FLYING BLIND

Benchmark standardisation seems to have passed the banking world by, laments *Andrew Colin* of StatPro

MEET ERIC: EVERY fixed-income manager has an 'Eric' somewhere. Eric tends to be an overworked quant in the middle or back office, with one of the most thankless yet vital tasks in the whole investment process – that of maintaining the manager's fixed income benchmarks (or indices).

Every day, Eric sits down and runs a group of programs developed in-house – often using spreadsheets and ad-hoc databases – to update benchmarks from assorted files supplied by the banks. This benchmark data consists of name codes for securities, face value holdings, coupon payments and price.

It may also include risk figures (duration, modified duration, convexity), coupon, yield, margin to swap, and other descriptive quantities.

Fiddly and tedious

Eric's job is time-consuming, fiddly and tedious. It needs repeated checking and verification. But it's absolutely vital to your institution's financial health, because without accurate, timely benchmarks and (most importantly) detailed data, you're effectively flying blind, with only the vaguest idea of where your portfolios have been and where they are going, relative to the rest of the marketplace.

Most modern applications rely on standardised software and data formats that make everyone's life easier. Imagine the reaction if each bank supplied its own research in its own custom format, instead of Word or PDF.

So can't Eric just download the data he needs from the internet? Curiously, no. When it comes to benchmark data, standardisation seems to have passed the banking world by. Every index provider supplies this mission-critical data in a different format. As a result, Eric's job is ten times harder than it need be. More importantly, it's easier for errors to creep in.

But it's actually much worse than this.

Security-level data

Eric has a hard job, and it is about to become a lot harder. The reason is the emergence of more sophisticated risk management and attribution systems that specifically require security-level benchmark data – in particular, returns.

As currently supplied, benchmark data is usually adequate for a snapshot of market risk. The supplied exposures, prices and risks come straight out of the index provider's computer and can be taken as definitive.

The tricky stuff starts when you try using this data to get to returns. This is where the traditional approach of supplying benchmark holdings files falls down, because they don't include security-level returns, nor is there any easy way to get them.

At StatPro, we've just completed a year-long project to reverse-engineer the fixed income benchmarks our customers use, using exactly the same data that index providers supply.

In particular, we wanted to be sure we could replicate the published returns of each benchmark from the supplied security-level holdings files. We eventually got there, but there were some major surprises along the way.

The first problem was that the published description of how the benchmark works was often quite different to the way that the numbers are actually combined together to produce the published returns.

A free-for-all

The project discovered that
• the pricing and treatment of cash flows were often unspecified (particularly for timing of coupons, coupon payments on holidays, and non-standard first and last coupons);
• the treatment of reinvested coupon amounts was murky;
• settlement periods were a free-for-all: some suppliers used same-day settlement for bond pricing, while others used country-specific conventions;

• and the only consistent factor was that nobody documented what they were doing.

Bizarrely, we couldn't even rely on the index provider that supplied the data to give a clear explanation of how the index fits together.

Working through a co-operative (and very large) European bank, I got it to ask the major US bank who publishes its benchmarks for a detailed account of how some of the more esoteric coupon reinvestment features worked.

The answer was: "We don't have time to tell you right now, but if you find out could you let us know?"

Clearly, it didn't understand the details of how its benchmark worked. Not an encouraging situation for anyone who wants to get from benchmark holdings files to returns data.

Substantial risk

So here's the problem: the data the index providers are supplying to Eric are not going to meet the requirements of your institution over the next few years.

In its current form, the way benchmark data is supplied is becoming a substantial business risk for the majority of fixed-income institutions, not just for post-hoc performance and attribution, but also for live risk management.

For example, what's the cost of finding a

THE DESCRIPTION OF HOW THE BENCHMARK WORKS WAS QUITE DIFFERENT TO THE WAY THE NUMBERS ARE COMBINED TO PRODUCE THE RETURNS

benchmark has changed its exposures and has shortened by a quarter of a year, the yield curve has moved substantially, and you haven't restructured your flagship portfolio in time because the benchmark holdings weren't up to date? Don't laugh, I've seen it happen.

The old days

In the old days, we got by using the daily holdings sheets for the one or two benchmarks required, and by using published returns for the benchmark as a whole.

If we wanted to match exposures by duration or maturity bucket, we'd do it by using published data. The more sophisticated fund manager might have a spreadsheet loaded with security-level benchmark holdings and calculate these numbers on the spot. And that was it.

Sadly, although the technology has moved on, the benchmarks haven't.

In the new world, we're running detailed attribution analyses and zeroing in on individual stocks rather than market sectors; we're using multiple benchmarks, over numerous currencies and asset types; and holdings are updated daily, rather than weekly or monthly.

For the level of results demanded by the marketplace, this all needs much, much more detailed information about the benchmark – ▶▶

BENCHMARKDATA

HOW MUCH EASIER LIFE WOULD BE IF ONLY THESE NUMBERS WERE AVAILABLE IN A STANDARD MACHINE-READABLE FORMAT

information that in many cases simply cannot be supplied in its current form.

In particular, the old approach is about to be rendered obsolete by the advent of powerful fixed-income attribution systems, for which it is vital to be able to complete the circle by moving from holdings to returns in a transparent manner. As we found, this just isn't possible without a huge amount of unnecessary effort.

Overworked index desks
Unnecessary? This is perhaps the strangest thing about the whole problem, because index providers already have this information.

Index desks do a great job, but are usually overworked and get their results from large legacy applications that need a programmer or a mathematician to understand in depth.

For example, a few years ago, I had some detailed questions for a major index provider in the Asia-Pacific region. The front-line staff couldn't answer my questions, but they passed me from desk to desk until I ended up talking to the original programmer who had written the system – who answered my questions.

But I was lucky: it's very difficult to get the same sort of feedback from other benchmark suppliers.

Security-level, verified weights and returns data is there – somewhere – inside the index

provider's software. But it's inaccessible, because making it available to customers will involve lengthy, expensive modifications to the application.

Daily security returns
Why is it inaccessible? Because, almost certainly, nobody asked for it before. Ten years ago, how many fixed income managers needed daily security returns?

So the customer has to reproduce and validate the same numbers, at a great deal of cost, expense and effort. How much easier life would be if only these numbers (weights, returns) were available in the same way as prices and coupons – in a standard machine-readable format.

So what should the asset management industry be demanding?

Three suggestions
Customers understand commercial pressures, and are willing to pay for a quality product – it's unrealistic to expect benchmark data to be available for nothing.

Instead, here are three suggestions for index suppliers:
• include security-level returns with raw benchmark data;
• fully document the algorithms that produce the

benchmark return – that is, show how to go from security-level returns to the published benchmark return. The documentation can be something as simple as a spreadsheet template or a list of instructions, but it has to work;
• and provide data in a standard, machine-readable format.

Easy and reliable
Index providers who can't show a transparent route from holdings to returns are going to experience growing reluctance to use their products. Business risks aside, is a fund manager going to put in the substantial effort required to reverse-engineer a benchmark when the alternative is so much easier and more reliable?

What if you're an index provider that puts these ideas into practice? You will find your market share will grow and your customers will be delighted – and Eric will be too.

• Andrew Colin is fixed income research director at the Statpro Group, managing director of Statpro's Australian office, and adjunct professor in the Faculty of Business, Queensland University of Technology, Brisbane

• Note: The opinions expressed in this article are those of the author