Portfolio optimisation: models and solution approaches

J.E. Beasley

Abstract

We review financial portfolio optimisation. 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 portfolio optimisation problems such as: index tracking; enhanced indexation; absolute return; market neutral.

We outline the mathematical optimisation models that can be adopted for portfolio problems such as these and solution approaches that can be used.

Key concepts

Different types of portfolios require different mathematical models

Even for portfolios intended for the same purpose the model to use is not uniquely defined

Adopt an optimisation mindset in building/choosing a model

Nonlinear models are more difficult to solve numerically than linear models

Back to school

Could you solve

$$2x + 3y = 7$$

$$4x - 7y = 1$$

simultaneous linear equations

Back to school

Could you solve

$$2(x/y) + 3xy = 7$$

$$4x - 7y = 1$$

simultaneous nonlinear equations

Markowitz mean-variance portfolio optimisation
Purpose: portfolios which balance risk and return

We need some notation, let:

N be the number of assets (e.g. stocks) available

μ_i be the expected (average, mean) return (per time period) of asset i

σ_{ij} be the covariance between the returns for assets i and j

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:

$$\text{minimise} \quad \sum_{i=1}^{N} \sum_{j=1}^{N} w_i w_j \sigma_{ij}$$

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 and σ_{ij} is positive semidefinite, 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 drawn from the UK FTSE index of 100 top companies.

Efficient frontier for the FTSE 100

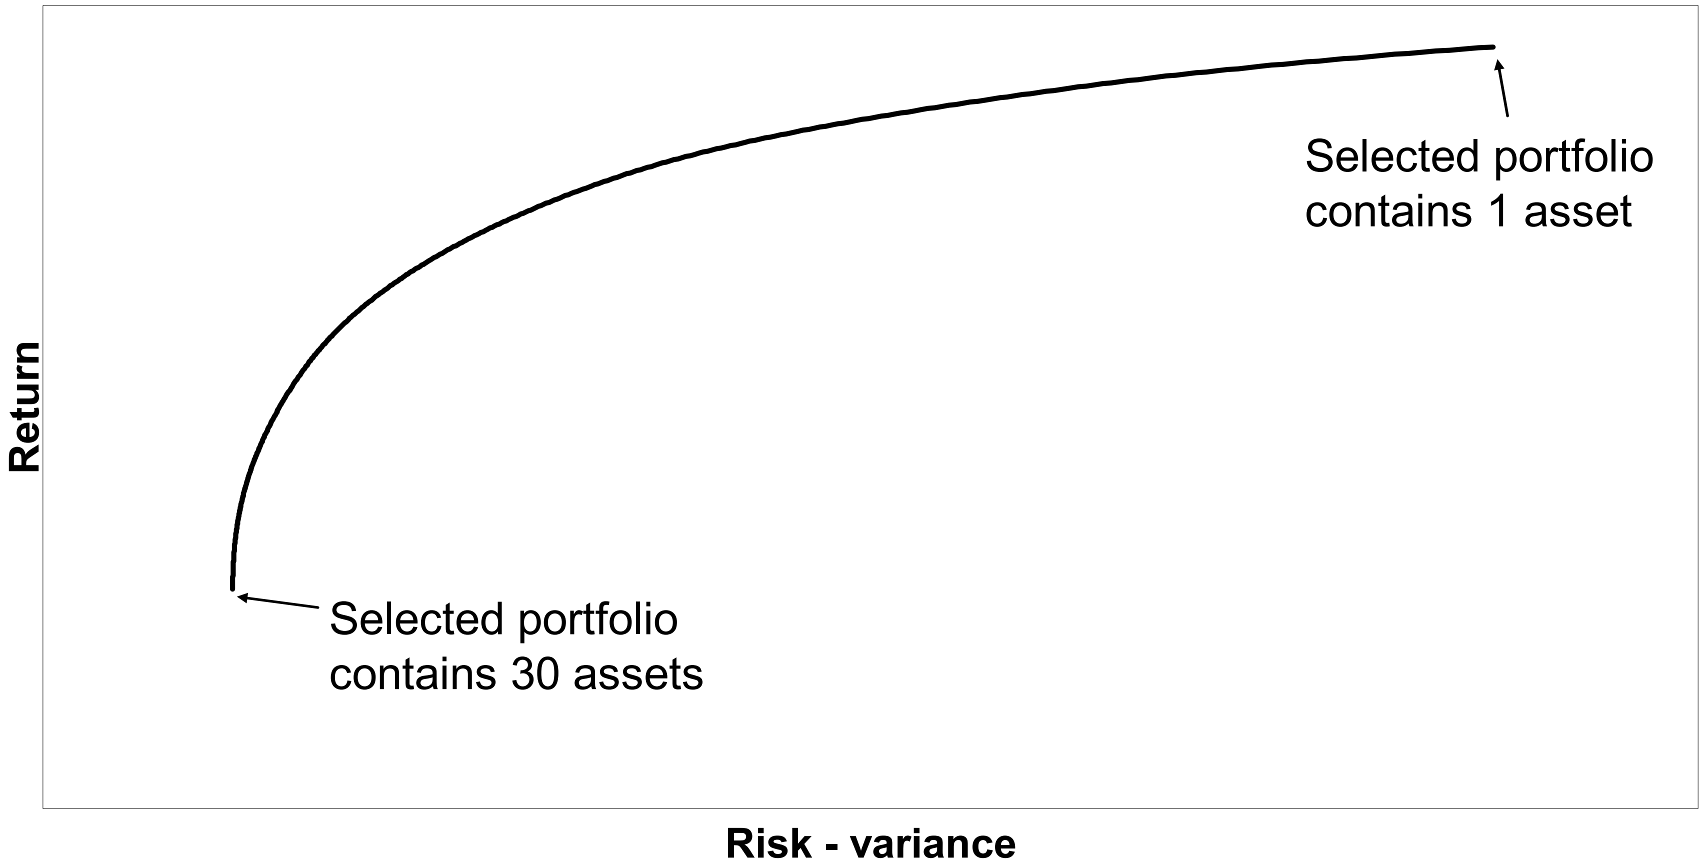

The approach followed in Markowitz mean-variance optimisation is:

- look into the (immediate) past for relevant data (in-sample data)
- use that data to form a portfolio, as outlined in the model above, where the in-sample data is used to produce values for μ_i and σ_{ij} that are used in the optimisation model
- hold that portfolio into the (near) future (out-of-sample)

The underlying logic is that, since accurately forecasting future prices/returns for assets is extremely difficult, data from the immediate past is our best guide to construct a portfolio to hold into the immediate future.

All of the mathematical models you will see here have been deliberately simplified because of time constraints. If you want to add in more realistic features such as:

- restricting the proportion invested in any asset
- restricting the proportion invested in sets of assets (class/sector constraints)
- rebalancing an existing portfolio
- incorporating transaction costs associated with trading
- incorporating both long and short positions

this can (in many cases) be done.

Transaction cost

In the context of Markowitz mean-variance portfolio optimisation the role of transaction cost is that it is the price we pay (now) to enable us to move from our existing portfolio to a new portfolio that will (on the basis of in-sample optimisation), have a better performance than our existing portfolio.

As such the role of the holding period, the number of time periods (H) for which we intend to hold the new portfolio, must be considered. In particular a consequence of the Markowitz model in the presence of transaction cost is that trading:

- reduces return; but also
- reduces risk.

As an indication as to how the holding period affects the efficient frontier the figure below shows efficient frontiers for a 31 asset example for a variety of holding periods.

Here the dotted nature of the efficient frontiers results from the fact that (for computational reasons) we have only examined a limited number of values of R in plotting the frontiers seen.

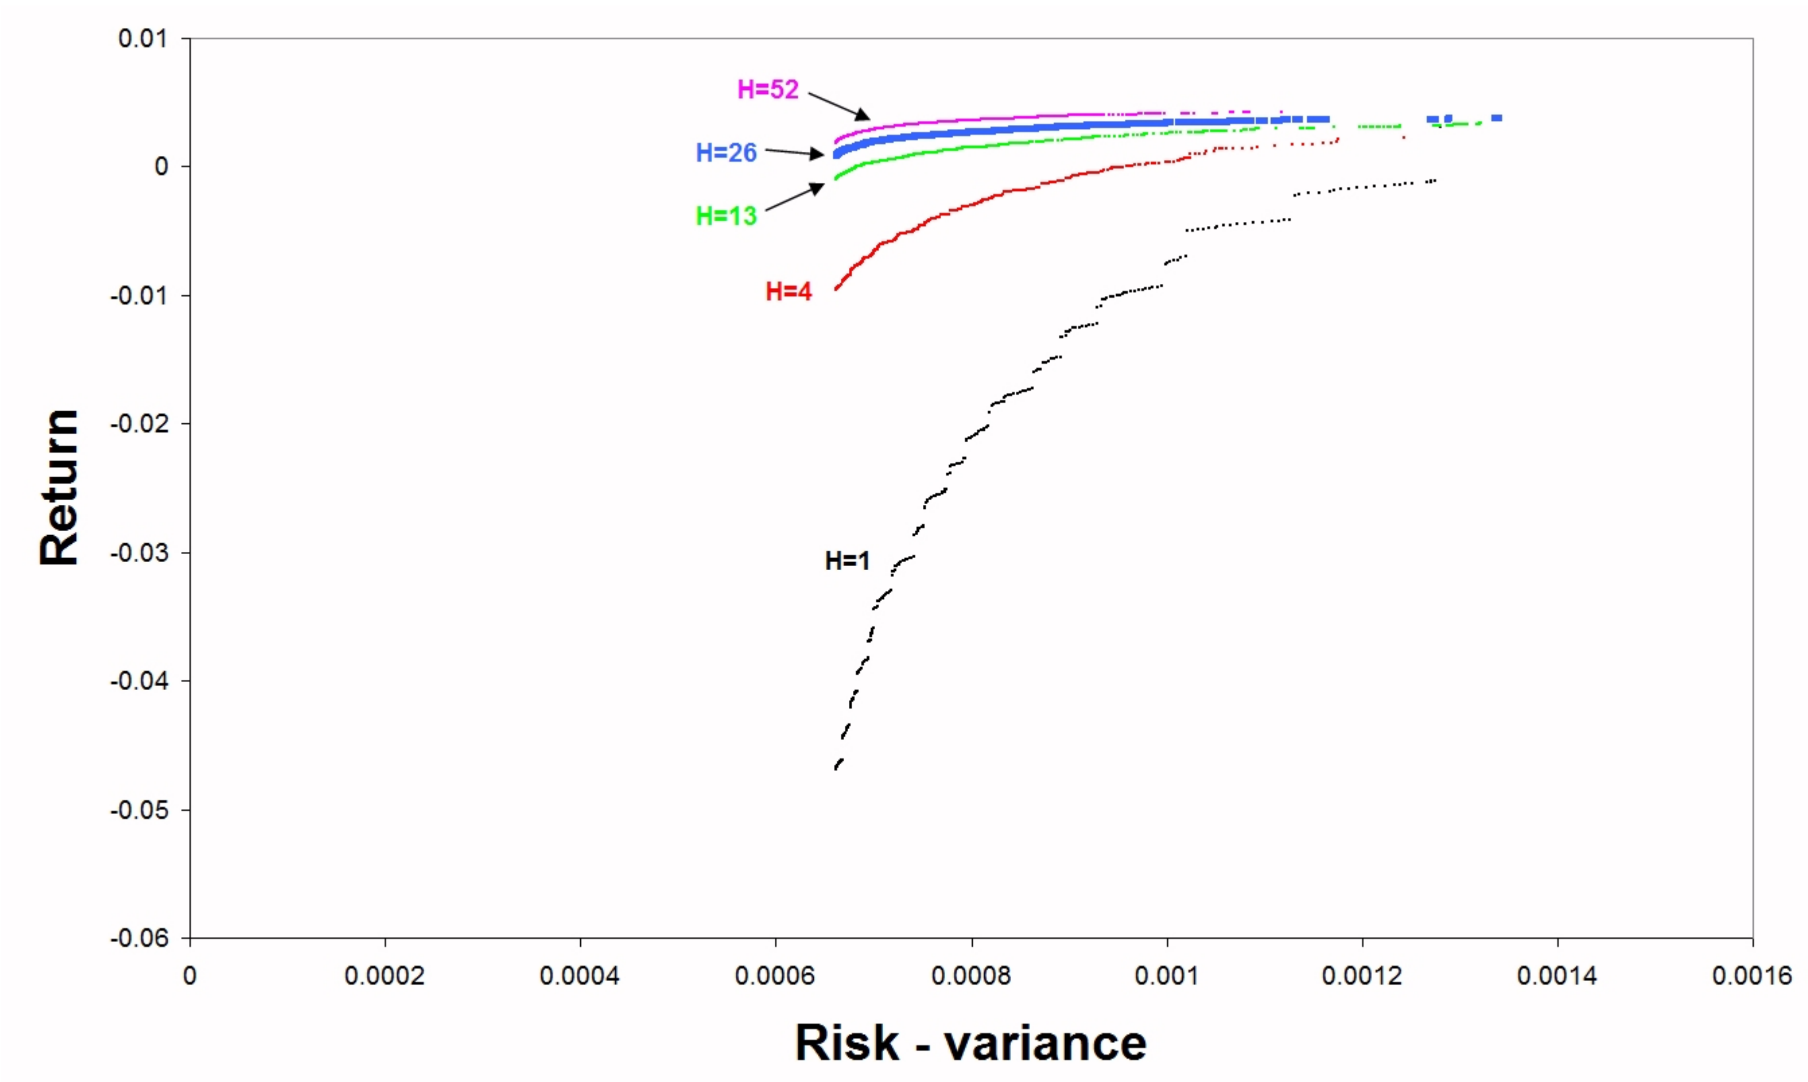

0.01
0
-0.01
-0.02
-0.03
-0.04
-0.05
-0.06
Return
0
0.0002
0.0004
0.0006
0.0008
0.001
0.0012
0.0014
0.0016
Risk - variance
H=52
H=26
H=13
H=4
H=1

Markowitz mean-variance portfolio optimisation with cardinality constraints

Purpose: portfolios which balance risk and return AND which allow us to control the number of assets held

Imposing a cardinality constraint to restrict the number of assets (K, say) in which we can invest can only (for a given level of return) increase 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 impose a cardinality constraint is that we may find it more convenient to have a portfolio with just a few assets, or simply that we desire a degree of control to shape the optimised portfolio with respect to the number of assets that it contains.

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 $\sum_{i=1}^{N} \sum_{j=1}^{N} w_i w_j \sigma_{ij}$

subject to

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

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

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

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

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

Algorithmically the above mathematical program is hard to solve (as it is a mixed-integer quadratic program). A mixed-integer program is one in which some variables take continuous (fractional) values, some take integer values.

Typically therefore in the literature metaheuristics such as:

- genetic algorithms (evolutionary algorithms)
- tabu search
- variable neighbourhood search
- simulated annealing

have been applied to the problem.

Definitions

An optimal algorithm is one which (mathematically) guarantees to find the optimal solution (e.g. to an optimisation problem)

A heuristic algorithm has no such guarantee

A metaheuristic is a "framework" within which you design a heuristic for the problem you are considering

For the unconstrained portfolio optimisation problem the frontier was a nice smooth continuous curve

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).

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 $4 \times 3/2 = 6$ possible choices of pairs of assets, each of which leads to a different curve in the figure below.

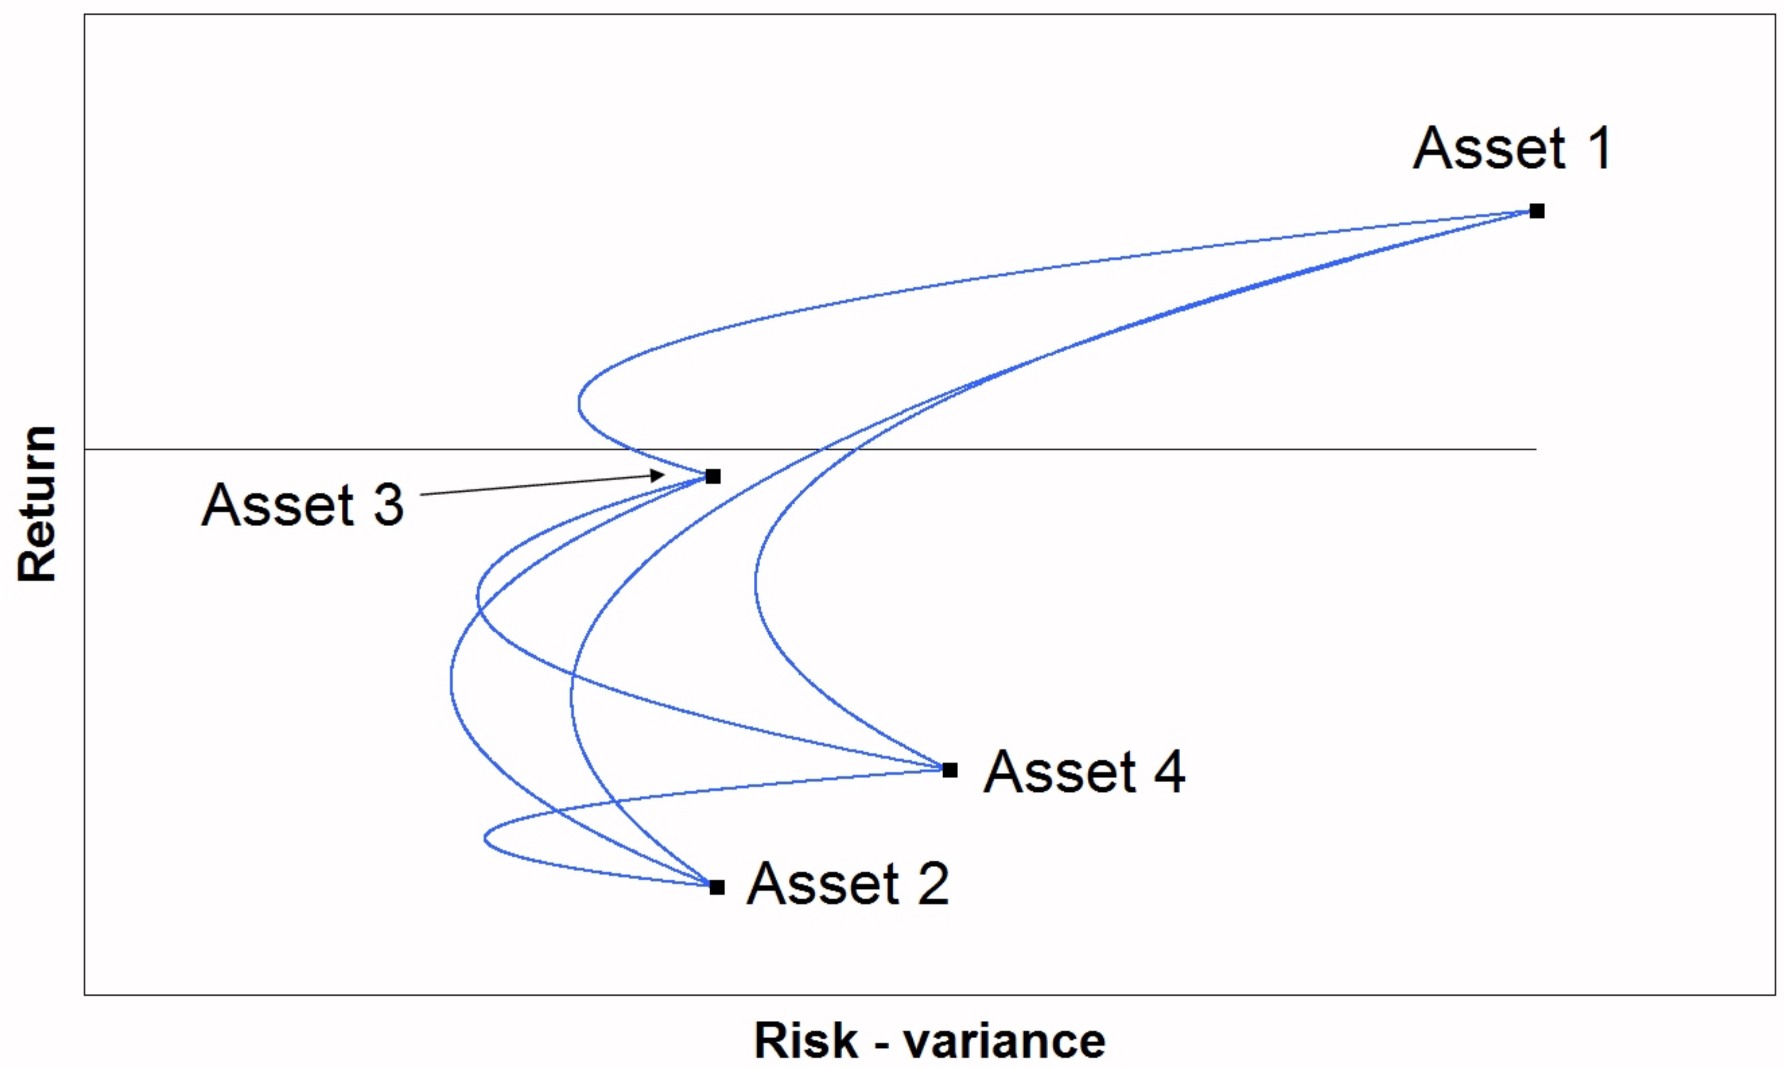
Asset 1
Asset 3
Asset 4
Asset 2
Return
Risk - variance

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.

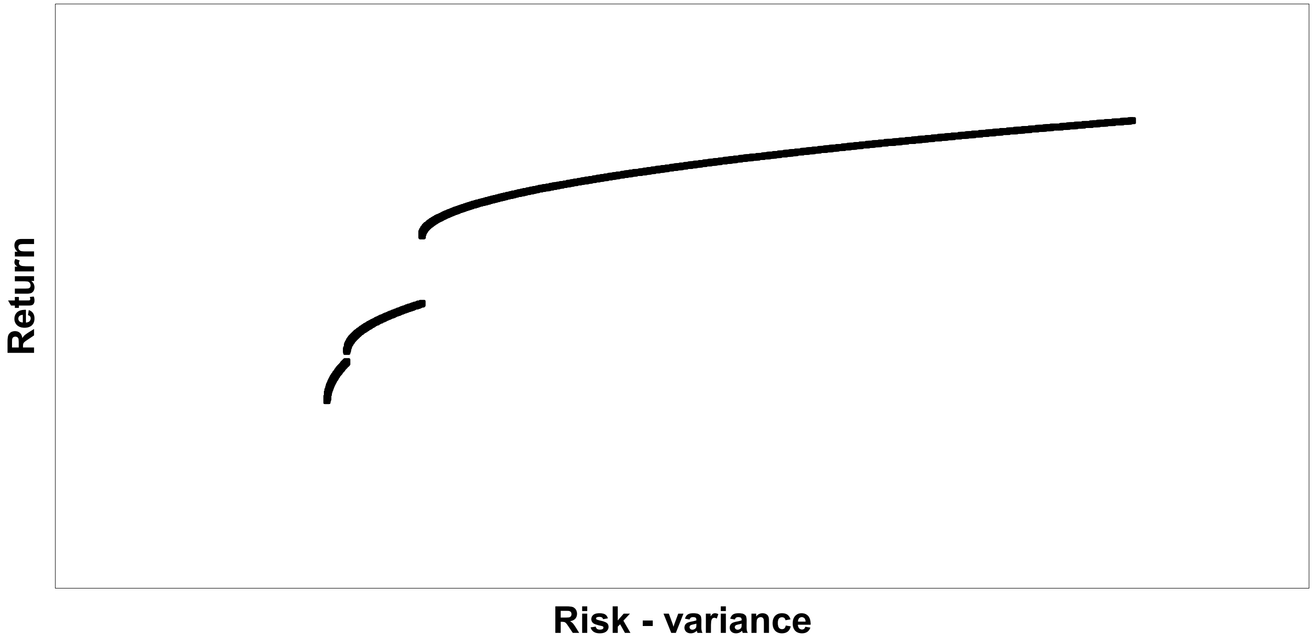

As an illustration of what can be achieved with cardinality constrained portfolio optimisation we show below some tradeoff curves for assets chosen from the DAX index.

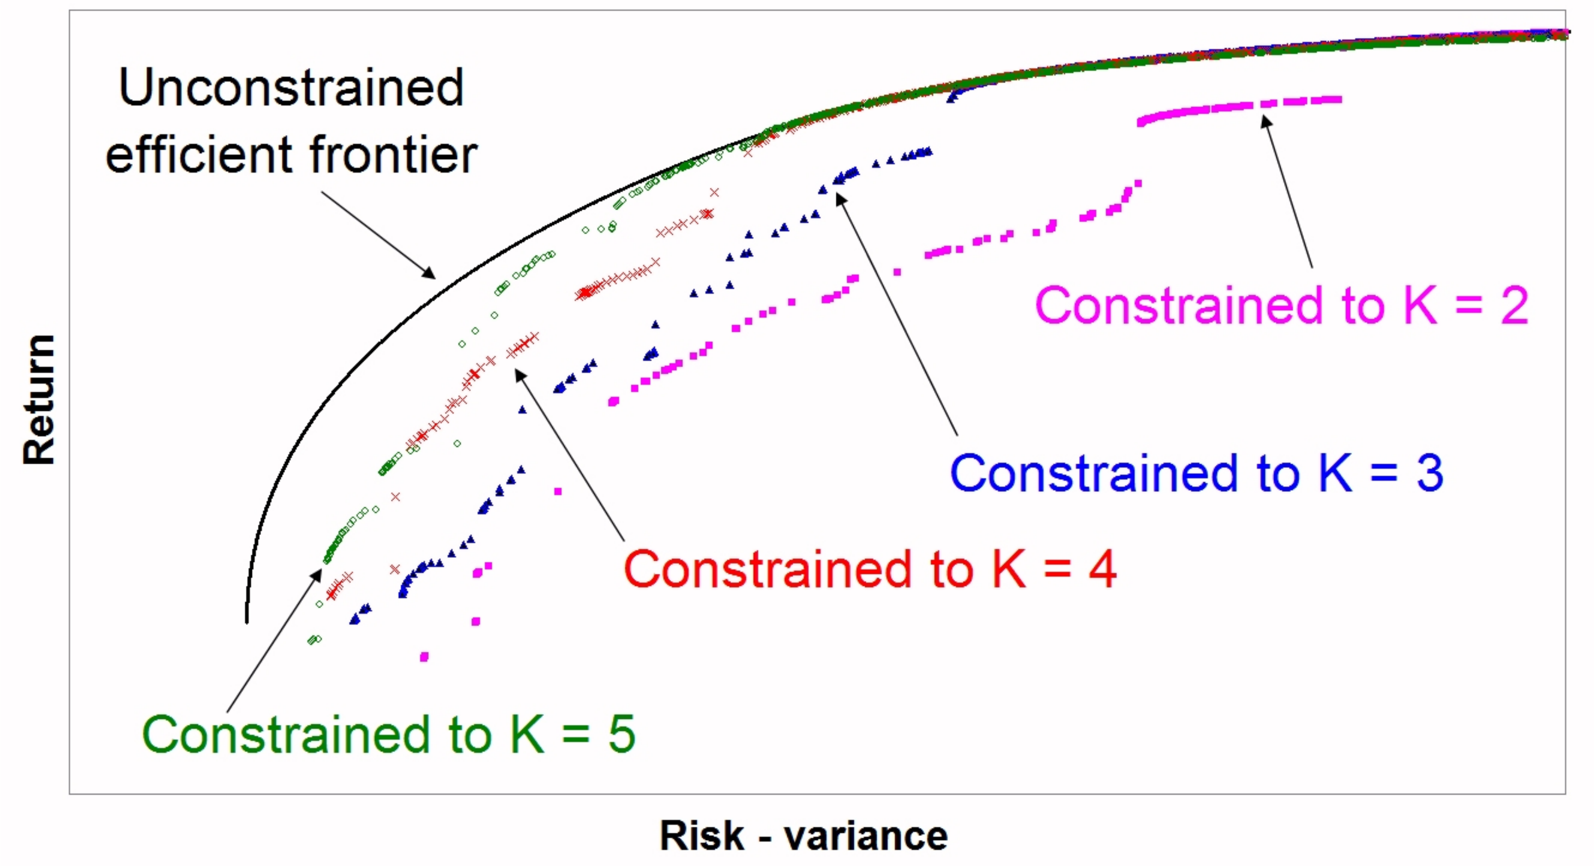

Consistent portfolios

Purpose: portfolios which perform out-of-sample in a fashion consistent with their in-sample performance

The logic behind Markowitz approaches is that we use known data (in-sample data) to construct frontiers of the types you have seen above and then choose a portfolio from the frontier to invest in. Our portfolio then varies in value as we hold it out-of-sample.

Criticisms of Markowitz approaches essentially often focus around the fact that, out-of-sample, portfolios chosen from the (in-sample derived) frontier do not behave in a fashion that their position on the frontier would imply.

Fact
If you decide to invest in a portfolio (e.g. one chosen from the Markowitz efficient frontier) you are ***explicitly*** investing in a particular in-sample distribution of returns.

Question
If you consider the returns you get out-of-sample from your chosen portfolio would you like those returns to be drawn from the same distribution as the in-sample distribution or not?

The picture below shows an approach that attempts to illuminate this.

In the green ***consistency region*** (shown using hollow squares) we have portfolios where we will typically get returns out-of-sample consistent with in-sample behaviour.

The red region, shown using solid squares, exhibits inconsistent behaviour.

Note that for this particular example most of the portfolios on the efficient frontier exhibit inconsistent behaviour.

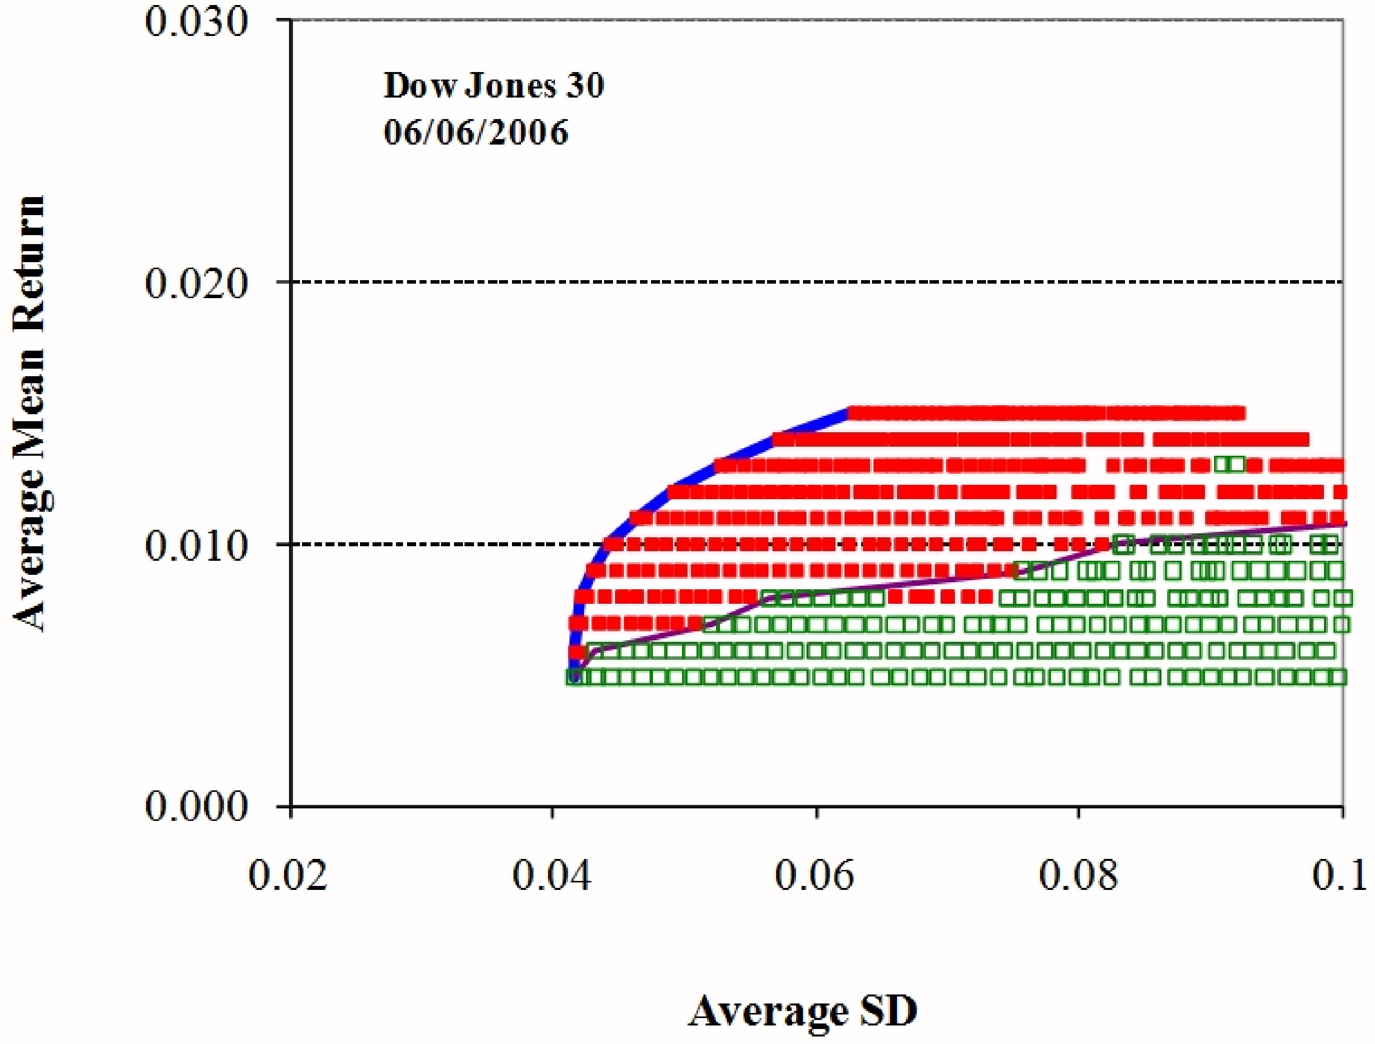
0.030
0.020
0.010
0.000
Dow Jones 30
06/06/2006
Average Mean Return
0.02
0.04
0.06
0.08
0.1
Average SD

Index tracking

Purpose: portfolios which give the same return as a specified benchmark market index

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 historic look-back approach. To ask the historic 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 immediate future.

Let:

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

I_t be the value of the index at time t

C be the amount we have to invest at time T

Then our decision variables are:

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

z_i = 1 if any of stock i is held in the 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:

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

$$V_{iT}x_i/C \leq z_i \qquad i=1,\ldots,N$$

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

$$x_i \geq 0 \qquad i=1,\ldots,N$$

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

These constraints are general in the sense that they also apply to the other portfolio problems we consider below, but these constraints will not (for space reasons) be repeated below.

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 mixed-integer nonlinear program, and so hard to solve. Algorithmically metaheuristics can be applied.

Alternatives

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 models that take ***different (but valid)*** views of the problem. Taking Markowitz mean-variance (for example, where we are seeking portfolios that balance risk and return) 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?

Answer: intercept $\alpha = 0$

slope $\beta = 1$

Answer: intercept $\alpha = 0$

slope $\beta = 1$

Optimisation model: minimise $|\alpha - 0|$

minimise $|\beta - 1|$

Optimisation model: minimise $|\alpha - 0|$

minimise $|\beta - 1|$

For example here we could adopt a two-stage approach and first minimise $|\beta - 1|$ to achieve an optimal value for β of β^*; then minimise $|\alpha - 0|$ but setting $\beta=\beta^*$ to retain the optimal value of β achieved at the first stage.

Although this optimisation model seems nonlinear in fact we can linearise it so as to end up with a (mixed-integer) linear program that computationally is easily solved.

Consider the two different mathematical optimisation models for index tracking above:

- one a nonlinear model based on minimising the average squared difference between tracking portfolio return and index return
- one a linear model based on regression, seeking to achieve regression coefficients for the tracking portfolio, when its return is regressed against index return, as close to ideal values (intercept zero, slope one) as possible

Each of these models had their own logic and each, by themselves, seem perfectly reasonable. However there is limited overlap between them. One uses regression coefficients for example, the other does not. This is an illustration of one of the key concepts we introduced at the start of this tutorial paper: ***even for portfolios intended for the same purpose the model to use is not uniquely defined.***

In fact there are other alternative models for index tracking.

For instance instead of using least-squares (mean) regression (which as a statistical technique dates from the late 1800's) use quantile regression (dating from the late 1970's).

More specifically adopt a regression model for index tracking, but based on median (50% quantile) regression.

Enhanced indexation – specified excess return

Purpose: portfolios which achieve a specified excess return with respect to a given benchmark market index

- 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 A_t

So here for example we could apply our first index tracking model above directly and minimise $\sum_{t=1}^{T}(r_t - A_t)^2/T$ to find an enhanced indexation portfolio.

But again alternatives exist, specifically:

1. use least-squares mean regression, as above for index tracking, but with respect to A_t, seeking a regression slope of one and a regression intercept of zero (so seeking $r_t = A_t$)

2. use median quantile regression, as above for index tracking, but with respect to A_t, seeking a regression slope of one and a regression intercept of zero

Question

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

Can you see how to use this to decide the composition of an exchange traded fund (ETF) portfolio that gives a multiplier λ of the return on a specified index?

Question

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

Can you see how to use this to decide the composition of an exchange traded fund (ETF) portfolio that gives a multiplier λ of the return on a specified index?

Answer Set $A_t = \lambda R_t$

e.g. λ=2, λ=-1

Enhanced indexation - outperform

Purpose: portfolios which do better than a given benchmark market index

Here one objective that can be used is a modified Sortino ratio, namely:

$$\text{maximise } \left(\sum_{t=1}^{T} r_t/T - R^{mean}\right)/\sqrt{\left[\sum_{t=1}^{T} (\min(0, r_t - R^{mean}))^2/T\right]}$$

$$\text{where } R^{mean}=\left(\sum_{t=1}^{T} R_t/T\right) \text{ is the mean return on the index}$$

The table below shows some results for this objective, solved using a metaheuristic algorithm. Here we consider the S&P Global 1200 index (over the period 1999-2006), where we are choosing a portfolio of 100 stocks and holding that portfolio for a specified period before rebalancing.

Holding period in weeks	Out-of-sample return (% per year) in excess of the index
4	10.5
12	12.3
24	11.2
36	12.0
48	9.5

My experience has been that (subject to certain qualifications) it is not difficult to develop decision models based on optimisation which can find portfolios that (out-of-sample) outperform an index.

Typically what you cannot do (**or at least I cannot do**) is to have precise control over the level of outperformance.

Absolute return/market neutral portfolios

Purpose: portfolios which do well irrespective of how a benchmark market index performs

The terms “absolute return portfolio” and “market neutral portfolio” tend to be used fairly interchangeably in finance. Here we adopt an academic perspective and distinguish between absolute return/market neutral in the sense that:

- we regard an absolute return portfolio as one that (ideally) gives a constant return per period
- we regard a market neutral portfolio as one that (ideally) has zero correlation between portfolio return and index return

Absolute return portfolio

Suppose we perform a linear regression of the return from the portfolio against time, i.e. the regression $r_t = \alpha+\beta t$. What intercept α and slope β would you like if you are seeking an absolute return portfolio that (ideally) gives a constant return per period?

Absolute return portfolio

Suppose we perform a linear regression of the return from the portfolio against time, i.e. the regression $r_t = \alpha + \beta t$. What intercept α and slope β would you like if you are seeking an absolute return portfolio that (ideally) gives a constant return per period?

Answer: slope $\beta = 0$
intercept α as large as possible

Absolute return portfolio

Suppose we perform a linear regression of the return from the portfolio against time, i.e. the regression $r_t = \alpha+\beta t$. What intercept α and slope β would you like if you are seeking an absolute return portfolio that (ideally) gives a constant return per period?

Answer: slope $\beta = 0$
intercept α as large as possible

Optimisation model: minimise $|\beta - 0|$
maximise α

This can be expressed as a (mixed-integer) linear optimisation model

The figure below shows some out-of-sample results for this approach, for the S&P Global 1200 index. Here (over the period 1999-2006) we are choosing a portfolio of 250 stocks and holding that portfolio for 26 weeks before rebalancing.

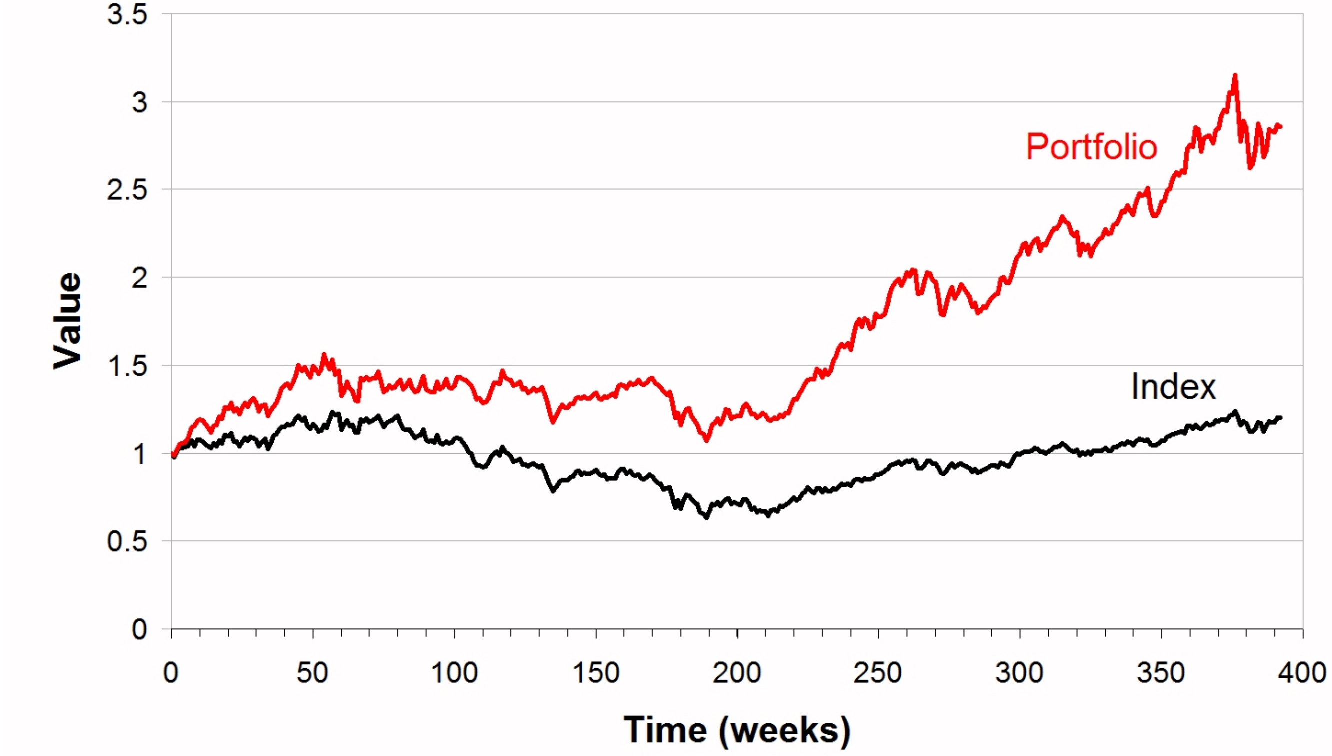

Market neutral portfolio

Suppose that we calculate the correlation coefficient Γ between portfolio return r_t and index return R_t. What value for the correlation coefficient would you like if you are seeking a market neutral portfolio?

Market neutral portfolio

Suppose that we calculate the correlation coefficient Γ between portfolio return r_t and index return R_t. What value for the correlation coefficient would you like if you are seeking a market neutral portfolio?

Answer: $\Gamma = 0$

Market neutral portfolio

Suppose that we calculate the correlation coefficient Γ between portfolio return r_t and index return R_t. What value for the correlation coefficient would you like if you are seeking a market neutral portfolio?

Answer: $\Gamma = 0$

Optimisation model: minimise $|\Gamma|$

This is a mixed-integer nonlinear program

The table below shows out-of-sample statistics for market neutral portfolios (derived using the model above) with respect to a number of benchmark indices.

In that table we show results for Long Only portfolios and for portfolios with a fixed mix of Long/Short positions (so for 130/30 we have 30% of the investment in short positions, 130% in long positions).

The table shows the out-of-sample values (based on a holding period of 13 weeks) for correlation and portfolio excess return (return over and above the index) over 30 rebalances for three indices.

	Out of sample	
Long Only	Correlation	Excess return (% p.a.)
S&P Europe 350	0.448	20.5
S&P US 500	0.494	11.5
S&P Global 1200	0.330	32.0
130/30		
S&P Europe 350	0.263	10.2
S&P US 500	0.319	17.1
S&P Global 1200	0.264	-9.5

Key concepts

Different types of portfolios require different mathematical models

Even for portfolios intended for the same purpose the model to use is not uniquely defined

Adopt an optimisation mindset in building/choosing a model

Nonlinear models are more difficult to solve numerically than linear models

Thank you for your attention

Any questions?