ANDREW COLIN

FIRST PRINCIPLES ATTRIBUTION IN FIA

FLAMETREE TECHNOLOGIES PTY LTD

Introduction

This paper describes in detail how Flametree FIA performs *first-principles* attribution on fixed income securities. While the presentation is based on a bond's cash flows, the same approach can be applied to any other type of fixed income security, such as a floating rate note, an MBS, or a swap.

The perturbational model for attribution

Most commercial attribution systems rely on the perturbational model (see Colin, 2016).[1] This approach uses externally sourced risk numbers as a proxy for a full pricing model, via the perturbational equation

[1] Colin, A, Chapter 8. *Mastering Attribution in Finance*, Pearson/FT Press, 2016

$$r = y \times \delta t - MD \times \delta y + \frac{1}{2} C \delta y^2 \qquad (1)$$

The perturbational approach does not require that any securities be priced. Instead, the return of the security is closely approximated by using equation (1) and its risk numbers

- y (yield to maturity)
- MD (modified duration)
- C (convexity)

When combined with the elapsed time δt and the change in yield δy, the first term in (1) gives the security's carry return, and the second and third terms provide the return generated by changes in its yield. Change in yield δy can be further decomposed into parallel and non-parallel movements in the overall yield curve, into risk-free and credit components, or some other decomposition that matches the investment process.

The advantages of this approach is that the same approach works for virtually all securities. If accurate risk numbers are available for all securities at all dates, this can be a straightforward way to run attribution. However, perturbational attribution can also be very data-intensive, and providing accurate risk numbers for some securities such as OTC swaps may be difficult.

FIA therefore offers a first-principles approach to attribution, in addition to the perturbational approaches described above.

The first-principles model for attribution

The first-principles approach does not require any risk numbers. Instead, a library of pricing functions is used to calculate the price, and hence return, of all securities at each stage of the attribution process.

In contrast to the perturbational approach, the first-principles approach only requires knowing

- the security's future cash flows, and
- the yield curve at each date for which attribution is to be calculated.

This typically reduce data requirements by orders of magnitude.

Pricing models must be provided for all security types, but in our experience calculation time is not an issue when suitably optimised code is used.

Prerequsites

The price P of the bond is given by

$$P = \sum_t \frac{CF_t}{(1 + y_t)^t} \tag{2}$$

where

- CF_t is the cash flow of the bond at time t
- y_t is the discount yield at time t, interpolated from the yield curve for that date

The return r of the bond over the interval $[t_0, t_1]$ is given by

$$r = \frac{p_1 - p_0 + CF}{p_0} \tag{3}$$

where

- p_0 is the price of the bond at time t_0
- p_1 is the price of the bond at time t_1
- CF is the value of any cash flows over the interval $[t_0, t_1]$

Carry return

Carry return is the return due to the passage of time, and is driven by both regular payments (coupons) and the proximity of the security's maturity payment. It is generated even if there is no change in market conditions, such as a curve movement.

In principle, it is possible to calculate carry return by pricing the security with the same yield curve at the start and end of the interval. However, this can lead to problems generated by the precise timing of coupon payments, which can be affected by weekends, public holidays and bond issuer conventions.

In practice, a more robust solution is to calculate the security's yield to maturity, and then using perturbational attribution to determine carry return.

FIA calculates carry by first deriving the security's YTM from its price, using an iterative calculation. The price of the security is either calculated from (2) or read, and the expression

$$P = \sum_t \frac{c_t}{(1 + YTM)^t} \tag{4}$$

is solved numerically for YTM. Carry is then given by

$$c = YTM \times \delta t \tag{5}$$

Curve return

This section covers all return from sovereign and credit curves. It also applies to IL (inflation linked) bonds, where a real yield curve is used for pricing rather than a zero coupon curve.

To run first-principles attribution, FIA sets up a set of yield curves that iterate through each attribution risk.

For instance, to model the effect of parallel and non-parallel shifts in the sovereign curve, the first curve is the yield curve at the start of the interval. The second curve is the same curve, plus a parallel shift over the interval has been added. The third curve is the curve at the end of the interval. The interim curves are completely artificial; they are are set up to measure the effects of different types of yield curve movements on the security's price.

In general, if there are n risks to measure, $n + 1$ curves must be set up.

By pricing the same security using these three curves, and finding the successive differences in its price, the return from parallel shift and non-parallel shift can be found.

The same approach can be used for any yield curve decomposition. For instance, the key rate duration (KRD) approach requires a set of curves where changes are added in at varying tenor points, such as 0, 1, 3, 5, 10 years. In this case, FIA creates a set of artificial curves with a perturbations at each tenor point added in successively. The returns generated by each variation can then be measured by pricing each security before and after this change has been made to the yield curve.

The same approach is also used for credit, by measuring the effects of spread changes between risk-free and credit curves.

Example

In the following example, the user wishes to see return due to the curve decomposed into returns from parallel and non-parallel movements in the yield curve.

Consider a bond with the following cash flows:

Year	0	1	2	3	4
Cash flows	\$5	\$5	\$5	\$5	\$105

Table 1: Bond cash flows

Over a calculation interval of 3 months, the yield curve changes as follows:

Tenor	0	1	2	3	4	5
START	5.1%	5.2%	5.3%	5.4%	5.5%	5.6%
END	5.0%	5.1%	5.1%	5.2%	5.2%	5.3%

Table 2: Yield curve changes

Carry return

The price of the bond at the start of the interval is \$103.2902, implying a YTM of of 5.4877%. Assuming that the calculation interval is 0.25 of a year, the return due to carry is therefore

$$r_{carry} = y \times \delta t = 0.054877 \times 0.25 = 1.372\%. \qquad (6)$$

Curve return

There are several ways to work out the parallel movement in a curve. Here, we use one of the simplest, which is arithmetic averaging. The average yield at the start of the interval is 5.35%, and the average yield at the end is 5.125%, so the difference between the two is -0.225%; this is the average downwards movement of the curve.

As there are two curve effects to calculate, we therefore need three curves: one at the start of the interval, one with parallel shifts added in, and one with parallel and non-parallel shifts - which is the curve at the end of the interval. The three curves are therefore

Tenor	0	1	2	3	4	5
START	5.100%	5.200%	5.300%	5.400%	5.500%	5.600%
PARALLEL	4.875%	4.975%	5.075%	5.175%	5.275%	5.375%
END	5.000%	5.050%	5.100%	5.150%	5.200%	5.250%

Table 3: Successive yield curves

First principles attribution requires that the security be priced using each curve. The prices, and hence returns, from each stage are shown in Table 4.

Curve	Price	Return
START	103.2901522	
START+PARALLEL	104.0740599	0.7589%
END	104.3156096	0.2321%
TOTAL		0.9910%

Table 4: Curve returns

Assuming that there are no credit effects, we now have an attribution decomposition for our bond.

Risk	Return
Carry	1.3719%
Parallel shift	0.7589%
Non-parallel shift	0.2321%
TOTAL	2.3629%

Table 5: Attribution analysis

Risk numbers

We have already shown how YTM is calculated. Other risk numbers are available using the expressions

Modified duration

$$MD = -\frac{p_1 - p_0}{2 p_0 \Delta}. \tag{7}$$

where

- $p_0 = p(y, d)$
- $p_1 = p(y + \delta, t)$
- $\Delta = 0.00001$

Convexity

$$C = -\frac{p_2 - 2p_1 + p_0}{p_0 \Delta^2}. \tag{8}$$

- $p_1 = p(y, d)$
- $p_1 = p(y + \delta, t)$
- $p_2 = p(y - \delta, t)$

- $\Delta = 0.00001$

As discussed previously, these risk numbers are not used in the first principles attribution calculation, but are available for other types of risk analysis.

Discussion

The main issue to bear in mind when using first-principles attribution is calculating the security's carry return correctly. For instance, if a corporate bond is priced using the sovereign curve, its YTM will be lower than the actual value and as a result its carry return will be understated.

Two solutions to this problem are

1. Provide an additional credit curve that matches the bond's rating. Then the bond will be priced from this credit curve and its YTM will be more accurate. Curve and credit returns will then be calculated from the supplied curves, and any residual return will be flagged as security-specific. In this case, there is no way to discriminate between residual due to carry differences and residual due to other security-specific effects, but the difference is likely to be small.

2. Provide a market YTM in the returns file. This will override any calculated YTM and as a result the carry return will be accurate. The other parts of the calculation will proceed as before, with all residual return assigned to credit.