

Random Portfolios: Practice and Theory



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

April 2009

Given at R/Finance 2009 in Chicago 2009 April 24-25.

You can find more information via the random portfolios page on the Burns Statistics website:

http://www.burns-stat.com/pages/Finance/random_portfolios.html

Outline

- **Basic Idea**
 - **Applications**
 - **Performance Measurement**
 - **Constraint Bounds**
 - **Theoretical Issues**
-


Basic Idea

CONSTRAI_{nt}

Random portfolios are intimately tied to constraints.

The idea of random portfolios is that we are sampling from the set of portfolios that satisfy all of the constraints.

It may be the case that whenever there are portfolio constraints that random portfolios can be of use somehow.




This is a sketch of the situation in a three asset case. We want to sample from the orange area.

There are two things wrong with the picture:

- 1) All the lines are straight, meaning there are only linear constraints. Constraints are not always linear.
- 2) The picture implies a continuous space. Real portfolios are discrete. You can't buy π shares of IBM.

Constrained weights



Because of the discreteness, we are really on a lattice like the red dots.

I learned at (the infamous) Jak's that this picture is a Rorschach test. There are those who see a coffin and those who don't.

Statistical Connections

- **Sometimes similar to statistical bootstrap**
- **Sometimes equivalent to random permutation test**

We'll see random portfolios used both in the spirit of the bootstrap and in the spirit of random permutation tests.

Random Portfolios: History

- **Goes back at least to Chicago 1965**

Applications

Applications

- Performance Measurement
- Testing Trading Strategies
- Evaluating Constraints
- Validating Risk Models
- ...

Performance measurement and testing trading strategies are conceptually similar except one is *ex post* and the other is *ex ante*. The use of random portfolios differs markedly in the two tasks, however.

Random portfolios can be used to validate risk models. The receiver of a model can send portfolios through the model that have characteristics similar to portfolios the receiver holds. The creator of the model can look for weak spots in the model.

There is certain to be a lot more applications that we haven't yet thought of.

Performance Measurement

- **Benchmarks**
 - **Peer Groups**
 - **Random Portfolios**
-

Benchmark Measurement

- **One-sample t-test on time series of difference of returns**

We have a time series of differences, but we ignore the time element of the data.

The key problem with this technique is that it takes a very long time to know if the fund is good or not.

Power with 3 Years Quarterly

Info ratio	5% sig.	1% sig.
0.5	20.3%	6.0%
1.0	49.0%	21.0%

**If population is 90% IR=0, 10% IR=1,
almost half of selected will be zero skill**

Focus on the 5% significance test and an information ratio of 1. We declare less than half of these very good funds to have skill after three years. Meanwhile we are letting through 5% of the funds that have zero skill.

Consider the case where we have 1000 funds: 900 have zero skill and the remaining 100 have information ratio 1. With our test we expect to declare 49 of the skilled funds to have skill. We also expect to declare 45 of the zero skill funds to have skill.

Power with 10 Years Quarterly

Info ratio	5% sig.	1% sig.
0.5	46.5%	21.2%
1.0	92.9%	76.6%

With 10 years of data the results for information ratio 1 have become reasonable. But how many funds are the same fund a decade later?

Results for the more likely information ratio one-half are still not good even after 10 years.

And these numbers are highly optimistic.

Model for Power Statistics

- **Normal Distribution**
- **Constant Outperformance**


The power statistics were created with a model. Models are wrong.

I assumed a normal distribution. Returns are not normal.

I redid the simulations using a t-distribution. The power was the same. That was a surprise to me.


The other key assumption was that outperformance is constant through time. That is not true, and not true in a very significant fashion.

I considered simulating the power statistics with varying outperformance, but I don't know how to parameterize that in a realistic manner.



This plot is of a set of 100 funds that track the S&P 500. Something is not constant through time in this plot – either the IQ of the fund managers or the ease of outperforming the benchmark.

I maintain that assuming non-constant outperformance is the more realistic possibility. Especially since we know a mechanism to explain it: if the assets with the biggest weights in the benchmark happen to do relatively well, then the benchmark will be hard to beat.



In case you doubt me, here is the same plot with 95% confidence intervals that indicate there really are differences.


Peer Groups

- **Assemble a cohort of “similar” funds**
- **Report the percentile of our fund within the cohort for one time period**
- **What does that mean?**

Note that “similar” is in quotation marks. “similar” could be a conference all on its own. For a lot of hedge funds there are no reasonable peers at all.

Also note that we are using only one time period now – that’s progress.


If we don’t think too hard, then we think we know what the percentile means. We don’t.



The vertical lines represent different funds. The x-axis is returns.

We only ever see skill plus noise (the top of the plot). The peer group technology wants the ordering of skill plus noise to be the same as the ordering of skill. Like the plot here.

Pure fantasy.



We're going to get something more like this. Actually I suspect that this is probably about as good as it gets. My guess is that usually it is much worse than this.

Consider the possibility of all the funds having the same skill – maybe it is zero skill, maybe it is a lot of skill. If they all have the same skill, then the percentile is the percentile of luck.

There is no way for us to tell from the percentiles whether they all have the same skill or not.

Perfect Performance Measurement

- **Look at all possible portfolios that the manager might have held**
- **Take the return of each of these portfolios over the time period**
- **Compare actual return to the distribution from all possibilities**

Let's take a detour and think about perfect performance measurement. What would this look like?

Fund managers have a cloud of portfolios that they MIGHT hold, and from that cloud they pick one. So we want to look at the distribution of returns from the cloud and see where the actual portfolio lies in that distribution.

Some Caveats

- **Should account for implicit as well as explicit constraints**
- **Return need not be the measure**
- **Trading is allowed**

Fund managers may have implicit constraints (such as growth-oriented) as well as explicit constraints. We should take the implicit constraints into account in addition to the explicit constraints.

I have talked of returns as the measure of performance. You can use whatever measure you like: risk-adjusted returns or whatever. For simplicity I will continue to just say “returns”.

When describing what fund managers do, I said they pick one portfolio out of the cloud. That implies that they do no trading throughout the period of interest. Fund managers do trade, so really we have a path through the cloud, not a single point. The point is easier to visualize, but the path doesn't complicate the analysis any.

Perfect Performance Measurement

- **Number of possible portfolios is finite but astronomical**

Perfection is hardly ever achievable.


We are not going to be able to deal with the whole cloud of portfolios that the fund manager might hold. The field of statistics says that taking a random sample is going to be practical and adequate in such a case. That is, we want to generate random portfolios with the constraints of the fund.

Random Portfolio Measurement

- **Take a random sample from the set of all portfolios that might have been**
- **Fraction of random portfolios with a larger return than the fund is a p-value**
- **Null Hypothesis is zero skill**

We can do a statistical hypothesis test (a la random permutation test) where the null hypothesis is that the fund exhibits zero skill. If the fund performs well enough relative to the random portfolios, then we may be in a position to declare that it exhibited skill during the period.

Alternatively if the fund performed quite poorly, we may declare that it exhibited negative skill for the period.



Now that we know how to really do performance measurement we can look at the other two methods in the new light.

But first, let's compare random portfolios to the matched portfolios that David Kane talked about in his presentation.

The idea of matched portfolios is that you make alternative portfolios look like the portfolio under question in as many respects as possible except the aspect that is supposedly adding skill.

David's example was of a system that picks stocks. The matching portfolios pick stocks that the system doesn't think are good, but otherwise tries to make the portfolios similar.

Use random portfolios when you know or can guess constraints. Use matching portfolios when there are no (or vague) constraints.

Peer Groups

- **Superficially like random portfolios**
- **Comparing against “random” portfolios of unknown skill**
- **Tension between lots of peers for power, few peers for similarity**

Peer groups look a lot like random portfolios: both use a single time period, both compare the fund of interest to a number of alternative portfolios.

The key difference is that we don't know the skill level of the peers. We do know the skill level of the random portfolios – it is zero skill.

With peer groups we want a lot of peers in order to get more precision with the percentile. On the other hand we want only a few peers because we want only the most similar funds.

There is no such problem with random portfolios. We can generate as many as we like.


Benchmark Measurement

- Testing against 1 “random” portfolio

Saying that we are comparing with one portfolio is actually not quite right. We are comparing at different time points, so the more accurate analogy is that the number of alternative portfolios is the number of time periods used. The composition of the alternative portfolios is highly correlated across time.

Some people might not be so pleased with the idea that the benchmark is random. I'm thinking of those who are keen on the idea of the market portfolio.

The next slide should be enough to convince you that the benchmark is random.




Benchmark Measurement

- Testing against 1 “random” portfolio
- Benchmark may not satisfy constraints


The benchmark is usually outside the explicit constraints of the fund.

Even if it obeys the explicit constraints, it will be outside the political constraints for the fund. A fund manager is not going to please clients by charging them a fee to hold a portfolio that the client could buy essentially for free.

Since the benchmark violates constraints, the comparison is unfair. We're not sure to whom it is unfair in any one period *ex ante*.



But wait there's more ...



Sometimes we know all or some of the positions of the portfolio at the start of the time period. This is information that neither benchmarks nor peer groups can take advantage of. But random portfolios can use that information.

The yellow line is the distribution of returns over a calendar year under a certain set of constraints. The red line is the distribution under the same set of constraints, plus it is starting at a certain portfolio and has slightly over 200% turnover (buys plus sells) throughout the year.

Upper tail probabilities

Return	uncond	cond
30%	.52	.19
31%	.43	.10
32%	.35	.06
33%	.28	.03
35%	.16	.004
40%	.03	0

We can see that there is a big advantage in this case for the inference of skill. Unconditionally we need to see a 40% return over the year in order to get a 3% p-value. Conditional on the starting portfolio, we only need to see a 33% return to get the 3% p-value.

Constraints

- **Why do we impose constraints?**
- **Insurance**
- **What protection are we buying?**
- **What price is the premium?**

We turn now to evaluating constraints.

The reason we impose constraints is so that the portfolio won't do anything too stupid. That is, we are buying insurance.

FTSE Example

- **FTSE 350 Data**
- **10,000 portfolios generated for each set of constraints**
- **Returns: 2006 Jan 01 – 2006 June 01**
- **Long-only**
- **90 – 100 assets in portfolio**
- **Nested set of linear constraints**

The key bit of this slide is that we have a nested set of constraints.

FTSE Linear Constraints

- **Large cap versus Mid cap**

- **10% - 30%** **70% - 90%**
- **13% - 27%** **73% - 87%**
- **17% - 23%** **77% - 83%**

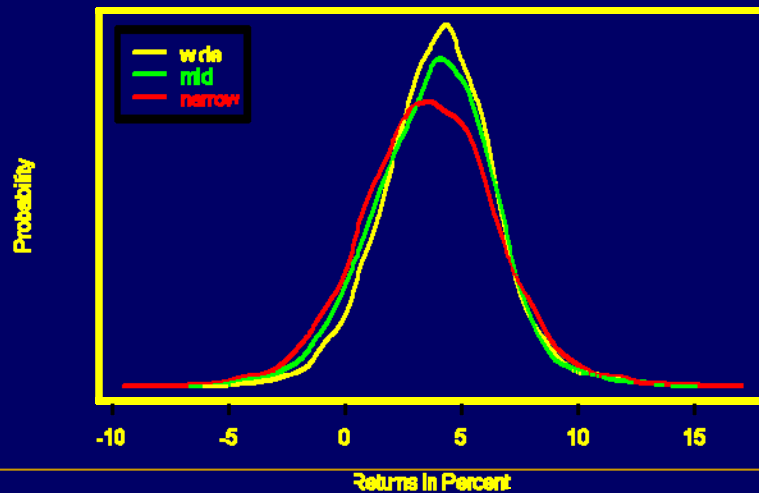
- **High yield versus Low yield**

- **50% - 70%** **30% - 50%**
- **53% - 67%** **33% - 47%**
- **57% - 63%** **37% - 43%**

FTSE Linear Constraints

- **5 Sectors**
 - **10% - 30%**
 - **13% - 27%**
 - **17% - 23%**

FTSE Return Distributions

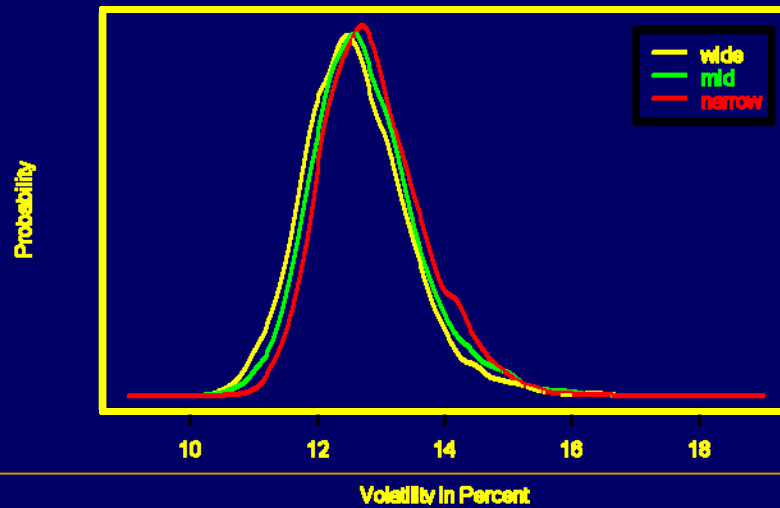


When we constrain more, we expect to be giving up some of the upside to protect against the down-side.

In this case we get just the opposite.

Quite a puzzling result. One idea is that we are constraining into a high volatility region.

FTSE Volatility Distributions




It does seem to be that we are constraining into a high volatility region.

The distributions here are based on the exact same random portfolios as last slide. The only difference is that we are looking at volatility instead of returns.

An idea for something to do is to constrain volatility to be no more than 12% -- a significant constraint but not egregious.

FTSE Return Distributions: Constrained Volatility (at most 12%)



Now we constrain away some of those nasty positive returns. And oh, by the way, we still have a wider distribution when we constrain more.

I'm not sure how common perverse examples like this are. But I'll be surprised if we aren't often surprised by what constraints are doing once we start looking.

Theory

Pat's Conjecture

Random portfolios are the most powerful (practical) method of performance measurement

If you find the conjecture not to be true, I certainly want to hear about that.

A Theoretical Project

- **Formal setting that justifies “perfect performance measurement” claim**

The theoretical task is to justify the “perfect performance measurement” claim that I made earlier (without any justification).

Random portfolios should be thrown in as the practical way of actually getting the measurement.

Interlude: Generating Random Portfolios

- **Rejection method**
 - **How many billion years do you have?**
- **Linear polytope methods**
 - **Often do not apply**
- **Random search technique**
- **New methods are possible**

The rejection method is often a handy way of generating random variables. It is going to be unworkable for random portfolios in all but the most trivial of cases.

(A linear polytope is the shape you get when you only have linear constraints – tetrahedra, cubes, dodecahedra are examples in 3-space.) People are working on random generation within linear polytopes, but this is not generally adequate in practice. We just saw an example of a volatility (quadratic) constraint, and we have used integer constraints on the number of names in portfolios.

Another way of generating random portfolios in a case where an optimization is being done is to permute the expected returns randomly and do the optimization. (You will want to pay attention to the standard errors of the expected returns.) This method is cheap to implement but expensive in execution time.

It makes sense to me to look for new methods of generation.

Random Search Technique

- **Create an objective function**
 - Penalize for broken constraints
 - Zero when all constraints satisfied
- **Minimize the objective**
 - Start at a totally random portfolio
 - Make random moves towards the goal

The method I use to generate random portfolios is a random search technique. We create a function that is zero when all the constraints are satisfied and positive elsewhere.

We then minimize that to get one random portfolio.



My image of this is of a volcanic crater with a lake in it.

We start at a random spot inside the crater. We then start kicking a rock downhill. Where it splashes into the lake is our random portfolio.

USGS image.


What Could Be Wrong?

- **We stop kicking when we get to the lake**
- **High probability of being close to a binding constraint**

For practitioners, the answer is that nothing at all can go wrong – just use random portfolios without question.

For the theoreticians the answer is that we stop kicking the rock once we get to the lake, so we are likely to be close to shore. We are likely to be close to at least one binding constraint.

Typical Non-uniformity



Here is a sketch of that situation. The random portfolios are likely to be clustered near the surface of the feasible* region and sparse in the interior.

* “feasible” means feasible to be one of the selected portfolios -- that is, satisfying all the constraints.

A Fix (Sort of)

- **Continue kicking once in the lake**

A way to get a more uniform distribution is to skip the rock some number of times across the lake. The 100th splash is likely to be farther from shore than the original splash.

I haven't yet implemented this, but it is on my to-do list.

Measuring Non-Uniformity

- **Easy if there are only a few feasible portfolios**
- **Hard if there are more feasible portfolios than protons in the universe**
- **A great research opportunity**

This brings up the question of how do we know if we are close to uniform or not.

If there are 10 feasible portfolios, then generate 1000 random portfolios and count to see that each one gets hit just about 100 times.

It gets harder from there. But some bright young thing can probably make progress on the problem.

Is Uniform Right?

- **Probably not**
- **Fund managers tend to be tight on some constraints**
- **Optimal distribution will depend on the application**

The uniform distribution is no doubt the “right” distribution for theoreticians. In practice uniform is probably not the right distribution.

Fund managers think they have constraints. They are often tight on some of the constraints. If they weren't, then they wouldn't think that they had constraints. Hence the fund manager distribution is likely to look a lot like the naïve random search distribution.

Also consider that in any time period the best portfolio is going to be outside the constraints of a fund with (basically) probability 1. In this light the job of the fund manager is to select which constraints to be tight on.



In summary random portfolios are a wide open field both for theory and for practice. We've seen a few ideas for theoretical work. Random portfolios can be thought of as a sort of statistical bootstrap which has generated hundreds of statistical theory papers.

In practice we know of several quite important applications of random portfolios, and there are roughly half a zillion still to be found.

There was a question from the floor about the connection between random portfolios and the resampled efficient frontier. Assuming there are no expected return or variance constraints, then all of the resampled portfolios will satisfy the constraints of the problem. That is, they satisfy the constraints just like random portfolios. But the resampled portfolios are in a corner of the feasible region because of the optimization they go through.