DISCUSSION PAPER PI-1008

Do Endowment Funds Select the Optimal Mix of Active and Passive Risk

Keith C. Brown and Cristian Tiu

October 2009

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University
106 Bunhill Row London
EC1Y 8TZ
UNITED KINGDOM

http://www.pensions-institute.org/

Do Endowment Funds Select the Optimal Mix of Active and Passive Risk?*

Keith C. Brown† Cristian Tiu‡

First draft: March 10, 2009
Current draft: October 7, 2009

Forthcoming in the *Journal of Investment Management*

*We thank Jessica Shedd and the National Association of College and University Business Officers for furnishing much of the data required by this study. We are also especially grateful for the contributions of Lorenzo Garlappi and the comments of several other colleagues, including Steve Dimmock, Elroy Dimson, John Griffin, Bing Han, Jennifer Huang, Bruce Lehmann, Bing Liang, Jeff Kubik, Clemens Sialm, Laura Starks, Sheridan Titman, Raman Uppal, Bas Werker, and Roberto Wessels. We have also benefitted from conversations with a number of people involved in the endowment fund industry: Bob Boldt and Andrea Reed (Perella Weinberg), Gary Hill, Cathy Iberg and Uzi Yoeli (University of Texas Investment Management Company), John Griswold (Commonfund), Larry Siegel (Ford Foundation), Jim Hille (TCU Endowment Fund) and Chance Burhance (International Fund Services). Finally, the comments of an anonymous referee have improved the exposition in the study greatly. All errors in the paper are solely our responsibility.

†McCombs School of Business, B6600, The University of Texas at Austin, Austin, TX, 78712. Email: kcbrown@mail.utexas.edu

‡Jacobs Management Center, University at Buffalo, Buffalo, NY 14260. Email: ctiu@buffalo.edu

Do Endowment Funds Select the Optimal Mix of Active and Passive Risk?

Abstract

The investment decision confronting managers of multi-asset class portfolios can be characterized in terms of the passive (i.e., benchmark, or policy) and active (i.e., market timing and security selection) strategies they adopt. In this paper, we investigate whether managers select the appropriate combination of active and passive allocations in their portfolios. Noting that this issue is ultimately a risk management question, we adapt and extend a simple framework for establishing what constitutes the optimal level of active and passive risk exposure. We then examine the question empirically using a database consisting of the allocation decisions and investment performance of a large set of university endowment funds over the period from 1989 to 2005. Our findings show that (i) the average endowment had too little active risk exposure in its portfolio, (ii) endowment funds could have significantly increased their risk-adjusted performance by enhancing the scale of the alpha-generating strategies they were already employing, and (iii) this tendency to under-utilize active management skills was more pronounced for larger endowments than for smaller ones. We conclude that the typical endowment fund could have improved its performance by increasing the commitment to its active management skills.

JEL Classification Codes: G11; G23

Keywords: Risk budgeting, Active-passive allocations, Investment performance, Endowment funds.

Do Endowment Funds Select the Optimal Mix of Active and Passive Risk?

1 Introduction

It is well established that the performance of an investment portfolio that is diversified across multiple asset classes can be thought of as being driven three distinct decisions that its manager makes: (i) long-term (i.e., strategic, or policy) asset allocations; (ii) temporary adjustments (i.e., tactical) to these strategic allocations in response to current market conditions; and (iii) the choice of a particular set of holdings to implement the investment in each asset class (i.e., security selection). The first of these performance components is commonly referred to as the passive portion of the portfolio while the latter two collectively represent the active positions the manager adopts. Trying to understand and disentangle the issues defining the myriad dimensions of the relationship that exist between these passive and active performance components has become, perhaps, the preeminent challenge confronting the portfolio management industry.

In this study, we extend this "active-passive" discussion by considering the following question: Given that both active and passive risks play a role in determining the overall performance of a portfolio, is it possible to establish what the optimal combination of these components should be? To address this problem, we create a simple analytical framework that extends the pioneering work of Treynor and Black (1973) (hereafter T-B) to a multi-asset class setting. Specifically, after showing how portfolio returns can be decomposed along the lines of the three managerial decisions outlined above, we adopt the central T-B notion that the active and passive portions of any collection of securities can themselves be viewed as separate assets that can be combined in an optimal manner. Assuming an appropriate risk-adjusted decision rule—the reward-to-variability and reward-to-downside risk ratios, in our case—then allows us to calculate the optimal passive and active allocations in the portfolio by finding the levels of those components that maximize the decision metric.

To demonstrate the insights gained by this optimization process, we analyze the investment performance of a sample of college and university endowments funds over the period from 1989 to 2005. Several salient characteristics of the endowment fund environment make these investment vehicles particularly well-suited to this purpose, including the fact that they have virtually infinite time horizons, they tend to invest across as many as a dozen different asset classes, and they have shown a historical commitment to pursuing both passive and active strategies (see, for instance, Swensen (2009)). Making use of this database allows us to pose two specific queries. First, based on an in-sample analysis of the entire database, did the level of active risk that the average endowment fund manager assumed constitute an optimal level of exposure? Second, using an out-of-sample analysis that splits the database in half chronologically, could the typical endowment manager have improved performance in the second half of the period by altering the portfolio's active-passive risk combination based on information revealed in the first half?

Our empirical analysis strongly suggests that the answers to these questions are "no" and "yes," respectively. That is, over the entire sample period, we document that the average endowment fund was over-allocated to the passive component in the portfolio by as much as 30 percent, meaning that it severely under-emphasized the active management skills that were already in place. Further, based on our out-of-sample analysis, we demonstrate that if the typical endowment manager had adjusted the portfolio's passive and active risk allocations based on the findings from the first half of the sample period, the risk-adjusted performance during the second sub-period would have improved by about 20 percent. We stress, however, that these reallocations from passive to active risk exposure do not necessarily require the average endowment to increase the net level of its transactions or take on a significantly greater level of total risk—particularly if the active and passive components are negatively correlated—nor does it suggest an increased proportional exposure to asset classes that are perceived to be riskier (or otherwise objectionable), such as hedge funds or private equity.

To assess the robustness of these findings, we replicated the analysis on various stratifications of our sample based on cross-sectional differences in characteristics that have been shown to matter for endowment funds. Specifically, following Brown, Garlappi, and Tiu (2009), we divide the fund sample into extreme values of assets under management and spending ratio, as well as distinguish those funds associated with public versus private academic institutions. The results indicate that while some of these characteristics appear to significantly influence the active-passive allocation decision, others do not. We show that the size of the fund makes an appreciable difference, with the optimal active risk allocation for the largest funds being almost double what they actually hold whereas small funds are shown to be exploiting their active management talents more completely. On the other hand, cross-sectional variations in endowment spending rates or the profile of the sponsoring university had no material impact on the findings for the overall sample. Overall, we conclude that the typical endowment fund—and particularly the typical large fund—could have improved its performance by taking more active risk while keeping its policy allocation unaltered.

The remainder of the paper is organized as follows. In the next section, we review the relevant literature as well as develop our simple framework for assessing the optimality of a manager's active-passive allocation decision. In Section 3, we describe our endowment fund database while Section 4 presents our main empirical findings on both a full-sample and cross-sectional basis. Section 5 then addresses several normative implications concerning the ability of investors to implement this framework in practice while Section 6 concludes the study.

2 Finding the optimal mix of active and passive risk

2.1 A brief overview of the risk budgeting process

One of the most important recent trends in the management of investment portfolios has been the development of formal procedures to help managers determine the appropriate combination of risk exposures that they should be holding. In particular, in what has come to be referred to as the risk budgeting process, considerable attention has been paid to establishing the conceptually correct mixture of passive (i.e., systematic) risk and active (i.e., unsystematic or residual) risk that defines the portfolio. Indeed, as characterized by Chow and Kritzman (2001), an investor's risk budget can ultimately be viewed as a means of converting her asset allocation decisions—expressed in monetary terms—into these risk assignments.[1]

Although they differ in their specific details, a common feature of the various extant risk budgeting methodologies is that the optimal combination of active and passive risk is established as the one that maximizes the total risk-adjusted expected returns for the portfolio. In this context, Clarke, de Silva, and Wander (2002) make the crucial observation that while the total investment return is simply the sum of the returns generated by the active and passive parts of the portfolio, the portfolio's total risk is not additive but rather depends on (i) the volatility of passive returns,

[1] The list of studies representing the risk budgeting literature is lengthy and space limitations prohibit providing an exhaustive catalogue here. However, Pearson (2002) offers an excellent introduction and overview of the topic.

(ii) the volatility of active returns, and (iii) the correlation that exists between the passive and active return components. While some risk allocation models ignore this interaction term (e.g., Waring and Siegel (2003) cite historical data to support their assumption that these correlations are insignificantly different from zero), others take it into account. For instance, Berkelaar, Kobor, and Tsumagari (2006) allow for non-zero active risk interactions in the development of their correlation-adjusted method for maximizing a portfolio's risk-adjusted performance.

A notable feature that appears to be pervasive throughout the risk budgeting literature is that there does indeed seem to be a role for taking active risk in the portfolio. This is not a trivial outcome given that, in equilibrium, the overall contribution from active management should be a zero-sum game.[2] However, the fact that the net expected return to the active risk-taking is zero does not mean that all active managers will fail to produce a positive alpha (i.e., the difference between the realized and the expected return). Thus, as Winkelman (2003) and Leibowitz (2005) note, the case for active portfolio risk rests on markets not always being in equilibrium and the investor's ability to identify genuinely skillful active managers in advance. With respect to the latter condition, Grinold (1989) and Grinold and Kahn (2000) have established a useful framework for assessing the expected risk-adjusted contribution of active managers through the skill they are perceived to possess and the breadth with which they can deploy those skills.[3]

Of course, such justifications for including active risk elements in an investment portfolio are not new. Often lost in the risk budgeting discussion is the fact that Treynor and Black (1973) addressed this topic almost four decades ago. Nominally an analysis of the role that security analysis can play in markets that are almost efficient, the T-B framework provides a compelling way to think about how passive and active allocations should be pooled. The essential insight in the T-B analysis is that the optimal combination of the active portfolio—which results from the application of security analysis to identify a limited number of undervalued assets—and a passive benchmark portfolio is itself a straightforward portfolio optimization problem. That is, T-B treats the active and passive portions of an investment portfolio as two separate "assets" and then calculates the mix of those

[2]For instance, both the capital asset pricing model of Sharpe (1964) and the arbitrage pricing model of Ross (1976) are well-known general models of risk and expected return in which systematic risk is rewarded but unsystematic risk is not.

[3]This relationship between risk-adjusted performance, manager skill, and the breadth of the investment opportunities is called "the law of active management," and it has been the basis for substantial additional analysis; see, for example, Qian and Hua (2004) and Clarke, de Silva, and Thorley (2006). Also, Alford, Jones, and Winkelmann (2003) as well as Harlow and Brown (2006) have explored the empirical benefits and limitations of attempting to identify superior active fund managers.

assets that maximizes the reward-to-variability (i.e., Sharpe) ratio. It is then demonstrated that the investment allocation assigned to the active portfolio strategy increases with the level of alpha it is expected to produce (i.e., the active "benefit"), but decreases with the degree of unsystematic risk it imposes on the investment process (i.e., the active "cost"). This is an important insight because it suggests that taking more active risk in a portfolio will not necessarily lead to an increase in total risk; if, for instance, the manager's active investment is negatively correlated with the passive component (e.g., an effective short position in an industry or sector benchmark) the overall risk in the portfolio might actually decline.

For all of its elegance, Kane, Kim, and White (2003) have noted that the T-B model has had a surprisingly low level of impact on the finance profession in the years since its publication. They attribute this neglect to the difficulty that investors have in forecasting active manager alphas with sufficient precision to use the T-B methodology as a means of establishing meaningful active and passive portfolio weights on an ex ante basis. However, this begs the question of whether the model offers a useful way of assessing ex post whether the proper active-passive allocation strategy was adopted by the portfolio manager. Said differently, using the T-B model, did the manager affect an appropriate combination of active and passive exposures, given the alphas that were actually produced? In the subsequent sections, we provide just such an empirical analysis using our sample of endowment funds.

2.2 Decomposing portfolio returns

In order to address the issue of whether fund managers deploy the various risks in their portfolio in an optimal manner, it is first necessary to split the returns they produce into their passive and active components. We follow a methodology similar to Brinson, Hood, and Beebower (1986) and Brinson, Singer, and Beebower (1991) and decompose the returns of a managed portfolio into their three fundamental components: (i) strategic asset allocation policy (i.e., benchmark), (ii) tactical allocation (i.e., market timing), and (iii) security selection. As noted earlier, this decomposition reflects closely the investment choices within a typical endowment fund; the first of these components represents a passive decision typically made by the endowment's Board while the latter two are active decisions made by the endowment's investment staff.[4]

[4]There is substantial survey evidence that decomposing returns in this manner is, in fact, a reasonable way to characterize the endowment management process. For instance, Griswold (2008) shows that endowments invest

Formally, let $R_{i,t}$ be the total realized return on fund i at the end of period t, $w_{i,j,t}$ the actual portfolio weight of fund i in asset class j at the end of period t and $r_{i,j,t}$ the period-t return on asset class j. The realized return can then be expressed as

$$R_{i,t} = \sum_{j=1}^{N} w_{i,j,t-1}\ r_{i,j,t}, \tag{1}$$

where N denotes the set of investable asset classes.

The *policy allocation* return of a fund is the part of the total return that can be achieved by passively applying the strategic policy asset allocation weights to the benchmark indices for each asset class. In other words, no skill in timing exposures to the various asset classes or selecting securities within each class is required to achieve this return. The benchmark return for fund i at the end of period t is formally defined as follows

$$R^B_{i,t} = \sum_{j=1}^{N} w^B_{i,j,t-1} r^B_{j,t} \tag{2}$$

where $w^B_{i,j,t-1}$ denotes the asset allocation policy weight in asset class j at the end of period $t-1$ and $r^B_{j,t}$ is the period t buy-and-hold return on the benchmark index for asset class j. Note that determining the level of (2) requires only knowledge of (i) the policy asset allocation weights and (ii) the benchmark index for each asset class.

We define the *market timing* component, the first of the two active return elements, as follows

$$R^T_{i,t} = \sum_{j=1}^{N} (w_{i,j,t-1} - w^B_{i,j,t-1}) r^B_{j,t}, \tag{3}$$

where, as before, $w_{i,j,t-1}$ denotes the actual weight in class j at the end of period $t-1$ and $w^B_{i,j,t-1}$ is the corresponding policy weight for the same investment horizon, decided one period in advance. As formulated, this expression captures the return that is achieved by over- or under-weighting the benchmark asset in an attempt to increase returns or reduce risk.

The last component of a portfolio's total return captures the effect of *security selection*, which is the result of the second active decision the portfolio manager can make. This is the excess return of

between 65-85 percent of the funds dedicated to a particular asset class in actively managed portfolios, with the remainder being passively indexed.

the managed portfolio in a given asset class over the hypothetical return achievable by an investor who allocates resources in the benchmarks according to the policy weights. Formally, we define the security selection return as

$$R_{i,t}^{S} = \sum_{j=1}^{N} w_{i,j,t-1}(r_{i,j,t} - r_{j,t}^{B}), \tag{4}$$

where $r_{i,j,t}$ is the return generated by a portfolio of securities chosen by the manager responsible for asset class j.

From the decompositions in (2), (3) and (4), it is clear that the portfolio's total return in any period t can be expressed as the sum of its passive and active components or

$$R_{i,t} = R_{i,t}^{B} + R_{i,t}^{T} + R_{i,t}^{S} \tag{5}$$

which, in turn, implies that the *total active* portion of the portfolio's return can be obtained as the difference between its overall return and the passive component

$$R_{i,t} - R_{i,t}^{B} = R_{i,t}^{T} + R_{i,t}^{S} = R_{i,t}^{A}. \tag{6}$$

Assuming that $R_{i,t}^{B}$ adequately expresses the portfolio's expected return during period t, then $R_{i,t}^{A}$ in (6) represents the manager's periodic realized alpha from his or her active decisions.[5]

2.3 Optimal active-passive allocation: A simple framework

Following the intuition of Treynor and Black (1973), we address the question of whether portfolio managers adopt an optimal active-passive risk allocation by seeing if they take full advantage of their alpha-generating capabilities or whether they "leave money on the table" by mixing their passive and active positions in a sub-optimal manner. In what follows below, we present a straightforward extension of the T-B model that allows us to assess this issue in the context of the multi-asset class problem faced by endowment fund managers who have the ability to make active decisions about broad market and sector exposures as well as for individual security positions.

[5] More recently, Lo (2008) has suggested an alternative method for decomposing active and passive portfolio returns based on the manager's ability to forecast *future* market trends. As described below, however, this method requires data that was not available for this study and hence that proposed method cannot be examined herein.

To formalize the evaluation process, assume as before that the total actual return R of a particular endowment contains a passive benchmark component R^B as defined in (2) and an active component R^A representing, without loss of generality, the collective impact of the timing and selection decisions the manager makes. In this simplification, the active component can be written as $R^A = R - R^B$, where the subscript denoting the investment period is suppressed for convenience. Notice that a fund can achieve exposure to its benchmark either by investing directly in the indices composing R^B or indirectly through the formation of the active portfolio that generates R. Consequently, we can always think of this active portfolio itself as being a (trivial) combination obtained by investing 100% in the assets that deliver R and 0% in the assets that delivers R^B.

It is important to realize, though, that this "all active" portfolio will have only *indirect* exposure to the benchmark through R. The crucial insight in this setting is that by rescaling the existing positions in R and R^B, the portfolio manager can construct an alternative portfolio that has the *same* monetary commitment to the benchmark assets but achieved with a *different* combination of asset class and security exposures. For example, consider the new return

$$R(\lambda) = \lambda R + (1 - \lambda)R^B. \tag{7}$$

The portfolio implied by the weighted return in (7) is the de facto result of a *swap* transaction between the existing active portfolio and the benchmark allocation. If, for instance, the actual investment weights in the original portfolio and in the benchmark are identical (i.e., in the absence of market timing), the resulting swap portfolio has the same asset allocation weights, but its exposure to the benchmark is achieved through different securities than those contained in the active portfolio delivering R.

To implement such an exchange, an endowment does not have to invest in asset classes with which it is not already familiar. To see this, consider the case of a portfolio comprising a single asset class, US domestic equities. Suppose that the manager initially holds \$100 million of US equity in Portfolio R and then decides to revise its position by choosing an allocation of 110% in R ($\lambda = 1.1$) while simultaneously shorting 10% of the benchmark. After this swap, the new portfolio will contain \$110 million of this new equity position and will be short \$10 million of the benchmark for US equity (e.g., the CRSP value-weighted portfolio, the Wilshire 5000 or the Russell 3000). The net overall US equity position is still \$100 million and so the exposure to this asset class is unaltered, although the actual securities held in the adjusted managed portfolio will be different. For the swap to be implementable in practice, it is important that it does not alter substantially the overall exposure to an asset class, since most institutional investors have policies limiting the variation of their actual portfolio weights around their benchmark weights (i.e., tactical ranges). Furthermore, the implementation of this swap does not necessarily require short selling the benchmark index; it merely requires the sale of 10% of a combination of securities in a portfolio that is close to the index while simultaneously investing proportionally the proceedings from this sale in those securities remaining in the portfolio. Hence, no new active management skills are required beyond those the manager already possesses.

Since by definition $R = R^B + R^A$, using (7) the passive component of the post-swap portfolio is $R^B(\lambda) = R^B$ while the active component is $R^A(\lambda) = \lambda R^A$. By extension, then, when $\lambda > 1$ the swap portfolio will have a higher emphasis on the active management component than did the manager's original portfolio. In other words, choosing how much emphasis is best placed on active risk is equivalent to choosing how to best rescale the existing managed portfolio R by swapping a fraction of it against the benchmark portfolio R^B. It is important to note that although the implementation of such a strategy requires the actual portfolio to be scalable, it does not require any additional alpha-generating abilities compared to the actual portfolio R.

With this background, the specific question we would like to consider can be stated as follows: For any particular endowment fund, is it possible to produce a portfolio with a more pronounced commitment to active risk that would be preferred to the portfolio currently held? If so, we can conclude that such an endowment is passing on a readily available opportunity to use its active management skills in a way that could potentially add value through the market timing or security selection return components.

To answer any such question concerning the optimality of the investment process, we first need to identify endowment preferences and then solve for the rescaling parameter λ that maximizes those preferences. We specify two different characterizations of a fund's objective function, both of which assume that the investor is best served by the portfolio position that maximizes risk-adjusted returns. In the first specification, we follow the T-B approach of assuming that endowments select an active and passive mix in order to maximize the reward-to-total variability (i.e., Sharpe) ratio of the portfolio, or:

$$\lambda = \arg\max \frac{E[R(\lambda) - R_f]}{\text{std}[R(\lambda) - R_f]} \tag{8}$$

where $E[\cdot]$ and std$[\cdot]$ are estimated by their sample counterparts.[6]

Of course, one challenge with using the objective function in (8) is that it is based on the implicit assumption the return standard deviation is an adequate characterization of risk. However, several of the asset classes frequently included in endowment fund portfolios are likely to produce returns that are asymmetrically distributed (e.g., hedge funds).[7] Accordingly, the second optimization problem we examine assumes that endowments select active and passive components by maximizing the reward-to-downside risk (i.e., Sortino) ratio, defined as:

$$\lambda = \arg\max \frac{E[R(\lambda) - R_f]}{DR(\lambda)} \tag{9}$$

where $DR(\lambda) = \sqrt{E\left[\max\{E[R(\lambda) - R_f] - (R(\lambda) - R_f), 0\}^2\right]}$ is the downside risk measure, also computed via its sample counterpart. This second criterion can also be justified as being relevant to a university endowment that is concerned with producing returns in the portfolio that meet or exceed its spending requirement. In this case, one can argue that the semi-deviation, or downside risk, is the more relevant measure of uncertainty.

3 Data description

Our primary database is the set of endowment studies produced by the National Association of College and University Business Officers (NACUBO), which are annual publications based on surveys that gather information about net-of-expense return performance, asset allocation patterns, spending rules and rates, and manager and custodial relationships of college and university endowments throughout the United States, Canada, and Puerto Rico. The basic set of data available to us from this source covers the period from 1984 to 2005.[8]

[6] Specifically, for a given λ we construct a time series of the swap portfolio returns $R_t(\lambda) = \lambda R_t + (1-\lambda)R_t^B$. $E[\cdot]$ and std$[\cdot]$ are the time series mean and standard deviation of $R_t(\lambda)$ in excess of the risk-free rate R_f.

[7] An excellent discussion of this topic can be found in Ineichen (2007). Also, Harlow (1991) has demonstrated the potential problems associated with asset allocation process when the potential for asymmetric returns is ignored.

[8] TIAA-CREF has administered the survey since 2000; from 1988 to 1999, the survey was conducted in partnership with Cambridge Associates and before 1988 by the NACUBO Investment Committee.

Although the NACUBO surveys began in 1984, the participating institutions were not identified during the 1984-1988 period meaning that, for these five years, the asset allocation survey data cannot be merged with the information on assets under management, endowment fund payout or investment performance. This in turn means that the two components of active returns described above cannot be measured accurately and, as a consequence, the majority of our analysis will be limited to the post-1988 period. For the 1989-2005 time frame, identification of member endowments is possible and we obtained this information from NACUBO directly. Although the NACUBO studies are publicly available, identification of the members is not.

An important adjustment in the NACUBO surveying process during our sample period involves the collection of information on both the *actual* as well as the *intended* (i.e., policy) asset allocation schemes. In their surveys during the 2002-2005 period, NACUBO asked participating endowments to report not only their actual asset allocation but also their target levels for the next year. In the work to follow, we interpret this target allocation as deriving from the fund's policy, inasmuch as it represents the desired exposure to the various asset classes that is usually decided by the board of the institution as a general mandate for the investment process. For years prior to 2002, we take as proxies for the target weights the *actual* portfolio weights for each fund lagged by one year (i.e., $w^B_{i,j,t-1} = w_{i,j,t-2}$).

Table 1 provides an overview of some of the more notable characteristics of our endowment fund sample. Panel A of the display lists summary statistics for the full sample while Panels B and C given separate summaries for the endowments of private and public institutions, respectively. As shown in Panel A, the endowment fund universe increased significantly over the sample period; the number of funds more than tripled (from 200 to 709) and there was there was a roughly ten-fold increase in the aggregate assets managed by the industry (from \$25.4 billion to \$249.3 billion). Second, the assets under management (AUM) for the largest and smallest funds (e.g., \$25.5 billion versus \$1.26 million in 2005) indicates the tremendous cross-sectional heterogeneity in the sample and suggests that endowments of different sizes may face very different asset management problems that require different mixtures of active and passive risk.

The data in the first panel of Table 1 also show that while the overall distribution of fund returns are not highly skewed (i.e., there is not a large discrepancy between the mean and median returns reported for most years), the differences between the best and worst performing funds is typically considerable. Further, the last four columns of the display also show that there is a considerable degree of variation across endowments in terms their spending needs. The average annual payout, which NACUBO only began reporting in 1994, is about 5.0%, but the spread of values in a given year ranges from virtually zero to almost 25%. This is important in the present context because fund spending policies *may* influence the manager's risk allocation decision for the portfolio.

Another way in which endowment funds appear to differ in a meaningful way is whether the sponsoring educational institution is private or public. Panels B and C of Table 1 show these data starting in 1989, which is when NACUBO began recognizing this distinction. For any given year, there are three-to-four times more private school endowments than public ones and the average AUM for the former group is always larger. Additionally, private school funds generated a higher average return in each sample year and the range between the extreme outcomes in each subsample varies considerably on a year-to-year basis. On the other hand, there does not seem to be a discernable difference in the spending policies of the two institutional types.

Finally, we also collected annual return data for benchmark indexes representing the investable asset classes covered by the NACUBO endowment surveys. We define twelve distinct asset classes: U.S. Equity, Non-U.S. Equity, U.S. Fixed-Income, Non-U.S. Fixed-Income, Public Real Estate, Private Real Estate, Hedge Funds, Venture Capital, Private Equity, Natural Resources, Other Investments and Cash.[9] The benchmark definitions for these categories are shown in Table 2, which also lists several summary statistics for the annual investment performance they produced during the sample period.[10]

4 Are endowments exhausting their active management skills?: Empirical analysis

4.1 Decomposing endowment fund performance

Before proceeding with our assessment of whether endowment managers adopt an optimal mixture of active and risk exposures in their portfolios, it is first useful to examine the investment perfor-

[9] Brown, Garlappi, and Tiu (2009) contains a more detailed description of these asset class designations as well as thorough examination of the specific allocation tendencies of the endowment fund universe in these classes.

[10] Note that although the most popular U.S. Equity benchmark used by endowments is the Russell 3000 index, we have used the Center for Research in Security Prices (CRSP) value-weighted market index portfolio instead. The reason is that many endowments also use the Wilshire 5000 index as a benchmark for their U.S. Equity portfolios and thus we chose to use a broad, representative index. The results we present in the following section are unchanged if the Russell 3000 index is used instead.

mance in the fund sample over time. That is, are these funds capable of producing the consistent levels of alpha that would justify *any* active risk allocation? We address this question by analyzing the realized values of the annual return components produced by our endowment universe. Table 3 summarizes two dimensions of the annual peer group comparisons for the total return (R) as well as the passive (R^B) and active elements (R^T and R^S). Specifically, the display reports (i) the mean return for each component and (ii) the median return and inter-quartile range (i.e., the difference in returns produced by the endowments falling at the 75th and 25th percentiles of a particular distribution) for the overall return and the sum of the two active components.[11]

The reported findings support three primary conclusions about endowment fund performance over this period. First, as shown in the second and fifth columns of the display, both the overall returns and passive return components that the endowment universe generated are considerable. The total return averaged just under 10% per annum over the entire time frame, with positive mean returns reported in 13 of the 15 years represented. Further, the average passive return component was 9.04% for the whole period, which amounts to more than 90% (i.e., 9.04/9.96) of the mean total return. Of course, this finding is consistent with the conclusion of both Brinson, Hood, and Beebower (1986) and Ibbotson and Kaplan (2000) who showed that, over time, the strategic asset allocation decision is the most important determinant of the investment performance for any given fund.

The second thing to note from Table 3 is that there is a strong indication that endowment fund managers do indeed possess tangible alpha-generating skills. This is the more relevant finding for the present investigation because in the absence of active management skill the question of the appropriate way to combine passive and active risk exposure in the portfolio becomes moot. From the third-to-last column of the exhibit, the overall mean alpha generated by the endowment fund sample is almost 100 basis points per year, with a range across time of -1.26% to 5.21%. (Comparable values are reported in the penultimate column for the median fund.) Additionally, notice from the Tactical and Selection return columns that the latter component is the larger and more reliable determinant of the total active returns produced by the fund universe. In fact,

[11] In our database, we do not observe the returns $r_{i,j,t}$ for each asset class that are necessary to calculate the value for R^S directly from (4). However, this security selection return component can be easily obtained by subtracting the benchmark and market timing returns in (2) and (3), respectively, from the total return in (1). Further, recall that for most years in the sample, the estimation of policy benchmark weights requires the use of lagged actual weights, which explains why 1991 is the first year for which component returns are reported in Table 3 despite the fact the sample period begins in 1989.

when viewed over the entire sample period, security selection skills accounted for roughly 82% (i.e., 0.76/0.93) of the observed alpha.

Finally, while the mean alpha component data indicate that the average endowment manager is a genuinely skillful active investor, it is not necessarily the case that all of the endowments in the sample are equally capable in this regard. In fact, from the size of the reported inter-quartile ranges for the total and active return components, we can infer that there is a large degree of cross-sectional variation in the management abilities represented in the endowment universe. In particular, the last column in the display shows that, in a given year, the spread in active management performance between the upper and lower quartiles of the fund distribution can exceed 800 basis points, an amount that is about nine times the size of overall mean alpha itself! Further, this inter-quartile active return spread is never less than 278 basis points (i.e., in 2005). Finally, juxtaposing the values of the Total Return and Active inter-quartile ranges, it is clear that the active return component is the primary determinant of deviations in the overall investment performance across the sample. Thus, it is likely that there are important fund-level characteristics that distinguish the investment decision-making abilities of these managers.

4.2 Determining the optimal mix of active and passive risk: Full sample results

From the discussion in Section 2.3, the problem of finding the optimal λ can be thought of as a standard maximization problem with only two composite assets: R and R^B. Suppose, for instance, that an endowment fund's objective is to maximize its Sharpe ratio. If such an endowment is already choosing λ optimally, then, by construction, the solution in (8) will deliver an optimal weight $\lambda = 1$. The same holds true if the endowment maximizes its Sortino ratio instead and so the solution of the optimization problem for λ in either (8) or (9) provides us with a natural metric with which to analyze whether endowments are already optimally exploiting their active investment abilities (i.e., adopting an appropriate active risk allocation). An optimal value of $\lambda > 1$ implies that the fund is not fully exploiting its active management capabilities, and may have benefited from tilting its portfolio away from the benchmark indices and more toward its specific tactical asset class weights and individual security holdings.

To establish empirically the optimal exposure λ to active risk for each fund in our whole sample, we estimated the moments of the two composite assets R and R^B by using the entire observed

NACUBO endowment fund return time series from 1989 to 2005. We then solved separately for the λ values in (8) and (9) that maximized the Sharpe ratio and the Sortino ratio, respectively.[12]

Table 4 summarizes our findings. When the fund's objective is to maximize the Sharpe ratio, Panel A shows that the weights λ resulting from the maximization in (8) have a cross-sectional mean of 1.29 (t-statistic of 4.40) and a median of 1.17. As per our previous discussion, this finding can be interpreted as follows: It would have been optimal over this investment horizon for the typical endowment to *sell* 29% of its passive benchmark portfolio and hold 29% more of its actual managed portfolio. Maximizing (9) for the Sortino ratio produces a similar conclusion but with even stronger findings; the resulting cross-sectional average of the distribution of estimated λ coefficients is 1.52 (t-statistic of 6.20) and the median is 1.44.[13]

Because these initial findings are based on an *in-sample* analysis, it is reasonable to ask whether an endowment could still exploit such a strategy after learning about the extent of its active management abilities. Thus, a more interesting exercise is to analyze the *out-of-sample* consequences for a fund that decides to alter the active-passive risk profile in its portfolio based on the information revealed by these in-sample findings. To examine this possibility, we separated the sample into two distinct sub-periods and used the earlier of the two (i.e., from 1989 until 1997) to generate estimates of the moments of the current (R) and the benchmark (R^B) portfolios in the same manner as before. With these sample statistics, we calculated the optimal rescaling factor λ as of the end of 1997, for both the Sharpe and Sortino decision criteria. We then assumed that the endowments for which the optimizer recommends a weight of $\lambda > 1$ in 1997 (210 out of 339 endowments in 1997 for the Sharpe ratio case and 247 for the Sortino ratio case) actually did rescale their portfolio accordingly for the second half of the sample (i.e., from 1998 to 2005). Endowments for which $\lambda < 1$ were assumed to not implement the swap.

The results of this out-of-sample analysis are reported in Panel B of Table 4. When the set of λ values is obtained from maximizing (8) during the period 1989–1997, the implied strategy of enhancing active risk in the fund by holding λ of the actual portfolio results in an average increase in the endowment's Sharpe ratio of 0.0521 (t-statistic of 2.65) over the period 1998–2005 (i.e., from

[12] To minimize the effect of estimation error, we eliminated from consideration both the outlier weights at the 5% levels in the cross section before analyzing the properties of these distributions of λ values as well as all information for funds with fewer than five annual observations.

[13] To assess the robustness of these findings with respect to estimation error in the means and covariance matrix, we have replicated this analysis using only subsamples of endowments that have between six and 15 years of data available, as well as endowments for which quarterly data are available. The results do not change and λ remains statistically significantly above 1.00 in all cases.

a ratio of 0.3012 without the swap to 0.3533 with the swap). Further, the median Sharpe ratio increase is 0.0271 (z-statistic of 2.65).[14] Similarly, when the λ estimates are obtained from (9), the resulting increase in the average Sortino ratio for the endowments that follow the strategy of enhancing the active part of their portfolio is equal to 0.1079 (t-statistic of 4.05), from 0.5228 to 0.6307; the median Sortino ratio increase is 0.0672 (z-statistic of 2.81). These findings suggest quite strongly that (i) the typical endowment fund did not take enough active risk over the sample period, and that (ii) increasing the scale of the same alpha-generating skills it was already deploying would have significantly increased the level of the portfolio's risk-adjusted performance. With respect to the latter point, it is important to point out that scaling up alpha-generating skills does not necessarily imply an increased level of trading; the endowment could simply hold more of its overall portfolio in positions that are less correlated with the benchmark.

Another potentially important concern related to this conclusion is the effect that the swap has on the overall portfolio risk. In Panel C of Table 4 we report the changes in overall portfolio risk following swap implementation, using both the standard deviation or downside deviation of returns as volatility measures. The main observation from these findings is that the level of overall portfolio risk does indeed increase somewhat if the swap is implemented and that the increase is statistically—if not necessarily economically—significant. When the swap resulting from (8) is implemented, the standard deviation of the portfolio increases from 10.21% to 11.23%. (We get a similar result when the swap of implied by (9) is implemented.) Thus, despite the fact that this increase in portfolio volatility is more than offset by the simultaneous increase in return, there may well be some endowments which would be allowed to implement the swap only if such a risk increase is acceptable to their investment committees.

Finally, it is also instructive to address the following question: Which endowments are likely to see an increase in overall fund volatility following the implementation of the recommended active-passive risk swap? Table 5 attempts to answer this issue. Specifically, the display reports the estimated parameters of a cross-sectional regression of the volatility increases associated with the swap on several endowment-specific characteristics, including the correlation between the actual and benchmark portfolio returns, payout level, AUM, allocation to various broad asset classes and the level of total fund risk prior to swap implementation. The only characteristic related to the

[14] The distribution of the differences in the Sharpe ratios between the rebalanced portfolio and the original one ranges from -0.50 to 0.69. The second highest Sharpe ratio differential is 0.68.

post-swap risk increase that proves to be statistically significant for both risk measures is the level of fund risk prior to the swap implementation. This is intuitive: If an endowment held an overall riskier portfolio before the swap, shorting the benchmark to buy even more of it should clearly translate into a post-swap portfolio with a larger risk increase relative to the rest of funds. Finally, it is interesting to note that, contrary to conventional wisdom, scaling up the active risk exposure in a fund with a larger allocation to the alternative asset classes does not result in an increased overall level of volatility relative to endowments with a smaller allocation to hedges funds and private equity.

4.3 Optimal active-passive risk combinations: Cross-sectional results

The preceding section provides strong evidence supporting the notion that the typical endowment fund under-utilizes its alpha-generating ability and could significantly improve its performance by scaling up the active risk exposure in its portfolio. What is true on average, however, is not necessarily true across the entire sample. In fact, the descriptive statistics on the fund universe presented in Table 1 suggest that there are at least three characteristics—the AUM of the fund, the spending rate of the endowment, and whether sponsoring institutional is a public or a private school—that could have a material impact on the investment decisions made by the fund manager. In this section, we extend our previous analysis on the optimality of endowment active-passive risk exposures by looking at the findings on a cross-sectional basis.

Table 6 repeats our in-sample and out-of-sample calculations for the rescaling parameter λ based on the following ways of dividing the endowment universe into tractable sub-samples: The largest and smallest quartiles based on fund AUM (Panel A), the largest and smallest quartiles based on endowment spending rate (Panel B), and public versus private endowments (Panel C). In particular, for each of the respective divisions of the endowment universe, we calculate the average λ value over the entire sample period, as well as the change in the Sharpe and Sortino ratios during the second half of the period that would accompany the implementation of an active-passive swap based on information from the first half. To get a better sense of the pervasiveness of original results, we then compute the cross-sectional variation in those λ and objective function statistics, along with the t-statistics affiliated with those differences.

The results tell a mixed story about the importance of characteristic variability across the endowment sample with respect to the active-passive risk decision. On one hand, Panel A1 of Table 6 indicates that there are significant differences in the rescaling parameters estimated for large and small endowments. For instance, using the Sharpe ratio to represent the risk-adjusted objective, the average λ coefficient for the quartile of the fund sample with the largest AUM levels is almost twice the size of that for the lowest-AUM funds (i.e., 1.93 versus 1.07). There are two interesting aspects to this finding. First, it suggests that to achieve its optimal combination of exposures, the average large endowment should virtually double its current portfolio by shorting almost all of its passive position. The out-of-sample performance metrics in Panel A2 show that affecting this adjustment would dramatically increase both the Sharpe and Sortino ratios for the large endowment, albeit at marginal levels of statistical reliability. Of course, it is highly unlikely that any fund would be allowed the freedom to actually implement a swap that requires such an extreme use of leverage; this is a topic that we consider in more detail in the next section. Nevertheless, the findings underscore the extent to which the managers of large endowment funds are not deploying their full active management skills.[15]

The second interesting aspect of the results shown in Panel A is that the estimated λ levels for the average small endowment are not significantly different from 1.00. There is no appreciable increase in the performance statistics when the implied swap transaction is implemented by these small-fund managers (e.g., the pre-swap and post-swap Sharpe ratios are 0.2661 and 0.2914, respectively, yielding an insignificant increase of 0.0253). Taken at face value, this indicates that small endowments are fully utilizing their alpha-generating talents and these funds would gain very little by scaling up their current positions in exchange for the passive portfolio. What is also likely to be true, however, is that small funds have less access to active investment strategies in the first place, meaning that they have less to "leave on the table" when making their active-passive allocation decisions. This interpretation is consistent with the result reported by Brown, Garlappi, and Tiu (2009) that large endowments invest about five-to-ten times more of their proportional assets in hedge funds and private equity than do small endowments.

[15] It is worth recalling from Section 2.1 that, in equilibrium, active management should be a zero-sum game. However, it need not be the case that active management *within the endowment fund industry* adds no net value. Indeed, it is unlikely that endowment funds trade only with one another; they may, for instance, have retail investors as their most likely counterparties. Thus, there is no reason to suspect that larger endowments are benefitting at the expense of smaller funds.

In contrast with these cross-sectional AUM findings, there is no apparent difference in the λ and objective function estimates for the spending rate or public-private sub-samples. The coefficients reported in both Panels B and C of the display show that there are no statistically meaningful differences in these values for either characteristic. That is, despite the heterogeneity documented across the endowment universe for these variables themselves, using them to sort the sample made no difference in how the optimality of the active-passive risk allocation findings should be interpreted. For instance, the average rescaling parameters for the high-spending (i.e., 1.25) and low-spending (i.e., 1.22) quartiles using the Sharpe ratio objective function are not significantly different from one another. Further, both values are quite comparable to the average λ level for the overall sample (i.e., 1.29). A strikingly similar pattern emerges when the universe is divided into public and private institutional sponsorship. Thus, while both the spending needs and the profile of the educational structure might impact other aspects of the endowment management process, they do not appear to influence on a cross-sectional basis the decision of how much active risk to take.

5 Can the active-passive swap strategy be implemented?: Normative implications

Based on the analysis in the previous section, a reasonable question to ask is: If endowment fund managers could improve their risk-adjusted performance by this active-passive swap, why are they not already implementing the strategy? Broadly speaking, two answers are possible. First, it could be the case that the managers and sponsors of the various funds are unaware of—or otherwise fail to acknowledge—the active management skills that they possess within the organization. However, given the historical track record of producing positive alphas documented in Table 3, this seems unlikely to be the case. Indeed, the incentive-based performance (i.e., bonus) portion of the contracts that typify compensation arrangements in this part of the asset management industry is based on the explicit premise that managers *do* have measurable active management skills.

The second and more plausible reason that endowment fund managers appear to adopt a consistently sub-optimal combination of active and passive risk exposures is that there are impediments preventing them from implementing the superior alternative. Said differently, although fund managers recognize that they are under-utilizing their alpha-generating talents, they are prevented from altering the portfolio in a way that would correct the problem. We consider the normative implications of two possible impediments along these lines: Market scalability limitations and agency problems that may exist within the endowment fund complex.

The first potential problem that may be preventing fund managers from utilizing their skills fully is related to the fact that the strategies and portfolio positions responsible for the levels of alpha we observe in the sample are not easily scaled up in size. Such an obstacle would be consistent with the evidence reported by Chen, Hong, Huang, and Kubik (2004), who showed that active returns in the mutual fund industry declined significantly as the size of the fund grew, particularly when the investment involved transactions in small and illiquid securities.[16] In fact, of the two sources of alpha—tactical strategies and security selection—the latter would seem to be the most susceptible to this sort of "economy of scale" constraint; managers have a far deeper and more liquid array of choices (e.g., ETFs, index-linked derivatives) when implementing market timing trades. Thus, given our earlier finding that the selection component is the dominant driver of total active endowment fund returns, this impediment becomes a relevant factor. Clearly, endowments can only implement the recommended active-passive swap to the extent that returns to the security selection component are not already completely exhausted.

Assessing the direct impact that this scalability impediment might have on the endowment universe would require complete information about fund-level holdings, which is not available in our database. However, we presented findings in the previous section that provide strong indirect evidence that this sort of market-driven liquidity constraint does influence investment behavior in the endowment industry. Recall from Panel A of Table 6 that there was a tractable difference in how large and small funds managed their active-passive risk exposures. In particular, it was the managers of the largest endowments who appeared to be deploying only half of their active management skills, whereas small-fund managers were shown to be exploiting their full capabilities. This outcome is entirely consistent with the notion that large-fund managers are simply not able to scale up their alpha-generating investments to a greater extent (e.g., it may not be possible for an endowment to increase a $25 million limited partnership position in a hedge fund vehicle to a $50 million position on a discretionary basis). Conversely, consistent with Brown, Fraser, and

[16]Berk and Green (2004) have developed a formal theoretical model for active portfolio management that makes a similar prediction when managerial talent is a scarce resource that is dissipated as the scale of fund's operations increases.

Liang (2008), small-fund managers may be unable to access hedge fund partnerships in a profitable manner because their due diligence and monitoring costs are too high.

A second explanation for the apparent sub-optimality we observe in the way endowments manage their active-passive risk allocations may lie in the agency problems that exist in any principal-agent relationship. Besley and Prat (2003) as well as Walter (2004) have examined these conflicts in the specific case of the asset management industry and have identified three areas of friction that might exist between boards and managers: The responsibility for monitoring the manager, risk exposures and asset allocation decisions, and portfolio funding levels. They argued that if the decision-makers (i.e., the portfolio managers) do not bear the full cost of their actions, inefficient actions that adversely impact the position of the residual claimant (i.e., the fund sponsors) can result. In fact, Brown, Harlow, and Starks (1996) have shown the conditions in which fund managers alter the overall risk level of their portfolios in a manner consistent with an attempt to maximize their personal compensation.

This potential for agency conflicts would certainly seem to be present in the endowment fund industry. One natural way for fund sponsors to address these problems is to design investment policy statements that limit activities viewed as being potentially self-serving for managers (e.g., excessive risk-taking, the use of leverage) on an ex ante basis. For instance, in their examination of policy constraints used by mutual funds, Almazan, Brown, Carlson, and Chapman (2004) demonstrated the relatively wide-spread use of these covenants, particularly with regard to leverage-related strategies. Further, they also showed that these restrictions were more likely to be implemented in situations where directly monitoring investment managers was more costly to the investor. Thus, there may be cases in which endowment portfolios are under-allocated to active risk because the managers are simply not permitted by their boards to do otherwise.[17]

While the possibility that policy restrictions impede the optimal use of active risk is certainly plausible, there is nothing in the cross-sectional evidence in Panels B and C of Table 6 that substantiates this notion. Assuming that endowments which face different investment problems (i.e., have different spending rates) or are sponsored by different institutional types (i.e., private versus

[17] At the time of his arrest in early 2009, Bernard Madoff was the Chairman of the Board of Trustees for a private university with a $1.3 million endowment fund. By the end of that year, the university reported that $110 million, or approximately 8%, of their fund was lost due to fraudulent investment schemes. Thus, although the oversight and monitoring provided by the Board via policy restrictions and other measures are important safeguards in the overall management process, they are hardly fool-proof tools that can be employed with perfect knowledge and foresight.

public universities) might also have different sets of agency conflicts, the estimated λ levels for the various sub-samples might also be expected to differ. For example, both the boards and managers of public school funds are often under the direct scrutiny of a legislative body, which might imply a tighter set of restrictions than those faced by an otherwise comparable private endowment. However, as discussed earlier, any such differences did not result in measurably different rescaling parameter estimates. Of course, rather than indicate that agency problems do not matter, these data might simply imply that the nature of these conflicts are effectively the same for all endowments, regardless of the distinct characteristics that seem to separate the funds.

6 Concluding comments

Managers of multi-asset class portfolios—such as pension or endowment funds—attempt to enhance the performance of their strategic benchmark allocations by deploying their active investment skills in two ways: Market timing and security selection decisions. Thus, a fundamental choice that every such manager must make involves the appropriate mixture of passive and active exposures to include in the portfolio, which is tantamount to deciding the appropriate combination of active and passive risk to assume. Do managers make the right decision or do they tend to either over- or under-utilize their alpha-generating talents?

In this study, we have addressed this question in two ways. First, building on the essential insight of Treynor and Black (1973) that the active-passive risk allocation decision can be viewed as a portfolio optimization problem, we present a simple framework for determining what represents an optimal way of combining the two. Using an extensive database consisting of the allocation decisions and investment performance of a large set of university endowment funds over the period from 1989 to 2005, we then examine empirically the extent to which these managers exploit their value-adding capabilities. We found that the average endowment under-utilizes its active management skills by about 30% and that it could have significantly increased its risk-adjusted performance by shifting the existing portfolio away from its passive benchmarks and toward more of the active investment strategies and positions it already holds. Further, we show that this tendency is not uniformly distributed across the entire endowment universe but most pronounced for funds with the largest amount of assets under management. Thus, the answer to the central question that frames our investigation certainly appears to be "no."

The ancillary question that remains unanswered is why do endowment fund managers stop short of using their skills to the fullest potential? Rather than simply concluding that these funds are simply investing inefficiently, we considered two possible impediments to implementing a superior combination of active and passive exposures. First, we argued that it is quite likely that the alphas we documented for our sample are decreasing in scale, or are otherwise cost-ineffective to replicate on a larger basis. Second, we also suggested that rescaling the endowment portfolio in the manner recommended by the implied swap strategy might require transactions and alter risk in a way that creates additional agency problems in the fund organization. While our cross-sectional findings only support the former explanation, both are plausible alternatives that merit further research.

Table 1: Overview of the educational endowment fund universe

This table reports summary statistics for the sample of endowments in the NACUBO endowment study database. Included are annual data on the number of participating funds, assets under management (AUM), net-of-fee investment performance, and endowment payout. Panel A lists summary statistics for the entire universe, while Panels B and C present data for the funds of private and public schools, respectively.

		AUM ($ mill.)						Return					Payout				
year	N. obs	Total	Mean	Median	Std	Min	Max	Mean	Median	Std	Min	Max	Mean	Median	Std	Min	Max
Panel A. Full sample																	
2005	709	249297.29	352.61	77.43	1294.33	1.26	25473.72	9.16	9.00	3.29	-11.40	22.20	4.78	4.85	1.40	0.24	17.10
2004	705	228466.66	324.53	71.88	1138.76	1.92	22143.65	15.02	15.80	4.42	-0.60	32.10	4.88	5.00	1.62	0.00	18.40
2003	665	196111.70	294.90	72.42	1003.83	0.32	18849.49	2.69	2.60	3.67	-14.70	28.10	5.21	5.00	1.51	0.10	22.00
2002	535	176545.39	336.92	88.89	1053.23	0.16	17169.76	-5.95	-6.20	4.23	-19.80	10.10	5.22	5.00	1.54	0.30	22.00
2001	568	218185.98	393.13	106.24	1172.01	1.15	17950.84	-3.24	-3.55	6.31	-26.90	24.80	5.17	5.00	1.43	0.30	14.20
2000	507	222928.50	462.51	119.99	1373.86	6.19	18844.34	12.67	10.70	10.12	-12.20	58.80	4.97	5.00	1.40	0.10	13.30
1999	467	181781.61	410.34	132.47	1084.99	7.19	14256.00	10.79	10.50	4.71	-15.80	29.30	4.82	5.00	1.24	0.10	13.00
1998	445	159620.93	371.21	114.15	991.14	5.17	13019.74	18.06	18.10	3.97	3.70	31.80	4.69	4.70	1.21	0.30	14.70
1997	456	145122.97	328.33	99.30	858.51	5.78	10919.67	20.10	19.90	4.44	6.80	46.90	4.80	4.60	1.77	0.30	24.00
1996	405	112603.94	282.22	81.23	738.88	3.65	8811.78	17.25	17.20	4.02	0.80	39.00	4.84	4.80	1.30	0.10	12.80
1995	422	99416.70	236.14	72.08	600.90	0.07	7045.86	15.06	14.80	4.15	0.50	37.90	4.95	5.00	1.48	0.10	15.00
1994	375	70354.31	189.12	61.23	406.42	0.47	3529.00	3.25	3.20	4.35	-13.00	43.00	5.24	5.00	1.75	1.80	18.20
1993	394	76838.21	196.02	60.36	468.95	1.46	5778.26	13.15	13.30	4.28	-2.80	36.30	n.a.	n.a.	n.a.	n.a.	n.a.
1992	318	62236.12	196.33	53.66	533.94	1.73	5315.47	13.00	13.00	4.94	1.10	65.00	n.a.	n.a.	n.a.	n.a.	n.a.
1991	328	56023.55	171.85	51.67	437.94	3.39	5118.12	7.25	7.10	4.89	-14.40	53.00	n.a.	n.a.	n.a.	n.a.	n.a.
1990	298	51420.16	174.31	58.89	410.97	2.01	4653.23	10.02	9.55	6.01	-10.30	75.00	n.a.	n.a.	n.a.	n.a.	n.a.
1989	281	45818.39	165.41	56.87	390.98	1.87	4478.98	13.58	13.40	4.65	-6.20	37.40	n.a.	n.a.	n.a.	n.a.	n.a.
1988	262	49467.06	157.04	47.64	381.40	1.65	4155.78	1.26	0.70	4.78	-14.70	17.40	n.a.	n.a.	n.a.	n.a.	n.a.
1987	276	45222.98	152.78	46.38	360.68	1.54	4018.27	13.63	13.82	5.69	-6.88	31.89	n.a.	n.a.	n.a.	n.a.	n.a.
1986	258	39317.86	150.64	46.29	354.88	1.10	3435.01	27.17	27.16	6.72	5.24	52.46	n.a.	n.a.	n.a.	n.a.	n.a.
1985	280	31511.82	112.54	31.19	295.30	1.02	2927.20	25.61	25.95	6.96	1.25	52.12	n.a.	n.a.	n.a.	n.a.	n.a.
1984	200	25399.06	127.00	40.12	289.12	0.97	2486.30	-1.40	-2.24	15.12	-17.00	209.00	n.a.	n.a.	n.a.	n.a.	n.a.

Continued on the next page . . .

Table 1 (cont.): Overview of the educational endowment fund universe

year	N. obs	AUM ($ mill.) Total	AUM Mean	AUM Median	AUM Std	AUM Min	AUM Max	Return Mean	Return Median	Return Std	Return Min	Return Max	Payout Mean	Payout Median	Payout Std	Payout Min	Payout Max
Panel B. Private endowments																	
2005	535	196009.71	367.06	77.96	1430.50	1.26	25473.72	9.25	9.00	3.19	-1.00	22.20	4.79	4.81	1.43	0.24	17.10
2004	531	179359.61	338.41	72.79	1258.52	1.92	22143.65	15.07	15.80	4.30	-0.60	25.30	4.86	5.00	1.56	0.00	18.40
2003	506	154276.44	304.89	72.93	1097.94	1.58	18849.49	2.78	2.80	3.65	-14.70	28.10	5.18	5.00	1.53	0.10	22.00
2002	406	138023.88	345.92	87.95	1147.49	1.17	17169.76	-5.98	-6.30	4.15	-19.80	9.70	5.16	5.00	1.34	1.50	13.00
2001	438	174938.83	408.74	104.54	1276.78	1.15	17950.84	-3.23	-3.50	6.17	-26.90	14.30	5.10	5.00	1.30	1.50	14.20
2000	392	185127.68	494.99	121.38	1506.86	6.19	18844.34	12.60	10.50	10.38	-9.00	58.80	4.91	5.00	1.35	0.10	13.30
1999	363	151356.26	436.19	130.10	1191.00	7.19	14256.00	10.78	10.50	4.81	-15.80	29.30	4.77	4.80	1.20	0.10	13.00
1998	349	128724.41	379.72	110.85	1069.44	5.17	13019.74	18.01	18.10	3.98	3.70	31.80	4.66	4.60	1.21	2.00	14.70
1997	354	117736.28	342.26	97.64	932.95	7.50	10919.67	20.14	20.05	4.48	6.80	46.90	4.81	4.60	1.88	2.10	24.00
1996	320	92732.84	293.46	83.20	796.17	3.65	8811.78	17.25	17.15	4.03	0.80	39.00	4.87	4.80	1.30	2.30	12.80
1995	330	80851.02	245.75	72.56	651.06	0.07	7045.86	15.11	14.85	4.17	0.50	37.90	4.97	5.00	1.56	1.00	15.00
1994	295	57444.37	196.73	60.96	429.88	0.47	3529.00	3.36	3.10	4.61	-13.00	43.00	5.16	5.00	1.78	1.80	18.20
1993	310	63407.14	205.87	60.18	509.77	2.20	5778.26	13.10	13.25	4.43	-2.80	36.30	n.a.	n.a.	n.a.	n.a.	n.a.
1992	249	51462.60	206.68	53.66	584.29	1.73	5315.47	13.09	13.00	5.27	1.50	65.00	n.a.	n.a.	n.a.	n.a.	n.a.
1991	263	47809.00	182.48	51.65	479.09	3.39	5118.12	7.36	7.20	5.08	-14.40	53.00	n.a.	n.a.	n.a.	n.a.	n.a.
1990	228	41276.38	183.45	60.72	450.13	2.01	4653.23	10.25	9.80	6.40	-10.30	75.00	n.a.	n.a.	n.a.	n.a.	n.a.
1989	221	39135.75	178.70	56.87	430.85	1.87	4478.98	13.79	13.60	4.78	-6.20	37.40	n.a.	n.a.	n.a.	n.a.	n.a.
1988	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
1987	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
1986	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
1985	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
1984	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
Panel C. Public endowments																	
2005	174	53287.58	308.02	72.40	729.82	1.77	5433.09	8.89	8.80	3.60	-11.40	18.10	4.75	4.92	1.29	0.42	11.19
2004	174	49107.04	282.22	66.70	651.15	2.90	4951.65	14.88	15.65	4.77	0.20	32.10	4.93	5.00	1.78	0.00	12.00
2003	159	41835.25	263.11	68.13	616.83	0.32	4527.90	2.42	2.40	3.75	-12.10	14.10	5.34	5.00	1.43	1.00	11.20
2002	129	38521.51	308.17	95.83	671.85	0.16	4343.72	-5.87	-5.90	4.49	-17.80	10.10	5.40	5.00	2.03	0.30	22.00
2001	130	43247.15	340.53	113.87	715.19	7.59	4857.27	-3.28	-3.85	6.82	-20.50	24.80	5.40	5.00	1.76	0.30	12.80
2000	115	37800.82	350.01	117.61	743.66	7.05	5798.53	12.91	11.40	9.18	-12.20	47.30	5.17	5.00	1.55	0.30	11.90
1999	104	30425.34	316.93	140.47	547.26	8.22	4445.60	10.83	10.55	4.37	-4.30	22.00	5.01	5.00	1.37	0.30	11.20
1998	96	30896.53	339.52	132.05	621.33	7.93	3787.88	18.24	17.95	3.97	9.00	31.60	4.82	5.00	1.25	0.30	10.00
1997	102	27386.69	279.46	104.52	519.56	5.78	3216.65	19.93	19.65	4.29	9.20	36.10	4.75	4.70	1.34	0.30	9.10
1996	85	19871.10	239.41	80.39	460.91	4.12	2635.55	17.23	17.20	3.98	2.40	28.90	4.73	4.85	1.29	0.10	8.70
1995	92	18565.68	201.80	67.12	370.35	3.31	2190.59	14.88	14.75	4.07	1.30	30.20	4.89	5.00	1.15	0.10	7.30
1994	80	12909.94	161.37	62.24	306.64	2.72	1918.15	2.86	3.30	3.20	-5.80	13.00	5.57	5.30	1.60	3.80	13.00
1993	84	13431.07	159.89	63.54	270.39	1.46	1846.60	13.30	13.60	3.72	4.00	22.40	n.a.	n.a.	n.a.	n.a.	n.a.
1992	69	10773.52	158.43	53.86	281.15	4.24	1683.01	12.66	13.00	3.50	1.10	19.80	n.a.	n.a.	n.a.	n.a.	n.a.
1991	65	8214.55	128.35	52.20	190.10	4.29	1151.32	6.81	6.90	4.02	-10.90	15.80	n.a.	n.a.	n.a.	n.a.	n.a.
1990	70	10143.78	144.91	52.59	246.37	3.07	1494.93	9.29	9.15	4.47	1.20	26.20	n.a.	n.a.	n.a.	n.a.	n.a.
1989	60	6682.64	115.22	55.20	164.23	2.69	973.70	12.83	12.65	4.09	1.30	22.10	n.a.	n.a.	n.a.	n.a.	n.a.
1988	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
1987	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
1986	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
1985	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.
1984	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	n.a.	n.a.	n.a.	n.a.	n.a.

Table 2: Benchmark indices

This table lists the indices used as benchmarks for the asset classes represented in the NACUBO endowment study sample. Summary statistics are listed in percent per annum. Data for all of the indices are obtained from Bloomberg, Cambridge Associates and the University of Texas Investment Management Company.

			Summary statistics			
	Asset class	Benchmark index	Mean	Std	Median	SR
1.	U.S. Equity	CRSP Value Weighted	11.47	13.09	13.13	0.56
2.	Non U.S. Equity	MSCI World (Excl. US)	4.46	13.31	5.82	0.02
3.	U.S. Fixed Income	Lehman Bond Aggregate	8.04	4.41	8.64	0.88
4.	Non U.S. Fixed Income	Salomon Brothers Non US Bond Index	7.89	9.08	7.60	0.41
5.	Public Real Estate	NAREIT	13.27	12.79	9.06	0.71
6.	Private Real Estate	NCREIF	7.81	6.55	8.07	0.56
7.	Hedge Funds	HFRI-all fund Composite	14.31	7.61	13.09	1.34
8.	Venture Capital	Cambridge Associate VC index	26.59	56.20	17.42	0.40
9.	Private Equity	Cambridge Associate PE index	14.76	13.74	15.38	0.77
10.	Natural Resources	AMEX Oil (before 1992), GSCI (after 1992)	7.02	16.87	1.33	0.17
11.	Other Investments	—	—	—	—	—
12.	Cash	30-day U.S. T-Bill	4.14	1.92	4.73	0.00

Table 3: Decomposition of historical endowment fund returns

This table presents mean annual values for the various components of endowment fund returns from the NACUBO database as well as their grand time series means over the entire sample period. "Total Return" represents the overall returns reported to NACUBO. "Policy" are the average returns of the policy portfolios. "Active" represents the returns from active investment strategies, which include the "Tactical" and "Selection" categories. The exhibit also reports the annual inter-quartile (i.e., P75-P25) return spreads for the Total and Active return categories. Rounding accounts for any apparent discrepancies across categories.

	Total Return			Policy	Tactical	Selection	Active		
Year	Mean	Median	p75-p25	Mean	Mean	Mean	Mean	Median	p75-p25
2005	9.67	9.50	3.40	9.93	-0.52	0.26	-0.26	-0.13	2.78
2004	15.63	16.30	4.15	15.70	-0.44	0.37	-0.06	0.25	3.23
2003	2.37	2.40	3.20	3.08	0.35	-1.06	-0.71	-0.55	3.27
2002	-6.13	-6.30	4.45	-8.26	0.70	1.43	2.13	1.99	4.43
2001	-3.42	-3.70	6.90	-8.63	-0.38	5.59	5.21	5.22	7.06
2000	13.02	11.00	10.15	11.70	0.90	0.42	1.32	0.50	8.15
1999	10.85	10.60	5.38	11.77	0.28	-1.19	-0.92	-0.98	5.70
1998	18.09	18.10	4.30	17.91	0.12	0.06	0.18	0.21	4.29
1997	20.14	20.10	5.00	17.46	1.10	1.58	2.68	2.35	5.33
1996	17.35	17.20	3.80	14.99	0.28	2.08	2.36	2.13	3.74
1995	15.26	15.00	4.10	16.51	-0.68	-0.57	-1.26	-1.48	4.55
1994	3.39	3.20	3.15	1.14	0.16	2.09	2.25	2.01	2.98
1993	13.46	13.60	4.80	13.13	0.29	0.04	0.33	0.59	4.92
1992	13.00	13.00	3.08	12.07	0.12	0.81	0.93	0.75	3.86
1991	6.77	6.80	5.00	7.05	0.29	-0.57	-0.28	-0.07	4.15
Grand Mean	9.96			9.04	0.17	0.76	0.93		

Table 4: Optimal active-passive allocation

Panel A reports the scale parameter λ that maximizes the in-sample Sharpe or reward-to-downside-risk (Sortino) ratio, computed as described in Section 4. Panel B reports the out-of-sample Sharpe and Sortino ratios obtained by implementing during the period 1998–2005 the swap between the active and passive parts of the portfolio according to the value of the scale λ determined in the period 1989–1997. Panel C reports the out-of-sample relevant risk statistics (standard deviation and downside deviation) calculated over the period 1998-2005 when the swap determined during the period 1989–1997 is implemented.

Panel A: In-sample values of λ

Optimization Criterion	Mean λ	Median λ
Sharpe ratio	1.29	1.17
(t-stat)	(4.40)	(—)
Sortino ratio	1.52	1.44
(t-stat)	(6.20)	(—)

Panel B: Out-of-sample performance

Portfolio	Average Sharpe ratio	Average Sortino ratio
Without swap	0.3012	0.5228
With swap	0.3533	0.6307
Difference	0.0521	0.1079
(t-stat)	(2.65)	(4.05)

Panel C: Out-of-sample risk change

Portfolio	Average standard deviation	Average downside deviation
Without swap	10.21%	6.68%
With swap	11.23%	6.92%
Difference	1.01%	0.24%
(t-stat)	(3.15)	(3.71)

Table 5: Cross-sectional determinants of risk change following swap implementation

This table considers the effects of swap implementation on fund risk by reporting the estimated parameters of a cross-sectional regression of the change in post-swap endowment risk on the following variables: the correlation between actual and policy returns, payout level, endowment AUM, allocation to alternative asset classes, allocation to fixed-income, allocation to public equity, and pre-swap risk level. T-statistics are reported beneath each respective coefficient estimate.

Dependent Variable	Intercept	Correlation (actual, policy)	Payout	log AUM	Allocation to Alternatives	Allocation to Fixed Income	Allocation to Equity	Past Risk	Adjusted R-square
Δ (standard deviation)	-0.1098	0.0798	0.0000	0.0005	0.0126	0.0392	0.0155	0.5274	
t-stat	-1.7477	2.8332	0.3558	0.2455	0.1445	0.6489	0.2553	3.0030	6.93%
Δ (downside deviation)	-0.0408	0.0053	-0.0000	-0.0013	0.0914	0.0667	0.0319	0.3576	
t-stat	-0.8841	0.3162	-0.8534	-0.9065	1.5174	1.5197	0.7107	1.8465	2.65%

Table 6: Optimal active-passive allocation for different endowment types

Panel A1 reports the scale parameters λ that maximizes the in-sample Sharpe ratio and reward-to-downside-risk (Sortino) ratio, computed as described in Section 4, for the top quartile $q4$ and the bottom quartile $q1$ of assets under management. Panel A2 then reports the out-of-sample Sharpe and Sortino ratios obtained by implementing the swap between the active and passive parts of the portfolio during the 1998–2005 period according to the value of λ estimated over the period 1989–1997. Large and small endowments are compared based on their assets under management at the time when λ is determined in 1997. Panel B1 reports the scale parameters λ that maximizes the in-sample Sharpe and Sortino ratios for the top quartile $q4$ and the bottom quartile $q1$ of fund payout ratios. Panel B2 then reports the out-of-sample Sharpe and Sortino ratios obtained by implementing the swap between the active and passive parts of the portfolio during the 1998–2005 period according to the value of λ estimated over the period 1989–1997. Large and small payout endowments are compared based on their reported payout ratios at the time when λ is determined in 1997. Panel C1 reports the scale parameters λ that maximizes the in-sample Sharpe and Sortino ratios for private and public endowments. Panel C2 then reports the out-of-sample Sharpe and Sortino ratios obtained by implementing the swap between the active and passive parts of the portfolio during the 1998–2005 period according to the value of λ estimated over the period 1989–1997.

Panel A: Small Vs. Large AUM Endowments

Panel A1: In-sample values of λ

Optimization	Average λ		
Criterion	q4-Large	q1-Small	q4 - q1
Sharpe ratio	1.93	1.07	0.86
(t-stat)	(7.41)	(0.48)	(4.53)
Sortino ratio	2.29	1.26	1.03
(t-stat)	(8.84)	(1.36)	(4.30)

Panel A2: Out-of-sample performance

	Average Sharpe ratio			Average Sortino ratio		
Portfolio	q4-Large	q1-Small	q4 - q1	q4 - Large	q1 - Small	q4 - q1
Without swap	0.3979	0.2661	0.1318	0.7672	0.3960	0.3712
With swap	0.4763	0.2914	0.1848	0.9780	0.4597	0.5184
Difference	0.0783	0.0253	0.0531	0.2108	0.0636	0.1472
(t-stat)	(1.06)	(0.25)	(0.90)	(1.63)	(0.46)	(1.87)

Continued on the next page . . .

Table 6 (cont.): Optimal active/passive allocation for different endowment types

Panel B: Large Payout Vs. Small Payout Endowments

Panel B1: In-sample values of λ

Optimization	Average λ		
Criterion	q4-Large	q1-Small	q4 - q1
Sharpe ratio	1.25	1.22	0.03
(t-stat)	(3.50)	(1.97)	(0.23)
Sortino ratio	1.56	1.66	-0.10
(t-stat)	(6.39)	(4.81)	(-0.60)

Panel B2: Out-of-sample performance

	Average Sharpe ratio			Average Sortino ratio		
Portfolio	q4-Large	q1-Small	q4 - q1	q4 - Large	q1 - Small	q4 - q1
Without swap	0.3076	0.3181	-0.0105	0.5417	0.6126	-0.0708
With swap	0.3657	0.4274	-0.0617	0.6615	0.8048	-0.1433
Difference	0.0581	0.1093	-0.0512	0.1197	0.1922	-0.0725
(t-stat)	(1.08)	(0.91)	(-1.02)	(1.54)	(1.04)	(-1.11)

Continued on the next page . . .

Table 6 (cont.): Optimal active/passive allocation for different endowment types

Panel C: Public Vs. Private Endowments

Panel C1: In-sample values of λ

Optimization	Average λ		
Criterion	Private	Public	Private - Public
Sharpe ratio	1.27	1.20	0.07
(t-stat)	(3.73)	(1.38)	(0.45)
Sortino ratio	1.54	1.51	0.03
(t-stat)	(6.00)	(2.59)	(0.13)

Panel C2: Out-of-sample performance

	Average Sharpe ratio			Average Sortino ratio		
Portfolio	Private	Public	Private - Public	Private	Public	Private - Public
Without swap	0.3044	0.2890	0.0154	0.5526	0.4310	0.1216
With swap	0.3676	0.3084	0.0592	0.6717	0.5015	0.1702
Difference	0.0632	0.0194	0.0438	0.1192	0.0705	0.0486
(t-stat)	(1.13)	(0.27)	(0.94)	(1.58)	(0.51)	(0.74)

References

Alford, A., R. Jones, and K. Winkelmann, 2003, Budgeting risk along the active risk spectrum, in R. Litterman, ed.: *Modern Investment Management: An Equilibrium Approach* (Hoboken, NJ: John Wiley & Sons).

Almazan, A., K. Brown, M. Carlson, and D. Chapman, 2004, Why constrain your mutual fund manager?, *Journal of Financial Economics* 73, 289–321.

Berk, J., and R. Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, 1269–1295.

Berkelaar, A., A. Kobor, and M. Tsumagari, 2006, The sense and nonsense of risk budgeting, *Financial Analysts Journal* 62, 63–75.

Besley, T., and A. Prat, 2003, Pension fund governance and the choice between defined benefit and defined contribution plans, *Institute for Fiscal Studies* working paper.

Brinson, G., L. Hood, and G. Beebower, 1986, Determinants of portfolio performance, *Financial Analysts Journal* 42, 39–48.

Brinson, G., B. Singer, and G. Beebower, 1991, Determinants of portfolio performance II: An update, *Financial Analysts Journal* 47, 40–48.

Brown, K., L. Garlappi, and C. Tiu, 2009, Asset allocation and portfolio performance: Evidence from university endowment funds, *Journal of Financial Markets* forthcoming.

Brown, K., V. Harlow, and L. Starks, 1996, Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry, *Journal of Finance* 51, 85–110.

Brown, S., T. Fraser, and B. Liang, 2008, Hedge fund due diligence: A source of alpha in a hedge fund portfolio strategy, *Journal of Investment Management* 6, 23–33.

Chen, J., H. Hong, M. Huang, and J. Kubik, 2004, Does fund size erode mutual fund performance? the role of liquidity and organization, *American Economic Review* 94, 1276–1302.

Chow, G., and M. Kritzman, 2001, Risk budgets, *Journal of Portfolio Management* 27, 56–60.

Clarke, R., H. de Silva, and S. Thorley, 2006, The fundamental law of active portfolio management, *Journal of Investment Management* 4, 54–72.

Clarke, R., H. de Silva, and B. Wander, 2002, Risk allocation versus asset allocation, *Journal of Portfolio Management* 29, 9–30.

Grinold, R., 1989, The fundamental law of active management, *Journal of Portfolio Management* 15, 30–37.

——— , and R. Kahn, 2000, *Active Portfolio Management* (McGraw-Hill: New York) 2nd edn.

Griswold, J. (ed.), 2008, Commonfund benchmarks study: Educational endowment report, *Commonfund Institute, Wilton CT.*

Harlow, V., 1991, Asset allocation in a downside-risk framework, *Financial Analysts Journal* 47, 28–40.

———, and K. Brown, 2006, The right answer to the wrong question: Identifying superior active portfolio management, *Journal of Investment Management* 4, 15–40.

Ibbotson, R., and P. Kaplan, 2000, Does asset allocation policy explain 40, 90, or 100 percent of performance?, *Financial Analysts Journal* 56, 26–33.

Ineichen, A., 2007, *Asymmetric Returns: The Future of Active Management* (New York: John Wiley & Sons).

Kane, A., T. Kim, and H. White, 2003, Active portfolio management: The power of the *t*reynor-*b*lack model, Working Paper, University of California, San Diego.

Leibowitz, M., 2005, Alpha hunters and beta grazers, *Financial Analysts Journal* 61, 32–39.

Lo, A., 2008, Where do alphas come from?: A measure of the value of active investment management, *Journal of Investment Management* 6, 6–34.

Pearson, N., 2002, *Risk Budgeting: Portfolio Problem Solving with Value-at-Risk* (New York: John Wiley & Sons).

Qian, E., and R. Hua, 2004, Active risk and information ratio, *Journal of Investment Management* 3, 20–34.

Ross, S., 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–360.

Sharpe, W., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425–442.

Swensen, D., 2009, *Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment, 2nd edition* (Free Press: New York).

Treynor, J., and F. Black, 1973, How to use security analysis to improve portfolio selection, *Journal of Business* 46, 66–86.

Walter, I., 2004, Conflicts of interest and market discipline among financial services firms, *European Management Journal* 22, 361–376.

Waring, M., and L. Siegel, 2003, The dimensions of active risk, *Journal of Portfolio Management* 29, 35–51.

Winkelman, K., 2003, Issues in strategic asset allocation, in R. Litterman, ed.: *Modern Investment Management: An Equilibrium Approach* (Hoboken, NJ: John Wiley & Sons).