



No need to service a smooth running engine



While accuracy of pricing and calculation is absolutely vital for calculating a portfolio's performance, it's much less important for calculating attribution. Many investment institutions are spending a disproportionate amount of time and effort in putting more accuracy into attribution analysis than it really merits. Is yours one of them? **ANDREW COLIN** of Flametree Research explains

I was recently talking to the CEO of an engineering company that supplies wing parts for a certain aircraft manufacturer. Much to my surprise, he told me that many of the components made in his factory were actually machined to a relatively low level of accuracy, despite the extremely close tolerances required in many areas of aeronautical engineering. He was not cutting corners, but simply providing precisely what the customer specified. (This was reassuring, as I'm writing this article on one of these very aircraft).

It turns out there are sound economic reasons for what appears, on the surface, to be sloppy engineering. In the 1940s, when the mean lifetime of a Lancaster bomber on combat operations was three weeks, it made sense for an aircraft's components to be just good enough — and no more. By cutting out unnecessary work, the factory could build more components, put more aircraft in the air, and so keep up with Bomber Command's enormous losses.

In fact, it's likely the Axis forces were placed at a disadvantage because their engineering was too good. They concentrated on quality rather than quantity, and it cost them dearly.

The same cost-cutting continues in industry today, because it makes sound economic sense and makes no appreciable difference to the eventual product. For instance, you bore an engine cylinder to very fine tolerances, because an accurate fit between the cylinder and the piston is critical to the efficiency of the engine. You don't worry so much about the finish on the engine block, because there's no point; the engine will work just as well whether it's polished and chromed, or left as it came out of the casting.

What's this got to do with performance measurement? The answer is: a great deal.

Is your benchmark too accurate?

On the surface, it is a silly question. Of course you want the performance contributions from benchmark sectors, and indeed individual securities, to add up to the published benchmark

return. If they don't, any analysis you produce using these figures is useless.

But as long as this condition is met, does a variation of a fraction of a percent error across sector boundaries actually matter?

Even in the industry-standard Brinson equity attribution model, which splits performance between stock selection and asset allocation, an interaction return term appears that doesn't fall into either category.

If the size of this interaction term is larger than the inaccuracies arising from uncertain data, it's unlikely the data issues will materially affect the results of the attribution analysis.

My experience with fixed income benchmarks is that you can usually get to within less than one per cent of the published returns over several months by using published weights, and making simple assumptions about cash flow sizes and timing. Rescaling the individual returns of the individual securities by a very small amount will then make the overall return equal the published returns.

When it comes to turning this data into attribution statistics, the uncertainties arising from this tiny modification will be no greater than those from your attribution calculations, and usually much less. By being willing to adjust security-level numbers very slightly, you can save a huge amount of effort, while making no appreciable difference to your final results.

The new security problem

Let me show you another example. At one of my client sites we have been looking at how to run attribution on a complex CDO (collateralised debt obligation), as part of a much larger portfolio of vanilla and inflation-linked bonds. We can source daily prices for the CDO from the market, but we don't yet have the complex mathematical machinery in place to run a full attribution analysis. How do we include this security in the daily attribution report?

The unhelpful answer is to tell the manager to hold off buying any

Over-engineering can be a big distraction from the rest of the investment management process. As long as you have a handle on the size of your data tolerance issues, you'll save time and effort

of the new instrument type until the technology is in place to fit it into an attribution system. Naturally, this won't be popular with the dealers. The performance team is there to support the front office, not the other way around.

A better, more pragmatic answer comes from assessing the uncertainties involved. Initially, the expected holdings of this security are going to be rather small, so its performance contribution will be relatively minor. Therefore, the difference made by this new instrument to an attribution analysis (which includes parallel yield curve shifts and credit spreads) will be small — so small that it won't matter for any realistic purposes.

Accordingly, we agreed to put the performance contribution of this new instrument into a separate category, called unattributed. This way we can continue to use the existing attribution reporting and provide feedback to the bank about the sorts of investment risks the managers are taking, while including this new security in the report.

Of course, if the performance contribution from the unattributed bucket becomes large, then we'll need to have the extra attribution machinery in place. But that won't happen for a while, by which time a tried and tested CDO attribution module will be available.

If you let it, over-engineering can be a big distraction from the rest of the investment management process. As long as you have a handle on the size of your data tolerance issues, you'll save time and effort. **IPMR**