RBC Global Asset Management
PH&N Institutional

IPS INSIGHTS

A new perspective on the importance of different investment decisions

FOR INSTITUTIONAL INVESTORS ONLY

Joanna Mejza, MSc, CFA, ASA
Senior Manager
Institutional Portfolio Solutions

Ivor Krol
Managing Director,
Head, Institutional
Portfolio Solutions

*The **PH&N Institutional Portfolio Solutions (IPS)** team provides analytics, research, and thought leadership to institutional investors, helping them tackle long-term investment challenges and supporting strategic decision making.*

*In their **IPS Insights series**, the team shares pertinent findings and observations derived from practical experience.*

Executive summary

There exists a significant body of research that attempts to determine the relative importance of different investment decisions with respect to return outcomes. The onset of the investigation into this topic can be traced back to a paper originally published almost 40 years ago: *Determinants of Portfolio Performance*, written by Brinson, Hood, and Beebower in 1986. The authors' analysis in that seminal paper demonstrated that long-term asset allocation drives 90 percent of a portfolio's performance, leading to a generalized – and quite prevalent – hypothesis: investment policy ***asset mix decisions*** are far more important than investment ***fund implementation decisions***.

In the ensuing decades, this hypothesis has been the subject of much discussion among investment professionals. Some have fully accepted it, while others have severely critiqued it; criticisms have ranged from arguing that the original analysis missed important nuances, to stating that it was outright flawed. The ongoing debate even induced one of the authors of the original paper to wade back into the discussion decades later, stating, "We would not have guessed that a six-page article would be the focal point of a 20-year discussion."[1]

There is no doubt that the subject is both contentious, and one that is fundamentally important to institutional investors. At the same time, the complexity and variety of the investment decisions faced by institutions have expanded significantly since much of the original research was done. In light of these two considerations, this article aims to take a fresh stab at the topic. Starting with a review of the main points and findings that have been put forth on the subject over the decades, we conclude that the biggest point of disagreement is the way in which asset mix decisions are isolated and measured on a relative basis. Despite multiple proposals on how to define an appropriate baseline, we argue that a crucial consideration has been overlooked: the uniqueness inherent to institutional investors' individual circumstances. To address this, we devise a framework that introduces the concept of a reference portfolio that can serve as a customized baseline to measure and guide asset allocations decisions, allowing us to reinvestigate the original question. Then, within our proposed framework, we use historical benchmark and fund manager returns to represent a broad selection of fixed income, equity, and alternative investments for a modern-day investor, test different objective-oriented cases, and compare our results to those previously found.

[1] Randolph L. Hood, 2005.

Our primary finding is that the relative weight of one type of decision's impact on risk/return outcomes can vary significantly depending on the following factors:

- The investor's return objectives and/or risk tolerance
- The breadth of the chosen asset class opportunity set
- The alpha potential and tracking error profile of active management (which varies by asset class)

Consequently, despite extensive past efforts and our own more recent attempt to conclusively establish which investment decision matters more – asset mix or fund implementation – we find there is no universal result that can be relied upon given the large amount of variables at play. Therefore, we would argue that the default mindset for institutional investors should be to ascribe ***equal weight*** to both when making ex-ante investment decisions, even if one will likely end up being more consequential than the other on ex-post outcomes.

WHY THIS MATTERS

An institutional investor – whether it be a pension plan, an endowment fund, or some sort of collective trust – is almost always a multi-stakeholder entity with no one true owner of the assets. Investment decisions related to a defined benefit pension plan, for example, will typically be undertaken by a committee (and potentially sub-committees) made up of diverse individuals, some of whom may not even have a stake in the outcomes. Therefore, the responsible management of an institutional pool of assets requires a governance framework for assigning responsibilities and evaluating the impact of each decision within the investment process.[2]

A good starting point for such a program is to formally define investment objectives, including an articulation of the investment beliefs underlying them. From there, the natural next steps are to decide on asset mix and fund implementation, ideally supported by a systematic process. Finally, creating a mechanism for monitoring and reviewing these decisions serves to complete the cycle – a cornerstone of fiduciary accountability that promotes transparency with asset beneficiaries. Since we have established that asset mix and fund implementation decisions can both materially impact outcomes, the implication from a governance standpoint is that they should receive equal attention when being determined, monitored, and reviewed.

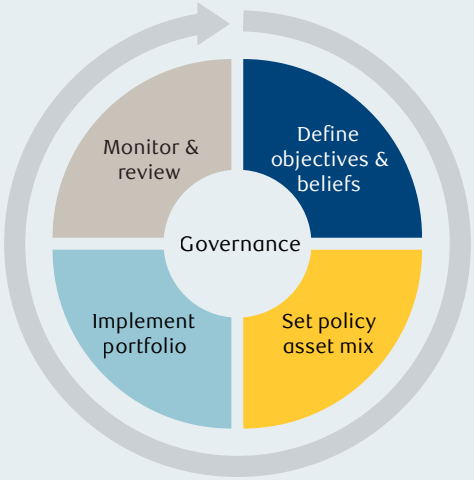

In practice, ***fund implementation decisions*** tend to be under constant surveillance due to the straightforward manner in which the performance of active funds is reported against their benchmarks. ***Asset mix decisions***, on the other hand, might not receive the same level of scrutiny given the challenge of defining a meaningful benchmark for their evaluation. However, the method we employ in our analysis – defining asset mix decisions as deviations from a customized reference portfolio – provides institutional investors with a template to explicitly tackle this challenge. Consequently, both types of investment decisions can be effectively monitored with equal opportunity to identify issues, implement corrective actions, and foster continuous improvements as part of a robust governance framework.

[2] Walter R. Good, 1984.

Introduction

This article revisits the topic of the relative importance of investment decisions that has been the subject of several prominent papers with very different conclusions. We will outline and comment on the main points that have been proffered on the subject over the decades, consider the merits (and potential issues) on both sides of this debate, and ultimately propose a closely related, but slightly different framework for our own attempt at investigating this question. Ultimately, we hope that a different perspective will lead to useful insights that could have practical applications for institutional investors navigating the governance challenges associated with their investment programs.

Historical background

Many investors and investment professionals alike may have heard the statement that long-term asset allocation is the single most important investment decision. In fact, it is often stated that 90 percent of a portfolio's performance is attributable to asset allocation, but what is the basis for this belief?

The genesis of the claim is a 1986 paper by Brinson, Hood, and Beebower. The authors' motivation for their study was the practical observation that many of their institutional pension clients tended to spend more time and effort on manager selection than on asset allocation policy, possibly because quantifying the contribution of the latter to investment returns was challenging. The authors' intention was therefore to devise a way to measure this effect and compare its impact and importance relative to that of the other types of investment decisions. Their finding was that "investment policy return [...] explained on average fully 93.6 percent of the total variation in actual plan return." And thus was born the statistic that is often quoted when discussing investment decision making.

We note that the original analysis involved certain assumptions and limitations that can be reasonably attributed to the era in which it was conducted. For example, the authors proxied the normal policy asset mixes using a combination of stocks, bonds, and cash, and excluded all other assets due to the limited availability of historical data.[3] However, several criticisms have emerged on more fundamental elements of their methodology that have stimulated a multi-decade debate on the subject, including a follow-up by one of the original authors 20 years later in which he addressed some of the most prevalent observations and reactions of both proponents and detractors. Below, we summarize what we consider to be the three main critiques that have emerged over the years, as well as a comment on each.

1. Use of ***variation*** in total portfolio returns

Their primary conclusion was based on how much of the portfolio return's variability could be explained by the different types of investment decisions, not the level of the realized returns themselves.

This detail is sometimes forgotten by those who support the original findings when invoking the statistic: asset allocation was found to explain 90% of a portfolio's *variance*, not its realized *return*. Some critics argue that investors are not primarily concerned with the variability of returns over time (Jahnke, 1997). Rather, they care about the range of possible outcomes at the end of their investment horizon, and focusing on variance could diminish the importance of security selection and market timing (i.e., active management).

It is true that not incorporating a return dimension into the analytical framework is to omit a critical piece of information. However, ignoring the risk side of the equation is equally problematic, especially since investors will inevitably have some level of risk tolerance for portfolio fluctuations. The variability of returns will dictate how smooth (or bumpy) the path will be, and this consideration is no less important than the long-term level.

[3] The other segment represented an average allocation of 9% and included "convertible securities, international holdings, real estate, venture capital, insurance contracts, mortgage-backed bonds, and private placements."

2. No explicit ***baseline*** for measuring asset allocation

The relative basis for measuring the return variance was a portfolio with zero return (i.e., an uninvested portfolio).

Some have pointed out that the 90% value might be overstated because the asset allocation decision had not been correctly isolated. They suggested that investors do not select a policy asset mix in lieu of holding cash – which is what the original study implicitly assumed – but in lieu of some type of "naïve alternative" portfolio. Any deviation from this portfolio would therefore represent the plan's specific policy decisions (Hensel, Ezra, and Ilkiw, 1991). In the literature, the alternative portfolio has been defined in several ways, such as the average asset allocation of a group of similar investors, or a broad equity market index. If asset allocation returns are redefined as only the excess over this baseline, then they would explain far less of total variability and contribute about the same as active management.

Although we agree with the points raised regarding the need for a deliberate baseline, in our opinion, the notion of a universal alternative portfolio against which any institutional investor can measure its asset allocation decisions is not consistent with the fundamental investment problem. It is not possible that a single portfolio reference could apply to all institutional investors unless all have the same return objectives, risk tolerances, time horizons, and constraints. For example, a generic market portfolio like a broad equity index completely misses any investment trade-off relative to cash flow obligations, such as pension benefit payments or endowment spending commitments.

3. Analysis considers each plan on an ***individual*** basis

The analysis only examines the return variation of a single fund and not the return variation between different funds.

Some have mistakenly interpreted the 90% figure to represent the impact of an investor's asset allocation decision relative to those of other investors (Nuttall, 2000), suggesting that context was of greater interest. This, however, represents a different framing of the asset allocation question, and yields a different answer. For instance, it was shown that only 40% of the variation of return among funds was explained by policy allocation, with the rest coming from active management (Ibbotson and Kaplan, 2010). Furthermore, strategies with higher degrees of active management (e.g., market-neutral vs. long-only stock picking) were also shown to further increase the relative contribution from alpha (Xiong, Ibbotson, Idzorek, and Chen, 2010). Therefore, in this context, the implication is that active management is far more consequential than previously thought.

Different institutional investors will make different asset allocation decisions, usually related to their distinct circumstances. For example, not all institutions have the same risk tolerance: a closed and mature pension plan in de-risking mode will not typically have the same exposure to return-seeking assets as one that is open with a young, active population. Drawing a comparison between the asset allocations of those two institutions would not provide meaningful insights because their risk levels are fundamentally different by design. Therefore, we believe that a comparison between investors provides a questionable basis for evaluating the relative importance of different investment decisions.

Our perspective

Both the original publication and all of the subsequent papers and observations have contributed valuable arguments and insights on this topic. However, what is an investor to make of the wide disparity in numbers found by the different researchers – from the original 90% to something significantly lower? In our opinion, we don't think any of the past findings can be overly relied upon because there is one underlying issue inherent to all of them: asset allocation decisions were measured in relation to baselines that had no connection to individual investor circumstances (e.g., cash, equity markets, or peers). Since it is highly unlikely that any two institutions will have the same set of objectives and beliefs, we should not expect there to exist a universal reference point for evaluating individual decisions. Finding an answer to the investment decision question therefore requires a meaningful baseline for explicitly measuring the asset mix component. This brings us to the fundamental purpose of the assets.

All institutions invariably have a set of objectives for their invested assets. For example, the goal of a defined benefit pension fund might be balancing member benefit security with employer contributions, whereas in the case of an endowment fund, it might be meeting current spending commitments while maintaining intergenerational equity. The way in which their respective assets are ultimately invested should, by definition, be supportive of these objectives from both an outcome and a risk-management standpoint. The starting point is normally to establish a desired long-term return target and an acceptable level of risk (e.g., volatility) unique to the investor's specific circumstances – their risk/return profile. The next step is often to decide on the policy asset mix intended to support the chosen profile. However, embedded within that process is the concept of a ***reference portfolio***.

In our experience, the use of a reference portfolio is more prevalent among very large institutional investors with more complex and delineated layers of governance and fund management; some notable examples include the New Zealand Superannuation Fund (NZSF), the University of Toronto Endowment Fund, and the Public Service Pension Plan. However, that shouldn't necessarily imply that reference portfolios don't have useful applications for institutional investors of all sizes. As NZSF states, "A reference portfolio approach is first and foremost a governance construct designed to facilitate clear decision making and accountability of decisions."[4] This suggests that the reference portfolio would be an ideal candidate for the "naïve alternative" portfolio described in the literature, and therefore an appropriate baseline for evaluating asset mix decisions.

While an investor can choose to set their asset mix to their implied reference portfolio, it is likely that they will want to optimize the risk/return trade-off. This process can vary significantly in terms of scope but generally involves incorporating additional asset classes with different risk, return, liquidity, and complexity profiles that would constitute deliberate decisions. That said, we want to emphasize that the process of articulating risk/return objectives and mapping them into a reference portfolio is a critically important exercise, as all investment decisions would inherently be functions of this initial step.

REFERENCE PORTFOLIO

A reference portfolio (RP) is typically comprised of a generic equity and fixed income allocation. It is meant to represent the easiest, cheapest, and most liquid portfolio consistent with the long-term expected return and/or risk aligned with an institution's investment objectives. The original concept is attributed to the Canadian Pension Plan (CPP), which introduced the RP as a way to create more delineation with respect to governance and decision making between the investment management team and the Board. The RP for base CPP was chosen to be 85% global equities and 15% Canadian government bonds, and it reflects the level of risk the Board is comfortable taking. The actual investment policy under the purview of the investment management team – a complex portfolio of fixed income, equities, derivatives, and private market alternatives – is expected to achieve better results than the RP on a risk-adjusted basis.

[4] New Zealand Superannuation Fund, *How We Invest White Paper - 2015 Reference Portfolio Review*, August 2, 2015.

Our proposed framework

Using the reference portfolio as a tool for measuring asset mix decisions, we propose a slightly modified framework comprised of three different portfolios, as illustrated below:

Figure 1: A three-layered approach

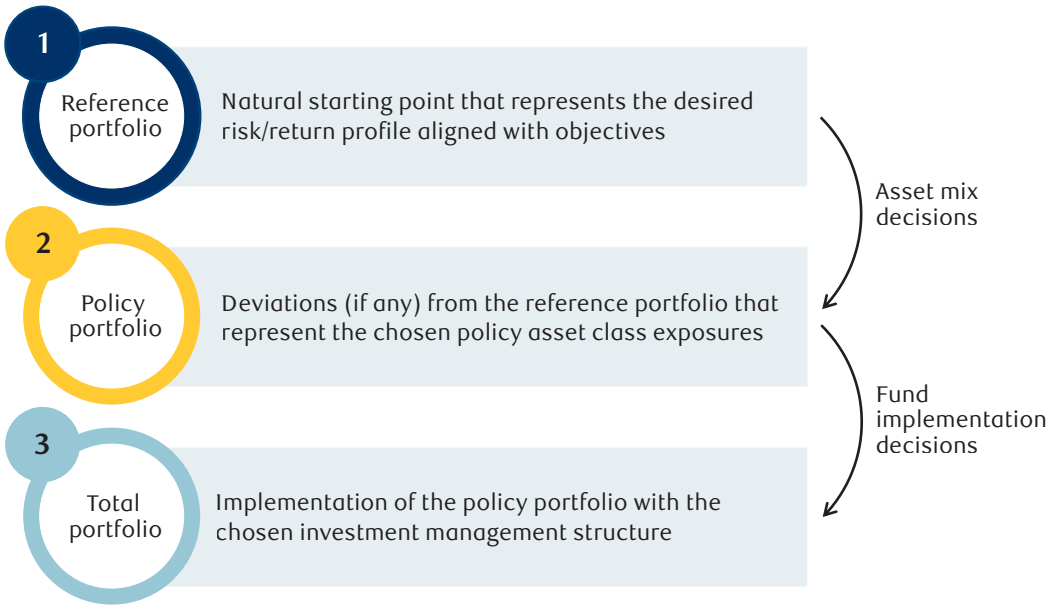

To better understand the structure of our proposed three-layered framework, we describe the basic edge cases for any given reference portfolio:

- Investor adopts the reference portfolio as the policy asset mix and then implements everything passively. In this situation, ***no investment decision has been made*** according to our framework.
- Investor creates a policy asset mix that is different from the reference portfolio and then implements everything passively. In this situation, even for minor deviations, ***100% of decision-related outcomes are attributable to asset mix***.
- Investor adopts the reference portfolio as the policy asset mix and then implements some or all components actively. In this situation, ***100% of decision-related outcomes are attributable to fund implementation***.

Considering our framework, we see that the weight of a particular decision can easily vary from one extreme to the other. This observation already reveals an important and original take on the question of which decision matters more: there is no universal answer that can constitute a general rule of thumb. That said, this does not mean that there is no further insight to be gained from an analytical investigation. As most investors will likely deviate from their reference portfolio and implement some or all of it actively, we may yet identify some tendencies associated with the decision types that would be of practical value to know. To investigate this more common case, we need to create a test framework that is as representative as possible for the current environment; the investment-decision landscape and its complexity have evolved significantly since both the original work and many of the follow-up papers on the subject.

Data

A significant hurdle to performing this type of analysis is the quality and availability of sufficient return data for both a representative cross-section of asset classes and a universe of fund managers that would be viable implementation options. For our purposes, we have sourced the required data from the eVestment database.

Our analysis is based on 10 years of monthly historical data using three broad categories of asset classes that vary by their degree of complexity and liquidity: traditional, specialty, and alternative. For each individual asset class, we have identified benchmarks that serve as both a representation of the generic risk exposure from an asset mix decision-making standpoint and a basis of comparison for relative manager performance in fund implementation. The details are provided in the following table:

Traditional	Specialty	Alternative
Investment grade bonds FTSE Canada Universe Bond Index	**High yield bonds** ICE BofA Global High Yield Index (CAD)	**Liquid** HFRI Macro (Total) Index (CAD)
Large-cap equities S&P/TSX Capped Composite Index MSCI World Index (CAD)	**Emerging market debt** J.P. Morgan Emerging Market Bond Index (CAD)	**Private market** Credit Suisse Leveraged Loan Index (CAD)
	Small-cap equities MSCI World Small Cap Index (CAD)	
	Low volatility equities MSCI World Minimum Volatility Index (CAD)	
	Emerging market equities MSCI Emerging Markets (EM) Index (CAD)	

We note that the list is by no means exhaustive, and some readers might observe notable omissions relative to their actual investment policies. The reason for this is a lack of fund return data for those particular asset classes in terms of sufficient length and sample size. That said, we believe this opportunity set is a realistic representation for our purposes of analyzing the impact of asset mix decisions, given the data limitations when trying to create an equally realistic representation of the associated fund implementation decisions.

For the purposes of creating a sufficiently deep universe of funds that should reasonably represent an institutional investor's manager selection pool, we also felt the need to apply some filters, especially since fund manager data from eVestment is self-reported and therefore subject to potential misclassifications and other errors. The results of the analysis will invariably be influenced by the constructed fund universes, which means that extreme outliers or misclassified strategies have the potential to distort the findings. The filters and their rationale are described as follows:

- A fund must have 10 years of return data and be benchmarked against the index used as a policy representation of the asset class in our analysis.[5]
- We applied a filter on assets under management (AUM) to remove funds below a threshold of 20% of the sample average AUM. This ensures that we exclude funds that would likely be too small to qualify as admissible for institutional investors. Furthermore, filtering the funds on AUM helps us avoid a potential distortion resulting from returns that might be a function of their smaller size and not achievable at an institutional scale.
- We excluded funds with excessive volatility (defined as more than twice the sample average volatility) under the assumption that the typical institutional investor would not consider a fund to be an admissible implementation option for a given policy exposure if it exhibited multiple orders of higher risk.

The resulting dataset is characterized by the 10-year quartile distributions of manager returns illustrated in *Figure 2*.[6]

[5] Because we have required a 10-year track record for all funds so that we have a sufficient time series of returns, there is a potential introduction of a survivorship bias which could favourably skew the manager return distributions.
[6] We do not consider investment management fees as they can vary significantly by asset class, fund, and investor AUM, making it very difficult to come up with a single generic estimate.

Figure 2: Quartile distribution of 10-year annualized manager returns

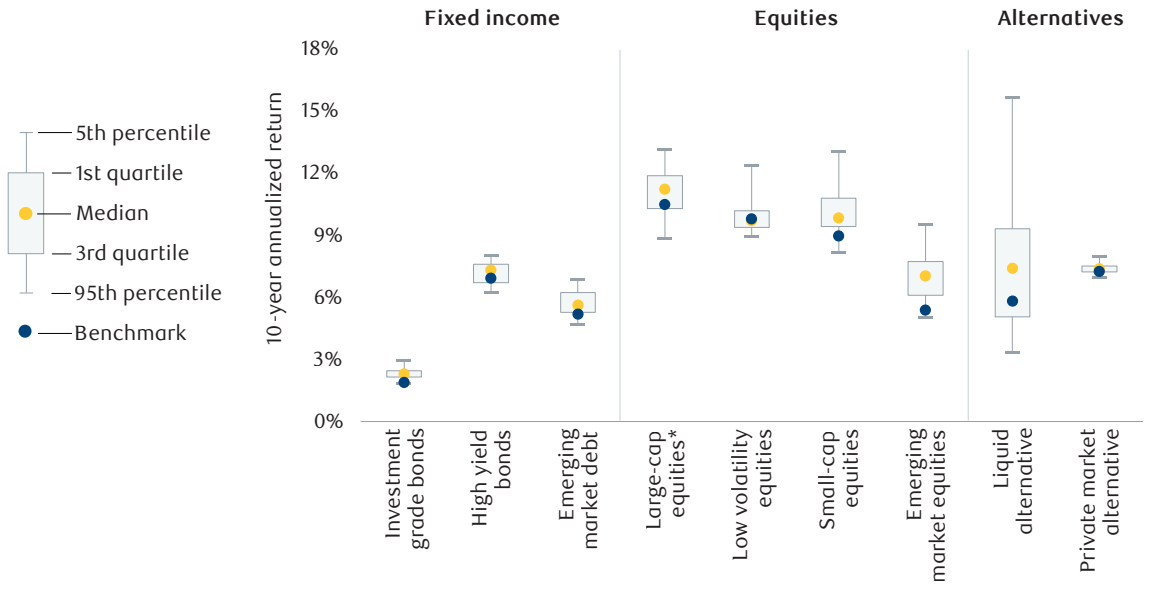

Average alpha	0.4%	0.3%	0.5%	0.6%	0.2%	1.3%	1.6%	1.9%	0.2%
Average tracking error	0.7%	1.7%	2.5%	3.3%	4.0%	4.4%	4.5%	9.2%	1.0%

*Split between Canadian (30%) and global equities (70%).
Source: RBC GAM, eVestment. 10 years ending June 30, 2024.

Methodology

A comprehensive analysis requires us to consider a variety of circumstances that broadly represent different investor situations. Recall that our main argument is that investors will have different return objectives and/or risk tolerances from one another, as well as different beliefs; each of these factors can influence the choice of admissible asset classes. We will therefore evaluate whether risk/return profile and/or breadth of opportunity set has any impact on asset mix and fund implementation decisions in the following ways:

- We will consider the inclusion of specialty asset classes for three different reference portfolios comprised of the traditional asset classes:
  - one for a low-risk investor (20% large-cap equities),
  - one for a balanced-risk investor (50% large-cap equities),
  - one for a high-risk investor (80% large-cap equities).
- We will consider the inclusion of alternatives in a second opportunity set for the balanced-risk investor only.

We now need a basis to establish illustrative policy portfolios for each reference portfolio. While there are many ways that an investor could deviate from their chosen reference portfolio in practice, we constructed the policy portfolios by maximizing the return for the same level of risk over the chosen historical period. We believe this is a reasonable assumption, given that practical applications for reference portfolios have typically been to serve as risk-equivalent performance benchmarks for policy portfolios. The resulting asset mixes (which have been constrained to avoid concentration risk in non-traditional classes)[7] are depicted in *Figure 3*.

[7] Refer to appendix for more details.

Figure 3: Policy portfolios

	Risk/return profile →			Breadth of opportunity set →	
	+ Specialty			+ Specialty	+ Alternative
	Low risk	Balanced risk	High risk	Balanced risk	Balanced risk
Reference portfolio	**20/80**	**50/50**	**80/20**	**50/50**	**50/50**
Fixed income	**83%**	**43%**	**10%**	**43%**	**30%**
Investment grade bonds	59%	26%	6%	26%	18%
High yield bonds	12%	9%	2%	9%	6%
Emerging market debt	12%	9%	2%	9%	6%
Equities	**17%**	**57%**	**90%**	**57%**	**56%**
Large-cap equities*	8%	29%	50%	29%	42%
Low volatility equities	8%	24%	28%	24%	12%
Small-cap equities	–	–	10%	–	2%
Emerging market equities	–	4%	2%	4%	–
Alternatives	–	–	–	–	**14%**
Liquid alternative	–	–	–	–	7%
Private market alternative	–	–	–	–	7%

*Split between Canadian (30%) and global equities (70%).

We also need a basis to construct testable arrays of different manager implementation options for each policy asset class. Ideally, we would like to test all possible fund combinations for a given policy asset mix across the entire 10-year period (assuming no manager turnover or asset mix changes),[8] as this would provide us with the maximum range of potential variation that could theoretically be associated with fund implementation decision making. However, this would produce an unmanageable number of combinations when considering all the asset classes listed above: over ten quadrillion! To reduce the dimensionality of the problem, we employ a stratified sampling technique based on the fund return quartiles:

1. For each asset class, funds are divided into groups based on the quartiles of their 10-year annualized returns.
2. All possible quartile combinations for all ten asset classes (with the large-cap equity split between Canadian and global) are then created – 1,048,576 in total when considering the full opportunity set. For example, at the extremes, we can have an implementation scenario where all asset classes have 1st quartile managers or 4th quartile managers.
3. For each quartile combination, one manager is selected at random from the specified quartile group for each asset class. The process is then repeated 50 times to ensure that the sample population is a good representation of the entire population.

Finally, having defined all components of our three-layered framework, the methodology for measurement purposes is depicted in *Equation 1*, on the following page. It involves decomposing the return deviations of the total portfolio from the reference portfolio (which captures the combined effect of asset mix and fund implementation decisions) into individual components.

[8] A necessary assumption for the purposes of the analysis but in practice, an ongoing governance and monitoring program would likely lead to periodic changes in both asset mix policy and investment manager structure, as well as potential tactical asset mix deviations.

Equation 1

$$\underbrace{(TP_{i,t} - RP_t)}_{\text{Total investment decisions}} = \underbrace{(PP_t - RP_t)}_{\text{Asset mix decisions}} + \underbrace{(TP_{i,t} - PP_t)}_{\text{Fund implementation decisions}}$$

where (assuming monthly rebalancing)
$TP_{i,t}$ = Total portfolio return for fund implementation *i* in month *t*
RP_t = Reference portfolio return in month *t*
PP_t = Policy portfolio return in month *t*

This formulation is similar to the one used by Xiong, Ibbotson, Idzorek, and Chen (2010) except that we introduce the investor-specific RP in lieu of a universal alternative portfolio. Furthermore, we do not look at the absolute return of TP, but rather its excess over RP, which we do not consider to be an explicit investment decision.

Equation 1 allows us to appropriately isolate the investment decisions, however, we must determine how to measure their effects. As previously mentioned, both return level and variability should be jointly considered; the former represents the return objective, and the latter represents the risk tolerance, the complete trade-off from a decision-making standpoint. While the impact on return level is straightforward, the impact on variability requires further consideration.

In the original work by Brinson, Hood, and Beebower (1986), the authors' approach consisted of performing a series of linear regressions, starting with the asset mix decision and then adding in the implementation decisions. This allowed them to quantify (using the R^2 statistic) how much *more* of the variance was explained by the inclusion of implementation decisions as their way of assessing relative importance. However, this methodology does not explicitly quantify the individual contribution of each component, which is our objective.

In our framework, one potential way to achieve this would be to calculate the individual R^2s of each component. However, the resulting values would not naturally sum to 100% without accounting for the interplay between them. This was addressed by Xiong, Ibbotson, Idzorek, and Chen (2010), who introduced a balancing term labelled the "interaction effect," but this term was not directly associated with any given component. Instead of working with the R^2 of components, we believe a more straightforward approach would be to work directly with their marginal contribution to risk (MCTR) where we can explicitly consider correlation and attribute it analytically.[9]

Results

Impact on return level

We begin our analysis by considering the level of realized returns. *Figure 4* illustrates the measured breakdown between the asset mix and fund implementation decisions' contributions to total returns over the 10-year period on both an absolute and relative basis.

Figure 4: Breakdown of 10-year annualized return by investment decision

	+ Specialty			+ Specialty	+ Alternative
	Low risk	Balanced risk	High risk	Balanced risk	Balanced risk
Portfolio returns					
Reference (RP)	3.6%	6.2%	8.8%	6.2%	6.2%
Policy (PP)	4.3%	7.3%	9.5%	7.3%	7.9%
Total (TP)	4.7%	7.7%	9.9%	7.7%	8.4%
Absolute contribution					
Asset mix (R_{PP}- R_{RP})	0.7%	1.0%	0.7%	1.0%	1.6%
Fund implementation (R_{TP}- R_{PP})	0.4%	0.4%	0.5%	0.4%	0.5%
Total (R_{TP}- R_{RP})	1.1%	1.4%	1.1%	1.4%	2.2%
Relative contribution					
% Asset mix	**64%**	**71%**	**58%**	**71%**	**75%**
% Fund implementation	**36%**	**29%**	**42%**	**29%**	**25%**

Totals might not sum exactly due to rounding.

[9] Refer to the appendix for the detailed formula.

With respect to the asset mix decision component, the absolute contribution to return varies depending on the reference portfolio: it is highest for the balanced-risk profile, which benefitted more from policy-level asset class decisions than the low-risk or high-risk profiles did (1.0% versus 0.7%). We attribute this observation to the likelihood that in the extremes of the risk/return spectrum, there is generally less room for an investor to make return deviations from their reference portfolio without fundamentally changing the risk profile from their stated tolerance. When we allow for a broader opportunity set with the addition of alternatives, the contribution increases further (1.6% versus 1.0%) because of the ability to deviate even more from the reference portfolio in a positive way.

With respect to the fund implementation decision component, the absolute contribution to return once again varies depending on the reference portfolio: it is greatest for the high-risk profile (0.5% versus 0.4%). This is attributable to its policy portfolio's greater emphasis on equity asset classes, especially small-cap and emerging markets, which exhibit higher average alpha in our historical sample (see *Figure 2*). We observe a similar occurrence when adding alternatives to the mix for the balanced-risk profile: the contribution increases (0.5% versus 0.4%) because the new policy portfolio is further tilted towards higher alpha asset classes.

Lastly, the resulting relative contributions by investment decision illustrate which one carried more weight. We observe values materially below the 90% figure for asset mix decisions, with significant variation between cases; this indicates that both risk/return profile and breadth of opportunity set influence the relative importance of the different decisions on return outcomes for the aforementioned reasons. That said, the trends are a bit harder to discern because the relative values are interdependent, meaning changes in one automatically affect the other.

Impact on return variability

Next, we consider return variability as represented by the standard deviation of returns (i.e., volatility). *Figure 5* illustrates the measured breakdown between the asset mix and fund implementation decisions' contribution to the total volatility associated with investment decisions ($\sigma_{TP\text{-}RP}$).

Figure 5: Breakdown of annual volatility by investment decision

	+ Specialty			+ Specialty	+ Alternative
	Low risk	Balanced risk	High risk	Balanced risk	Balanced risk
Portfolio volatility					
Reference (RP)	5.8%	7.5%	9.9%	7.5%	7.5%
Policy (PP)	5.8%	7.5%	9.9%	7.5%	7.5%
Total (TP)	5.8%	7.7%	10.1%	7.7%	7.7%
Absolute contribution					
Asset mix ($MCTR_{PP\text{-}RP}$)	1.5%	1.7%	1.2%	1.7%	1.7%
Fund implementation ($MCTR_{TP\text{-}PP}$)	0.3%	0.7%	1.5%	0.7%	1.1%
Total ($\sigma_{TP\text{-}RP}$)	1.8%	2.4%	2.7%	2.4%	2.8%
Relative contribution					
% Asset mix	**84%**	**70%**	**44%**	**70%**	**62%**
% Fund implementation	**16%**	**30%**	**56%**	**30%**	**38%**

Totals might not sum exactly due to rounding.

With respect to the asset mix decision component, the marginal contribution to risk (MCTR) also varies by reference portfolio: it is highest for the balanced-risk profile (1.7%) and lowest for the high-risk profile (1.2%). Given that the reference and policy portfolios have the same volatility, the observed differences are related to the degree of overlap in broad risk exposures between them, with more similarity resulting in lower MCTRs. An example of an overlap is when different asset classes share a strong exposure to a common risk factor despite having distinct risk/return trade-offs. This is the case in the high-risk profile, where both the reference and policy portfolios are dominated by equity market risk even though the latter introduces a variety of different flavours of equities, hence the lowest MCTR. One might therefore expect the inclusion of alternatives (which tend to add diversification) to increase the MCTR, however this is not what we observe for the balanced risk profile. The alternatives free up risk capacity, allowing returns to be pushed higher by increasing the allocation to global equities – the highest-returning asset class in sample. This, by extension, brings the new policy portfolio closer to the reference portfolio, and ultimately results in a similar MCTR as the policy portfolio without alternatives.

With respect to the fund implementation decision component, we observe a notable increase in the MCTR when moving from the low-risk (0.3%) to the high-risk profile (1.5%). Recall that the high-risk policy portfolio is tilted towards equity asset classes that, in addition to exhibiting higher realized alpha, also have more material volatility versus their benchmarks. Therefore, the MCTR is higher when the policy portfolio is implemented with higher tracking error strategies. This effect likewise explains the higher MCTR that accompanies the addition of alternatives for the balanced-risk profile (1.1% versus 0.7%).

Finally, like with returns, the relative contribution of each investment decision to portfolio volatility can vary widely depending on the risk/return profile and the universe of asset classes. For example, the impact of asset mix decisions ranges from very high levels (84%) to less than half (44%). We note that the magnitude of this impact is different from a volatility perspective than from a return perspective, highlighting the value of considering both bases.

Comments and additional considerations

Our analysis was intended to identify potential trends in, or relationships between, the two types of investment decisions, rather than to produce precise values that could be cited as a rule. Recall from the edge cases that a hard number upon which every investor can rely to rank decision importance does not actually exist; this is confirmed by the results. In fact, which decision has more impact changes based on each investor's unique combination of circumstances that define their risk/return objectives and opportunity set. Furthermore, we would expect all of these specific values to vary significantly if initial inputs were changed; for example, if we were to work with a different fund universe, select a different time period, or work with net-of-fee returns. That said, we did observe two general tendencies across the more common cases that we consider to be reliable inferences:

1. *The impact of asset mix decisions can be more consequential when an investor has increased ability to deviate from their reference portfolio (in terms of return and risk factor exposures) while respecting their risk/return objectives.*
2. *The impact of fund implementation decisions can be more consequential when an investor allocates to asset classes that have higher alpha potential and/or higher tracking error.*

It bears repeating that both decisions will ultimately interact with each other, meaning that exact relative outcomes will vary by situation.

A final point of interest worth mentioning is that there is an additional layer of variability that can greatly affect the impact of fund implementation decisions: the quartile in which performance falls and the related tracking error. In our analysis, we looked at averages, but the implementation experience for a given policy portfolio could vary based on the risk/return profile of the alpha and the dispersion between managers in different asset classes. *Figure 6* illustrates the average alpha ($R_{TP} - R_{PP}$) by quartile ranking as well as the associated average tracking error (σ_{Tp-PP}) for the balanced risk profile with alternatives.

Figure 6: Alpha vs. tracking error by quartile ranking

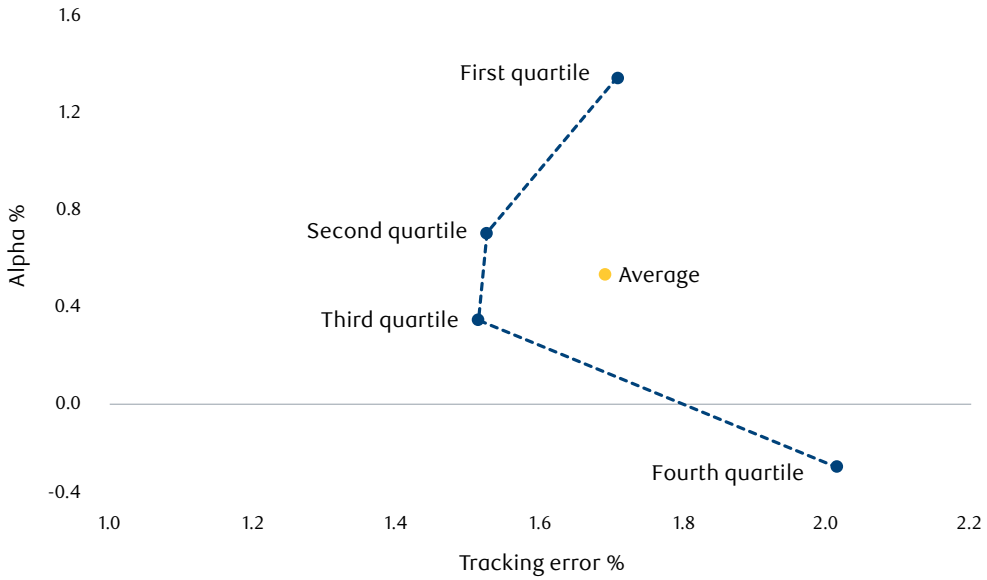

Focusing on the extremes – that is, the first and fourth quartiles – we see the extent to which the alpha's return and tracking error diverge significantly from the average. With respect to the return, if the outcome was a first quartile performance (1.3% instead of 0.5%), then fund implementation decisions would go from a contribution of 25% to a more substantial 45%. With respect to tracking error, if the outcome was in the fourth quartile (2.0% instead of 1.7%), then fund implementation decisions would explain 46% of the volatility as opposed to 38%. Therefore, the greater the dispersion between quartiles in terms of alpha or tracking error, the more consequential fund implementation decisions can become.

Conclusion

The question of which type of investment decision carries more weight on outcomes for institutional investors has been the subject of much research and debate. On the one hand, there is the original finding from Brinson, Hood, and Beebower (BHB) that policy decisions drive 90% of a portfolio's performance, and on the other hand, a bevy of articles that question elements of their work and come to very different conclusions. In reviewing a significant body of the relevant research, we believe there to be good points and overlooked aspects on both sides of the debate. BHB's desire to quantify the contributory effects of asset mix decisions on total portfolio performance is commendable, especially when it was not widely considered by others at the time. However, their methodology does not use a baseline for measuring these decisions. Subsequent work attempts to address this omission, but introduces different issues, such as establishing an explicit baseline, but not one that could be considered generally relevant to all institutional investors. We therefore sought to devise a framework that would address this point and believe that our proposed three-layered approach integrates BHB's original idea and addresses its critics. Namely, we introduce the reference portfolio as an objective baseline for any institutional investor given their unique objectives and situation.

With our framework defined, we set out to answer the same question as our predecessors: What is the relative importance of asset mix versus fund implementation decisions on investment outcomes? Our objective was to determine whether we could analytically derive the "right" number, and our conclusion is that no such number exists. There are too many dimensions, and too much variation within those dimensions, for a definitive rule of thumb that could be universally applied to all investor situations. As we saw in our analysis, it is very difficult to properly isolate the breakdown between individual decisions because of the interplay between them. For instance, we found that the more an investor can deviate from their reference portfolio through larger and more complex opportunity sets, the more impact asset mix decisions can have. But in that same scenario, there will most likely be an increase in the alpha and tracking

error potential of the implementation decisions associated with the larger opportunity set which would in turn increase their potential impact, particularly if the risk-adjusted return is especially good (Q1) or especially poor (Q4).

Therefore, the conclusion of our analysis is that both types of decisions should be treated as ***equally important*** and receive the same amount of attention, as well as scrutiny. Instead of trying to rank each investment decision in some order of importance, institutional investors would be better served by understanding the factors that can influence their individual impact and measuring them as part of a formal system of monitoring and review. This way, adjustments can be made to both asset mix and fund implementation as needed to continually support positive investment outcomes. The three-layered framework underlying our analysis can be used to accomplish this in practice, likely leading to enhanced governance and accountability for decision making that was arguably the motivation behind BHB's original paper.

References (by publication date)

Walter R. Good. 1984. "Accountability for Pension Fund Performance," *Financial Analysts Journal*, vol. 40, no. 1 (January/February): 39-42.

Gary P. Brinson, Randolph L. Hood, and Gilbert L. Beebower. 1986. "Determinants of Portfolio Performance," *Financial Analysts Journal*, vol. 42, no. 4 (July/August): 39–44.

Gary P. Brinson, Brian D. Singer, and Gilbert L. Beebower. 1991. "Determinants of Portfolio Performance II: An Update," *Financial Analysts Journal*, vol. 47, no. 3, (May/June): 40–48.

Chris R. Hensel, D. Don Ezra, and John H. Ilkiw. 1991. "The Importance of the Asset Allocation Decision," *Financial Analysts Journal*, vol. 47, no. 4 (July/August): 65-72.

William W. Jahnke. 1997. "The Asset Allocation Hoax," *Journal of Financial Planning*, vol. 10, no. 1 (February): 109-113.

Roger G. Ibbotson and Paul D. Kaplan. 2000. "Does Asset Allocation Explain 40, 90, or 100 Percent of Performance?" *Financial Analysts Journal*, vol. 56, no. 1 (January/February): 26-33.

Randolph L. Hood. 2005. "Determinants of Portfolio Performance - 20 Years Later," *Financial Analysts Journal*, vol. 61, no. 5 (September/October): 6-8.

Roger G. Ibbotson. 2010. "The Importance of Asset Allocation," *Financial Analysts Journal*, vol. 66, no. 2 (March/April): 18-20.

James X. Xiong, Roger G. Ibbotson, Thomas M. Idzorek, and Peng Chen. 2010. "The Equal Importance of Asset Allocation and Active Management," *Financial Analysts Journal*, vol. 66, no. 2 (March/April): 22-30.

John Nuttall. 2010. "The Importance of Asset Allocation." https://publish.uwo.ca/~jnuttall/asset.pdf

Appendix

Marginal contribution to risk (MCTR)

The investment decision formula is given as:
(TP-RP)=(PP-RP)+(TP-PP)

For simplicity, we define:
A=TP-RP
B=PP-RP
C=TP-PP

This allows us to rewrite the formula as:
A=B+C

The marginal contribution of B and C to the volatility of A can be expressed as:

$$MCTR_B = \frac{(\sigma_B^2 + \rho_{BC}\,\sigma_B\,\sigma_C)}{\sigma_A}$$

$$MCTR_C = \frac{(\sigma_C^2 + \rho_{BC}\,\sigma_B\,\sigma_C)}{\sigma_A}$$

Portfolio asset mix constraints

	Maximum
Specialty fixed income	**40% of total fixed income**
Specialty equities	**50% of total equities**
Small-cap equities	25% of specialty equities
Emerging market equities	25% of specialty equities
Alternatives	**20% of total portfolio**

IC2506758 (06/2025)
025PIM027 (05-2025) IPS INSIGHTS INVESTMENT DECISIONS EN V10 ACC 07/02/2025