PART
1

Asset Allocation and Institutional Investors

CHAPTER 1

Asset Allocation Processes and the Mean-Variance Model

This is the first of two chapters discussing asset allocation, with a focus on the decision-making process of asset allocators who consider portfolios consisting of traditional as well as alternative asset classes. This chapter describes the basic steps of the asset allocation process followed by a typical asset allocator. The objectives and constraints that apply to different types of asset owners are presented, and the important features of strategic and tactical asset allocation approaches are discussed. The chapter then explains the mean-variance approach, which is the best-known quantitative approach to allocation. Finally, some important limitations of the mean-variance approach are discussed.

1.1 IMPORTANCE OF ASSET ALLOCATION

Asset allocation refers both to the process followed by a portfolio manager to determine the distribution of an investor's assets to various asset classes and to the resulting portfolio weights. The allocation is determined to meet one or more objectives subject to a set of constraints set by the investor or dictated by the markets. An objective might be to maximize the expected value of a portfolio at a certain date subject to a set of constraints either established by the investor, such as a maximum level of return volatility or a maximum exposure to certain sectors, or dictated by the markets, such as no short selling of certain assets and a minimum investment level demanded by hedge fund managers.

While asset allocation refers to composition of an investor's portfolio in terms of different asset classes, we define **security selection** as the process through which holdings within each asset class are determined. For example, the asset allocation process may suggest that 20% of an investor's portfolio should be allocated to hedge funds, while security selection in this case is concerned with the hedge fund managers that are eventually selected for the investment purpose.

The importance of asset allocation versus security selection has been the subject of a long-running and controversial debate. The basic question is: Which of these two decisions has a larger impact on a portfolio's performance? As it turns out, the answer to this seemingly simple question is not that simple and, in some sense, it is impossible to provide.

First, we must specify whether the performance of a diversified or a concentrated portfolio is being measured. Clearly, the performance of a concentrated portfolio that consists of some allocation to cash and the rest to a single stock is mostly determined by the security selection decision. A significant portion of the characteristics of this portfolio's performance through time will depend on the choice of the single stock that constitutes the risky part of the portfolio. The choice of allocating a portion of the portfolio to cash will have some impact on the portfolio's performance, but it will be relatively small. In contrast, security selection is likely to have only a minor impact on the portfolio's performance if its equity portion consists of several thousand stocks that are listed around the world.

Second, we need to specify what is meant by portfolio performance. Is the impact of asset allocation on expected monthly return the sole criterion for evaluating the importance of asset allocation? How about higher moments of the return distribution or the beta of the portfolio with respect to some benchmark? As will be discussed, what is meant by performance will have an impact on the importance of asset allocation.

One of the most notable studies on the importance of asset allocation was published in 1986 by Brinson, Hood, and Beebower (BHB). The authors regressed the quarterly rates of return reported by a group of U.S. pension funds against passively managed benchmarks that were created using the weights proposed by the investment policy statements of the pension funds. The goal was to examine the relationship between the actual performance of the funds and the performance that would have been realized had the funds invested their capital in passively managed market indices according to the weights set forth in their investment policy statements. The average *r*-squared of these regressions exceeded 90%. Although BHB were clear in presenting their results, the rest of the investment community took the reported *r*-squared figure and made the blanket statement that more than 90% of the performance of these pension funds could be explained by the asset allocation decision described in the investment policy and that less than 10% of the performance could be explained by the active management decisions of the portfolio managers, such as security selection and tactical tilts. This would be the right conclusion if by performance one means the return *volatility* of the portfolio through time. However, this would be an incorrect conclusion if by performance one means the average return itself through time. In other words, BHB never claimed that 90% of the average return on diversified portfolios could be explained by the asset allocation decision.

As discussed in the CAIA Level I book, the *r*-squared of the regression tells how much of the variation in the dependent variable can be explained by variations in the independent or explanatory variables. In other words, the BHB study only confirmed that more than 90% of variability in the realized returns of fully diversified portfolios could be explained by the asset allocation decision. More important, it did not say anything about the impact of asset allocation on the average return on those pension funds. The study had a lot to say about the second moment of the funds' return distribution and very little about the first moment of their return distribution. Further, the sample included fully diversified portfolios and therefore could not consider the importance of security selection because the portfolio managers had already decided to fully diversify and not to hold concentrated positions in securities that they considered to be undervalued. In short, the study was not meant to answer some of the most important questions faced by asset allocators, but it did spur a large set of studies that have gradually provided answers to practitioners.

Three important questions that could be asked and answered regarding the importance of asset allocation for the performance of diversified portfolios are:

1. How much of the variability of returns across time is explained by the asset allocation framework set forth in the investment policy? That is, how many of a fund's ups and downs are explained by its policy benchmarks? The impact of asset allocation on time variation was studied in BHB. Since then, a number of studies have reexamined this question (Ibbotson and Kaplan 2000). These studies generally agree that a high degree (85% to 90%) of the time variation in diversified portfolios of traditional assets is explained by the overall asset allocation decisions of asset owners and portfolio managers. Therefore, if an asset allocator wants to evaluate the expected volatility of two diversified portfolios, then the asset allocation policies of the two funds will be very informative.
2. How much of the difference in the average returns among funds is explained by differences in the investment policy? That is, if the average returns of two diversified funds are compared, how much of the difference in relative performance can be explained by differences in asset allocation policies? The answer depends greatly on the sample, but most studies show that less than 50% of the difference in average returns can be explained by differences in asset allocation. Other factors—such as asset class timing, style within asset classes, security selection, and fees—explain the remaining differences. Therefore, if an asset allocator wants to evaluate the expected returns of two diversified funds, asset allocation policies of the two funds will be useful, but other factors should be taken into account.
3. What portion of the average return of a fund is explained by its asset allocation policy? In this case, we are considering the absolute performance of a fund. That is, suppose the realized average return on a fund is compared with the return on the fund if the manager had implemented the proposed asset allocation using passive benchmarks. How do these two performances compare? Does the manager outperform the passive implementation of the asset allocation policy? This appears to be the most relevant question, because it directly tests the active management of the portfolio. It turns out that this is the most difficult question to answer, and the available results are highly dependent on the sample and the period they cover. Most studies find that asset allocation has little explanatory power in predicting whether a manager will outperform or underperform the asset allocation return. In fact, available studies covering samples of mutual funds and pension funds conclude that 65% to 85% of them underperform the long-run asset allocation described in their investment policy statements or their passive benchmarks (Ibbotson and Kaplan 2000).[1]

Given the importance of asset allocation, the rest of this chapter focuses on the asset allocation process, the role of asset owners in determining the objectives and constraints of the process, and the difference between strategic and tactical asset allocation programs.

1.2 THE FIVE STEPS OF THE ASSET ALLOCATION PROCESS

This section describes the typical steps that must be taken to implement a systematic asset allocation program.[2] A systematic approach enables the asset allocator to design and implement an investment strategy for the sole benefit of the asset owners. Such an approach needs to focus on the objectives and the constraints that are relevant to the asset owner. We begin with a discussion of the first of the five steps in the asset allocation process: identifying the asset owners and their potential objectives and constraints. In most cases, assets are managed to fund potential liabilities. In some instances, these liabilities represent legal obligations of the asset owner, such as the assets of a defined benefit (DB) pension fund. In other cases, assets are not meant to fund legal obligations but to fund essential needs of the asset owners or their beneficiaries. For example, a foundation's assets are managed to fund its future philanthropic and grant-giving activities. The nature of these potential needs or liabilities is a major determinant of the objectives and constraints of each asset owner.

The second step involves developing an overall approach to asset allocation. A critical step is preparing the investment policy statement. The **investment policy statement** includes the asset allocator's understanding of the objectives and constraints of the asset owners, the menu of asset classes to be considered, whether active or passive approaches will be used, and how often and under what circumstances the allocation will be changed. Such changes arise because of fundamental changes in economic conditions or changes in the circumstances of the asset owner.

The third step is implementing the overall asset allocation policy described in the investment policy statement. This step will require applications of both quantitative and qualitative techniques to determine the weight of each asset class in the portfolio. Since allocations to alternative investments typically involve selection and allocation to managers (e.g., hedge fund and private equity managers), this step will need to have built-in flexibility, as extensive due diligence on managers must be completed, and thus planned allocations may turn out to be infeasible. For instance, the planned allocation may turn out to be less than the minimum investment level accepted by the manager who has emerged on top after the due diligence process.

The fourth step is allocating the capital according to the optimal weights determined in the previous step based on the due diligence and manager evaluation already conducted by the portfolio manager's team or outside consultants.

The final step is monitoring and evaluating the investments. Inevitably, the realized performance of the portfolio will turn out to be different than expected. This will happen because of unexpected changes in the market and because selected fund managers did not perform as expected. As previously stated, the investment policy statement should anticipate circumstances under which the allocation will be revised. This chapter focuses on the first four steps of the asset allocation process. The final step, which deals with benchmarking, due diligence, monitoring, and manager selection, was covered in CAIA Level I (benchmarking) and the rest of this book (due diligence, monitoring, and manager selection).

1.3 ASSET OWNERS

A systematic asset allocation process starts with the asset owners. Chapters 3 through 6 of this book provide detailed descriptions of major types of asset owners and their investment strategies. This section briefly describes major classes of asset owners. Although the list of asset owners will not be exhaustive, it should be sufficient to highlight the differences that exist among major types of asset owners and how their characteristics influence their asset allocation policies. The following sections discuss four categories of asset owners:

1. Endowments and foundations
2. Pension funds
3. Sovereign wealth funds
4. Family offices

1.3.1 Endowments and Foundations

Endowments and foundations serve different purposes but, from an investment policy point of view, share many characteristics. **Endowments** are funds established by not-for-profit organizations to raise funds through charitable contributions of supporters and use the resources to support activities of the sponsoring organization. For example, a university endowment receives charitable contributions from its supporters (e.g., alumni) and uses the income generated by the fund to support the normal operations of the university. Endowments could be small or large, but since they have long investment horizons and are lightly regulated, the full menu of assets is available to them. In fact, among institutional investors, endowments are pioneers in allocating to alternative assets.

Foundations are similar to endowments in the sense that funds are raised through charitable contributions of supporters. These funds are then used to fund grants and support other charitable work that falls within the foundation's mandate. Most foundations are long-term investors and are lightly regulated in terms of their investment activities. However, in order to enjoy certain tax treatments, they are required to distribute a minimum percentage of their assets each year. Foundations are able to invest in the full menu of assets, including alternative asset classes.

1.3.2 Pension Funds

Pension funds are set up to provide retirement benefits to a group of beneficiaries who typically belong to an organization, such as for-profit or not-for-profit businesses and government entities. The organization that sets up the pension fund is called the plan sponsor. There are four types of pension funds (Ang 2014):

1. National pension funds. **National pension funds** are run by national governments and are meant to provide basic retirement income to the citizens of a country. The U.S. Social Security program, South Korea's National Pension Service, and the Central Provident Fund of Singapore are examples of such funds. These types of funds may not operate that differently from sovereign wealth funds,

which are described later in this chapter and in Chapter 5 of this book. The investment allocation decisions of these large funds are controlled by national governments, which makes their management different from private pension funds. Given the size and long-term horizons of these funds, the menu of assets that are available for potential investments is large and includes various alternative assets.

2. Private defined benefit funds. **Private defined benefit funds** are set up to provide prespecified pension benefits to employees of a private business. The plan sponsor promises the employees of the private entity a predefined retirement income, which is based on a set of predetermined factors. Typically, these factors include the number of years an employee has worked for the firm, as well as his or her age and salary. The plan may include provisions for changes in retirement income, such as a cost-of-living adjustment or a portion of the retirement income to be paid to the employee's surviving spouse or young children. The plan sponsor directs the management of the fund's assets. While these funds may not match the size or the length of time horizon of national funds, they are still large long-term investors, and therefore the full menu of asset classes, including alternative assets, are available to them.
3. Private defined contribution funds. **Private defined contribution funds** are set up to receive contributions made by the plan sponsor into the fund. The pension plan specifies the contributions that the plan sponsor is expected to make while the firm employs the beneficiary. The contributions are deposited into accounts that are tied to each beneficiary, and upon retirement, the employee receives the accumulated value of the account. The employee and the plan sponsor jointly manage the fund's assets, in that the sponsor decides on the menu of asset classes available, and the employee decides the asset allocation. The menu of asset classes available to these funds is smaller than both national funds and defined benefit funds. Lumpiness of alternative investments, lack of liquidity, and government regulations typically prevent these funds from investing in a full range of alternative asset classes. Historically, real estate is one alternative asset class that has been available to these funds. In recent years, liquid alternatives have slowly become available as well.
4. Individually managed accounts. **Individually managed accounts** are no different from private savings plans, in which the asset allocation is directed entirely by the employee. Since the funds enjoy tax advantages, they are not free from regulations, and therefore the list of asset classes available to the beneficiary will be limited. In particular, privately placed alternative investments are not normally available to these funds.

1.3.3 Sovereign Wealth Funds

Sovereign wealth funds (SWFs) are funds set by national governments as a way to save and build on a portion of the country's current income for use by future generations of its citizens. SWFs are similar to national pension funds in the sense that they are owned and managed by national governments, but the goal is not to provide retirement income to the citizens of the country.

SWFs have become major players in global financial markets because of their sheer size and their long-term investment horizons. Most SWFs invest a portion of

JWBT1968-c01 JWBT1968-CAIA Printer: Courier Westford August 10, 2016 14:51 Trim: 7in × 10in

their assets in foreign assets. SWFs are relatively new, and their growth, especially in emerging economies, has been tied to the rise in prices of natural resources such as oil, copper, and gold. In some cases, SWFs are funded through the foreign currency reserves earned by countries that enjoy a significant trade surplus, such a China.

SWFs are large and have very long horizons; therefore the full menu of assets should be available to them. However, because national governments manage them, they may not invest in all available asset classes.

1.3.4 Family Offices

Family offices refer to organizations dedicated to the management of a pool of capital owned by a wealthy individual or group of individuals. In effect, it is a private wealth advisory firm established by an ultra-high-net-worth individual or family.

The source of income for a family office can be as varied as the underlying family that it serves. In some cases, the capital is spun off from an operating company, while in other cases, it might be funded with what is known as legacy wealth, which refers to a second or third generation of family members that have inherited their wealth from a prior source of capital generation. The financial resources of a family office can be used for a variety of purposes, from maintaining the family's current standard of living to providing benefits for many future generations to distributing all or a portion of it through philanthropic activities in the current generation. Family offices tend to have relatively long time horizons and are typically large enough to invest in a full menu of assets, including alternative asset classes.

1.4 OBJECTIVES AND CONSTRAINTS

As already discussed, different asset owners have their own particular objectives in managing their assets and face various constraints, which could be internal or external. An **objective** is a preference that distinguishes an optimal solution from a sub-optimal solution. A **constraint** is a condition that any solution must meet. Internal constraints are imposed by the asset owner and may be a function of the owner's time horizon, liquidity needs, and desire to avoid certain sectors. External constraints result from market conditions and regulations. For instance, an asset owner may be prohibited from investing in certain asset classes, or fees and due diligence costs may prevent the owner from considering all available asset classes. The next sections describe the issues that must be considered while attempting to develop a systematic understanding of asset owners' objectives and constraints.

1.5 INVESTMENT POLICY OBJECTIVES

Asset owners' objectives must be expressed in terms of consistent risk-adjusted performance values. In other words, it is safe to assume that asset owners would prefer to earn a high rate of return on their assets. However, higher rates of return are associated with higher levels of risk. Therefore, asset owners should present their objectives in terms of combinations of risks and returns that are consistent with market conditions and their level of risk tolerance. For instance, the objective of earning 30%

per year on a portfolio that has 8% annual volatility is not consistent with market conditions. Such a high return would require a much higher level of volatility. Also, if the asset owner states that her objective is to earn 25% per year with no reference to the level of risk that she is willing to assume, then it could lead the portfolio manager to create a risky portfolio that is entirely inconsistent with her risk tolerance. Therefore, asset owners and portfolio managers need to communicate in a clear language regarding return objectives and risk levels that are acceptable to the asset owner and are consistent with current market conditions.

1.5.1 Evaluating Objectives with Expected Return and Standard Deviations

Consider the following two investment choices available to an asset owner:

- Investment A will increase by 10% or decrease by 8% over the next year, with equal probabilities.
- Investment B will increase by 12% or decrease by 10% over the next year, with equal probabilities.
- The expected return on both investments is 1% (found as the probability weighted average of their potential returns); however, their volatilities will be different (see Equation 4.9 of CAIA Level I).
- Investment A: Standard deviation $= \sqrt{0.5 \times (0.10 - 0.01)^2 + 0.5 \times (-0.08 - 0.01)^2} = 9\%$
- Investment B: Standard deviation $= \sqrt{0.5 \times (0.12 - 0.01)^2 + 0.5 \times (-0.10 - 0.01)^2} = 11\%$

If an asset owner expresses a preference for investment A over investment B, then we can claim that the asset owner is risk averse. Although it is rather obvious to see why a risk-averse asset owner would prefer A to B, it will not be easy to determine whether a risk-averse investor would prefer C to D from the following example:

- Investment C will increase by 10% or decrease by 8% over the next year, with equal probabilities.
- Investment D will increase by 12% or decrease by 9% over the next year, with equal probabilities.

In this case, compared to investment C, investment D has a higher expected return (1.5% to 1%) and a higher standard deviation (10.5% to 9%). Depending on their aversion to risk, some asset owners may prefer C to D, and others, D to C.

1.5.2 Evaluating Risk and Return with Utility

Different asset owners will have their own preferences regarding the trade-off between risk and return. Economists have developed a number of tools for expressing such preferences. Expected utility is the most common approach to specifying the preferences of an asset owner for risk and return. While a utility function is typically used to express preferences of individuals, there is nothing in the theory or application that would prevent us from applying this to institutional investors as well. Therefore, in the context of investments, we define **utility** as a measurement of the satisfaction that an individual receives from investment wealth or return. **Expected utility** is the probability weighted average value of utility over all possible outcomes. Finally, in the context of investments, a **utility function** is the relationship that converts an investment's financial outcome into the investor's level of utility.

Suppose the initial capital available for an investment is W and that the utility derived from W is $U(W)$. Thus, the expected utilities associated with investments A and B can be expressed as follows:

$$E[U(W_A)] = 0.5 \times U(W \times 1.10) + 0.5 \times U(W \times 0.92) \quad (1.1)$$

$$E[U(W_B)] = 0.5 \times U(W \times 1.12) + 0.5 \times U(W \times 0.90) \quad (1.2)$$

The function $U(\bullet)$ is the utility function. The asset owner would prefer investment A to investment B if $E[U(W_A)] > E[U(W_B)]$.

Suppose the utility function can be represented by the log function, and assume that the initial investment is \$100. Then:

$$E[U(W_A)] = 0.5 \times \ln(100 \times 1.10) + 0.5 \times \ln(100 \times 0.92) = 4.611 \quad (1.3)$$

$$E[U(W_B)] = 0.5 \times \ln(100 \times 1.12) + 0.5 \times \ln(100 \times 0.90) = 4.609 \quad (1.4)$$

In this case the asset owner would prefer investment A to investment B because it has higher expected utility. Applying the same function to investments C and D, it can be seen that $E[U(W_C)] = 4.611$ and $E[U(W_D)] = 4.615$. In this case, the asset owner would prefer investment D to investment C.

APPLICATION 1.5.2

Suppose that an investor's utility is the following function of wealth (W):

$$U(W) = \sqrt{W}$$

Find the current and expected utility of the investor if the investor currently has \$100 and is considering whether to speculate all the money in an investment with a 60% chance of earning 21% and a 40% chance of losing 19%. Should the investor take the speculation rather than hold the cash?

The current utility of holding the cash is 10, which can be found as $\sqrt{100}$. The expected utility of taking the speculation is found as:

$$E[U(W)] = (0.60 \times \sqrt{121}) + (0.40 \times \sqrt{81}) = 10.2$$

Because the investor's expected utility of holding the cash is only 10, the investor would prefer to take the speculation, which has an expected utility of 10.2.

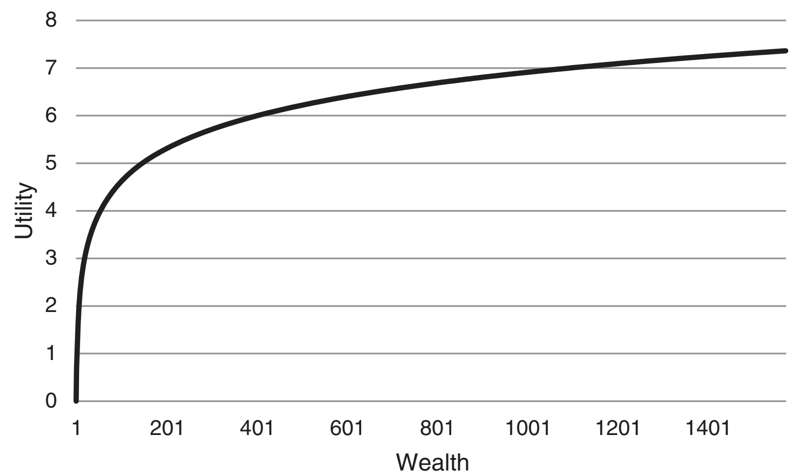

EXHIBIT 1.1 Logarithmic Utility Function

1.5.3 Risk Aversion and the Shape of the Utility Function

We are now prepared to introduce a more precise definition of risk aversion. An investor is said to be **risk averse** if his utility function is concave, which in turn means that the investor requires higher expected return to bear risk. Exhibit 1.1 displays the log function for various values of wealth. We can see that the level of utility increases but at a decreasing rate.

Alternatively, a risk-averse investor avoids taking risks with zero expected payoffs. That is, for risk-averse investors, $E[U(W)] > E[U(W + \tilde{e})]$, where $\tilde{e}$ is a zero mean random error that is independent from W.

1.5.4 Expressing Utility Functions in Terms of Expected Return and Variance

The principle of selecting investment strategies and allocations to maximize expected utility provides a very flexible way of representing the asset owner's preferences for risk and return. The representation of expected utility can be made more operational by presenting it in terms of the parameters of the probability distribution functions of investment choices. The most common form among institutional investors is to present the expected utility of an investment in terms of the mean and variance of the investment returns. That is,

$$E[U(W)] = \mu - \frac{\lambda}{2} \times \sigma^2 \tag{1.5}$$

Here, μ is the expected rate of return on the investment, σ^2 is the variance of the rate of return, and λ is a constant that represents the asset owner's degree of risk aversion. It can be seen that the higher the value of λ, the higher the negative effect of variance on the expected value. For example, if λ is equal to zero, then the investor is said to be risk neutral and the investment is evaluated only on the basis of its expected return. A negative value of λ would indicate that the investor is a risk seeker and actually prefers more risk to less risk.

The **degree of risk aversion** indicates the trade-off between risk and return for a particular investor and is often indicated by a particular parameter within a utility function, such as λ in Equation 1.5. The fact that the degree of risk aversion is divided by 2 will make its interpretation much easier. It turns out that if Equation 1.5 is used to select an optimal portfolio for an investor, then the ratio of the expected rate of return on the optimal portfolio in excess of the riskless rate divided by the portfolio's variance will be equal to the degree of risk aversion.

Example: Suppose $\lambda = 5$. Calculate the expected utility of investments C and D.

$$\mu_C - \frac{\lambda}{2} \times \sigma_C^2 = 1.0\% - (5/2) \times (9.0\%)^2 = -0.01025$$

$$\mu_D - \frac{\lambda}{2} \times \sigma_D^2 = 1.5\% - (5/2) \times (10.5\%)^2 = -0.01256$$

In this case, the expected utility of investment C is higher than that of investment D; therefore, it is the preferred choice. It can be verified that if $\lambda = 1$, then the expected utility of investments C and D will be 0.0059 and 0.00949, respectively, meaning that D will be preferred to C.

APPLICATION 1.5.4

Suppose that an investor's expected utility, $E[U(W)]$, from an investment can be expressed as:

$$E[U(W)] = \mu - \frac{\lambda}{2} \times \sigma^2$$

where W is wealth, μ is the expected rate of return on the investment, σ^2 is the variance of the rate of return, and λ is a constant that represents the asset owner's degree of risk aversion.

Use the expected utility of an investor with $\lambda = 0.8$ to determine which of the following investments is more attractive:

Investment A: $\mu = 0.10$ and $\sigma^2 = 0.04$
Investment B: $\mu = 0.13$ and $\sigma^2 = 0.09$

The expected utility of A and B are found as:

Investment A: $E[U(W)] = 0.10 - \frac{0.8}{2} \times 0.04 = 0.084$

Investment B: $E[U(W)] = 0.13 - \frac{0.8}{2} \times 0.09 = 0.094$

Because the investor's expected utility of holding B is higher, investment B is more attractive.

EXHIBIT 1.2 Properties of Two Hedge Fund Indices

Index	Annualized Mean	Annualized Std. Dev.	Skewness
HFRI Fund of Fund Defensive	7.09%	5.70%	0.24
HFRI Fund Weighted Composite	9.77%	6.76%	–0.61

Source: HFR and authors' calculations.

1.5.5 Expressing Utility Functions with Higher Moments

When the expected utility is presented as in Equation 1.5, we are assuming that risk can be measured using variance or standard deviation of returns. This assumption is reasonable if investment returns are approximately normal. While the normal distribution might be a reasonable approximation to returns for equities, empirical evidence suggests that most alternative investments have return distributions that significantly depart from the normal distribution. In addition, return distributions from structured products tend to deviate from normality in significant ways. In these cases, Equation 1.5 will not be appropriate for evaluating investment choices that display significant skewness or excess kurtosis. It turns out that Equation 1.5 can be expanded to accommodate asset owners' preferences for higher moments (i.e., skewness and kurtosis) of return distributions. For example, one may present expected utility in the following form:

$$E[U(W)] = \mu - \frac{\lambda_1}{2} \times \sigma^2 + \lambda_2 \times S - \lambda_3 \times K \tag{1.6}$$

Here, S is the skewness of the portfolio value; K is the kurtosis of the portfolio; and $\lambda_1, \lambda_2,$ and λ_3 represent preferences for variance, skewness, and kurtosis, respectively. It is typically assumed that most investors dislike variance ($\lambda_1 > 0$), like positive skewness ($\lambda_2 > 0$), and dislike kurtosis ($\lambda_3 > 0$). Note that the signs of coefficients change.

Example: Consider the information about two hedge fund indices in Exhibit 1.2.

If we set $\lambda_1 = 10$ and ignore higher moments, the investor would select the HFRI Fund Weighted Composite as the better investment, as it would have the higher expected utility (0.075 to 0.055). However, if we expand the objective function to include preference for positive skewness and set $\lambda_2 = 1$, then the investor would select the HFRI Fund of Fund Defensive as the better choice, because it would have a higher expected utility (0.29 to –0.54).

1.5.6 Expressing Utility Functions with Value at Risk

The preceding representation of preferences in terms of moments of the return distribution is the most common approach to modeling preferences involving uncertain choices. It is theoretically sound as well. However, the investment industry has developed a number of other measures of risk, most of which are not immediately comparable to the approach just presented. For instance, in the CAIA Level I book, we learned about value at risk (VaR) as a measure of downside risk. Is it possible to use this framework to model preferences in terms of VaR? It turns out that in a rather ad hoc way, one can use the preceding approach to model preferences on risk and return when risk is measured by VaR. That is, we can rank investment choices by calculating the following value:

$$E[U(W)] = \mu - \frac{\lambda}{2} \times \text{VaR}_\alpha \tag{1.7}$$

Here, λ can be interpreted as the degree of risk aversion toward VaR, and VaR_α is the value at risk of the portfolio with a confidence level of α.

We can further generalize Equation 1.7 and replace VaR with other measures of risk. For instance, one could use risk statistics, such as lower partial moments, beta with respect to a benchmark, or the expected maximum drawdown.

1.5.7 Using Risk Aversion to Manage a Defined Benefit Pension Fund

To complete our discussion of objectives, we now consider an application of the previous framework to present the objectives of a defined benefit (DB) pension fund. The following information is available:

- Current value of the fund: €V billion
- Number of asset classes considered: N
- Return on asset class i: R_i
- Weight of asset class i in the portfolio: w_i
- Return on the portfolio: $R_p = \sum_{i=1}^{N} w_i R_i$

Assuming that the preferences of the DB fund can be expressed as in Equation 1.5, the portfolio manager will select the weights, w_i, such that the expected utility is maximized. That is, Equation 1.8 expresses the objective function that is maximized by choosing the values of w_i. Of course, the portfolio manager must ensure that the weights will add up to one and some or all of the weights will need to be positive.

$$E[U(W)] = V \times E[R_p] - \frac{\lambda}{2} \times \text{Var}[V \times R_p] \tag{1.8}$$

1.5.8 Finding Investor Risk Aversion from the Asset Allocation Decision

As mentioned previously, the value of the risk aversion has an intuitive interpretation. The expected excess rate of return on the optimal portfolio ($E[R_P] - R_f$) divided by its variance, σ_P^2, is equal to the degree of risk aversion, λ:

$$\lambda = \frac{E[R_P] - R_f}{\sigma_P^2} \tag{1.9}$$

The value of the parameter of risk aversion, λ, is chosen in close consultation with the plan sponsor. There are qualitative methods that can help the portfolio manager

EXHIBIT 1.3 Hypothetical Risk Returns for Two Portfolios

Portfolio	Annualized Mean	Annualized Std. Dev.
Aggressive	15%	16%
Moderate	9%	8%

select the appropriate value of the risk aversion. The portfolio manager may select a range of values for the parameters and present asset owners with resulting portfolios so that they can see how their level of risk aversion affects the risk-return characteristics of the portfolio under current market conditions.

Example: Consider the information for two well-diversified portfolios shown in Exhibit 1.3. The riskless rate is 2% per year.

Assuming that these are optimal portfolios for two asset owners, what are their degrees of risk aversion?

We know from Equation 1.9 that the expected excess return on each portfolio divided by its variance will be equal to the degree of the risk aversion of the investor who finds that portfolio optimal.

Aggressive investor: $(15\% - 2\%)/(16\%^2) = 5.1$

Moderate investor: $(9\% - 2\%)/(8\%^2) = 10.9$

As expected, the aggressive portfolio represents the optimal portfolio for a more risk-tolerant investor, while the moderate portfolio represents the optimal portfolio for a more risk-averse investor.

APPLICATION 1.5.8

Suppose that an investor's optimal portfolio has an expected return of 10%, which is 8% higher than the riskless rate. If the variance of the portfolio is 0.04, what is the investor's degree of risk aversion, λ?

Using Equation 1.9, λ can be expressed as:

$$\lambda = \frac{E[R_P] - R_f}{\sigma_P^2} = \frac{0.08}{0.04} = 2$$

1.5.9 Managing Assets with Risk Aversion and Growing Liabilities

As mentioned earlier in the chapter, most asset owners are concerned with funding future obligations using the income generated by the assets. In the previous example, the DB plan has liabilities that will need to be met using the fund's assets. Suppose the current value of these liabilities is L euros. Further, suppose the rate of growth in liabilities is given by G, which could be random. In this case, the objective function of Equation 1.8 can be restated as:

$$E[U(W)] = V \times E[R_p] - \frac{\lambda}{2} \times \text{Var}[V \times R_p - L \times G] \tag{1.10}$$

In this case, the DB plan wishes to maximize the expected rate of return on the fund's assets, subject to its aversion toward deviations between the return on the fund and the growth in the fund's liabilities. In other words, the risk of the portfolio is measured relative to the growth in liabilities. Later in this chapter, we will demonstrate how this problem can be solved.

One final comment about evaluating investment choices: Although the framework outlined here is a flexible and relatively sound way of modeling preferences for risk and return, the presentation considered only one-period investments and decisions. Economists have developed methods for extending the framework to more than one period, where the investor has to withdraw some income from the portfolio. These problems are extremely complex and beyond the scope of this book. However, in many cases, the solutions that are based on the single-period approach provide a reasonable approximation of the solutions obtained under approaches that are more complex.

1.6 INVESTMENT POLICY CONSTRAINTS

The previous section introduced the expected utility approach as a simple and yet flexible approach to modeling risk-return objectives of asset owners. This section discusses the typical set of constraints that must be taken into account when trying to select the investment strategy that maximizes the expected utility of the asset owner.

1.6.1 Investment Policy Internal Constraints

Internal constraints refer to those constraints that are imposed by the asset owner as a result of its specific needs and circumstances. Some of these internal constraints can be incorporated into the objective function previously discussed. For example, we noted how the constraint that allocations with positive skewness are preferred could be incorporated into the model. However, there are other types of constraints that may be expressed separately. Some examples of these internal constraints are:

- Liquidity. The asset owner may have certain liquidity needs that must be explicitly recognized. For example, a foundation may be anticipating a large outlay in the next few months and therefore would want to have enough liquid assets to cover those outflows. This will require the portfolio manager to impose a minimum investment requirement for cash and other liquid assets. Even if there are no anticipated liquidity events where cash outlays will be needed, the asset owner may require maintaining a certain level of liquidity by imposing minimum investment requirements for cash and cash-equivalent investments, and maximum investment levels for such illiquid assets as private equity and infrastructure.
- Time horizon. The asset owner's investment horizon can affect liquidity needs. In addition, it is often argued that investors with a short time horizon should take

less risk in their asset allocation decisions, as there is not enough time to recover from a large drawdown. This impact of time horizon can be taken care of by changing the degree of risk aversion or by imposing a maximum limit on allocations to risky assets. Time horizon may impact asset allocation in other ways as well. For instance, certain asset classes are known to display mean reversion in the long run (e.g., commodities). As a result, an investor with a short time horizon may impose a maximum limit on the allocation to commodities, as there will not be enough time to enjoy the benefits of potential mean reversion.

- Sector and country limits. For a variety of reasons, an asset owner may wish to impose constraints on allocations to certain countries or sectors of the global economy. For instance, national pension plans may be prohibited from investing in certain countries, or a university endowment may have been instructed by its trustees to avoid investments in certain industries.

Asset owners may have unique needs and constraints that have to be accommodated by the portfolio manager. However, it is instructive to present asset owners with alternative allocations in which those constraints are relaxed. This will help asset owners understand the potential costs associated with those constraints.

1.6.2 Investment Policy and the Two Major Types of External Constraints

External constraints refer to constraints that are driven by factors that are not directly under the control of the investor. These constraints are mostly driven by regulations and the tax status of the asset owner.

- Tax status. Most institutional investors are tax exempt, and therefore allocation to tax-exempt instruments are not warranted. Because these investments offer low returns, the optimization technique selected to execute the investment strategy should automatically exclude those assets. In contrast, family offices and high-net-worth investors are not tax exempt, and therefore the impact of taxes must be taken into account. For example, constraints can be imposed to sell asset classes that have suffered losses to offset realized gains from those that have increased in value.
- Regulations. Some institutional investors, such as public and private pension funds, are subject to rules and regulations regarding their investment strategies. In the United States, the Employee Retirement Income Security Act (ERISA) represents a set of regulations that affect the management of private pension funds. In the United Kingdom, the rules and regulations set forth by the Financial Services Authority impact pension funds. In these and many other countries, regulations impose limits on the concentration of allocations in certain asset classes.

1.7 PREPARING AN INVESTMENT POLICY STATEMENT

The next step in the process is to develop the overall framework of the asset allocation by preparing an investment policy statement (IPS).[3]

1.7.1 Seven Common Components of an Investment Policy Statement

The policy may include a recommended strategic allocation as well. The following is an outline of a typical IPS based on seven common components.

1. Background. A typical IPS begins with the background of the asset owner and its mission. It reminds all parties who the beneficiaries of the assets are.
2. Objective. The overall goals of the asset owner are described. For instance, the IPS of a foundation may state that the broad objectives are to (1) maintain the purchasing power of the current assets and all future contributions, (2) achieve returns within reasonable and prudent levels of risk, and (3) maintain an appropriate asset allocation based on a total return policy that is compatible with a flexible spending policy while still having the potential to produce positive real returns. The IPS may also provide additional details about the level of risk tolerance, the investment horizon, and the level of expected return that is needed to meet certain liabilities.
3. Asset classes. This segment will include a list of asset classes that the portfolio manager is allowed to consider for allocation. It may provide additional information about how each asset class will be accessed. For instance, the asset owner may decide to use a passive approach to allocations to traditional asset classes and then use active managers for alternative asset classes.
4. Governance. The organizational structure of the fund is described here. The responsibilities of various parties who are involved in the investment process (e.g., the portfolio manager, investment committee, administrator, and custodian) are carefully explained.
5. Manager selection. This section describes the basic framework that the asset owner will follow in selecting outside managers. For example, it may state that all hedge fund managers will need to have three years of experience with at least $100 million in assets under management.
6. Reporting and monitoring. The IPS describes the reporting requirements for the portfolio manager (e.g., frequency, type of reports, and disclosures).
7. Strategic asset allocation. The IPS may include the long-run allocation of the fund during normal periods. The statement may include upper and lower limits for each asset class as well. Further details about strategic asset allocation are discussed in the next section.

1.7.2 Strategic Asset Allocation: Risk and Return

The central focus of strategic asset allocation (SAA) is to create a portfolio allocation that will provide the asset owner with the optimal balance between risk and return over a long-term investment horizon. The SAA not only represents the long-run normal allocation of the investors' assets but also serves as the basis for creating a benchmark that will be used to measure the actual performance of the portfolio. The SAA also serves as the starting point of the tactical asset allocation process, which will adjust the SAA based on short-term market forecasts.[4] (Tactical asset allocation will be discussed in the next chapter.)

SAA is based on long-term risk-return relationships that have been observed in the past and that, based on economic and financial reasoning, are expected to persist under normal economic conditions into the future. While historical risk-return relationships are used as the starting point of generating the inputs needed to create the optimal long-run allocation, these historical relationships should be adjusted to reflect fundamental and potentially long-lasting economic changes that are currently taking place. For example, although long-term historical returns to investment-grade corporate bonds were once high, the prevailing yields on those instruments would indicate that the long-run return from this asset class should be adjusted down.

In developing long-term risk-return relationships for major asset classes, it is important to begin with fundamental factors affecting the economy. Macroeconomic performance of the global economy is the driving force behind the performance of various asset classes. The expected return on all asset classes can be expressed as the sum of three components:

$$\begin{aligned}\text{Asset Class Return} = {} & \text{Short-Term Real Riskless Rate} \\ & + \text{Expected Inflation} + \text{Risk Premium}\end{aligned} \tag{1.11}$$

The real short-term riskless rate of interest is believed to be relatively stable and lower than the real growth rate in the economy.[5] Typically, there is a lower bound of zero for this rate. Therefore, if the global economy is expected to grow at 3% per year going forward, the short-term real riskless rate is expected to be somewhere between zero and 1%. In turn, population growth and increases in productivity are known to be the major drivers of economic growth. Long-term expected inflation is far less stable, as it depends on central banks' policies as well as long-term economic growth. Historically, it was believed that long-term expected inflation would depend on the growth rate in the supply of money relative to the real growth rate in the economy. For instance, it was believed that long-term inflation would be around 5% if the money supply were to grow at 8% in an economy that is growing at 3%. However, this long-term relationship has been challenged by empirical observations following the 2008–9 financial crisis.

Once long-term estimates of the short-term real riskless rate and expected inflation have been obtained, the next step involves the estimation of the long-term risk premium of each asset class. At this stage, one may assume that historical risk premiums would prevail going forward. This would be particularly appropriate if we believe that historical estimates of volatilities, correlations, and risk exposures of various asset classes are likely to persist into the future. For instance, if the long-term historical risk premium on small-cap equities has been 5%, then, assuming 2% expected inflation and a 1% short-term real riskless rate, one could assume an 8% expected long-term return from this asset class.

For several reasons, long-term returns from alternative asset classes could be more difficult to estimate. First, while alternative assets such as real estate and commodities have a long history, some of the more modern alternative asset classes (e.g., hedge funds or private equity) do not have a long-enough history to obtain accurate estimates of their risk exposures and risk premiums. Second, to the degree that alpha was a major source of return for alternative asset classes in the past, the same level of alpha may not be available going forward if there is increased allocation to this asset class by investors. That is, the supply of alpha is limited, and increased competition is bound to reduce it. Third, the alternative investment industry has shown to be quite

EXHIBIT 1.4 Hypothetical Strategic Asset Allocation for an Endowment

	Minimum Allocation	Strategic Allocation	Maximum Allocation
Cash and Short-Term Treasuries	2%	5%	10%
Long-Term Investment-Grade Bonds	10%	10%	20%
High-Yield Bonds	0%	5%	7%
Large-Cap Equities	5%	15%	30%
Small-Cap Equities	0%	10%	15%
Emerging and Frontier Equities	0%	10%	15%
Commodities	0%	5%	10%
Real Estate and Other Real Assets	5%	15%	20%
Private Equity	5%	15%	20%
Hedge Funds	5%	10%	20%
		100%	

innovative and adaptive in response to changing economic conditions. Therefore, we should expect to see new classes of alternative assets going forward, with their potential place in investors' strategic asset allocations unknowable at this point.

1.7.3 Developing a Strategic Asset Allocation

Given the risk-return preference of an asset owner and estimates of expected long-term returns from various asset classes, the portfolio manager and the asset owner can develop an SAA. Exhibit 1.4 displays a hypothetical SAA for a U.S. endowment.

A typical IPS contains a strategic asset allocation and describes the circumstances under which the strategic asset allocation could change; for example, due to fundamental changes in the global economy or changes in the circumstances of the asset owner, the SAA could be revised.

1.7.4 A Tactical Asset Allocation Strategy

Related to SAA is tactical asset allocation (TAA), which is a dynamic asset allocation strategy that actively adjusts a portfolio's SAA based on short- to medium-term market forecasts. TAA's objective is to systematically exploit temporary market inefficiencies and divergences in market values of assets from their fundamental values. Long-term performance of a broadly diversified portfolio is driven mostly by its SAA over time. TAA can add value if designed based on rigorous economic analysis of financial data so it can overcome the headwinds created by the costs associated with portfolio turnover and the fact that global financial markets are generally efficient. The next chapter will provide further details about TAA and the more recent developments based on factor allocation and economic regime-driven investment strategies.

1.8 IMPLEMENTATION

After the completion of the IPS, the next step is its implementation. A variety of quantitative and qualitative portfolio construction approaches are available for this stage. We will focus our attention on the mean-variance approach, as it is the best-known approach, and most of the subsequent developments in this area have attempted to improve on its shortcomings. Some of these approaches are discussed in the next chapter.

Earlier, this chapter discussed how the general expected utility approach could be used to represent preferences in terms of moments of a portfolio's return distribution. In particular, we noted that optimal portfolios could be constructed by selecting the weights such that the following function is maximized:

$$\mu - \frac{\lambda}{2} \times \sigma^2 \tag{1.12}$$

where μ is the expected return on the portfolio, λ is a parameter that represents the risk-aversion of the asset owner, and σ^2 is the variance of the portfolio's return. The next section provides a more detailed description of this portfolio construction technique and examines the solution under some specific conditions. Later sections will discuss some of the problems associated with this portfolio optimization technique and offer some of the solutions that have been proposed by academic and industry researchers.

1.8.1 Mean-Variance Optimization

The portfolio construction problem discussed in this section is the simplest form of mean-variance optimization. The universe of risky investments available to the portfolio manager consists of N asset classes. The single-period total rate of return on the risky asset i is denoted by R_i, for $i = 1, \ldots, N$. We assume that asset zero is riskless, and its rate of return is given by R_0. The weight of asset i in the portfolio is given by w_i. Therefore, the rate of return on a portfolio of the $N + 1$ risky and riskless asset can be expressed as:

$$R_p = w_0 R_0 + w_1 R_1 + \cdots + w_N R_N \tag{1.13}$$

$$w_0 + w_1 + \cdots + w_N = 1 \tag{1.14}$$

For now, we do not impose any short-sale restriction, and therefore the weights could assume negative values.

From Equation 1.14, we can see that $w_0 = 1 - \sum_{i=1}^{N} w_i$. If this is substituted in Equation 1.13 and terms are collected, the rate of return on the portfolio can be expressed as:

$$R_p = w_1(R_1 - R_0) + \cdots + w_N(R_N - R_0) + R_0 \tag{1.15}$$

The advantage of writing the portfolio's rate of return in this form is that we no longer need to be concerned that the weights appearing in Equation 1.15 will add up to one. Once the weights of the risky assets are determined, the weight of the riskless asset will be such that all the weights would add up to one.

Next, we need to consider the risk of this portfolio. Suppose the covariance between asset i and asset j is given by σ_{ij}. Using this, the variance-covariance of the N risky assets is given by:

$$\sum = \begin{bmatrix} \sigma_{11} & \cdots & \sigma_{1N} \\ \vdots & \sigma_{ij} & \vdots \\ \sigma_{N1} & \cdots & \sigma_{NN} \end{bmatrix} \tag{1.16}$$

The portfolio problem can be written in this form, where the weights are selected to maximize the objective function:

$$\max_{w_1,\ldots w_N} E\left[\sum_{i=1}^{N} w_i(R_i - R_0) + R_0\right] - \frac{\lambda}{2} \times \text{Var}\left[\sum_{i=1}^{N} w_i(R_i - R_0) + R_0\right] \tag{1.17}$$

This turns out to have a simple and well-known solution:

$$\begin{bmatrix} w_1 \\ \vdots \\ w_N \end{bmatrix} = \frac{1}{\lambda} \sum{}^{-1} \times \begin{bmatrix} E[R_1 - R_0] \\ \vdots \\ E[R_N - R_0] \end{bmatrix} \tag{1.18}$$

The solution requires one to obtain an estimate of the variance-covariance matrix of returns on risky assets. Then the inverse of this matrix will be multiplied into a vector of expected excess returns on the N risky assets. It is instructive to notice the role of the degree of risk aversion. As the level of risk aversion (λ) increases, the portfolio weights of risky assets decline. In addition, those assets with large expected excess returns tend to have the largest weights in the portfolio.

1.8.2 Mean-Variance Optimization with a Risky and Riskless Asset

To gain a better understanding of the solution, consider the case of only one risky asset and a riskless asset. In this case, the optimal weight of the risky asset using Equation 1.18 will be:

$$w = \frac{1}{\lambda} \frac{E[R - R_0]}{\sigma^2} \tag{1.19}$$

The optimal weight of the risky asset is proportional to its expected excess rate of return, $E[R - R_0]$, divided by its variance, σ^2. Again, the higher the degree of risk aversion, the lower the weight of the risky asset.

For example, with an excess return of 10%, a degree of risk aversion (λ) of 3, and a variance of 0.05, the optimal portfolio weight is 0.67. This is found as (1/3) × (0.10/0.05). Note that Equation 1.19 may be used to solve for any of the variables, given the values of the remaining variables.

APPLICATION 1.8.2

Consider the case of mean-variance optimization with one risky asset and a riskless asset. Suppose the expected rate of return on the risky asset is 9% per year. The annual standard deviation of the index is estimated to be 13% per year. If the riskless rate is 1%, what is the optimal investment in the risky asset for an investor with a risk-aversion degree of 10?

The solution is:

$$w = \frac{1}{10} \times \frac{0.09 - 0.01}{0.13^2} = 47.3\%$$
$$w_0 = 1 - 47.3\% = 52.7\%$$

That is, this investor will invest 47.3% in the risky asset and 52.7% in the riskless asset. By varying the degree of risk aversion, we can obtain the full set of optimal portfolios.

1.8.3 Mean-Variance Optimization with Growing Liabilities

Equation 1.10 displayed the formulation of the problem when the asset owner is concerned with the tracking error between the value of the assets and the value of the liabilities. Similar to Equation 1.18, a general solution for that problem can be obtained as well. Here we present a simple version of it when there is only one risky asset. The covariance between the rate of growth in the liabilities and the growth in assets is denoted by δ, and L is the value of liabilities relative to the size of assets:

$$w = \frac{1}{\lambda}\frac{E[R - R_0]}{\sigma^2} + L\frac{\delta}{\sigma^2} \quad (1.20)$$

It can be seen that if the risky asset is positively correlated with the growth in liabilities (i.e., $\delta > 0$), then the fund will hold more of that risky asset. The reason is that the risky asset will help reduce the risk associated with growth in liabilities. For instance, if the liabilities behaved like bonds, then the fund would invest more in fixed-income instruments, as they would reduce the risk of the fund.

Example: Continuing with the previous example, suppose the covariance between the risky asset and the growth rate in the fund's liabilities is 0.002, and the value of liabilities is 20% higher than the value of assets. What will be the optimal weight of the equity allocation?

$$w = \frac{1}{10}\frac{0.09 - 0.01}{0.13^2} + 1.2\frac{0.002}{0.13^2} = 61.5\%$$

It can be seen that, compared to the previous example, the fund will hold about 14% more in the risky asset because it can hedge some of the liability risk.

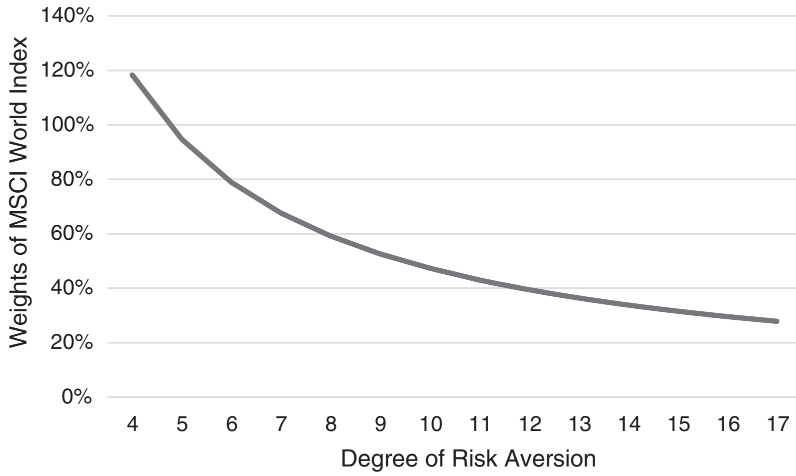

EXHIBIT 1.5 Optimal Weights of Risky Investment and Degree of Risk Aversion

By changing the degree of risk aversion in the first example, we can obtain a set of optimal portfolios, as shown in Exhibits 1.5 and 1.6.

It can be seen that at low degrees of risk aversion (e.g., 4), the investor will be investing more than 100% in the MSCI World Index, which means a leveraged position will be used. In addition, we can see the full set of expected returns and volatility that the optimal portfolios will assume, which is referred to as the efficient frontier.

The points appearing in Exhibit 1.6 correspond to various degrees of risk aversion. For instance, the optimal risk-return trade-off for an investor with a degree of

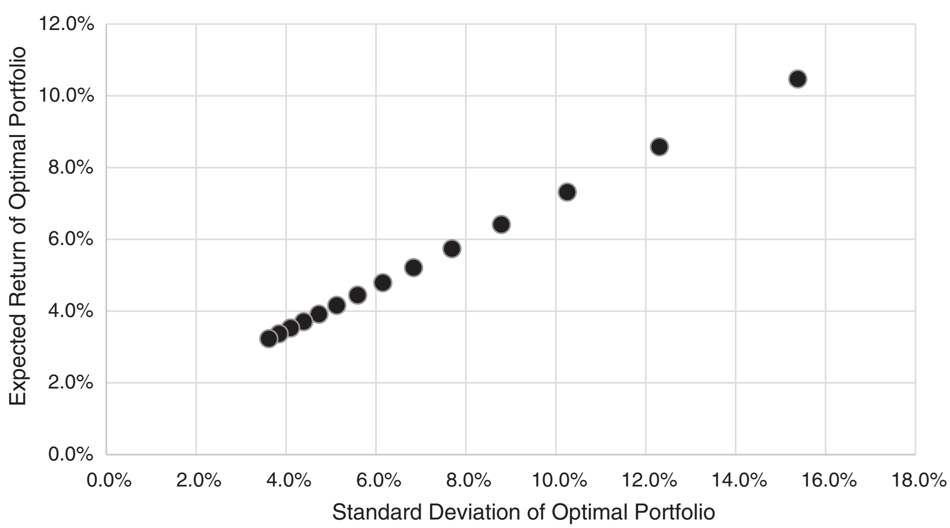

EXHIBIT 1.6 Expected Returns and Standard Deviations of Optimal Portfolios

risk aversion of 4 is represented by a portfolio that is expected to earn 10.5% with a volatility of about 15.4%.

1.8.4 Mean-Variance Optimization with Multiple Risky Assets

It turns out that a similar graph will be obtained even if the number of asset classes is greater than one. In that case, the graph will be referred to as the efficient frontier. The **efficient frontier** is the set of all feasible combinations of expected return and standard deviation that can serve as an optimal solution for one or more risk-averse investors. Put differently, no portfolio can be constructed with the same expected return as the portfolio on the frontier but with a lower standard deviation, or, conversely, no portfolio can be constructed with the same standard deviation as the portfolio on the frontier but with a higher expected return.

Example: In this example, the set of risky asset classes is expanded to three. The necessary information is provided in Exhibit 1.7. The figures are estimated using monthly data in terms of USD. The annual riskless rate is assumed to be 1%. Note that these estimates are typically adjusted to reflect current market conditions. This example is meant to illustrate an application of the model.

Using the optimal solution that was displayed in Equation 1.18, the optimal weights of a portfolio consisting of the three risky assets and one riskless asset can be calculated for different degrees of risk aversion. The results are displayed in Exhibit 1.8.

A number of interesting observations can be drawn from these results. First, notice that the optimal weights are not very realistic. For example, for every degree of risk aversion, the optimal portfolio requires us to take a short position in the MSCI World Index. Second, the optimal investment in the HFRI index exceeds 100% for some degrees of risk aversion considered here. Third, unless the degree of risk aversion is increased beyond 40, the optimal portfolio requires some degree of leverage (i.e., negative weight for the Treasury bills). Finally, the bottom two rows display annual mean and annual standard deviation of the optimal portfolios. These represent points on the efficient frontier.

EXHIBIT 1.7 Statistical Properties of Three Risky Asset Classes

			Annual Variance-Covariance		
1990–2015	**Annual Mean**	**Annual Standard Deviation**	**MSCI World**	**Barclays Global Aggregate**	**HFRI Fund Weighted Composite**
MSCI World	8%	15%	0.0234	0.0025	0.0073
Barclays Global Aggregate	6%	5%	0.0025	0.0029	0.0006
HFRI Fund Weighted Composite	10%	7%	0.0073	0.0006	0.0046

Source: Bloomberg and authors' calculations.

EXHIBIT 1.8 Optimal Weights and Statistics for Different Degrees of Risk Aversion

	Degree of Risk Aversion			
	10	15	20	40
MSCI World	−86%	−57%	−43%	−22%
Barclays Global Aggregate	186%	124%	93%	46%
HFRI Fund Weighted Composite	307%	205%	153%	77%
Treasury Bills	−307%	−171%	−103%	−2%
Optimal Portfolio Statistics				
Annual Mean	31%	21%	16%	9%
Annual Standard Deviation	17%	12%	9%	5%

Source: Authors' calculations. (Note that because of rounding errors, the weights do not add up to one.)

As we just saw, mean-variance optimization typically leads to unrealistic weights. A simple way to overcome this problem is to impose limits on the weights. For instance, in the example just provided, we can impose the constraint that the weights must be nonnegative. Unfortunately, when constraints are imposed on the weights, a closed-form solution of the type presented in Equation 1.18 can no longer be obtained, and we must use a numerical optimization package to solve the problem.[6]

If we repeat the example but impose the constraint that weights cannot be negative, the resulting optimal portfolios will reflect those displayed in Exhibit 1.9.

It can be seen that the weight of the MSCI World Index is constantly zero for all degrees of risk aversion. This means that portfolios that are on the efficient frontier in this case do not have any allocation to the MSCI World Index. Another important point to consider is that the optimal portfolios do not have the same attractive risk-return properties. By imposing a constraint, the resulting portfolios are not as optimal as they were when there were no constraints.

The mean-variance optimization approach discussed in this section can be presented in different forms. The Appendix at the end of this book provides two alternative methods that have appeared in the literature. The advantages of the approach presented in this section are twofold. First, as Equations 1.18 and 1.19 showed,

EXHIBIT 1.9 Optimal Weights without Short Sale Constraints

	Degree of Risk Aversion			
	10	15	20	40
MSCI World	0%	0%	0%	0%
Barclays Global Aggregate	138%	92%	69%	34%
HFRI Fund Weighted Composite	176%	117%	88%	44%
Treasury Bills	−214%	−109%	−57%	22%
Optimal Portfolio Statistics				
Annual Mean	23%	16%	12%	7%
Annual Standard Deviation	15%	10%	7%	4%

simple closed-form solutions can be obtained when there are no constraints. Second, the approach can be easily expanded to take into account preferences for higher moments of the probability distributions of asset returns. This will turn out to be important for our purpose, as alternative investments tend to have return distributions that deviate significantly from the normal distribution; therefore, their higher moments will be of interest to investors.

1.8.5 Hurdle Rate for Mean-Variance Optimization

An interesting implication of mean-variance portfolio optimization when there is a riskless asset is that the benefits of diversification can be shown to cause low-return assets to be desirable for inclusion in a portfolio. It is easy to show that the addition of a new asset to an already optimal portfolio will improve its risk-return properties (i.e., increases the expected utility) if the expected rate of return on this new asset exceeds a hurdle rate. The **hurdle rate** is an expected rate of return that a new asset must offer to be included in an already optimal portfolio. An asset being considered for addition to a portfolio should be included in the portfolio when the following expression is satisfied:

$$E[R_{\text{New}}] - R_f > (E[R_P] - R_f) \times \beta_{\text{New}} \tag{1.21}$$

Here, $E[R_{\text{New}}]$ is the expected rate of return on the new asset, $E[R_p]$ is the expected rate of return on the optimal portfolio, R_f is the riskless rate, and β_{New} is the beta of the new asset with respect to the optimal portfolio.[7]

Equation 1.21 states that the addition of the new asset to an optimal portfolio will improve the risk-return properties of the portfolio if the expected excess rate of return on the new asset exceeds the expected excess rate of return on the optimal portfolio times the beta of the new asset. If the new asset satisfies Equation 1.21, then its addition to the optimal portfolio will move the efficient frontier in the northwest direction. For example, if the beta of the new asset is zero, then the new asset will improve the optimal portfolio as long as its expected rate of return exceeds the riskless rate. If the new asset has a negative beta, then it could improve the optimal portfolio even if its expected rate of return is negative. In other words, assets that can serve as hedging instruments could have negative expected returns and still improve the performance of a portfolio.

APPLICATION 1.8.5

Suppose that an investor is using mean-variance optimization, the expected annual rate of return of an optimal portfolio is 16%, and the riskless rate is 1% per year. What is the hurdle rate for a new asset that has a beta of 0.5 with respect to the optimal portfolio?

Given the formula of Equation 1.21, the hurdle rate would be:

$$E[R_{\text{New}}] - 1\% > (16\% - 1\%) \times 0.5 \quad [R_{\text{New}}] > 8.5\%$$

What if the new asset has a beta of –0.3, which means that it can hedge some of the portfolio's risk?

$$E[R_{\text{New}}] - 1\% > (16\% - 1\%) \times (-0.3) \qquad E[R_{\text{New}}] > -3.5\%$$

In this case, even if the new asset is expected to lose some money (i.e., less than 3.5%), its addition to the optimal portfolio could still improve its risk-return properties.

1.8.6 Issues in Using Optimization

We have already seen that even in the case of three risky asset classes and the riskless rate, reasonable estimates of the weights could not be obtained unless short sale restrictions were imposed, and even in that case, no allocation to the MSCI World Index was recommended. This lack of allocation to an asset class was driven mostly by the portfolio's high volatility and relatively low return. In other words, using the past as an unbiased forecast of the future and using global equities as an asset class would have meant no allocation would be made to equities, and the portfolio would have focused on the remaining two assets. In practice, implementing an optimization method (mean-variance optimization, in particular) for portfolio allocation decisions raises major challenges.

1.8.7 Optimizers as Error Maximizers

Portfolio optimizers are powerful tools for finding the best allocation of assets to achieve superior diversification, given accurate estimates of the parameters of the return distributions. When the mean-variance method is used, there is a need for accurate estimates of expected returns and the variance-covariance matrix of asset returns. However, portfolio optimizers that use historical estimates of the return distributions have been derogatorily called "error maximizers" due to their tendency to generate solutions with extreme portfolio weights. For example, very large portfolio weights are often allocated to the assets with the highest mean returns and lowest volatility, and very small portfolio weights are allocated to the assets with the lowest mean returns and highest volatility. It is then argued that assets with the highest estimated means are likely to have the largest positive estimation errors, whereas assets with the lowest estimated means are likely to have the largest negative estimation errors. Hence, mean-variance optimization is likely to maximize errors. Therefore, if an analyst overstates mean returns and understates volatility for an asset, then the weights that the model recommends are likely to be much larger than an institutional investor would consider reasonable. Further, other assets are virtually omitted from the portfolio if the analyst supplies low estimates of mean returns and high estimates of volatility.

A typical attempt to use a mean-variance optimization model for portfolio allocation is this: (1) The portfolio manager supplies estimates of the mean return, volatility, and covariance for all assets; (2) the optimizer generates a highly unrealistic solution that places very large portfolio weights on what are considered the most attractive assets, with high mean return and low volatility, and zero or minuscule portfolio weights on what are considered the least attractive assets, with low mean return and high volatility; and (3) the portfolio manager then modifies the model by adding constraints or altering the estimated inputs—including mean, variance, and covariance—until the resulting portfolio solutions appear reasonable.

The problem with this process is that the portfolio weights become driven by the subjective judgments of the analyst rather than by the analyst's best forecasts of risk and return. The remaining sections discuss a variety of challenges that emanate from the tendency of mean-variance portfolio optimizers to select extreme portfolio weights.

It is important to point out that while higher-frequency data tends to improve the accuracy of the estimated variance-covariance matrix, it will do nothing to improve the accuracy of the estimated means; only a longer history has the potential to do so. To see this, assume that we have five years of annual data on the price of an asset. The annual rate of return on the asset is calculated to be $R_{t+1} = \ln(P_{t+1}/P_t)$. Now consider an estimate of the average return using the observed four annual returns:

$$\bar{R} = \frac{1}{4}\{\ln(P_2/P_1) + \ln(P_3/P_2) + \ln(P_4/P_3) + \ln(P_5/P_4)\} = \frac{1}{4}\{\ln(P_5/P_1)\} \qquad (1.22)$$

Notice that all the intermediate prices cancel out, and only the first and the last prices matter. This result will not change even if one could use daily or even high-frequency data. The accuracy of the mean depends on the length of data and not on the frequency of the observations.

This observation regarding mean accuracy leads to the following dilemma. To obtain accurate estimates of the mean, it is necessary to have a very long history of prices. However, firms, industries, and economies go through drastic changes over long periods, and it would be highly unlikely that all observed prices would have come from the same distribution. In other words, of all the estimated parameters, the estimated mean is most likely to be the least accurate, yet it is the one with the most influence on the outputs of the mean-variance optimization.

The final difficulty in deriving estimates of return and risk for each asset class is that return and risk are nonstationary, meaning that the levels of risk and return vary substantially over time. Therefore, the true risk and return over one period may be substantially different from the risk and return of a different period. Thus, in addition to traditional estimation errors for a stationary process, estimates for security returns may include errors from shooting at a moving target.

1.8.8 Data Issues for Illiquid Assets

As noted, mean-variance optimizers can be error maximizers. Therefore, erroneous forecasts of the mean, variance, and covariance can result in extreme portfolio weights, with a resulting portfolio concentration in a few assets with estimated high means and estimated low volatilities. Most institutions view such concentrated positions as unacceptable speculation on the validity of the forecasted mean and volatility.

Although higher frequency of observed data can improve the accuracy of the estimated variance and covariance, for most alternative assets, high-frequency data is not available. More important, the assets whose prices cannot be observed with high frequency tend to be illiquid, and the reported quarterly returns are based on appraisals such as those used in real estate and private equity. These prices tend to be smoothed and therefore can substantially understate the variance and covariance of returns. Because volatility and covariance are key inputs in the optimization process, asset classes with low estimated correlation and volatility receive relatively large weights in the optimal portfolio. If smoothing has caused the reported volatility and correlation of an asset to substantially underestimate the true volatility, then a traditional mean-variance optimizer would overweight the asset. In this case, and to prevent extremely large allocations to assets with smoothed returns, the time series of returns may be unsmoothed, as discussed in Chapter 15, before being added to the optimization routine. But unsmoothing is imperfect, and other issues with accurately forecasting volatility and correlation remain.

1.8.9 Data Issues for Large-Scale Optimization

The problems with covariance estimation include the potentially large scale of the inputs required. This is typically not an issue when working at the asset class level, at which the investor may consider 10 asset classes for inclusion in the portfolio. However, optimizing an equity portfolio selected from a universe of 500 stocks has very large data requirements. A 500-asset optimization problem requires estimates of covariance between each pair of the 500 assets. Not only does this problem require $n(n-1)/2$, or 124,750 covariance estimates, but it is also difficult to be confident in these estimates, especially when there are too many to analyze individually. Also, notice that to estimate 124,750 covariance terms, we need more than 124,750 observations, or more than 10,000 years of monthly data or more than 340 years of daily data.

The problem of needing to calculate thousands of covariance estimates can be reduced with factor models. Rather than estimating the relationships between each pair of stocks in a 500-stock universe, it can be easier to estimate the relationship between each stock and a limited number of factors. While some investors simply choose to estimate the single-factor market model beta of each stock, others use multifactor models. To see how a factor model can reduce the data requirement, suppose the return on each asset can be expressed as a function of one common factor and some random noise:

$$R_i = a_i + b_i F + \varepsilon_i \quad (1.23)$$

where F is the common factor and a_i, b_i are the estimated parameters. It is assumed that for two different assets, the error terms ε_i and ε_j are uncorrelated with each other. Under this assumption, the covariance between two assets is given by:

$$\text{Cov}[R_i, R_j] = b_i \times b_j \times \text{Var}[F] \quad (1.24)$$

This means that to estimate the covariance matrix of 500 assets, we need 500 estimates of b_i and one estimate of Var[F].

1.8.10 Mean-Variance Ignores Higher Moments

A problem that is especially acute with alternative investments is that the mean-variance optimization approach considers only the mean and variance of returns.

This means that the optimization model does not explicitly account for skewness and kurtosis. Investors' expected utility can be expressed in terms of mean and variance alone if returns are normally distributed. However, when making allocations to alternative investments and other investments with nonzero skewness and nonzero excess kurtosis, portfolio optimizers tend to suggest portfolios with desirable combinations of mean and variance but with highly undesirable skewness and kurtosis. In other words, although mean-variance optimizers can identify the efficient frontier and help create portfolios with the highest Sharpe ratios, they may be adding large and unfavorable levels of skewness and kurtosis to the portfolio.

For example, two assets with returns that have the same variance may have very different skews. In a competitive market, the expected return of the asset with the large negative skew might be substantially higher than that of the asset with the positive skew to compensate investors willing to bear the higher downside risk. A mean-variance optimizer typically places a much higher portfolio weight on the negatively skewed asset because it offers a higher mean return with the same level of variance as the other asset. The mean-variance optimizer ignores the unattractiveness of an asset's large negative skew and, in so doing, maximizes the error.

There are three common ways to address this complication. First, as we saw earlier, it is possible to expand our optimization method to account for skewness and kurtosis of asset returns. Second, we can continue with our mean-variance optimization but add the desired levels of the skewness and kurtosis as explicit constraints on the allowed solutions to the mean-variance optimizer, such as when the excess kurtosis of the portfolio returns is not allowed to exceed 3, or when the skewness must be greater than –0.5. A problem with incorporating higher moments in portfolio optimization is that these moments are extremely difficult to predict, as they are highly influenced by a few large negative or positive observations. In addition, in the second approach, a portfolio with a desired level of skewness or kurtosis may not be feasible at all. Finally, the analyst may choose to explicitly constrain the weight of those investments that have undesirable skew or kurtosis. For example, the allocation to a hedge fund strategy that is known to have large tail risk (e.g., negative skew) might be restricted to some maximum weight.

1.8.11 Other Issues in Mean-Variance Optimization

The results from a mean-variance optimization can be extremely sensitive to the assumptions, as small changes in the mean return or covariance matrix (i.e., the set of all variances and covariances) can lead to enormously different prescribed portfolio weights. The high sensitivity of portfolio optimizers to the input data has led to approaches that attempt to harness the power of optimization to identify diversification potential without generating extreme portfolio weights. In addition, in most cases, portfolio managers want to adjust the historical estimates to reflect their views about the estimated parameters going forward. For instance, a portfolio manager may want to incorporate her view that the health care sector is likely to do better than indicated by its historical track record. Perhaps the most popular modification to account for views and obtain reasonable estimates of weights is described by Black and Litterman.

The first problem addressed by the Black-Litterman approach is the tendency of the user's estimates of mean and variance to generate extreme portfolio weights in a mean-variance optimizer. Note that if a security truly and clearly offered a large expected return, low risk, and high diversification potential, then demand for the security would drive its price upward and its expected return downward until the demand for the security equaled the quantity available. In competitive markets, securities prices tend toward offering a perceived combination of risk and return in line with other assets.

The key to understanding the Black-Litterman approach is to understand that if a security offers an equilibrium expected return, then the demand for the asset will equal the supply. Further, the optimal allocation of the asset into every well-diversified portfolio will be equal to the weight of the asset in the market portfolio (i.e., the market weight). Thus, an equilibrium expected return for a security is the expected return that causes the optimal weight of that security in investor portfolios to equal its market weight.

This observation means that if the portfolio manager has no views about the future performance of a particular asset class, then its market weight should be used. For instance, a market-cap-weighted portfolio of global equities would be optimal. However, since market cap weights are not well defined for some asset classes, the Black-Litterman approach will need to be adjusted for application to alternative assets.

The primary innovation of the Black-Litterman approach is that it allows the investor to blend asset-specific views of each asset's expected return with views that would be consistent with market weights in a market equilibrium model.

Whereas some asset allocators employ advanced techniques such as the Black-Litterman approach to reduce the sensitivity of the weights to the expected risks and returns, a much larger number of asset allocators choose to add additional constraints to the optimization model to circumvent the difficulties and sensitivities of mean-variance optimization. Common additional constraints include the following:

- Limits on estimated correlation between the return on the optimal portfolio and the return on a predefined benchmark
- Limits on divergences of portfolio weights from benchmark weights
- Limits or ranges on the prescribed portfolio weights

The last constraint, limits on portfolio weights, is the most popular. These constraints can prescribe upper or lower weights, outside of which the asset allocator will not invest. The portfolio optimizer is forced to generate weights within those ranges. In practice, however, many investors use so many constraints that the constraints have more influence on the final asset allocation than does the mean-variance optimization process. While each of the added constraints may help the asset allocator avoid extreme weights, the approach may ultimately lead to having the constraints define the allocation rather than the goal of diversification.

1.9 CONCLUSION

This chapter has introduced the asset allocation process, with a focus on using the mean-variance approach to create optimal portfolios. The asset allocation process discussed in this chapter consists of five steps, of which four were discussed.

Step 1 focuses on understanding who the asset owners are and their mission in managing assets. Step 2 examines the asset owner's objectives and constraints. Here we discussed the expected utility and its mean-variance version as a flexible way of quantifying an asset owner's objectives. Two types of constraints, internal and external, were explained.

Step 3 deals with preparing the investment policy statement, which will provide a general framework for the actual asset allocation. One of the key features of this statement is to develop a list of asset classes to be considered. Step 4 is implementation, which was covered with a focus on mean-variance optimization and its potential problems.

NOTES

1. Hundreds of studies have attempted to determine if active managers outperform passive strategies. S&P Dow Jones Indices publishes SPIVA® U.S. Scorecard on a regular basis. It reports on the relative performance of U.S. mutual funds.
2. More detailed discussions of asset allocation processes can be found in Maginn et al. (2007) and Ang (2014).
3. For further details, see Maginn et al. (2007).
4. For a detailed discussion of strategic asset allocation, see Eychenne, Martinetti, and Roncalli (2011) and Eychenne and Roncalli (2011).
5. In a simple equilibrium model, the short-term real riskless rate is shown to equal the real growth in the economy minus a premium that depends on the volatility of the economy's real growth rate and the degree of risk aversion. See Cox, Ingersoll, and Ross (1985).
6. The problem is still a rather standard optimization program and can be solved using Solver from Excel or similar packages.
7. From linear regression and the CAPM we know that $\beta_{New} = Cov[R_P, R_{New}]/Var[R_p]$.

REFERENCES

Ang, A. 2014. *Asset Management: A Systematic Approach to Factor Investing*. Oxford: Oxford University Press.

Brinson, G. P., L. R. Hood, and G. L. Beebower. 1986. "Determinants of Portfolio Performance." *Financial Analysts Journal* 42 (4): 39–44.

Cox, J., J. Ingersoll, and S. A. Ross. 1985. "An Intertemporal General Equilibrium Model of Asset Prices." *Econometrica* 53 (2): 363–84.

Eychenne, K., S. Martinetti, and T. Roncalli. 2011. "Strategic Asset Allocation." Lyxor Asset Management, Paris.

Eychenne, K., and T. Roncalli. 2011. "Strategic Asset Allocation—An Update Following the Sovereign Debt Crisis." Lyxor Asset Management, Paris.

Ibbotson, R. G., and P. D. Kaplan. 2000. "Does Asset Allocation Policy Explain 40, 90, or 100 Percent of Performance?" *Financial Analysts Journal* 56 (1): 26–33.

Maginn, J. L., D. L. Tuttle, J. E. Pinto, and D. W. McLeavey. 2007. *Managing Investment Portfolios*. 3rd ed. Hoboken, NJ: John Wiley & Sons.