

PART

One

Alpha/Beta Building Blocks of Portfolio Management

COPYRIGHTED MATERIAL

CHAPTER 1

The Modern Endowment Allocation Model

Over the past two decades, many institutional investors, beginning with the larger ones, adopted a broadly diversified asset allocation with greatly reduced allocations of traditional U.S. equities and bonds. Endowments and foundations, in particular, adopted this allocation model.

TRULY LONG-TERM ORIENTATION

The traditional and widely used institutional asset allocation benchmark or policy portfolio commonly consisted of a majority of U.S. equities and a reciprocal proportion of U.S. bonds. Beginning in the 1990s, endowments increasingly took advantage of the nature of their liabilities to adopt a purposeful diversification directed toward long term outcomes and a much-reduced focus on the short term.

NOVEL ASSET CLASSES AND SPECIAL ACCESS

This (re)evolution was largely driven by a growing awareness of investable assets containing return premia derived from such nontraditional features as illiquidity, longer investment horizons, less-than-transparent valuation, and other factors. Leading endowments turned to nonstandard alternative assets and to managers with niche expertise, special flexibility, and unique market access. As shown in Exhibit 1.1, between 1992 and 2008, college and university endowment allocations to nonstandard assets—real estate, hedge funds, private equity, natural resources, venture capital, and other alternatives—rose from 3 percent to more than 25 percent, with a




EXHIBIT 1.1 College and University Endowment Asset Allocation
(equally weighted)

Source: 1992–2008 NACUBO/TIAA-CREF Endowment Surveys

commensurate decrease in allocations to fixed income, cash, and public U.S. equity (National Endowment Surveys, hereafter NES).¹

Dramatically, by the end of this period, the largest endowments and foundations—those over \$1 billion in assets—had increased their commitment to nonstandard assets, on average, to 50 percent of the total portfolio.

REMAKING THE INVESTMENT MANAGER RELATIONSHIP

Endowments and foundations have traditionally depended on a successful selection of internal and external investment managers. The growing commitment to alternative assets as opposed to traditional assets has required a reworking of many of these processes, including far more intensive manager screening, vetting, and monitoring; significantly increased sensitivity to alignment of interests between the institution and its managers; and the need to incubate and nurture special investment talent (internally as well as externally). Specific mechanisms include direct support for investment startups and spinoffs, selected acceptance of lockups and performance fees, clawbacks, and a regular intensive analysis of manager performance and risk.


MORE MARKET-SENSITIVE ALLOCATIONS

The tradition of the long term policy portfolio with relatively fixed asset categories was at one point ubiquitous in the endowment and foundation world. The strategic policy portfolios and accompanying asset class buckets were intended to act as benchmarks against which actual allocations could be gauged. This system promoted rigid allocations, regular rebalancing, and adherence to the underlying asset buckets. However, increased market volatility and the appearance of attractive new asset classes called into question this tradition of overly rigid allocations and fixed asset buckets. Institutions were urged to work in a more flexible fashion with allocations (for example, establish wider allocation bands and more frequent policy portfolio reviews) and to use assets that did not necessarily fit into the traditional categories (Bernstein 2003; Leibowitz and Hammond 2004). The late Peter L. Bernstein was one of the earliest and most articulate authors arguing for a rethinking of the policy portfolio concept. Such an approach would, it was hoped, enable the endowment to deal with increased investment uncertainty as well as take advantage of new opportunities.

Beginning in the 1990s, this modern approach to endowment management paid off handsomely, providing returns that far surpassed traditional equity-and-bond portfolios. As shown in Exhibit 1.2, the largest endowment portfolios averaged an equally weighted return of 12.1 percent for the 19 years ending in June 2008, in contrast to 9 percent average annual return for a portfolio of 60 percent equities and 40 percent bonds (S&P 500 and Lehman Aggregate indexes, respectively). Over this period, the larger endowments' 12.1 percent return also far exceeded the 8.2 percent earned by the less diversified smaller endowments. These realized returns for the larger endowments were also far greater than the theoretical returns projected from the standard expected risk-and-return (covariance) models.

This success of the modern allocation model did not go unnoticed. The value of this approach was underscored in 2000, and then revised in 2009 by the publication of David Swenson's groundbreaking treatise, *Pioneering Portfolio Management*. Many endowments, foundations, and pension funds began to look for ways to emulate these allocations.

The apparent attractions of the modern endowment allocation model for institutions and individuals include the well-known benefits of diversification (for example, low correlations with traditional stocks and bonds), risk control (closer to bond-like volatility), and return enhancement (aiming for stock-like returns). For many, these attractions seemed obvious, especially during a period with some of the lowest interest rates on record and diminished expectations for the standard equity risk premia. As a



*1990–1997 Size > \$400 million; 1998–2008 Size > \$1 billion
**1990–1997 Size < \$25 million; 1998–1999 Size , \$75 million; 2000–2001 Size < \$100 million; 2002–2008 Size < \$25 million

EXHIBIT 1.2 College and University Endowments' Annual Investment Returns
Source: 1990–2008 NACUBO/TIAA-CREF Endowment Surveys

result, a fundamental shift seemed to be under way in what it meant to be well diversified. The new definition closely matched the modern allocation model, which offered upside return and a possibility for downside protection.

However, as dramatically illustrated by the 2008–2009 investment market meltdown, when equities underperformed, some of the diversification benefits of the modern allocation model seemed to evaporate. For the fiscal year ending in June 2009, many of the endowment returns fell as much as 30 percent, compared to –14 percent for a traditional 60/40 portfolio. These losses were far larger than anticipated, especially for fully diversified portfolios that included a significant percentage of assets that were supposed to provide an absolute return.

With the most recent results in mind, is the modern allocation model just an illusion in which anticipated returns can be undermined by additional risk compared to the more traditional model? Or, are the model's benefits a function of the larger institutions' ability to get in early, but ultimately prove transitory as the investment herd subsequently crowds in, thus driving down returns? Or, did the endowment model experience a stress period unlikely to be repeated when markets return to normal?

This book focuses on the modern allocation model, first to understand the source of its value to investors, then to analytically examine its theoretical and actual behavior, and finally to reassess where and when it should be used given its benefits and limitations. It does so by adopting a new approach to evaluating the risk-and-return characteristics of standard and nonstandard asset classes and their use in portfolios. This analysis is based on the observation that, despite the inclusion in institutional portfolios of a variety of asset classes well outside the traditional stock-bond orbit, *U.S. equities continue to act as the overwhelmingly dominant risk factor for most institutional portfolios.*

Our analysis indicates that the modern allocation model does not fit the textbook definition of portfolio diversification, whereby additional asset classes are used to bring the portfolio closer to the efficient frontier by reducing risk or increasing expected return. Instead, because of the dominance of U.S. equities as a risk factor, the risk of additional asset classes and portfolios containing them can be described using a *beta with respect to equities*. Similarly, asset class and portfolio returns can be decomposed into a component associated with the underlying equity exposure and a *beyond-beta alpha* that is more uniquely associated with the asset class itself.

Building on this insight, the book develops analytical tools for evaluating institutional portfolios and applies those tools to develop a deeper understanding of the risk-and-return dimensions of these portfolios.

ASSET ALLOCATION

The modern allocation model follows in the general tradition of portfolio theory first developed in the 1950s (particularly, but certainly not restricted to Markowitz 1952, 1991; Sharpe 1963, 1964; Ross 1976), but it poses certain challenges as well. The Capital Asset Pricing Model or CAPM (Sharpe 1964) initially focused primarily on individual securities and implications for portfolio formation. If the expected return of a single security is given by $E(R_i)$, the return of the market by $E(R_e)$, and the risk-free rate by r_f , then

$$E(R_i) = r_f + \beta_i[E(R_e) - r_f]$$

In this formulation, the CAPM says that, in equilibrium, idiosyncratic risk will be diversified away, leaving only the systematic effect of the market tempered by how much the individual security responds to market movements (that is, its beta, β_i , or the covariance of the individual security return and the market return divided by the market return variance.²

Since the relationship between the market return and the individual security return is linear, it is not surprising that subsequent practical and academic developments pointed out that the CAPM could be scaled up to the portfolio level and hence to asset selection and allocation (Brinson et al., 1991) with the following form:

$$E(R_p) = r_f + [E(R_e) - r_f] \sum_{i=1}^n \omega_i \beta_i$$

in which R_p is the portfolio return and ω_i and β_i are, respectively, the portfolio weight and market-related beta of the i th asset class.³ Now portfolio formulation is done at the level of asset classes rather than individual securities, leaving security selection to a subsequent step. There is a large literature of helpful and appropriate approaches to portfolio construction in response to asset-class allocation, for example (Ross and Roll 1984; Campbell and Viceira 2002).

There are a couple of well-known implications of this simple characterization of asset allocation. First, it is the foundation of mean-variance optimization, which is the search for the set of portfolios—the efficient frontier—in which each portfolio maximizes return for a specified level of risk.

Second, asset allocation requires that asset classes under consideration are well defined, behave in predictable ways, and are widely available.

Inputs—expected future asset returns, volatilities, and covariances—must be specified in advance in order to do a formal static or dynamic mean-variance optimization.

The input estimation challenge is ubiquitous and ever-present, but its problems are magnified when using newer, nonstandard asset classes. Unlike traditional asset classes with century-long (or more) return histories (Ibbotson 2004; Dimson et al., 2002), most nonstandard asset classes do not enjoy a lengthy history of well-documented returns. The Goldman Sachs Commodities Index was created in 1991, but subsequently estimated back to 1970. The historical performance data on hedge funds, which now number over 8,000 in a wide variety of strategies, suffers from survivor bias, self-reporting, portfolio illiquidity, return backfilling, and secular versus cyclical trends in returns (Lo 2005; Schneeweis and Pescatore 1999; Rhodes-Kropf et al., 2004). Similar issues are said to affect published venture capital (Jones and Rhodes-Kropf 2002) and private equity returns (and in these latter cases, long holding periods and the use of IRR-based returns can result in return smoothing). Real estate performance suffers from all these problems and is notoriously unreliable as a guide to the future.

Modern allocation modeling can also be challenged by other characteristics of nonstandard assets, including asymmetric and fat-tail distributions, returns that are relatively more dependent on manager skill, illiquidity effects, and evolving return distributions. (In a subsequent chapter, we call these and other challenges that are exogenous to the model itself examples of *dragon risk*.)

The question for our purpose is not whether a given nonstandard asset is attractive or unattractive. Rather, the issue is how to value its relative attractiveness in a specific portfolio context in light of the special risks involved.

Furthermore, the increase in the number of asset classes considered in any portfolio formulation exercise can affect the confidence we have in the results. With only a few asset classes, especially when there is considerable confidence about expected returns, portfolio optimization modeling is fairly straightforward and stable. A change in the inputs that describe an asset will have a fairly predictable effect on the proportion of the portfolio assigned to that asset, so stress-testing an allocation model in those circumstances can proceed in a straightforward manner. On the other hand, as the number of asset classes proliferates, results become unstable and harder to anticipate. The inclusion or exclusion of a single asset class or change in that asset's inputs could have unpredictable or nonintuitive effects on the allocation assignments, not just for one asset class but for many.

For example, consider the following set of asset classes in Exhibit 1.3. Some, such as U.S. equities, U.S. bonds, and cash, are the standard assets

EXHIBIT 1.3 Asset Class Expected Real Return and Volatility

	Return	Volatility (Sigma)	Correlation with U.S. Equity
U.S. Equity	7.25	16.50	1
International Equity	7.25	19.50	0.65
Emerging Mkt Equity	9.25	28.00	0.45
Absolute Return	5.25	9.25	0.5
Equity Hedge Funds	5.75	12.75	0.85
Venture Capital	12.25	27.75	0.35
Private Equity	10.25	23.00	0.7
REITS	6.50	14.50	0.55
Real Estate	5.50	12.00	0.1
Commodities	5.25	19.00	−0.25
U.S. Bonds Govt	3.50	7.00	0.35
U.S. Bonds All	3.75	7.50	0.3
U.S. Bonds TIPS	3.25	6.50	0.35
Cash	1.50	2.00	0.35

Source: Morgan Stanley Research

used in what is often called *traditional asset allocation*, while others, such as hedge funds, private equity, and real estate, are considered to be nonstandard assets that are added in modern asset allocation.

These inputs—expected returns, sigmas, and an accompanying covariance matrix for the full set of standard and nonstandard assets—are supplied by a consulting firm that is deeply involved in institutional asset allocation.

Using these inputs, Exhibit 1.4 illustrates the challenges associated with modern asset allocation using nonstandard classes. The simplest allocation, Portfolio A, uses an allocation of just two standard assets, U.S. equities and cash. This point lies on the efficient frontier with an expected volatility (sigma) of 9.90 percent and an expected overall return of 4.95 percent. The Sharpe ratio for this portfolio is 0.35.

Portfolio B is similar to Portfolio A, except that U.S. bonds are substituted for the cash allocation. The resulting 60/40 mix is often referred to as the *traditional allocation*. Again, the results are plausible and unsurprising: the expected return rises to 5.85 percent, with an expected volatility (sigma) of 11.17 percent and a slightly improved Sharpe ratio of 0.39.

When we turn to the portfolio B', things begin to change. For this portfolio, we allow the optimizer to include venture capital along with equities, bonds, and cash. At a preselected expected risk level chosen to

EXHIBIT 1.4 Portfolio Risk-and-Return Characteristics

	A	B	B'	C	C'	C''
U.S. Equity	60	60	19	20	0	0
International Equity	▲	▲	▲	15	0	0
Emerging Mkt Equity	▲	▲	▲	5	10	16
Absolute Return	▲	▲	▲	10	0	0
Equity Hedge Funds	▲	▲	▲	▲	0	0
Venture Capital	▲	▲	35	10	22	▲
Private Equity	▲	▲	▲	10	8	31
REITS	▲	▲	▲	▲	30	20
Real Estate	▲	▲	▲	10	22	21
Commodities	▲	▲	▲	▲	8	12
U.S. Bonds All	▲	40	46	20	0	0
Cash	40	▲	▲	▲	0	0
Expected Return	4.95	5.85	7.37	7.08	8.07	7.73
Standard Deviation	9.90	11.17	11.17	10.83	10.83	10.83
Sharpe Ratio	0.35	0.39	0.53	0.52	0.61	0.58

Source: Morgan Stanley Research

match the sigma of portfolio B, the resulting unconstrained allocation seems nonsensical, to say the least. Portfolio B' allocates a whopping 35 percent to the nonstandard asset. This result is clearly unacceptable in the real world, even though the expected return and Sharpe ratio represent significant improvements over the first two portfolios.

With the trend toward using a broader array of asset classes in an attempt to diversify risk, a number of institutions have embraced the more modern type of allocation represented by Portfolio C. In Portfolio C, the direct exposure to U.S. equities is reduced to only 20 percent. There is a significant 15 percent exposure to international equities, as well as a 5 percent exposure to emerging markets. Absolute return, reflecting certain categories of hedge funds, amounts to 10 percent. This 10 percent weighting is also applied to venture capital, private equity, and real estate. As with equities, bonds have been reduced to 20 percent, far lower than in traditional portfolios. Given the assumptions contained in the covariance matrix, Portfolio C turns out to have a volatility of 10.83 percent and an expected return of about 7.08 percent, surprisingly close to the 11.17 percent volatility of the traditional Portfolio B. Portfolio C appears more diversified in regard to asset classes and sources of return—but in fact, is not really much different from the traditional Portfolio B in regard to this form of risk.

For the fifth portfolio, Portfolio C', we throw out all the constraints and throw in the kitchen sink, that is, all of the available asset classes. We then pick the point on the efficient frontier where the sigma matches the volatility level of Portfolio C. Note that this unconstrained optimization produces an allocation with *all* nonstandard assets and *no* standard assets. Most of the portfolio statistics are highly attractive—compared to Portfolio C, expected return rises by more than a percentage point and the Sharpe ratio improves by nearly 0.10. However, the elimination of standard assets from the mix is counterintuitive, to say the least.

Finally, Portfolio C'' is similar to Portfolio C', except that it removes venture capital from the mix. Once again, the unconstrained optimizer assigns no weight to standard assets. Even more, changes to the asset weights relative to Portfolio C' are difficult to understand without deeper analysis. In sum, neither B', C', nor C'' would make sense in most institutional settings, nor would they be easy to explain to an investment committee.

As we will see in subsequent chapters, one way to deal with the challenges of asset-class inputs, portfolio stability, and other issues is to torture the modeling process by imposing piecemeal constraints, adjusting inputs, and other fixes that make the resulting portfolios more palatable. It would be far more satisfying and understandable to find a simple, transparent approach to the problem, one that reveals the critical risk and return characteristics of the underlying assets as well as the portfolio itself. For many institutional portfolios, the total beta approach can provide greater clarity as well as a more intuitive perspective on the most important determinants of total portfolio risk.

BETA-BASED RISK AND RETURN: THE SIGMA AND BETA LINES

Returning to the standard CAPM model, the third and most important implication is that the portfolio return depends on the sum of the various asset class betas scaled by the asset class weights.

What follows from this notion is that the beta sensitivity to equities is the parameter that captures about 90 percent or more of the volatility risk for most allocations seen in the U.S. institutional market. This single parameter is a value that lurks hidden within virtually every asset class, and that, in aggregate, accumulates to become the portfolio's overall exposure to the equity market. Once these underlying beta values are uncovered, it becomes clear that while the traditional 60/40 appears quite different from the highly diversified endowment model, they in fact share certain common risk characteristics. This analysis suggests that portfolio betas can




EXHIBIT 1.5 Expected Returns and the Cash-Equity Sigma Line

Data Source: Cambridge Associates Data

be used, within limits, to determine the likelihood of adverse events that fundamentally set the risk limits for a wide range of the asset allocations seen in practice.

A simple way to understand these conclusions is to start with the more familiar volatility measure and then see how introducing beta can add to our understanding. Exhibit 1.5 uses the asset class inputs from Exhibit 1.4 to display each asset's expected return relative to a sigma line that intersects the cash and U.S. equity returns.

The majority of asset risk-and-return dots lie near or on the line, indicating that their return assumptions appear to be in direct linear proportion to their relative expected volatility.⁴

At the heart of this book lies the proposition that by looking at asset class and portfolio risk in a different way—through the beta lens—we can simplify the analysis of portfolio risk and find hidden meaning that will improve our understanding of complex asset class and portfolio behavior and help us build better tools for asset allocation and portfolio construction. It can also help us better measure the sources of volatility, how risk relates to return, and the decomposition of asset class and portfolio return in their fundamental building blocks.

To begin this journey, we can review the beta risk measure and compare it to the total volatility measure. In the most basic terms, beta indicates how much we would expect a particular asset to move in response to a 1 percent

change in the overall equity market. One way to think about it is that it is the coefficient of the regression of asset class returns on the equity market.⁵

An equivalent way to describe beta is that it is the correlation between the asset (or portfolio) return and the market return, multiplied by the ratio of their volatilities,

$$\beta_p = \frac{\text{Cov } p, e}{\sigma_e^2} = \rho_{pe} \frac{\sigma_p}{\sigma_e}$$

where p is still the individual asset class or investment portfolio, e is the market portfolio (or an equivalent), and ρ_{pe} is the correlation between the asset class/investment portfolio and the market portfolio.⁶

In this formulation, we can see that, other things being equal, the higher the correlation between the asset class (or portfolio) returns and market returns, the higher the beta. But in thinking about beta, we should not neglect the other terms. Like the correlation, the relationship between portfolio or asset class volatility and beta is linear and positive. Not so with *market* volatility, however, which has a powerful inverse and nonlinear effect on beta. As market volatility rises, beta will fall, all other things being equal (which itself would probably be a rare event!).⁷

So what does this mean in practical terms? First, in our terms, each asset class beta is a measure of that asset's risk with respect to domestic equities. Second, a widely cited advantage of nonstandard assets is their supposedly low correlations and betas with respect to traditional equities. Whether or not this is true in all cases, we can see that a low beta can result from three nonexclusive conditions: (1) low correlation between an asset class and the market; (2) low asset class volatility; or (3) high equity market volatility, or any combination of 1, 2, or 3. So, we are reminded that asset class beta and correlation are not the same thing, and the respective volatilities also play a major role. Thus, an asset class may have a low correlation with U.S. equity, but still have a relatively significant beta sensitivity.

The beta measure is useful for a number of reasons. First, it incorporates both the correlation and volatility effects described earlier. Unlike sigma, betas can be related to a common risk factor such as the S&P 500 and therefore become comparable, scalable, and additive. Moreover, within the portfolio context, it is the total beta that is typically the overwhelmingly dominant source of risk. As long as the residuals are not themselves highly correlated, the total portfolio beta is simply a weighted sum of the individual asset class or security betas. Thus, an asset's beta provides a far more accurate measure of its contribution to the overall portfolio risk than the asset's volatility alone.

EXHIBIT 1.6 Structural Betas and Alphas

	Return	Sigma	Beta	Alpha
U.S. Equity	7.25	16.50	1.00	0.00
International Equity	7.25	19.50	0.77	1.39
Emerging Mkt Equity	9.25	28.00	0.76	3.42
Absolute Return	5.25	9.25	0.28	2.32
Equity Hedge Funds	5.75	12.75	0.66	0.56
Venture Capital	12.25	27.75	0.59	7.47
Private Equity	10.25	23.00	0.98	3.15
REITS	6.50	14.50	0.48	2.35
Real Estate	5.50	12.00	0.07	3.82
Commodities	5.25	19.00	-0.29	5.73
U.S. Bonds Govt	3.50	7.00	0.15	1.36
U.S. Bonds All	3.75	7.50	0.14	1.69
U.S. Bonds TIPS	3.25	6.50	0.14	1.18
Cash	1.50	2.00	0.04	0.00

Data Source: Cambridge Associates Data

Exhibit 1.6 now uses the market assumptions in Exhibit 1.3 but displays the asset class betas and beyond-beta alpha returns.

With these derived betas and the associated returns, Exhibit 1.7 replaces Exhibit 1.5's sigma volatility with a beta risk measure on the horizontal axis and displays a *beta line* intersecting cash and U.S. equities.

In contrast to the sigma cash-equity line, most of the asset class returns now lie above the beta line. As such, expected total return consists of three components as illustrated in the exhibit by the decomposition of the expected return assumed for venture capital (12.25 percent). First, all asset classes build on the foundation represented by the risk-free rate ($r_f = 1.50$ percent), the cash-equivalent return. The second return component (3.27 percent)—the return between the risk-free return line and the beta line—is a direct linear function of an asset class's beta and could, in theory, be replicated by a combination of equities and cash.

The third component is the return in excess of the beta line (7.47 percent). This last component, which we call beta-based or *structural alpha*, is a function of the unique risk-and-return profile of the individual asset class. The three return components sum to the total return for the asset class (1.50 percent + 3.27 percent + 7.47 percent = 12.25 percent). And, as we will see in the next chapter, each component has its own associated risk.




EXHIBIT 1.7 The Beta Line and Structural Alphas

Data Source: Cambridge Associates Data

The remainder of this book builds on these simple concepts of a structural beta and beta-based alpha to illuminate the endowment model and to handle the inclusion of newer forms of asset classes in institutional portfolio construction. First, we examine the nature and use of structural beta and alpha, where it becomes clear that many portfolios that appear quite different on the surface in fact share common risk characteristics. The consequence is that many portfolios containing nonstandard assets seem to be designed to pursue extra return rather than reduced risk. In contrast, we find that identifying and using structural alphas to design portfolios can enhance returns with only modest levels of additional risk at the fund level. These hidden sources of excess return can be uncovered by analyzing the fundamental risk structure of standard return assumptions.

We then turn more directly to the role of structural alpha and beta in asset allocation and portfolio construction. By reversing the usual asset allocation process in which traditional assets form the base, an *alpha core* of nonstandard assets can be used along with supplementary traditional *swing* assets to better relate a portfolio's risk to its expected return.

Drawing on our theoretical analysis of beta and alpha-based asset allocation, we then consider actual portfolio behavior in selected regimes. Of particular interest are the implications of *stress betas* during significant market declines, when betas can rise significantly and portfolios can fail, at least for a time, to provide the expected risk protection.

In a concluding chapter, we offer suggestions for the future of the endowment model in light of our theoretical and applied findings. The modern endowment model is not a magic potion that will smooth returns and lower short-term volatility, but rather a strategy for accumulating incremental returns and achieving more divergent outcomes—over the long term. One critical implication is that the endowment model should continue to be an attractive option for long term investors that are able to ride out periods of significant short term volatility. Another is that institutional leaders and individual investors should not be lulled by the prospect of incremental excess returns into thinking that these will be available during each and every investment period. A third implication is the unexpected challenge for any modern portfolio of obtaining excess returns above the risk-free rate. Finally, a deeper and more exact analysis of the sources of portfolio risk and return can assist institutions in constructing and managing assets in a way that will better reflect their investment objectives and related needs.

NOTES

1. We use college and university endowments as an example and recognize that they are a subset of the world of nonprofit institutional asset management, including independent foundations, public pensions, and, arguably, private pensions. Unless specified, all annual figures here refer to fiscal years rather than calendar years.
2. Or $\beta_i = \frac{\text{cov}(R_i, R_e)}{\sigma_e^2}$.
3. At the time, this shift in the level of analysis was a significant step (and is still being debated today (Ibbotson and Kaplan 2000; Kritzman and Page 2003)).
4. Notable exceptions are commodities and emerging market equities, both of which lie below the line. These are the result of the inputs, which were supplied by an independent source.
5. In formal terms, $y_i = \alpha_i + \beta_i x_e$, where y_i is the return of an individual security or asset class, β_i is the coefficient that describes the effect of the market return, x_m , on the individual security or asset class return, and α_i is the excess return of the security or asset class over the market portfolio. For a portfolio of securities or asset classes, $y_p = \sum_{i=1}^n \alpha_i \omega_i + x_m \sum_{i=1}^n \beta_i \omega_i$, in which ω_i is the weight of the i th security or asset class in the portfolio.
6. In recent years, the term *beta* has also been used to refer to gaining exposure to individual asset classes of all types through index-like portfolios, such as how well a fixed-income fund represents the behavior of the overall bond market. Moreover, in theory, the market can be defined most broadly as the sum of all available world investment opportunities (or more narrowly as all of a country's investable securities). For this analysis, the reference market is U.S. equity.

7. In other words, a change in beta with respect to correlation is given as $\frac{d\beta_p}{d\rho_{pe}} = \frac{\sigma_p}{\sigma_e}$, while a change in beta with respect to sigma is given as $\frac{d\beta_p}{d\sigma_e} = -\frac{\rho_{pe}\sigma_p}{\sigma_e^2}$.

REFERENCES

- Bernstein, P. L. 2003. "Which policy do you mean?" *Economics and portfolio strategy*. New York: Peter L. Bernstein, Inc.
- Brinson, G. P., B. D. Singer, and G. L. Beebower. 1991. "Determinants of portfolio performance II: An update." *Financial Analysts Journal* 47 (3): 40–48.
- Campbell, J. Y., and L. M. Viceira. 2002. *Strategic asset allocation: Portfolio choice for long-term investors*. New York: Oxford University Press.
- Dimson, E., P. Marsh, and M. Staunton. 2002. *Triumph of the optimists: 101 years of global investment return*. Princeton, NJ: Princeton University Press.
- Ibbotson, R. 2004. *Stocks, bonds, bills and inflation yearbook: 1983*. Chicago: Ibbotson Associates.
- Ibbotson, R., and P. Kaplan. 2000. "Does asset allocation explain 40, 90 or 100 percent of performance?" *Financial Analysts Journal* 56 (1): 26–33.
- Jones, C., and M. Rhodes-Kropf. 2002. *The price of diversifiable risk in venture capital and private equity*. Columbia University Working Paper.
- Kritzman, M., and S. Page. 2003. "The hierarchy of investment choice." *Journal of Portfolio Management* 29 (4): 11–23.
- Leibowitz, M. L., and P. B. Hammond. 2004. "The changing mosaic of investment patterns." *Journal of Portfolio Management* 30 (3): 10–25.
- Lo, A. 2005. *The dynamics of the hedge fund industry*. Charlottesville, VA: Research Foundation of the CFA Institute.
- Markowitz, H. M. 1952. "Portfolio selection." *Journal of Finance* 7 (1): 77–91.
- . 1991. "Foundations of portfolio theory." *Journal of Finance* 46 (2): 469–477.
- National Endowment Surveys (NES). National Association of College and University Business Officers 1992–2008.
- Rhodes-Kropf, M., A. Ang, and R. Zhao. 2004. *Do funds of funds deserve their extra fees?* Columbia University Working Paper.
- Ross, S. A. 1976. "The arbitrage theory of capital asset pricing." *Journal of Economic Theory* 13: 341–360.
- Ross, S. A., and R. Roll. 1984. "The arbitrage pricing theory approach to strategic portfolio planning." *Financial Analysts Journal*, May–June: 14–26.
- Schneeweis, T., and J. F. Pescatore. 1999. *The handbook of alternative investment strategies*. New York: Institutional Investor, Inc.
- Sharpe, W. F. 1963. "A simplified model for portfolio analysis." *Management Science* 9 (2): 277–293.
- . 1964. "Capital asset prices: A theory of market equilibrium under conditions of risk." *Journal of Finance* 19 (3): 425–442.
- Swenson, D. F. Revised 2009. *Pioneering portfolio management: An unconventional approach to institutional investment*. New York: The Free Press.