ANDREW COLIN

MACRO ATTRIBUTION IN FIA

FLAMETREE TECHNOLOGIES PTY LTD

Motivation

In a multi-manager portfolio, there are two main drivers of active return: the asset owner and the individual managers.

Micro attribution evaluates the decisions of the individual investment managers, but it cannot offer a comprehensive analysis of the total fund, as it does not consider

- the allocation of funds within the plan's hierarchy
- the performance of other investment managers within the plan.

In addition to conventional attribution tools, asset owners therefore require a *macro level perspective* for high level decisions, and in particular the ability to run macro attribution. This complements conventional performance analysis by assessing decisions made by the asset owner with respect to specific benchmarks.

Macro attribution evaluates decisions made by the asset owner by attributing active returns to three core effects:

- Allocation effect: describes how well the asset owner allocated funds with respect to a benchmark
- Selection effect: describes how well the asset owner selected individual managers within the fund
- Benchmark misfit: the degree of mismatch between benchmarks at different levels in the hierarchy

Here, benchmark misfit represents the difference in benchmark returns between an asset or asset class and its enclosing group. Mismatch between benchmarks at different levels of the hierarchy represents an additional source of return in an attribution report.

Mismatch return appears in multiple places:

- At the manager level: when manager tracking benchmarks do not match the tracking benchmarks of the sub-group in which they belong
- At allocation levels: where the sub-group policy benchmarks do not aggregate to the benchmark of the parent group in which they belong

The aim of this section is

- to shown how to modify aggregated weights and returns of sectors within a benchmark (the manager benchmark) so that their values match an externally imposed set of weights and returns (the policy benchmark).
- to illustrate how to integrate benchmark mismatch return into the context of Brinson attribution.

Terminology

The *asset owner* is the overall owner of all investment assets within a plan (pension, foundation, endowment, etc.), which are invested with individual managers and rolled up to form an aggregate portfolio. This aggregated portfolio is called the *total plan* or the *fund*.

The asset owner's investment team or board of directors make decisions regarding the investments held within the plan. These decisions fall into two classes:

- The relative proportions invested at the asset class or sub asset class level. This is called the *asset allocation* decision.
- The specific market value invested in individual managers or funds. This is called the *stock selection* decision.

Examples of asset classes are 'Private equity', 'Absolute return', 'Public equity, Return seeking fixed income'.

The asset owner also specifies an asset class-level custom benchmark (the *policy benchmark*) to which the return of invested asset classes is compared. The policy benchmark is made up of weights and returns of various asset classes. The policy benchmark can also be defined at lower levels such as sub asset class, but is not defined at the individual manager/fund level.

Critically, the instruments in which the asset owner invests may have different benchmarks to the policy benchmark, as the instrument benchmark is defined by the portfolio/fund manager. Therefore, the weights and returns of the benchmark implied by these investments may

be different to those of the asset owner's externally imposed policy benchmark.

A managed fund can therefore be seen as having two benchmarks:

- an asset-level manager benchmark, and
- a sector-level policy benchmark.

The returns generated by this benchmark mismatch are generated by differences between the weights and returns of a manager's benchmark compated to the weights and returns of the original policy benchmark. The subject of this paper is the measurement of these mismatch effects in an attribution context.

In the following sections, we use the convention that security-level weights and returns are shown in lower case (w, r), while sector-level weights and returns are shown in upper case (W, R).

The *manager benchmark* B is a set of asset-level weights w_i and returns r_i supplied at the asset level:

$$B = \{(w_i, r_i); i \in B\} \tag{1}$$

where

$$\sum_{i \in B} w_i = 1 \tag{2}$$

The *policy benchmark* P is a set of sector-level weights W_S and returns R_P:

$$P = \{(W_s, R_s); s \in P\} \tag{3}$$

where

$$\sum_{s \in P} W_s = 1 \tag{4}$$

By construction, $W_s \neq 0$, as sectors with zero weight are not included in the policy benchmark P.

Modifying the manager benchmark

From a functional perspective, macro attribution requires the ability to overwrite aggregated sector weights and

returns at various levels within a benchmark B, so that its sector-level weights and returns match those of the policy benchmark P.

The most straightforward way to do this is to manually overwrite the weights and returns of the given sectors within the benchmark. However, this means that the ability to calculate overall benchmark weights and returns from constituent security weights and returns is lost, which greatly complicates reporting. In addition, the ability to run smoothing using standard Carino or geometric algorithms is lost.

For these reasons, FIA instead changes sector statistics by adding one or more dummy securities to affected sectors in the manager benchmark, so that its sector-level weights and returns match those in the policy benchmark. The details of this process are shown in the next section.

Matching the policy benchmark

Dummy security weights

The weight of dummy security w_D for sector S is given by

$$w_D = w_P - w_B \tag{5}$$

where

- w_P is the (supplied) weight of sector S in the policy benchmark P;
- w_B is the (calculated) weight of sector S in the manager benchmark B, given by $w_M = \sum_{i \in S \subset B} w_i$.

Dummy security returns

The return of the dummy security r_D for sector S is given by

$$r_D = \frac{1}{w_D}[r_P(w_D + w_B) - c_B] \tag{6}$$

where

- w_D is the weight of the dummy security for sector S, as calculated above;
- r_P is the (supplied) return of sector S in the policy benchmark;
- w_B is the (calculated) weight of sector S in the manager benchmark, calculated as above;
- c_B is the (calculated) contribution of sector S in the manager benchmark, given by $c_B = \sum_{i \in S \subset B} w_i r_i$.

Use of future-type securities

A special case to consider is the treatment of a policy sector that has the same weight as the corresponding sector in the manager benchmark:

$$w_D = w_P - w_B = 0 \tag{7}$$

Here, a single dummy security cannot be used to modify the sector's return, as its weight must remain unchanged. Instead we must use a future-type security, which has return but no weight.

Futures are modeled in FIA by adding *two* dummy securities to the sector:

- a leg with a weight $w = 1$ and a return of $r_P(w_D + w_B) - c_B$;
- an offset leg with a weight of $w = -1$ and a return of 0.

The net effect of these two legs is to change the return of the sector to the desired value without affecting its weight.

Order of calculation

FIA sets up these dummy securities by starting at the lowest tier of the benchmark, adding dummy securities and working upwards. This ensures that exposures and returns of dummy securities at the lower levels are included when dummy security holdings are calculated at higher levels.

Brinson analysis

This section assumes that dummy securities have been added to the benchmark, so that its sector-level weights and returns match those of the policy benchmark.

Asset allocation

The asset allocation contribution c_S^{AA} for sector i is calculated as follows:

$$c_i^{AA} = (w_i^P - \left[\frac{w_S^P}{w_S^B}\right] \times w_i^B) \times (r_i^B - r_S^B) \tag{8}$$

where sector i is a member of (higher level) sector S, and

- w_i^P is the weight of sector i in the portfolio
- w_S^P is the weight of sector S in the portfolio
- w_i^B is the weight of sector i in the benchmark
- w_S^B is the weight of sector S in the benchmark
- r_i^B is the return of sector i in the benchmark
- r_S^B is the return of sector S in the benchmark

Note that if the attribution analysis is run at the top-most level (that is, sector S is the entire portfolio or benchmark), expression [8] simplifies to

$$c_i^{AA} = (w_i^P - w_i^B) \times (r_i^B - r_S^B) \tag{9}$$

since $w_S^P = w_S^B = 1$.

Stock selection

Stock selection contribution c_i^{SS} for asset i is calculated at the asset level:

$$c_S^{SS} = w_i^P \times (r_i^P - r_i^B) \tag{10}$$

where w_i^P is the weight of asset i in the portfolio, and r_i^P and r_i^B are the returns of asset i in the portfolio and benchmark, respectively.

Manager mismatch return

Manager mismatch contribution c_i^M is also calculated at the asset level:

$$c_i^M = w_i^P \times (r_i^B - R_i^B) \tag{11}$$

where w_i^P is the weight of asset i in the portfolio, r_i^B is the return of asset i in the benchmark, and r_S^B is the return of the benchmark sector S to which asset i belongs.[1]

[1] Note that we have used w_i^P in place of w_i^B in expressions [10] and [11], so that interaction return is combined with both terms. For more on this topic, refer to Colin (2016).

Hierarchy mismatch return

Hierarchy mismatch return is, strictly speaking, an asset allocation effect. It is defined as asset allocation contribution from a give sector when the weight of that sector in the benchmark is zero. FIA moves any such contributions from asset allocation returns to a separate category, labelled as hierarchy mismatch.

Appendix

Note: This document refers to FIA version 2.8, build 264, and later versions.

Introduction

For a conventional Brinson analysis, FIA requires three sources of data:[2]

- a *portfolio file*, containing security-level weights and returns for the managed portfolio
- a *benchmark file*, containing security-level weights and returns for the manager benchmark
- a *security master file*, containing security-level identifiers, classification data and currency.

[2] In this section, we assume that data is supplied in CSV files and FIA is run from the command line. The same approach applies if data is supplied as in-memory arrays and FIA is called via the API.

Formats for all these files are supplied on the Flametree wiki at `http://www.flametree.wiki`.

To run macro attribution, FIA also uses a *policy benchmark* file, which contains dates, sector identifiers, weights, local and base returns. For each sector in the policy file, FIA adds dummy securities to the manager benchmark data. The weights and returns of these dummy securities are set internally so that the weights and returns of each sector in the benchmark match those in the policy benchmark.

A report run with macro attribution will show additional sources of return in the attribution reports, due to mismatches between the manager benchmark and the policy benchmark.

The policy benchmark file

The policy benchmark file contains five columns: dates, sector names, weights, and base and local currency returns.

Records can be supplied in any order in the file, and the dates must correspond to dates in the manager benchmark file.

However, no validation checks are performed on the sector names in the manager benchmark. If they are not already in the existing classification scheme, FIA adds a new sector to FIA's reports.

Column 2 contains data on the sector to which the correction is to be applied. This information is supplied in the form

$$SECTOR = VALUE1|SUBSECTOR = VALUE2|SUBSUBSECTOR = VALUE3$$

Note that

- this format is identical to the way that classification data is supplied in the security master file. The 'sector equals value' model is used for all classification, and sector assignments are delineated by a pipe symbol.
- The order in which sector assignments are provided must follow the classification hierarchy. In other words, the top-level classification must come first, followed by sucessively finer levels of detail.
- Sector names can be specified for different levels.
- Sector assignments can be nested.

Configuration file settings

There are five new macro-attribution related settings available to the user in the configuration file:

- **MacroFile**: contains name, and path if required, of the file containing data for the policy benchmark.
- **MacroDateFormat**: specifies the format of dates in the policy benchmark file
- **MismatchLabel**: specifies the label given to manager mismatch effects in FIA's reports. The default value is 'Manager misfit'.
- **HierarchyLabel**: specifies the label given to hierarchy mismatch effects in FIA's reports. The default value is 'Hierarchy misfit'.
- **HideDummySecurities**: Removes all system-generated dummy securities from FIA's interactive report.[3]

[3] With some securities removed from the interactive report, contributions and exposures at lower levels may not aggregate to the various totals displayed. FIA displays a warning at the bottom of each interactive report when this switch is active.

Worked example

The reader is referred to the attached files in the **macro** demo dataset, available with this paper or from the author:

macro_s.csv: (security master data)

macro_p.csv: (portfolio data)

macro_b.csv: (manager benchmark data)

macro_m.csv: (policy benchmark data)

macro.cnf: (configuration file)

The policy benchmark data looks as follows:

EFFECTIVE_DATE	SECURITY_ID_CODE	W	BASE
28/02/2025	CLASS=FIXED_INCOME|SECTOR=CORPORATE	0.3	0.06
28/02/2025	CLASS=FIXED_INCOME|SECTOR=GOVERNMENT	0.2	0.02
28/02/2025	CLASS=EQUITY|SECTOR=SMALL_CAP	0.3	-0.0843
28/02/2025	CLASS=EQUITY|SECTOR=LARGE_CAP	0.2	-0.0843
28/02/2025	CLASS=FIXED_INCOME	0.5	0.03

Results are similar to the table shown below. Note that benchmark weights and returns are identical to the values supplied in the policy benchmark data:

Sectors	Weight	Weight	Weight	Total return	Total return	Total return
CLASS/SECTOR/SECURITY	Portfolio	Benchmark	Active	Portfolio	Benchmark	Active
LARGE_CAP	29.200%	20.000%	9.200%	(6.064%)	(8.430%)	2.366%
SMALL_CAP	29.200%	30.000%	(0.800%)	(6.064%)	(8.430%)	2.366%
EQUITY	58.400%	50.000%	8.400%	(6.064%)	(8.430%)	2.366%
(Unclassified)	0.000%	0.000%	(0.000%)	0.000%	(0.000%)	0.000%
CORPORATE	16.600%	30.000%	(13.400%)	6.205%	6.000%	0.205%
GOVERNMENT	25.000%	20.000%	5.000%	3.500%	2.000%	1.500%
FIXED_INCOME	41.600%	50.000%	(8.400%)	4.579%	3.000%	1.579%
SUMMARY	100.000%	100.000%	0.000%	(1.637%)	(2.715%)	1.078%

Notes

Dummy security definitions and holdings are generated internally by FIA, and should not be set up by the

user. Dummy securities are identified with the prefix 'DUMMY' in FIA's reports.

The calculated returns of these dummy securities are often large, and this can prove confusing for consumers of the reports. If this is the case, we recommend only displaying contributions on FIA's reports, rather than raw returns.

Macro attribution is only run if a macro file is provided using the **MacroFile** setting. If this parameter is not set, or set to an empty string, macro attribution will not be run.

Further information

The files for this example, together with sample output, are available from the author.

This paper was written in collaboration with Bank of New York in preparation for joint development of a macro attribution capability for client funds.

For more information on Flametree Technologies' attribution capabilities, please contact the author at andrew.colin@flametree.global.