

Attribution for Asset Owners

While much has been written about attribution in the context of individual asset classes, the asset owner confronts a much more complex scenario. In the interest of maximizing return/risk payoff (often subject to external mandates, liability streams, and risk constraints), asset owners frequently diversify their total fund amongst multiple asset classes: equity, fixed income, hedged and un-hedged currency, and alternative and illiquid investments. Frequently, a proprietary investment process—tailored to the objectives and constraints of the fund mandate—is designed and employed to determine allocations to classes. The allocation processes can be complex, detailed, and highly specific, comprising multiple interim steps and decisions. Attributing the value-added to such a process is, therefore, a challenge that carries a high degree of difficulty.

The most apparent stumbling block is that of clashing methodologies. Equity, fixed income, currency, and hedging all—to some degree—can be analyzed using ‘standard’ attribution templates. But our objective is not merely to create separate, class-specific analyses; rather, we must find a way to merge these into a coherent, total-fund picture.

Even more severe incompatibilities are encountered when illiquid asset classes—private equity, real estate, infrastructure, natural resources, et alia—are added to the mix. For a variety of reasons, performance of these assets is typically measured using an internal rate of return—mathematically irreconcilable with the time-weighted method used for traditional classes. Further complications arise from the vagaries of illiquid asset valuation: lack of comparables, low frequency, and lagged valuations are difficult to integrate rationally with attributions sourced from daily public-market pricing.

Disparate classification schemes—or “drill-downs”—present another hurdle. Clearly, we want to have a drill-down in the equity sleeve—say, by sector—that is different than that of the fixed income sleeve, perhaps by rating. But can our attribution engine support this asymmetry?

Benchmark definition presents difficulties as well. Meaningful benchmarks for a MAC portfolio rarely exist in the wild. We are, therefore, tasked with constructing a bespoke benchmark, often blending a variety of single-class indices into a relevant baseline. Out-of-benchmark bets must also be considered when appropriate indices are unavailable or otherwise not included.

Finally, integrating risk into the attribution framework is mandatory. The core imperative of the MAC asset owner is to manage return/risk of variably correlated risky assets; measuring this strategic source of value is critical.

‘CUT-AND-SEW’: THE TRADITIONAL APPROACH

Our first MAC attribution method is one that’s been employed many times over many years—in fact, the basic concept will be recognizable to anyone with experience in ‘balanced’ portfolio attribution.

The cut-and-sew method begins by measuring the allocation effect derived from top-level asset class weights. Drilling down into each class, we proceed with distinct methodologies suited specifically to that class. For example, Brinson-Fachler attribution for domestic equity, Karnosky-Singer for international, and key-rate duration for fixed income. Within illiquid classes, we might even present IRR results (albeit without tying it up to the fund-level TWR).

Though vanilla versions of this approach have long been in service, more sophisticated asset owners have recently expanded beyond traditional practice, developing finer and more precise ways of tuning cut-and-sew attribution.

Attribution for Asset Owners

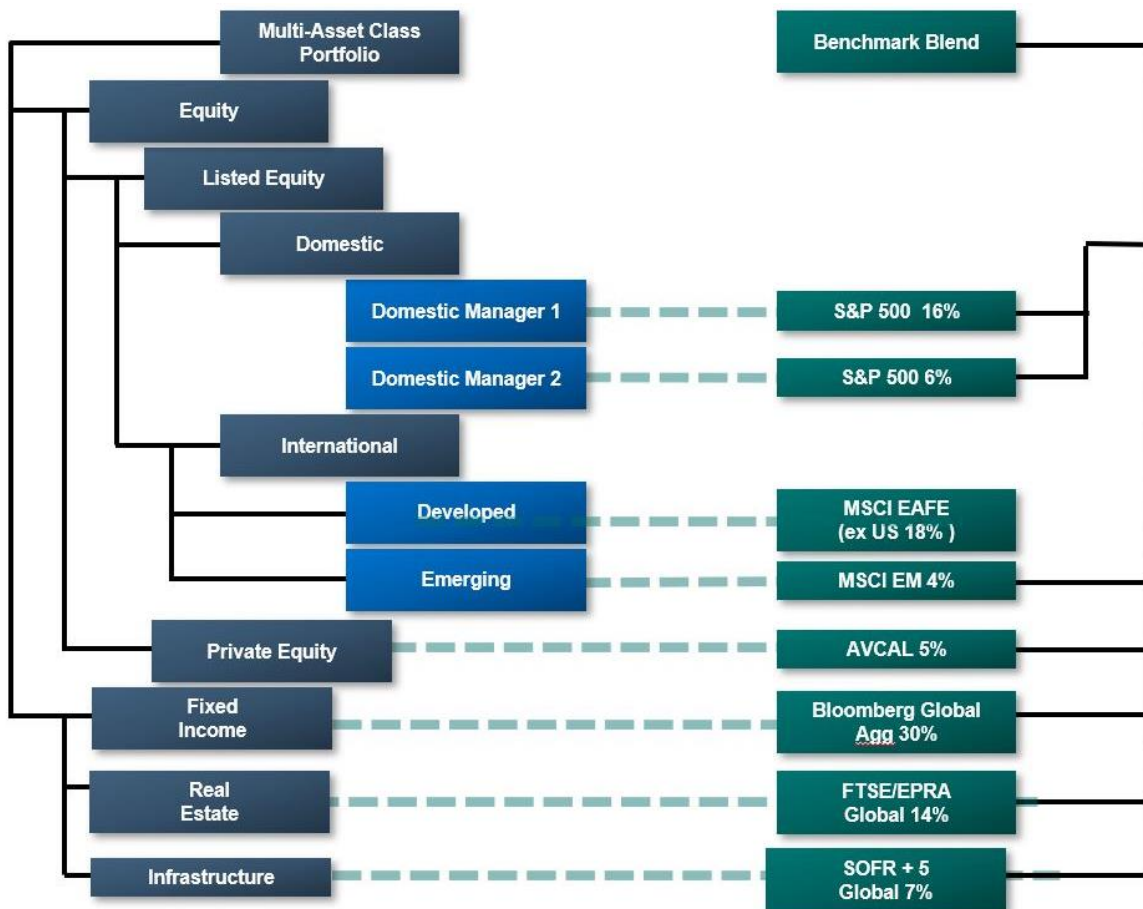
'CUT-AND-SEW': THE TRADITIONAL APPROACH (continued)

The drill-down capabilities of any MAC attribution must be sophisticated, supporting multiple asymmetric levels with methodology-specific schemes. Going even deeper, a great number of choices can be made at the class level to further attune the analysis to specific strategies:

- Specification of distinct methodology and return-decomposition models – such functionality implies distinct report columns by class
- Utilization of strategy, or decision-based attribution methodologies within the class
- Weight scaling, to sum to 100% or to the parent node weight
- Attribution effect scaling to explain the total management effect or only the parent selection effect
- Look-through, by depth or by class type

Other innovations have arisen in benchmark construction and tuning. Highly complex blends, with zero- or partially-weighted members, mapped to specific fund class nodes can be designed to isolate relevant attribution effects within the portfolio. An example of such a benchmark construct and its relationships to the portfolio might appear as the below:

Creative Benchmarking



Clearly, developing a benchmark incorporating some or all of these node-by-node features is a formidable task. Even using a tool capable of supporting such complexity, care must be taken to ensure meaningful, intentional results are produced.

(continued)

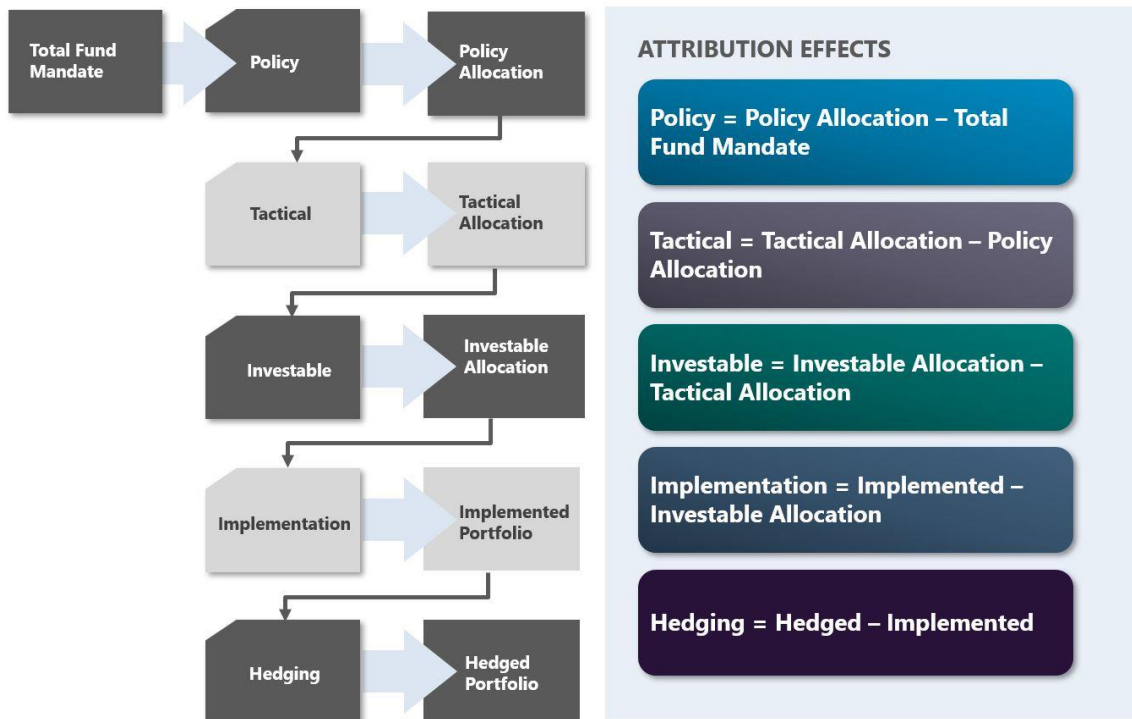
Attribution for Asset Owners

PROCESS-BASED ATTRIBUTION: A DIFFERENT TAKE

Lately, we've been working with asset owners, taking a more direct approach to MAC attribution. Process-based attribution begins by analyzing the benchmark construction process itself rather than the ex-post total-fund portfolios resulting from it. First, we break the process down into steps: distinct, progressive modifications to allocation, which together—most often, serially, though parallel steps can be accommodated—comprise the transformation of the mandate into an investible, risk-optimized policy benchmark for the total fund. For many asset owners, these process steps are already well documented internally in a formal construction policy.

Next, we measure the value added by each step as the difference between the before- and after-returns of the notional portfolio achieved by executing that step's process. An example of a typical process-based attribution might appear as the below:

Process-based Attribution Example



The process-based methodology is extremely flexible and readily incorporates hybrids with other attribution types. Each step can be further decomposed using methods including strategy-, thematic- or decision-based attribution. As an example, individual decisions driving shifts in the tactical allocation notional can be tagged—think geo-political/economic, ESG, risk balancing—and the performance effects of these measured separately in a drill-down.

CONCLUSION

Asset owners are frequently compelled by their mandate to construct and allocate multi-asset class benchmarks and funds. The design and implementation of methodologies and systems that attribute performance and risk to these presents a formidable set of challenges.

During our time spent developing successful solutions with asset owner clients, we've found that a structured, inclusive approach to designing and implementing MAC attribution solutions not only increases their change and level of success but delivers a strategically vital perspective of value to the entire enterprise. In all cases, the collaboration between Meradia's broad experience, and asset owners' comprehensive understanding of their construction process has been the key to successful achievement of their MAC attribution objectives.



Mark David, CFA

Co-leading Meradia's Performance, Risk & Analytics practice, Mark is an expert in the design and implementation of methodologies underlying complex investment analytics – and the data required to support them. He has assisted front office and operations/IT teams at dozens of the world's most complex asset managers and owners. Designing, specifying, building, and integrating performance and risk systems, and datasets to drive them—Mark consistently delivers solutions that produce accurate and insightful analysis across all asset classes – listed, derivative, hedge fund, private equity, real estate and infrastructure.