THE JOURNAL OF FINANCE • VOL. LXIII, NO. 4 • AUGUST 2008

Hedge Funds: Performance, Risk, and Capital Formation

WILLIAM FUNG, DAVID A. HSIEH, NARAYAN Y. NAIK, and TARUN RAMADORAI*

ABSTRACT

We use a comprehensive data set of funds-of-funds to investigate performance, risk, and capital formation in the hedge fund industry from 1995 to 2004. While the average fund-of-funds delivers alpha only in the period between October 1998 and March 2000, a subset of funds-of-funds consistently delivers alpha. The alpha-producing funds are not as likely to liquidate as those that do not deliver alpha, and experience far greater and steadier capital inflows than their less fortunate counterparts. These capital inflows attenuate the ability of the alpha producers to continue to deliver alpha in the future.

HEDGE FUNDS ARE LIGHTLY REGULATED active investment vehicles with great trading flexibility. They are believed to pursue highly sophisticated investment strategies, and promise to deliver returns to their investors that are unaffected by the vagaries of financial markets. The assets managed by hedge funds have grown substantially over the past decade, increasingly driven by portfolio allocations from institutional investors.[1] Hedge funds have also been attracting attention from academics, who have recently documented several interesting facts. First, a large proportion of the variation in hedge fund returns can be explained by market-related factors (see, for example, Fung and Hsieh (1997, 2001, 2002, 2004a, 2004b), Agarwal and Naik (2004), and Hasanhodzic and Lo (2006)).[2] This suggests that hedge fund fees (which are often substantial fractions of the total returns earned by funds, not their alphas) provide compensation for taking on systematic risk rather than for exploiting arbitrage

*Fung and Naik are at London Business School, Hsieh is at Duke University, and Ramadorai is at Said Business School at the University of Oxford, the Oxford-Man Institute for Quantitative Finance, and CEPR. We thank Omer Suleman and especially Aasmund Heen for excellent research assistance. We also thank an anonymous referee, Robert Stambaugh (the editor), John Campbell, Mila Getmansky, Harry Markowitz, Ludovic Phalippou, Tuomo Vuolteenaho, and seminar participants at the London School of Economics, University of Oxford, Stockholm Institute for Financial Research, Warwick University, the IXIS-NYU hedge fund conference, University of Massachusetts at Amherst, the 2006 Western Finance Association annual meetings, the 2006 INQUIRE-UK conference, the 2006 Northwestern Conference on Hedge Funds, and the 2007 American Finance Association annual meetings for comments. We gratefully acknowledge support from the BSI Gamma Foundation and from the BNP Paribas Hedge Fund Centre at the London Business School.

[1] According to the TASS Asset Flows report, aggregate hedge fund assets under management have grown from US $72 billion at the end of 1994 to more than US $670 billion at the end of 2004.

[2] See Agarwal and Naik (2005) for a comprehensive survey of the hedge fund literature.

opportunities.[3] Second, capital inflows to hedge funds are positively related to their total returns (see Agarwal, Daniel, and Naik (2004)). Finally, the evidence indicates that larger funds perform worse than smaller funds (see Getmansky, Lo, and Makarov (2004)).

These recent findings raise some important questions. First, does the average fund have any ability to deliver alpha? If not, then can we find any funds that are capable of delivering alpha? Are there distinct differences between the attributes of such alpha-producing funds, and those of funds that deliver returns solely on account of systematic risk exposures? Do investors treat these groups of funds differently when making capital allocation decisions, and do capital inflows adversely affect the ability of funds to deliver alpha in the future? In this paper, we employ data on a large cross-section of hedge funds to shed light on these questions.

Before using data on hedge funds, one must minimize the well-documented biases in the data. These biases arise from a lack of uniform reporting standards (hedge funds have not historically been regulated). For example, hedge fund managers can elect whether to report performance at all; if they do, they can decide the database(s) to which they report. They can also elect to stop reporting at their discretion. This ability to self-report biases hedge fund returns upward.[4] Furthermore, real-life constraints make reported hedge fund returns difficult to obtain in practice. For example, some hedge funds may report their returns to a database despite being closed to new investments, or returning money to investors. Exiting hedge fund investments may be similarly problematic—funds often impose constraints on the withdrawal of capital, using lockup periods, redemption, and notice periods. Finally, investors experience significant costs when accessing hedge funds, such as search and due diligence costs, which are hard to measure (see Ang, Rhodes-Kropf, and Zhao (2005) and Brown et al. (2007)).

Fung and Hsieh (2000) suggest a technique to mitigate the biases in hedge fund data. They recommend using data on funds-of-funds (hedge funds that invest in portfolios of other hedge funds) rather than data on individual hedge funds, arguing that fund-of-fund returns are a more accurate representation of the returns earned by hedge fund investors. As an example, consider the case of a hedge fund that is about to liquidate. Such a fund has an incentive to stop reporting to data vendors several months before the liquidation event. Consequently, the hedge fund's return history in the database will not reflect the full extent of losses incurred by its investors. In contrast, a fund-of-funds investing in the same hedge fund has a higher chance of surviving the collapse of one of the investments in its relatively more diversified portfolio. Thus, the fund-of-fund's return more accurately reflects the losses experienced by investors in the underlying hedge fund (albeit indirectly). Furthermore, fund-of-fund returns reflect the cost of real-life constraints involved in hedge fund investment,

[3] Alpha measures the average return accrued over and above compensation for exposure to different sources of systematic risk. See Berk and Green (2004) for one theoretical model in which the ability of managers is measured by their alpha.

[4] See, for example, Fung and Hsieh (2000) and Liang (2000) for more in-depth insights into the potential measurement errors that can arise as a result of voluntary reporting.

such as the constraints or delays that hedge funds impose on the withdrawal of capital. Finally, fund-of-fund returns reflect the costs of managing a portfolio of underlying hedge funds, as they are reported net of an additional layer of fees.[5] Accordingly, we consolidate data on funds-of-funds from the major database vendors. After carefully removing duplicates and filtering the data, our final sample comprises 1,603 funds-of-funds over the period from January 1995 to December 2004. These data constitute the most comprehensive set of funds-of-funds used in the academic literature to date.

Using these data, we first examine whether the average fund delivers alpha. Clearly, our estimates of alpha will be inaccurate if the risk exposures of funds change over time and we do not account for this fact (see Admati and Ross (1985)). Therefore, we extend the standard results in the hedge fund literature by testing for the presence of structural breaks in hedge fund risk exposures. We identify two such structural breaks, which we associate with major market events: the Long Term Capital Management (LTCM) crisis in 1998, and the NASDAQ crash in early 2000. Using these breakpoints, we estimate the average alpha of fund-of-funds over three subperiods demarcated by the months of September 1998 and March 2000. We find that over our 120-month sample period, the average fund-of-funds delivered positive and statistically significant alpha only in the 18-month subperiod between October 1998 and March 2000.

Despite the apparent scarcity of alpha at the average level, are there funds capable of delivering alpha in the cross-section? Using robust bootstrap techniques and the Fung and Hsieh (2004a) seven-factor model, we show that such funds do exist in our data. We segregate these alpha producers (have-alpha funds) from the remainder (beta-only funds). We find that the have-alphas possess quite different properties from beta-only funds. In particular, the have-alpha funds exhibit far lower liquidation rates than beta-only funds. On average, 5 years after classification, only 7% of have-alpha funds are liquidated, compared to 22% of beta-only funds. This difference is highly statistically significant and suggests that alpha-producing funds are more likely to stay in business than the remainder. Furthermore, the behavior of capital inflows to the two groups of funds suggests that hedge fund investors discriminate on the basis of alpha when allocating capital to funds. Have-alphas receive far greater inflows of capital than beta-only funds. Over the period from 1997 to 2004, capital flows into have-alpha funds grew at an average annual rate of approximately 30%, compared to the 8% growth of capital flows to beta-only funds. Furthermore, the capital flows into have-alphas are steady and do not significantly respond to recent past returns, while the flows into beta-only funds are characterized by return-chasing behavior. This suggests that there may be different groups of investors providing capital to have-alpha and to beta-only funds, a possibility that we discuss later in the paper.

This leads to our final question: Do capital inflows adversely affect the ability of alpha producers to deliver alpha in the future? We find strong evidence in

[5] The recent availability of hedge fund investable indices illustrates this point: These indices have returns that are far lower than those of the reported average hedge fund indices. In comparison, the reported average fund-of-funds indices are quite close in magnitude to the investable indices, and have a high correlation with them in recent years.

support of this conjecture. Have-alpha funds that experience relatively high capital inflows are less likely to be subsequently reclassified as have-alpha funds, while those experiencing lower capital inflows have a better chance of delivering alpha in the future. Capital inflows also affect the information ratio of funds, not just the propensity of funds to deliver alpha: Have-alpha funds that experience high (low) capital flows have a significantly lower (higher) t-statistic of alpha in the future.[6] Our evidence also suggests that capital inflows adversely impact alpha at the aggregate level: In the last few years of our sample, there has been a substantial increase in capital flows to the hedge fund industry. Simultaneously, the magnitude of alpha delivered by the average have-alpha fund has experienced a statistically significant decline.

Our findings are in line with the recent theoretical literature on active portfolio management. Berk and Green (2004) present a rational model with two key building blocks. First, managers have differential ability to generate risk-adjusted returns, but face decreasing returns to scale in deploying their ability. Second, investors learn about managerial ability from past risk-adjusted performance and direct more capital toward funds with superior performance. Our results indicate that the assumptions of their model are an accurate description of the prevailing conditions in the hedge fund industry. This has potentially important consequences: In Berk and Green's equilibrium, actively managed funds deliver zero risk-adjusted, after-fee returns to their investors. Although reality is far more complicated than the premises of any theoretical model, the decline in the aggregate level of alpha toward the end of our sample period suggests that the hedge fund industry may be headed in this direction.

The organization of the paper is as follows. Section I presents the data. Section II describes our methodology. Section III reports the results, and Section IV presents our conclusions.

I. Data

The main databases with data on funds-of-funds are HFR, CISDM, and Lipper TASS. We merge and consolidate data from all three databases and eliminate duplicate funds, as well as different fund share classes, which are created for regulatory and accounting reasons but are virtually identical to one another. Our final data set consists of 1,603 funds from January 1995 to December 2004.

Following the data vendors, we classify funds into three categories: alive and reporting, alive but stopped reporting, and liquidated. The data vendors provide this information for the majority of the funds in the databases. However, occasionally funds are classified as "defunct," which means that they either liquidated or stopped reporting, but the vendors do not provide information about which one of these events occurred. In such cases we inspect the assets under management (AUM) and returns of the funds in question. If the final AUM reported by a fund is very low relative to the maximum AUM over the fund's lifetime, and if the returns in the final months of the fund's history are below the industry average return, we classify the fund as liquidated; otherwise we classify

[6] The t-statistic of alpha is also known as the "information ratio" of a fund, a commonly employed performance measure in the investment management industry.

Table I

Summary Statistics

For each year represented in a row, the columns show the total number of funds-of-funds in the data at the end of the year, the number of funds that entered the data during the year, the number that were liquidated during the year, the number that stopped reporting during the year, the total AUM in billions of U.S. dollars of the funds alive at the end of each year, and the mean, median, and *SD* of the annual return at the end of the year across all funds.

Year	Number of Funds	Born	Liquidated	Stopped Reporting	Total AUM (U.S.$ BN)	Mean Return	Median Return	*SD* Return
1995	248	57	9	4	18.4	0.14	0.13	0.16
1996	336	107	13	6	26.0	0.15	0.15	0.09
1997	415	103	19	5	43.0	0.17	0.16	0.11
1998	487	111	22	17	37.6	0.00	0.02	0.15
1999	575	124	23	13	42.3	0.24	0.20	0.20
2000	657	127	26	19	49.4	0.08	0.10	0.14
2001	763	167	36	25	60.5	0.05	0.06	0.08
2002	898	174	25	14	77.8	0.02	0.02	0.07
2003	1036	208	48	22	123.1	0.12	0.10	0.11
2004	1158	203	35	46	194.6	0.07	0.07	0.04

it as a fund that is alive but has stopped reporting. For each fund classified using our procedure, we cross-check our classification with industry sources.

Table I presents descriptive statistics on our consolidated data. (Note that all the return data we employ is net of all fees and management costs.) First, mirroring the growth in AUM in the hedge fund industry, the AUM in funds-of-funds has grown from US $18 billion at the end of 1995 (around 25% of total AUM in the hedge fund industry according to the TASS asset flows report) to around US $190 billion in 2004 (close to 30% of the industry). Second, the data exhibit time variation in birth, liquidation, and closing rates. The average birth rate is 27%, the average liquidation rate is 4.7%, and the average rate of funds that stopped reporting despite being alive is 2.7% per year. Third, the equally weighted mean return across funds is 10.3% over our sample period. However, these returns vary substantially both within and across years. For example, in 1998, the average return of the funds in our data is zero, which is unsurprising given the cataclysmic events that occurred during that year.

II. Methodology

A. Risk-Adjusted Performance Evaluation

Throughout our analysis, we model the risks of funds using the seven-factor model of Fung and Hsieh (2004a). These seven factors have been shown to have considerable explanatory power for fund-of-fund and hedge fund returns.[7] The set of factors comprises: the excess return on the S&P 500 index (*SNPMRF*); a small minus big factor (*SCMLC*) constructed as the difference between the

[7] See Fung and Hsieh (2001, 2002, 2004a, 2004b). Agarwal and Naik (2004) present a factor model that includes some of the same factors as the Fung-Hsieh model.

Wilshire small and large capitalization stock indices; the excess returns on portfolios of lookback straddle options on currencies (*PTFSFX*), commodities (*PTFSCOM*), and bonds (*PTFSBD*), which are constructed to replicate the maximum possible return to trend-following strategies on their respective underlying assets;[8] the yield spread of the U.S. 10-year Treasury bond over the 3-month T-bill, adjusted for the duration of the 10-year bond (*BD10RET*); and the change in the credit spread of Moody's BAA bond over the 10-year Treasury bond, also appropriately adjusted for duration (*BAAMTSY*).

B. Time Variation and Structural Breaks

A static analysis of the risk structure of fund returns is not appropriate if funds change their strategies over the sample period that we investigate. Fung and Hsieh (2004a) study vendor-provided fund-of-fund indices, and perform a modified version of the CUSUM test (Page (1954)) to find structural breaks in fund factor loadings. Following their analysis, we identify breakpoints with major market events, namely, the collapse of LTCM in September 1998 and the peak of the technology bubble in March 2000. We rigorously test for the validity of these prespecified break points in our data, using a version of the Chow (1960) test in which we replace the standard error covariance matrix with a heteroskedasticity-consistent covariance matrix (White (1980), Hsieh (1983)). In particular, we estimate the following regression:

$$R_t = \alpha_1 D_1 + \alpha_2 D_2 + \alpha_3 D_3 + (D_1 X_t)\beta_{D1} + (D_2 X_t)\beta_{D2} + (D_3 X_t)\beta_{D3} + \varepsilon_t$$

where $X_t = [SNPMRF_t \quad SCMLC_t \quad BD10RET_t \quad BAAMTSY_t \quad PTFSBD_t$

$$PTFSFX_t \quad PTFSCOM_t]. \tag{1}$$

Here, R_t is the (equally weighted) average excess return across all funds in month t, D_1 is a dummy variable set to one during the first period (January 1995 to September 1998) and zero elsewhere, D_2 is set to one during the second period (October 1998 to March 2000) and zero elsewhere, and D_3 is set to one during the third period (April 2000 to December 2004) and zero elsewhere. The X matrix comprises the seven factors in the Fung and Hsieh (2004a) model, described in detail in the previous subsection. Note that there are a total of 24 regressors in equation (1), including the dummy variables. This framework allows us to perform the Chow test and to estimate the alpha of the average fund-of-funds in the three subperiods after accounting for time-varying risk exposures. The next subsection describes the methodology we use to detect whether funds in the cross-section have the ability to produce alpha.

C. Cross-Sectional Differences in Funds

The previous subsection uses the time series of average fund-of-fund returns to estimate alpha. In this subsection, we seek to identify cross-sectional

[8] See Fung and Hsieh (2001) for a detailed description of the construction of these primitive trend-following (PTF) factors.

differences in fund alpha. We adopt the perspective of a hypothetical investor contemplating a hedge fund investment. This investor infers the ability of a manager by evaluating a fund's performance over the past 2 years, selects funds that exhibit superior performance, and directs capital toward them. At annual intervals, he rebalances the portfolio, reassessing the available investment opportunity set and selecting funds afresh. We implement this exercise as follows: Each year, we select all funds in the data with a full return history over the previous 2-year period. (For example, at the end of 1996, we select all funds with a complete return history between January 1995 and December 1996.) Using Fung and Hsieh's (2004a) seven-factor model and the nonparametric procedure of Kosowski et al. (2006) (see the Appendix for details), we identify the funds that deliver significantly positive alpha and segregate them from those that do not.[9] As previously mentioned, we denote the former set of funds as have-alpha funds and the remainder as beta-only funds. We then repeat the alpha estimation exercise at the end of each year, using data from the most recent 2-year period. We choose a 2-year window in an effort to both ensure sufficient degrees of freedom to estimate alpha and to capture the propensity of fund risk exposures to vary over time. Note that our selection procedure could result in a change in the identities of the have-alpha funds and beta-only funds each year, depending on births, deaths, and the risk-adjusted performance of funds over the prior 2 years.

D. Capital Flow Analysis

We follow Sirri and Tufano (1998) and others, and construct the quarterly net flow of capital into each of the funds in our sample as:[10]

$$F_{iq} = \frac{AUM_{iq} - AUM_{iq-1}(1 + R_{iq})}{AUM_{iq-1}}. \tag{2}$$

Here, F_{iq}, AUM_{iq}, and R_{iq} are respectively, the flows for a fund i in quarter q expressed as a percentage of lagged AUM; the AUM of the fund i in quarter q; and the returns of fund i in quarter q. We winsorize the quarterly flows across all funds each quarter at the 1st and 99th percentiles to attenuate the effect of outliers. We then estimate the relationship between flows, past flows, and past returns using the following regression:

$$F_{gq} = \gamma_{g0} + \gamma_{gr} R_{gq-1} + \gamma_{gf} F_{gq-1} + u_{gq}. \tag{3}$$

9 We also experiment with imposing parametric structure on the serial correlation of the residuals (we do this nonparametrically using the Politis and Romano (1994) stationary bootstrap) by applying the Getmansky, Lo, and Makarov (2004) correction to undo any potential autocorrelation in fund returns. The results of our bootstrap experiments are qualitatively unaffected by the use of this procedure. All of these results are available on request.

10 Flows are computed under the assumption that they come in at the end of each quarter, once fund returns are accrued. We also experiment with the assumption that flows come in at the beginning of each quarter; our results are invariant to this assumption.

The quarterly flow measure F_{gq} is regressed on lagged quarterly flows F_{gq-1} and lagged quarterly returns R_{gq-1}. This regression is estimated separately for each subgroup g of funds-of-funds (g can be have-alpha or beta-only). We employ a Newey and West (1987) covariance matrix using four quarterly lags to account for any possible autocorrelation and heteroskedasticity in the residuals.

In another exercise, we investigate the relationship between capital flows and different measures of performance separately for have-alpha and beta-only funds. While Sirri and Tufano (1998) investigate the return–flow relationship for mutual funds, we have an additional dimension of performance to consider in the case of hedge funds, namely, alpha. Therefore, in addition to returns, we examine the relationship between capital flows, the magnitude of alpha, and the t-statistic of alpha. We implement this exercise as follows. Each year, we sort the three different measures of performance (returns, alpha, and the t-statistic of alpha) into quintiles across all funds in each group. We then estimate the following regression specification within these quintiles separately for have-alpha and beta-only funds:

$$F_{iy}^{q} = Base^{q} + \phi^{q} PerfRank_{iy-1}^{q} + u_{iy}^{q}. \tag{4}$$

Here, the annual flow measure F_{iy}^{q} for a fund i in year y and performance quintile q is computed as a percentage of end-of-previous-year AUM. The regression intercept $Base^{q}$ is the baseline capital flow for the quintile, and $PerfRank_{iy-1}^{q}$ is the relative performance ranking of a fund within its performance quintile q in year $y-1$, computed in ascending order, and ranges between zero and one. For example, if a fund i is the top-ranking (median) fund in the top-performing quintile of have-alpha funds in year $y-1$, $PerfRank_{iy-1}^{5}= 1(0.5)$.

E. Capacity Constraints

In order to uncover the relationship between capital flows and subsequent risk-adjusted performance, we condition the future performance of have-alpha and beta-only funds on the level of capital flows that they receive. In particular, in each classification period we compute the average quarterly flow experienced by all funds in the final year of that classification period. We then sort and group funds into two subcategories, based on whether they receive above-median or below-median capital flows. We examine the future performance of these subcategories of funds (in each of the have-alpha and beta-only groups) in three ways. First, we inspect the transition probabilities of these funds, that is, the probability that above-median-flow and below-median-flow funds are subsequently reclassified as have-alpha funds and beta-only funds in the next nonoverlapping 2-year period. Second, we compute the average level of alpha in the subsequent 2-year nonoverlapping period for the two subcategories of funds. Finally, and analogously, we compute the average t-statistic of alpha for the two subcategories of funds to ascertain whether the information ratio for a fund is affected by its level of capital flows. We estimate the statistical significance of the difference in these metrics for the two subcategories using the

Wald test and a cross-correlation and heteroskedasticity-consistent covariance matrix, which is computed using the method of Rogers (1983, 1993).

F. Is Alpha Changing over Time for Have-Alpha and Beta-Only Funds?

Based on the results from our analysis of capital flows in subsection D and capacity constraints in subsection E, we might expect to find changes in alpha production for the average have-alpha member over time if capacity constraints are beginning to bite at the industry level. We therefore construct equally weighted portfolios of have-alpha and beta-only funds from January 1997 to December 2004, which track their returns in the year *after* they were classified.[11] This means that the composition of the portfolio could change from year to year as our selection procedure picks different have-alpha and beta-only funds each year. Furthermore, all performance evaluation is completely out-of-sample. For example, some of the funds selected in the 1995 to 1996 period may become defunct during the performance evaluation period of 1997, in which case the 1997 returns would not incorporate their returns from the point at which they exit the data. We then rerun equation (1), successively replacing the average fund return on the left-hand side with the have-alpha and beta-only portfolio returns:

$$\begin{aligned} R_{gt} &= \alpha_{g1}D_1^1 + \alpha_{g2}D_2 + \alpha_{g3}D_3 + \left(D_1^1 X_t\right)\beta_{gD1} + (D_2 X_t)\beta_{gD2} \\ &\quad + (D_3 X_t)\beta_{gD3} + \nu_{gt} \\ \text{where } X_t &= [\mathit{SNPMRF}_t \quad \mathit{SCMLC}_t \quad \mathit{BD10RET}_t \quad \mathit{BAAMTSY}_t \\ &\quad \mathit{PTFSBD}_t \quad \mathit{PTFSFX}_t \quad \mathit{PTFSCOM}_t]. \end{aligned} \tag{5}$$

Here, the subscript g denotes the group, namely, have-alpha or beta-only. The main difference between equations (1) and (5) is that D_1^1 is a dummy variable that is now set to one between January 1997 and September 1998, and zero elsewhere. This difference reflects the fact that we employ data between 1995 and 1996 to identify the constituents of the have-alpha and beta-only portfolios in 1997. The dummy variables D_2 and D_3 are as previously defined in equation (1).

III. Results

A. Risk-Adjusted Performance Evaluation and Time Variation

Table II reports the results from estimating equation (1). The rows of Table II list the explanatory variables, and the columns report the subperiods over which they are estimated. We test whether the vectors of

[11] Fung and Hsieh (2006) conduct a similar exercise using a rolling regression. While their focus is on minimizing the one-step ahead forecast errors of the seven-factor model, ours is on comparing the differences between have-alpha funds and have-beta funds.

Table II

The Changing Risks of Funds-of-Hedge-Funds

The top panel of this table contains estimates of:

$$R_t = \alpha_1 D_1 + \alpha_2 D_2 + \alpha_3 D_3 + (D_1 X_t)\beta_{D1} + (D_2 X_t)\beta_{D2} + (D_3 X_t)\beta_{D3} + \varepsilon_t$$

where $X_t = [SNPMRF_t\ \ SCMLC_t\ \ BD10RET_t\ \ BAAMTSY_t\ \ PTFSBD_t\ \ PTFSFX_t\ \ PTFSCOM_t]$.

Here, R_t is the (equally weighted) average annualized excess return across all funds in month t, D_1 is set to one during the first period (January 1995 to September 1998) and zero elsewhere, D_2 is set to one during the second period (October 1998 to March 2000) and zero elsewhere, and D_3 is set to one during the third period (April 2000 to December 2004) and zero elsewhere. The regressors X are described in the text. The bottom panel contains estimates of Chow structural break test Chi-squared statistics. White heteroskedasticity-consistent standard errors are reported below the coefficients. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Period I	Period II	Period III
Constant	0.0009	0.0093**	0.0006
	0.0011	*0.0018*	*0.0008*
SNPMRF	0.2866***	0.1314***	0.1388***
	0.0323	*0.0404*	*0.0177*
SCMLC	0.1371***	0.2993***	0.1289***
	0.0637	*0.0309*	*0.0199*
BD10RET	−0.0169	0.4799***	0.1603***
	0.1203	*0.1288*	*0.0303*
BAAMTSY	0.7160***	0.6376***	0.1421**
	0.1503	*0.1630*	*0.0622*
PTFSBD	0.0062	0.0576***	−0.0018
	0.0119	*0.0170*	*0.0027*
PTFSFX	0.0111**	−0.0199***	0.0140***
	0.0047	*0.0092*	*0.0050*
PTFSCOM	0.0269***	−0.0140**	0.0120
	0.0078	*0.0044*	*0.0075*
Adjusted R^2	0.737		
Number of months	120		
Period for Chow Structural Break Test			
Test for Period I & II Break			
Chi-sq(14)			248.42***
Test for Period I Break			
Chi-sq(7)			86.70***
Test for Period II Break			
Chi-sq(7)			82.75***

coefficient estimates $\hat{\beta}_{D1}$ and $\hat{\beta}_{D2}$ are jointly different from $\hat{\beta}_{D3}$ using the heteroskedasticity-consistent covariance matrix. The χ^2 test statistic with 14 degrees of freedom is 248.4, indicating a strong rejection of the null hypothesis that the slope coefficients are the same across the three subperiods. This confirms that the risk exposures of funds change over time. Furthermore, the structural break points that we use, namely, September 1998 and March 2000, are strongly supported by the data, and we can easily reject the null hypotheses of no structural break in periods I and II.[12] The results in Table II also indicate that the average fund-of-funds only exhibits statistically significant alpha during the second subperiod ($\widehat{\alpha}_2$ is the only statistically significant intercept), which spans the bull market period between October 1998 and March 2000. Furthermore, the adjusted-R^2 statistic is around 74% for the returns of the average fund. The magnitude of the R^2 statistic suggests that funds take on a significant amount of factor risk. This confirms the results extensively documented in the literature.

We check the robustness of our results in a number of ways. First, we replace the three PTF factors with the Agarwal and Naik (2004) out-of-the-money put option on the S&P 500. We expand the set of factors, incorporating the excess returns on the NASDAQ technology index. As in Asness, Krail, and Liew (2001), we add in lagged values of the factors, one at a time. Finally, we correct individual fund returns for return-smoothing using the Getmansky et al. (2004) correction. None of these robustness checks qualitatively affects our conclusions.

Our results underscore the fact that identifying time variation in factor loadings is important when evaluating the risk-adjusted performance of hedge funds. The results show that the average fund does not deliver alpha either in period I or in period III. However, inferences drawn from the average return series potentially hide important heterogeneity in the set of funds. Perhaps there are funds in our sample that consistently generate alpha, which we are unable to detect when we analyze the average return across all funds. We now turn to the results from our cross-sectional analysis.

B. Cross-Sectional Differences in Funds

The first three columns of Table III report the number of funds included in the bootstrap experiment in each 2-year period (all funds with 2 complete years of return history in each of the selection periods), and the percentage of the total number of funds in the have-alpha and beta-only groups. The first feature to note is that the number of funds selected in each 2-year period is steadily increasing over time. This is a reflection both of the increasing availability of data, and of the growth in the hedge fund industry. Second, on average across our sample period, 22% of the funds are classified as have-alpha funds, while a much larger percentage of funds are classified as beta-only funds.[13] Third, the percentage of total funds allocated to the have-alpha group fluctuates over time, ranging from a low of 10% at the end of 1998 to a high of 42% at the end

[12] Specifically, we separately estimated the results in Table II in incremental form from subperiod to subperiod. Here we find that the third period's factor loading estimates are statistically different from that of the first period in five of the seven factors. A similar comparison to the second period shows that six out of the seven factor loading estimates are statistically different. This incremental version of Table II is available from the authors on request.

[13] We check whether there were any funds that delivered statistically negative alpha in the set. In each classification period, we find that fewer than 5% of the funds in the set had this property. This number is lower than the significance level of our test. Therefore, we cannot reject the hypothesis that there are no "negative alpha" funds in our data.

Table III

Transition Probabilities of Have-Alpha and Beta-Only Funds

The rows show the 2-year period in which funds are classified as have-alpha and beta-only funds. The columns are, in order, the total number of funds with 2 full years of return history in each of the classification periods; the percentage of the total classified as have-alpha funds; the percentage of the total classified as beta-only funds; and the percentages of have-alpha and beta-only funds that are classified in the subsequent nonoverlapping period as have-alpha, *beta-only*, liquidated, or stopped reporting. For example, in 1996 to 1997, of 259 total funds, 34% and 66%, respectively, were classified as have-alpha and beta-only, 17% of these have-alpha funds were reclassified as have-alpha funds in the 1998 to 1999 period, 74% as beta-only funds, 5% as liquidated, and 5% as stopped reporting (numbers are rounded to the nearest percent). In contrast, 7% of the 1996 to 1997 beta-only funds were classified as 1998 to 1999 have-alpha funds, 73% were reclassified as beta-only funds, 13% as liquidated, and 6% as stopped reporting. The final rows report the Wald test statistics (using a cross-correlation and heteroskedasticity-robust covariance matrix) for the hypothesis that the have-alpha and beta-only transition probabilities are the same. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

		Proportion			P(2-Year Transition)			
Classification Period	Number of Funds	Have-Alpha	Beta-Only	From/To	Have-Alpha	Beta-Only	Liquidated	Stopped Reporting
1995–1996	195	0.21	0.79	Have-alpha	0.24	0.68	0.02	0.05
				Beta-only	0.04	0.82	0.10	0.04
1996–1997	259	0.34	0.66	Have-alpha	0.17	0.74	0.05	0.05
				Beta-only	0.07	0.73	0.13	0.06
1997–1998	307	0.10	0.90	Have-alpha	0.81	0.16	0.00	0.03
				Beta-only	0.26	0.58	0.09	0.07
1998–1999	374	0.17	0.83	Have-alpha	0.27	0.65	0.05	0.03
				Beta-only	0.18	0.62	0.12	0.09
1999–2000	448	0.42	0.58	Have-alpha	0.24	0.64	0.06	0.06
				Beta-only	0.09	0.70	0.12	0.09
2000–2001	506	0.22	0.78	Have-alpha	0.30	0.65	0.03	0.02
				Beta-only	0.10	0.77	0.08	0.04
2001–2002	584	0.17	0.83					
2002–2003	700	0.15	0.85					
All Cases	422	0.22	0.78	Have-alpha	0.28	0.65	0.04	0.03
				Beta-only	0.14	0.69	0.11	0.06
				Wald statistic	26.82***	3.27*	87.04***	8.65***

of 2000. This suggests that the ability of funds to deliver alpha is sensitive to market conditions.

The final four columns in Table III report transition probabilities for have-alpha and beta-only funds. The row headings indicate the 2-year period over which the funds are classified, while the column headings indicate the percentage of funds that are classified as have-alpha or beta-only in each nonoverlapping 2-year period, as well as the percentage of funds that have liquidated or stopped reporting over the period.[14] The results indicate that there is a greater chance for a fund to deliver alpha in the subsequent period if it is a member

[14] Note that the final period we consider is 2002 to 2003, since we require at least 1 year of out-of-sample data for our performance analysis, and several of the funds in the databases do not provide complete return histories for the year 2004. This data limitation also prevents us from reporting 2-year transition probabilities for the classification period 2001 to 2002.

of the have-alpha group to begin with. In particular, the last few rows of Table III show that the overall average transition probability for a have-alpha fund into the subsequent have-alpha group is 28%, while that for a beta-only fund is 14%. This difference is highly statistically significant. However, the average result hides the fact that the year-by-year alpha-transition probability for a have-alpha fund is always higher than that for a beta-only fund. In some years, the transition probability differential is much higher than the average. For example, in the 1997 to 1998 classification period (which includes the LTCM crisis), the 2-year transition probability for a have-alpha fund is 81%, in contrast to the 26% probability for a contemporaneously classified beta-only fund. Overall, there appears to be greater alpha persistence among members of the have-alpha group.

Table IV reports the percentage of have-alpha funds and beta-only funds that are liquidated at the end of each year over a 5-year postclassification period.

Table IV

Liquidation Probabilities of Have-Alpha and Beta-Only Funds

The rows correspond to the 2-year period in which funds are classified as have-alpha and beta-only funds. The columns indicate the proportion of have-alpha funds and beta-only funds that were liquidated 1, 2, 3, 4, and 5 years after the classification period. The top panel shows the liquidation probabilities for the have-alpha funds and the bottom panel for the beta-only funds. For example, in 1996 to 1997, of 259 total funds, 34% and 66%, respectively, were classified as have-alpha funds and beta-only funds (from the previous table), of these have-alpha funds, 3% were liquidated in 1998, 5% were liquidated by the end of 1999, and 10% were liquidated at the end of 2002. In contrast, for the beta-only funds classified in 1996 to 1997, 8% were liquidated in 1998, 13% were liquidated by the end of 1999, and 24% were liquidated at the end of 2002. The final rows report the Wald test statistics (using a cross-correlation and heteroskedasticity-robust covariance matrix) for the hypothesis that the have-alpha and beta-only liquidation probabilities are the same. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Classification Period	Year 1	Year 2	Year 3	Year 4	Year 5
		Have-Alpha Liquidation Probabilities			
1995–1996	0.00	0.02	0.02	0.05	0.07
1996–1997	0.03	0.05	0.09	0.10	0.10
1997–1998	0.00	0.00	0.00	0.03	0.03
1998–1999	0.00	0.05	0.11	0.11	
1999–2000	0.04	0.06	0.09		
2000–2001	0.03	0.03			
2001–2002	0.02				
Average	0.02	0.03	0.06	0.07	0.07
		Beta-Only Liquidation Probabilities			
1995–1996	0.06	0.10	0.16	0.19	0.22
1996–1997	0.08	0.13	0.19	0.23	0.24
1997–1998	0.06	0.09	0.16	0.18	0.19
1998–1999	0.08	0.12	0.15	0.16	
1999–2000	0.11	0.12	0.14		
2000–2001	0.05	0.08			
2001–2002	0.05				
Average	0.07	0.11	0.16	0.19	0.22
Wald statistic	36.15***	87.04***	42.39***	100.41***	53.57***

On average, we find that only 7% of have-alpha funds are liquidated 5 years after classification, while for the beta-only funds the comparable number is 22%. The difference is again highly statistically significant for every postclassification year. Clearly, have-alpha funds have a greater ability to avoid liquidation, regardless of the length of the postclassification period. These results are unchanged if we also control for the length of any individual fund's history prior to classification, suggesting that they are not driven by backfill bias. Taken together, the results in Tables III and IV suggest that there are significant differences in ability in the cross-section of funds, as measured by alpha. Alpha producers distinguish themselves both by their higher propensity to persistently deliver alpha, as well as their lower liquidation rates. But do investors treat alpha producers differently from the remainder of funds? The next subsection examines the relationship between the flow of capital and performance.

C. Capital Flow Analysis

Table V reports the equally weighted average annual flow into have-alpha and beta-only funds in each year following their classification. On average, the have-alpha funds experience a statistically significant inflow of 29.7% per annum in the year following classification, in contrast to the far lower inflows experienced by the beta-only funds. Indeed, the overall average level of flow for the beta-only funds is not statistically different from zero at the 10% level of significance. Figure 1 confirms this analysis. The figure is created by indexing

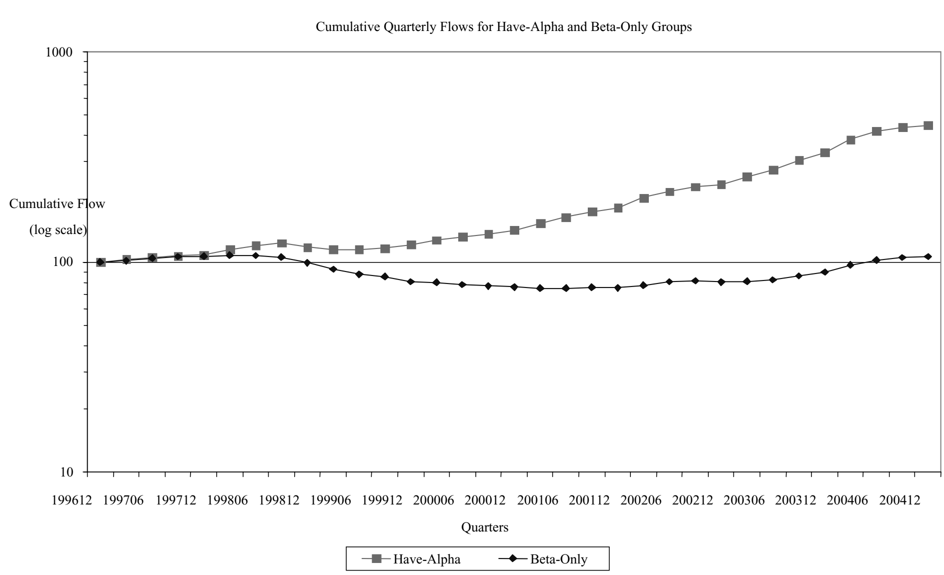

Figure 1. Cumulative flows for have-alpha and beta-only funds. The X-axis shows the month for which the flow index is plotted on a logarithmic scale on the Y-axis. The index begins at a value of 100 in December 1996 and successive values are given by $Index_{gq} = Index_{gq-1} * (1 + F_{gq})$, where F_{gq} is the flow percentage for group g (g is have-alpha or beta-only) for quarter q.

December 1996 to 100 and multiplying this level by the compounded growth in out-of-sample, equally weighted, quarterly flows each year for each group. To take a specific example, at the end of 1997, the have-alpha flow index takes on a value of 106.8, which is the product of the four quarterly, equally weighted flow observations in 1997 for the have-alpha funds selected in the period 1995 to 1996. The figure is shown on a logarithmic scale to accommodate the significant differences between the two groups, and shows that the have-alpha flow index reaches a level of 448 at the end of December 2004. In sharp contrast, the beta-only flow index ends up at a level of 106.

While Figure 1 shows a dramatic difference between capital inflows to the two groups of funds, it masks a more intriguing set of time patterns. Table V shows that have-alpha funds and beta-only funds experienced significant inflows in 1997, of 9.1 and 10.5% per annum, respectively. In 1998, the year of the

Table V

Flows into Have-Alpha and Beta-Only Funds

The rows correspond to the years in which funds are classified as have-alpha funds and beta-only funds. The columns report the average annual flow for the subsequent year across all funds in the group indicated in the column heading. Annual flows are computed as the product of quarterly flows, and quarterly flows are computed as increase in AUM less accrued returns, under the assumption that flows came in at the end of the quarter. For example, in 1996 to 1997, we classify funds as have-alpha funds and beta-only funds. The columns reveal that have-alpha funds experience an average inflow of 9.7% of end-1997 AUM over the subsequent year, 1998. White (1980) heteroskedasticity-consistent standard errors are reported below yearly estimates, and cross-correlation and heteroskedasticity-robust standard errors are reported below the pooled estimate. The final column reports the results from a hypothesis test that the have-alpha and beta-only flows are the same for each time period denoted in rows. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Classification Period	Have-Alpha Flows(t+1)	Beta-Only Flows(t+1)	Stat. Sig. Diff.?
1995–1996	0.091*	0.105**	—
	0.047	0.041	
1996–1997	0.097**	−0.061***	**
	0.040	0.021	
1997–1998	0.032	−0.174***	***
	0.045	0.020	
1998–1999	0.187***	−0.021	***
	0.050	0.019	
1999–2000	0.324***	0.041	***
	0.042	0.039	
2000–2001	0.349***	0.085***	***
	0.062	0.020	
2001–2002	0.483***	0.161***	***
	0.072	0.028	
2002—2003	0.404***	0.244***	***
	0.060	0.025	
Overall	0.297***	0.082	***
	0.044	0.050	

LTCM crisis, we see that the beta-only funds experienced significant outflows of 6.1%, as compared to the statistically significant 9.7% inflows experienced by the have-alpha funds. In the year following the LTCM crisis, beta-only funds continued to experience dramatic outflows of 17.4%, while for the have-alpha funds, there seems to be sufficient continuing interest to offset the impact of capital flight induced by the LTCM crisis. On net, this results in positive but statistically insignificant flows to the have-alpha funds in 1999. These results may be due to the fact that the LTCM crisis forced investors to look more carefully at the quality of funds. The same pattern continued in 2001, after the NASDAQ crash. Have-alpha funds experienced 32.4% inflows, while beta-only funds received 4.1% inflows in this year. Finally, between 2002 and 2004, although both groups saw significant inflows, the have-alpha funds on average received three times the inflows received by beta-only funds.

There are other important differences between the flows into have-alpha funds and beta-only funds. Table VI inspects the flow–return relationship for each group. The results here show that the flows into have-alpha funds show no evidence of return-chasing behavior—the coefficient of quarterly flows on lagged quarterly returns is not statistically significant. However, this is not true for the flows into the beta-only funds. For the beta-only funds, high (low) returns over a quarter precede statistically significant increases (decreases) in capital flows in the subsequent quarter. These results are consistent with a scenario in which unsophisticated positive-feedback investors are attracted to beta-only funds, and sophisticated investors who are able to detect the presence of alpha provide a steady stream of capital to have-alpha funds. According to press reports, the primary capital providers to the hedge fund industry can be divided into two distinct groups: high net worth individuals, and, more recently, institutional investors such as defined-benefit pension funds and university endowments.[15] There is a growing literature that suggests that institutional investors may be more sophisticated than individual investors (two recent examples are Cohen, Gompers, and Vuolteenaho (2001) and Froot and Ramadorai (2008)). This may account for the very different behavior of capital flows into have-alpha and beta-only funds.

Table VII analyzes the flow–performance relationship for have-alpha funds and beta-only funds in a more detailed fashion, presenting estimates of equation (4) separately for the two groups of funds. Panel A of the table reveals that the have-alpha funds in the top four return quintiles appear to experience greater baseline inflows than the bottom quintile of funds ranked on returns. In contrast, for the beta-only funds, the bottom quintile of funds experiences a baseline outflow of 15%, while the top quintile experiences a baseline inflow of 15.7%. This is a reiteration of the results from Table VI that capital flows

[15] The National Association of College and University Business Officers (NACUBO) shows that university endowments have increased their allocation to hedge funds from 6.1% of their endowment (US $14.4 BN) in 2001 to 16.6% in 2005 (US $49.6 BN). Over the same period, the top 200 defined benefit pension plans increased their allocation from US $3.2 BN to US $29.9 BN (source: www.pionline.com).

Table VI

Return Chasing in Have-Alpha and Beta-Only Funds

This table presents estimates of

$$F_{gq} = \gamma_{g0} + \gamma_{gr} R_{gq-1} + \gamma_{gf} F_{gq-1} + u_{gq}$$

estimated separately for each subgroup g of funds (g is have-alpha or beta-only). The quarterly flow measure F_{gq} in each case is expressed as a percentage of end-of-previous quarter AUM, and is regressed on lagged quarterly flows F_{gq-1} and lagged quarterly returns R_{gq-1}. The column headings indicate the subgroup g for which the equation is estimated. Newey-West autocorrelation and heteroskedasticity-consistent standard errors are presented below the coefficients, estimated using four quarterly lags. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Have-Alpha Flows	Beta-Only Flows
Intercept	−0.002	−0.005
	0.009	0.004
Ret (L1)	0.203	0.325***
	0.160	0.095
Flow (L1)	0.778***	0.809***
	0.078	0.074
Adjusted-R^2	0.466	0.717
Number of quarters	32	32

to beta-only funds exhibit trend-chasing behavior, while those into have-alpha funds do not. The added advantage of estimating equation (4) is that it allows us to control for the performance–flow relationship within each quintile at the same time that we estimate the baseline level of flows to each quintile. However, Table VII reveals that this within-quintile performance–flow relationship is almost always statistically insignificant for both have-alpha and beta-only funds. There are two exceptions to this finding: the bottom quintile of beta-only funds ranked on returns, in which the performance–flow relationship is positive, and the bottom quintile of have-alpha funds ranked on the t-statistic of alpha, in which the performance–flow relationship is negative.

Finally, Panel B of Table VII ranks have-alpha funds by the level and t-statistic of alpha. Inspecting the baseline flow coefficients, it appears that the response of flows into have-alpha funds may be nonlinearly related to both the level of alpha and the t-statistic of alpha. For example, the top quintile of funds ranked by the level of alpha receives higher baseline inflows than the fourth quintile of funds, although baseline flows to quintiles two, three, and four are not statistically distinguishable from one another.

D. Evidence on Capacity Constraints

Does the observed behavior of capital flows affect the ability of have-alpha funds to deliver alpha in the future? Table VIII conditions the 2-year transition

Table VII

The Flow–Performance Relationship in Different Performance Quintiles

Table VII presents estimates of $F_{iy}^{q} = Base^{q} + \phi^{q} PerfRank_{iy-1}^{q} + u_{iy}^{q}$. Panel A (B) presents results when the performance measure is returns (t-statistic of alpha and the level of alpha). The annual flow measure F_{iy} for a fund i in year y is computed as a percentage of end-of-previous-year AUM. The regression is run separately for each performance quintile. The $PerfRank_{iy-1}$ is computed in ascending order *within* each performance quintile, and ranges between zero and one. For example, in the column labeled "Top" in Panel B, in rows labeled "t-statistic of Alpha," the regression is run only for funds that are members of the highest t-statistic of alpha quintile of have-alpha funds in year $y-1$; and $PerfRank_{iy-1}$ is the relative rank of funds *within* this highest t-statistic of alpha quintile of have-alpha funds in year $y-1$. White (1980) heteroskedasticity-robust standard errors and the adjusted-R^2statistics are reported below the coefficient estimates. The rows labeled "Stat. Sig. Diff." report results from a Wald test of the hypothesis that the coefficient is the same as that in the previous quintile. For example, in column Q2, the symbol under the baseline row is for the test that the baseline flow estimated in Q1 is the same as in Q2, computed using a panel regression and a cross-correlation and heteroskedasticity-robust estimator. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Panel A:		Bottom	Q2	Q3	Q4	Top
Have-alpha Returns	Baseline	0.147**	0.403***	0.402***	0.271**	0.449***
		0.082	0.097	0.113	0.134	0.132
	Stat. sig. diff		*			***
	Perf. rank	−0.013	−0.126	−0.095	0.025	−0.125
		0.112	0.126	0.247	0.185	0.204
	Stat. sig. diff					
	Adjusted R^2	−0.007	−0.003	−0.005	−0.007	−0.005
Beta-only Returns	Baseline	−0.150***	0.032	0.085	0.125**	0.157**
		0.035	0.041	0.106	0.057	0.067
	Stat. sig. diff		***			
	Perf. rank	0.127***	0.013	0.029	0.099	0.049
		0.049	0.040	0.079	0.080	0.056
	Stat. sig. diff					
	Adjusted R^2	0.009	−0.002	−0.002	0.001	−0.002

Panel B:		Bottom	Q2	Q3	Q4	Top
Have-Alpha t-statistic of alpha	Baseline	0.484***	0.323***	0.157**	0.365***	0.297***
		0.148	0.071	0.073	0.078	0.084
	Stat. sig. diff			*	**	
	Perf. rank	−0.343*	−0.057	0.252	−0.104	−0.017
		0.178	0.133	0.192	0.112	0.089
	Stat. sig. diff					
	Adjusted R^2	0.027	−0.006	0.008	−0.004	−0.007
Have-alpha Level of alpha	Baseline	0.324***	0.359***	0.252***	0.238***	0.511***
		0.073	0.100	0.077	0.061	0.147
	Stat. sig. diff					**
	Perf. rank	−0.053	−0.038	0.034	0.037	−0.355
		0.118	0.127	0.113	0.102	0.257
	Stat. sig. diff					
	Adjusted R^2	−0.006	−0.007	−0.007	−0.007	0.018

Table VIII

Transition Probabilities for Above- and Below-Median Flow Have-Alpha Funds

The rows correspond to the 2-year period in which funds are classified as have-alpha funds. The columns are, in order: the group affiliation, that is, whether the average fund classified as a have-alpha experienced inflows (in the second year of the classification period) that were above or below the median have-alpha inflow in that year; the number of have-alpha funds in each group; the percentage of the flow group members classified as have-alpha funds in the subsequent classification period; the percentage of the flow group members classified as beta-only funds; the percentage of group members that were liquidated; and the percentage that stopped reporting. For example, in 1996 to 1997, 43 have-alpha funds had above the median inflow in 1997. Of these, 12% were classified as have-alpha funds, 79% as beta-only funds, 5% were liquidated, and 5% stopped reporting in 1998 to 1999. Percentages may not add up to 100 because of rounding error. The final rows report the Wald test statistics (using the cross-correlation and heteroskedasticity-robust covariance matrix) for the hypothesis that the above- and median-flow transition probabilities are identical. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

			P(2-Year Transition)			
Classification Period	Flow Group	Number of Funds	Have-Alpha	Beta-Only	Liquidated	Stopped Reporting
1995–1996	Above median	20	0.20	0.65	0.05	0.10
	Below median	21	0.29	0.71	0.00	0.00
1996–1997	Above median	43	0.12	0.79	0.05	0.05
	Below median	44	0.23	0.68	0.05	0.05
1997–1998	Above median	15	0.93	0.07	0.00	0.00
	Below median	16	0.69	0.25	0.00	0.06
1998–1999	Above median	31	0.26	0.65	0.10	0.00
	Below median	32	0.28	0.66	0.00	0.06
1999–2000	Above median	93	0.17	0.77	0.03	0.02
	Below median	93	0.30	0.51	0.10	0.10
2000–2001	Above median	56	0.16	0.84	0.00	0.00
	Below median	56	0.45	0.46	0.05	0.04
Overall	Above median	258	0.22	0.72	0.03	0.02
	Below median	262	0.34	0.55	0.05	0.06
Wald statistic			6.25***	6.03***	0.60	4.01**

probabilities of have-alpha funds on the inflows experienced in the final year of the classification period. The results indicate that above-median-flow funds have lower (higher) transition probabilities to the have-alpha (beta-only) group in the subsequent classification period. Across all years, an above-median-flow have-alpha fund has a 22 (72)% probability of being classified as a have-alpha (beta-only) fund in the subsequent nonoverlapping classification period. In contrast, for the below-median-flow *have-alpha* funds, there is a 34 (55)% probability of being classified as a have-alpha (beta-only) fund in the subsequent nonoverlapping classification period. These differences are statistically significant at the 1% level.

Table IX repeats the same analysis for the beta-only funds, and shows that there is little evidence to suggest that capacity constraints are relevant

Table IX

Transition Probabilities for Above- and Below-Median Flow Beta-Only Funds

The rows correspond to the 2-year period in which funds are classified as beta-only funds. The columns are, in order: the group affiliation, that is, whether the average beta-only fund experienced inflows (in the second year of the classification period) that were above or below the median beta-only inflow in that year; the number of beta-only funds in each group; the percentage of the flow-group members classified as have-alpha funds in the subsequent classification period; the percentage of the flow-group members classified as beta-only funds; the percentage of group members that were liquidated, and the percentage that stopped reporting. For example, in 1996 to 1997, 86 have-alpha funds had above the median inflow in 1997. Of these, 7% were classified as have-alpha funds, 76% as beta-only funds, 8% liquidated, and 9% stopped reporting in 1998 to 1999. Percentages may not add up to 100 because of rounding error. The final rows report the Wald test statistics (using the cross-correlation and heteroskedasticity-robust covariance matrix) for the hypothesis that the above- and median-flow transition probabilities are identical. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

			P(2-Year Transition)			
Classification Period	Flow Group	Number of Funds	Have-Alpha	Beta-Only	Liquidated	Stopped Reporting
1995–1996	Above median	77	0.04	0.79	0.09	0.08
	Below median	77	0.04	0.84	0.12	0.00
1996–1997	Above median	86	0.07	0.76	0.08	0.09
	Below median	86	0.07	0.71	0.19	0.03
1997–1998	Above median	138	0.27	0.58	0.09	0.06
	Below median	138	0.25	0.59	0.09	0.07
1998–1999	Above median	155	0.20	0.64	0.08	0.08
	Below median	156	0.15	0.60	0.15	0.10
1999–2000	Above median	130	0.08	0.82	0.08	0.03
	Below median	131	0.10	0.60	0.16	0.15
2000–2001	Above median	197	0.15	0.75	0.06	0.05
	Below median	197	0.06	0.79	0.11	0.04
Overall	Above median	783	0.11	0.75	0.08	0.06
	Below median	785	0.10	0.68	0.14	0.08
Wald statistic			4.05**	0.63	18.32***	0.16

for these funds. While there is a statistically significant difference in the ability of above-median- and below-median-flow beta-only funds to transition to the have-alpha group, the magnitudes of these transition probabilities are extremely similar (at 11% and 10%, respectively), suggesting that this difference is of limited economic importance.

We also condition the future t-statistic of alpha and the future level of alpha on the level of capital flows experienced by the have-alpha and beta-only funds. Table X reveals that for the have-alpha funds, the adverse effects of high capital flows on future risk-adjusted performance manifest themselves in reductions in the average t-statistic of alpha. Above-median-flow have-alpha funds exhibit an average alpha t-statistic of 1.47, while for the below-median-flow have-alpha funds, the comparable number is 1.84. This

Table X

Quantitative Measures of Alpha for Above- and Below-Median Flow Have-Alpha Funds

The rows correspond to the 2-year period in which funds are classified as have-alpha funds. The columns are, in order: the group affiliation, that is, whether the average have-alpha fund experienced inflows (in the second year of the classification period) that were above or below the median have-alpha inflow in that year; the number of have-alpha funds in each group; the average t-statistic of alpha for these funds in the subsequent classification period; and the annual average magnitude of alpha for these funds in the subsequent classification period. For example, in 1996 to 1997, 43 have-alpha funds had above the median inflow in 1997. For the ones that survived (see Table VI for details), the average t-statistic of alpha in the 1998 to 1999 classification period was 0.929, and the average annual alpha magnitude was 3.2% over the risk-free rate. The final rows report the Wald test statistics (using the cross-correlation and heteroskedasticity-robust covariance matrix) for the hypothesis that the above- and median-flow average t-statistic of alpha and average magnitude of alpha are identical. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Classification Period	Flow Group	Number of Funds	t-statistic of Alpha	Level of Alpha
1995–1996	Above median	20	1.61	0.047
	Below median	21	1.44	0.045
1996–1997	Above median	43	0.93	0.032
	Below median	44	1.45	0.068
1997–1998	Above median	15	4.42	0.109
	Below median	16	4.30	0.108
1998–1999	Above median	31	1.38	0.044
	Below median	32	1.39	0.018
1999–2000	Above median	93	1.03	0.023
	Below median	93	1.52	0.036
2000–2001	Above median	56	1.72	0.031
	Below median	56	2.59	0.062
Overall	Above median	258	1.47	0.035
	Below median	262	1.84	0.047
Wald statistic			5.97***	2.31

difference is statistically significant at the 1% level. While the difference also shows up in the future level of alpha, that is, have-alpha funds experiencing high inflows appear to have lower future levels of alpha, the high variance in the level of alpha in the cross-section renders this difference statistically insignificant. Finally, Table XI reveals, akin to the results for the transition probabilities in Table IX, that there are no real effects of capital flows on the future risk-adjusted performance of beta-only funds.

The results in this section indicate that conditioning the future performance of a fund on its current level of capital inflows provides useful information. Taken together, our findings thus far indicate that capital flows have primarily been directed toward have-alpha funds, and that these inflows have had an adverse effect on their risk-adjusted performance. The next section refines the analysis of averages that we conducted in Table II, highlighting intertemporal

Table XI

Quantitative Measures of Alpha for Above- and Below-Median Flow Beta-Only Funds

The rows correspond to the 2-year period in which funds are classified as beta-only funds. The columns are, in order: the group affiliation, that is, whether the average beta-only fund experienced inflows (in the second year of the classification period) that were above or below the median beta-only inflow in that year; the number of beta-only funds in each group; the average t-statistic of alpha for these funds in the subsequent classification period; and the annual average magnitude of alpha for these funds in the subsequent classification period. For example, in 1996 to 1997, 86 beta-only funds had above the median inflow in 1997. For the ones that survived (see Table VII for details), the average t-statistic of alpha in the 1998 to 1999 classification period was 0.804, and the annual average alpha magnitude was 5.3%. The final rows report the Wald test statistics (using the cross-correlation and heteroskedasticity-robust covariance matrix) for the hypothesis that the above- and median-flow average t-statistic of alpha and average magnitude of alpha are identical. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Classification Period	Flow Group Final Year	Number of Funds	t-statistic of Alpha	Level of Alpha
1995–1996	Above median	77	−0.09	−0.028
	Below median	77	−0.17	−0.040
1996–1997	Above median	86	0.80	0.053
	Below median	86	0.89	0.061
1997–1998	Above median	138	1.69	0.050
	Below median	138	1.60	0.042
1998–1999	Above median	155	0.78	0.018
	Below median	156	0.52	−0.011
1999–2000	Above median	130	−0.02	−0.005
	Below median	131	0.18	0.001
2000–2001	Above median	197	1.10	0.029
	Below median	197	0.78	0.024
Overall	Above median	783	0.82	0.020
	Below median	785	0.73	0.013
Wald statistic			1.13	2.00

variation in the performance of the average have-alpha and average beta-only fund.

E. Intertemporal Variation in the Alpha of Have-Alpha Funds and Beta-Only Funds

We estimate equation (5) and report the results in Table XII. The first feature to note in Table XII is that there are differences in the systematic risk exposures of the two groups. For example, the beta-only funds seem to have consistently greater exposure to *SNPMRF*. Second, the adjusted-R^2 statistics confirm that the Fung and Hsieh (2004a) seven-factor model continues to offer good explanatory power for the two groups of funds. Third, the structural breakpoints used for the analysis of the average fund are confirmed to exist for the two groups of funds as well. The intercepts from the regression reveal that in the first

Table XII

The Changing Risks of Have-Alpha and Beta-Only Funds

The top panel of this table contains estimates of:

$$R_{gt} = \alpha_{g1} D_1^1 + \alpha_{g2} D_2 + \alpha_{g3} D_3 + \left(D_1^1 X_t\right) \beta_{gD1} + (D_2 X_t)\beta_{gD2} + (D_3 X_t)\beta_{gD3} + \nu_{gt}$$

where $X_t = [SNPMRF_t \;\; SCMLC_t \;\; BD10RET_t \;\; BAAMTSY_t \;\; PTFSBD_t \;\; PTFSFX_t \;\; PTFSCOM_t]$.

Here, R_{gt}, the dependent variable, is the (equally weighted) average annualized excess return across all funds in group g in month t (g is have-alpha or beta-only), D_1^1 is a dummy variable set to one during the first period (January 1997 to September 1998) and zero elsewhere, D_2 is set to one during the second period (October 1998 to March 2000) and zero elsewhere, and D_3 is set to one during the third period (April 2000 to December 2004) and zero elsewhere. The regressors X are described in the text. The columns report the estimates for the regression for each group. The bottom panel contains estimates of Chow structural break test Chi-squared statistics for each group. White (1980) heteroskedasticity-consistent standard errors are reported below the coefficients. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Period I		Period II		Period III	
Variable	Have-Alpha	Beta-Only	Have-Alpha	Beta-Only	Have-Alpha	Beta-Only
Constant	0.0047***	−0.0017	0.0160***	0.0066***	0.0018*	−0.0002
	0.0013	*0.0022*	*0.0016*	*0.0024*	*0.0010*	*0.0009*
SNPMRF	0.1449***	0.3896***	−0.0514	0.1916***	0.1069***	0.1535***
	0.0187	*0.0356*	*0.0347*	*0.0570*	*0.0245*	*0.0193*
SCMLC	0.1438***	0.0901	0.2871***	0.3173***	0.1192***	0.1433***
	0.0528	*0.0814*	*0.0229*	*0.0410*	*0.0339*	*0.0234*
BD10RET	−0.0500	−0.3517***	0.5764***	0.4707***	0.1678***	0.1685***
	0.0768	*0.1293*	*0.0557*	*0.1774*	*0.0391*	*0.0331*
BAAMTSY	0.7828***	0.4475	0.2945***	0.8022***	0.2062***	0.1394**
	0.1822	*0.2992*	*0.0740*	*0.2304*	*0.0775*	*0.0680*
PTFSBD	0.0013	0.0336	0.0631***	0.0663***	−0.0052	−0.0008
	0.0089	*0.0188*	*0.0132*	*0.0218*	*0.0038*	*0.0034*
PTFSFX	0.0081	0.0104	−0.0278***	−0.0209	0.0092	0.0159***
	0.0059	*0.0098*	*0.0048*	*0.0133*	*0.0070*	*0.0045*
PTFSCOM	0.0029	0.0622***	−0.0266***	−0.0101	0.0169	0.0117
	0.0162	*0.0202*	*0.0048*	*0.0069*	*0.0102*	*0.0072*
Adjusted-R^2	0.733	0.752				
Number of months	96	96				

Period for Chow Structural Break Test	Have-Alpha	Beta-Only
Test for Period I & II Break		
Chi-sq(14)	55.21***	113.88***
Test for Period I Break		
Chi-sq(7)	159.51***	50.50***
Test for Period II Break		
Chi-sq(7)	358.35***	212.77***

subperiod, the have-alpha funds delivered a statistically significant alpha of 47 basis points per month (on an out-of-sample basis), or 5.6% per annum in excess of the risk-free rate. In contrast, the beta-only funds did not produce any detectable alpha over this period. The imprecisely estimated negative coefficient suggests that the fees and costs of beta-only funds destroyed any alpha they may have produced. During the second subperiod, although both groups delivered statistically significant alpha, the alpha of the have-alpha funds (at 160 basis points per month) was almost 2.5 times that delivered by the beta-only funds (66 basis points per month). In the final subperiod, the alpha of the have-alpha funds deteriorated. The have-alpha funds generated an alpha of 18 basis points a month, or 2.2% per annum. This could be attributed to the significant capital inflows experienced by the have-alpha funds, and the attendant declines in alpha that these inflows presage.

IV. Conclusions

In this paper, we employ a comprehensive data set of funds-of-funds to investigate performance, risk, and capital formation in the hedge fund industry over the decade from 1995 to 2004. We find that fund-of-fund returns are largely driven by their exposure to the seven risk factors of Fung and Hsieh (2004a). Controlling for these factor exposures, we find that the average fund-of-fund did not generate alpha, except in the period between October 1998 and March 2000. However, the average masks interesting cross-sectional variation in alpha generation. We find that, on average, 22% of the funds deliver positive and statistically significant alpha. We separate these alpha-producing funds (have-alphas) from the remainder (beta-only), and analyze the differences between these two groups.

Have-alpha funds are less likely to be liquidated, and have a higher propensity to persistently deliver alpha than beta-only funds. They also receive far greater capital inflows than beta-only funds. While capital flows into have-alpha funds do not exhibit trend-chasing behavior, those into beta-only funds significantly respond to past returns. Furthermore, have-alpha funds that experience high capital inflows have lower probabilities of being classified in the future as have-alpha funds, and have lower future information ratios.

Our findings can be interpreted in several ways. The different behavior of capital flows into have-alpha and beta-only funds suggests that there may be differences between the investors in the two groups, with relatively less sophisticated return-chasing investors investing in beta-only funds, and more sophisticated investors with an ability to detect the presence of alpha investing in have-alpha funds. Our results are also consistent with the assumptions of the Berk and Green (2004) rational model of active portfolio management, namely, that there exist significant differences in the ability of funds to generate alpha, that these funds face diminishing returns to scale in deploying their ability, and that investors rationally direct capital toward alpha-generating funds. In Berk and Green's equilibrium there is zero alpha available to investors. Our findings suggest that the hedge fund industry may be headed in this direction.

Appendix: Bootstrap Experiment

The following discussion closely follows Kosowski et al. (2006). Consider, for example, the 307 funds-of-funds in the 1997 to 1998 group. In this set, we could use a t-test to find funds with positive and statistically significant alpha. We would estimate alpha fund-by-fund, regressing fund returns on risk factors. We would then inspect the t-statistic of each fund's alpha, and if it were greater than a critical value of approximately 1.64, we would reject the null hypothesis of no alpha at the 5% level (for a one-sided test).

However, when we use this critical value, we rely on the assumptions that the residuals from the factor regressions are homoskedastic, serially uncorrelated, and cross-sectionally independent, which would result in normally distributed alphas and t-distributed alpha t-statistics. However, these assumptions about the residuals are very likely to be violated, given the much-noted nonnormality of hedge fund returns, which would result in an unknown distribution of alphas. Thus, the critical values obtained from the normal distribution would be incorrect. The bootstrap helps us to relax the assumptions of independence, normality, and zero serial correlation and enables us to find the correct critical values.

A description of the bootstrap experiment follows:

Cross-sectional Bootstrap

Step 1. For each fund i, regress the excess return on risk factors:

$$r_{i,t} = \hat{\alpha}_i + x_t'\hat{\beta}_i + \hat{\varepsilon}_{i,t}, t = 1, \ldots, T. \tag{A1}$$

Save $\hat{\beta}_i$, $\hat{\varepsilon}_{i,t}$, and the t-statistics of $\hat{\alpha}_i$, $\hat{t}(\hat{\alpha}_i)$, which can be calculated using standard OLS, Newey–West, or other standard errors. Do this for all funds $i = 1, \ldots, I$. Save the t-statistics, as well as the quintiles of the cross-sectional distribution of the t-statistics, for example, the 95$^{\text{th}}$ quintile, $\hat{t}_{0.95}$.

Step 2. Draw T periods with replacement from $t = 1, \ldots, T$. Call the resampled periods $\{t = s_1^b, \ldots, s_T^b\}$, where $b = 1$ is bootstrap number 1. For each fund, create the resampled observations:

$$r_{i,t}^b = x_t'\hat{\beta}_i + \hat{\varepsilon}_{i,t}, \quad \text{for} \quad t = s_1^b, \ldots, s_T^b. \tag{A2}$$

These draws impose the null that the alpha is zero, preserve the cross-sectional correlation of the residuals $\hat{\varepsilon}_{i,t}$ across funds, and preserve the higher-order correlation of the regressors and the residuals. For each fund i that has data for all resampled periods (this is true of all funds in the 1997 to 1998 period), run the regression:

$$r_{i,t}^b = \hat{\alpha}_i^b + x_t'\hat{\beta}_i^b + \hat{\varepsilon}_{i,t}^b \quad \text{for} \quad t = s_1^b, \ldots, s_T^b. \tag{A3}$$

Save the t-statistic of $\hat{\alpha}_i^b$, $\hat{t}^b(\hat{\alpha}_i^b)$. In each resample, save all the simulated t-statistics of the constant terms, $\hat{t}^b(\hat{\alpha}_i^b)$, across all funds (307 in the 1997 to 1998 period). Next, in each resample, inspect the cross-sectional distribution of the t-statistics. Suppose we are interested in the 95$^{\text{th}}$ percentile of the cross-sectional

t-statistics. Then after each resample, we can look at the 95[th] percentile of $\{\hat{t}^b(\hat{\alpha}_i^b)\}$ over all 307 funds. Call this $\hat{t}^b_{0.95}$.

Step 3. Repeat Step 2 for $b = 1, \ldots, B$. This gives two distributions, one for $\{\hat{t}^b(\hat{\alpha}_i^b)\}$ and one for $\{\hat{t}^b_{0.95}\}$.

Step 4. For each fund i, if $\hat{t}(\hat{\alpha}_i)$ is in the upper decile of the distribution of the simulated t-statistics, $\{\hat{t}^b(\hat{\alpha}_i^b)\}$, we call it a have-alpha fund. Otherwise, we call it a beta-only fund. We also inspect where $\hat{t}_{0.95}$ is in the distribution of $\{\hat{t}^b_{0.95}\}$.

Stationary Bootstrap

We use the stationary bootstrap of Politis and Romano (1994) to allow for weakly dependent correlation over time. Here, replace Step 2 as follows. First, draw randomly from the sample $t = 1, \ldots, T$. For the second resample observation, draw a uniform random variable from (0,1). If it is less than Q, then use the next observation. If we are at the end of the sample, start from the beginning again. If greater than Q, draw a new observation. We do this for $Q = 0$, 0.1, 0.5. We report the results for $Q = 0.5$ in the main body of the paper. The other two choices for Q do not materially affect our results.

REFERENCES

Admati, Anat, and Stephen Ross, 1985, Measuring investment performance in a rational expectations equilibrium model, *Journal of Business* 58, 1–26.

Agarwal, Vikas, Naveen Daniel, and Narayan Y. Naik, 2004, Flows, performance and managerial incentives in hedge funds, Working paper HF-016, London Business School.

Agarwal, Vikas, and Narayan Y. Naik, 2004, Risks and portfolio decisions involving hedge funds, *Review of Financial Studies* 17, 63–98.

Agarwal, Vikas, and Narayan Y. Naik, 2005, Hedge funds, *Foundations and Trends in Finance* 1, 103–170.

Ang, Andrew, Matthew Rhodes-Kropf, and Rui Zhao, 2005, Do funds-of-funds deserve their fees-on-fees?, Working paper, Columbia University.

Asness, Clifford, Robert Krail, and John Liew, 2001, Do hedge funds hedge? *Journal of Portfolio Management* 28, 6–19.

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

Brown, Stephen J., William N. Goetzmann, Bing Liang, and Christopher Schwarz, 2007, Optimal disclosure and operational risk: Evidence from hedge fund registration, Yale ICF Working Paper No. 06-15.

Chow, Gregory, 1960, Test of equality between sets of coefficients in two linear regressions, *Econometrica* 28, 591–605.

Cohen, Randolph, Paul Gompers, and Tuomo Vuolteenaho, 2001, Who underreacts to cashflow news? Evidence from trading between individuals and institutions, *Journal of Financial Economics* 66, 409–462.

Froot, Kenneth. A., and Tarun Ramadorai, 2008, Institutional portfolio flows and international investments, *Review of Financial Studies*, 21, 937–972.

Fung, William, and David A. Hsieh, 1997, Empirical characteristics of dynamic trading strategies: The case of hedge funds, *Review of Financial Studies* 10, 275–302.

Fung, William, and David A. Hsieh, 2000, Performance characteristics of hedge funds and CTA funds: Natural versus spurious biases, *Journal of Financial and Quantitative Analysis* 35, 291–307.

Fung, William, and David A. Hsieh, 2001, The risk in hedge fund strategies: Theory and evidence from trend followers, *Review of Financial Studies* 14, 313–341.

Fung, William, and David A. Hsieh, 2002, The risk in fixed-income hedge fund styles, *Journal of Fixed Income* 12, 6–27.

Fung, William, and David A. Hsieh, 2004a, Hedge fund benchmarks: A risk based approach, *Financial Analysts Journal* 60, 65–80.

Fung, William, and David A. Hsieh, 2004b, The risk in hedge fund strategies: Theory and evidence from long/short equity hedge funds, Working Paper, London Business School.

Fung, William, and David A. Hsieh, 2006, Hedge funds: An industry in its adolescence, *Federal Reserve Bank of Atlanta Economic Review* fourth quarter, 1–34.

Getmansky, Mila, Andrew W. Lo, and Igor Makarov, 2004, An econometric model of serial correlation and illiquidity in hedge fund returns, *Journal of Financial Economics* 74, 529–610.

Hasanhodzic, Jasmina, and Andrew W. Lo, 2006, Can hedge-fund returns be replicated? The linear case, Working Paper, MIT.

Hsieh, David, 1983, A heteroskedasticity-consistent covariance matrix estimator for time series, *Journal of Econometrics* 22, 281–290.

Kosowski, Robert, Alan Timmerman, Halbert White, and Russell Wermers, 2006, Can mutual fund stars really pick stocks? New evidence from a bootstrap experiment, *Journal of Finance* 61, 2551–2595.

Liang, Bing, 2000, Hedge funds: The living and the dead, *Journal of Financial and Quantitative Analysis* 35, 309–326.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Page, E.S., 1954, Continuous inspection schemes, *Biometrika* 41, 100–114.

Politis, Dimitris. N., and Joseph P. Romano, 1994, The stationary bootstrap, *Journal of the American Statistical Association* 89, 1303–1313.

Rogers, William H., 1983, Analyzing complex survey data, Working Paper, Rand Corporation, Santa Monica, CA.

Rogers, William H., 1993, Regression standard errors In clustered samples, Stata Technical Bulletin Reprints STB-13: STB-18, 88–94.

Sirri, Eric R., and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance*, 53, 1589–1622.

White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817–838.