Optim***IS***ation Modelling ***A***pplications, Brunel University, UK
http://people.brunel.ac.uk/~mastjjb/jeb/jeb.html

Abstract
In this talk we review portfolio optimisation, with a focus on financial applications. Here the problem is to decide the assets (a portfolio) to hold that have desired characteristics. Markowitz mean-variance portfolio optimisation is relatively well known, but has been extended in recent years to encompass cardinality constraints. Less considered in the scientific literature are problems such as:

- index tracking – where the objective is to match the return achieved on a benchmark index such as the S&P500
- enhanced indexation - where the objective is to exceed the return achieved on a benchmark index; here we may have a desired specified excess return, or we simply wish to do better than the benchmark
- absolute return – where the objective is to achieve a desired return (irrespective of how the market, as represented by the benchmark index, performs)

We will outline the mathematical optimisation models that can be adopted for portfolio problems such as these and review the results achieved to date.

Markowitz mean-variance portfolio optimisation
To proceed with Markowitz mean-variance portfolio optimisation we need some notation, let:

N be the number of assets (e.g. stocks) available
μ_i be the expected (average, mean) return of asset i
ρ_{ij} be the correlation between the **returns** for assets i and j ($-1 \leq \rho_{ij} \leq +1$)
s_i be the standard deviation in **return** for asset i
R be the desired expected return from the portfolio chosen

Then the decision variables are:
w_i the proportion of the total investment associated with (invested in) asset i ($0 \leq w_i \leq 1$)

Using the standard Markowitz mean-variance approach we have that the unconstrained portfolio optimisation problem is:

Minimise $$\sum_{i=1}^{N} \sum_{j=1}^{N} w_i w_j \rho_{ij} s_i s_j$$

subject to

$$\sum_{i=1}^{N} w_i \mu_i = R$$

$$\sum_{i=1}^{N} w_i = 1$$

$$0 \leq w_i \leq 1 \qquad i=1,\ldots,N$$

Here we minimise the total variance (**risk**) associated with the portfolio whilst ensuring that the portfolio has an expected **return** of R, and that the proportions sum to one. This formulation is a simple nonlinear programming problem.

Usually nonlinear problems are difficult to solve but in this case because the objective is **quadratic**, computationally effective algorithms exist so that there is (in practice) little difficulty in calculating the optimal solution for any particular data set.

The point of the above optimisation problem is to construct an ***efficient frontier***, (**unconstrained efficient frontier, UEF**) a smooth non-decreasing curve that gives the best possible tradeoff of **risk** against **return**, i.e. the curve represents the set of **Pareto-optimal** (**non-dominated**) portfolios.

One such efficient frontier is shown below for assets (shares) drawn from the UK FTSE (Financial Times Stock Exchange) index of 100 top companies.

Efficient frontier for the FTSE 100

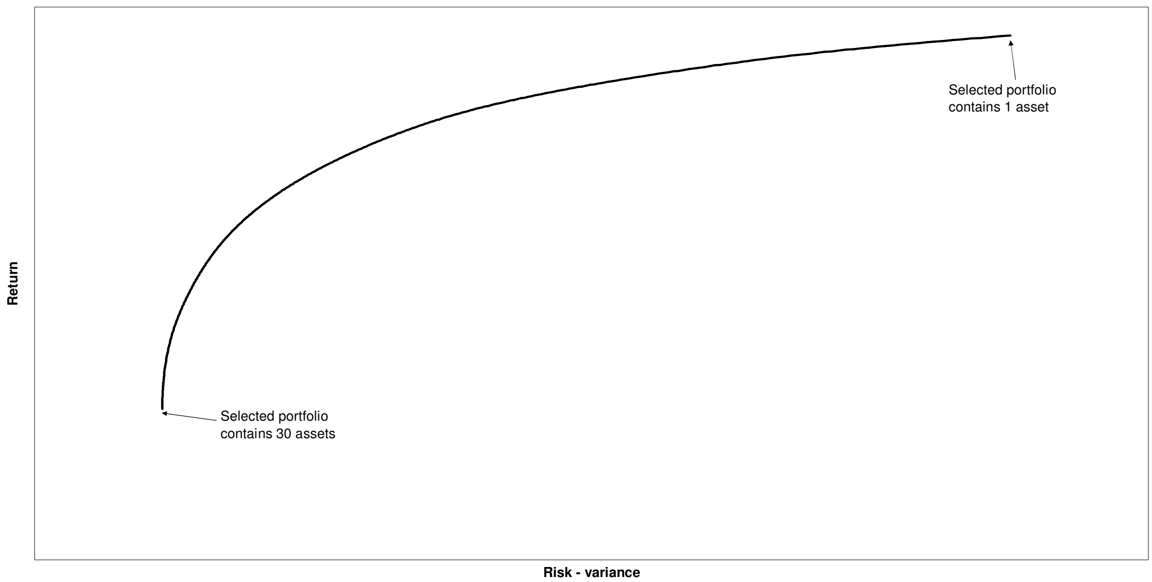

One thing to note here is that the formulation given above can be easily extended to add linear constraints that reflect a number of investment restrictions that we might have in any particular case (e.g. sector constraints constraining the proportion invested in particular market sectors such as energy, telecommunications, etc).

Markowitz mean-variance portfolio optimisation with cardinality constraints
Imposing a cardinality constraint to restrict the number of assets in which we can invest can only (for a given level of return) increase the risk. This is because in the UEF above, for a given level of return, the risk is at a minimum (as the mathematics for the Markowitz model explicitly requires). The practical reason why we might allow an increase in risk is that we may find it more convenient/less costly in terms of transaction cost to have a portfolio with just a few assets.

In order to extend our formulation to the cardinality constrained case let:
K be the desired number of assets in the portfolio

Introducing zero-one decision variables:

z_i = 1 if any of asset i is held
= 0 otherwise

the cardinality constrained portfolio optimisation problem is:

minimise $$\sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j \rho_{ij} s_i s_j$$

subject to

$$\sum_{i=1}^{N} w_i \mu_i = R$$

$$\sum_{i=1}^{N} w_i = 1$$

$$\sum_{i=1}^{N} z_i = K$$

$$0 \le w_i \le z_i \qquad i=1,...,N$$

$$z_i \in [0,1] \qquad i=1,...,N$$

You should note here that the material presented here relating to cardinality constrained portfolio optimisation is relatively up to date (such material only being reported in the academic literature in 2000). By contrast the (unconstrained) Markowitz model dates from the 1950's. In terms of practical application the main changes since the 1950's for Markowitz are:

- better (cheaper and faster) computers
- better software for solving the quadratic program involved in the Markowitz model
- improved data availability and better quality data

For the unconstrained portfolio optimisation problem the frontier was a nice smooth continuous curve. In particular note that for every possible return (between the minimum return achievable and the maximum return achievable) we had a portfolio on the efficient frontier with an associated risk.

In the presence of cardinality constraints the efficient frontier may become **discontinuous**, where the discontinuities imply that there are certain returns which no rational investor would consider (since there exist portfolios with less risk and greater return). This means the efficient frontier in the cardinality constrained case is distinctly different from that in the unconstrained case.

To illustrate this point the figure below shows four assets (stocks) drawn from the FTSE 100. In that figure all possible portfolios involving **exactly** two assets (K=2) are shown.

For four assets there are 4×3/2 = 6 possible choices of pairs of assets, each of which leads to a different curve in the figure below.

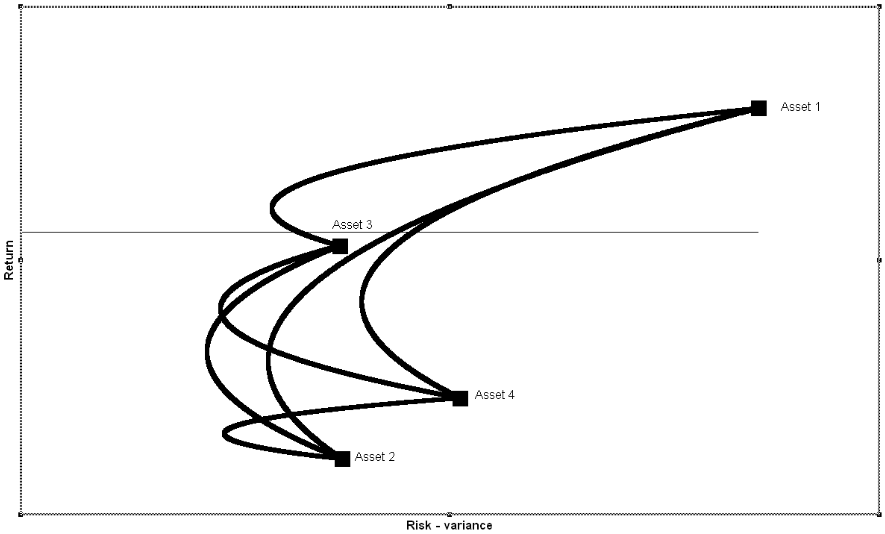

Suppose we are interested in the return shown on the horizontal line in the diagram above. There are only three possible portfolios that achieve this return:

- one where the horizontal line intersects the curve between assets 3 and 1 (so a portfolio consisting of those two assets in some proportion)
- one where the horizontal line intersects the curve between assets 2 and 1 (so a portfolio consisting of those two assets in some proportion)
- one where the horizontal line intersects the curve between assets 4 and 1 (so a portfolio consisting of those two assets in some proportion)

Are these three portfolios dominated or not?

The figure below shows the efficient frontier as derived from the figure above. Note the discontinuities – where there are return values for which there is no portfolio having that return which is not dominated. We refer to the cardinality constrained efficient frontier as the CCEF.

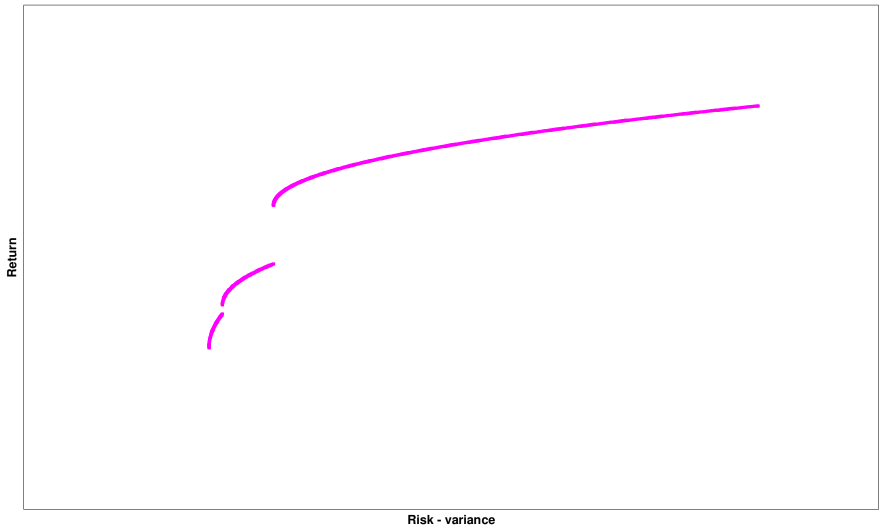

Algorithmically the above mathematical program for the CCEF is hard to solve (as it is a mixed-integer quadratic program). Typically therefore in the literature metaheuristics such as genetic algorithms, simulated annealing and tabu search have been applied to the problem.

As an illustration of what can be achieved with cardinality constrained portfolio optimisation we show below some tradeoff curves for assets chosen from the DAX index.

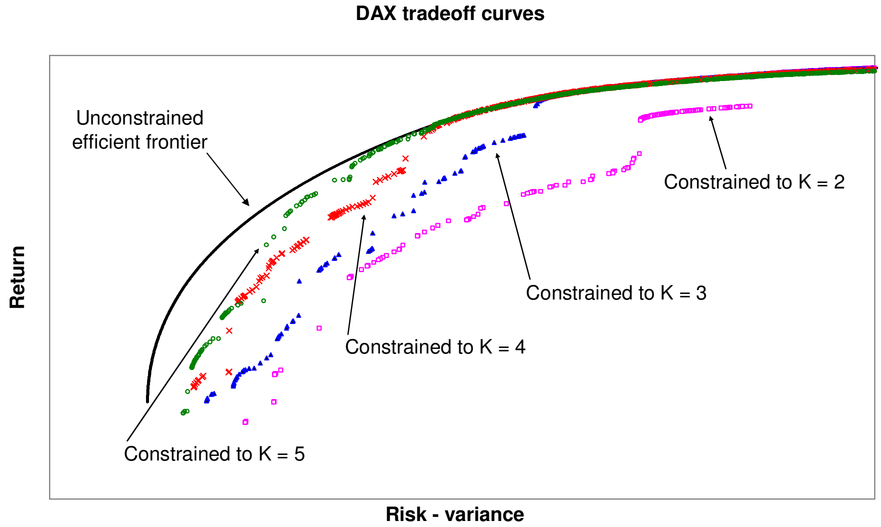

Index tracking

Stock market indices such as: Dow Jones, FTSE 100, FTSE All Share, S&P500 tell you, in an easy to digest form, how the stocks (companies) represented in the index have changed in value over time.

If you want you can invest your money in index tracking funds ("trackers") which aim to reproduce the performance of the index over time, perhaps by investing in all of the stocks that make up the index. If an index tracking fund invests in all of the stocks in the index in such a way that its investment in each stock mirrors index composition (e.g. if a stock makes up 10% of the index then it makes up 10% of the investment) then the fund is said to be following a **full/complete replication** strategy.

Full replication is possible – but as the number of stocks in the index grows it can be an expensive strategy in terms of transaction cost. This is because:

- stocks typically enter/leave the index at regular intervals (as the composition of the index changes) and so the entire fund must be **rebalanced** as this occurs to mirror the index as it changes
- any new money that is invested in (or money taken out of) the fund must be spread across all stocks to mirror the index

Suppose therefore that we do not wish to adopt full replication. Then in essence we can view the index tracking problem as a ***decision problem***, namely to decide the ***subset*** of stocks to choose so as to (hopefully perfectly) mirror/reproduce the performance of the index over time. We call the subset of stocks we choose a **tracking portfolio (TP)**.

Suppose that we observe over time 0,1,2,...,T the value of N stocks, as well as the value of the index we want to track. Further suppose that we are interested in deciding the best set of K stocks to hold (where K<N), as well as their appropriate quantities. In index tracking we want to answer the question:

*"what will be the best set of K stocks to hold, as well as their appropriate quantities, so as to best track the index in the **future** (from time T onward)?"*

Our basic approach in index tracking is a historical look-back approach. To ask the historical question:

"what would have been the best set of K stocks to have held, as well as their appropriate quantities, so as to have best tracked the index in the ***past*** *(i.e. over the time period [0,T])?"*

and then hold the stocks that answer this question into the future.

Let:

N be the total number of distinct stocks (companies) in which we can invest

K be the desired number of distinct stocks in the TP

ε_i be the minimum proportion, and

δ_i be the maximum proportion, of the TP that must be held in stock i if any of stock i is held

X_i be the number of units of stock i in the current TP

V_{it} be the value of one unit of stock i at time t

T be such that we have observed historical values for stocks and the index over the time period 0,1,2,...,T. The time T represents a decision point, a time at which we **may** switch from our current TP $[X_i]$ to a new TP.

I_t be the value of the index at time t

C be the total value of the current TP $[X_i]$ at time T plus any **cash change** in the portfolio (either new cash available for investment or cash being withdrawn at time T), i.e. $C = \sum_{i=1}^{N} V_{iT}X_i$ + cash change

γ be the limit on the proportion of C that can be consumed by transaction cost

Then our decision variables are:

x_i the number of units of stock i that we choose to hold in the new TP

z_i = 1 if any of stock i is held in the new TP
= 0 otherwise

Without significant loss of generality we allow $[x_i]$ to take fractional values.

The constraints associated with the index tracking problem are:

$$\sum_{i=1}^{N} z_i = K$$

$$\varepsilon_i z_i \le V_{iT}x_i/C \le \delta_i z_i \qquad i=1,...,N$$

$$C_{trans} = \sum_{i=1}^{N} \text{transaction cost}(X_i \rightarrow x_i \text{ at time T})$$

$$C_{trans} \le \gamma C$$

$$\sum_{i=1}^{N} V_{iT}x_i = C - C_{trans}$$

$$x_i \ge 0 \qquad i=1,...,N$$

$$z_i \in [0,1] \qquad i=1,...,N$$

In time period t we get a return associated with the index, $R_t = \log_e(I_t/I_{t-1})$, where we define return using continuous time. If, in each and every time period, the return associated with the TP, $r_t = \log_e[[\sum_{i=1}^{N} V_{it}x_i]/[\sum_{i=1}^{N} V_{it-1}x_i]]$, was EXACTLY equal to R_t then this might seem ideal. A possible objective in terms of index tracking is therefore: minimise $\sum_{t=1}^{T} (r_t - R_t)^2/T$, i.e. minimise average squared error.

The index tracking problem, as formulated above, is a nonlinear integer program, and so hard to solve. Algorithmically metaheuristics can be applied. The figure below shows some results with this objective for the S&P500.

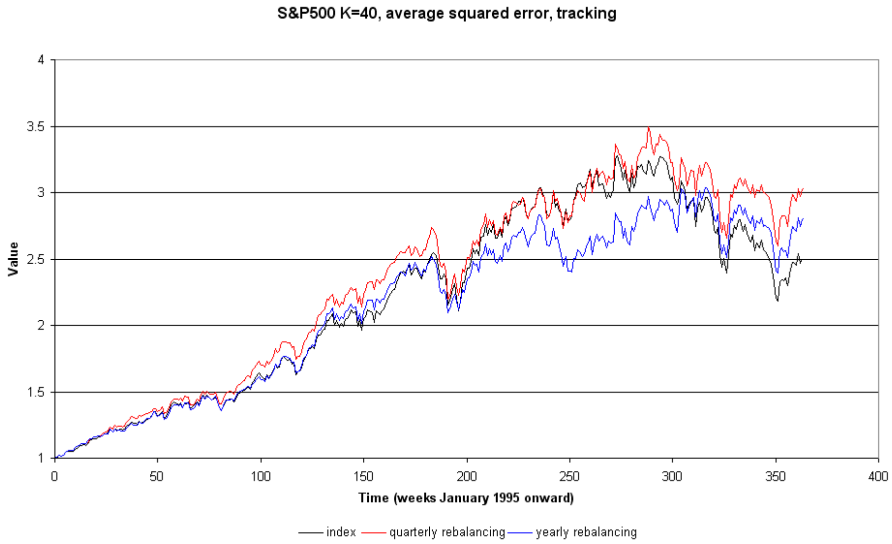

Alternatives

Note here that one consideration when working in the finance area as it relates to portfolio optimisation is that ***problems are (mathematically) typically not uniquely defined***. That is, there are different mathematical formulations that take different (but valid) views of the problem. Taking Markowitz mean-variance (for example) could we not define risk using some measure other than variance?

With regard to index tracking an alterative view relates to regression. Suppose we perform a linear regression of the return from the tracking portfolio against the return from the index, i.e. the regression $r_t = \alpha+\beta R_t$. What intercept α and slope β would you expect to get if you perfectly track the index?

Enhanced indexation

Suppose we are interested in excess return, return over and above the return on the index. Here we seek to out-perform the index (***enhanced indexation***).

How then can we construct TPs that "both track the index and exceed it"? One way is:

- take the return R_t given by the current index
- create an **artificial (enhanced return) index** whose return is $A_t = R_t+R^*$ where R^* is the desired excess return per time period
- track this enhanced return index

The advantage of this approach is:

- the constraints we have given for the index tracking problem automatically apply to the enhanced indexation problem

- any algorithm developed to find TPs for the index tracking problem can "automatically" be applied to the problem of deciding a portfolio of stocks (a TP) so as to outperform the index.

Here we again have alternatives, in particular in relation to the objective, one of which is semi-specified out-performance where we require $r_t \geq A_t \; \forall t$, leading to the objective: minimise $\sum_{t=1}^{T} (\min[0, r_t - A_t])^2/T$

The figure below shows some results with this objective for the S&P500.

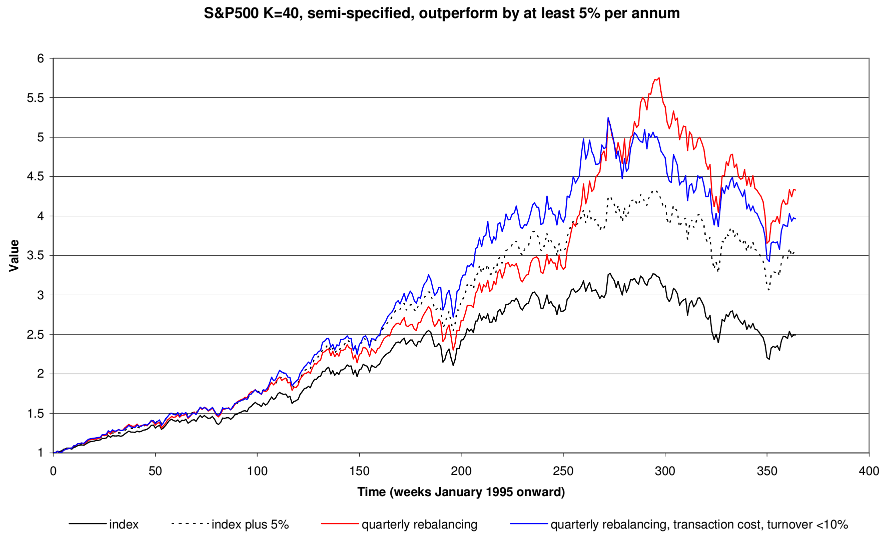

One issue with trying to out-perform the index is that sometimes one need not specify a desired excess return (R^* above), rather one just wishes ***to do better than the index***. Here one objective that can be used is a modified Sortino ratio, namely:

maximise $(\sum_{t=1}^{T} r_t/T - R^{mean})/\sqrt{[\sum_{t=1}^{T} (\min(0, r_t - R^{mean}))^2/T]}$, where $R^{mean}=(\sum_{t=1}^{T} R_t/T)$ is the mean return on the index.

Absolute return

Here the objective is to choose a portfolio so as to achieve a desired return (irrespective of how the market, as represented by the benchmark index, performs).

One simple way to do this is simply to replace the index return R_t in the index tracking objective by the desired (constant) return Δ, so the objective becomes

minimise $\sum_{t=1}^{T} (r_t - \Delta)^2/T$

Conclusions

We have reviewed a variety of mathematical models that can be applied to problems concerned with portfolio optimisation. One key point to note is that problems are (mathematically) typically not uniquely defined and so different mathematical formulations potentially exist.

References

1. Chang, T.-J., Meade, N., Beasley, J.E. and Sharaiha, Y.M., “Heuristics for cardinality constrained portfolio optimisation” Computers & Operations Research 27 (2000) 1271-1302.
2. Beasley, J.E., Meade, N. and Chang, T.-J., “An evolutionary heuristic for the index tracking problem” European Journal of Operational Research 148 (2003) 621-643.
3. Canakgoz, N.A. and Beasley, J.E., “Mixed-integer programming approaches for index tracking and enhanced indexation” European Journal of Operational Research 196 (2009) 384-399.
4. Meade, N. and Beasley, J.E., “Detection of momentum effects using an index out-performance strategy” Quantitative Finance 11 (2011) 313-326.