



Overview of FIA

EXECUTIVE SUMMARY

FIA is a software component that calculates attribution on portfolios of equities and fixed income securities.

The system models any investment process currently in use, including all those available in Barclays Point and Citi YieldBook.

Although it can be used as a stand-alone program, FIA is also designed to integrate seamlessly into PMAR workflows, allowing best-of-breed analytics to be deployed within existing workflows.

Designed by a team of industry thought-leaders, FIA is a highly cost-effective solution to the problems of high analytics software costs, business risk, and lengthy integration projects.



INTRODUCTION

This brief document provides an overview of FIA, Flametree Technologies' portfolio attribution software. The overview can only give a very general impression of FIA's capabilities. For more information, please visit our documentation wiki at www.flametreetechnologies.com.

DEPLOYMENT OPTIONS

FIA's core engine has been written to allow flexible deployment while providing an exact match to our users' analytics requirements. FIA can be run

- as a stand-alone program under Windows, using our GUI (graphical user interface)
- via a batch process, under any operating system
- as white label software, allowing FIA's functionality to be supplied under the client's brand
- via a RESTful API over the internet.

Source code is also available. Please contact us for more information on this option.

INSTALLING FIA LOCALLY

Under Windows, FIA is installed using a simple stand-alone installation utility. This sets up all libraries and other system requirements, as well as a library of test examples.

The process typically takes less than five minutes, after which FIA is ready to use.

PASSING DATA TO FIA

Flametree does not supply market data, and this must be sourced by the user.


However, FIA is designed to be 'data-light' and has much lower data requirements than competing attribution systems. In particular, the program does not require daily risk numbers (yield, duration, convexity) although these can be used if available. Typical data requirements are shown in Table 1.

Data type	Comments
Security level weights and returns (base and local currency)	Typically extracted from performance returns system (portfolio) or vendor files (benchmark)
Security master data	Names, IDs, classification data, credit rating, associated yield curves, maturity, coupon. This data only needs to be set up once and can be shared between portfolios.
Yield curve data (optional)	Zero curves, par curves, real curves, credit/sector curves
Index data (optional)	CPI data for IL bonds, LIBOR/EURBOR for FRNs, paydown schedules for sinking bonds
FX data (optional)	If supplied, allows calculation of base currency returns from local returns, and vice versa

Table 1



Detailed documentation on the formats of all files is available on our wiki.



Data can be supplied in three ways:

- As CSV files. These are read and parsed by the fia64 harness program and passed to the core attribution engine.
- Via in-memory arrays. If the library is called via FIA's API, data is passed via two-dimensional string arrays.
- As SQL tables. If data is already available in an external SQL database, FIA can read this information directly via an ODBC connection to Oracle, SQL Server, Informix and many other databases.

Much of this data can be supplied automatically. For instance, security master and yield curve data can be extracted from Bloomberg, and benchmark returns and security definitions from benchmark holdings files.

USING DATA FROM FIA

FIA makes its results available in the following ways:

- Reports in Excel format. These include drill-down summaries, security-level risks, positioning by duration and maturity categories, and many others



- JSON and CSV reports, containing the same information as the Excel reports
- SQL database link. FIA can write directly to an SQL database, using an ODBC link

CONFIGURABILITY

FIA allows extensive configurability. For instance, the same portfolio can be analysed in multiple ways to produce reports for widely different audiences.

SPEED

FIA is extremely fast. On a high-end PC, FIA can run attribution on a month's worth of daily data for the Barclays Global Aggregate benchmark (approximately 15,000 securities) in 30 seconds.

This speed also enables ad-hoc, interactive attribution analysis from a dealer's desk.

USE CASES

LOW VOLUME

For small numbers of portfolios, analysis may be performed manually. FIA's architecture gives immediate, targeted feedback on any data issues so that, for instance, incorrectly set up securities may be identified and corrected. FIA supports batching from the GUI.

HIGH VOLUME

For high volumes (hundreds of portfolios or more), we recommend workflow automation. FIA supports this by

- Providing extensive diagnostic information on missing data
- Definition and reuse of attribution model templates
- Batching portfolio processing as command-line processes, allowing deployment across multiple machines and time zones

At this scale, the user will ideally have set up suitable databases of security master data, yield curves, and other market quantities. Providing this information is available, integration with FIA is straightforward.

EXPANDABILITY

FIA can be extensively modified and reconfigured by the user, if required.

- *OpenPricing* allows new types of security to be defined, or new pricing functions for existing securities to be integrated. This allows proprietary or third party pricing code to be used within FIA.
- *OpenRisk* allows new types of market risk to be defined so that the return they generate is shown in FIA's attribution reports.



Full examples and documentation are provided.

While most users will never use these features, their availability ensures that FIA can always keep pace with new developments in the fixed income and derivative markets.

ATTRIBUTION MODELS

A brief overview of FIA's attribution models is shown in Table 2.

Risk type	Specific risk
Bottom-up	Decomposition of individual security returns due to Carry (none, aggregated, risk-free/credit carry, pull-to-par/running yield) Risk-free curve movements (none, aggregated, duration/non-parallel movements, shift/twist/curvature, principal component analysis, key rate durations, CCB model) Credit (shifts between one or more curves, including risk-free, swap, Treasury, sector and credit curves) Capital (return due to changes in YTM only) Security-specific yield changes - Movements in OAS (option-adjusted spread) and Z-spread Convexity Rolldown Cash deposits Inflation Paydown for sinking securities Optionality Leverage Price effect Residual effects
Top-down	Decomposition of returns due to asset allocation and stock selection, using the Brinson-Fachler and Brinson-Hood-Beebower models. Asset allocation return can be measured at multiple sector levels hybrid: combination of asset allocation returns (eg, allocation by country against benchmark), plus fixed income effects; carry and credit allocation
Duration allocation	Duration allocation return, where overweighting of duration is measured by sector and return is generated by movements in the risk-free curve. This generates a market direction return, a duration allocation return, and a security selection return, which can be further decomposed into bottom-up effects. The van Breukelen attribution algorithm is a specialised form of this approach.



	Spread duration allocation return, which is similar to duration allocation return but measures return due to credit spread changes rather than curve movements DTS (duration times spread) return. This is similar to spread duration allocation but uses DTS as a proxy for volatility.
FX	FIA implements naïve currency attribution. Karnosky-Singer attribution is in development. Hedging returns are also available

Table 2

Industry-standard attribution models such as Campisi, Tim Lord, key rate durations and DTS can all be expressed in the above framework, and are included as standard templates for new attribution analyses.

RETURNS DATA

FIA does not calculate returns on security holdings as this is a functionally distinct requirement to attribution.

However, we have available several returns calculation engines which have been developed for internal use. Please contact us for more information.

TECHNOLOGY

FIA is written entirely in cross-platform C++. Our clients have run the program under Windows, Linux and Solaris.

The program does not use a database and runs its calculations entirely in memory. However, data links are available to read and write results to and from SQL databases if required.

Harnesses are available to call the program from other languages, including C#, Java and Python.

DEVELOPMENT

Flametree are committed to continual improvement and total quality management of our products. We review market developments and listen to our users to ensure that our software matches our users' requirements, both current and future.

FURTHER INFORMATION

For more information, please contact

Andrew Colin PhD +61 413 385 057 andrew.colin@flametree.global

Graeme Condie +44 7960 866 039 graeme.condie@flametree.global