

ANDREW COLIN AND KATALIN KISS

A GENERALIZED HYBRID
FIXED INCOME ATTRI-
BUTION MODEL

FLAMETREE TECHNOLOGIES PTY LTD

Copyright © 2017 DST Global Solutions Limited and Flametree Technologies Pty Limited

Revised and corrected November 2017

All Rights Reserved. No part of this paper or any of its contents may be reproduced, copied, modified or adapted, without the prior written consent of both DST Global Solutions Limited and Flametree Technologies Pty Limited.

Introduction

THE REQUIREMENT for fixed income attribution continues to rise dramatically as the value it brings to the investment process becomes more widely appreciated.

The attribution model described in this paper addresses both user demand and the limitations of existing commercially available systems. It incorporates extensive input from industry practitioners, portfolio managers, software vendors and technology experts to provide a workable, cost-effective solution to many problems that have previously affected this area.

The model has been implemented by Flametree Technologies and is currently licensed to several major industry vendors, with further partnerships under discussion.

The FIA concept

THE FLAMETREE APPROACH arose as a response to the shortcomings of commercially available attribution systems, including requirements for large volumes of daily risk data, long implementation times, high costs and inflexible modeling capabilities.

Our aim has been to provide a fixed income attribution *engine* that addresses all of these issues. This approach has many advantages such as

- no requirement for risk numbers, which are calculated internally
- use of user-supplied returns to ensure consistency between other performance data and attribution reports
- very high calculation speed
- widest possible range of attribution models, allowing an exact match between investment process and attribution reporting
- rapid deployment, since only three data files are required to run attribution.

Design considerations

THE DESIGN of a successful fixed income attribution model is a difficult problem that has defeated many practitioners. Perhaps the biggest - but least appreciated - obstacle is that attribution requires integration of many apparently unrelated mathematical, computing and financial factors such as the underlying data model and workflow, depth and configurability of analysis, cost and data tradeoffs, and simplicity of deployment.

The guiding principles in the design of our model are:

1. *The assumption that weights, base and local currency returns are available at the security level.* Virtually any portfolio manager with a requirement for attribution will have performance systems in place that can provide this information, while benchmark data is generally available at the same level of detail. We see the provision of accurate returns data as an entirely separate issue from the provision of attribution analysis, especially since performance calculation packages are widely available.

The use of externally calculated returns ensures that attribution reports are always consistent with other performance reports, so no reconciliation of results between systems is required.

2. *Low data requirements.* Historically, a major obstacle to the successful implementation of attribution models has been an over-reliance on perturbational models which require large volumes of high-frequency risk numbers (yields, modified durations, convexities) as proxies for exact pricing functions.

The model described here allows the perturbational approach to be replaced or supplemented by a first-principles approach to pricing, which only requires yield curves and limited static security definitions. Instead of requiring daily risk numbers, we allow securities to be priced by discounting their cash flows under a range of interest rate scenarios. This typically reduces data volume requirements by orders of magnitude.

3. *Flexibility.* Attribution reporting requirements vary greatly in complexity and scope, and it is critical for user acceptance that the attribution model be configurable to match whatever investment process is in use.

Many existing fixed income attribution systems are constrained by design to a small number of attribution models that may not

match a manager's changing requirements. For instance, an investor whose investment process involves extensive spread duration allocation will find only limited value in bottom-up curve analysis.

4. *The ability to mix top-down and bottom-up attribution models to generate hybrid models*, thus blurring the distinction between equity, fixed income and balanced (or mixed) attribution requirements.
5. *A 'plug and play' approach*. We believe it should not be necessary to replace a manager's core performance system solely to introduce an attribution capability. By deploying the attribution system as a stand-alone module with clearly defined, self-contained IT and data requirements, the user can continue to use an existing performance engine as a feed for their new attribution capabilities. The performance software may then be replaced at a later date if need be without affecting the attribution workflow.
6. *Recognition that FI attribution is significantly more complex than equity attribution*, and that it requires entirely new sources of data, as well as new expertise. We see simplicity of deployment as critical to meeting this requirement.

Security coverage

OUR MODEL is designed to be able to model all security types traded in the marketplace, either currently or in the future. This is achieved in three ways:

1. *Use of a first-principles pricing library*. The Flametree engine uses a core library of building block pricing routines that allows the vast majority of securities to be accurately modeled. These include government and corporate bonds (both investment grade and high yield), agencies, securitized debt such as ABS and MBS, CMBS, money market securities, callable bonds, sinking securities, various types of FRNs, futures, forwards, FX options, IRS, inflation-linked securities, emerging market debt, and others.

In addition, the program's nested portfolio and lookthrough capabilities allow the definition of new security classes that are portfolios of individual sub-securities. For instance, a vanilla interest-rate swap can be represented as a portfolio containing +1 unit of a bond, and -1 unit of an FRN, where the bond and FRN are modeled using the standard building blocks for these instruments.

2. *Mix-and-match first-principles and perturbational attribution.* Our model allows both first-principles pricing and the use of risk numbers as a pricing function proxy. For cases where no pricing model is available (such as a complex CMBS), or where there are sources of accurate risk numbers, it may be preferable to use perturbational attribution. This feature can be configured from the security level upwards, and can change over time using the program's effective date functionality.

3. *Use of security-level customizable output buckets.* In some cases the ability to direct return to a particular risk source may be needed. For instance, a credit default swap may either have its return allocated to the generic *Credit* sector or to a custom sector such as *CDS return*.

In practice, we find that this combination of approaches allows flexible and complete coverage of all security types without the requirement for complex workarounds.

The mathematics behind first-principles and perturbational attribution is covered in more detail below.

Attribution models

NO INDUSTRY STANDARDS exist for fixed income attribution, and in our view this will remain the case indefinitely, due to the wide range of investment approaches active in the market.

Rather than imposing a particular attribution methodology on the user, our approach has been to acknowledge this lack of standardization and to make the attribution model as flexible and configurable as possible. It may then be adapted to measure the specific returns of whichever investment process is in use.

Although the Flametree engine has been designed as a fixed income attribution system, the program's top-down attribution capabilities allow it to generate equity attribution reports as well, simply by allocating all non-allocation returns to a custom *stock selection* category.

Top-down attribution

Market allocation

Top-down attribution measures the effects of allocation decisions by market weight, duration contribution, or other measures such as duration times spread. The well-known Brinson models are market weight allocation algorithms, which decompose active return from each sector S into return contributions from asset allocation and stock selection c_S^{AA} and c_S^{SS} , using

$$c_S^{AA} = (w_S^P - w_S^B) \times (r_S^B - r^B) \quad (1)$$

and

$$c_S^{SS} = (r_S^P - r_S^B) \times w_S^P \quad (2)$$

where w_S^P and w_S^B are the weight of the sector in portfolio and benchmark, r_S^P and r_S^B are the returns of the sector in the portfolio and benchmark, and r^B is the overall return of the benchmark.¹ The sector weights and returns are given by

¹ For simplicity, we have aggregated interaction return with stock selection return.

$$w_S^P = \sum_{i \in S} w_i^P \quad (3)$$

$$w_S^B = \sum_{i \in S} w_i^B \quad (4)$$

$$r_S^P = \frac{\sum_{i \in S} w_i^P r_i}{\sum_{i \in S} w_i^P} \quad (5)$$

$$r_S^B = \frac{\sum_{i \in S} w_i^B r_i}{\sum_{i \in S} w_i^B} \quad (6)$$

and the benchmark return by

$$r^B = \sum_{i \in B} w_i^B r_i \quad (7)$$

where the sums are over individual securities in the given sector or portfolio.

Market allocation attribution is a useful a tool to measure the effectiveness of equity management. The reason it is not widely used in

the fixed income world is that it ignores some important risks. For instance, if you choose to invest 50% of your funds into fixed income, the Brinson model will ignore the difference between a T-bill and a 30-year T-bond, despite the two having radically different levels of interest rate risk. This is why duration allocation attribution is preferred by bond fund managers: it is a much closer match to the way that portfolio managers actually implement investment strategies.

For instance, if a contraction in credit spreads in the 5-10 year region of the curve is expected for US corporates, it makes sense to overweight duration in that part of the portfolio, and duration allocation attribution will measure the return made by this decision.

Duration allocation

Duration allocation models measure return made by duration mismatches between portfolio and benchmark at various levels. This results in three categories of attribution return (market direction return, market allocation return, security selection return) rather than the two from the market allocation model (asset allocation, stock selection).

Duration asset allocation requires the following data for each security i in portfolio and benchmark:

w_i^P, w_i^B : weights of security in portfolio and benchmark;

MD_i : modified duration;

y_i : yield to maturity;

δy_i : total change in yield to maturity;

δt : interval over which the calculation is to be performed, as a fraction of a year

To perform duration attribution, a yield is first calculated for each sector S in the benchmark:

$$\delta y_S = \frac{\sum_{i \in S} (w_i^B \times MD_i \times \delta y_i)}{\sum_{i \in S} (w_i^B \times MD_i)} \quad (8)$$

Note that benchmark weights are used for calculation of all aggregated yield changes.

The *market direction return contribution* is given by

$$c^{MD} = -(MD^P - MD^B) \times \delta y^B \quad (9)$$

where δy^B is the overall change in yield for the entire benchmark, and MD^P and MD^B are the modified durations of portfolio and benchmark respectively.

For each sector S , the *contribution to market allocation return* is given by

$$c_S^{MA} = -(w_S^P - w_S^B) \times MD_S^B \times (\delta y_S^B - \delta y^B) \quad (10)$$

where δy_S^B is the change in benchmark yield for sector S .

Lastly, individual security selection returns for security i are given by

$$c_i^{SS} = -(w_i^P - w_i^B) \times MD_i \times (\delta y_i - \delta y_S^B) \quad (11)$$

where quantities with a suffix S refer to the sector of that name.

The sum of returns over all risks from equations (9), (10) and (11) reduces to

$$r_i = - \sum_i (w_i^P - w_i^B) \times MD_i \times \delta y_i \quad (12)$$

which is the security-level active return generated by each source of risk, as expected.

Note that it is possible to mix and match attribution models in the same analysis. For instance, carry return may be analyzed using market value attribution, but spread returns using a duration allocation model. This point is illustrated in the worked example below.

Multiple level allocation attribution

Allocation decisions are frequently made at multiple levels within a portfolio. For instance, a manager might firstly decide to overweight certain countries, and then to overweight particular industrial sectors within each country. Transparent attribution requires that the excess return generated by each type of decision be reported separately.

To achieve this result at the sector level requires that the benchmark be temporarily reweighted so that so that, *at the country level*, it becomes identical to the portfolio. With this adjustment in place, there is no country overweighting or underweighting in play. Returns from country allocation will be zero, so any remaining asset allocation returns will be due to industry allocation.

Our model provides this functionality as standard for both market weight and duration allocation attribution, for any number of nested allocation decisions.

Hybrid models

A CRITICAL REQUIREMENT for many users is the ability to implement a hybrid model, which incorporates both top-down and bottom-up returns. For instance, we have seen credit desks that require Brinson analysis for carry return, spread duration allocation attribution for credit spreads, and key rate duration analysis for sovereign curve movements. This is easily achieved using the building blocks described above.

A particular example of a hybrid attribution model is a mixed attribution model. In this case, asset allocation returns are calculated using asset type weights. Where appropriate, the remaining stock selection weights are then decomposed in terms of bottom-up effects.

For instance, a mixed portfolio might contain both equities and bonds, and in this case a mixed attribution model is appropriate.

The first decision made will be whether to over- or under-weight entire countries or asset classes relative to benchmark. The top level asset allocation return measures the return made by this decision.

The next level of decision may have been at the level of industry or sector weights within each asset type. Here, a second asset allocation return provides the return generated by these subsidiary decisions.

Lastly, stock selection return measures individual stock contributions to return. For fixed income securities, this return contribution may be decomposed further into returns from carry, curve and credit effects.

Mixed and hybrid attribution models can be combined if required, to measure return from asset type weighting, spread duration allocation decisions, and individual fixed income returns.

Bottom-up attribution

BOTTOM-UP ATTRIBUTION is the decomposition of a single security's return in terms of its sources of risk. For equities, the user must select a known set of risks and run a complex statistical study to measure the effects of these factors on the portfolio's return. For fixed income, the factors are usually assumed to be already known (passage of time, movements in curves) and the decomposition is run in these terms. Here, fixed income attribution is in essence a specialized form of risk adjusted attribution.

Single-security return decomposition is typically run in one of two ways:

Pricing securities from first principles

The most direct way to price a security is to calculate its individual cash flows, to price them using the appropriate discount rate, and to add them together:

$$p = \sum_i \frac{C_i}{(1 + r_i)^{t_i}} \quad (13)$$

where p is the security's price, C_i is the cashflow, r_i the interest rate, and t_i the time to maturity (in years) of the i th cash flow.

The security is priced with and without the effect of the current risk (such as a parallel curve movement, or a change in spread due to a particular credit factor), and the return due to that risk is then given by the difference between the two prices, divided by the starting price.

Calculating return using the perturbational equation

Assuming that the price p of an arbitrary security is a function of time t and yield y , we can express δp in terms of a Taylor expansion, and write

$$\delta p = \frac{\partial p}{\partial t} \delta t + \frac{\partial p}{\partial y} \delta y + \frac{1}{2} \frac{\partial^2 p}{\partial y^2} \delta y^2 + O(\delta t^2, \delta y^3) \quad (14)$$

If we divide throughout by p and write

$$r = \frac{\delta p}{p} \quad (15)$$

$$y = \frac{1}{p} \frac{\partial p}{\partial t} \quad (16)$$

$$MD = -\frac{1}{p} \frac{\partial p}{\partial y} \quad (17)$$

$$C = \frac{1}{2p} \frac{\partial^2 p}{\partial y^2} \quad (18)$$

equation (14) becomes

$$r \approx y \delta t - MD \delta y + \frac{1}{2} C \delta y^2 \quad (19)$$

where r is the security's local return, y its yield to maturity, MD its modified duration, C its convexity, δt the elapsed time (in years), and δy the security's change in yield over the calculation interval. y , MD and C are often collectively referred to as the security's *risk numbers*.

It is tempting to view equation (19) as a 'one size fits all' approach to attribution, and several commercial systems have been built on this basis. Unfortunately, the assumption is seldom valid. Many securities have other sources of return, such as inflation for TIPS and inflation-linked gilts; others (such as FRNs) have multiple risk sensitivities; the model is only exact for a security with a single cashflow; and some specialized types of securities, such as Australian and New Zealand bond futures, do not generate carry. Any system that offers perturbationally-based attribution should therefore offer the ability to customize the perturbation equation according to the type of security.

Unfortunately, supplying daily risk numbers can be a surprisingly difficult (and expensive) problem. It can take many man-months to set up reliable, robust feeds for risk numbers. Even after this point, risk numbers for some security types such as OTC derivatives may still need to be calculated in-house, and a single incorrect value can distort the entire analysis.

Sources of fixed income return

THE MAIN SOURCES of fixed income return are carry, curve and credit, although there are many others depending on the level of analysis required and the securities held.

Carry return

Carry return is the return generated by the passage of time, due to the payment of coupons and the approach of maturity, when a fixed income security must be redeemed at par. Carry return is closely approximated by

$$r_{carry} = y \times \delta t \quad (20)$$

where y is the security's yield to maturity, and δt is the elapsed time.

Carry return may be decomposed further in two ways:

Pull-to-par and running yield

A manager who has purchased a bond at a price below par will show positive returns from pull-to-par effects as it approaches maturity.

To view carry return broken down in this way, calculate the running yield, which is given by

$$r_{\text{running yield}} = C/P \quad (21)$$

where C is the security's coupon, and P is the (clean) price. This will give the instantaneous return of the security, ignoring any long-term capital gain effects. The pull-to-par yield is then the yield to maturity, minus the running yield.

Risk-free and credit carry

Another way to break down yield to maturity is to regard it as a sum of two yields: a *risk-free yield* and a *credit spread yield*. The carry return is then given by the sum of the risk-free carry and the credit carry:

$$r_{\text{risk free}} = y_{\text{risk free}} \times \delta t \quad (22)$$

$$r_{\text{credit spread}} = y_{\text{credit spread}} \times \delta t = (y - r_{\text{risk free}}) \times \delta t \quad (23)$$

$$r_{\text{carry}} = r_{\text{risk free}} + r_{\text{credit spread}} \quad (24)$$

This type of decomposition is of particular interest to credit traders, who may add value by investing in high-yield stocks without taking interest rate risk. If the strategy is successful, the credit carry of the portfolio will exceed that of the benchmark and the value added will be the difference between the credit carry for the portfolio and the credit carry for the benchmark.

Sovereign curve return

THE BULK OF RETURN in many portfolios is generated by parallel movements in the risk-free curve. This is often referred to as dura-

tion return, since its magnitude is equal to the (negative) modified duration, times the size of the parallel curve shift.

The calculation of this return requires knowing the size of any such parallel shift. As there is no industry-wide agreement on the definition of this quantity, we allow the average curve level to be defined using one of

- Arithmetic averaging (simple but overweights the short end of the curve, where sampling is more dense)
- Trapezoidal integration, which approximates the area under the curve and divides by the longest maturity
- The level of the curve at a maturity or modified duration point equal to that of the benchmark.

The parallel shift is then taken to be the difference between the average curve level at successive dates.

Other curve decomposition algorithms may be used as required. For instance,

- The shift/twist/butterfly (STB) model measures parallel shift as above, twist as the change in the slope of the curve between the 3 and 10 year points, and curvature as any remaining curve movement after parallel and twist movements have been removed.
- The principal component analysis (PCA) model calculates the eigenfunctions of the curve and allocates movement to movements of order 0, 1 and 2.
- The key rate duration (KRD) model perturbs the curve at user-supplied tenor points and uses this curve to calculate the effect of a change in the curve at this maturity only.

Credit movements

CREDIT EFFECTS are driven by changes in the spread between the sovereign curve and the sector curve for a particular security. Country curve allocation is a particular type of credit return, where (for instance) alpha may be generated by contracting credit spreads between debt issued by different countries in the Euro-zone.

Our model allows families of credit curves to be associated with particular securities, so that (for instance) the return made by changes in the AAA-AA, AA-A, and A-B spreads may be measured. Alternatively, a sector-specific or industry curve may be associated with the security, so that the return made by changes in the sector curve (sector return), and changes between the security's market yield and that curve (security-specific return) measured .

Additional effects

DEPENDING UPON the type of analysis required, other sources of return may be active. For instance, a portfolio with many MBS may generate substantial return from convexity, and in this case it makes sense to report convexity returns.

In addition to the main three sources of return, any of the following effects may be included in the system's outputs:

- Rolldown return
- Convexity return
- Inflation and break-even return (for inflation-linked securities)
- FX return
- Paydown return (for sinking securities such as amortizing bonds and MBS)
- Cash deposit return
- Price return

Other effects can be defined by the user, depending on the pricing model and the security's configuration parameters.

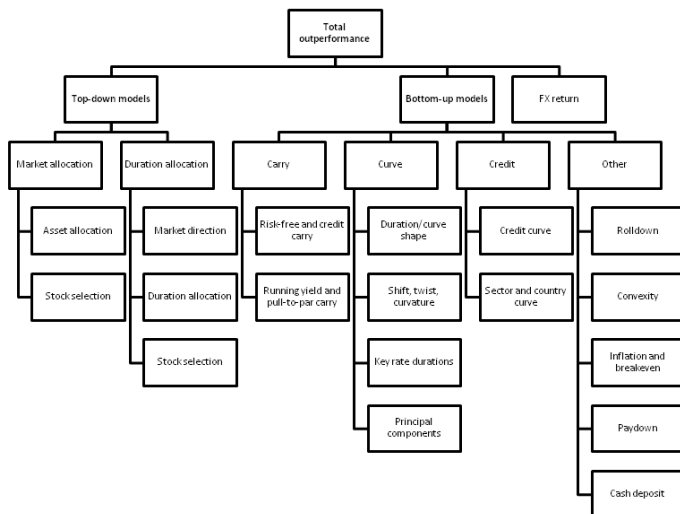


Figure 1: Sources of fixed income attribution returns

Implementation

The importance of implementation issues is often overlooked when selecting an attribution platform, but they can form one of the largest (and costliest) barriers to successful provision of an attribution capability. This rapid calculation ability also simplifies generation of attribution reports over long time periods. Figure 2 shows the result of a contiguous attribution calculation on a corporate bond portfolio over a two-year period. This capability is particularly useful when, for instance, a new investment mandate requires the provision of historical attribution reports in a particular format.

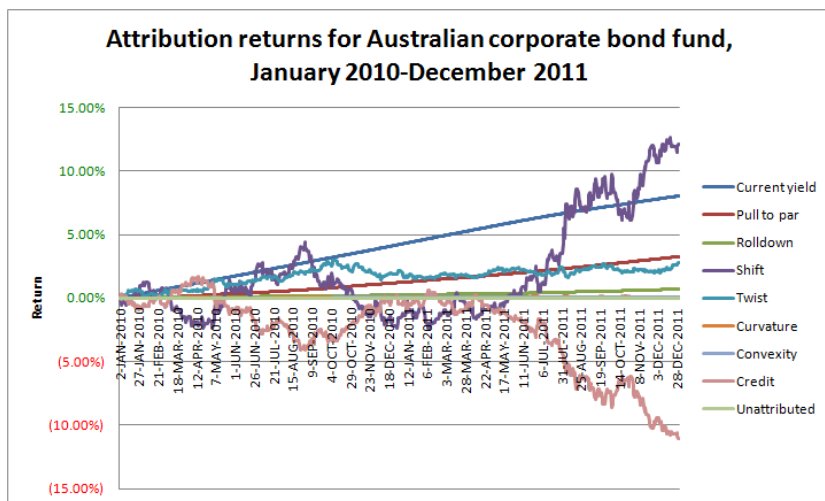


Figure 2: Cumulative attribution returns for an Australian corporate bond fund, 2010-2011

Future trends in fixed income attribution

Strategy attribution

A growing requirement for many managed funds is the ability to provide strategy attribution capabilities, in which holdings within a portfolio are assigned to one or more investment strategies, such as *duration bet, curve steepening, Latin American credit spread play*.

Although this is predominantly a data management issue, our model has the ability to support such analyses in a natural way by assigning returns from different securities to subportfolios. The full range of attribution and reporting capabilities is then available on each strategy.

Liability-driven investment (LDI)

An LDI strategy is driven by the requirement to fund current and future liability cash flows, rather than to beat a known benchmark. However, the requirement for attribution remains the same in both cases. Much of the responsibility of an LDI manager is to ensure that the portfolio is sufficiently hedged against different types of market movements, and an attribution analysis will supply clear and unambiguous feedback on whether this aim was achieved.

The Flametree model can be applied to LDI portfolios in exactly the same way as conventional managed portfolios, with the liability cashflows modeled in terms of conventional and inflation-linked securities.

Reporting

ALTHOUGH NOT strictly part of the attribution model, we regard reporting as an important part of the attribution process. The volume of data generated by an attribution analysis can easily overwhelm the user, and the provision of reporting tools and techniques to generate insight from this firehose of data is a vital part of the overall workflow.

Suitable reporting techniques range from the simple (rolling up performance contribution from benchmark stocks that are not held in the portfolio) to the widely used (generation of Excel reports using drill-down and roll-up capabilities, allowing the user to identify areas of interest) to the more exotic, such as interactive treemaps to

summarize the sources of active return.

The reporting requirements of portfolio managers and client reporting staff will usually differ widely. For a flagship fund, a simple statement of carry, curve and credit returns may be all that is needed by a marketing team, but the front office may require much more detailed analysis. The model described handles both cases easily by using different configuration sets. For instance, predefined templates are provided for duration/curve reshaping, Campisi, Tim Lord, key rate duration and top-down attribution models.

Summary

TEST DEPLOYMENTS and feedback strongly indicate that this model meets the vast majority of fixed income managers' attribution requirements, and that it can be implemented quickly with minimal business risk and at reasonable cost.

Worked example

ALL PERFORMANCE in this example is presented as performance contribution, which is the product of a security or a sector's weight and its return. Performance contributions can then be aggregated to sector or portfolio level.

Consider the following portfolio and benchmark. Each security lies in one of two sectors $\{S_1, S_2\}$.

Sector	Security	w_i^P	w_i^B	MD_i	y_i	δy_i
S1	A	13%	5%	1.97	3.30%	-0.70%
S1	B	13%	0%	2.33	3.40%	-0.60%
S1	C	22%	44%	2.89	3.25%	-0.40%
S1	D	6%	8%	3.05	4.40%	-0.20%
S2	E	8%	13%	3.43	4.40%	-0.10%
S2	F	10%	5%	4.80	4.90%	0.00%
S2	G	11%	10%	5.20	5.10%	0.10%
S2	H	17%	15%	5.80	5.10%	0.20%

Table 1: Weights, risks and yields for sample portfolio and benchmark

For security i ,

- w_i^P and w_i^B are security weights in portfolio and benchmark, respectively;
- MD_i is modified duration;
- y_i is yield to maturity;
- δy_i is aggregate change in yield.

The portfolio and benchmark have virtually identical modified durations (3.6902 vs 3.6900 years respectively). The period over which the returns are measured is 0.25 of a year.

The changes in yield may be decomposed further by source of risk. Assuming that yield changes are due to parallel, non-parallel and credit shifts, Table 1 can be supplemented by the data in Table 2.

Bottom-up attribution

One way to interpret these results is to take a bottom-up view, and to regard all return contributions as having been made at the security level. The displayed contributions are *active* (portfolio performance contribution minus benchmark performance contribution):

Security	$\delta y_i^{Parallel}$	$\delta y_i^{Non-parallel}$	δy_i^{Credit}
A	-0.20%	-0.50%	0.00%
B	-0.20%	-0.40%	0.00%
C	-0.20%	-0.30%	0.10%
D	-0.20%	-0.20%	0.20%
E	-0.20%	-0.10%	0.20%
F	-0.20%	0.00%	0.20%
G	-0.20%	0.10%	0.20%
H	-0.20%	0.20%	0.20%

Table 2: Changes in yield, decomposed by risk

Security	c_i^{Carry}	$c_i^{Parallel}$	$c_i^{Non-parallel}$	c_i^{Credit}
A	0.0660%	0.0315%	0.0788%	0.0000%
B	0.1105%	0.0606%	0.1212%	0.0000%
C	-0.1788%	-0.1272%	-0.1907%	0.0636%
D	-0.0220%	-0.0122%	-0.0122%	0.0122%
E	-0.0550%	-0.0343%	-0.0172%	0.0343%
F	0.0613%	0.0480%	0.0000%	-0.0480%
G	0.0128%	0.0104%	-0.0052%	-0.0104%
H	0.0255%	0.0232%	-0.0232%	-0.0232%
Subtotal	0.0203%	0.0000%	-0.0485%	0.0285%
Total				0.0002%

Table 3: Security-level performance contributions from bottom-up risks

Active carry contribution is calculated as

$$c_i^{Carry} = (w_i^P - w_i^B) \times y_i \times \delta\tau \quad (25)$$

and the active return contribution due to individual sources of risk is calculated as

$$c_i^{risk} = -(w_i^P - w_i^B) \times MD_i \times \delta y_i^{risk} \quad (26)$$

For instance, active carry return from security A was given by $(13\% - 5\%) \times 3.69\% \times 0.25 = 0.0660\%$, while active parallel shift return was $-(13\% - 5\%) \times 1.97 \times -0.20\% = 0.0315\%$.

At the aggregate level, carry generated an active return of 2.03 bp. The modified duration of portfolio and benchmark were virtually identical, so overall return due to parallel curve movements was zero, as expected. In addition, the portfolio lost 4.85 bp from non-parallel movements in the curve, but made back 2.85 bp from credit shifts. Overall, the portfolio's active return was very close to zero.

Top-down attribution

Suppose that the portfolio had instead been managed using a top-down duration allocation strategy, in which risk is apportioned to portfolio and benchmark sectors.

In this case, allocation decisions will have been made with a view to generating both excess carry return and excess return from duration. The attribution analysis will allow the effects of both decisions to be compared.

To see the effects of the allocation decision on carry return, break down the carry contribution as follows. Note that we are using a Brinson analysis for carry, since carry returns are driven by absolute yield rather than changes in yield over an attribution interval.

Allocation return for carry

The sector-level contribution to allocation return is given by

$$c_S^{AA} = (w_S^P - w_S^B) \times (r_S^B - r_B) \quad (27)$$

Sector	$w_S^P - w_S^B$	r_S^B	r_B	c_S^{AA}
Sector 1	-3%	0.8539%	1.0098%	0.0047%
Sector 2	3%	1.2163%	1.0098%	0.0062%
				0.0109%

Table 4: Asset allocation performance contributions for carry

Asset allocation return is a sector-level effect, so we only show returns at the sector level.

Stock selection return for carry

Stock selection is a security-level effect. Including interaction returns, it is given by

$$c_i^{SS} = (w_i^P - w_i^B) \times (r_i - r_S^B) \quad (28)$$

Overall, the weighting decisions made by the manager generated outperformance from carry, with a roughly equal contribution from allocation to sectors (1.09 bp) and from individual stock carry contributions (0.94 bp).

Security	$w_i^P - w_i^B$	r_i^{carry}	r_S^B	c_i^{SS}
A	8%	0.8250%	0.8539%	-0.0023%
B	13%	0.8500%	0.8539%	-0.0005%
C	-22%	0.8125%	0.8539%	0.0091%
D	-2%	1.1000%	0.8539%	-0.0049%
E	-5%	1.1000%	1.2163%	0.0058%
F	5%	1.2250%	1.2163%	0.0004%
G	1%	1.2750%	1.2163%	0.0006%
H	2%	1.2750%	1.2163%	0.0012%
				0.0094%

Table 5: Stock selection performance contributions for carry

Duration allocation return

Carry returns have been analyzed using a Brinson analysis, since these are driven by market weights.

By contrast, returns due to duration allocation allocation should be analyzed using a non-Brinson approach, since active returns from curve effects are driven by duration contributions.

For a duration allocation analysis, recall that there are *three* sources of active return rather than the two from Brinson attribution: *market direction*, *duration allocation*, and *duration selection* return.

Market direction return

Performance contribution from market duration effects c^{MD} is a global source of return, and is calculated at the portfolio level:

$$c^{MD} = -(MD^P - MD^B) \times \delta y^B \quad (29)$$

MD^P	MD^B	δy^B	c^{MD}
3.6902	3.6900	-0.2000%	0.0000%

Table 6: Market direction performance contribution

This zero value for market return reflects the virtually equivalent modified durations of portfolio and benchmark, reflecting the manager's decision to be neutral interest rate risk at the portfolio level.

In fact, the portfolio had several active duration decisions in play, but at the sector level rather than the overall portfolio level. The return made from these lower level duration allocation decisions are reflected in the next source of return.

Duration allocation return

Just as for carry allocation return, duration allocation return contribution c_S^{DA} is calculated by sector, using

$$c_S^{DA} = -(w_S^P - w_S^B) \times MD_S^B \times (\delta y_S^B - \delta y^B) \quad (30)$$

where δy_S^B is the change in benchmark yield for sector S .

Sector	$w_S^P \times MD_S$	$w_S^B \times MD_S$	δy_S^B	δy^B	c_S^{DA}
Sector 1	1.3778	1.6141	-0.3982%	-0.2000%	-0.0468%
Sector 2	2.3124	2.0759	0.0628%	-0.2000%	-0.0622%
Total	3.6902	3.6900			-0.1090%

Table 7: Duration allocation performance contributions

These results show that the manager made poor duration allocation decisions in both sectors.

For instance, in Sector 1 the fund was short duration by 0.2363 of a year, making the fund less sensitive to yield changes than the benchmark. Yields in this sector fell by 39 bp compared to the benchmark's decrease of 20 bp. This decrease in yields generated a rise in prices, but because the portfolio was shorter duration than the benchmark, it underperformed, generating a net active return of $-(1.3778 - 1.6141) \times (-0.3982\% - -0.2000\%) = -4.68$ bp.

Conversely, in Sector 2 the fund was 0.2365 years long, making it more sensitive to yield changes than the benchmark. The yield of this sector rose while that of the benchmark fell, driving down prices and hence returns. The net result was again negative for the fund, generating a performance contribution of -6.22 bp.

Duration selection return

Duration selection return is calculated on a per-security basis, using

$$c_i^{SS} = -(w_i^P - w_i^B) \times MD_i \times (\delta y_i - \delta y_S^B) \quad (31)$$

Security	$w_S^P \times MD_S$	$w_S^B \times MD_S$	δy_i^B	δy_S^B	c_i^{SS}
A	0.2561	0.0985	-0.7000%	-0.3982%	0.0476%
B	0.3029	0.0000	-0.6000%	-0.3982%	0.0611%
C	0.6358	1.2716	-0.4000%	-0.3982%	-0.0011%
D	0.1830	0.2440	-0.2000%	-0.3982%	0.0121%
E	0.2744	0.4459	-0.1000%	0.0628%	-0.0279%
F	0.4800	0.2400	0.0000%	0.0628%	0.0151%
G	0.5720	0.5200	0.1000%	0.0628%	-0.0019%
H	0.9860	0.8700	0.2000%	0.0628%	-0.0159%
Total					0.0889%

Table 8: Duration selection performance contributions

For instance, security A was overweight in the portfolio, and contributed 0.1576 years to the portfolio's active duration position. Its yield fell by 70 bp compared to the 39.82 bp decrease within Sector 1. The security's net contribution to performance (as distinct to the contribution of the sector, or the portfolio as a whole) was therefore $-0.1576 \times (-0.7000\% - -0.3982\%) = 4.76$ bp.

Summarizing the duration allocation analysis, we have

Risk	Return
Carry allocation	0.0109%
Carry selection	0.0094%
Market direction	0.0000%
Duration allocation	-0.1090%
Duration selection	0.0889%
TOTAL	0.0002%

Table 9: Summary results for duration allocation attribution

Comparing the top-down and bottom-up views

In both cases, the carry return aggregates to the same total, although in the top-down model carry can be decomposed in terms of allocation to particular sectors.

The negative curve return can either be assigned to adverse twist movements in the curve (bottom-up attribution) or to poor allocation to market sectors. The two risks are closely related, since Sector 1 held securities with shorter durations, while Sector 2 held securities with longer durations. In other words, this is exactly the type of behaviour we would have expected if the curve had flattened or steepened. However, we did not need a security-level breakdown to supply this information.

Which analysis should be used depends on how the portfolio was

managed. For a bottom-up manager, a duration allocation analysis would be inappropriate, as it would decompose returns in terms of risks that had never been examined or managed.

Spread duration allocation attribution

Many portfolios, particularly for emerging-market debt (EMD) are managed in terms of spread duration allocation rather than modified duration.

The techniques described in this paper are equally applicable to portfolios managed in this way.

Hybrid attribution

To construct a *hybrid analysis*, the duration selection term can be decomposed further. Table 8 shows the duration selection return due to the aggregated return for each security. If we use the data in Table 3 and construct similar tables, one for each source of risk, the result is a report that shows return from both top-down sources (market direction, duration allocation) and bottom-up sources, which replace the duration selection term. The result is shown in Table 10:

Risk	Return
Carry allocation	0.0109%
Carry selection	0.0094%
Market direction	0.00004%
Duration allocation	-0.1090%
Parallel curve return	0.0000%
Twist curve return	0.0380%
Credit return	0.0509%
TOTAL	0.0002%

Table 10: Summary results for duration allocation attribution

Note that the sum of the last three terms equals the duration selection return, and that the parallel curve return is zero. This is expected, as any parallel curve return will be appear in the *market direction* term. We have left the term in place as other curve breakdowns may not have a parallel curve shift term, such as a principal component or a key rate duration analysis. In this case, the term will not be zero.

Biographies

ANDREW COLIN is founder of Flametree Technologies, a company that provides innovative performance and attribution software to fund managers of all sizes. He was previously Head of Fixed Income Research at StatPro Ltd, and has held positions in finance, academia and defence in the UK and Australia.

Andrew holds a PhD in applied mathematics from the University of St Andrews. He is a Fellow of the Institute of Mathematics and its Applications, and holds Chartered Mathematician (C.Math) accreditation. He is also Adjunct Professor in the School of Business at the University of Tasmania.

Andrew may be contacted at andrew.colin@flametree.global.

KATALIN KISS is Global Solutions Manager - Performance at DST Global Solutions. Prior to joining DST, Katie was UK Head of Performance at UBS Global Asset Management where she was responsible for performance reporting and analysis.

Katie also led the strategic implementation of a new global performance tool across the organisation.

Prior to UBS Katie began her career as a performance analyst at Russell Mellon CAPS in 2002.

Katalin may be contacted at katie.kiss@dstgs.com.