JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 49, No. 3, June 2014, pp. 797–815

doi:10.1017/S0022109014000465

Recovering Delisting Returns of Hedge Funds

James E. Hodder, Jens Carsten Jackwerth, and Olga Kolokolova*

Abstract

Numerous hedge funds stop reporting each year to commercial databases, wreaking havoc with analyzing investment strategies that incur the unobserved delisting return. We use estimated portfolio holdings for funds-of-funds to back out estimated hedge-fund delisting returns. For all exiting funds, the estimated mean delisting return is insignificantly different from the average monthly return for live hedge funds. However, funds with poor prior performance and no clearly stated delisting reason had a significantly negative estimated mean delisting return of −5.97%, suggesting that a shock to their returns "tips them over the edge" and leads to delisting.

I. Introduction

Each year, a substantial fraction of hedge funds stop reporting their results to commercial databases. For example, the data used in this paper exhibit an average annual "delisting" rate of 14.55%.[1] These data are a combined database created from six major commercial databases (ALTVEST, BarclayHedge, Center for International Securities and Derivatives Markets (CISDM), Eurekahedge, HFR, and TASS) for Jan. 1994–June 2009.[2] Delisted funds are often described as "dead funds," but many of them continue to exist. In our data, only 23.59% of delisting funds indicated they were being liquidated, and some 0.97% state that they were merged with another hedge fund. Another 2.08% indicate that they stopped providing their returns because they closed to further investments (potentially due to stellar performance and large inflows of investment capital).

*Hodder, jhodder@bus.wisc.edu, School of Business, University of Wisconsin, 975 University Ave, Madison, WI 53706; Jackwerth, jens.jackwerth@uni-konstanz.de, Department of Economics, University of Konstanz, PO Box 134, Konstanz 78457, Germany; and Kolokolova, olga.kolokolova@mbs.ac.uk, Manchester Business School, University of Manchester, Booth St W, Manchester M15 6PB, United Kingdom. We thank the following for helpful comments on earlier versions of the paper: Hendrik Bessembinder (the editor), Kostas Iordanidis, Veronika Krepely Pool, Ingmar Nolte, Winfried Pohlmeier, Tarun Ramadorai (the referee), and seminar participants at the University of Cyprus-Nikosia, Humboldt University, Instituto de Estudios Superiores de la Empresa (IESE), the University of Konstanz, Universitat Pompeu Fabra, and the University of Zurich.

[1]In what follows, we will use the terms "delist" and "exit" to equivalently indicate that the fund has stopped reporting its performance to database providers.

[2]Our versions of the respective databases cover somewhat differing time periods; but in the aggregate, the combined data span the Jan. 1994–June 2009 period. There is also overlapping coverage of some funds, and we adjust for that overlap.

Moreover, the remaining 73.36% of delisted funds either did not indicate why they ceased reporting or provided noninformative statements such as "requested by manager."

When studying hedge-fund performance, one faces the issue of what return should be attributed to the period when a fund stops reporting. Simply dropping that period ignores the fact that fund investors will actually experience the delisting return. In contrast, Posthuma and Van der Sluis (2004) used 0%, −50%, and −100% to cover a wide range of possibilities for the unknown delisting return. This drew a strong response from two practitioners, Van and Song ((2005), p. 7), who call the assumption of a −50% delisting return "outrageous." However, if a fund has suffered massive losses and is being liquidated, a large negative delisting return is possible. Particularly if a fund had large illiquid positions that would be difficult to value and sell, its mark-to-market valuation prior to delisting could seriously underestimate the extent of losses from liquidation under adverse circumstances.

We develop a methodology for estimating delisting returns based on a fund-of-funds (FoF) being a portfolio of positions in individual hedge funds.[3] If we had direct information on FoF portfolio positions, it would be straightforward to back out returns for delisting funds using that information plus the FoF returns and the returns of live hedge funds for the delisting month. Unfortunately, we do not have FoF portfolio positions. Instead, we estimate those portfolio holdings through a matching algorithm related to principal component analysis. We then obtain delisting returns based on the difference between the observed return for each FoF and the return from its estimated portfolio holdings in live (still reporting) hedge funds during a period where one of its hedge-fund holdings delists. More details are provided in the next section.

Our estimated mean delisting return for all exiting funds is negative but not significantly different from the mean monthly return of 0.56% for all hedge funds in our sample during Jan. 2000–June 2009. We document some return persistence, with hedge funds that delist after positive average returns over 6 months tending to have higher delisting returns than the average hedge fund. Symmetrically, hedge funds that delist after having negative average returns tend to have negative delisting returns. Digging deeper, we find that a negative mean delisting return is largely due to funds that did not state a clear reason for delisting. Those funds had an estimated mean delisting return of −5.97%. On the other hand, funds that stated they were being liquidated after a negative average return over the previous 6 months had an estimated mean delisting return of −0.59%, which is not significantly different from the average monthly return for live hedge funds. This is one example of a more general pattern where liquidated funds have average exit returns substantially better than the set of funds with no clear stated reason for delisting. We explore this issue further in Section IV.

[3]Fung and Hsieh (2000) as well as Fung, Hsieh, Naik, and Ramadorai (2008) have also noted that FoF returns implicitly incorporate the delisting returns of individual hedge funds; however, they do not use the portfolio connection to actually back out the delisting returns. Nevertheless, Fung et al. ((2008), p. 1778) do point out that the absence of delisting returns leads to a situation where a "fund-of-fund's return more accurately reflects the losses experienced by investors in the underlying hedge fund (albeit indirectly)."

There is a literature that explores hedge-fund performance prior to delisting.[4] However, there have been few attempts to examine performance after delisting. Ackermann, McEnally, and Ravenscraft (1999) used a combined data set with underlying data from two providers, Managed Account Reports, Inc. (MAR) and Hedge Fund Research, Inc. (HFR). During 1993–1995, their combined data included 37 "terminated" funds (liquidated, restructured, or merged into another fund) plus an additional 104 funds that stopped reporting without a clear indication as to why. Out of this total of 141 delisting funds, those authors were able to obtain delisting returns for some fraction of the 37 terminated funds via a special request to HFR. The response from HFR indicated an average delisting return of -0.7%, with a surprisingly rapid final redemption averaging only 18 days after delisting. It would appear that some of the terminating funds were in the process of liquidating while still reporting returns. Unfortunately, those data are rather early (1993–1995), predating the boom in the hedge-fund industry, and are based on a relatively small sample (at most, 37 terminating funds).

Agarwal, Fos, and Jiang (2013) also perform some analysis of delisting returns, although the authors focus largely on an attempt to estimate "self-reporting bias" using information from 13F filings with the U.S. Securities and Exchange Commission (SEC) that primarily address U.S. equity and some option positions. The analysis covers a longer period from 1980 to 2007; however, their reported number of delisting hedge funds is still rather limited (only 187 instances). Moreover, these 187 funds are not liquidated or merged but continuing to operate. Even so, their estimated mean delisting return of -0.72% is quite similar to that reported for terminated funds by Ackermann et al. (1999). It should be noted that 13F filings are quarterly and involve sizable management firms (assets under management (AUM) of over $100 million) rather than individual hedge funds, which suggests their estimated returns are for management firms and may involve multiple hedge funds.

There is also a recent paper by Aiken, Clifford, and Ellis (2013) that estimates hedge-fund returns based on reported quarterly hedge-fund holdings during 2004–2009 by each of 80 FoFs that were registered with the SEC. That paper also focuses on self-reporting bias but does report some results for delisting hedge funds. Those results indicate delisting funds underperform funds that remain listed by approximately 0.45% monthly during the quarter after delisting. That estimate is on a risk-adjusted basis using the Fung and Hsieh (2001) 7-factor model. However, the delisting funds in this paper (as in Agarwal et al. (2013)) are not liquidated or merged but continuing to operate. Moreover, a potentially important issue with this paper is that it checks for listing (delisting) in only two databases (TASS and BarclayHedge).

The remainder of the paper is as follows: Section II provides details on the matching algorithm and the econometric model of FoF returns. In Section III, we describe our empirical design and basic characteristics of the data sample. Results are contained in Section IV, with several robustness checks described in Section V. Section VI concludes.

[4]See, for example, Brown, Goetzmann, and Ibbotson (1999), ter Horst and Verbeek (2007), as well as Liang (2000).

II. The Basic Model

Since we do not have precise information on portfolio holdings for each FoF in our sample, we need a procedure for estimating those holdings. We use a matching algorithm described below that is conceptually related to principle components. As a preliminary step, we need to "gross up" the reported FoF returns to a pre-fee level (i.e., to the return level before management and incentive fees were extracted by the FoF). That pre-fee FoF return is the return on a portfolio of post-fee hedge-fund returns (management and incentive fees having already been extracted by the respective hedge funds). As our FoF and hedge-fund return data are all post-fee, we transform the FoF returns to a pre-fee basis using an algorithm closely related to Brooks, Clare, and Motson (2008) and detailed in Kolokolova (2011).

In our implementation, we use a 36-month rolling window and consider only FoFs and hedge funds that report returns for all months in the relevant window. As with many other implementation choices for our basic methodology, we have examined robustness to variations in the choice of a 36-month window. To avoid cluttering the exposition, we defer discussion of such robustness checks until Section V. As a general statement, our qualitative results are robust, but there can be some variation in point estimates.

For each FoF, we find the hedge fund whose (post-fee) returns are most highly correlated with the (pre-fee) returns of that FoF. Then, we regress the FoF returns on the chosen hedge fund and obtain the residual returns. In these regressions, we impose upper and lower limits on the estimated weights (more details below) to assure a reasonable level of portfolio diversification and avoid highly concentrated holdings that would be rather unlikely in FoF portfolios. Next, we find a second hedge fund that is now the most highly correlated with the residual returns for that FoF. We add that hedge fund to the portfolio, find new residual returns, and proceed in this fashion until we have 15 hedge funds in the portfolio.[5] Additionally, after having added the 10th hedge fund, we require the estimated portfolio weights in all subsequent portfolios to sum up to unity.

Once we work out the set of matched hedge funds for each FoF, we are ready to model the pre-fee returns of the FoF as a portfolio of the (post-fee) returns on the matched hedge funds. Hedge funds within each match are indexed by j. The (pre-fee) FoF returns are always indicated with an uppercase R, and the live hedge-fund returns (post-fee) are denoted with a lowercase r_L. We use $T = 36$ consecutive returns to estimate the following regression model for each FoF, with those FoFs indexed by i and time periods (months) by t:

$$(1) \quad R_{it} = [r_{Lt}]\beta_i + \varepsilon_{it}, \qquad t = 1, \ldots, T, \qquad \text{and} \quad i = 1, \ldots, N_{\text{FoF}},$$
$$\text{s.t.} \quad \beta_{\min} \leq \beta_i \leq 0.10, \qquad \sum_j \beta_{ij} = 1,$$

[5]We reexamine this restriction as well as others in Section V and the Internet Appendix (available at www.jfqa.org).

where N_{FoF} is the number of all possible subsamples of T consecutive returns for the FoFs reporting to our database. We do not make any assumptions concerning the distribution of the error term ε_{it} except that it has a zero mean.

Since equation (1) implicitly has unlevered returns for the FoFs, our main results utilize only those FoFs that report not using leverage. These FoFs attempt to remain close to fully invested, and we do not include the riskless asset as one of the potential investments. In order to ensure economically sensible portfolio positions, we restrict the loadings β_i (portfolio weights for FoF_i) on the matched hedge funds to be smaller than 0.10 and larger than some minimal value $\beta_{\min}$. For the main part of our analysis, $\beta_{\min}$ is set at 0.02. We further assume that each FoF is fully invested in its set of matched hedge funds.[6]

We now turn to the fitted return of the FoF in period $T + 1$. If all the hedge funds in that particular FoF portfolio are still alive, then the fitted return is simply calculated with the portfolio weights that were estimated using equation (1) coupled with the observed returns of the matched hedge funds for period $T + 1$:

$$\hat{R}_{i,T+1} = [r_{L,T+1}]\hat{\beta}_i. \tag{2}$$

Now consider the situation where a hedge fund delists and does not report its return for period $T + 1$. We denote that unreported return as $r_{E,T+1}$. The econometrics and computations turn out to be much simpler if we base our estimates on matched FoF portfolios where there is a single delisting hedge fund. That situation represents approximately 89% of our matched sample, and we drop matches with multiple delisting hedge funds from the estimation procedure. Note that with one delisting fund in the portfolio, the vector of live returns $r_{L,T+1}$ will be one shorter than in the above situation, where all hedge funds for a given FoF portfolio remained alive. In period $T + 1$, an FoF with a (single) delisting hedge fund in its portfolio will have an actual return that can be expressed as

$$R_{i,T+1} = [r_{L,T+1}, r_{E,T+1}]\beta_i + \varepsilon_{i,T+1}. \tag{3}$$

We approximate the true betas with the estimated betas from equation (1), and estimate the delisting return as

$$\begin{aligned} \hat{r}_{E,T+1} &= \frac{R_{i,T+1} - [r_{L,T+1}]\hat{\beta}_{L,i}}{\hat{\beta}_{E,i}}, \\ t &= 1, \ldots, T \quad \text{and} \quad i = 1, \ldots, N_{\mathrm{FoF}}, \end{aligned} \tag{4}$$

where $\hat{\beta}_{L,i}$ and $\hat{\beta}_{E,i}$ are the estimated betas, respectively, for the 14 hedge funds staying alive and the one delisting hedge fund in period $T + 1$ for the matched portfolio of FoF_i. The numerator of equation (4) contains an estimation error that

[6]There is a potential omitted variables problem in that a given FoF may be invested in one or more hedge funds that are not in our database. Our procedure implicitly approximates such missing funds by a linear combination of hedge funds that are in our database. Simulation studies discussed in Section V indicate that our methodology works relatively well, even with a hypothetically large number of missing funds. As a practical matter, our combined database is large and should have a substantial portion of the relevant hedge funds, further mitigating the potential omitted variables problem.

is amplified when dividing by a fractional $\hat{\beta}_{E,i}$ (which is also estimated with error). Particularly when $\hat{\beta}_{E,i}$ is low, this calculation can result in large errors. We mitigate this problem by discarding matches where $\hat{\beta}_{E,i} < 0.05$ as well as trimming (in each tail) the most extreme 1% of remaining estimates from equation (4).

We also consider the fact that several FoFs might invest in the same hedge fund. If that hedge fund delists, then the associated delisting return $r_{E,T+1}$ should be the same for all FoFs with that hedge fund in their portfolios. To ensure that result, we add up the relevant equations (3) while keeping the $r_{E,T+1}$ constant.[7] The estimated realization of the delisting return in this case is

$$(5) \qquad \hat{r}_{E,T+1} = \sum_i \frac{R_{i,T+1} - [r_{L,T+1}]\hat{\beta}_{L,i}}{\sum_i \hat{\beta}_{E,i}}, \qquad t = 1,\ldots,T,$$

where the sum is taken across all FoF matches i that include the delisted hedge fund of interest.

We estimate the mean delisting return by averaging the individual realizations calculated above. Our matching procedure does not require precise hedge-fund identification, and the returns of funds truly included in an FoF portfolio can be proxied by returns of different (but correlated) funds in the matching portfolio. Nevertheless, the estimate of μ_E is unbiased only if an FoF truly invests into k delisted hedge funds and the corresponding matched portfolio also has exactly k delisted funds. One cannot guarantee that exact correspondence regarding the number of delisted funds while constructing the matching portfolios, and hence, we need to adjust the estimated μ_E for potential bias.

Since we use only matches that have exactly one delisted fund, the following biases can occur. First, consider an FoF that did not actually invest in any delisted fund, but the estimated matching portfolio erroneously contained a single delisted fund. Using this match, one would estimate not an unobserved delisting return (on average μ_E) but the return of a hedge fund that was still alive. The higher the share of such matches, the more the estimated μ_E will be biased toward the average return of hedge funds that were reporting to the database, which we denote by μ_{HF}. Second, if an FoF truly invested into one delisted hedge fund and the estimated matching portfolio also has one delisted fund, then the match has perfect correspondence and does not bias the estimate of μ_E. Third, consider an FoF that actually had investments in two or more hedge funds that delisted, but that FoF was matched with a portfolio having only one delisted fund. If the number of truly delisted funds was 2, the resulting average estimate would be $\mu_E + (\mu_E - \mu_{\text{HF}})$ instead of μ_E. Simulation results described below indicate that the probability is only 0.09% that an FoF with 3 or more truly delisting hedge funds is matched with a single delisting fund. Consequently, our adjustment procedure does not consider cases with three or more truly delisting hedge funds in a single FoF portfolio.

[7]We have also tested our methodology by comparing estimated delisting returns where the same hedge fund is part of different FoFs with estimated delisting returns where different hedge funds are part of different FoFs. As expected, the mean absolute differences are lower in the former case (see Section IA.1 of the Internet Appendix).

The biases due to the above mismatches can be corrected if one knows the share of matches for each type. Let us denote by p_k the probability that an FoF truly invested in k delisted funds, and the estimated matching portfolio indicates the existence of only one delisted fund. Then the estimated biased delisting return $\mu_E^{\text{Estimated}}$ is a weighted average of the unbiased estimate μ_E^{Unbiased} and the average return of hedge funds in the database μ_{HF}.[8] That is

$$(6) \qquad \mu_E^{\text{Estimated}} = p_0\mu_{\text{HF}} + p_1\mu_E^{\text{Unbiased}} + (1 - p_0 - p_1)\left(2\mu_E^{\text{Unbiased}} - \mu_{\text{HF}}\right),$$

and we can solve for μ_E^{Unbiased}:

$$(7) \qquad \mu_E^{\text{Unbiased}} = \frac{\mu_E^{\text{Estimated}} - (2p_0 + p_1 - 1)\mu_{\text{HF}}}{2 - 2p_0 - p_1}.$$

The probabilities p_k are not known but can be estimated using a simulation. For each FoF in the database observed in a given month, we construct a hypothetical FoF from 15 randomly selected (with replacement), existing hedge funds. The portfolio weights are uniformly and randomly selected in the interval 0.02 to 0.10 and are required to sum to 1. Some hedge funds will have all 37 returns, and some will delist in month 37. We thus obtain a database of simulated FoFs of the same dimension as the original database, but where we know the number of delisting hedge funds for each FoF.

We next employ our usual matching procedure. Based on those estimated matches, we compute the frequencies for matches in which one estimated delisting fund (using our matching procedure) corresponds to 0, 1, 2, and 3 or more true delistings in the simulated FoFs. We repeat the complete simulation 100 times and compute the estimated probabilities p_k as averages of the corresponding frequencies. In $p_0 = 59.97\%$ of all cases, there is no delisting hedge fund in an FoF. In $p_1 = 37.64\%$ of all cases, there is 1 delisting hedge fund, and in $p_2 = 2.30\%$ of all cases, there are 2 delisting hedge funds. Three or more delisting hedge funds is very rare, occurring in only 0.09% of all cases.

III. Data Characteristics and Implementation

We begin this section with a description of the data before proceeding to a discussion of our bootstrap procedure for estimating standard errors.

A. The Data

We have constructed a joint database using a union of six major databases (ALTVEST, BarclayHedge, CISDM, Eurekahedge, HFR, and TASS) from which we deleted duplicates and different share classes of the same fund. That joint database is large, containing more than 20,000 hedge funds and about 6,000 FoFs that reported sometime during the Jan. 1994–June 2009 period. Those funds are

[8] In our adjustment, we use the average monthly return of all reporting hedge funds in the sample. This also includes funds that were alive during a portion of the Jan. 2000–June 2009 period but eventually died.

classified into dead and live hedge funds plus dead and live FoFs. We use only funds that report in U.S. dollars (USD) and have a performance record after Jan. 2000. This leaves us with 16,398 individual hedge funds and 5,031 FoFs. Panel A of Table 1 reports descriptive statistics for those funds over the period from Jan. 2000 to June 2009. A fund being designated as live or dead in that table refers to its status as of June 2009. Note that the monthly returns are post-fee for both hedge funds and FoF in Panel A, just as they are reported in the database.

TABLE 1

Descriptive Statistics

Table 1 reports descriptive statistics for funds from the union of six databases (ALTVEST, BarclayHedge, CISDM, Eurekahedge, HFR, and TASS). Panel A is based on all unique funds reporting in U.S. dollars during Jan. 2000–June 2009. Panel B is based on the funds used in our analysis, after we dropped the first 12 observations for all hedge funds and eliminated any hedge fund and FoF that did not have at least 36 consecutive remaining observations between Jan. 1997 and June 2009. The performance of these funds is reported between Jan. 2000 and June 2009. We also eliminate FoFs that report using leverage. Return statistics are based on monthly returns in percentages. Note that all returns in Panel A are post-fee. In Panel B, the FoF returns are grossed up to a pre-fee basis, while the hedge-fund returns remain post-fee. All values except Number of Funds are averages of the corresponding statistics for the individual funds.

Panel A. All Funds (Jan. 2000–June 2009)

	Hedge Funds (post-fee)			Funds of Funds (post-fee)		
	All	Live	Dead	All	Live	Dead
No. of funds	16,398	8,847	7,551	5,031	3,625	1,406
Lifetime in years	3.27	4.72	2.00	4.12	4.82	2.53
Mean return	0.55	0.70	0.37	0.25	0.22	0.31
Median return	0.50	0.79	0.16	0.46	0.51	0.34
STD	4.60	4.33	4.92	2.45	2.48	2.37
Min return	−10.18	−11.01	−9.21	−6.71	−7.34	−5.09
Max return	11.68	11.78	11.56	5.32	5.27	5.45
Skewness	−0.07	−0.22	0.11	−0.63	−0.81	−0.18
Kurtosis	5.04	5.71	4.25	5.56	5.97	4.47
Sharpe ratio	0.14	0.22	0.05	0.13	0.12	0.16

Panel B. Funds with at Least 36 Returns (Jan. 2000–June 2009)

	Hedge Funds (post-fee)			Funds of Funds (pre-fee, no leverage)		
	All	Live	Dead	All	Live	Dead
No. of funds	7,910	4,716	3,194	1,348	921	427
Lifetime in years	5.37	6.44	3.81	5.56	6.33	3.91
Mean return	0.56	0.75	0.26	0.56	0.62	0.45
Median return	0.51	0.84	0.02	0.71	0.84	0.41
STD	4.55	4.17	5.11	2.54	2.46	2.72
Min return	−11.66	−12.19	−10.87	−7.51	−8.06	−6.32
Max return	13.34	13.06	13.74	6.92	6.71	7.38
Skewness	−0.08	−0.26	0.18	−0.54	−0.77	−0.04
Kurtosis	6.00	6.67	4.99	6.65	7.17	5.52
Sharpe ratio	0.11	0.17	0.04	0.20	0.20	0.20

We eliminate the first 12 returns for each hedge fund in order to mitigate backfill bias. Our matching procedure requires funds that report returns for at least 36 consecutive months, and we eliminate all funds that do not satisfy that requirement (after deleting the first 12 monthly returns for hedge funds). Except for robustness tests discussed in Section V, we utilize only FoFs that indicate they never use leverage.

When one looks at delisting events before Jan. 2000, nearly half are reported as occurring at year end; however, in many cases, the last several months of reported returns were all zeros. Thus, we believe that monthly delisting dates before Jan. 2000 are not reliable. Consequently, we use only funds that report at least 36 returns after Jan. 1997, such that their reported delisting occurs no earlier than

Jan. 2000. Panel B in Table 1 reports descriptive statistics for those funds, and we have 7,910 hedge funds, of which 3,194 delisted (died) at some time prior to the end of June 2009. Among the 1,348 FoFs in our restricted sample, 921 are classified as live funds; however, we can still use the 427 dead FoFs for windows of time when they were alive. For the FoF statistics in Panel B, we report pre-fee returns computed using the algorithm of Kolokolova (2011) mentioned previously. When implementing that algorithm, we use the reported fee structure for each FoF; however, as a point of information, the typical FoF in our data charges a management fee of 1% and an incentive fee of 10% per year.

B. Bootstrapped Standard Errors

Calculating standard errors for our analysis is potentially problematic due to the multiple-layer estimation procedure and the consequent accumulation of errors from the potential mismatch of FoF portfolios and the estimation of betas. Moreover, the different FoF matches will typically have overlapping time series. Because of these issues, we use a bootstrap approach to estimate standard errors. In particular, we utilize a two-stage procedure that bootstraps over the FoFs and the hedge funds. For the first stage, define an FoF instance to be a sequence of 37 returns for the relevant FoF. From the original data, we randomly draw with replacement the same number of FoF instances as in that original data to create a bootstrapped FoF instance set. In the second stage, we begin by identifying the set of hedge funds that provide 36 returns in parallel to the first 36 returns of an FoF instance. Some of these hedge funds have a 37th return in parallel with the FoF instance, whereas others exit and have just 36 returns. We then draw with replacement out of this set of hedge funds a bootstrapped hedge-fund universe of the same size and potentially containing both live and exiting hedge funds. We use that bootstrapped hedge-fund universe when we run our matching procedure for the associated FoF instance.

We employ our matching method with each FoF instance and its hedge-fund universe in order to generate bootstrapped matches. This approach allows us to have bootstrapped matches that contain differing hedge funds as well as portfolio weights that differ from our original match. We obtain a new estimate for μ_{E} using this bootstrapped set of matches and beta estimates. Finally, we use our bias correction described above to adjust for a mismatched number of delisting funds and obtain unbiased estimates for μ_{E}. We repeat this entire procedure 1,000 times to obtain bootstrapped standard errors that allow for potential mismatch of FoF portfolios, estimation error in the portfolio weights, overlapping time series, and small sample effects.

We also considered a three-stage bootstrap, where we resample the 36 months of an FoF instance and the associated hedge-fund universe's returns by time slice (keeping the actual returns aligned by month). Results change little, but this approach destroys the time-series return properties that will be important to our interpretation of the results. More information and results using the three-stage bootstrap are provided in Section IA.2 of the Internet Appendix. In the main body of this paper we utilize the two-stage bootstrap in order to preserve time-series properties for returns.

IV. Results

The results discussed in this section are based on FoF matches using our standard procedure described above. For the entire initial set of matches, the average holdings of individual hedge funds (β_i) are estimated to be 0.067, with the standard deviation across matches of 0.033. The average loadings on the delisting funds are estimated to be 0.062, which increases to 0.087 after we discard matches where $\hat{\beta}_{E,i} < 0.05$. Discarding matches with low estimated betas has two effects. For one, it reduces the estimated delisting return variability by avoiding division in equation (4) using very low betas. On the other hand, it introduces a bias by avoiding large absolute returns, which pushes estimated delisting returns toward 0. In a simulation study we find that, given the typical magnitudes of our estimated delisting returns, the simulated bias is only a fifth of the simulated standard deviation of our delisting returns. Thus, we feel comfortable proceeding with our choice of discarding matches with low estimated betas. Details on the simulation results can be found in Section IA.3 of the Internet Appendix.

In Table 2, we report estimated mean delisting returns for "All" matches as well as for funds that stated they were being "Liquidated" or provided "No Reason" that was informative regarding their reason for delisting.[9] For the set of All delisting hedge funds, we find an estimated average monthly delisting return (bias-corrected) of -1.61%. Although negative, that estimate is rather noisy and not significantly different from the average return for all hedge funds of 0.56% reported in Panel B of Table 1. Moreover, this result is quite different from a very large negative delisting return such as -50%, and the bootstrapped standard deviation (STD) is such that we can be quite confident that the average delisting firm does not have such a large negative delisting return. That conclusion is further supported by a simulation test reported in Section V that indicates our procedure (albeit noisy) would reliably find a mean delisting return that was large and negative if the process generating the data had such a large negative mean.

TABLE 2

Mean Delisting Returns

In Table 2 we report the monthly delisting returns (bias-corrected) based on matched portfolios of FoFs where the adjusted R^2 of the main regression model is at least 25% and the portfolio weight of the delisting fund is at least 5%. We use 1% tail trimming. Mean delisting returns and their standard deviations are in percentages per month.

	No. of Matches	Mean Delisting Return	Bootstrapped STD of Mean Delisting Return	Nonparametric *p*-Value for Difference with Average HF Return
All	1,392	−1.61	0.94	0.16
Liquidated	358	2.69	1.75	0.23
No Reason	998	−3.18	1.09	0.06

Turning to the Liquidated and No Reason fund categories considered separately, the situation changes. Funds in the No Reason category have a negative

[9]Other self-reported categories such as "merged" and "closed to further investment" were too small to have reliable mean estimates. Among all delisted hedge funds, only 0.87% of funds report delisting because of being merged, and some 2.08% because of being closed to further investment.

average delisting return of −3.18%, which is significantly different from the 0.56% average monthly return for all hedge funds with a *p*-value of 0.06. The Liquidated funds have a positive estimated average delisting return of 2.69%, which is significantly different from the estimate for No Reason funds (*p*-value of 0.06) but not significantly different from the average monthly return for all hedge funds. This pattern seems a bit surprising.

One tends to think that funds being liquidated were presumably poor performers and likely to have negative delisting returns rather than positive. In contrast, it seems plausible that the mean delisting return of funds that did not state a clear reason for delisting could be similar to the average monthly return of all (live) hedge funds. It might be that a substantial fraction of those No Reason funds were doing fairly well and delisted for other (unstated) reasons. Perhaps they merged or even were closed to further investment but did not bother to state that reason. Reporting to a database can be characterized as a form of advertising, and there could be a variety of reasons to stop advertising. Moreover, poor past performance should not necessarily indicate a negative delisting return if the fund's assets have been properly marked to market. Yet, we find a significantly negative average delisting return for the No Reason funds.

A. Top and Bottom Funds

To investigate this issue further, we sorted the exiting hedge funds into Top and Bottom groups, such that Top funds exhibit positive average returns over the 6 months prior to delisting, whereas Bottom funds exhibit negative average returns. Mean delisting returns for these subcategories are reported in Table 3.

TABLE 3

Mean Delisting Returns: Top versus Bottom

In Table 3 we report the monthly delisting returns (bias-corrected) based on matched portfolios of FoFs where the adjusted R^2 of the main regression model is at least 25% and the portfolio weight of the delisting fund is at least 5%. We use 1% tail trimming. Mean delisting returns and their standard deviations are in percentages per month. Top (Bottom) funds have positive (negative) average returns over the 6 months prior to the delisting event.

	No. of Matches	Mean Delisting Return	Bootstrapped STD of Mean Delisting Return	Nonparametric *p*-Value for Difference with Average HF Return	Nonparametric *p*-Value for Difference between Top and Bottom Funds
Panel A. Top Funds					
All	807	0.31	1.17	0.38	0.05
Liquidated	194	5.46	2.35	0.16	0.27
No Reason	593	−1.28	1.29	0.43	0.05
Panel B. Bottom Funds					
All	585	−4.25	1.42	0.04	—
Liquidated	164	−0.59	2.49	0.44	—
No Reason	405	−5.97	1.77	0.01	—

There is evidence of return persistence, with the Top funds having higher mean delisting returns than the Bottom set of funds. The *p*-value of that difference for all funds is 0.05. Top funds have a modestly positive mean delisting return of 0.31%, whereas Bottom funds have a relatively large negative mean delisting return of −4.25%. The estimate for Bottom funds is significantly different from the average return of the all (live) hedge funds (p-value of 0.04). Note that the estimate for Bottom funds is quite large on an annualized basis, with −4.25% monthly equating to −51% annually (without compounding).

In identifying Top versus Bottom performing funds, we also used three alternative metrics: returns relative to the Standard & Poor's (S&P) 500 and two measures of drawdown. We further define Top and Bottom in two ways, namely, with our usual cutoff (positive versus negative average returns over the 6 months prior to delisting) and alternatively, with Top being the best 30% and Bottom being the worst 30% of hedge funds when assessed based on the above metrics. The results are very consistent across these different approaches, with the main message being that Bottom funds selected by a variety of plausible approaches have low delisting returns. Those results for Bottom funds are reported in Table 4.

TABLE 4

Mean Delisting Returns of Bottom Funds Selected by Alternative Metrics

In Table 4 we report the monthly delisting returns (bias-corrected) based on matched portfolios of FoFs where the adjusted R^2 of the main regression model is at least 25% and the portfolio weight of the delisting fund is at least 5%. We use 1% tail trimming. Mean delisting returns and their standard deviations are in percentages per month. We use several ways to define Bottom funds, which are described in the first column. NAV is the net asset value, and HF is hedge fund.

Bottom Funds	No. of Matches	Mean Delisting Return	Bootstrapped STD of Mean Delisting Return	Nonparametric p-Value for Difference with Average HF Return
Negative average 6-month return	585	−4.25	1.42	0.04
30% lowest average 6-month return	427	−3.60	1.71	0.06
Average 6-month return below the corresponding return on the S&P 500 index	576	−3.76	1.42	0.04
30% largest difference in the 6-month fund return and the return on the S&P 500 index	397	−3.95	1.62	0.06
Drawdown based on the highest to lowest fund NAV is below the median	786	−2.84	1.26	0.06
30% largest drawdown based on the highest to lowest fund NAV	442	−2.79	1.89	0.13
Drawdown based on the highest to last fund NAV is below the median	789	−3.08	1.26	0.06
30% largest drawdown based on the highest to last fund NAV	439	−3.34	1.86	0.08

The persistence of poor results for Bottom funds is consistent with Getmansky, Lo, and Makarov (2004), who found persistence among live funds. Also, it is probable that some funds are exiting because their strategy and/or implementation is performing poorly in the then-current economic environment. Most such funds would presumably be in the Bottom set, and assuming the environment continued to be unfavorable as they exited, return persistence seems reasonable. The negative mean delisting returns for both Bottom Liquidated and Bottom No Reason funds are consistent with that story; however, the result for Top Liquidated funds in Table 3 is somewhat counterintuitive. Return persistence itself is not surprising, but if a fund is apparently doing well, why is it being liquidated?

Looking at returns, we find positive average returns over the half-year prior to delisting of 1.18% per month for the Top Liquidated group and 1.38% per month for the Top No Reason group. So the average returns are positive but actually lower for the Top Liquidated group compared with the Top No Reason funds.

The size of the Top Liquidated funds was on average a relatively small USD 57 million 6 months prior to delisting, compared with USD 123 million for the Top No Reason funds and USD 172 million for live funds with positive average returns over 6-month periods. Also, the Top Liquidated funds had an average net outflow of USD 73 thousand per month during the 36 months prior to exit. In contrast, the Top No Reason group over the comparable period had an average inflow of USD 1.36 million per month (very similar to the average monthly inflow for all live funds of USD 1.30 million). This pattern suggests that Top Liquidated funds might have delisted because of their inability to attract enough capital, potentially not covering their fixed costs, and almost certainly not generating the personal profits for which their managers had hoped.

We also looked at other characteristics of Liquidated and No Reason funds beyond the Top and Bottom classifications (based on positive or negative average returns for 6 months prior to delisting). We found that Liquidated funds have lower returns than No Reason funds during the previous 12- and 36-month periods (p-values of 0.01 and 0.00, respectively). Furthermore, alpha based on the Fung and Hsieh (2001) model estimated over the 36-month period is significantly lower for Liquidated funds, which would likely make it difficult to attract capital.

If we also use the Top and Bottom classification, we find Bottom Liquidated funds had an average monthly return of −0.003% over the 3 years prior to liquidation. These funds were actually losing money on average for 3 years, and it is not surprising they decided to liquidate. The Top Liquidated funds did have positive average returns of 0.54% over the 36 months prior to liquidation but suffered from small size and weak fund flow, as discussed previously.

In contrast, both Top and Bottom No Reason funds had positive average monthly returns (0.77% and 0.24%, respectively) over the 3 years prior to delisting. It seems likely that some of the Top No Reason funds were performing well, simply decided to "stop advertising," and had an unremarkable delisting return (not significantly different from the average return of a live fund). On the other hand, it may well be that many of the Bottom No Reason funds "blew up" and suddenly stopped reporting. Particularly for funds with illiquid positions, this would be consistent with the relatively large negative delisting returns we estimate for the Bottom No Reason group. Both these scenarios contrast with the Liquidated funds (Top and Bottom), many of which appear to have been slowly strangling prior to announcing their liquidation.

B. Sorting on Other Fund Characteristics

Sorting on variables other than performance does not yield significant differences in delisting returns when we try offshore versus onshore, audited versus unaudited, fund size, fund flow, styles, leverage, fees, serial correlation (proxying for investment liquidity), and loadings on various Fung and Hsieh (2001) factors. The corresponding tables are Tables IA.1–IA.8 in the Internet Appendix. To put these results in perspective, some of the variables may be mismeasured or missing (e.g., AUM is not reported consistently by many funds), resulting subgroups can turn out to be small, or variables might not be clearly related to performance (e.g., large positive as well as negative returns would both generate large variances). One exception, where we find some significance, is the sort with respect to estimated alpha from the Fung and Hsieh (2001) 7-factor model. The average delisting return of No Reason funds with high alpha is -1.03%, which is significantly higher at the 10% level than -5.87%, the average delisting return of No Reason funds with low alpha (see Table IA.6).

We tried combining the above variables with past performance in economically sensible ways. Such double sorting can give us additional information about economically important variables, but the significance may decline due to the smaller number of funds in each bin of the double sort. For example, it is reasonable to expect that hedge funds with positive past returns that also have relatively low volatility would have better delisting returns compared with hedge funds that historically had negative average returns and higher volatility. Indeed, we find such results when we consider Top funds with 30% lowest return volatility versus bottom funds with 30% highest return volatility (p-value of 0.13). The results are reported in Panel A of Table 5.

TABLE 5

Mean Delisting Returns: Top versus Bottom Funds with Double Sorts

Table 5 reports the monthly delisting returns (bias-corrected) based on matched portfolios of FoFs where the adjusted R^2 of the main regression model is at least 25% and the portfolio weight of the delisting fund is at least 5%. We use 1% tail trimming. Mean delisting returns and their standard deviations are in percentages per month. Panel A reports the results for Top funds having low return volatility and Bottom funds having high return volatility. Panel B reports the results for Top funds having high estimated alpha based on the Fung and Hsieh (2001) 7-factor model and Bottom funds having low estimated alphas. Panel C reports the results for Top offshore funds and Bottom onshore funds.

	No. of Matches	Mean Delisting Return	Bootstrapped STD of Mean Delisting Return	Nonparametric p-Value for Difference with Average HF Return	Nonparametric p-Value for Difference between Top and Bottom Funds
Panel A. Funds Sorted on Return STD					
Top, 30% Lowest Return STD					
All	298	0.42	1.47	0.49	0.13
Liquidated	70	4.47	3.20	0.21	0.37
No Reason	220	−0.89	1.70	0.30	0.13
Bottom, 30% Highest Return STD					
All	243	−5.43	2.52	0.09	—
Liquidated	59	−2.10	4.56	0.42	—
No Reason	176	−6.74	2.97	0.06	—
Panel B. Funds Sorted on Alphas					
Top, 30% Highest Alphas					
All	303	2.17	2.01	0.24	0.06
Liquidated	60	4.36	4.43	0.26	0.26
No Reason	230	2.24	2.28	0.28	0.06
Bottom, 30% Lowest Alphas					
All	254	−6.28	2.31	0.08	—
Liquidated	69	−2.96	4.10	0.38	—
No Reason	179	−7.63	2.86	0.07	—
Panel C. Funds Sorted on Offshore/Onshore					
Top, Offshore					
All	386	1.99	1.58	0.26	0.12
Liquidated	116	4.91	3.02	0.26	0.50
No Reason	259	0.72	1.81	0.36	0.08
Bottom, Onshore					
All	274	−4.79	2.14	0.15	—
Liquidated	50	4.91	4.40	0.32	—
No Reason	216	−7.37	2.46	0.08	—

In a similar vein, we tried separating out the best of the Top funds and the worst of the Bottom funds based on estimated alpha using the Fung and Hsieh (2001) 7-factor model over a longer (36-month) horizon. The results are in Panel B of Table 5 and show that the added conditioning on alpha increases the difference between All Top and All Bottom average delisting returns compared with Table 3 (p-value of 0.06).

Finally, we condition on being domiciled in the United States or not. At least some offshore hedge funds are presumably less regulated and more capable of exploiting profitable investment strategies (thus performing better). We double sort so that we compare offshore, Top funds versus onshore, Bottom funds. The better No Reason funds outperform the worse No Reason funds in terms of delisting returns (p-value of 0.08) (see Panel C of Table 5).

V. Robustness

In this section, we first evaluate the general quality of our matching algorithm, and then discuss the stability of our basic results to implementation changes in the estimation procedure.

A. Quality of the Matching Algorithm

We investigate the quality of our matching algorithm by constructing hypothetical FoF returns from reported hedge-fund returns using the basic simulation method described above. However, we introduce a fictitious delisting return drawn from a normal distribution with known mean and standard deviation. We use three settings for the simulated delisting returns: 1% mean and 5% standard deviation, which is quite close to the sample values in Panel B of Table 1; −10% mean and 5% standard deviation; and −50% mean and 10% standard deviation. We also address the issue that in practice some hedge funds in an FoF portfolio might not be observed in our database. For each simulation of delisting returns, we create three subscenarios: one where all 15 hedge funds can be found in the database; another where only 10 can be found; and a rather extreme situation where only 5 can be found. Details and further discussions can be found in Section IA.4 of the Internet Appendix.

We show in Table 6 that our procedure does a good job of recovering large negative mean delisting returns of −10% and −50%, and it does not mistakenly find large negative mean returns when the true mean delisting return is 1%. This is true even when only 33% of hedge funds are visible in the database. Thus, we are rather confident that our procedure would not miss a large and negative mean delisting return even if the database contained only a modest fraction of the hedge-fund universe.

Lockups, gates, and notice periods all make it difficult for an FoF manager to quickly alter the fund's portfolio; however, we recognize that an FoF may not be held constant for 36 months. To examine potential implications of this issue, we implemented a simulation using a monthly turnover rate for all FoFs of 1.8% (equivalent to 20% annually, which would correspond to roughly half of each FoF portfolio turning over in a 3-year period). Further details are in Section IA.5 of the Internet Appendix. If the delisting return was from a distribution with a

TABLE 6

Simulated Performance Results

Table 6 reports mean delisting returns as well as the bootstrapped standard deviations of the mean delisting return for simulated samples of FoF returns. Each FoF is modeled as a portfolio of 15 individual hedge funds. For simulated delisting funds, the hypothetical delisting return is drawn from a normal distribution with given mean (μ_E) and standard deviation (σ_E), expressed in percentages per month. The reported estimates are obtained using our standard procedure with a subset of the hedge funds used to generate the FoF returns being visible to our matching algorithm. We vary the fraction of visible funds using 100%, 67%, and 33% of the total generating set. We consider three possible delisting return distributions for hedge funds, characterized by pairs (μ_E, σ_E) of (1, 5), (−10, 5), and (−50, 10). Values are in percentages per month.

No. of Visible Funds	No. of Matches	Mean Delisting Return	Bootstrapped STD of Mean Delisting Return
Panel A. $(\mu_E, \sigma_E) = (1, 5)$			
15	2,590	−0.05	0.48
10	1,935	1.12	0.68
5	1,035	−0.71	0.52
Panel B. $(\mu_E, \sigma_E) = (-10, 5)$			
15	2,574	−8.43	0.51
10	1,914	−8.94	0.54
5	1,005	−6.34	0.54
Panel C. $(\mu_E, \sigma_E) = (-50, 10)$			
15	2,576	−39.84	0.89
10	1,884	−38.66	0.91
5	1,033	−33.62	0.96

mean of 1% and a standard deviation of 5%, our procedure finds a mean return of 0.86%. Even if the delisting return was from a distribution with a −10% monthly mean return and a standard deviation of 5%, or with a mean return of −50% and a standard deviation of 10%, the estimated mean delisting returns are also relatively accurate at −7.68% and −38.81%, respectively. This suggests that the estimated mean delisting returns reported in Table 2 are not very sensitive to the possibility of turnover in the FoF portfolios.

We also examined the accuracy of the matching algorithm and estimated portfolio weights by comparing the forecasted FoF portfolio return in the 37th month with the actual FoF return in those matches where we have no delisting funds (consequently, having a full set of returns for the 37th month). Our average forecast error is only 0.052% with a standard error of 1.76% for matches with R^2 above 25%.

B. Stability of the Empirical Results

To assess result stability, we also implemented our procedure using variations on the basic methodology. Tables with results using these variations on our standard approach are provided in the Internet Appendix as indicated below. Most resulting changes relative to the estimated mean delisting returns reported in Table 2 are substantially less than one bootstrapped standard deviation from the original estimate, and we interpret them as minor differences.

The variations on our basic methodology included:

a. Allowing investment in a riskless asset with a beta between 0.02 and 0.10. Given the small variability of the riskless rate, this is also essentially equivalent to adding a constant term when estimating equation (1). Results are in Table IA.9.

b. Using rolling windows of 30 and 42 months in Table IA.10.

c. Altering the minimum beta limit to 0.01 and to 0.04 in Table IA.11.

d. Increasing the maximum beta to 0.20 in Table IA.12.

e. Increasing the minimum R^2 to 0.50 in Table IA.13.

f. Employing 0.05 as the trimming level for excluding outliers from estimated delisting returns in Table IA.14.

g. Including only FoFs where we cannot reject the hypothesis of no serial correlation in returns at the 1% significance level in Table IA.15.

h. Reducing the number of hedge funds in the FoF portfolio to a lower limit of 10 in Table IA.16.

i. Allowing up to 26 hedge funds in a match, where 26 corresponds to the average reported number in FoF portfolios for our data. Results are in Table IA.17.

j. Employing a procedure that allows up to 32 hedge funds (as an upper limit) in a match, using a rolling window of 42 months. Results are in Table IA.18.

Another potential issue for our results concerns the possibility that an FoF manager identifies a hedge fund that seems likely to exit and seeks to unwind the FoF's position in that hedge fund before the exit takes place. As mentioned previously, gates, notice periods, etc. make it difficult for the FoF manager to quickly adjust and get out of a potentially exiting hedge fund. In the case of merged hedge funds or hedge funds closed to new investments, it might not even be desirable for the FoF to eliminate its positions in those funds. Moreover, predicting delisting is difficult. Thus, it is hard for FoFs to get out prior to an exit event. Nevertheless, we implemented a robustness check and reestimated the delisting returns assuming that in month 37, the actual holding of an FoF in the delisting fund is half of the estimated weight (beta). In effect, we are assuming the FoF was successful in identifying the exiting hedge fund and was able to unwind half its position prior to the exit. That half of the estimated weight was equally distributed among the surviving hedge funds in that FoF portfolio. Not surprisingly, halving the portfolio weight of exiting funds doubles (roughly) mean delisting returns (see Table IA.19). Furthermore, their bootstrapped standard errors similarly increase due to the smaller portfolio weight, and No Reason delisting returns are no longer significantly different from the average hedge-fund return.

We need a reasonably lengthy estimation period (e.g., 36 months) to get an estimated weight for the exit month and are thus limited to constant weights for FoF positions in hedge funds. Authors such as Bollen and Whaley (2009) and Patton and Ramadorai (2013) have shown that hedge fund and FoF risk exposures vary greatly over time. However, time-varying FoF exposure results from two effects: the time variation of FoF holdings of hedge funds and the time variation of risk exposures within those hedge funds. Our paper is concerned only with the former, and we argue that this effect is much smaller than the within hedge-fund variation.

Also, a number of features in our methodology attenuate the problem further. First, we reestimate FoF holdings for each sequential 36-month window. So weights are constant only within a window, not for the life of the FoF. Second, we used 30-, 36-, and 42-month windows, implying constant weights for different period lengths, with little change to the results. Third, we allow for simulated turnover within the FoF holdings, holding the weights constant but changing the selected hedge fund. That simulation indicates our methodology does a good job. Thus, our method appears quite robust to time variation in FoF holdings of hedge funds.

VI. Concluding Comments

Relatively little has been known about returns after hedge funds delist from a database. We examine the situation by modeling the econometric relationship between funds of funds and the portfolios of hedge funds into which they invest. This structure allows us to estimate the average delisting return of −1.61% for all delisting hedge funds. That estimate is not significantly different from the 0.56% average monthly return for all (live) hedge funds and nowhere near a disastrously negative number such as −50%. Our procedure for inferring FoF portfolio holdings is noisy, but with a large number of matches (nearly 1,400 in Table 2), we obtain enough precision to have confidence in our average estimates.

We also find that returns of delisting hedge funds are somewhat persistent, with the results for Bottom funds being quite pronounced. We divided funds with negative performance over the previous 6 months into those that also stated they were being liquidated (Bottom Liquidated) and those that did not provide a clear reason for exiting (Bottom No Reason). The Bottom Liquidated funds had an unremarkable mean delisting return of −0.59%, but the Bottom No Reason funds had a strongly negative delisting return of −5.97%, which is significantly below 0.56% (average monthly return of all hedge funds) with a p-value of 0.01. As discussed earlier, it seems likely that many of the Bottom No Reason funds may have been forced to exit suddenly under adverse circumstances.

It is straightforward that funds with negative prior returns might decide to delist, but why are a substantial number of Top funds opting to exit? We drilled deeper into this issue and found that Top Liquidated funds had a relatively large mean delisting return of 5.46%; however, this estimate also has a relatively large bootstrapped STD of 2.35%. One tends to think that funds are liquidated because of poor performance, and this result seems inconsistent with that view. However, those funds are small and have been experiencing weak fund flow. Thus, it appears that the Top Liquidators are a set of small funds that did not perform well enough to attract substantial inflows and achieve a critical mass, where they became sufficiently profitable for their managers. Hence after struggling for a time, those managers may simply have opted to liquidate the funds and move on to more promising endeavors.

In summary, we find that most exiting funds are in categories that have mean delisting returns that are not significantly different from the average monthly return of live hedge funds. In marked contrast, Bottom No Reason funds have a mean delisting return (-5.97%) that may result from being forced to exit suddenly under adverse circumstances. Even so, -5.97% is a long way from a disastrous number such as -50%.

References

Ackermann, C.; R. McEnally; and D. Ravenscraft. "The Performance of Hedge Funds." *Journal of Finance,* 54 (1999), 833–874.

Agarwal, V.; V. Fos; and W. Jiang. "Inferring Reporting-Related Biases in Hedge Fund Databases from Hedge Fund Equity Holdings." *Management Science,* 59 (2013), 1271–1289.

Aiken, A. L.; C. P. Clifford; and J. Ellis. "Out of the Dark: Hedge Fund Reporting Biases and Commercial Databases." *Review of Financial Studies,* 26 (2013), 208–243.

Bollen, N. P. B., and R. E. Whaley. "Hedge Fund Risk Dynamics: Implications for Performance Appraisal." *Journal of Finance,* 64 (2009), 987–1037.

Brooks, C.; A. Clare; and N. Motson. "The Gross Truth about Hedge Fund Performance and Risk: The Impact of Incentive Fees." *Journal of Financial Transformation,* 24 (2008), 33–42.

Brown, S. J.; W. N. Goetzmann; and R. G. Ibbotson. "Offshore Hedge Funds: Survival and Performance, 1989–95." *Journal of Business,* 72 (1999), 91–117.

Fung, W., and D. A. Hsieh. "Performance Characteristics of Hedge Funds and CTA Funds: Natural versus Spurious Biases." *Journal of Financial and Quantitative Analysis,* 35 (2000), 291–307.

Fung, W., and D. A. Hsieh. "The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers." *Review of Financial Studies*, 14 (2001), 313–341.

Fung, W.; D. A. Hsieh; N. Y. Naik; and T. Ramadorai. "Hedge Funds: Performance, Risk, and Capital Formation." *Journal of Finance,* 63 (2008), 1777–1803.

Getmansky, M.; A. W. Lo; and I. Makarov. "An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns." *Journal of Financial Economics,* 74 (2004), 529–610.

ter Horst, J., and M. Verbeek. "Fund Liquidation, Self-Selection and Look-Ahead Bias in the Hedge Fund Industry." *Review of Finance,* 11 (2007), 605–632.

Kolokolova, O. "Strategic Behavior within Families of Hedge Funds." *Journal of Banking and Finance,* 35 (2011), 1645–1662.

Liang, B. "Hedge Funds: The Living and the Dead." *Journal of Financial and Quantitative Analysis,* 35 (2000), 309–326.

Patton, A., and T. Ramadorai. "On the High-Frequency Dynamics of Hedge Fund Risk Exposures." *Journal of Finance,* 68 (2013), 597–635.

Posthuma, N., and P. J. van der Sluis. "A Critical Examination of Historical Hedge Fund Returns." In *Intelligent Hedge Fund Investing: Successfully Avoiding Pitfalls through Better Risk Evaluation*, Chap. 13, B. Schachter, ed. London, UK: Risk Books (2004), 365–388.

Van, G. P., and Z. Song. "Hedge Fund Commentary from VAN." Working Paper, Van Hedge Fund Advisors International, www.edge-fund.com/VanSong2005.pdf (2005).