Optimising and Constraining Portfolio Distances

Patrick Burns
http://www.burns-stat.com

June 2010

The boring title.

This talk was given at the London Quant Group Investment Technology Day on 2010 June 21.

Goldilocks and the 3 Little Pigs

Patrick Burns
http://www.burns-stat.com

June 2010

The real title.

Fundamental Fund Manager

The 2010 Spring Seminar of LQG included a debate on optimisation versus not optimising. I'll provide a synopsis of that debate, but I want to do that from the perspective of a fundamental fund manager.

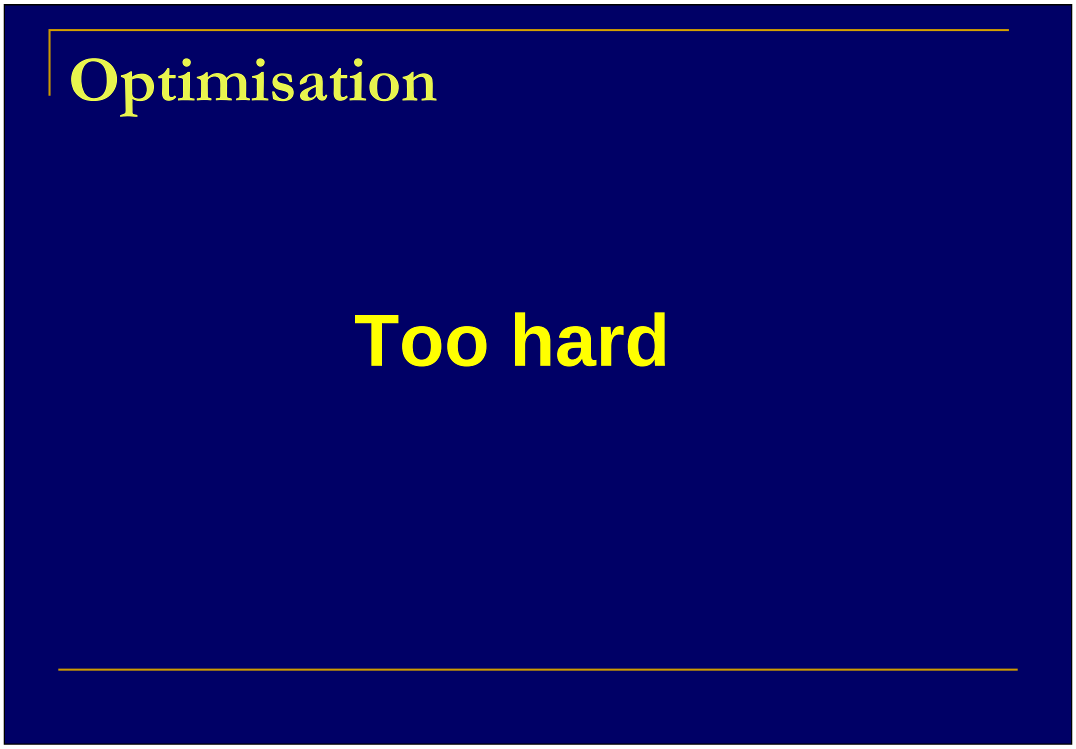

Sebastian Ceria of Axioma argued for optimisation.

His position can be summarised as: we put everything that we know into the machine and press the button, how can that not be the best that we can do?

We fundamental fund managers are likely to say that it’s too hard. In particular, we don’t know how to create expected returns. (And we don’t think that quants can do it either – they just won’t admit it.)

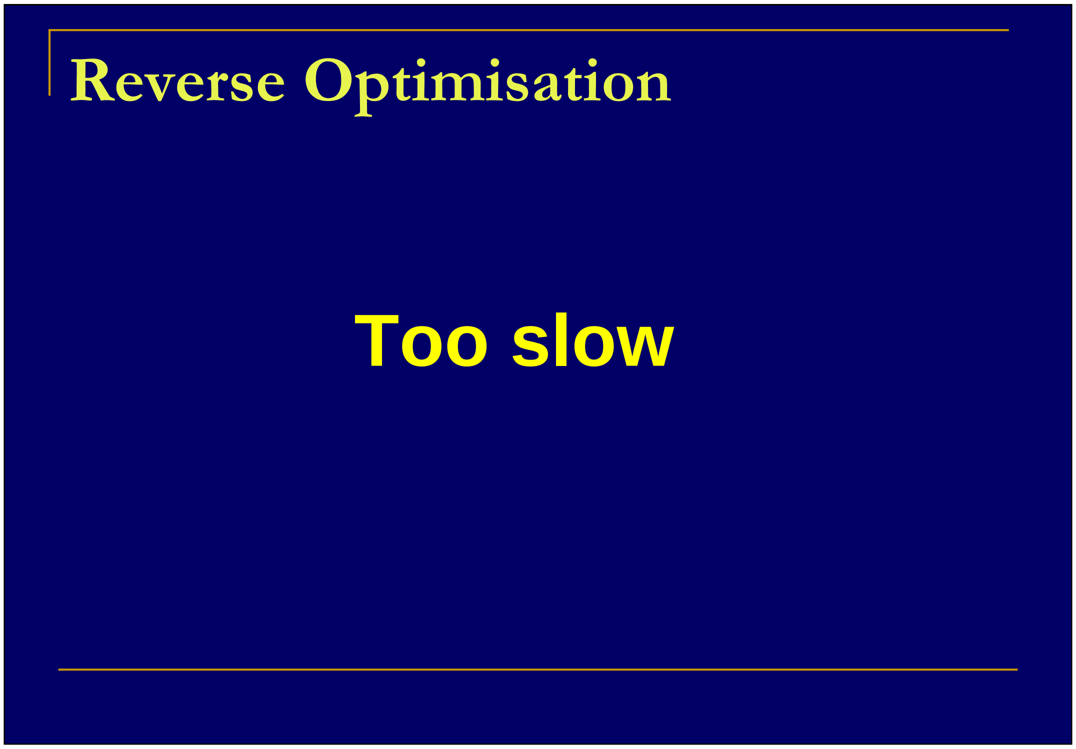

Jason MacQueen of R-squared argued against optimisation.

He agrees that creating expected returns is too hard. His alternative is to use the positions of the portfolio and a variance matrix to create a set of implied expected returns. If those implied returns don't match with our view, then we submit proposed changes to our portfolio until we get implied expected returns that we do like.

We fundamental fund managers are likely to think this is too slow. It is an iterative process, an iterative process that involves contact with nerds. (And we hear rumour that quants are fighting among themselves about how well the procedure actually works.)

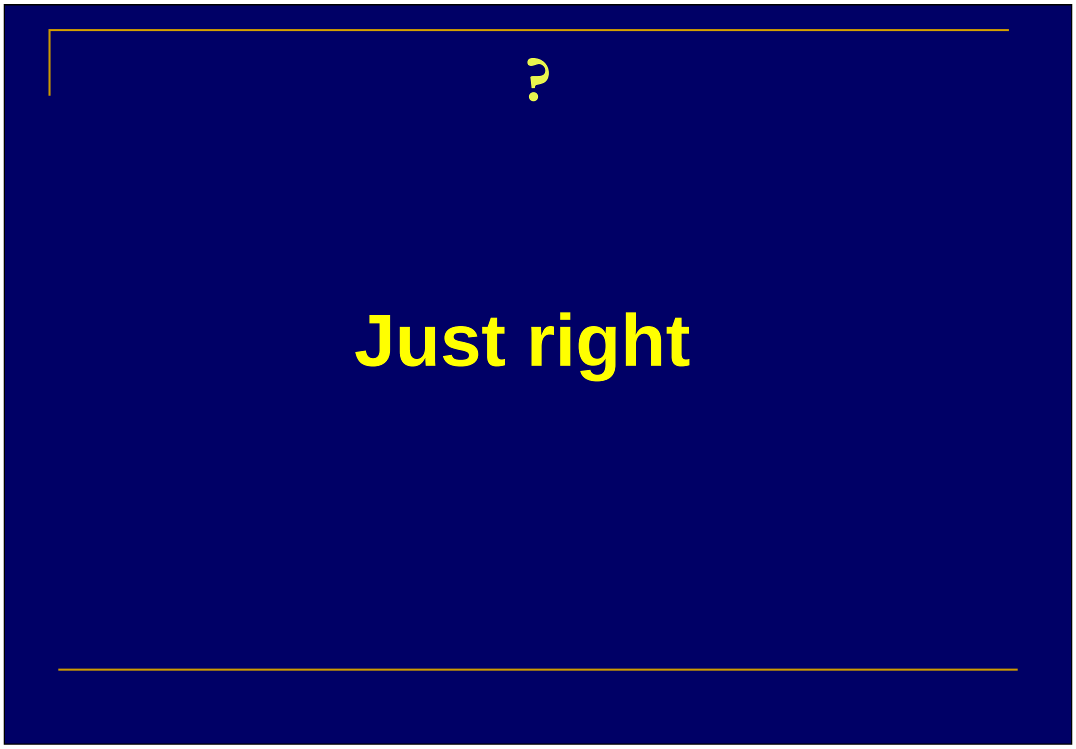

We are looking for something that is “just right”. What might that something look like?

I don’t know, but I do have a proposal.

Value Distance

Amount of money in the trade to get from one portfolio to the other

We need a couple of definitions before we get to the proposal.

We need a sense of distance between two portfolios.

One way of thinking of the distance is the amount of money (buys plus sells) that it takes to get from one portfolio to the other.

Weight Distance

Sum of absolute differences of weights

An alternative definition is the sum of the difference in weights between the two portfolios.

If the gross value of the two portfolios are equal, then these two definitions are equivalent.* But they are different for portfolios that differ in size.

In particular, the weight distance can be zero but the value distance positive. For example, if the second portfolio has double each of the positions of the first, then the weights are equal but the value distance is the gross value of the first portfolio. This will be important later.

* For long-short portfolios I believe the equivalence depends on the definition of weight to be position value divided by gross value.

Utility-free Optimisation

Minimise distance to Target Portfolio while still obeying constraints

Software is: Portfolio Probe

Fundamental fund managers know the assets that they like. If pushed, they can probably be persuaded to put weights on those assets to create an ideal Target Portfolio.

My proposal is to find the portfolio that is closest to the target portfolio subject to the constraints being obeyed. One of the constraints might be (or probably should be) turnover. Or perhaps it would be a constraint on the transaction costs.

The software that does this, more standard optimisation and generates random portfolios is Portfolio Probe.

That application is what I think might generate some sales. But that is not why I thought of portfolio distances.

It might be thought that I sit around with my Planning Committee (pictured above) and we come up with ideas that are most likely to create sales. Actually I sit around with my Planning Committee and we try to come up with the toys that I'd most like to play with. Once I've played with them, I'm willing to share.

In order to tell you about the toy I wanted to play with, I need to tell you a story.

A story that involves pigs.

Random Portfolios

- …
- Performance Measurement

The story also involves random portfolios.

The idea of random portfolios is very simple. We are just taking a sample from the population of portfolios that obey a given set of constraints.

There are numerous applications of random portfolios – many of which have yet to be discovered.

The killer application for random portfolios is performance measurement.

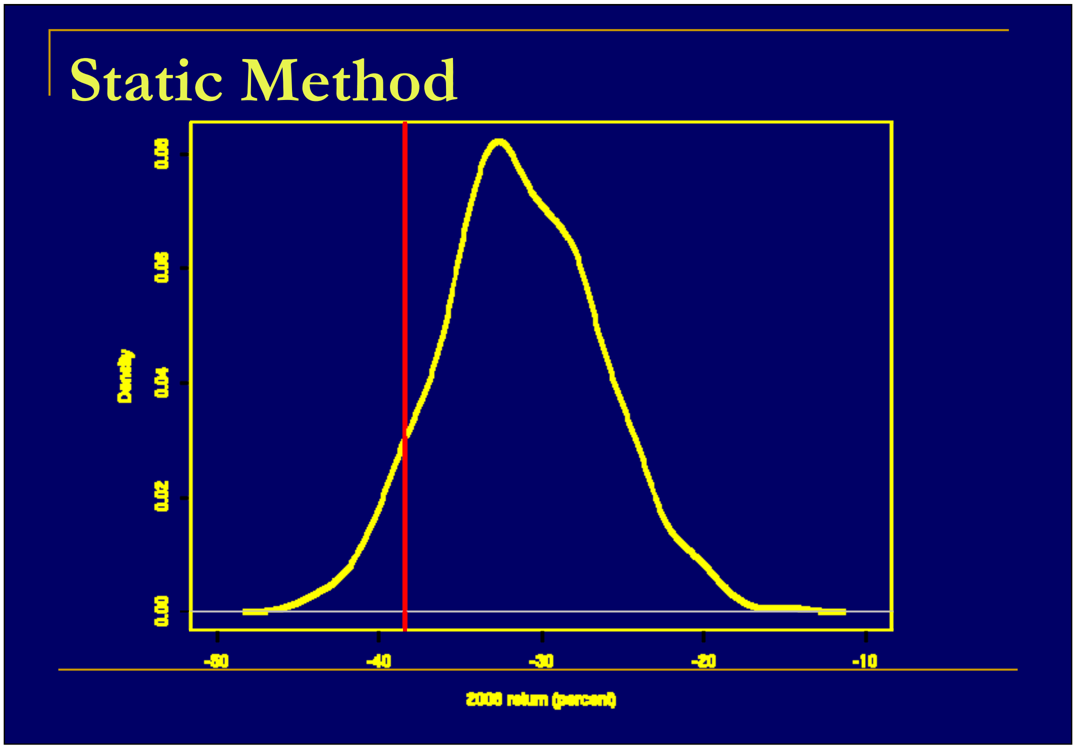

The static method of using random portfolios for performance measurement is very easy to do.

We have a time period of interest – in this case it is the year 2008.

We generate a number of random portfolios that obey the constraints of the fund at the beginning of the time period. We hold those portfolios throughout the time period. We find the distribution of returns of the random portfolios, and see where the return of the actual fund fits in.

In this case it doesn't do very well. This is a momentum strategy so we wouldn't expect it to do very well in 2008.

Note that there is no problem using risk-adjusted returns or some other utility instead of returns.

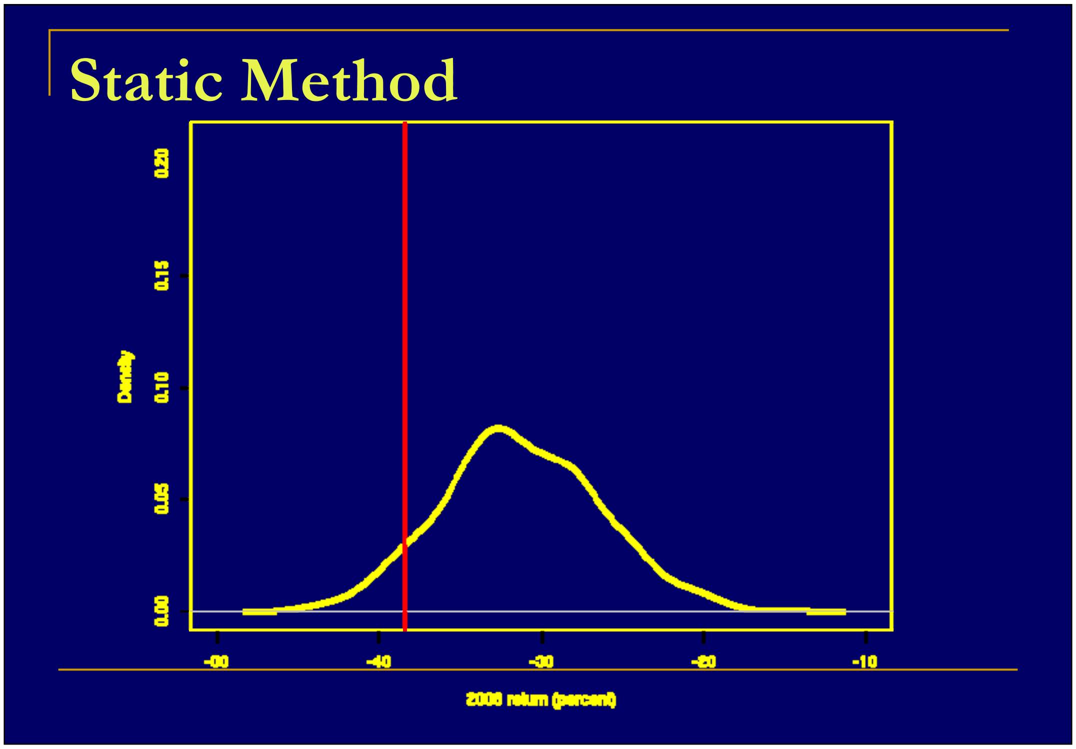

There is a second approach that is more powerful, so we'll rescale the picture.

The second approach is more powerful because we use more information.

Sometimes we know some or all of the positions of the fund at the beginning of the time period, and we have a sense of the turnover of the fund throughout the period.

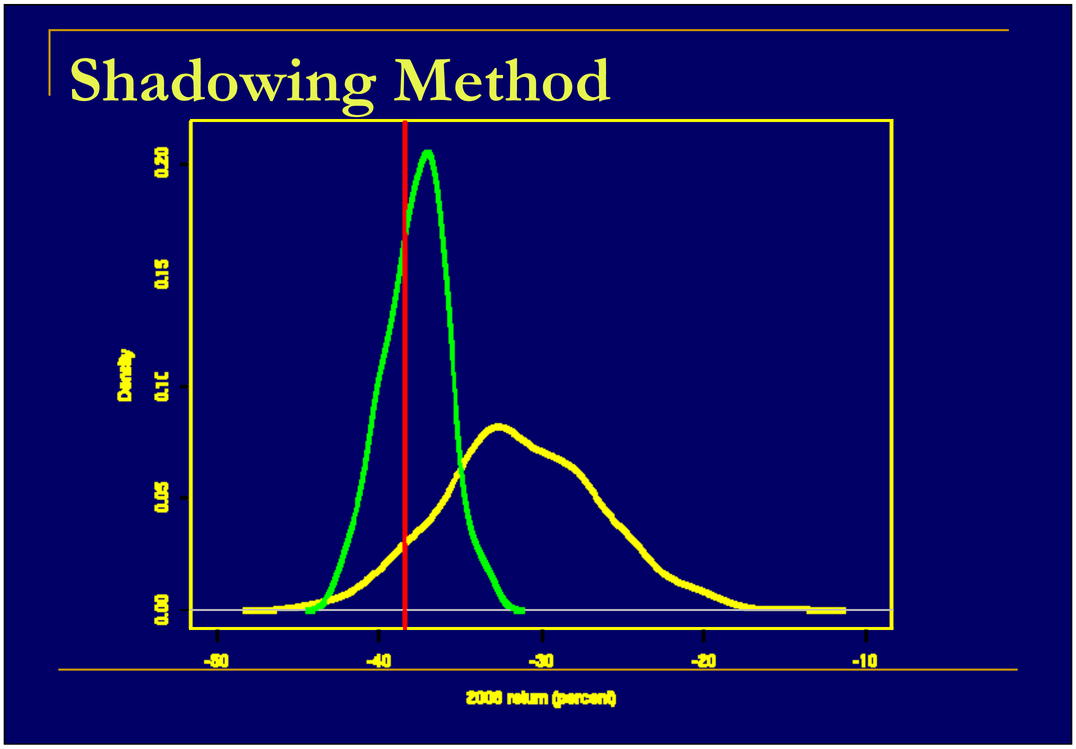

In this case we know all of the positions at the start of the period, and the turnover is 400% (buys plus sells) evenly spaced throughout the year.

To get one random path we start where the fund does at the beginning of the period, make random trades throughout the year obeying the constraints (and the turnover).

We do that a number of times to get the distribution of returns.

In this case the fund still doesn't do very well, but it is better in this distribution than in the static distribution.

The green distribution is the most important thing in the talk. This is what gives really useful performance measurement. This is what can transform fund management.

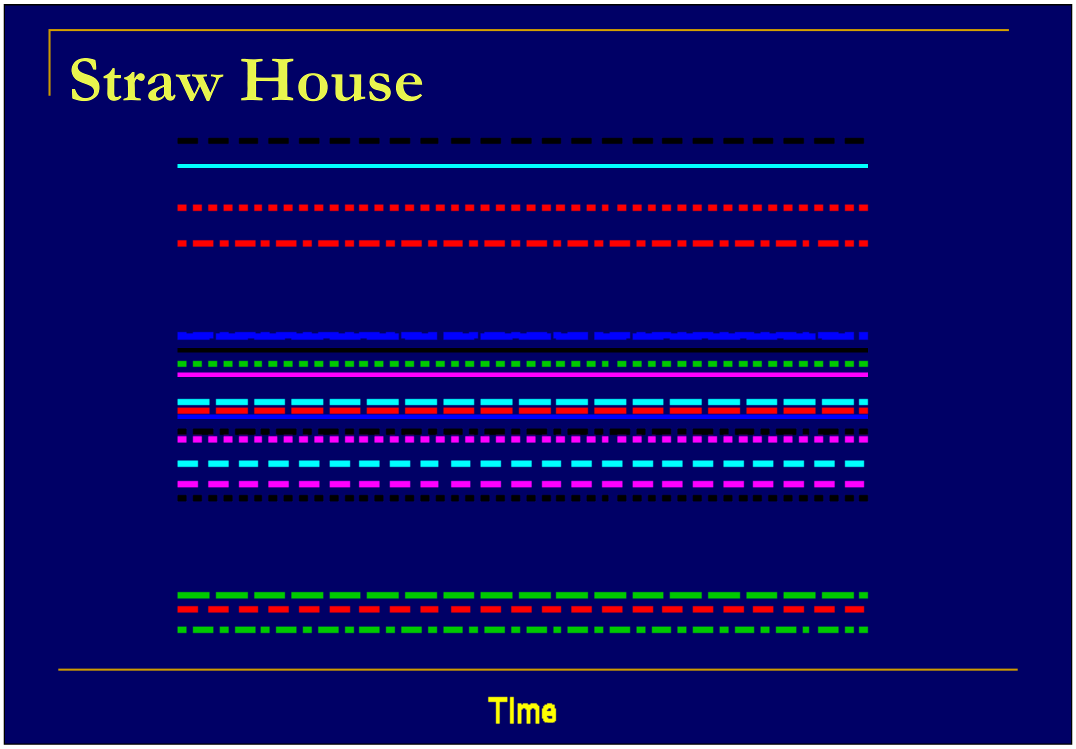

Let's look at this from a different perspective.

This is a sketch of the static method. The x-axis is the time in the time period of interest. Think of the y-axis as the position size of a typical asset. The different lines are different random paths.

Note that the y-axis is position size, not weight. If we plotted weight, then the lines would drift about because the prices change throughout the period.

The shadowing method starts at one particular position and then drifts away.

This picture works whether we think of the y-axis as position size or as weight.

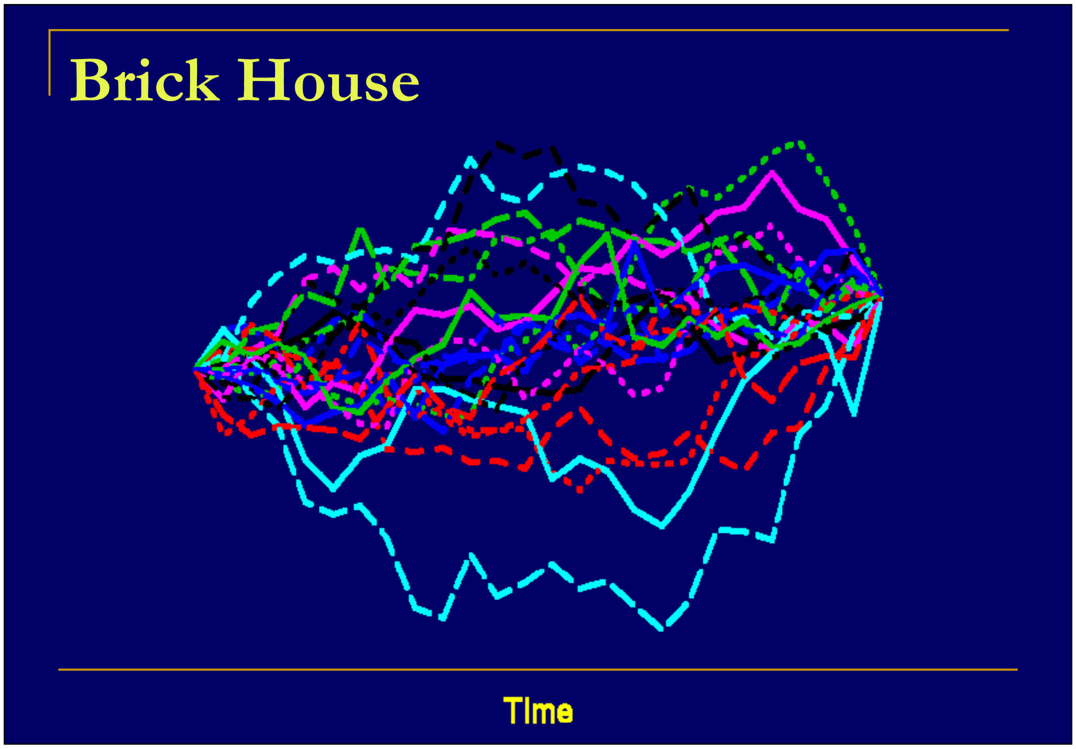

The toy I wanted to play with is depicted here. What if we knew not only the starting portfolio but also the ending portfolio. How would we generate random paths that start at a particular spot and end with a particular set of weights?

Note that we've changed the meaning of the y-axis for this plot, it is now weights and not position.

I realised that the key thing in generating such paths was that the distance to the final portfolio had to converge to zero.

Pictured here is the distance profile for our example. It is of interest relative to what we'll see later that the distance is relatively large until the end, and then goes down steeply.

We can generate paths just like with the shadowing method except with one additional constraint: the distance at each point in time is no more than the distance profile (plus epsilon which in the present case was 20 basis points).

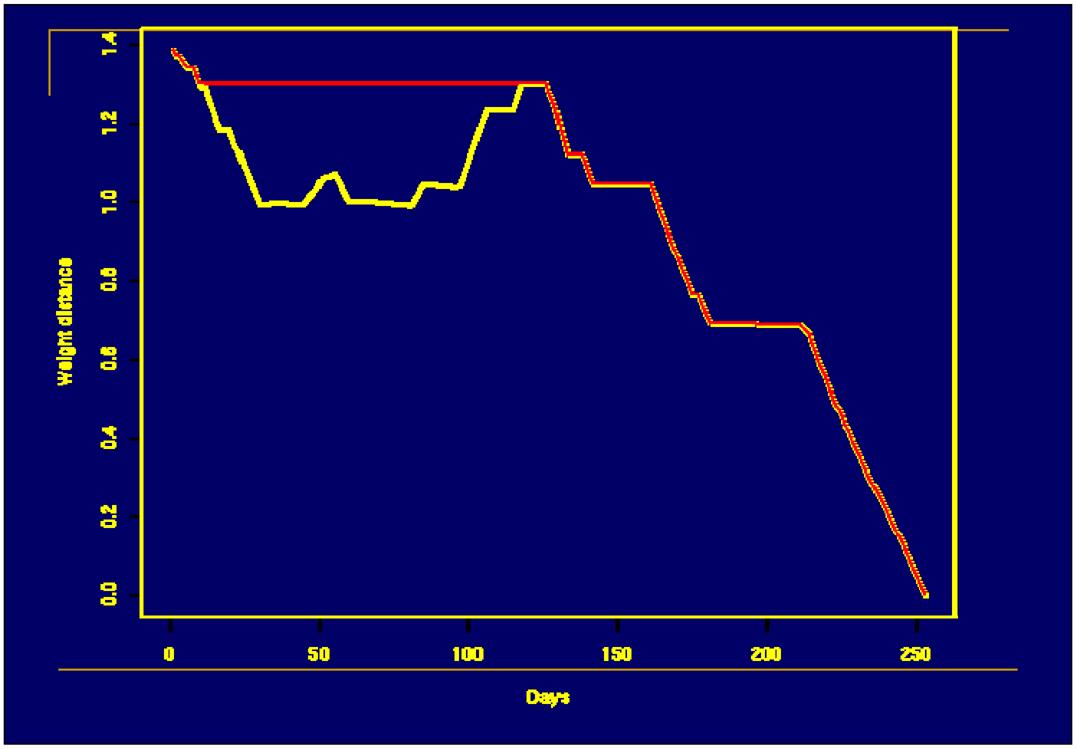

It doesn’t make sense to constrain the distance to be smaller than it will be later, so we really want to use the convex hull of the distance profile for our constraints.

This assumes that we know the distance profile when we are generating our random paths. I have some ideas about what to do when we don’t know the profile, but suggested approaches are welcome.

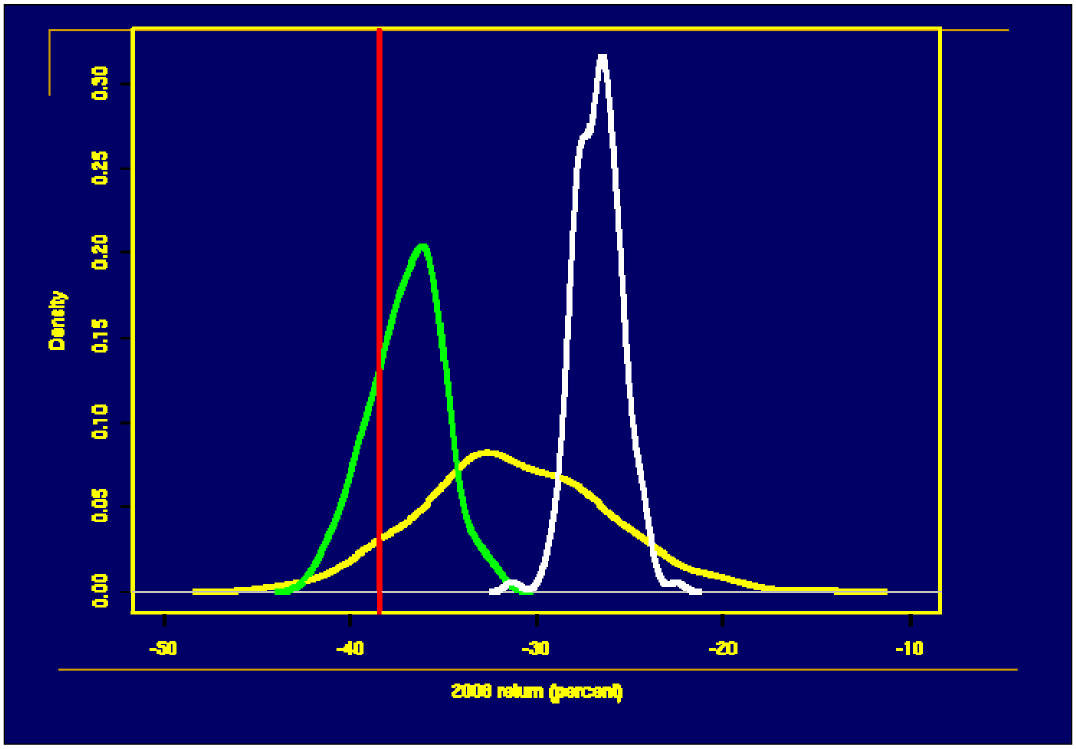

The white distribution is that of the random paths with both the starting and ending portfolios known.

This is doing some sort of performance measurement. But what is it measuring?

It is sort of market timing, but I don't think quite market timing.

If we had selected the ending portfolio in advance (say, 60% equities and 40% fixed income), then it would be measuring market timing.

But the ending portfolio was just where we happened to end, there was no desire to actually end up there.

As far as I can see the white distribution really is just a toy. Perhaps people will come up with cogent uses of the technique, but I don't think I have.

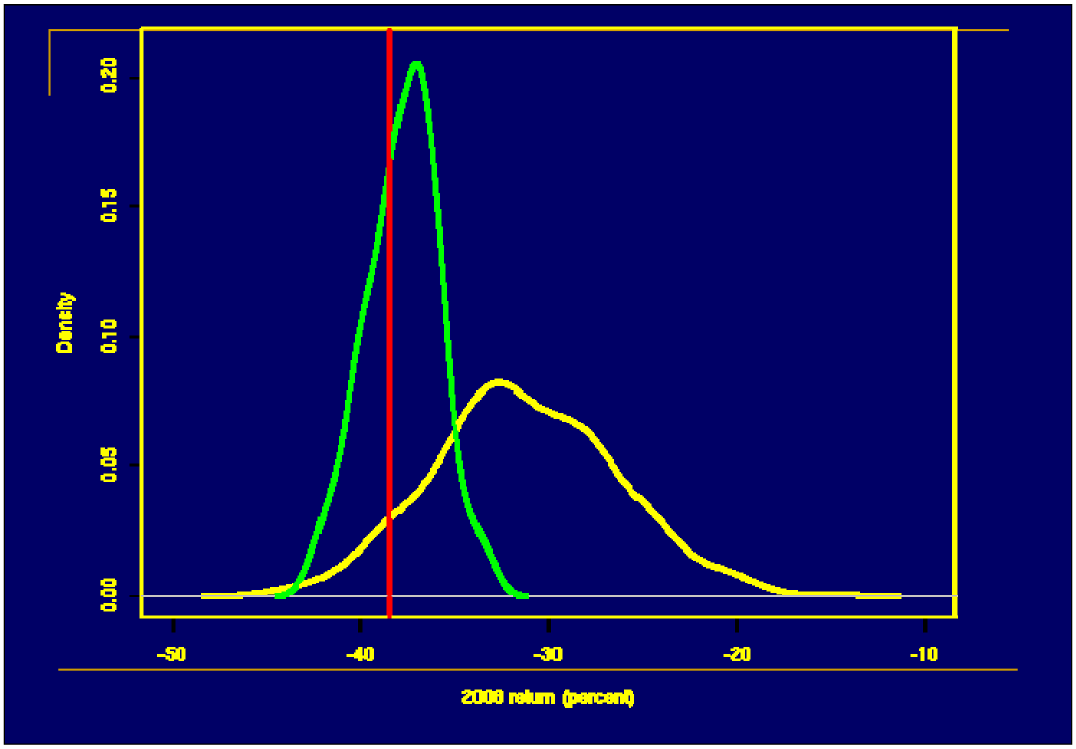

In this picture the red line shows the return of the strategy.

But for the moment pretend that it is the return of the original portfolio for the period. If that were true, then the picture is just what we should expect:

The original portfolio has a return that is in the lower tail of the static distribution. If we deviate slightly from that portfolio, we are likely to go towards the centre of the static distribution. There is a regression effect. So the centre of the green distribution is between the return for the original portfolio and the centre of the static distribution.

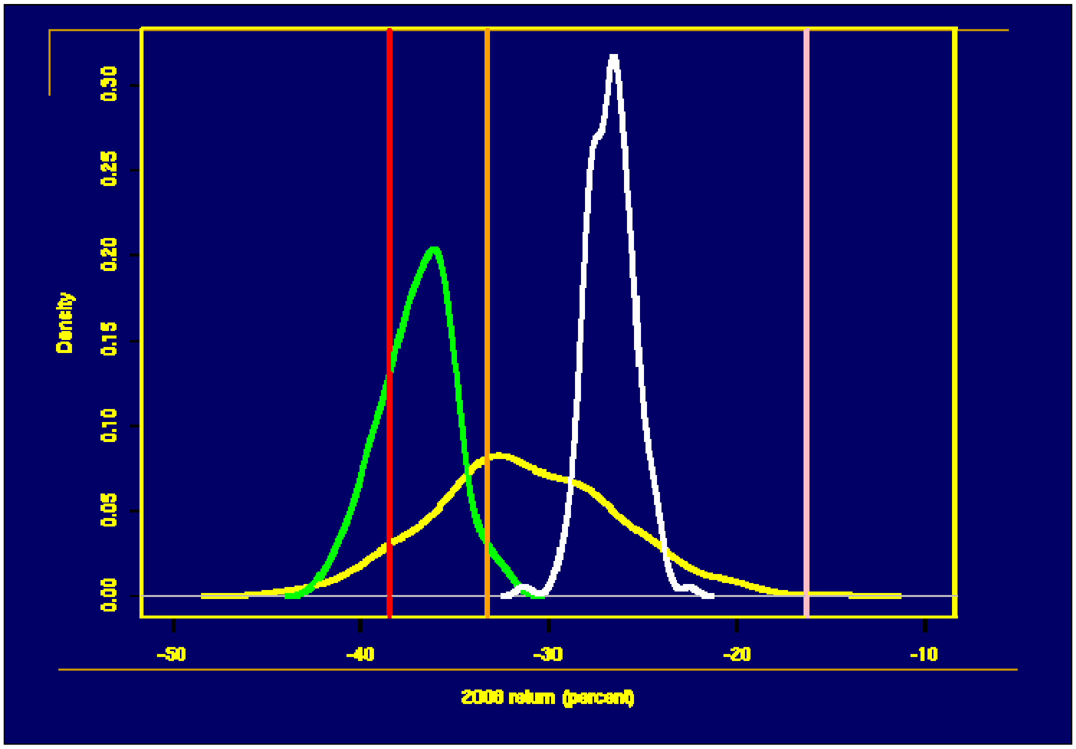

Back to reality. The return of the original portfolio is actually the orange line, and the return of the final portfolio is the pink line.

So the white distribution is in a quite logical location between the return for the original portfolio and the return of the final portfolio.

But the green distribution is weird. The return of the original portfolio is roughly in the centre of the static distribution. So we would expect the green distribution to be sort of centred on the orange line. It isn't.

I don't have any more of an explanation of this than anyone else. All I know to say is that 2008 was very volatile so there is a lot more latitude for unusual results.

A couple people from the audience suggested that it makes sense that the final portfolio has a return in the upper static tail because it is a momentum strategy. The momentum time frame plays a part though.

And they all lived happily ever after

Except that Goldilocks is in a spot of bother.

What is Optimal?

Target: $wt_A = 50\%$, $wt_B = 50\%$

Constraints: Long-only

$wt_A + wt_B <= 40\%$

The “right” answer for this is to allocate 20% to each of A and B.

But with the distance we have, an allocation of 40% A and 0% B is optimal as well. As is 0% A and 40% B. Actually anything in between these is optimal. (This is precisely the same as the result that Least Absolute Deviation regression need not have a unique answer.)

We can ensure the “right” answer for this problem by minimising the sum of squared weight differences.

What is Optimal?

Target: wt_A = 75%, wt_B = 25%

Constraints: Long-only

$wt_A + wt_B \leq 40\%$

Here is another problem with the same constraints but a different target. The “right” answer is 30% A and 10% B.

The least squares answer is 40% A and 0% B.

To get the “right” answer in this case we need weighted least squares with the weights equal to one over the target weight.

And they all lived happily ever after

Really