Portfolio Optimisation Inside Out

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

Given 2011 December 19 at the Computational and Financial Econometrics conference, held jointly with the conference for the European Research Consortium for Informatics and Mathematics, in London.

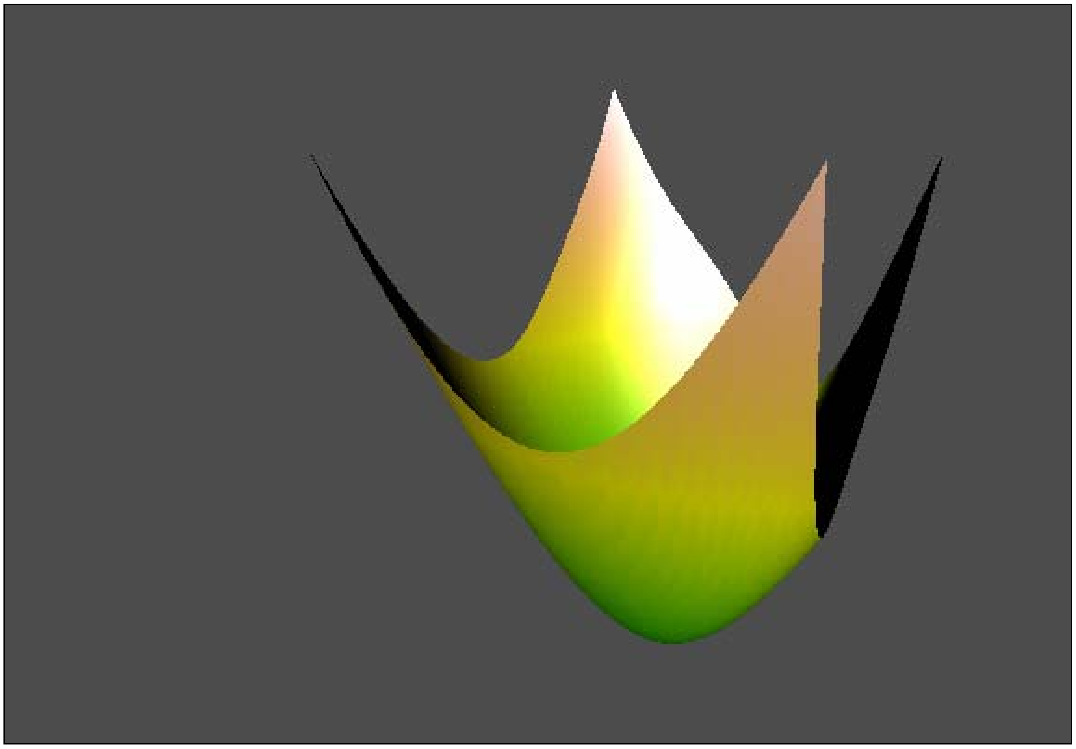

The usual way of approaching portfolio optimization is to think of a utility.

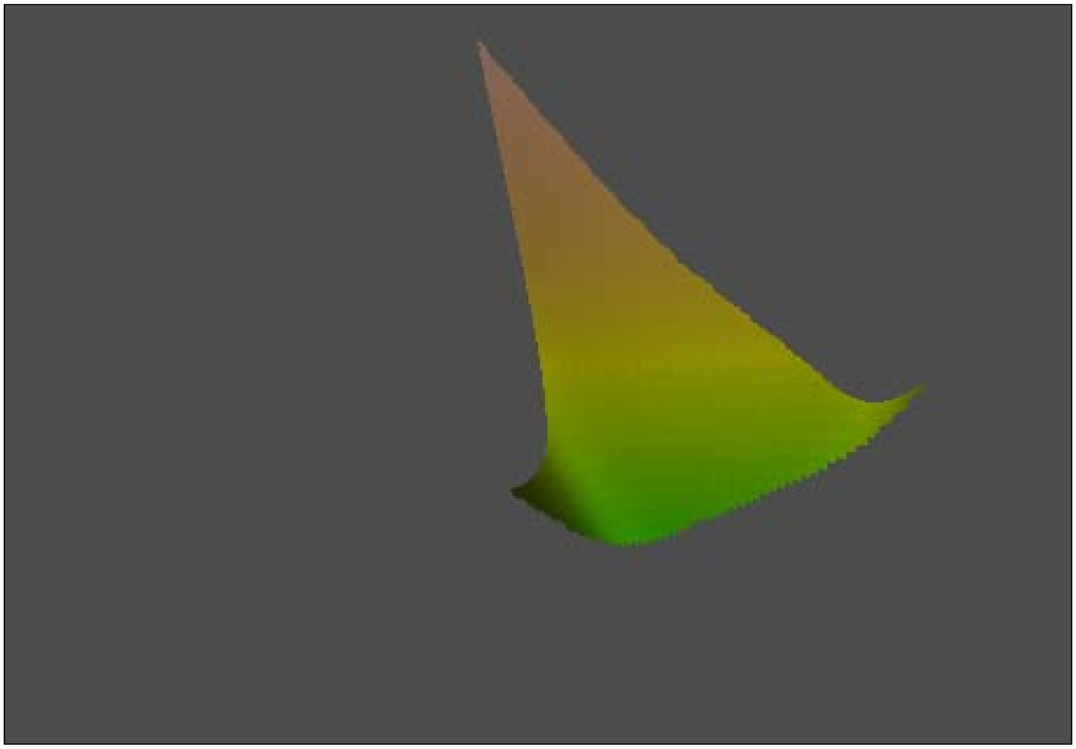

Then pare away what is outside the constraints.

Utility is king, constraints are secondary.

That’s a perfectly good way to look at the problem. But not the only way.

Constrained weights

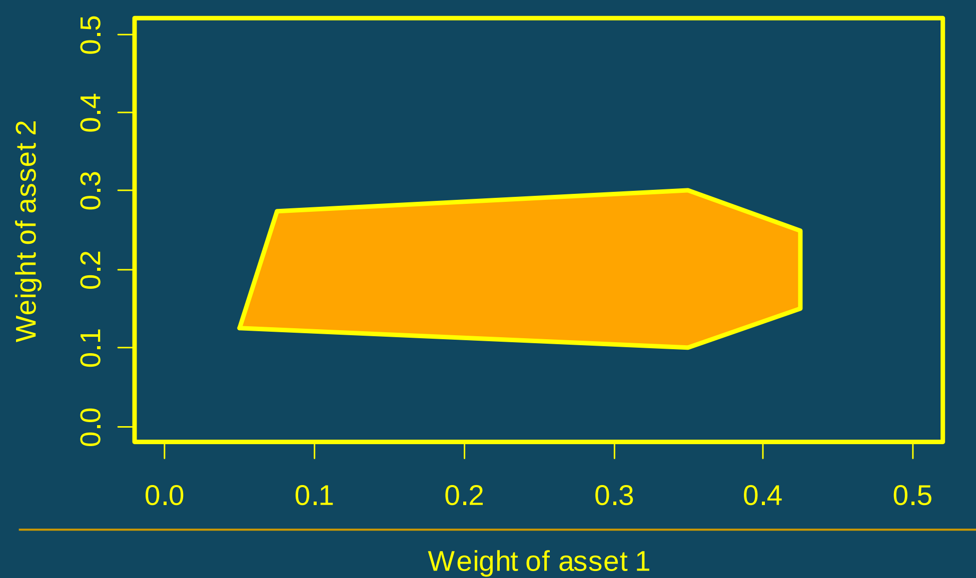

We're going to take a different view – we'll be in a world where the constraints are always met, and the utility is secondary. Anything outside the constraints will be invisible to us.

We'll have a local view rather than a global view.

Because we take a different point of view, we will see different things.

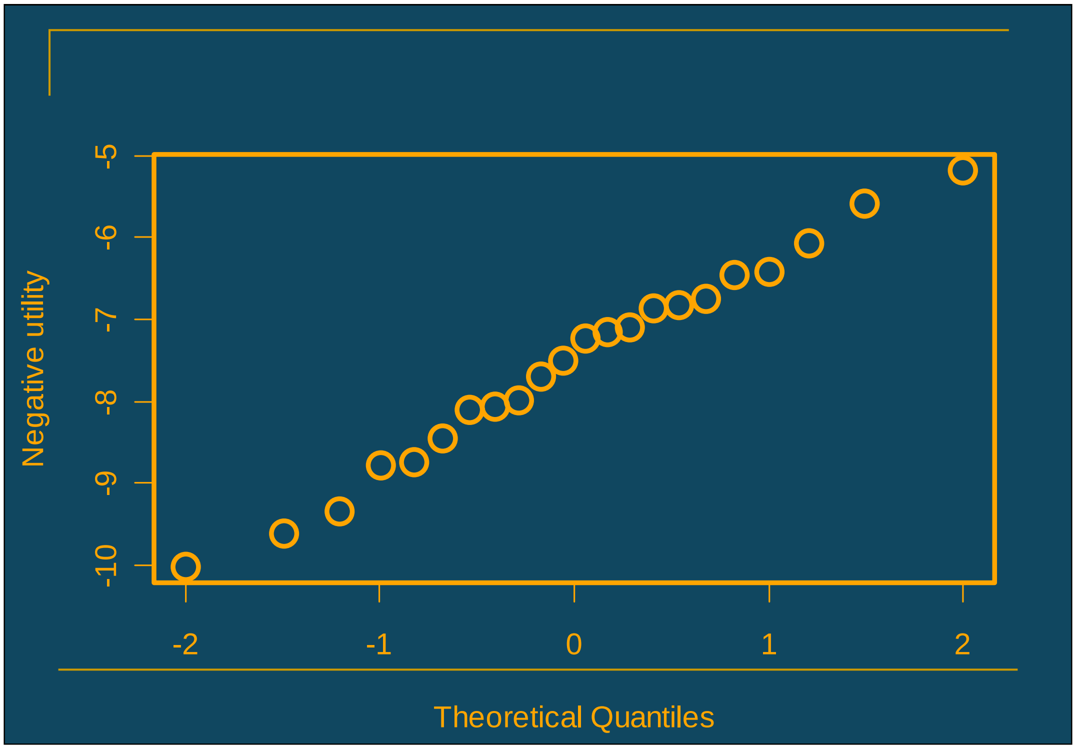

Here is a complete portfolio optimization. We have 22 trades to choose from. (Some people might think this problem is over-constrained but what do they know.)

We have been trained to think of optimization as being about gradients and moving through space. But really in portfolio optimization what we are doing is going through the store with our trolly looking on the shelves for the best trade to select.

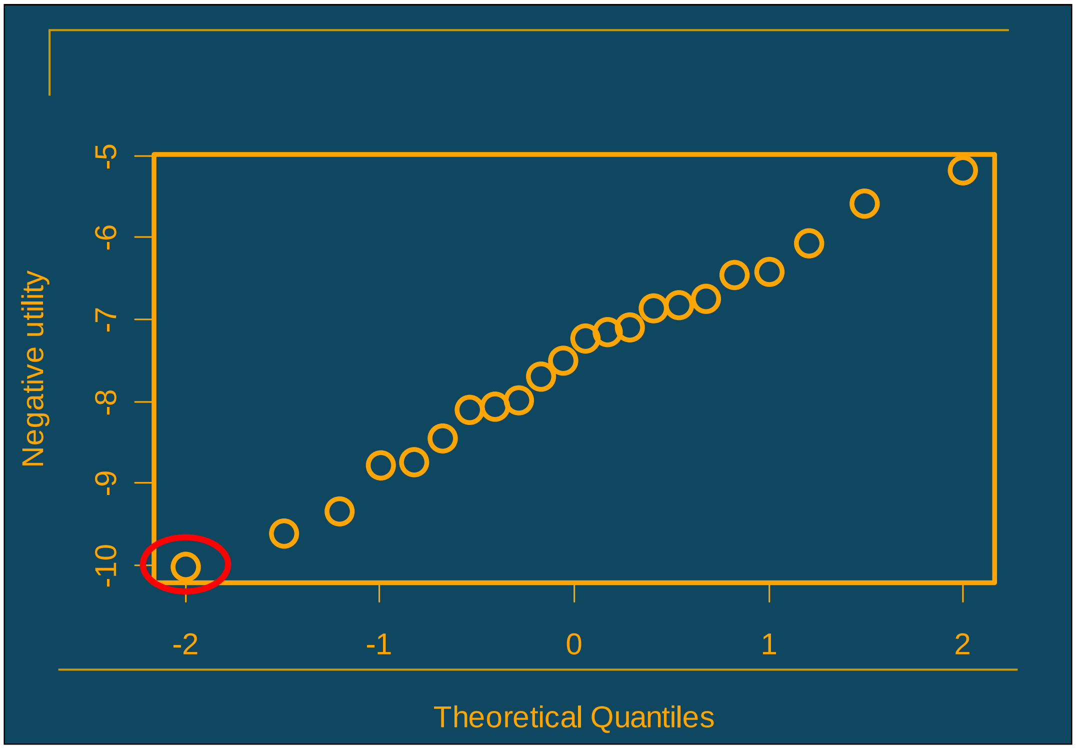

The trade we are going to select is the one circled in red.

Optimization by convention is minimizing, so we are minimizing negative utility. In this problem we are maximizing the information ratio.

We are quite optimistic – we think we'll get an information ratio of 10.

The x-axis on the plot is arbitrary, but in this case it is the quantiles from a Gaussian distribution.

Once we have the trade in our trolly, we head for the cashier and then home.

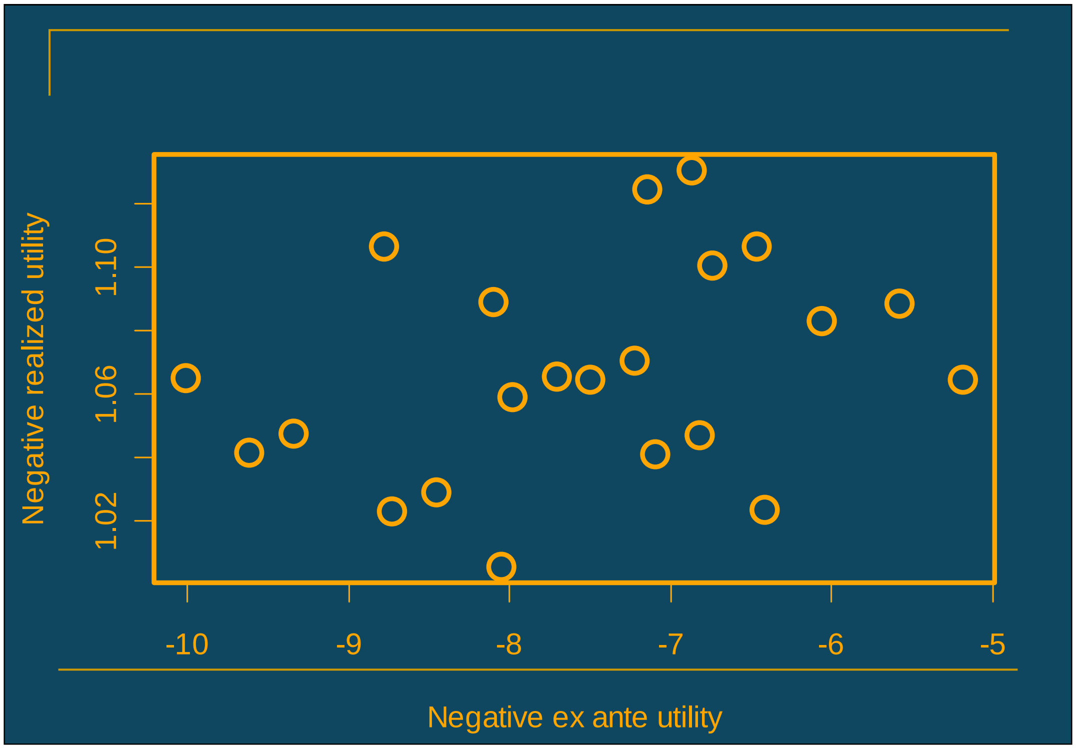

Only once we're home do we get to see how good of a shopper we were in the store.

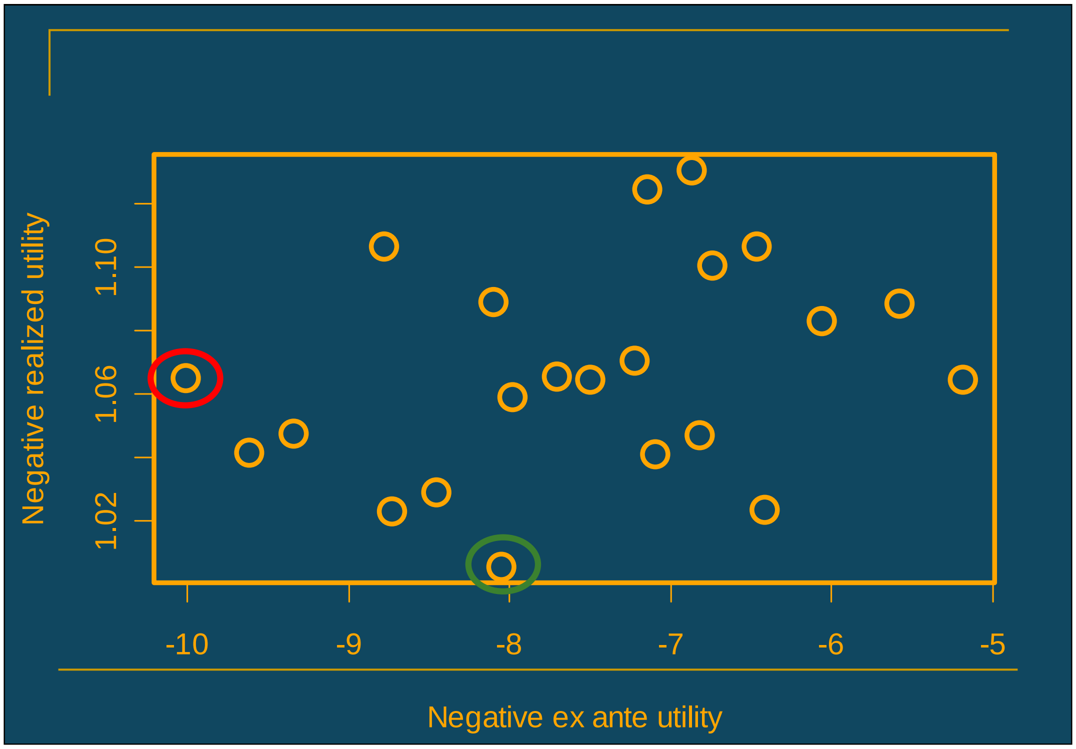

We picked the one circled in red. It turned out to be quite mediocre among our choices. We should have picked the one circled in green.

We see that indeed we were optimistic – we were looking for an information ratio of ten, we got one near negative one.

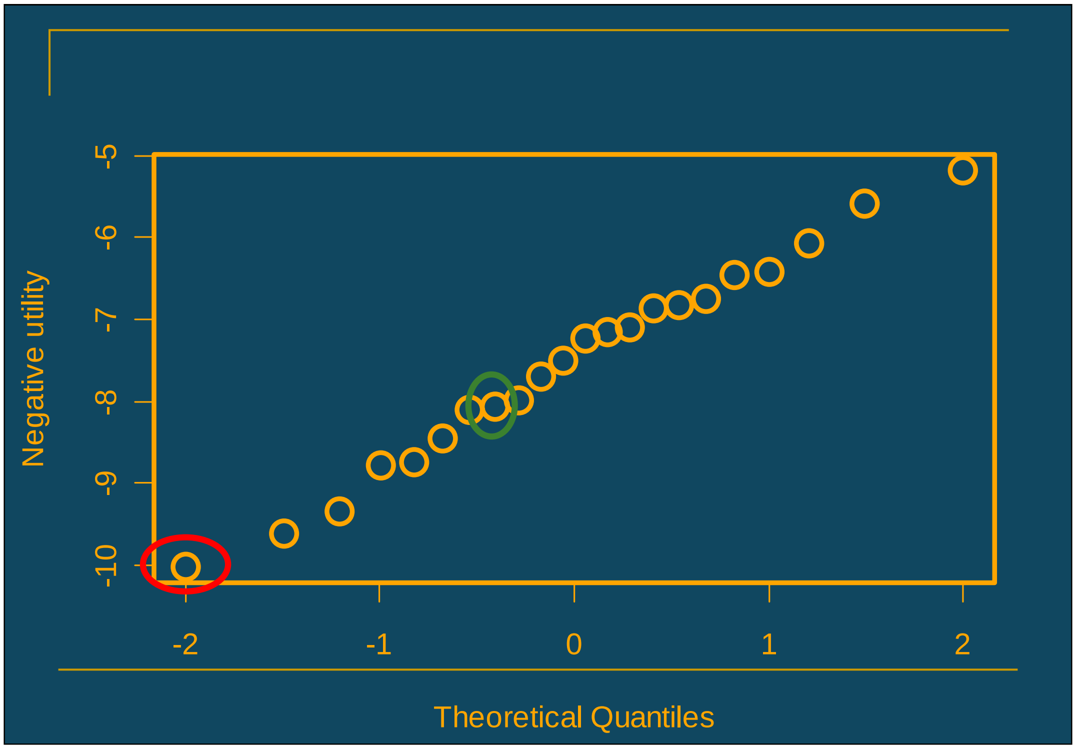

If when we were in the store, the salesperson had said “You should get that one circled in green.” We would have said “No thank you, we don’t need any help.”

The problem is that we have a low correlation between the ex ante estimate of utility and the realized utility.

Low correlation necessarily means difficult (that is, bad) optimization.

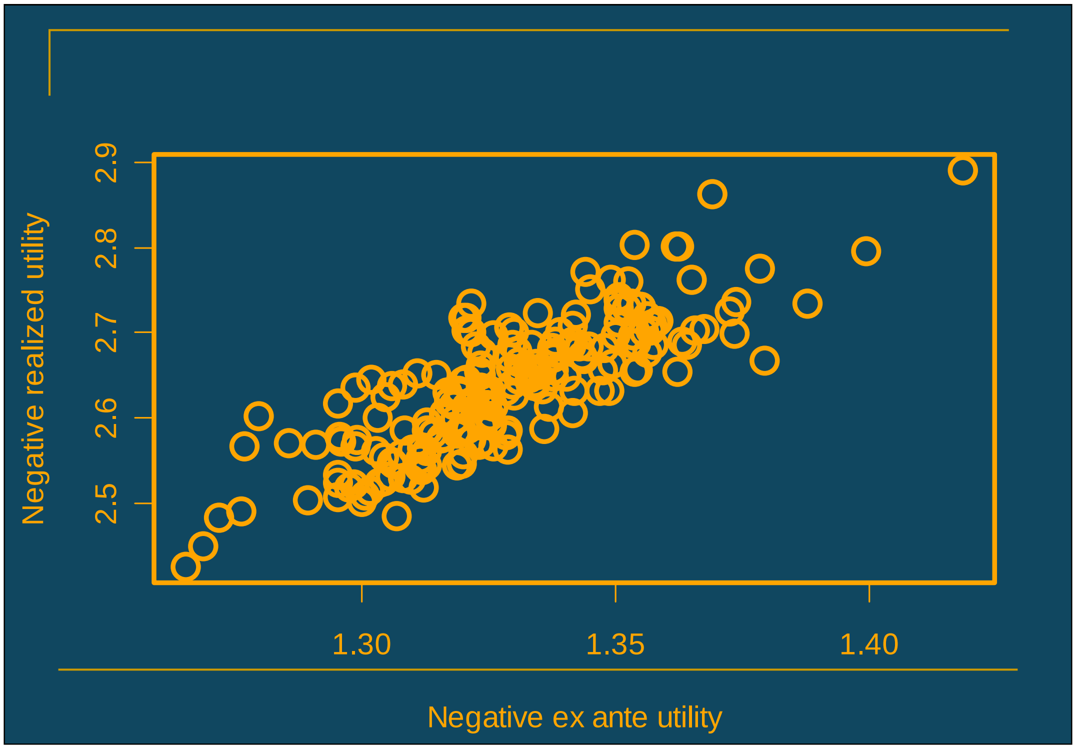

Here is an easier optimization.

It is easier because we have changed the utility. Instead of maximizing the information ratio, we are minimizing variance.

But still in this case we see we were optimistic. We thought we'd get a variance of something like 1.3 but we actually got about 2.5. This was in the first quarter of 2008, so things were beginning to hot up.

Even though we got the level wrong, we still have quite a nice correlation between ex ante and realized.

Note that the smallest ex ante value is also the smallest realized value. So we have perfect optimization in this case. Well, we would have if this were all of our choices, but actually we are just looking at a sample of our choices.

We used the general technique of random portfolios to get the data for the last plot.

The idea of random portfolios is that we have a set of portfolio constraints. We then sample from the population of portfolios that obey all of the constraints.

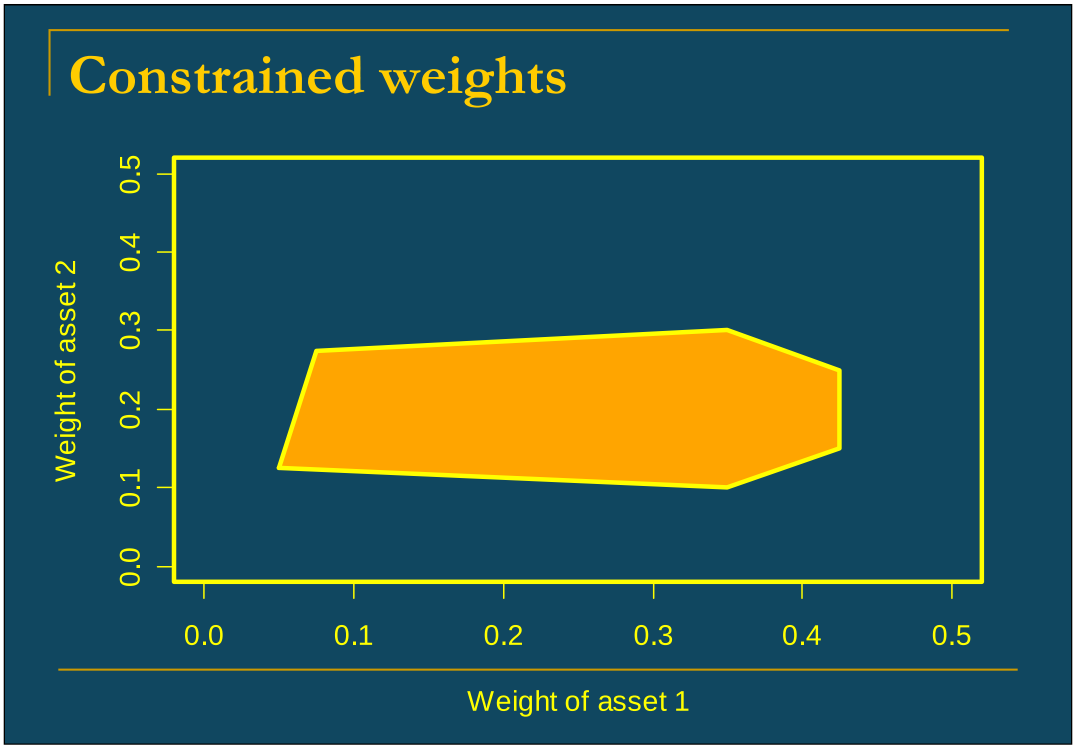

We are sampling from the orange area.

You can think of it as drawing portfolios from an urn. The portfolios that are in the urn are determined by the constraints.

Change the constraints and you change the urn.

TANGENT

We go off on a tangent here and explain how to actually generate random portfolios.

There are many ways of doing it, I'll tell you my strategy.

In order to do that, I need to tell you my approach to portfolio optimization.

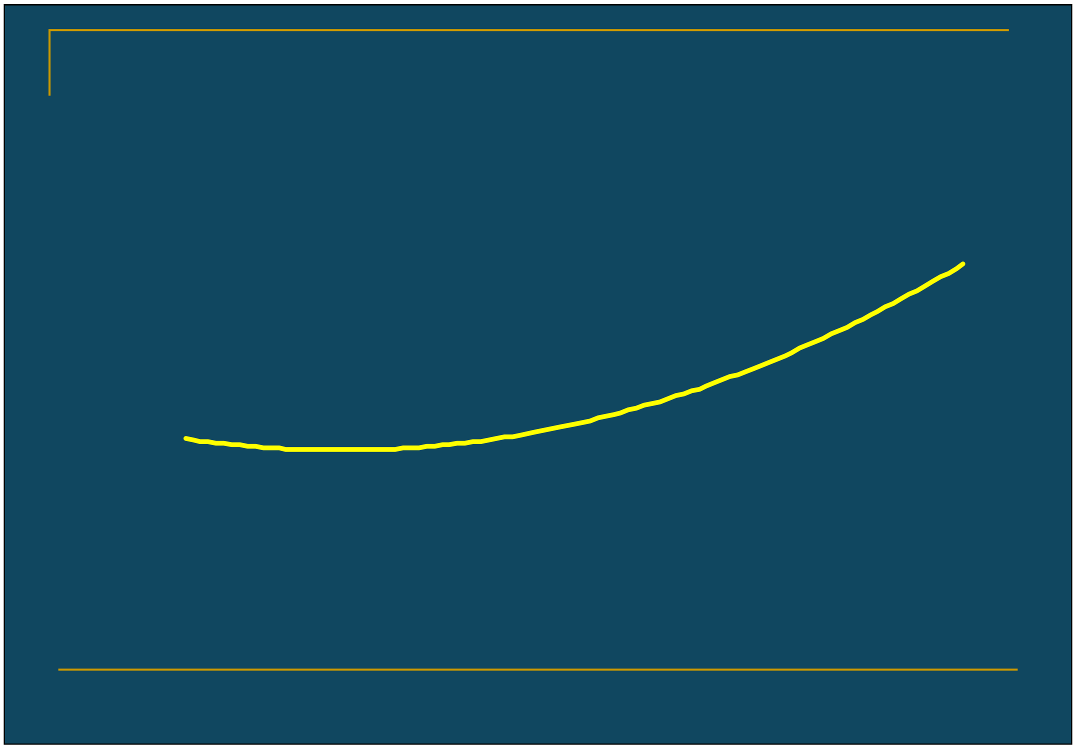

We are back to the global view.

Optimization has a negative utility that we want to minimize.

But there are also constraints to consider. How do we do that?

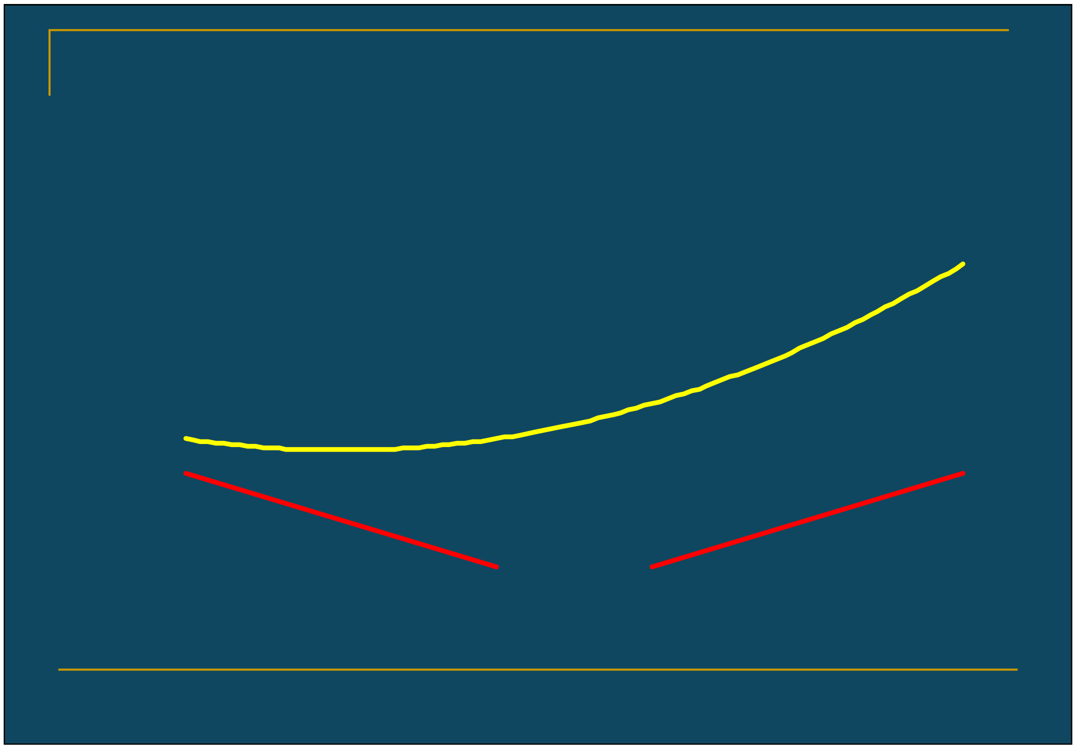

The way I handle most constraints is to penalize violations of them.

The bigger the violation, the bigger the penalty.

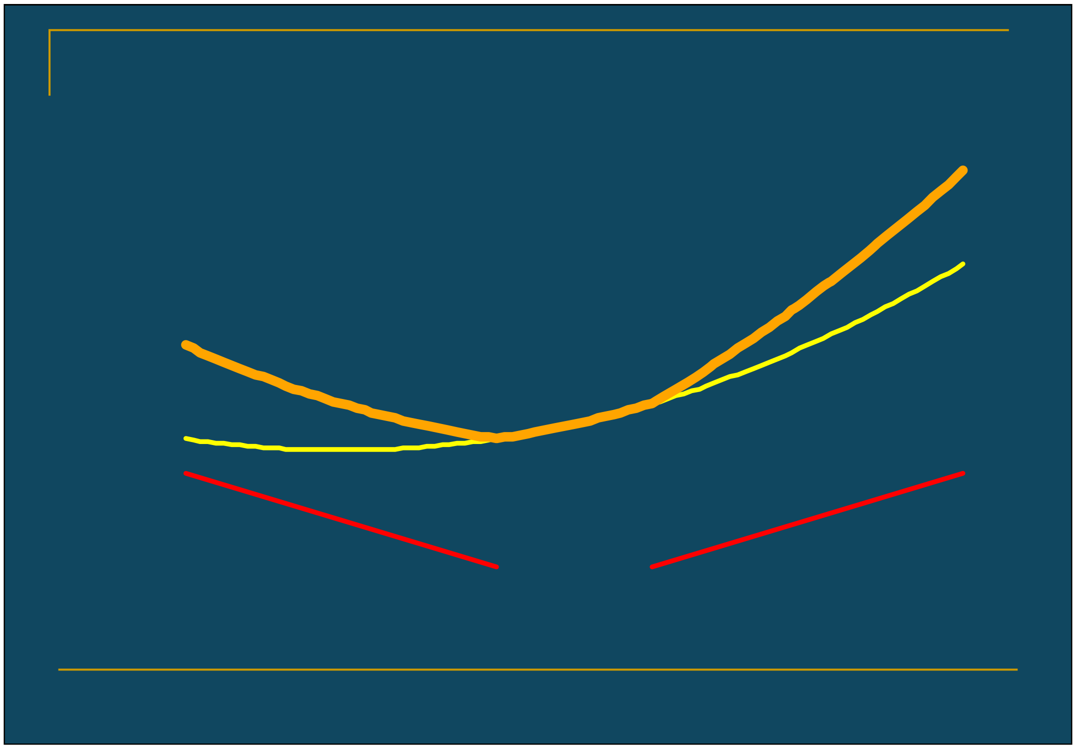

What the optimizer sees is negative utility plus constraint penalties, and that is what it minimizes.

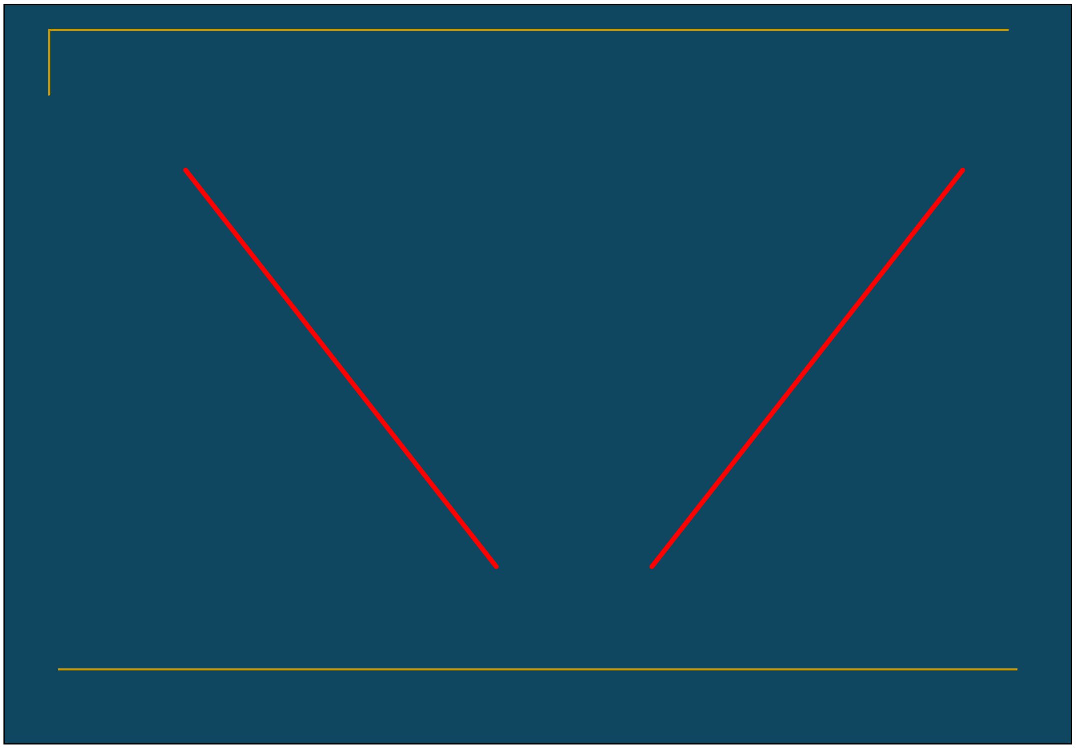

With random portfolios there is no utility. (Sometimes that is the whole point of using random portfolios – that we are ignoring utility.)

But there are still constraints that we can penalize.

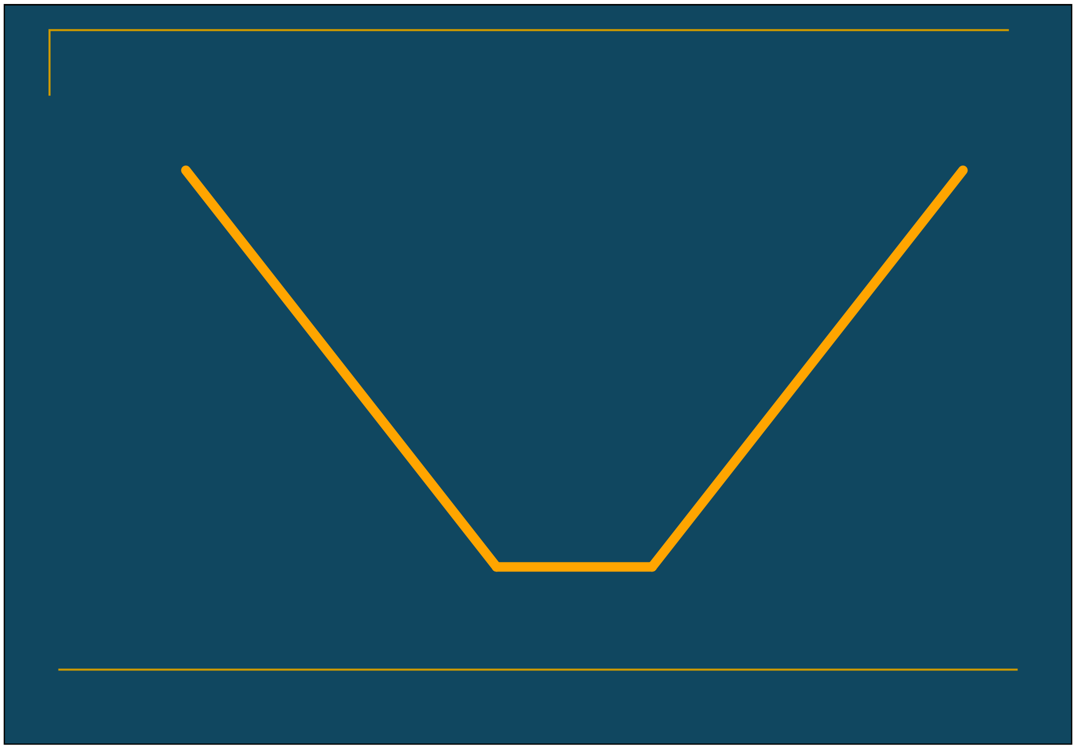

If we define the utility as identically zero, then we have something to minimize. When we get to zero, we know we are done. We can't go lower, and we know that all the constraints are satisfied.

We can increase the dimension of that image by 1 if we think of a volcanic crater with a lake in it.

The process is to parachute onto a random spot in the crater. Once we hit ground we start kicking a rock downhill. Where the rock splashes into the lake is our random portfolio.

We then typically do that hundreds or thousands more times.

Photo by USGS of Mount Pinatubo

END TANGENT

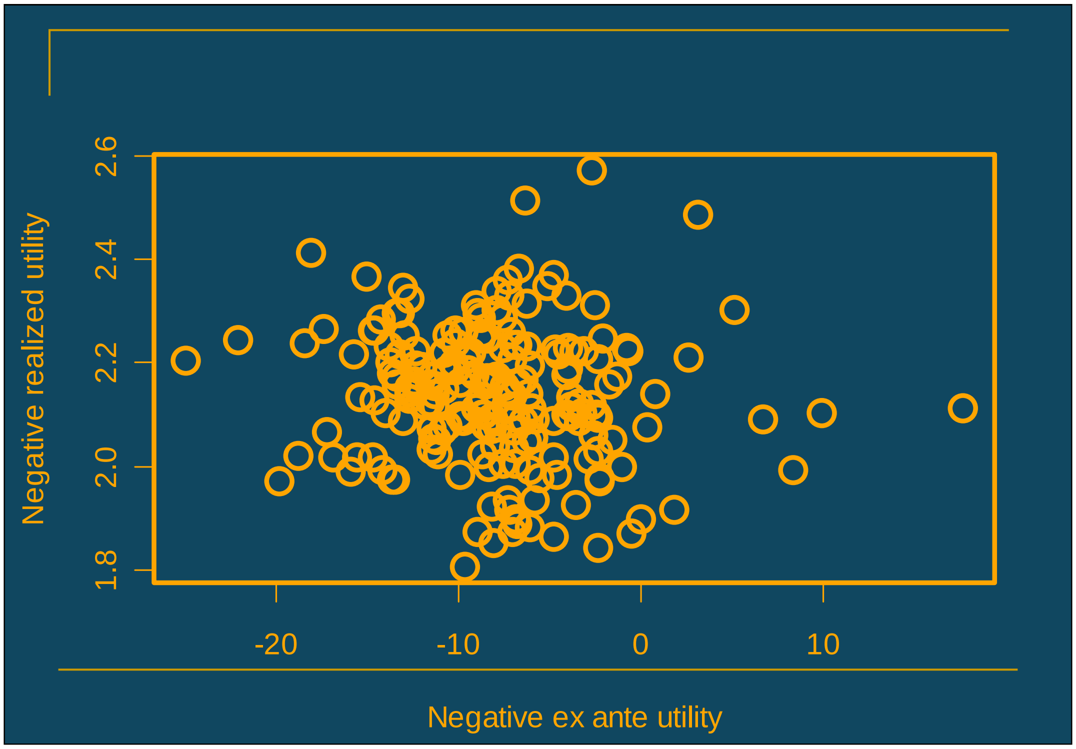

We're back in our local view.

And we're back doing a hard problem.

We're maximizing information ratio here, and the correlation between our ex ante estimate and the realized is negative. This does not lead to good optimization.

We want high correlation.

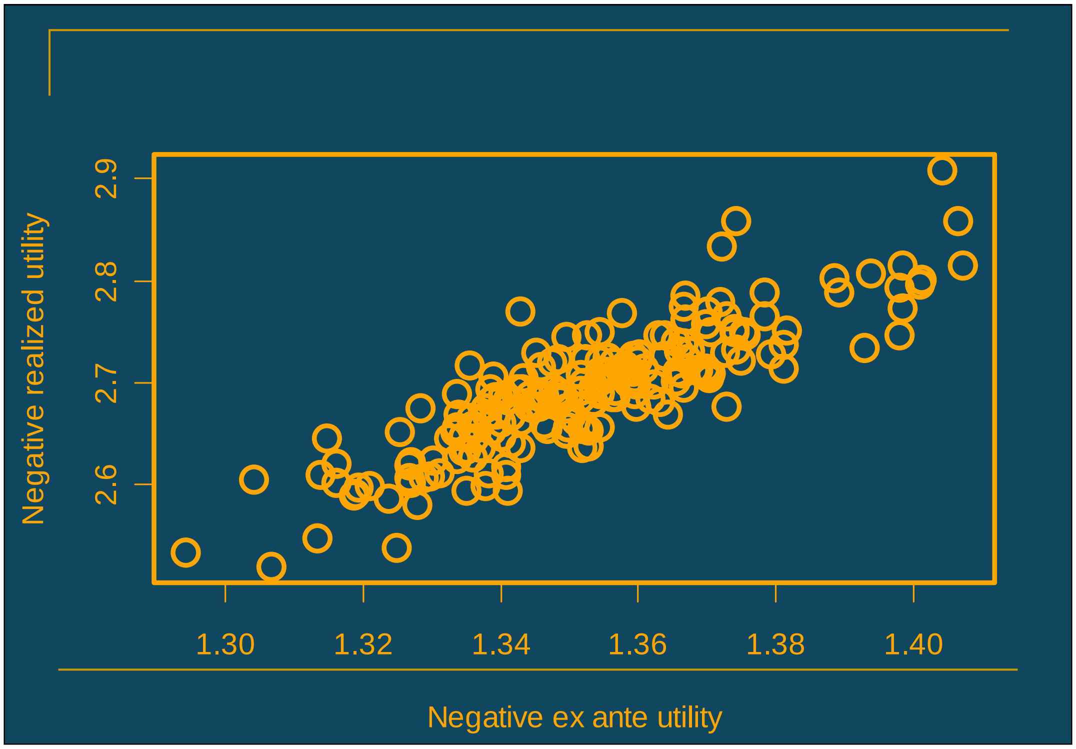

Here we have high correlation.

It is even higher correlation than what we saw before when minimizing variance.

By the numbers you might guess that this is a very similar problem to the one we saw before.

You would be correct. The difference is that a constraint has been swapped for a similar one.

Second problem: no asset contributes more than 5% weight

First problem: no asset contributes more than 5% to the variance

The more recent plot uses a weight constraint: the maximum weight of any asset is 5%.

The original problem said that the maximum fraction of the portfolio variance from any asset is 5%.

I claim that when people impose the weight constraint, they are subconsciously thinking they are doing the variance fraction constraint. We didn’t used to have the technology to do variance fraction constraints, so weight constraints were substituted.

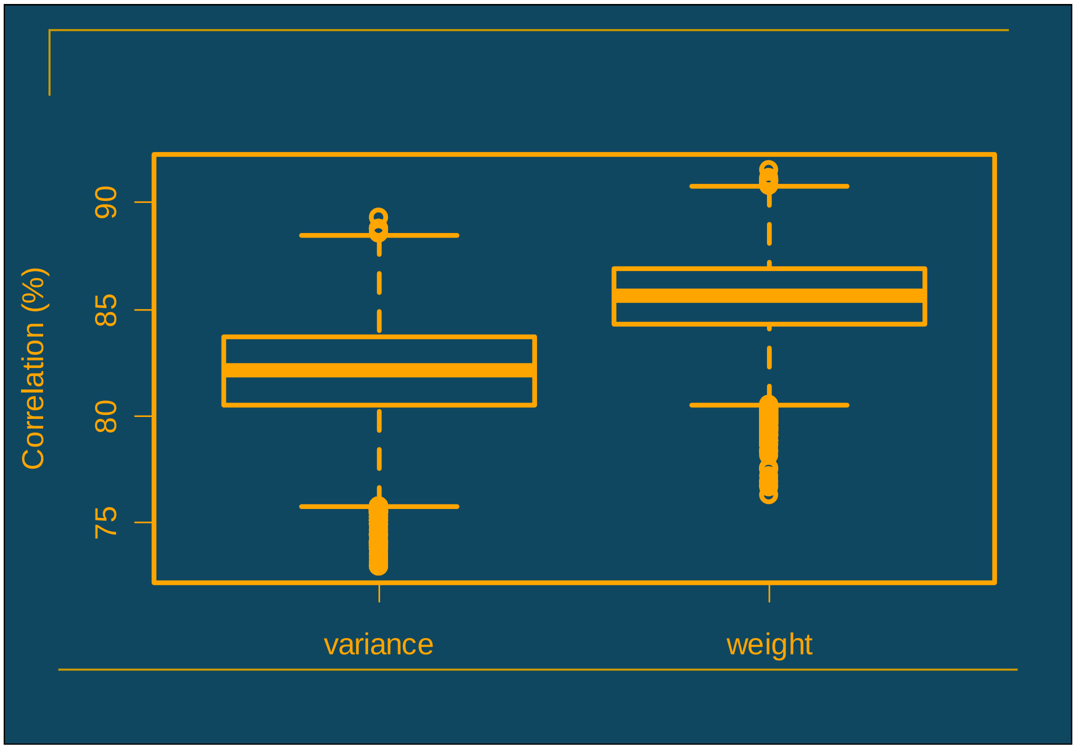

The plot shows boxplots of bootstrapped correlations in the two cases.

The weight constraint really does seem to have a higher correlation of ex ante to realized than the variance fraction constraint.

This makes the weight constraint appear better for optimization: high correlation means good optimization.

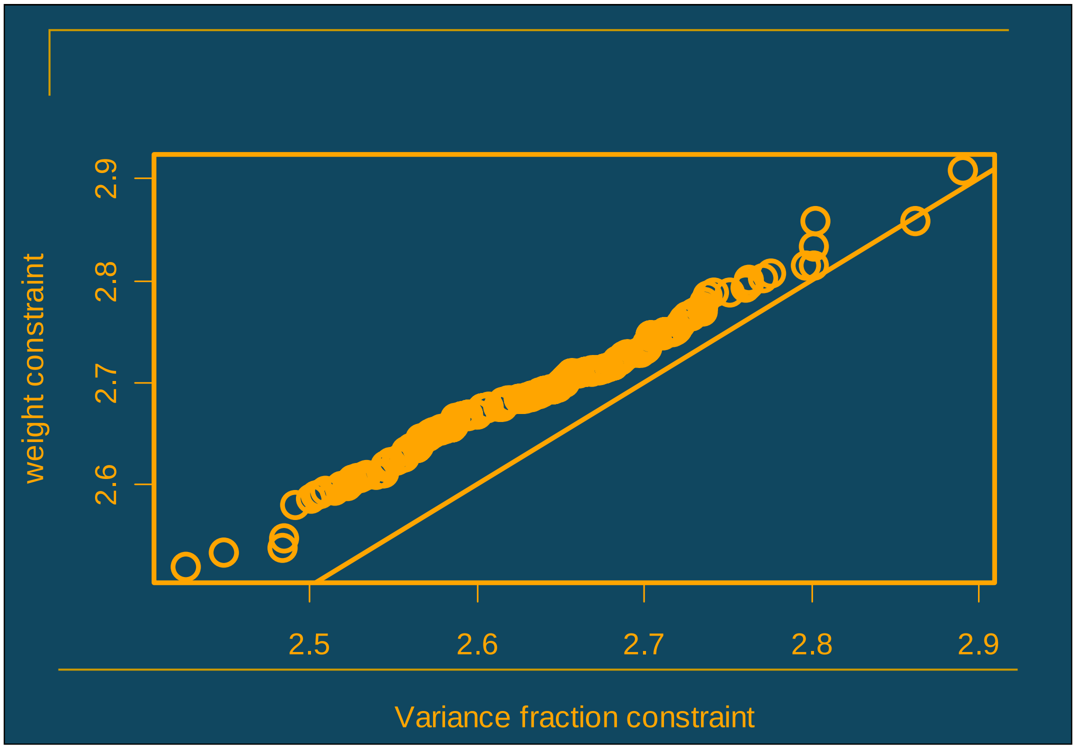

But really what we want in optimization is good utility.

We see that in this case the variance fraction constraint induces smaller variances than the weight constraint. This is especially true at the good end. The higher correlation with the weight constraints seems unlikely to overcome the difference in utility.

Of course to really know anything about the weight versus variance fraction constraints we would need to look at this over many time periods, and probably with several sets of constraints (other than weight versus variance fraction).

What have we seen?

A technique has been suggested to us of how to find constraints that will improve optimization. Of course it still needs to prove itself. If it works, it would be most useful for cases where we are using expected returns.

I hypothesize that the most helpful constraints (should there be any) would be contingent on how the expected returns are estimated.

There are brute force ways of using random portfolios to examine the value of constraints, but I quite like the subtlety of this idea.