Detecting Performance Persistence in Fund Managers

Book benchmark alpha analysis.

Rajesh K. Aggarwal, Galin Georgiev, and Jake Pinato

RAJESH K. AGGARWAL is associate professor of finance in the Carlson School of Management of the University of Minnesota, Minneapolis.
raggarwal@csom.umn.edu

GALIN GEORGIEV is director at Pacific Alternative Asset Management Company in Irvine, CA.
ggeorgiev@paamco.com

JAKE PINATO is Associate at The Provident Group & Longship Capital in New York City.
jpinato@provident-group.com

In active portfolio management, determining whether a fund manager's performance is likely to persist is a critical question. Is performance due to skill or to luck?

For hedge funds in particular, there are few if any firmly established quantitative ways to aid a portfolio manager of a fund of hedge funds in assessing the likelihood of performance persistence. While one may have past returns, given the dynamic nature and high turnover of the typical hedge fund's portfolio and the high dependence on market conditions, past returns are often a poor predictor of future returns.

There is a considerable literature on hedge fund returns predictability. Some authors find instances of performance persistence, while others do not (see, e.g., Brown, Goetzmann, and Ibbotson [1999], Agarwal and Naik [2001], Edwards and Caglayan [2001], Schneeweis, Kazemi, and Martin [2002, 2003], Amenc, El Bied, and Martellini [2003], Kat and Menexe [2003], and Jagannathan, Malakhov, and Novikov [2006]). These researchers analyze and decompose hedge fund returns using single- or multifactor models based on market and style factors. They address the question of performance persistence by focusing on the balance between hedge fund style and manager-specific risk and return. None of them use data on underlying hedge fund positions, so they are limited to inferring hedge fund risks from returns, combined in some cases with knowledge of hedge fund styles or strategies.

We establish a new metric for detecting performance persistence in hedge fund managers in a fund of funds context. Using hedge fund positions as they evolve over time, we decompose hedge fund manager performance to generate what we call *book benchmark alphas*.

This approach requires knowledge of underlying hedge fund positions, which are highly proprietary, although for large institutional investors and funds of funds, position-level transparency from managers is achievable. We demonstrate that for the particular fund of hedge funds to which we have access, these alphas are relatively time horizon-invariant and show persistence out-of-sample.

Alphas generally have higher predictive power than returns, and significantly higher power in shorter time periods. Book benchmark alphas thus provide a powerful additional tool for fund manager evaluation despite only a few months of positions and returns data.

While our analysis is developed in the context of a fund of hedge funds because of data considerations, the book benchmark analysis is more general. The approach can be used in any case involving manager selection as long as there is position-level transparency. For example, university endowments and pension funds have portfolios allocated across multiple managers. The approach is equally useful in evaluating institutional (long-only) equity fund managers.

TRADITIONAL PORTFOLIO MANAGEMENT FOR A FUND OF FUNDS

Quantitatively evaluating managers' past performance as a stand-alone investment vehicle is not an easy task, despite the multitude of return-to-risk measures such as the Sharpe ratio and its many variations. It is even harder to evaluate a manager's value-added to an already established large portfolio of funds.

How do we define relative or marginal return and risk and, more important, how do we combine them in a conceptually meaningful way? These are standard questions in portfolio construction, and in the fund of hedge funds context, they present some unique challenges.

Return and Risk: The Underlying Data Source

Traditional quantitative active portfolio management usually computes both returns and risk from the same source—the historical returns of the underlying individual securities. This makes sense because most individual securities usually have sufficient historical data points to provide a meaningful statistical analysis. This is not the case when individual securities are replaced by actively managed portfolios (e.g., hedge funds). In this case, there are typically only a small number of monthly returns available for the hedge funds with inherently dubious origin (managers themselves) for the more illiquid ones. More important, these returns give no information about the actual portfolio holdings or even general exposures and risks.

Traditional quantitative active portfolio management of equities assumes a sizable universe of stocks for trading (in the thousands) so one can run multidimensional optimizers to achieve the maximum information ratio (based on historical returns). A manager of a portfolio of hedge funds does not have this luxury; one usually has to choose among only a few candidates to hire or replace at any given time.

Benchmarks: History and Transparency

Traditional active portfolio management evaluation focuses on external passive benchmarks, usually traded indexes. Traded indexes have an extended history (with daily frequency) and are transparent. This allows for quantitative portfolio construction and beta/alpha analysis based on historical returns.

This is not the case in the hedge fund universe; investable hedge fund benchmarks are still in their infancy. The actual securities underlying the hedge fund index are typically not available, so it is impossible to understand the nature of the underlying risks.

It is of course widely accepted that traditional market indexes, while useful, do not adequately describe the dynamics of hedge fund returns (although there is growing evidence that for certain hedge fund strategies, market and other style indexes can provide some level of predictability; see Amenc, El Bied, and Martellini [2003] and Fung and Hsieh [2004], for example).

Marginal Risk (VaR): No Straightforward Relation to Actual Returns

The concepts of marginal or incremental risk (VaR) were born out of the recognition that in an established portfolio, what really matters is not so much the absolute risk of the individual positions but their relative risk in terms of the overall portfolio. What VaR does not address is the question of incremental performance. How does one link marginal performance and risk attribution?

DECOMPOSING FUND RETURNS: BOOK BENCHMARK ALPHAS AND BETAS

We propose an approach to address the concerns of relative return and risk. The overall idea is to use the information content in fund managers' actual positions to assess their skill. First, we separate the information sources for the manager's return and risk. We use actual contemporaneous positions for understanding the manager's marginal risk (beta), but continue to use actual manager returns to quantify managers' idiosyncratic realized alpha. This approach strikes a reasonable balance between the two current standards: either looking at position-based marginal risk measures with no relation to return, or looking at performance-based betas, which provide very limited information about the underlying risk.

Second, we use the overall fund of hedge funds portfolio (book) as the benchmark. The largest universe for which a fund of funds can practically obtain position transparency is of course its own portfolio. An added benefit of this choice—at least in theory—is that a manager with zero alpha and beta equal to one to the book can be fired, and the proceeds reinvested pro rata in the current portfolio.

The advantage gained from this approach is the intuitive and straightforward transition from book betas to alphas based on actual returns. This immediately yields the desired link between marginal risk (book betas) and marginal returns (book alphas).

EXHIBIT 1

Net Market Value as % of NAV (left axis) versus Book Benchmark Beta (right axis)

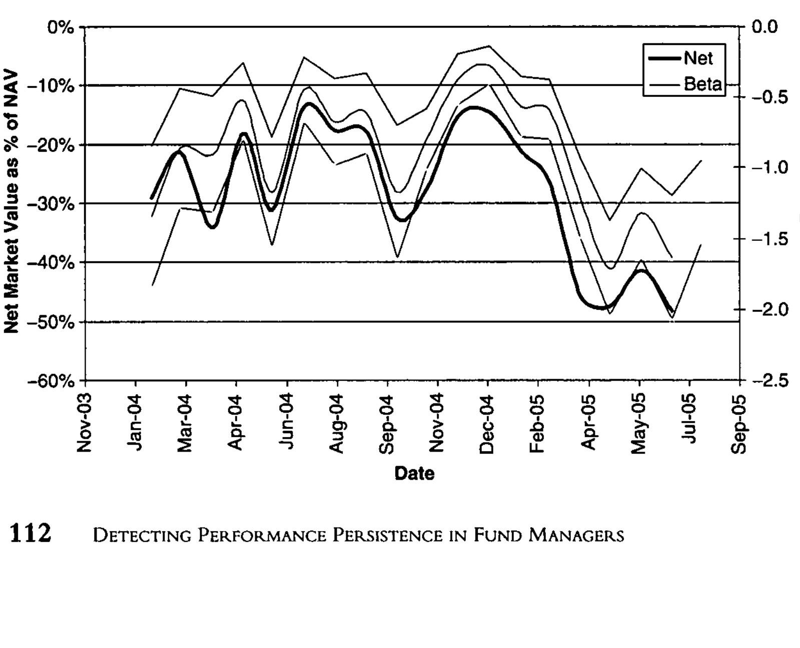

We define β_{it} as the book benchmark beta of manager i in time period t (monthly in our case), computed as follows. At the beginning of the time period, we take the full set of positions in the manager's portfolio and in the overall book portfolio. We generate historical returns for both portfolios going back in time and assuming that the positions have been held constant. We then regress the generated return series from the manager's portfolio on the return series from the book portfolio to compute the manager's book beta for that period.

In the next time period, we again take the full set of positions (which have typically changed), generate return series as before (holding the portfolio constant), and regress manager returns on book returns to generate the book beta for this time period. We do this for every time period and every manager.

The frequency of the position sampling (time period) is monthly in our case. The choice of time period is important and should vary depending upon the nature of the book. Ideally, the frequency of position sampling should match the frequency with which managers make meaningful changes in their positions.

The weighted average beta for the overall book for a given time t should equal one:

$$\sum_{i=1}^{N} w_{it}\beta_{it} = 1 \qquad (1)$$

where the index i runs across all the N managers in the book, and w_{it} is the weight of manager i in the book at time t.

These monthly betas computed from the manager's positions perform quite well in capturing a manager's overall exposure. Exhibit 1 graphs a specific hedge fund's book betas in conjunction with its net market exposure (as a percentage of its net asset value). The hedge fund is an equity-only short-seller, so its net market exposure is negative.

The changes in the betas fairly accurately reflect the significant monthly changes of the fund's net equity market exposure. The month-to-month changes could not have been captured if one had used the actual 18 monthly returns to calculate a fixed beta.

We can now define the realized alpha α_{it} for manager i at time t as follows:

$$\alpha_{it} = R_{it} - \beta_{it}R_t^B \qquad (2)$$

where R_{it} and R_t^B are the actual returns of manager i and the book benchmark at time t. We do not net out the risk-free rate from the book and manager returns because hedge fund returns tend to be judged on an absolute return basis. The book benchmark alpha in Equation (2) has to be distinguished from the in-sample alpha (the intercept of the regression used for computing the corresponding beta) because the latter would have been realized in practice only if the portfolio had not changed throughout the in-sample period.

Note that the alphas we are using are not risk-adjusted. Similar to the traditional in-sample alphas, the realized alphas defined above still have the nice property that their average across the whole book is by construction zero for any given time period t:

$$\sum_{i=1}^{N} w_i \alpha_{it} = \sum_{i=1}^{N} w_i R_{it} - \sum_{i=1}^{N} w_i \beta_{it} R_t^B = R_t^B - R_t^B = 0 \quad (3)$$

The book benchmark alphas are therefore a tool to rank managers within the current portfolio, but have no meaning in absolute terms. This is perfectly acceptable and perhaps desirable as the best—if not the only—reference point for a fund manager is his or her own peers. Most important, hiring and retaining managers with superior book alphas will improve the overall return-risk profile of the book. We demonstrate this in the appendix under some simplifying assumptions.

While these definitions are fairly intuitive and do not represent a radical departure from current active portfolio management, they do make use of different information and can potentially lead to improvements in performance attribution. It is also natural to go a step farther and investigate whether to create a book benchmark information ratio measure. In other words, we could scale the alphas by some measure of residual risk (volatility).

It turns out that in a typical fund of hedge funds this would not add a lot of value. As we discuss in the appendix, under some standard assumptions, the improvement it may bring to the return-to-risk ratio of the overall portfolio is of second order to the manager's weight. Thus, we confine ourselves to scrutinizing the alphas alone.

Another natural extension of the book benchmark approach is to ask whether managers have beta—timing skill—are they able to anticipate movements in market or book returns? We have started some preliminary investigation of this issue, but have not yet found significant evidence of such skill. We view this as a promising area for future research.

We take the analysis one step farther when we examine whether book benchmark alphas have predictive power for future or out-of-sample alphas (alpha persistence) and how alpha predictability relates to return predictability. This is the key test of our approach.

DATA

Throughout the study, we analyze a portfolio of hedge funds consisting of all equity managers who had an investment from Pacific Alternative Asset Management Company (PAAMCO) in any month from December 2003 through July 2006. In hedge fund strategy terms, they are either long-short equity, equity market-neutral, or short bias funds. They are equally weighted and form what we refer to as the book.

As we naturally cannot discuss the overall PAAMCO portfolio, this is a natural candidate for a subportfolio allowing for historical and other simulations on equity positions with the availability of historical returns. With equity market managers, we do not need to worry about marking illiquid positions to market as the price of equities is immediate and transparent.

While our approach is general and can be used across different strategies, for the purposes of this article we confine ourselves to equity-based strategies. This avoids issues of whether styles or strategies are driving differences in performance (see Amenc, El Bied, and Martellini [2003] and Fung and Hsieh [2004]).

If and when a hedge fund exits the PAAMCO portfolio within the time period, it disappears from the book because its positions are not available any longer. (We discuss the selection and survivorship bias naturally created from this choice and how we account for it later.)

In order to depend as little as possible on the specific choice of time periods during our 31-month interval, we designate various subperiods as the in-sample period and subsequent subperiods as the out-of-sample period. For a given month, a manager's beta relative to the book is computed by regressing the returns generated from the manager's positions at the beginning of the month (with the first month starting on January 1, 2004) against the returns generated from the overall book portfolio's positions at the same time. Then, using manager monthly book betas and the actual manager and book net returns (after fees), we compute the realized monthly alphas as described above.

Exhibit 2 describes the data. While the complete sample has 53 managers, 45 have data for the first 15 months (January 2004 through March 2005), and 44 have data for the last 16 months (April 2005 through July 2006). In total, there are 900 monthly observations (betas and returns for an individual manager). In principle, with 53 managers and 31 months, we should have a total of 1,643 manager-month observations, but in practice we have 900 observations. Of the missing observations, 353 are missing because the managers were hired after January 2004, and 390 are missing because the managers were either fired or closed down prior to July 2006. Both sets of missing observations could impart a bias to the results.

By construction, the average book benchmark beta for the funds is 1 and the average book benchmark alpha is 0. These hold for both in- and out-of-sample periods. Over the sample period, average monthly raw returns are 71 basis points.[1] This is comparable to the CSFB/Tremont Investable Long/Short Equity Index average monthly return (the closest available index) of 52 basis points over the sample period. The standard deviations of both returns and alphas are relatively high—2.3% and 1.9%, respectively, for the entire sample.

In order to establish whether there is alpha persistence (and how it relates to return persistence), we use various combinations of in-sample and out-of-sample periods, determined as follows: 1) The in-sample periods are the same for all managers, starting in January 2004, with a length of at least six months, increasing by one quarter and rolling forward by as many months as the out-of-sample period length; 2) the out-of-sample periods are no longer than half the length of the in-sample period; 3) every in-sample/out-of-sample pair should have at least the original number of monthly observations minus at most three data points for the in-sample period; in the out-of-sample period, there should be at least the number of out-of-sample monthly observations as in the previous grouping in order to avoid duplicate observations in the out-of-sample collection of data points.

Application of these criteria produces five groupings of in-sample/out-of-sample pairs:

1. In-sample (IS) = 6 months with a minimum of 4 monthly observations, out-of-sample (OS) = 3 months, minimum of 2 monthly observations.[2]
2. IS = 9 months (minimum of 6), OS = 4 months (minimum of 3).
3. IS = 12 months (minimum of 9), OS = 6 (minimum of 4).
4. IS = 15 months (minimum of 12), OS = 7 months (minimum of 6).
5. IS = 18 months (minimum of 15), OS = 9 months (minimum of 7).

Note that group 5 includes only one in-sample and one out-of-sample period per manager, while groups 1 through 4 have multiple observations per manager.

The numbers of in-sample/out-of-sample pairs for each group are shown in Exhibit 3. For example, the first group has 199 pairs of corresponding in-sample and out-of-sample periods. We then compute the average alpha and the average return for every in-sample and for every out-of-sample period for every manager (so for the first group, for example, there are 199 in-sample average alphas and 199 out-of-sample average alphas and the same number of corresponding returns).

Exhibit 3 shows the correlations between those alphas and returns. They are all significant except for the correlations for group 5. Correlations typically fall between 40% and 75%.

While the correlations are generally high, it is also clear that alphas potentially

EXHIBIT 2
Descriptive Statistics

	All (Jan '04-Jul '06)	Jan '04-Mar '05	Apr '05 – Jul '06
Number of Managers	53	45	44
Number of Observations	900	451	449
Beta (Book)	1.00 (0.93)	1.00 (0.98)	1.00 (0.89)
Alpha (Book)	0.0% (1.9%)	0.0% (1.8%)	0.0% (2.0%)
Average Raw Monthly Returns	0.71% (2.3%)	0.67% (2.0%)	0.76% (2.5%)
CSFB/Tremont Investable Long/Short Equity Index Average Monthly Returns	0.52%	0.73%	0.33%

Means are reported with standard deviations in parentheses. Betas are measured relative to the book and by construction average one. Alphas are computed month-by-month by subtracting the manager's monthly beta times the book return from the manager's realized return. Alphas, by construction, average zero. Returns are monthly.

convey unique information beyond simply looking at returns. Thus, while we will also examine the question of return persistence, we focus primarily on alpha persistence to see if there is incremental information one can learn by looking at alphas.

RESULTS

We start by asking whether our decomposition provides meaningful insight. Are the book benchmark alphas really predictive? Second, does our decomposition improve upon alternatives such as looking solely at past returns (a usual metric in the hedge fund universe)? We also compare our results to those from a naive book benchmark analysis where the betas and the alphas are computed using actual returns, not positions (i.e., one static beta and alpha per manager), and we examine the robustness of the results.

To address the first question, we consider again the set of the average alphas for every in-sample and every out-of-sample period (for every manager) and compute the correlation between the in-sample average alphas and the out-of-sample average alphas. Exhibit 4 presents the results.

There are two notable findings. First, whatever the definition of in-sample versus out-of-sample periods, there is a high degree of correlation between in- and out-of-sample alphas, ranging from 0.39 to 0.61, suggesting a high degree of persistence. All are statistically significant at the 5% level. Performance persistence is both high and economically significant. Using 12 months as the in-sample period, a 1.0 basis point improvement in in-sample monthly performance (alpha) is associated with a 0.92 basis point increase in out-of-sample monthly alpha.

Second, whether we use multiple observations or a single observation per manager, we find a high level of persistence. Thus, our results are not driven by performance persistence for a small number of managers with multiple observations.

So how does book alpha predictability compare with return predictability? Correlations between in-sample average returns and out-of-sample average returns for the five in-sample and out-of-sample groups are also in Exhibit 4. In all cases, the in-sample versus out-of-sample return correlations are weaker than those based on alphas, suggesting that the alpha decomposition provides better inference than do returns. We demonstrate this in graphs in Exhibit 5 for the first group where the outperformance of alphas is most visible.

Only the return correlation for the second group is statistically significant, although this may be a power issue for the third through fifth groupings. To confirm

EXHIBIT 3
Correlations Between Alphas and Returns

Groupings	Sample Period (Minimum)	# IS-OS Pairs	Correlation of IS alphas and IS returns	Correlation of OS alphas and OS returns	Correlation of IS+OS alphas and IS+OS returns
1	IS 6 Month (4) OS 3 Month (2)	199	0.58* (0.48 - 0.67)	0.69* (0.61 - 0.76)	0.66* (0.60 - 0.71)
2	IS 9 Month (6) OS 4 Month (3)	118	0.53* (0.38 - 0.65)	0.65* (0.53 - 0.74)	0.61* (0.52 - 0.68)
3	IS 12 Month (9) OS 6 Month (4)	60	0.52* (0.31 – 0.68)	0.76* (0.63 - 0.85)	0.70* (0.59 - 0.78)
4	IS 15 Month (12) OS 7 Month (6)	32	0.59* (0.31 - 0.78)	0.37* (0.02 - 0.63)	0.43* (0.21 – 0.61)
5	IS 18 Month (15) OS 9 Month (7)	13	0.46 (-0.13 – 0.8)	0.05 (-0.51 – 0.59)	0.13 (-0.27 – 0.49)

A minimum number of months of data must be present to be included (listed in parentheses). Two standard deviation confidence intervals in parentheses undernealth correlations. Confidence intervals are computed using Fisher's Z-transformation.

EXHIBIT 4
Alpha and Return Predictability for All Managers

Groupings	Sample Period (Minimum)	# IS-OS Pairs	Correlation of IS alphas and IS returns	Correlation of OS alphas and OS returns	Correlation of IS+OS alphas and IS+OS returns
1	IS 6 Month (4) OS 3 Month (2)	199	0.58* (0.48 - 0.67)	0.69* (0.61 - 0.76)	0.66* (0.60 - 0.71)
2	IS 9 Month (6) OS 4 Month (3)	118	0.53* (0.38 - 0.65)	0.65* (0.53 - 0.74)	0.61* (0.52 - 0.68)
3	IS 12 Month (9) OS 6 Month (4)	60	0.52* (0.31 – 0.68)	0.76* (0.63 - 0.85)	0.70* (0.59 - 0.78)
4	IS 15 Month (12) OS 7 Month (6)	32	0.59* (0.31 - 0.78)	0.37* (0.02 - 0.63)	0.43* (0.21 – 0.61)
5	IS 18 Month (15) OS 9 Month (7)	13	0.46 (-0.13 – 0.8)	0.05 (-0.51 – 0.59)	0.13 (-0.27 – 0.49)

EXHIBIT 5

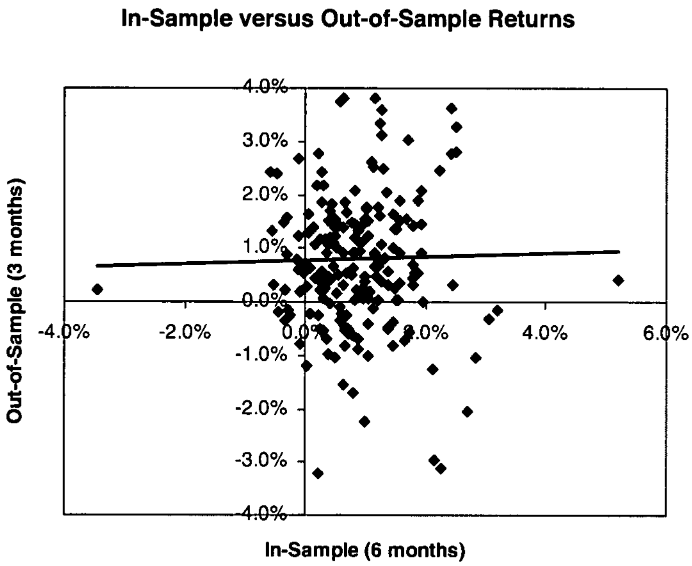

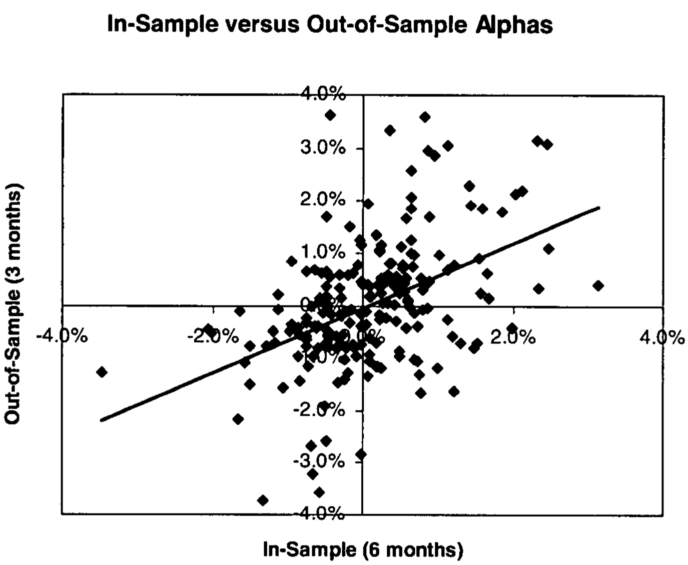

EXHIBIT 6

Regression (Static) Alphas and Betas

Grouping	Sample Periods (Minimum)	# IS-OS Pairs	Correlation of Static IS and OS alphas	Correlation of IS and OS Book Benchmark alphas	Correlation of IS and OS returns
5	IS 18 Month (15) OS 9 Month (7)	13	0.06 (-0.51 – 0.59)	0.56* (0.01 – 0.85)	0.51 (-0.06 – 0.83)

that alphas provide better predictability than returns, the last column of Exhibit 4 tests for differences in the correlations of alphas and returns. In three of the five groups, the alpha correlations are significantly higher than the return correlations.[3]

We next consider whether the book benchmark alphas computed using dynamic betas (based on actual positions month-by-month) are better than static alphas and betas, i.e., alphas and betas based on regressions of actual manager returns on book returns without any knowledge of positions. This is hard to accomplish in real time (if only monthly returns are available) because one obviously needs a long history of past returns to make the beta calculations statistically meaningful. Nevertheless, we compute static alphas and betas using the fifth group (18 monthly observations in the in-sample period), and then examine alpha predictability.

We regress each manager's actual returns on the overall book returns for the in-sample period and then for the out-of-sample period. We take the alphas from these regressions, and then compute the in-sample and out-of-sample static alpha correlation. The results are presented in Exhibit 6.

The in-sample versus out-of-sample static alpha correlation is low and not statistically significant. It is much weaker than the corresponding book benchmark alpha correlation of 0.56. Further, the static alpha correlation is lower than the in-sample and out-of-sample return correlation, which is also not statistically significant.

Given the high uncertainty in the static alpha and beta calculations, one probably needs at least three years of past monthly returns to carry out a meaningful analysis based solely on returns (see Fung and Hsieh [2004]). Over a three-year interval, the portfolio exposures and risks would typically change significantly—an effect that the static beta obviously cannot catch.

BIAS

One natural concern in interpreting the results is the special and relatively small basket of hedge funds we are dealing with (compared to the whole universe) and the relatively short time period for which we have data. When managers are not in the PAAMCO portfolio (and hence not in the book), we do not have their positions, and hence cannot compute their betas and alphas.

Managers who are hired have been heavily screened, and those performing poorly (after being hired) are more likely to be fired and therefore eliminated from the book.

We should note that many of the data biases that traditionally plague hedge fund performance research are not an issue for us. For example, selection bias in hedge fund databases or indexes is not relevant here—we have all the managers in the fund of funds. Survivorship bias in hedge fund databases is also not an issue here, although we do worry about retention bias in the fund of funds. Similarly, there is no backfill bias here—we use only actual returns and positions while the hedge fund managers are in the book.

We can divide our bias concerns into two categories, which we address separately.

Hiring Bias

The hedge funds in our portfolio are a very special and small subset of the hedge funds universe; they have survived due diligence and have been hired (at least for a while) by a major institutional fund of funds. We first note that this is not a bias as far as the book benchmark alphas are concerned because they are by construction relative to the book, and the individual manager alphas add up to zero. Within every subgroup of hedge funds—even the very best or the very worst—one would have on average positive or negative alpha managers.

On the other hand, the hiring bias should manifest itself in the predictability of returns. If one manages to create a priori a portfolio with high and consistent subsequent positive returns, the correlations between the in-sample and out-of-sample returns are obviously biased upward (past returns are predictive). In our results, we find fairly weak evidence of return predictability. Further, by using the predictability of returns as a benchmark for judging the predictability of alphas, we are being conservative as far as any hiring bias goes.

A less significant but still an important concern is that some managers (new managers) were hired after the beginning of our sample (January 2004), and these managers may differ in some respects from the managers who had been hired before (old managers). As noted above, 353 monthly data points are missing because managers' data were not present in January 2004.

We deal with this by noting that some of our analysis focuses solely on the old managers. For example, our results in Exhibit 4 for managers with 18 months of in-sample data and 9 months of out-of-sample data implicitly use only old managers. These results also show alpha persistence.

Retention Bias

Within our special small subset of hedge funds, the longer a hedge fund remained a part of the fund of funds (i.e., did not exit), the more likely it is to have consistent positive returns (and possibly alphas). We note above that there are 390 monthly observations missing because managers were fired or went out of business. This potentially creates a bias in the predictability of returns, and may account for the increased correlations of in-sample versus out-of-sample returns of the managers when the length of the periods increases (see Exhibit 4). For example, when the in-sample period increases from six months to nine months, the correlation between in-sample and out-of-sample returns increases from about zero to about 20%.

Our analysis has two features that address the retention bias. First, some of our time horizons are deliberately chosen to be short because for shorter in-sample and out-of-sample intervals, retention bias is less of a concern. Specifically, retention bias occurs when the measurement interval we choose for our empirical study is longer than the interval the fund of funds uses in deciding whether to fire managers. While we do not know the exact interval chosen by the fund of funds, we are less likely to exceed it when we use shorter in-sample and out-of-sample periods. We find alpha predictability at even a six-month in-sample period.

Second, we use as a benchmark for alpha predictability the predictability of the corresponding returns. There is not only a highly significant correlation of in-sample and out-of-sample alphas, but this correlation is higher than the correlation of the corresponding in-sample and out-of-sample returns (see Exhibits 4 and 5). Thus, even if retention bias causes return predictability, we demonstrate alpha predictability over and above return predictability.

An additional feature of our results is worth noting. We have shown that alphas render superior information and are more predictive for shorter time periods—i.e., for in-sample periods of six to nine months—than returns. These are the relevant time scales within which managers of funds of funds would normally like to make their decisions about retaining hedge fund managers.

From a practitioner's point of view, book alphas have significant value-added as additional predictive information. Having managers' explicit positions compensates for the lack of a long return history.

CONCLUSION

We develop a new approach for the evaluation of fund managers within a portfolio (book). Using the explicit positions of the funds and the positions of the overall portfolio, we decompose each fund's return into a book beta and book alpha component. We test this book benchmark analysis on a subportfolio of equity-based hedge funds in the PAAMCO proprietary portfolio during a 31-month period.

Our results indicate that book benchmark alphas are significantly more predictive than returns when the in-sample periods are short (six to nine months). This suggests that book benchmark alphas are a valuable quantitative tool for managing a portfolio of hedge funds with position-level transparency.

While our analysis is developed for a fund of hedge funds because of data considerations, the book benchmark analysis is more general. It can be used in any situation involving manager selection as long as there is position-level transparency. For example, the approach could be equally useful in evaluating long-only fund managers.

At this stage, we have only scratched the surface of what might be done with the book benchmark analysis. Important future steps would be to fully develop the formal theoretical framework for book benchmark alphas and to extend the analysis empirically. This could help us understand how much dispersion in the alphas is needed in order to draw meaningful conclusions. Also, it would be natural to incorporate style factors and investigate whether market environments affect the predictability of book benchmark alphas. Finally, it would be interesting to study whether managers have predictable beta-timing skill in the context of the book benchmark analysis.

APPENDIX

We use the standard notation $E[X], \sigma[X], IR[X] = \frac{E[X]}{\sigma[X]}$ to denote expected value, standard deviation (volatility), and the information ratio of a random variable of returns X. Consider a portfolio of hedge funds (the book) with (a random variable of) returns R^B. Our problem is to hire (or fire) a manager with (a random variable of) returns R and predefined weight w, which is positive if we hire the manager (and negative if we fire the manager). Assume that based on the manager's and the book's recent positions, we have been able to quantify the manager's beta and alpha relative to the book:

$$R = \beta R^B + \alpha + \varepsilon$$

where $E[\varepsilon] = 0, \varepsilon \perp R^B$.

In this notation, α is a constant equal to the expected value of the manager's alpha to the book, and ε is the random term with volatility $\sigma[\varepsilon]$. We will first establish a simple approximation for the change in the information ratio of our book $IR[R^B + wR] - IR[R^B]$ as a result of the hiring/firing decision.

Proposition

If

1. $\sigma[R^B] \sim \sigma[R]$ *(i.e., the two volatilities are of the same order of magnitude),*
2. $w \ll 1$,
3. $|\beta| \sim 1$ *(i.e., beta is not an order of magnitude higher than 1 in absolute value), and β is a constant, then:*

$$IR[R^B + wR] - IR[R^B] = \frac{1}{\sigma(R^B)}\left[w\alpha + O(w^2)\right]$$

That is, in first order of the manager's weight, the change of the Sharpe ratio of the book is proportional to α and does not depend on $\sigma[\varepsilon]$.

Proof

The assumptions imply that either $\sigma[\varepsilon] \ll \sigma[R^B]$ or $\sigma[\varepsilon] \sim \sigma[R^B]$. Hence, from the definition of α and ε, one has:

$$\begin{aligned} IR[R^B + wR] &= \frac{E[R^B + wR]}{\sigma[R^B + wR]} \\ &= \frac{(1+w\beta)E[R^B] + w\alpha}{\sqrt{(1+w\beta)^2\sigma[R^B]^2 + w^2\sigma[\varepsilon]^2}} \\ &= \frac{(1+w\beta)E[R^B] + w\alpha}{(1+w\beta)\sigma[R^B]\sqrt{1 + w^2\dfrac{\sigma[\varepsilon]^2}{(1+w\beta)^2\sigma[R^B]^2}}} \\ &= IR[R^B] + \frac{1}{\sigma(R^B)}(w\alpha + O(w^2)) \end{aligned}$$

Q.E.D.

The assumptions are in place for a typical diversified portfolio of hedge funds (except perhaps the constancy of beta): The volatility of individual managers should not exceed three to four times the volatility of the overall book; the target weight should not exceed 5%; and the absolute value of beta to the book should not exceed 3. For less diversified portfolios, as in a multistrategy hedge fund portfolio, the higher-order terms will be important and need careful consideration.

If there is anything surprising about this proposition, it is that under these assumptions, the volatility of the manager's alpha enters the information ratio differential as a second-order term, and one can hence focus on alphas' expected value alone.

ENDNOTES

The authors thank Jim Berens, Jane Buchan, Masahiro Fujimoto, Philippe Jorion, and Dennis Logue for suggestions and comments. A special thanks to Clark Kibler for assistance with the data.

[1]These are equal-weighted returns on a small subset of PAAMCO's portfolio, and hence the book described here is not representative of either the PAAMCO equity managers or the overall PAAMCO fund of funds.

[2]For example, if January 2004 through June 2004 is the first in-sample period and July 2004 through September 2004 is the first out-of-sample period, then the second in-sample period would be April through September 2004 and the second out-of-sample period would be October through December 2004, and so on. For the first group, we have tightened our third requirement so that the in-sample period has at least four observations to minimize the possibility that our results are contaminated by using observations with a small amount of data.

[3]For groups 1–4, the in-sample data periods overlap, possibly spuriously increasing autocorrelations. This effect should be most pronounced for group 1 because it uses six months of data for the in-sample period, and as many as three months may overlap with another observation. The first-order autocorrelation of the in-sample alphas for group 1 is 47%. If we eliminate the overlap by requiring each new observation to start at the end of the previous observation, the first-order autocorrelation drops to 26%. For the in-sample returns, the first-order autocorrelations with and without overlapping are 26% and 9%, respectively. Hence the impact of the data period overlap on the in-sample alphas and in-sample returns is comparable—around a 20% increase in first-order autocorrelation for both. This suggests that the better predictability of alphas is not a spurious result of the particular choices of data periods.

REFERENCES

Agarwal, V., and N. Naik. "Multi-Period Performance Persistence Analysis of Hedge Funds." *Journal of Financial and Quantitative Analysis*, 35 (2001), pp. 327-342.

Amenc, N., S. El Bied, and L. Martellini. "Predictability in Hedge Fund Returns." *Financial Analysts Journal*, 59 (2003), pp. 32-46.

Brown, S., W. Goetzmann, and R. Ibbotson. "Offshore Hedge Funds: Survival and Performance 1989-1995." *Journal of Business*, 72 (1999), pp. 91-117.

Edwards, F., and M. Caglayan. "Hedge Fund Performance and Manager Skill." *Journal of Futures Markets*, 21 (2001), pp. 1003-1028.

Fung, W., and D. Hsieh. "Hedge Fund Benchmarks: A Risk-based Approach." *Financial Analysts Journal*, 60 (2004), pp. 65-80.

Jagannathan, R., A. Malakhov, and D. Novikov. "Do Hot Hands Persist Among Hedge Fund Managers? An Empirical Evaluation." NBER Working Paper 12015, 2006.

Kat, H., and F. Menexe. "Persistence in Hedge Fund Performance: The True Value of a Track Record." *Journal of Alternative Investments*, 5 (2003), pp. 66-72.

Schneeweis, T., H. Kazemi, and G. Martin. "Understanding Hedge Fund Performance: Research Issues Revisited—Part I." *Journal of Alternative Investments*, 4 (2002), pp. 6-22.

——. "Understanding Hedge Fund Performance: Research Issues Revisited—Part II." *Journal of Alternative Investments*, 5 (2003), pp. 8-30.