


 A full-page background image showing a hiker in mid-air, jumping over a large, natural rock arch. The hiker is wearing a backpack and using trekking poles. The scene is set in a desert-like environment with orange-brown rock formations under a clear blue sky.

Taking the random route

Measuring performance against a benchmark has been the standard methodology across the buy side for a number of years, but as **Andrew Colin** explains, not all benchmarks are equally challenging to outperform

Joe Micawber was uneasy. The back office had just sent through the monthly figures for his firm's flagship fund. With a billion dollars under management in treasuries and gilts, the year-to-date numbers looked healthy, if not spectacular.

His hand-picked team, each with an average of 20 years' investment experience, had again outperformed the benchmark for the quarter, and the institutional investors were happy.

But the boutique investment fund down the hall, recently started up by two aggressive MBAs, was showing spectacular returns – three times higher than Joe's flagship fund. How could he prove that his bond team had better investment skills?

Measuring investment skill is a hot topic these days. Rather than just asking how much return their funds made, astute investors are looking a level deeper and trying to measure the skills of their managers.

Traditional ways of measuring investment skill range from the straightforward (looking at the portfolio's performance relative to the rest of the market over the last one, five and 10 years)

to the quant-driven (tracking error, information ratios and other statistical measures), through to the full machinery of attribution analysis, which measures returns made on separate sources of risk in the portfolio.

These all have strengths and weaknesses, but their common factor is that they measure return and risk relative to a benchmark. Performance above a benchmark is *prima facie* evidence of true investment skill.

Or is it?

This has been a widely held view for many years, but it's starting to show cracks. Not all benchmarks are equally hard to outperform. When large cap stocks are rising, it will be hard to beat the benchmark; when they are falling, it will be relatively easy. Yet few skill assessment schemes take account of this.

There is, however, a powerful but relatively little-known technique for measuring the skill of an investment manager without going anywhere near a benchmark. Counter-intuitively, it is based on the idea of building a random portfolio.

Suppose we could form the set of all the possible investment portfolios, subject to the same constraints as the fund. Out of this huge number, your manager has used his skill and judgment to select just one of them. Was that a good choice?

For real markets, it would be impossibly time-consuming to work out the returns of every possible portfolio. Fortunately, we don't have to – statistical sampling theory means that we can get a good idea of how these returns are distributed by building a much smaller number of randomly selected portfolios, either using a dartboard or (more realistically) a random-number generator in software.

Clearly, the returns of these portfolios are not all going to be the same. Most will be bunched up around the average, but a few will be a long way above or below the norm, and the overall returns will follow a bell-shaped curve – a normal distribution.

Once you have this curve, it becomes very easy to determine whether your manager showed real skill. Are the returns of your

manager's portfolio clustered in with the bulk of the returns, or do they lie out on the right, beyond the pack?

Grouping

What's important here is not just the average level of returns, but how tightly they are grouped. If 99% of all possible portfolios had returns between 0% and 1%, then a return of 2% is clearly exceptional, given the constraints under which the manager worked.

On the other hand, if 99% of all returns lie between 0% and 4%, then that same 2% return looks much less impressive. The first situation is most unlikely to have come about due to chance, and is strongly suggestive of real investment skill; the second is not. And this is what we're trying to measure.

This type of skill assessment is of course independent of any benchmark. But suppose we repeat the same analysis with the benchmark stocks. If the benchmark's return is below the average performance figure, then any fund measured against that benchmark is

internally, by the compliance manager or customer mandates.

Others will be due to external factors, such as cases in which diverging substantially from the benchmark may just not be physically possible, due to market liquidity. It is much harder to buy \$1million of stock in a company with a \$10million market capitalisation than in a company with a \$10billion market cap – and even if you could, taking such a large stake in a company brings a host of time-consuming legal issues in its wake.

Narrowing options

For the larger manager, this means that the range of possible investment portfolios will narrow substantially, and it will be much harder to take the extreme positions that the smaller manager can hold.

Realistically, the weights of each stock in the portfolio for this manager will have to be fairly close to those in the benchmark. With this smaller universe of possible positions, the spread of returns will contract, and it will be much harder to make very high returns.

and transparent performance analytics.

This is one reason why making a high return on a large fund can be much harder than with a small fund, as many managers of funds of funds discover; the boutique approach based on a single great idea just doesn't scale up.

In fact, an investor in a fund working with fewer operational constraints should demand that the extra range of opportunities be balanced by extra performance.

Random portfolios can give you an idea of just how much more return you should be expecting for the risks the manager is taking.

So next time you are comparing returns against that hedge fund down the hall, consider the total amounts under management, and the constraints under which you are both working.

It may well be that although you did not make as much – in relative terms – as your neighbours, your skill is demonstrably greater. And that is worth knowing. >

Andrew Colin is director of fixed income research for the StatPro Group and author of 'Fixed Income Attribution' (Wileys, 2005). He is also adjunct professor in the Faculty of Business, Queensland University of Technology. He can be contacted on andrew.colin@statpro.com.

Measuring investment skill is a hot topic these days. Rather than just asking how much return their funds made, astute investors are looking a level deeper and trying to measure the skills of their managers

getting a free ride. For instance, if a portfolio made of random stocks from the benchmark shows a 10% return, but the benchmark returns 8%, then the free ride is 2%.

Arguably, this amount should be deducted from all published active returns before claiming superior investment skills. Nobody (to my knowledge) currently does so.

Calculating the spread of benchmark returns is such a simple idea that one wonders why benchmark vendors do not supply this data, along with daily returns. In fact, there is a fairly good reason, and it involves looking at what investments the manager can actually make.

The large fund manager will typically be subject to numerous constraints in the positions he takes. Some will be imposed

The small fund manager, on the other hand, typically has far fewer constraints on what positions can be held, so he will be able to diverge much more widely from the benchmark, using a mixture of physical holdings and derivatives. In fact, this manager may find they can generate a year's outperformance with a few well-placed, lucky bets.

But the educated investor should treat these sorts of results with extreme caution. One or two winning streaks are not necessarily evidence of real investment skill – as every casino manager knows. Impressive results may be genuine, but should always be placed in context with a clear, articulated investment process, risk control, realistic back-testing,

