Hedge Funds: The Good, the (not-so) Bad, and the Ugly

Yong Chen
Pamplin College of Business
Virginia Tech
(540) 231-4377
ychen75@vt.edu

Michael T. Cliff
Pamplin College of Business
Virginia Tech
(540) 231-5061
mcliff@vt.edu

October 20, 2008

PRELIMINARY AND INCOMPLETE.
PLEASE DO NOT QUOTE WITHOUT PERMISSION.

Abstract

We propose a new methodology to evaluate the skill of fund managers. We assume that a manager's skill, characterized by alpha, is drawn from one of several distributions based on their ability (e.g., Good, Neutral, or Bad). Each of these distributions has a mean and variance of alphas. The composite distribution of a sample of fund managers is therefore a mixture of Normal distributions. We estimate these distributional parameters by fitting the composite distribution to the sample of alphas. We are also able to calculate the probability a manager is of a given type conditional on their alpha. We illustrate our approach in a sample of hedge funds. We find that about half of the managers in our sample have skill, even though only 20% have alphas that are significantly positive.

1 Introduction

The hedge fund industry has grown dramatically in the recent two decades, generating tremendous interest from academics and practitioners in hedge fund performance evaluation. Several existing studies [e.g., Ackermann, McEnally, and Ravenscraft (1999); Brown, Goetzmann, and Ibbotson (1999); Liang (1999)] provide evidence that hedge funds, on average, deliver positive abnormal performance and outperform mutual funds. Recently, Kosowski, Naik, and Teo (2007) employ a bootstrap analysis to show that top hedge fund performance cannot be attributed to pure luck and hedge fund performance persists at an annual horizon. Despite such evidence on the existence of superior performance in hedge funds, an important question is still unanswered: what fraction of hedge funds actually possesses the ability to deliver alpha? In other words, is the documented abnormal return driven by only a small number of top-performing funds, or do most hedge funds generate positive risk-adjusted returns? This question is of great interest to hedge fund investors since its answer should have direct implications for how to successfully identify and choose skilled funds to realize investment profits. Our paper attempts to answer the question for the first time in the hedge fund literature by applying a new statistical technique to a large sample of hedge funds.

Detection of skill among hedge fund managers is a challenging task. Since skill is not directly observable, a proxy such as alpha—the excess return to a portfolio that has no ability but the same factor exposures as the fund—is often used. The challenges of using alpha to detect skill fall into two broad categories. The first is the problem of false positives that arises in any setting where hypothesis tests are applied to an entire set of estimates. At a stated nominal significance level such as 5%, we expect to see that fraction of the sample with significant t-statistics by chance alone. If 25 managers in a group of 1000 have "significant" alphas it would be incorrect to conclude that these "top managers" have skill. This issue would be even more important if the distribution of alphas is non-Normal.

The second broad challenge in detecting skill is the fact that estimates of alpha may be noisy or biased. Noisy alphas make it difficult to statistically measure economically meaningful skill. The problem begins with the fact that the fund's underlying investments are noisy. Compounding the problem is the relatively short history of returns for most funds. Biased alphas will arise due to several flavors of sample selection bias, as well as possible misspecification of the fund's risk exposure. Among the selection biases are an "instant history" bias as new funds entering the database are tilted toward those with unusual prior success. The voluntary nature of hedge fund reporting means that there is also a "delisting bias" since funds that do very poorly will often not report these returns to the database. Biases in alpha related to the misspecification of the risk factors can be due to the failure to include the proper risk factors to complications arising from funds' dynamic trading behavior. For example, Foster and Young (2008) note that a manager can use simple call-writing strategy to deliver fake alpha. We calculate that a manager following such a strategy can earn a return of 7.8% with 90% probability. The other 10% of the time investor's are wiped out, but this –100% return does not show up in the database due to the aforementioned delisting bias.

In this paper we develop a new approach to evaluate the prevalence of skilled managers. In this regard, our paper is similar in spirit to the recent paper on skill among mutual fund managers by Barras, Scaillet, and Wermers (2007). Our approach extends theirs in two important ways. First, we are able to infer not only the fraction of funds that have skill, but also the magnitude and variability of skill. Second, *for each fund* we are able to make probabilistic statements about the likelihood that the manager has skill, given their estimated alpha.

The intuition of our approach is to view each fund manager i, characterized by his skill α_i, as coming from one of several distributions with mean alpha equal to μ_j. As a stylized example, we can view managers as being either "Good" (say $\mu_G = 50$ basis points per month), "Bad" (e.g., $\mu_B = -50$ bp), or "Neutral" ($\mu_N = 0$). Therefore the observed distribution of fund alphas comes from a mixture of distributions and we back out the parameters governing these distributions using the EM algorithm. The estimated mixing probabilities tell us the fraction of funds that are Good, Bad, or Neutral and we also can quantify the magnitude and variability of alpha for each skill-type. These parameters can then be used with the estimated alpha for an individual fund to determine the likelihood that the particular manager comes from each skill-type.

The key input to our approach is an estimate of alpha for each fund. We propose several refinements to the alphas that are commonly used to evaluate hedge fund performance. First, we address the instant history bias by including a dummy variable for the first twelve (or twenty) months of returns. The standard approach is to drop these months which is problematic given the already short history for most funds. Second, we extend the standard Fung and Hsieh (2004) seven-factor model to include the book-to-market (HML) and momentum (UMD) factors. These factors are known to generate positive average returns, so a model excluding them would reward exposure to these factors with alpha.

We take our approach to the data to assess performance among a sample of 4,965 hedge funds from TASS. We find that about one third of all funds, and nearly 40% of live funds, have seven-factor alphas that are statistically greater than zero. The average fund has an alpha of 39 basis points per month, similar to the findings of Kosowski, Naik, and Teo (2007). Controlling for the instant history bias with our dummy variable reduces the average alpha by ten basis points, leaving a quarter of funds with significantly positive alphas. The "luck" during the incubation period as measured by the dummy variable is 45 bp per month. Including the HML and UMD factors leads to a similar, but separate, reduction in alphas. With our final "full" specification for alpha we find that 16% of funds have statistically significant evidence of good performance. Under three percent have significantly negative alphas. The average fund has alpha of 12 bp with a t-statistic of 0.7. However, as noted earlier it would be inappropriate to conclude from this that a sixth of our managers have skill.

When we apply our skill-detection technique to the subset of "pure" (not fund-of-funds) funds, we find that 50% have positive skill (39% "Good" $\mu = 24$ bp, and 11% "Excellent" $\mu = 134$ bp). Of the remaining funds, 8% are deemed to be "Bad" ($\alpha = -147$ bp) and 42% Neutral. The frequency of skilled or unskilled managers is much greater than the fraction of alphas that are significantly positive or negative in this subsample, 20% and 2% respectively. To further illustrate the richness of our approach, we note that a fund with an alpha of 0.56% (in the "middle" of μ_G and μ_E in terms of their standard deviations) would be deemed more likely to be Good—or even Neutral—than Excellent. This seemingly counter-intuitive result is due to the rarity of Excellent managers.

The remainder of the paper proceeds as follows. In Section 2 we outline our approach to inferring the parameters governing the family of distributions of managerial type. Section 3 discusses our data. In Section 4 we discuss our extensions to the seven-factor alphas and present the estimates of alpha. Section 5 implements our procedure for estimating the distributional parameters of the skill types. We offer some remarks on refining and extending our approach in Section 6 and conclude in Section 7. Appendix A has some details on alternative estimation approaches.

2 Detecting Manager Type

As a motivating example, suppose there are three types of fund managers: Good, Neutral, and Bad. Each type j is characterized by its distribution of managerial skill, Manager i belongs to one of the three types j and has ability $\alpha_i \sim \mathcal{N}(\mu_j, \sigma_j^2)$. We define the meaning of the Neutral type by setting $\mu_N = 0$.[1] The composition of the manager types is defined by their frequency in the population of managers, π_j with $\pi_G + \pi_N + \pi_B = 1$.

To illustrate our approach, suppose the distributions of funds types are parameterized as

[1]Berk and Green (2004) argue that with declining returns to scale, competition will drive after-fee alphas to zero. Our "neutral" managers have skill equal to their fees, so from the investor's perspective these managers are neither good nor bad.

shown in the table below. The figure shows the distribution for each manager type as well as the composite distribution. It is important to note that the composite distribution is highly non-Normal.

Skill Type	μ_j	σ_j	π_j
Good	3	1	.05
Neutral	0	1	.70
Bad	–2	1	.25

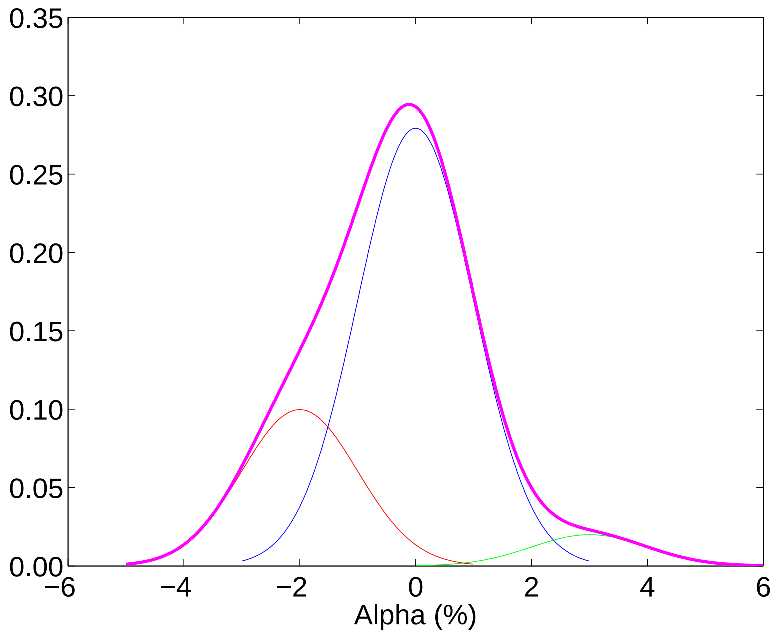

2.1 The General Setup

The econometrician observes fund returns, from which he can estimate the skill of each manager, $\hat{\alpha}_i$. We defer discussion of these non-trivial details to Section 4. The issue we focus on here is whether it is possible to estimate the parameters $\{\pi_j, \mu_j, \sigma_j\}$ from the composite distribution of the estimated alphas.

Barras, Scaillet, and Wermers (2007, hereafter BSW) propose a methodology to control for "false positives" in mutual fund performance evaluation. In particular, they correct for the fact that some significantly positive alphas are due to neutral managers with good luck (and some negative alphas are neutral managers with bad luck). In essence, they measure the fraction of alpha t-statistics larger than a specified value (e.g., 1.96). Under the null of no skill, 2.5% of the sample should reside in this tail. The "true" number of skilled managers is then the actual fraction of the sample above the critical value, less the 2.5% that are presumably there by chance.

While their approach is quite innovative, it suffers from two shortcomings. First, they assume that luck only affects neutral managers. They do not allow for skilled managers with bad luck (i.e., observed alphas that are insignificant or even significantly negative) or unskilled (bad) managers with good luck.[2] Basically, they are operating under the null of

[2]More precisely, they assume that luck is not sufficiently important so as to cause a non-neutral manager to have an alpha with p-value above some critical value such as 0.6.

no abnormal performance. Under the alternative that some managers have skill (good or bad), the observed performance of these skilled managers would also be affected by luck. Second, their approach allows estimation only of parameters similar to π; there is no way to comment on important issues such as the magnitude of the alphas for the Good managers.

Our methodology is similar in spirit to BSW but we aim to address these shortcomings. We allow for luck to affect all managers. Following our motivating example, we posit that there are J manager types (e.g., Good, Neutral, and Bad). Each type j is characterized by a Normal distribution of its alphas with mean μ_j and standard deviation σ_j. The composite distribution is then a mixture of Normals with mixing parameters π. We then choose the parameters $\{\hat{\mu}_j, \hat{\sigma}_j, \hat{\pi}_j\}$ by fitting this composite distribution to the the empirical distribution of $\hat{\alpha}$. We use a goodness-of-fit criteria to select the number of distributions J.

Once we have estimates of the parameters governing the distribution of α, we can then calculate the posterior probability that manager i belongs to each type j,

$$\Pr(\text{Mgr } i \text{ is type } j|\hat{\alpha}_i) = \frac{\hat{\pi}_j \phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)}{\sum_j \hat{\pi}_j \phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)} \tag{1}$$

where $\phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)$ is the Normal probability density with mean $\hat{\mu}_j$ and volatility $\hat{\sigma}_j$ evaluated at $\hat{\alpha}_i$. We emphasize that these posterior probabilities are not possible in the BSW approach.

2.2 Estimation Approach

We need a way to estimate the parameters for each family in the mixture of distributions. We focus our discussion on the EM algorithm, though we briefly mention several alternative estimation approaches (see Appendix A for additional details).

2.2.1 The EM Algorithm

The Expectation-Maximization (EM) algorithm iterates the following two steps until convergence. The iteration $\ell + 1$ estimates are

1. E-step:

$$p(\alpha_i \in \mathcal{A}_j |\ \alpha_i, \boldsymbol{\theta}_\ell) = \frac{\pi_{j,\ell}\ \phi(\alpha_i; \mu_{j,\ell}, \sigma_{j,\ell})}{\sum_{j=1}^{J} \pi_{j,\ell}\ \phi(\alpha_i; \mu_{j,\ell}, \sigma_{j,\ell})}$$

2. M-step:

$$\pi_{j,\ell+1} = \frac{1}{N}\sum_{i=1}^{N} p(\alpha_i \in \mathcal{A}_j|\ \alpha_i, \boldsymbol{\theta}_\ell)$$

$$\mu_{j,\ell+1} = \frac{1}{N}\sum_{i=1}^{N} \frac{p(\alpha_i \in \mathcal{A}_j|\ \alpha_i, \boldsymbol{\theta}_\ell)}{\pi_{j,\ell+1}}\ \alpha_i$$

$$\sigma^2_{j,\ell+1} = \frac{1}{N}\sum_{i=1}^{N} \frac{p(\alpha_i \in \mathcal{A}_j|\ \alpha_i, \boldsymbol{\theta}_\ell)}{\pi_{j,\ell+1}}\ (\alpha_i - \mu_{j,\ell+1})^2$$

The E-step uses equation (1) to calculate the probability that each data point comes from distribution j given the current parameter estimates. These probabilities are then used in the M-step to obtain maximum likelihood estimates of the parameters. The unconditional probability that a data point comes from a particular distribution (π_j) is the average across all data points of these probabilities. The mean and variance are weighted averages of the data or squared deviations from the mean, respectively. The weights for distribution j, $p(\alpha_i \in \mathcal{A}_j|\ \alpha_i, \boldsymbol{\theta}_\ell)/\pi_{j,\ell+1}$, pay more attention to observations deemed likely to be members of that distribution. As noted above, we impose the restriction $\mu_N = 0$ when estimating the parameters.

To gauge the performance of the EM procedure we evaluate the parameters in a simulated dataset and compare the estimates to the true parameters. Our simulated dataset contains alphas for 2,000 funds using the parameters shown in the earlier example.[3] This simulated dataset is in fact used to generate the graph above.

We compare our parameter estimates from the EM procedure to the true parameters in the first two columns of Table 1. Our estimated mean alpha for the Good manager type is 3.2, slightly above the true value of 3. For Bad managers, we estimate a mean of –1.9 which is close to the true value of 2. Estimated volatilities are all within about 0.1 of the true value of 1. Finally, the unconditional probabilities of manager type conform closely to the true values. We estimate that 4.1% of managers are Good (vs 5.0% truth), 70.7% are Neutral (vs 70.0%), and 25.2% Bad (vs 25.0%).

To help interpret whether our estimates are economically close to the true parameters, Table 2 evaluates equation (1) to measure the probability of each manager type conditional on observing an alpha of 3 (= μ_G), 1.5, 0 (= μ_N), –1, or –2 (= μ_B). Suppose a manager's estimated alpha is 3, exactly equal to the mean of the Good type. Intuitively, we would

[3] Since we currently examine only a single run of the simulation we adjust the randomly generated alphas to match the true parameters exactly. In future work we plan to leave this Monte Carlo sampling variation in tact but examine results across a large number of simulated datasets.

expect that this manager is most likely of type Good, though we would not be certain that the manager was lucky but Neutral skill. Panel A confirms this intuition. The true probability that a manager with $\hat{\alpha} = 3$ is actually Good is 86.5%. There is a 13.5% chance a Neutral manager got lucky, and a negligible chance a Bad manager was so fortunate. Our estimated probabilities are 78.5% Good and 21.5% Neutral, so we are off by about 8 percentage points. Moving to Panel B, we see a better correspondence between the estimated and true probabilities. If $\hat{\alpha} = 1.5$ we estimate a 2.7% chance of Good (vs actual 6.7%), 96.7% Neutral (vs 93.1%), and 0.6% Bad (vs 0.2%). Looking across the remaining panels we see a pattern whereby our estimated probabilities are "shrunk" toward the middle. When estimated alpha is a large positive value we understate the probability of the Good type. On the other side, when estimated alpha is a large negative value we underestimate the probability of the Bad type.

Overall, we conclude that the EM algorithm yields results that, while imperfect, are at least useful for characterizing the likelihood that a manager has skill. We now examine the performance of some alternative estimation approaches to see how our EM approach performs on a relative basis.

2.2.2 Alternative Estimation Approaches

We also consider three estimation approaches that are based on the moment generating function (MGF). We assume our data come from Normal distributions which have MGFs

$$M_x(\theta) = \exp\left(\theta\mu + \frac{1}{2}\theta^2\sigma^2\right).$$

The MGF for a mixture of Normals with mixing parameters π is a weighted average of the MGFs $\sum_{j=1}^{J} \pi_j M_j(\theta)$. In our experiments, these approaches are less robust than the EM algorithm. We briefly introduce each here and compare the estimation results to the EM algorithm. The interested reader can find details on these approaches in Appendix A.

The first approach uses the method of moments (MM) to match sample moments with derivatives of the moment generating function,

$$\left.\frac{\partial^n M_x(\theta)}{\partial\theta^n}\right|_{\theta=0} = \frac{1}{T}\sum_{t=1}^{T} x_t^n.$$

The second approach (MGF) uses the moment generating function directly,

$$\exp\left(\theta\mu + \frac{1}{2}\theta^2\sigma^2\right) = \frac{1}{T}\sum_{t=1}^{T}\exp(\theta x_t).$$

The third approach uses the characteristic function (CF). For a Normal variable, this means

$$\exp\left(i\theta\mu - \frac{1}{2}\theta^2\sigma^2\right) = \frac{1}{T}\sum_{t=1}^{T}\exp(i\theta x_t).$$

Tables 1 and 2 show the results from these approaches in the final three columns. The parameter estimates in Table 1 show that the MM and MGF approaches are perhaps slightly better than EM in estimating the distributional means, but worse when it comes to variance or unconditional probabilities. For example, the EM estimate of π_B was only 0.2% above the true value of 25%, while the MM estimate is 14.2% and MGF is 10.0%. On the other hand the CF estimates seem comparable in quality to the EM results.

Since it is hard to know whether it is more important to correctly estimate $\boldsymbol{\mu}$, $\boldsymbol{\sigma}$, or $\boldsymbol{\pi}$, we look to Table 2 to measure the combined effect. Here it becomes clear that reliable estimates of the conditional probabilities are not predominantly driven by the quality of the $\boldsymbol{\mu}$ estimate. The MM and MGF estimates, which provided better estimates of μ_G than the EM approach, do a very poor job at estimating conditional probabilities when the estimated alpha departs from zero. For example, when $\hat{\alpha} = 3$, the MGF probability that the manager is Good is only 22.9%, compared to a true probability of 86.5%. The CF estimates of the conditional probabilities, perhaps not surprisingly, are quite good. However they are slightly worse than our EM estimates. This fact, combined with the results in Table 1, reinforces the idea that main weakness of the EM approach—estimating $\boldsymbol{\mu}$ correctly—is the least important factor.

2.2.3 Selecting the Number of Distributions

When estimating the parameters in actual data we can be less sure that the composite distribution of alphas comes from a mixture of three Normal distributions. Therefore we compare estimates assuming different numbers of distributions and chose the one that fits the data the best. We use the Bayesian Information Criterion (BIC) to define the model fit,

$$BIC = q\ln(MSE) + k\ln(q)$$

where q is the number of points at which the density is evaluated (we use 500, spaced equally between ± 3), MSE is the mean-squared error between the "empirical" kernel density estimate (KDE) and our parametric density, and k is the number of parameters.

3 Data

Our sample of hedge funds comes from TASS. We start with 6,874 live or dead funds from 1994 through February 2008.[4] We only retain funds that report US Dollar returns net of fees and have at least 24 months of returns. This leaves 5,106 funds, from which we drop 141 funds with alphas we deem to be outliers (monthly $|\hat{\alpha}| > 3\%$).[5] Our final sample contains 4,965 funds, of which 2,444 (49.2%) are live and 2,521 are dead.

Since a fund of funds (FoF) has an additional layer of fees and is more diversified than pure hedge funds, we also segment our sample on that basis. Table 3 shows the frequency of live and total funds across each fund category. We have 1,050 FoFs (21.1% of the sample), of which 639 are live (60.9% of all FoFs). Non-FoFs have a much higher mortality rate, as only 46.1% are still alive. In part the lower mortality rate for FoFs is driven by the fact that they are younger and therefore have had less chance to fail. Table 4 measures the annual attrition rate as the fraction of funds living at the start of a year that do not survive to the end of the year. Non-FoFs fail at roughly double the rate as FoFs, consistent with the diversification offered by FoFs.

In addition to monthly fund returns, the TASS database provides a number of other variables measuring current fund characteristics. We have the dates of fund inception and when they start reporting to TASS (which can be later than their first return if the fund backfills its history). We also have the stated incentive and management fees (*Ifee* and *Mfee*) as well as a dummy variable set to one for funds with a high-water mark (HWM). For some, but not all, funds we have information on their investment terms such as lockup period, notice period, and minimum investment. We also have incomplete information on whether the manager invests his own capital, whether the fund is registered, whether the fund uses leverage, and the size of the fund (current and average). For funds that are designated as dead, we have some information on the reason they are no longer in the database (e.g., liquidated, merged, no report/answer).

Table 5 reports summary statistics for our sample. The first three columns show the number of non-missing observations, mean, and median for the full sample. The next three

[4]The TASS database contains funds back to 1977. They do not retain dead funds before 1994 so the earlier period contains a strong survivor bias.

[5]Kosowski, Naik, and Teo (2007) report monthly alphas range from –3.5% to 4.1%. Several of the funds we drop have alphas larger in magnitude than 10% per month.

columns repeat these calculations for the Live funds and the final three columns pertain to the Non-FoF sample. Only about 6% of funds are registered. Most do not have a lock-up period; among those that do, it is roughly one year. The typical fund requires a $4 million investment and requires about one month advance notice for withdrawals. Surprisingly, only about a third of managers report that they invest their own capital. Many of those indicating they do invest do not disclose the amount. Among the managers indicating an investment amount it is a rather small $4 million. The median fund has an incentive fee of 20% but the mean is somewhat lower at 16%. The mean or median management fee is about 1.5%, somewhat below the 2% rule of thumb. About two-thirds of fund have high water marks and a bit over half claim to use leverage.

The final few rows provide some information about the returns. Funds in our sample typically have five years of available returns. The average fund earns 0.87% per month and, not surprisignly, live funds earn higher returns than dead funds. The average fund has volatility of 3.25% per month (about 11% annualized). Live funds have lower volatility while non-FoFs have higher volatility. Finally, we see fairly high autocorrelation of monthly returns, about 0.14. This persistence has received considerable attention in the literature [e.g., Getmansky, Lo, and Makarov (2004)] and has a bearing on the proper estimation of the fund's systematic risk. Figure 1 shows these results graphically, as well as the Sharpe ratio. Dead funds have lower average returns but higher volatilities, so their Sharpe ratios are lower. Funds-of-funds tend to have lower average returns and volatilities, especially among dead funds, though Sharpe ratios are comparable.

4 Estimation of Alpha

The challenge of estimating abnormal performance for hedge funds is well-known in the literature. We broadly categorize these issues into two groups which we address in turn. First are a series of selection biases that, in one form or another, mean the econometrician does not have a representative sample of the fund's returns. Second is the issue of how to benchmark the fair return on the fund given its risk. Even if the econometrician does have a fund's entire history, the fact that hedge funds employ highly dynamic trading strategies and invest in a wide spectrum of asset classes means that a traditional Jensen's alpha would be inappropriate.

4.1 Selection Biases

The voluntary nature of hedge fund reporting creates a number of potential biases in the data [see, for example, Fung and Hsieh (2000) and Liang (2000).] The problem that we are best able to address is the "incubation" bias that arises due to backfilling of returns as new funds are added to the database. The literature has long recognized this problem. The typical solution is to drop the first 12 months of data [e.g., Kosowski, Naik, and Teo (2007)]. The problem with this approach is that most funds have relatively short histories to begin with, so it is preferable to avoid discarding precious data. Our approach is simply to include an incubation dummy variable in the regressions when estimating alpha. Specifically, we set the dummy variable to one during the first 12 (or twenty) months of the fund's history. This coefficient estimate on this dummy variable captures the incremental return the fund earns during this window.

A second bias we envision addressing is the standard "survivor" bias arising because poor funds die leaving a sample of survivor funds that overstates an investor's ex ante investment opportunities. In part, we can address this since we do have some returns for dead funds. However, the sample of dead funds suffers from a "delisting" bias. It is quite likely that funds stop reporting when they are having poor performance. If the fund recovers, it can then backfill these missing returns. If the fund eventually collapses, those particularly bad returns never make it into the database. One potential solution is to directly model the live or dead classification as a function of alpha and other controls, then use these estimated probabilities to infer the distribution of alphas as a truncated regression model. Cochrane (2005) uses a similar approach to modeling the return to venture capital, where returns are generally not observable unless the fund is successful. Section 6 briefly outlines our planned approach.

4.2 Benchmark Risk Factors

Numerous authors argue that hedge fund performance evaluation needs to account for the variety of asset classes in which funds may invest. Fung and Hsieh (2004) suggest a seven-factor model that adds the following factors to the standard CAPM

- Difference in returns on small and big stocks
- Change in yield on 10-year Treasuries
- Change in credit spread
- Bond lookback option
- Currency lookback option

- Commodities lookback option

To the extent that their model captures the risk of these "style bets," alphas from their model will not reward a passive position in these styles. It is worth noting that their factors certainly do not encompass all possible passive strategies. It is well documented that value stocks or past winners earn large average returns relative to our current understanding of the risk they entail. Thus, a lazy hedge fund manager could produce "alpha" simply by investing in HML or UMD. To address this point, we augment their model to include HML and UMD.

Another issue we address is the nonsynchroneity between fund returns and the factors. Hedge funds often have positions in highly illiquid assets, making it challenging to get accurate pricing. Funds must often rely on using stale historical prices or "mark-to-model" pricing of these illiquid positions. To partially mitigate this issue, we include lags of the market return and selected other factors (based on the significance of their slopes) in our expanded models.

4.3 Results

We begin by presenting alphas from the seven-factor model. Figure 2 shows the key properties of these alphas. Funds that exit the TASS database have very low average alphas, in the case of FoFs, or highly variable alphas, in the case of non-FoFs. These patterns are consistent with the extra layer of fees and diversification associated with FoFs. Dead funds also tend to have negatively skewed alphas. Finally, alphas for all fund categories have large kurtosis. Figure 3 shows the empirical distribution (KDE) of alphas for these four subsamples. It is plain to see that the composite distributions are non-Normal. Later we will implement our EM procedure to fit these composite distributions as mixtures of Normals.

Figure 4 provides a look at the distribution of alphas by fund investment strategy. The top panel reports boxplots of the alpha distributions for live funds only. The boxes represent the interquartile range, so all fund categories have positive alphas (among live funds) at the 25^{th} percentile. Given the dispersion in alphas it does not appear that a particular investment strategy is clearly superior. The bottom panel shows the fraction of funds still alive. Here the variation across categories is a bit more noticeable. Both FoF and multi-strategy (MS) have higher survival rates, consistent with their diversified nature.[6]

We now examine the effect of controlling for the incubation bias and variations on the factors in the model to estimate alpha. Tables 6 through 9 use the same basic layout. Each

[6] A caveat of this analysis is that it does not control for the age of the funds. A "new" strategy would have a high survival rate as the member funds have not yet had a chance to fail.

table compares the properties of alpha from the base (seven-factor) model to the alternative (Alt) under consideration. There are three sets of comparisons, comprising the rightmost six columns of these tables. The first pair is the average alpha under the base and alternative models. Next is the fraction of funds with alphas that are significantly negative (t-stat below –2). The final two columns indicate the fraction of alphas that are significantly positive ($t > 2$). The first column(s) of the table report average coefficient estimates unique to the alternative model under consideration. For each of these tables, the rows show results for the full sample (All) or the four subgroups formed on Live/Dead and non-FoF/FoF.

The alternative model under consideration in Table 6 adds a dummy variable for the "incubation" period to the seven-factor base model. In Panel A we assume a twelve month incubation period and Panel B uses 20 months. The first column shows the point estimate on the incubation dummy. It is clear that funds that make it into the TASS database have an unusually strong early period. Across all funds, for each month during the first year the typical fund earns 45 bp more than normal. This effect is especially strong for non-FoFs that die. These funds have an incubation effect of 66 bp per month. The fact that these funds die suggests that these managers were unskilled but initially lucky. The next two columns show the impact the incubation dummy has on the alphas. Across all funds, alphas drop from 39 bp to 29 bp. This 10 bp drop is smaller than the magnitude of the incubation dummy because the base model spreads the incubation performance over the funds life which is about five times as long as the incubation period. While the average alphas are now lower than before, they are still all positive. The final sets of columns in Panel A show that very few funds (under either model) have negative alphas. On the other hand, a sizable number of funds have significantly positive alphas. For example, across All funds, nearly one third have t-statistics bigger than 2. The incubation adjustment reduces this fraction somewhat, but it is still over one-quarter. Results in Panel B with the 20-month incubation period complement those in Panel A. The magnitude of the incubation dummy is a bit smaller, but the overall alphas and the fraction that are significant decline a bit from Panel A.

Next we add HML and UMD to the seven-factor model, but do not include the incubation dummy. Table 7 shows the results. The first two columns indicate that hedge funds tend to make bets on HML and UMD. The average loadings on these factors are not huge, reaching only 0.13 on the high end. But since passively following these strategies does not require (much) skill, fund managers should not get credit for doing so (at least not 2+20 type of credit!). The next pair of columns shows that the overall effect on alpha across all funds is about 6 bp. Within the the live funds the effect is a bit larger at around 10 bp. Yet even after making this adjustment, 27% of funds still have significantly positive alphas.

One possibility is that fund managers may have changed their exposures to HML and

UMD over time. For example, adjusting for momentum in mutual fund performance became more common following Carhart (1997). The seven-factor model for hedge funds was proposed several years later by Fung and Hsieh (2004). We allow for the fund's exposure to HML and UMD to shift following 2004. These coefficients, shown in the second and fourth columns of Table 8, indicate that funds reduced their HML exposure (to roughly neutral) but increased their UMD bets (from nearly neutral). Figure 5 shows the UMD loadings (overall and change) by fund investment type for the live funds. The shift toward UMD was consistently large across all investment categories. The remainder of the columns in the table are similar in message to the earlier versions: overall alphas decline a bit (about 7 bp), almost no funds have significantly negative alphas, and a large fraction have significantly positive alphas even after allowing changes in factor loadings.[7]

In Table 9 we combine these effects and also include lagged factor returns to partially address the nonsynchronous trading effect. In particular, the alphas in Panel A are the intercept from the regression of excess fund returns on the seven Fung-Hsieh factors, the 20-month incubation dummy, HML, UMD, and two lags of the market return. The model in Panel B expands this list of variables to include lagged values of any other factors that have statistically significant slope coefficients. Results are similar for both panels so we focus on those in Panel B. Across all funds, the average alpha is 12 bp under this model versus 39 bp from the seven-factor model. Thus, while the individual effects of our modifications were not huge, the cumulative effect is substantial. It remains the case, however, that there is very little evidence of bad performance. There is some slight variation across our classifications of funds, but the fraction of funds with significantly negative alphas is not much different than what we would expect by chance. The fraction of funds with significantly positive is now 16% overall, still far above the null of 2.5%, but about half the amount indicated by the seven-factor model. With the dead funds, our alternative model now yields average alphas that are close to zero, and about one in eight have significantly positive alphas. Figure 6 plots the distributions of these "full" alphas against those from the base seven-factor model. The plots make clear the reduction in alpha from our full model.

Our estimated alphas are consistent with the notion that a meaningful number of fund managers have skill above and beyond their fees. However, it is important to bear in mind some of the sample selection issues noted in Section 4.1. Most of these effects would cause the observed sample to overstate the ability in the full population. Particularly important is the fact that funds are unlikely to report the very bad returns that lead to their collapse.

[7] The Base results here do not match the prior tables since we lose some funds who do not have returns before and after 2004.

5 Estimation of Distributional Parameters

We now use our "full" alpha estimates from Panel B of Table 9 as the input to our EM algorithm and estimate the distributional parameters. As we do not know the actual number J of fund types in the data, we estimate the model using $J = 2, 3, 4, 5$ and measure model fit using the BIC. Using four or five distributions yields similar results so we focus on the case of $J = 4$ in the interest of parsimony.

The parameter estimates are reported in Table 10. For convenience we refer to the managerial types as "Bad" ($\hat{\mu} = -1.5\%$ per month), "Neutral" ($\mu = 0$ by assumption), "Good" ($\hat{\mu} = 0.24$), and "Excellent" ($\hat{\mu} = 1.3$). The unconditional probabilities π_j tell us the estimated fraction of the managers belonging to each type: 8% Bad, 42% Neutral, 39% Good, and 11% Excellent. The variability of alphas within each distribution lies between about 0.29 and 0.75. These results again are consistent with the claim that many hedge fund managers do have skill, though the caveats about selection biases remain.

To assess the statistical significance of the parameter estiamtes we employ a bootstrap analysis. We form 100 alternative datasets by drawing from our original sample with replacement. We estimate the parameters in each artificial sample, then calculate the standard deviation of the parameter estimates across the 100 samples. These standard deviations, reported in parentheses in the table, are a measure of the standard error of the parameter estimates. The parameters are estimated with reasonable precision. The important mixing probabilities have standard errors ranging from 1.1 to 2.8 percentage points.

Figure 7 gives a graphical depiction of the composite empirical density (solid line) along with the corresponding estimate (heavy dashed line) and the estimates from each of the underlying families (thin dashed). To the eye, the fitted composite distribution matches the empirical distribution fairly well.

We then examine the question about the posterior probability of a manager's type given an estimate of alpha. Using equation (1) for a range of estimated alphas, we plot these probabilities for each skill level separately in Figure 8. Each panel corresponds to the probability the manager is of a particular type: Bad (Panel A), Neutral (Panel B), Good (Panel C), or Excellent (Panel D). The circles in each plot correspond to the mean skill levels for the four groups, with the shaded dot indicating the skill group whose probability is plotted. For example, Panel B plots the probability the manager is Netural and the solid circle is $\mu_N = 0.0$. The shading in the plots represent the 95% confidence band from the bootstrap.

The conditional probabilities are combined together in Figure 9. The results are intuitive. Extremely large or small alphas come from the Bad or Excellent family with near certainty. Intermediate values are more likely from the Neutral or Good groups. Interestingly, a manager with alpha of 1% per month is more likely to be Neutral than a manager with alpha of 0.25% per month. This is because it is common for Good managers to be near 25 bp, but rare for them to be near 100 bp. It is also worth noting that the probability a manager is Neutral actually increases as observed alpha drops below zero (Panel B of Table 8). This occurs because there are a lot of Good managers with alpha near zero, but the Neutrals have less "competition" with the Goods when alpha is somewhat below zero because of the low σ_G.

To highlight the usefulness of our approach, consider a manager with an alpha of 0.56% per month. This level of alpha is halfway between Good and Excellent, in the sense that it is 1.1 standard deviations away from each mean. We estimate that there is a 57.4% chance this manager is merely Good and only a 6.6% chance they are Excellent. The reason for the low likelihood of being Excellent is that there simply aren't that many Excellent managers in the population. This might be a Excellent manager with bad luck, but more likely is that he is merely Good with good luck, or Neutral with really good luck. The latter case, which is 1.0 standard deviations above the mean, occurs with 35.7% probability due to the large fraction of the population that is Neutral (42%). For this example, our results suggest a cautious interpretation of what looks like good performance. Given the scarcity of truly outstanding managers, it is more likely that we are seeing a lucky manager.

6 Refinements and Extensions

- Funds with short histories tend to have extreme alphas, due in part to noise. See Figures 10 and 11.
- We have a truncated distribution. Can we infer the full population from our truncated sample? [See Cochrane (2005) for an application with venture capital.]
  - Poor performers exit the sample as they tend to fail. Very good performers may elect to stop reporting, especially if they are closed to new investors.
  - Model the live/dead funds using a probit model (see Table 11).
  - Use this probit model with the distribution of live alphas to infer alphas from the full population.

7 Conclusion

We present a new approach to evaluating the performance of an investment manager. Our approach focuses on the alternative hypothesis that some managers have good (or bad) skill.

Under this alternative, the composite distribution of fund alphas will be non-Normal. We estimate the distributional parameters by fitting the observed distribution of alphas.

Our approach offers two main advantages. First, it more realistically entertains the role that luck may play in the observed distribution of alphas. The conventional stance in the literature is to focus on the null hypothesis that all managers are neutral and to recognize that some of these managers will experience good or bad luck and wind up with a "significant" alpha. We allow luck to also affect managers with good or bad skill. A fund with an estimated alpha of zero may have no skill, or it may be a good manager with bad luck or a bad manager with good luck. Second, our approach yields estimates of the magnitudes and variability of alpha for the various skill-types. These parameter estimates can then be used in conjunction with the estimated alpha from an individual fund to make probabilistic statements about the likelihood that manager has positive, negative, or neutral skill.

Since alphas are the key input to our procedure, we offer several adjustments to the standard seven-factor model to estimate alpha. First, we explicitly control for the "instant history" bias by adding a dummy variable for the early months of the fund's life. Second, we add the HML and UMD factors since these are known to generate large average returns. Third, we include lagged values of the factors in order to address the nonsynchronous trading problem prevalent in hedge fund returns.

We take our procedure to a sample of hedge funds. We find that our refinements to the estimation of alpha reduce its magnitude by one-half to two-thirds depending on the subsample of funds. The fraction of funds with statistically significant alphas also shrinks by half. We find very few funds with significantly negative alphas.

We implement our skill-detection procedure on the subset of pure (non-FoF) funds and obtain a much different picture. We find much higher fraction of the sample with skill (positive or negative) than indicated by the frequency of significant t-statistics. We estimate that 8% of funds have negative skill, about four times the fraction of alphas that are significantly negative. 50% are deemed to have positive performance, more than double the fraction of alphas that are significantly positive. Overall, our findings confirm that it is important to allow for luck affecting all skill levels, rather than just neutral managers.

Appendix A Alternative Estimation Methods

A.1 Method of Moments

The method of moments dates back to Pearson (1894). The idea is quite simple: choose parameters so that model-implied moments are equal (or at least close) to sample moments. Model-implied moments are provided by the derivatives of the *moment generating function* (MGF).

For a univariate variable $x \sim \mathcal{N}(\mu, \sigma^2)$ the MGF is

$$M_x(\theta) = \exp\left(\theta\mu + \frac{1}{2}\theta^2\sigma^2\right)$$

Derivatives of the MGF evaluated at $\theta = 0$ give the moments of x. For example,

$$\begin{aligned} M_x'(\theta) &= M_x(\theta)(\mu + \theta\sigma^2) \\ M_x''(\theta) &= M_x(\theta)\sigma^2 + M_x'(\theta)(\mu + \theta\sigma^2)^2 \end{aligned}$$

Therefore, $M_x'(0) = \mu$ and $M_x''(0) = \sigma^2 + \mu^2$.

The method of moments estimate is obtained by choosing parameters to match moments,

$$E[x^n] = \partial^n M_x(\theta)/\partial\theta^n$$

In our example,

$$\begin{aligned} T^{-1}\sum_{t=1}^{T} x_t &= \mu \\ T^{-1}\sum_{t=1}^{T} x_t^2 &= \sigma^2 + \mu^2 \end{aligned}$$

which are the standard estimate of mean and variance for a Normal random variable.

To estimate the seven parameters for a mixture of three Normals, we need to use the first seven moments.

A.2 Moment Generating Function

Another approach to makes direct use of moment-generating functions. Instead of taking derivatives, just solve $E[\exp(\theta x)] = M_x(\theta)$. As an example, suppose $\mu = 0$ but we need to estimate σ. Then you can set $\theta = 1$ and choose σ by picking the value that balances

$\frac{1}{T}\sum_{t=1}^{T}\exp(x_t)$ with $\exp(\sigma^2/2)$. That is,

$$\hat{\sigma} = \sqrt{2\ln\left(\frac{1}{T}\sum_{t=1}^{T}\exp(x_t)\right)}$$

For a mixture of distributions, the MGF is the weighted average of the individual MGFs

$$M_\alpha(\theta) = \sum_{k=1}^{K}\pi_k M_{\alpha_k}(\theta)$$

If we assume $K = 3$ and $\mu_N = 0$ then there are seven remaining parameters (recall $\pi_N = 1 - \pi_B - \pi_G$). So we can pick seven values of θ and match the moments. θ need not be positive or integers, but the choice of the θs is consequential.

A.3 Characteristic Function

A more general version of the MGF method uses the characteristic function. Stated most simply, this replaces $i\theta$ with θ in $M_x(\theta)$. For a Normal variable, this means

$$E\left[\exp\left(i\theta x\right)\right] = \exp\left(i\theta\mu - \frac{1}{2}\theta^2\sigma^2\right).$$

Since $\exp(iz) = \cos(z) + i\sin(z)$ a single choice of θ gives two pieces of information to identify the parameters: the real part (cosine) and the imaginary part (sine). For example,

$$\begin{aligned} T^{-1}\sum_{t=1}^{T}\cos(\theta x_t) &= \cos(\theta\mu)\exp(-\theta^2\sigma^2/2) \\ T^{-1}\sum_{t=1}^{T}\sin(\theta x_t) &= \sin(\theta\mu)\exp(-\theta^2\sigma^2/2) \end{aligned}$$

Setting $\theta = 1$ and taking the ratio of these equations,

$$\frac{\sin(\mu)}{\cos(\mu)} = \frac{\sum_{t=1}^{T}\sin(x_t)}{\sum_{t=1}^{T}\cos(x_t)} = \tan(\mu)$$

and

$$\sigma = \sqrt{2\ln\left(\frac{\cos(\mu)}{T^{-1}\sum_{t=1}^{T}\cos(x_t)}\right)} = \sqrt{2\ln\left(\frac{\sin(\mu)}{T^{-1}\sum_{t=1}^{T}\sin(x_t)}\right)}$$

This gives rise to a GMM estimator using moments from the characteristic function. For k choices of θ there are $2k$ moments corresponding to the real and imaginary components.

A.4 Non-Parametric Kernel Density Estimate

The probability density function (pdf) of a random variable x characterizes the distribution of the variable. We are interested in understanding the properties of fund managers' skill, measured by α.

The kernel density estimate (KDE) estimates the density function as

$$\hat{f}(\alpha_i) = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{h}K\left(\frac{\alpha_i - \alpha_n}{h}\right).$$

The function K is known as the kernel; a standard Normal is a common choice and results are usually similar for alternative choices. What is the KDE doing? The density at data point α_i is a measure of the fraction of the sample that is "close" to α_i. The formula can be read as an average "probability" that a data point α_n is "near" α_i. "Near" is measured with respect to the bandwidth $h = 1.06\sigma\sqrt[5]{N}$. The bandwidth controls how much attention is paid to local observations versus the broader sample. When h is large the KDE is smooth, when it is small the plot is quite jagged.

With this in mind, it is instructive now to compare the KDE estimate to the estimated density from the EM procedure. That density is also a sum of Normal densities, but the key difference is that we impose much more parametric structure. The EM procedure yields estimates $\{\hat{\alpha}_j, \hat{\sigma}_j, \hat{\pi}_j\}$ and the density is

$$\hat{f}(\alpha_i) = \sum_{k=1}^{K}\hat{\pi}_k\phi(\alpha_i;\hat{\alpha}_k,\hat{\sigma}_k) = \sum_{k=1}^{K}\hat{\pi}_k\frac{1}{\sqrt{2\pi}\hat{\sigma}_k}\exp\left(-\frac{1}{2}\left(\frac{\alpha_i - \hat{\alpha}_k}{\hat{\sigma}_k}\right)^2\right)$$

The EM estimate builds the composite density from (say) $K = 3$ Normal distributions while the KDE estimate builds it from $N = \#$ data points densities. This structure allow us to make economic statements about the data that are less direct with the KDE estimate.

References

Ackermann, Carl, Richard McEnally, and David Ravenscraft, 1999, The performance of hedge funds: Risk, return, and incentives, *Journal of Finance* 54, 833–874.

Barras, Laurent, Olivier Scaillet, and Russell Wermers, 2007, False discoveries in mutual fund performance: Measuring luck in estimated alphas, University of Maryland Working Paper.

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

Brown, Stephen, Will Goetzmann, and Roger Ibbotson, 1999, Offshore hedge funds: Survival and performance 1989-95, *Journal of Business* 72, 91–117.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.

Cochrane, John H., 2005, The risk and return of venture capital, *Journal of Financial Economics* 75, 3–52.

Foster, Dean P., and H. Peyton Young, 2008, The hedge fund game: Incentives excess returns, and piggy-backing, Working Paper.

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

——— , 2004, Hedge fund benchmarks: A risk-based approach, *Financial Analysts Journal* 60, 65–80.

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, 529–609.

Kosowski, Robert, Narayan Naik, and Melvyn Teo, 2007, Do hedge funds deliver alpha? a bayesian and bootstrap analysis, *Journal of Financial Economics* 84, 229–264.

Liang, Bing, 1999, On the performance of hedge funds, *Financial Analysts Journal* 55, 72–85.

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

Pearson, Karl, 1894, Contributions to the mathematical theory of evolution, *Philosophical Transactions of the Royal Society, A* 185, 71–110.

Table 1: Comparison of Distribution Parameter Estimates

This table compares parameter estimates from several approaches: EM algorithm (EM), Method of Moments (MM), Moment Generating Function (MGF), and Characteristic Function (CF). The Truth column shows the true parameters used to simulate the data. The simulated sample consists of 2000 observations from the good (G), neutral (N), and bad (B) distributions. Each of these three distributions is assumed $\mathcal{N}(\mu_j, \sigma_j^2)$ with mixing probability π_j. $\mu_N = 0$ and all estimates impose this restriction.

	Truth	EM	MM	MGF	CF
μ_G	3.00	3.2445	3.1846	3.0913	3.1135
μ_B	−2.00	−1.9016	−1.8986	−1.8863	−1.8856
σ_G	1.00	0.8996	1.1045	1.0425	1.0736
σ_N	1.00	1.0571	1.3860	1.1713	1.0822
σ_B	1.00	1.0721	1.1992	0.9710	1.0807
π_G	0.05	0.0408	0.0193	0.0281	0.0403
π_N	0.70	0.7073	0.8383	0.8715	0.7085
π_B	0.25	0.2519	0.1424	0.1004	0.2513

Table 2: Comparison of Fund-Type Probabilities

This table shows the probability that a manager with a given $\hat{\alpha}$ is of type Good, Neutral, or Bad. These probabilities are calculated according to

$$\Pr(\text{Mgr } i \text{ is type } j|\hat{\alpha}_i) = \frac{\hat{\pi}_j \phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)}{\sum_j \hat{\pi}_j \phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)}. \quad (1)$$

The parameter estimates are from Table 1.

	Truth	EM	MM	MGF	CF
Panel A: $\hat{\alpha} = 3.00$					
Good	0.8654	0.7854	0.2290	0.4896	0.7264
Neutral	0.1346	0.2145	0.7706	0.5103	0.2735
Bad	0.0000	0.0001	0.0004	0.0000	0.0002
Panel B: $\hat{\alpha} = 1.50$					
Good	0.0665	0.0273	0.0159	0.0250	0.0459
Neutral	0.9312	0.9666	0.9779	0.9743	0.9476
Bad	0.0022	0.0061	0.0062	0.0007	0.0065
Panel C: $\hat{\alpha} = 0.00$					
Good	0.0008	0.0001	0.0004	0.0004	0.0008
Neutral	0.9532	0.9320	0.9465	0.9790	0.9273
Bad	0.0461	0.0679	0.0530	0.0206	0.0719
Panel D: $\hat{\alpha} = -1.00$					
Good	0.0000	0.0000	0.0000	0.0000	0.0000
Neutral	0.7368	0.7217	0.8387	0.8835	0.7199
Bad	0.2632	0.2783	0.1613	0.1165	0.2801
Panel E: $\hat{\alpha} = -2.00$					
Good	0.0000	0.0000	0.0000	0.0000	0.0000
Neutral	0.2748	0.3232	0.6435	0.6279	0.3392
Bad	0.7252	0.6768	0.3565	0.3721	0.6608

Table 3: Counts

This table summarizes the number of funds by type and indicates the number and frequency of funds that are still live at the end of our sample. Our sample consists of funds that report monthly returns net of fees on a US dollar basis during the period 1994-2008:12. We exclude funds with fewer than 24 monthly returns. The fund categories are: Long-Short Equity Hedge (LSEH), Fund-of-Funds (FOF), Event Driven (ED), Managed Futures (MF), Equity Market Neutral (EMN), Emerging Markets (EM), Multi-Strategy (MS), Global Macro (GM), Fixed Income Arbitrage (FIA), and Miscellaneous (which includes Convertible Arbitrage, Dedicated Short Bias, and Other).

	Total	Live	%Live
LSEH	1525	681	44.7
FOF	1050	639	60.9
ED	433	188	43.4
MF	423	167	39.5
EMN	293	135	46.1
EM	291	154	52.9
MS	284	208	73.2
GM	249	103	41.4
FIA	217	94	43.3
Misc	200	75	37.5
Non-FOF	3915	1805	46.1
Total	4965	2444	49.2

Table 4: Annual Attrition

This table shows the number of funds alive as of the start of the year (Live) and the number within that group that die by the end of the year (Dead). The tabulations are done separately for pure hedge funds (Non-FoF) and funds-of-funds (FoF). The final two columns show the fraction of those funds live at the start of the year that do not survive to the end of the year. The full sample contains 4,965 funds from TASS between 1994 and February 2008.

	Non-FoF		FoF		Mortality Rate (%)	
Year	Live	Dead	Live	Dead	Non-FoF	FoF
1994	616	0	157	0	0.00	0.00
1995	797	0	198	0	0.00	0.00
1996	1027	29	245	9	2.82	3.67
1997	1208	29	268	4	2.40	1.49
1998	1394	51	304	4	3.66	1.32
1999	1594	65	341	14	4.08	4.11
2000	1741	106	397	13	6.09	3.27
2001	1927	117	465	17	6.07	3.66
2002	2113	97	524	26	4.59	4.96
2003	2326	117	604	23	5.03	3.81
2004	2527	144	678	19	5.70	2.80
2005	2550	192	738	29	7.53	3.93
2006	2433	200	753	33	8.22	4.38
2007	2184	182	710	12	8.33	1.69

Table 5: Summary Statistics

This table presents summary statistics for our sample of 4,965 hedge funds during the period 1994 through 2/2008. "Avg Size" and "Size Now" are the average and most recent assets under management, in millions. "%Registered" is the percentage of funds that have voluntarily registered as investment advisers with the SEC. "%Lockup" is the fraction of funds that have a lockup period, and "Lockup" is the average lockup period (in months) for those that have a lockup. "Min Inv" is the minimum investment (in millions). "Notice Pd" is the notice period for withdrawals (in days). "%Own" and "Own Cap" indicate the fraction of fund managers investing in their own funds, as well as their dollar investment for those managers that do invest. "Ifee" and "Mfee" are the incentive and management fees, expressed as percentages. "%HWM" is the fraction of funds that have a high-water mark policy. "$Levered" is the fraction of funds that indicate they use leverage. The final four rows summarize the time series properties of each fund: T is the number of monthly returns available for a fund, $\bar{r}$ is the average monthly return, $\sigma(r)$ is the standard deviation of monthly returns, and $\rho(r)$ is the first-order autocorrelation.

	All			Live			Non-FoF		
	N	Mean	Median	N	Mean	Median	N	Mean	Median
Avg Size	4333	190.57	33.99	1955	329.91	61.75	3456	164.16	33.25
Size Now	4333	282.35	31.58	1955	504.26	86.40	3456	249.16	30.00
%Registered	4965	6.49	0.00	2444	7.53	0.00	3915	5.98	0.00
%Lockup	4965	24.41	0.00	2444	24.22	0.00	3915	25.11	0.00
Lockup	1212	11.92	12.00	592	12.03	12.00	983	11.73	12.00
Min Inv	4965	4.18	0.25	2444	7.75	0.25	3915	4.95	0.50
Notice Pd	4965	34.91	30.00	2444	39.15	30.00	3915	33.90	30.00
%Own	4965	31.72	0.00	2444	23.28	0.00	3915	32.98	0.00
Own Cap	612	4.37	0.00	387	4.13	0.00	483	3.30	0.00
Ifee	4952	16.29	20.00	2438	15.55	20.00	3904	17.04	20.00
Mfee	4952	1.46	1.50	2438	1.47	1.50	3904	1.45	1.50
%HWM	4959	61.95	100.00	2438	72.64	100.00	3911	62.03	100.00
%Levered	4965	58.73	100.00	2444	54.91	100.00	3915	60.23	100.00
T	4965	70.13	60.00	2444	74.15	64.00	3915	69.55	60.00
$\bar{r}$	4965	0.87	0.78	2444	0.99	0.86	3915	0.88	0.79
$\sigma(r)$	4965	3.25	2.40	2444	2.59	1.98	3915	3.44	2.55
$\rho(r)$	4965	0.14	0.14	2444	0.16	0.16	3915	0.14	0.14

Table 6: Instant History Analysis

This table presents compares the Fung and Hsieh (2004) seven-factor models (Base) to our expanded model that includes an incubation dummy variable (Alt). In Panel A the incubation dummy applies to the first twelve months of the fund's returns; in Panel B the incubation period is 20 months. The first column in each panel reports the average dummy variable coefficient estimate. The second and third columns report the average alpha under each model. The next two columns report the fraction of estimated alphas that are are significantly negative at the 5% significance level. The final pair or columns report the fraction of alphas that significantly positive. The rows of the table correspond to various subsets of the sample. The sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

Panel A: 12 Month Incubation

	CS Avg	CS Avg α		% of $t(\alpha) < -2$		% of $t(\alpha) > 2$	
	Inst12	Base	Alt	Base	Alt	Base	Alt
Live Non-FoF	0.33	0.49	0.43	0.72	0.55	39.72	32.85
Live FoF	0.27	0.52	0.46	0.94	0.47	39.75	33.18
Dead Non-FoF	0.66	0.33	0.16	3.32	3.98	25.78	19.24
Dead FoF	0.19	0.10	0.06	2.68	3.41	23.36	19.71
All	0.45	0.39	0.29	2.01	2.24	32.45	26.02

Panel B: 20 Month Incubation

	CS Avg	CS Avg α		% of $t(\alpha) < -2$		% of $t(\alpha) > 2$	
	Inst20	Base	Alt	Base	Alt	Base	Alt
Live Non-FoF	0.25	0.49	0.42	0.72	0.78	39.72	29.70
Live FoF	0.19	0.52	0.46	0.94	0.63	39.75	31.14
Dead Non-FoF	0.64	0.33	0.07	3.32	3.89	25.78	15.83
Dead FoF	0.19	0.10	0.02	2.68	4.87	23.36	19.22
All	0.40	0.39	0.24	2.01	2.42	32.45	23.12

Table 7: Alphas Controlling for HML and UMD

This table presents compares the Fung and Hsieh (2004) seven-factor models (Base) to our expanded model that includes the HML and UMD factors (Alt). The first two columns report the average HML and UMD coefficients. The third and fourth columns report the average alpha under each model. The next two columns report the fraction of estimated alphas that are are significantly negative at the 5% significance level. The final pair or columns report the fraction of alphas that significantly positive. The rows of the table correspond to various subsets of the sample. The sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

	CS Avg		CS Avg α		% of $t(\alpha) < -2$		% of $t(\alpha) > 2$	
	HML	UMD	Base	Alt	Base	Alt	Base	Alt
Live Non-FoF	0.12	0.08	0.49	0.40	0.72	0.66	39.72	32.02
Live FoF	0.13	0.08	0.52	0.42	0.94	1.56	39.75	31.61
Dead Non-FoF	0.06	0.01	0.33	0.30	3.32	3.22	25.78	22.70
Dead FoF	0.07	0.03	0.10	0.04	2.68	3.16	23.36	20.19
All	0.09	0.05	0.39	0.33	2.01	2.07	32.45	27.03

Table 8: Alphas Controlling for Change in HML and UMD

This table presents compares the Fung and Hsieh (2004) seven-factor models (Base) to our expanded model that includes dynamic exposures to the HML and UMD factors (Alt). The first two columns report the average HML and ΔHML coefficients. The latter represents the incremental HML loading since 2004. The third and fourth columns report similar average coefficients for the UMD factor. Columns five and six report the average alpha under each model. The next two columns report the fraction of estimated alphas that are are significantly negative at the 5% significance level. The final pair or columns report the fraction of alphas that significantly positive. The rows of the table correspond to various subsets of the sample. The sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

	CS Avg				CS Avg α		% of $t(\alpha) < -2$		% of $t(\alpha) > 2$	
	HML	Δ HML	UMD	Δ UMD	Base	Alt	Base	Alt	Base	Alt
Live Non-FoF	0.12	−0.10	0.02	0.14	0.51	0.44	0.73	0.81	49.96	40.44
Live FoF	0.15	−0.11	0.01	0.19	0.53	0.45	1.08	1.51	48.39	39.14
Dead Non-FoF	0.09	−0.10	0.02	0.06	0.46	0.40	1.23	1.50	39.53	34.06
Dead FoF	0.07	−0.03	0.03	0.06	0.19	0.14	0.64	1.27	32.48	26.75
All	0.11	−0.10	0.02	0.12	0.48	0.41	0.93	1.16	45.68	37.58

Table 9: Alphas from Full Specification

This table presents compares the Fung and Hsieh (2004) seven-factor models (Base) to our "full" model (Alt). In Panel A, alpha is the intercept from the regression of fund excess returns on the seven-factors, the 20-month incubation dummy, HML, UMD, and two lags of the market return. In Panel B, the regression uses all these variables plus lags of any other factors with significant slope coefficients. The first column in each panel reports the average t-statistic for alpha. The second and third columns report the average alpha under each model. The next two columns report the fraction of estimated alphas that are are significantly negative at the 5% significance level. The final pair or columns report the fraction of alphas that significantly positive. The rows of the table correspond to various subsets of the sample. The sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

Panel A: Lagged Market Only

	CS Avg $t(\alpha)$	CS Avg α		% of $t(\alpha) < -2$		% of $t(\alpha) > 2$	
		Base	Alt	Base	Alt	Base	Alt
Live Non-FoF	1.02	0.49	0.24	0.72	1.22	39.72	21.11
Live FoF	1.15	0.52	0.28	0.94	1.88	39.75	22.22
Dead Non-FoF	0.34	0.33	0.03	3.32	4.27	25.78	12.89
Dead FoF	0.30	0.10	−0.09	2.68	4.87	23.36	12.90
All	0.69	0.39	0.13	2.01	2.90	32.45	17.08

Panel B: Lagged Market and Other Significant Factors

	CS Avg $t(\alpha)$	CS Avg α		% of $t(\alpha) < -2$		% of $t(\alpha) > 2$	
		Base	Alt	Base	Alt	Base	Alt
Live Non-FoF	0.97	0.49	0.24	0.72	1.16	39.72	19.61
Live FoF	1.10	0.52	0.27	0.94	2.03	39.75	21.44
Dead Non-FoF	0.33	0.33	0.03	3.32	3.79	25.78	12.32
Dead FoF	0.24	0.10	−0.10	2.68	5.60	23.36	11.68
All	0.66	0.39	0.12	2.01	2.76	32.45	16.09

Table 10: Fitted Parameters of Alpha Distribution

This table presents the estimates of the distributional parameters for the Non-FoF subsample. This subsample consists of 3,915 funds for which we have at least 24 monthly returns between 1994 and 2/2008. We model the distribution of estimated alphas as a mixture of four Normal distributions which we label "Bad", "Neutral", "Good", and "Excellent". The parameters μ_j and σ_j are the mean and variability of alpha for each family and π_j is the mixing probability. We estimate the parameters using the EM algorithm. Standard errors, shown in parentheses, are estimated as the standard deviation of parameter estiamtes across 100 bootstrapped samples.

	$\hat{\mu}_j$	$\hat{\sigma}_j$	$\hat{\pi}_j$
Bad	-1.4685	0.7505	0.0809
	(0.1061)	(0.0377)	(0.0105)
Neutral	0.0000	0.5699	0.4194
		(0.0392)	(0.0276)
Good	0.2394	0.2899	0.3904
	(0.0264)	(0.0280)	(0.0174)
Excellent	1.3382	0.7025	0.1093
	(0.1072)	(0.0413)	(0.0161)

Table 11: Cross-Sectional Determinants Fund Mortality

This table reports results from a logit regression where the dependent variable is one for live funds and zero for dead funds. The dependent variables are: α is the alpha from our full regression (see Table 9); RegDum is a dummy for registered funds; LockDum is a dummy for funds with lockup periods; MinInv is the minimum investment amount (in millions); NoticePd is the notice period for withdrawals (in days); OwnDum is a dummy for funds whose managers are investors; Ifee and Mfee are the percentage incentive and management fees; HWMDum is a dummy for funds with highwater marks; LevDum is a dummy for funds that use leverage; σ is the volatility of monthly fund returns, and ρ is the autocorrelation of monthly returns. The regressions are run separately for pure hedge funds (Non-FoF) and funds-of-funds (FoF). Funds in the sample have at least 24 monthly returns from 1994 to 2/2008.

	Non-FoF			FoF		
	Coef	se	t-stat	Coef	se	t-stat
Int	0.9919	0.1529	6.49	−1.2544	0.3024	−4.15
α	0.4909	0.0582	8.43	1.2721	0.1543	8.24
α^2	−0.2314	0.0407	−5.68	0.0885	0.1239	0.71
RegDum	−0.1101	0.1499	−0.73	0.7695	0.3009	2.56
LockDum	−0.4143	0.0894	−4.64	−0.2204	0.2002	−1.10
MinInv	0.0268	0.0125	2.14	0.0202	0.0212	0.95
NoticePd	0.0090	0.0015	6.15	0.0021	0.0028	0.75
OwnDum	−0.3540	0.0817	−4.33	−0.9225	0.1773	−5.20
Ifee	−0.1002	0.0066	−15.20	0.1082	0.0116	9.36
Mfee	0.2949	0.0547	5.39	−0.3107	0.1171	−2.65
HWMDum	0.9637	0.0895	10.76	0.8993	0.1721	5.22
LevDum	−0.0392	0.0771	−0.51	0.0949	0.1571	0.60
σ	−0.1767	0.0178	−9.96	0.1371	0.0501	2.74
ρ	0.3085	0.2014	1.53	−0.9098	0.4740	−1.92
R^2		0.2332			0.3152	

Figure 1: Average Hedge Fund Returns

This figure plots cross-sectional averages of four time series properties of hedge fund returns. Within each subplot, the cross-sectional averages are done for subsamples formed based on whether the fund is a fund-of-fund or dead. The sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

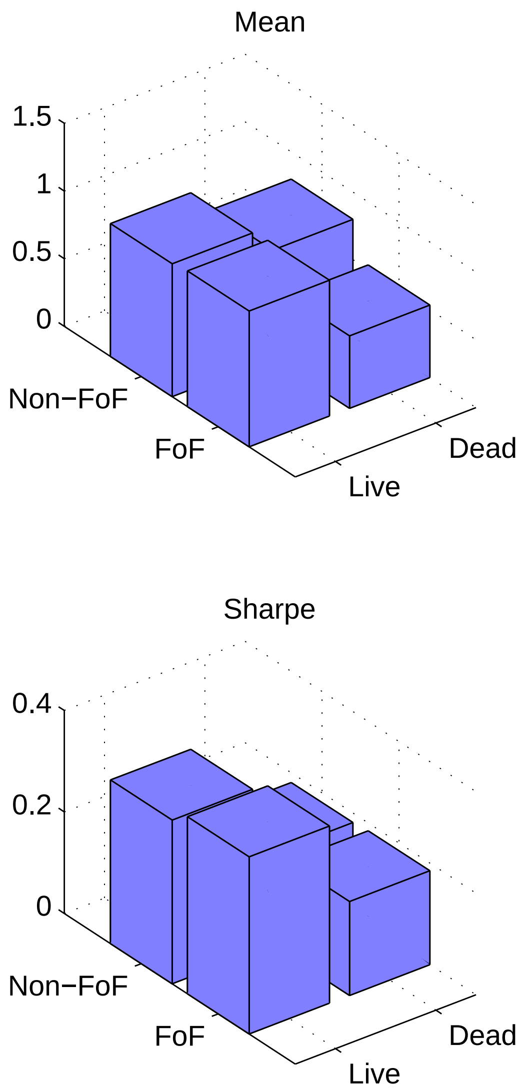

Figure 2: Average Hedge Fund Alphas

This figure plots cross-sectional averages of four time series properties of hedge fund alphas. The alphas come from the Fung and Hsieh (2004) model. Within each subplot, the cross-sectional averages are done for subsamples formed based on whether the fund is a fund-of-fund or dead. The sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

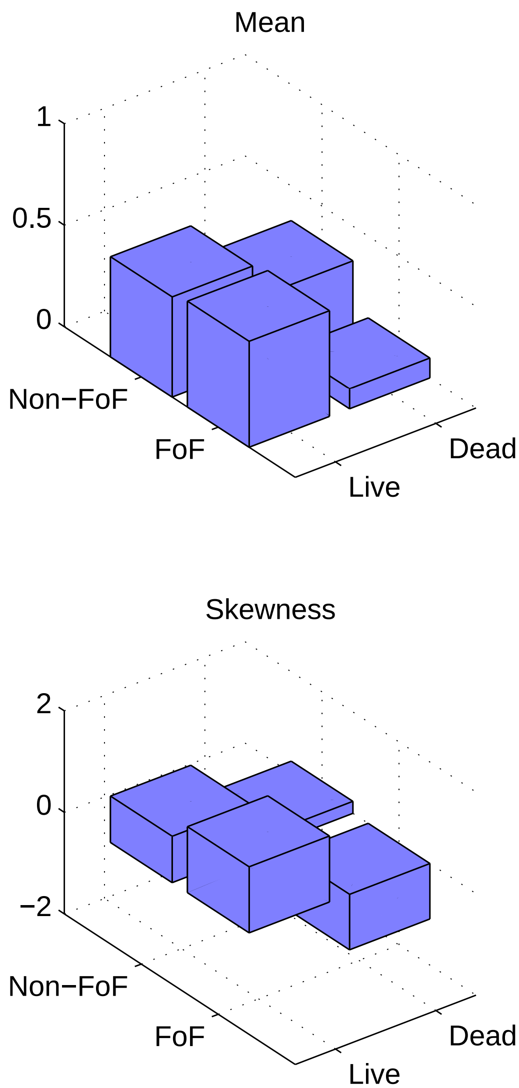

Figure 3: Distribution of Seven-Factor Alphas

This figure plots the distribution (KDE) of hedge fund alphas. The alphas come from the Fung and Hsieh (2004) model. Each subplot is based on a subsample formed based on whether the fund is a fund-of-fund or dead. The full sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

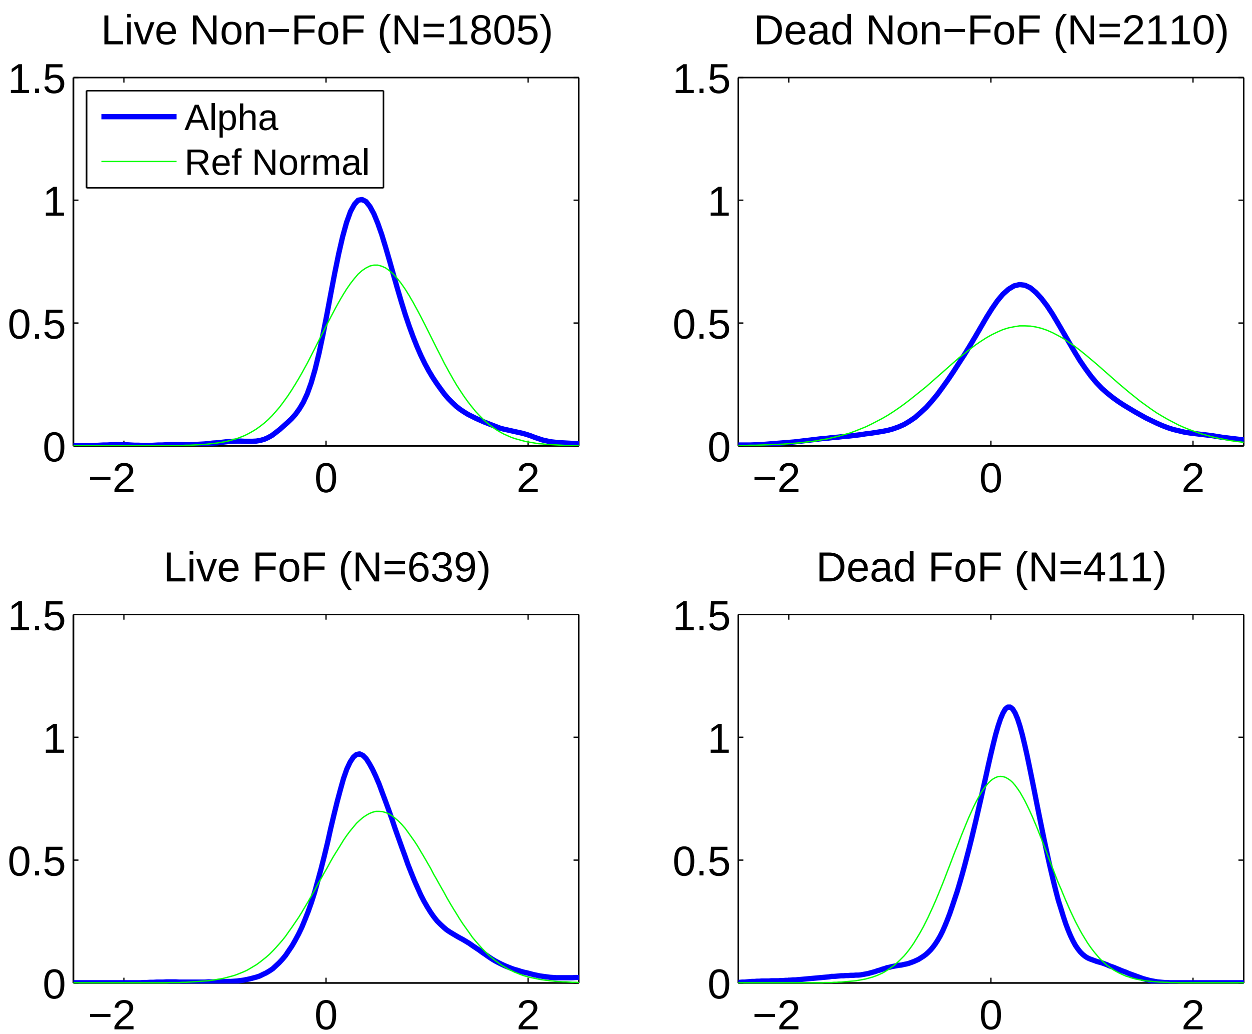

Figure 4: Seven-Factor Alphas by Fund Type

Panel A presents boxplots showing the distribution of alphas among live funds by fund type. Alphas come from the Fung and Hsieh (2004) model. The fund categories are: Event Driven (ED), Emerging Markets (EM), Equity Market Neutral (EMN), Fixed Income Arbitrage (FIA), Fund-of-Funds (FOF), Global Macro (GM), Long-Short Equity Hedge (LSEH), Managed Futures (MF), Multi-Strategy (MS), and Miscellaneous (which includes Convertible Arbitrage, Dedicated Short Bias, and Other). Panel B shows the fractions of each fund group that are still alive. The full sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

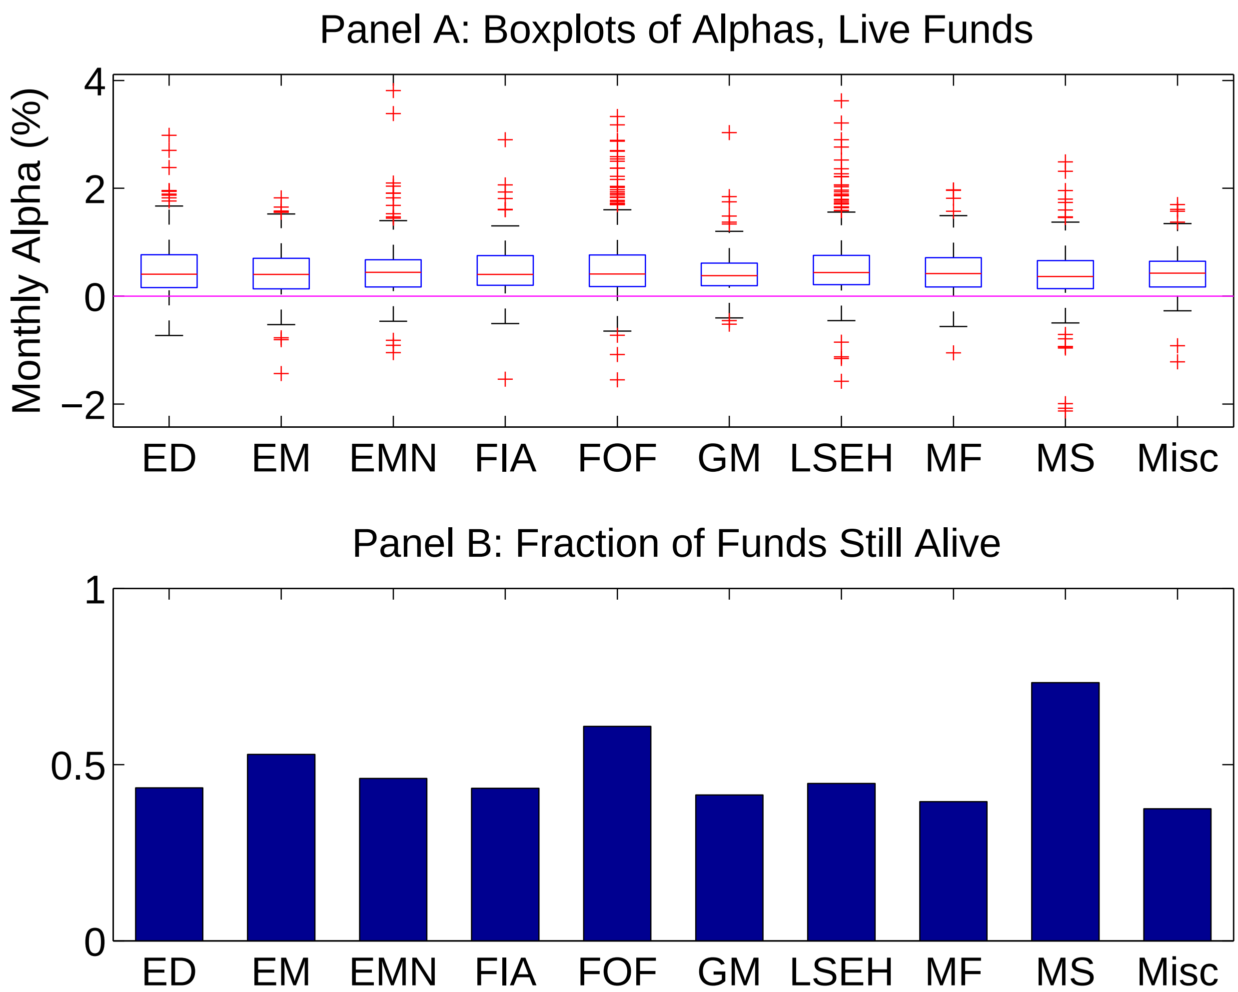

Figure 5: Change in UMD Loadings

This plot shows the average slope coefficients on UMD and Δ UMD from our model for alpha in Table 8. The UMD coefficient is the factor loading up until 2004. The factor loading after 2004 is the sum of UMD and ΔUMD, depicted by the total bar height. The fund categories are: Event Driven (ED), Emerging Markets (EM), Equity Market Neutral (EMN), Fixed Income Arbitrage (FIA), Fund-of-Funds (FOF), Global Macro (GM), Long-Short Equity Hedge (LSEH), Managed Futures (MF), Multi-Strategy (MS), and Miscellaneous (which includes Convertible Arbitrage, Dedicated Short Bias, and Other). The full sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

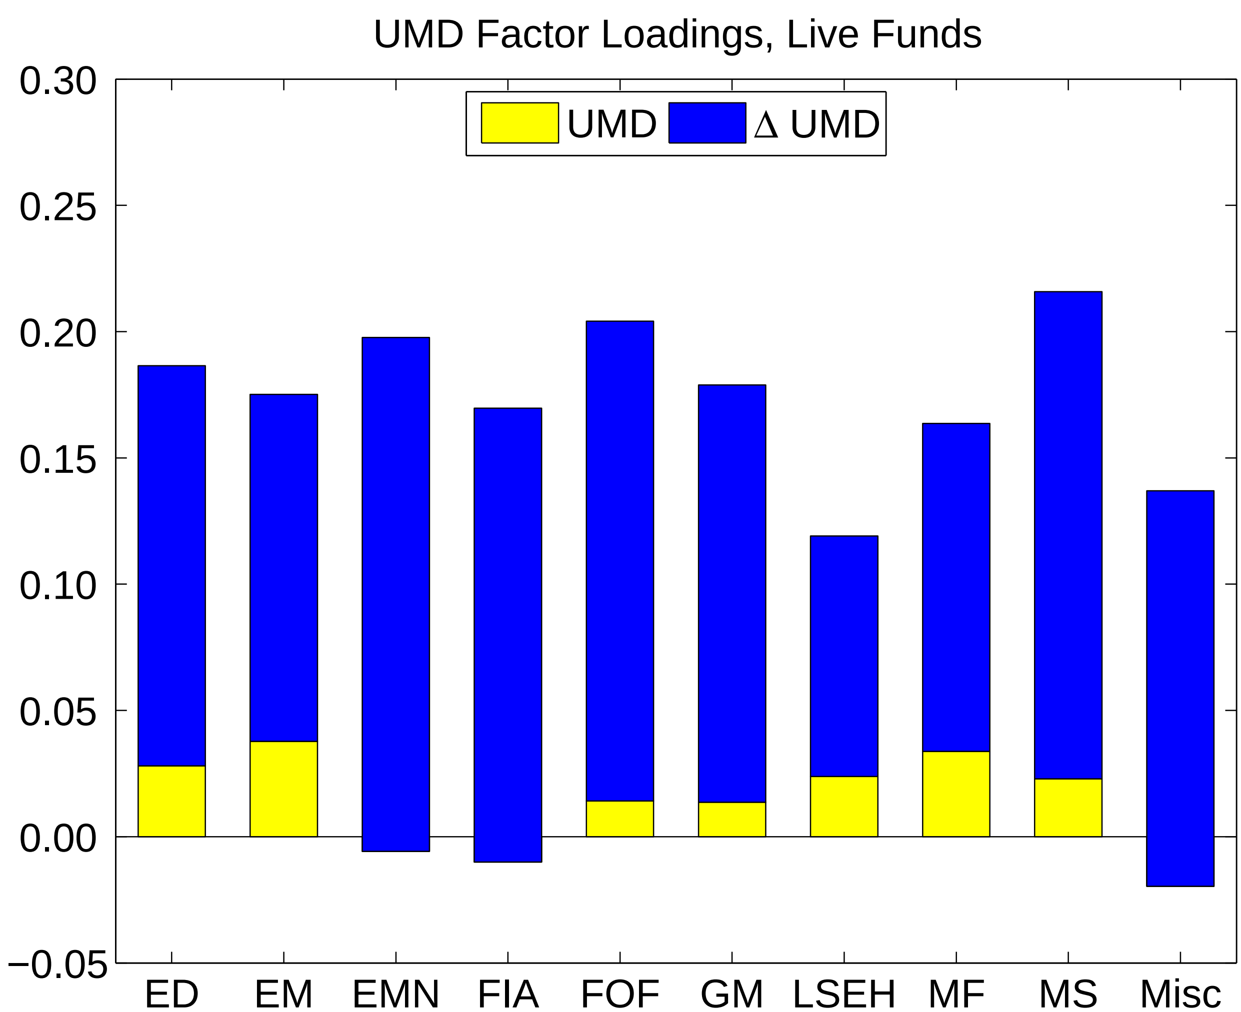

Figure 6: Distribution of Alphas from Full Specification

This figure plots the distribution (KDE) of hedge fund alphas. The heavy line uses the alphas from our "full" model (see Table 9). For comparison, we show alphas from the Fung and Hsieh (2004) as a thin line. Each subplot is based on a subsample formed based on whether the fund is a fund-of-fund or dead. The full sample consists of 4,965 funds with at least 24 monthly returns between 1994 and 2/2008.

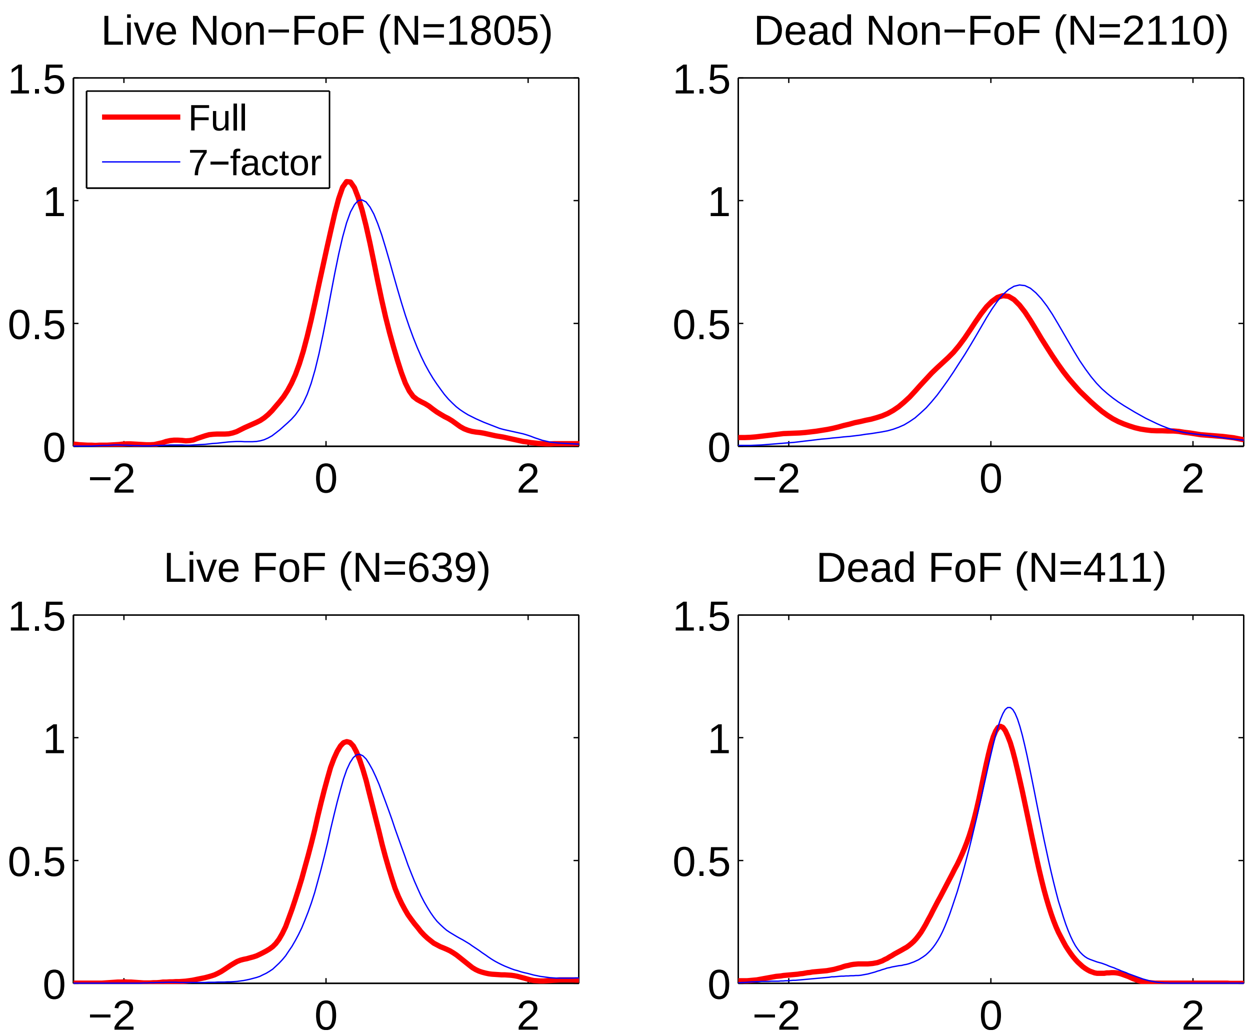

Figure 7: Empirical Estimate of Distributions of Alpha

This figure shows the estimated composite density (heavy dashed line) of alphas among Non-FoFs. Using the parameter estimates in Table 10, we construct the composite density as a mixture of the four underlying Normal densities (thin dashed lines). We plot the empirical KDE for comparison as the thin solid line. The Non-FoF subsample consists of 3,915 funds that have at least 24 monthly returns between 1994 and 2/2008.

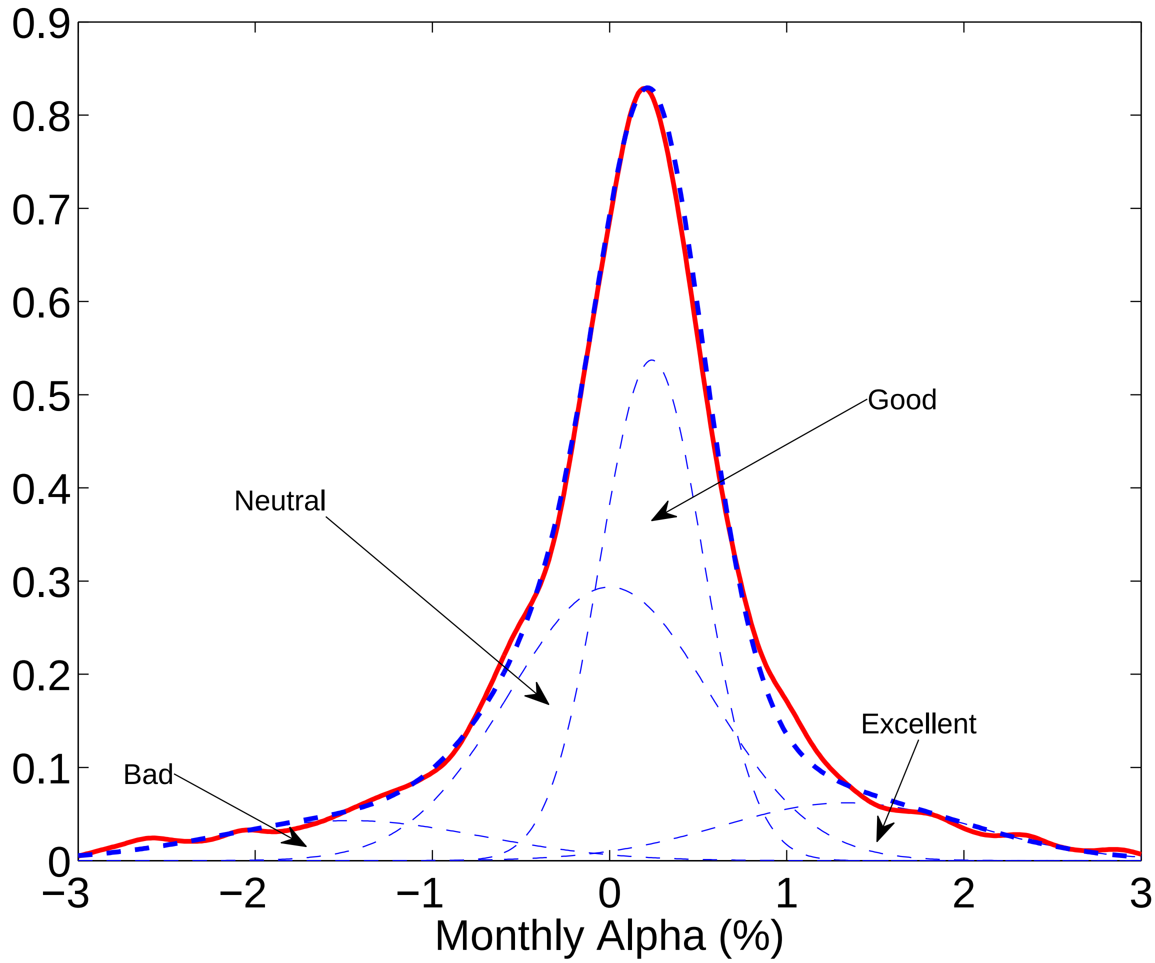

Figure 8: Probability of Fund Type

This figure plots the probability a manager is of the type specified in each panel, conditional on their alpha (expressed in annual percent). Parameter estimates come from Table 10.

$$\Pr(\text{Mgr } i \text{ is type } j|\hat{\alpha}_i) = \frac{\hat{\pi}_j \phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)}{\sum_j \hat{\pi}_j \phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)}. \quad (1)$$

The circles indicate the probabilities evaluated at the average alpha μ_j of each skill level. The circles are solid when the skill level being evaluated (the x coordinate) matches the skill level of the subplot. Shading indicates the 95% confidence band based on 100 bootstrapped samples. These plots are based on estimates from the subsample 3,915 Non-FoFs with at least 24 monthly returns between 1994 and 2/2008.

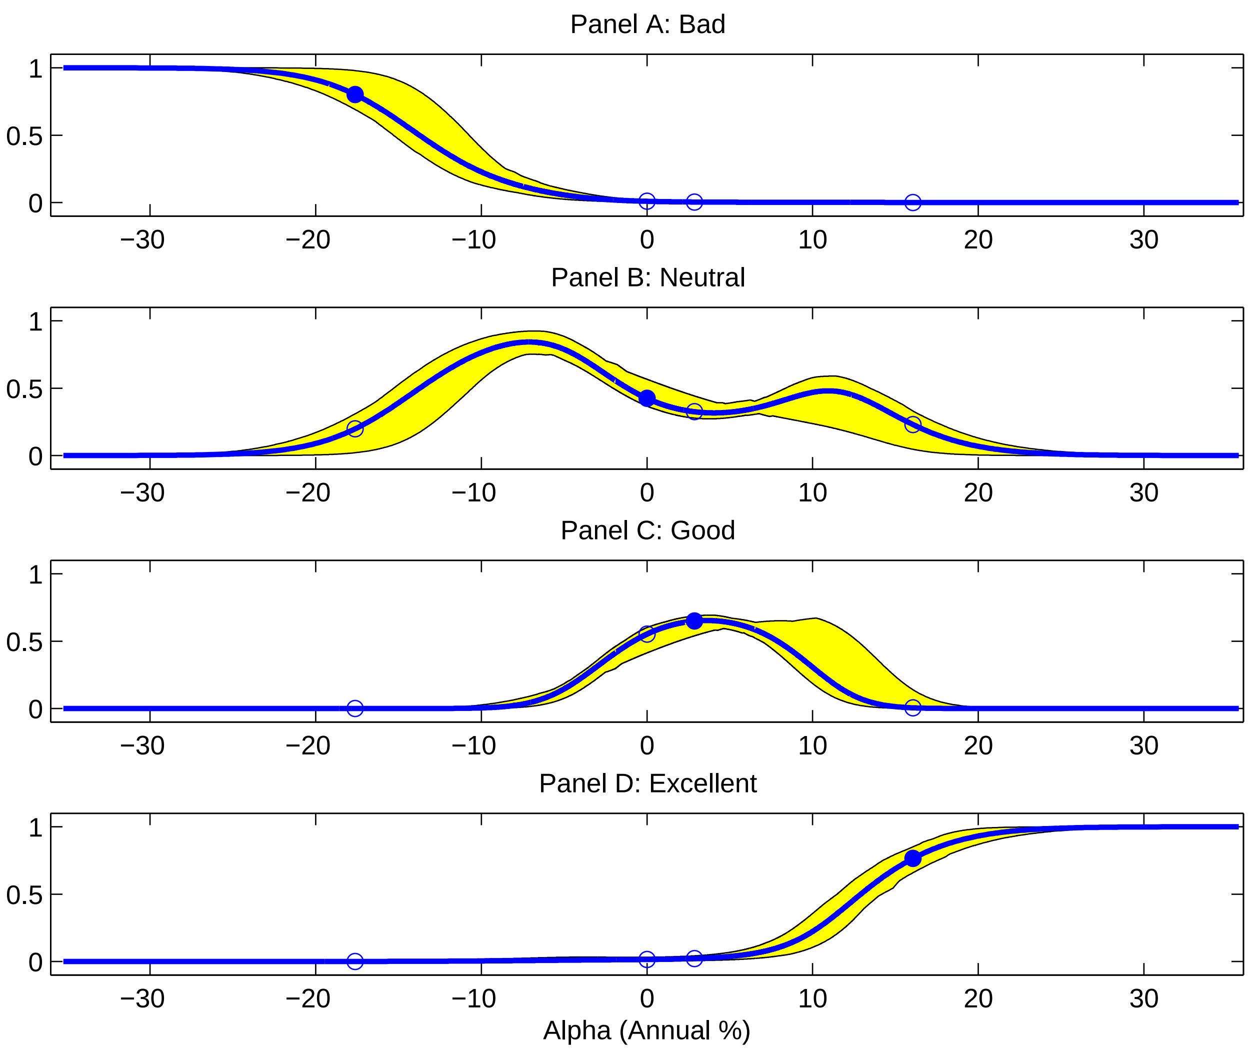

Figure 9: Area Chart of Fund-Type Probabilities

This figure plots the probability a manager is of the type specified in each panel, conditional on their alpha (expressed in annual percent). Parameter estimates come from Table 10.

$$\Pr(\text{Mgr } i \text{ is type } j|\hat{\alpha}_i) = \frac{\hat{\pi}_j \phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)}{\sum_j \hat{\pi}_j \phi(\hat{\alpha}_i; \hat{\mu}_j, \hat{\sigma}_j)}. \tag{1}$$

The figure aggregates the probabilities of each of the four skill-levels: "Bad" (red), "Neutral" (blue), "Good" (green), and "Excellent" (gold). This plots are based on estimates from the subsample 3,915 Non-FoFs with at least 24 monthly returns between 1994 and 2/2008.

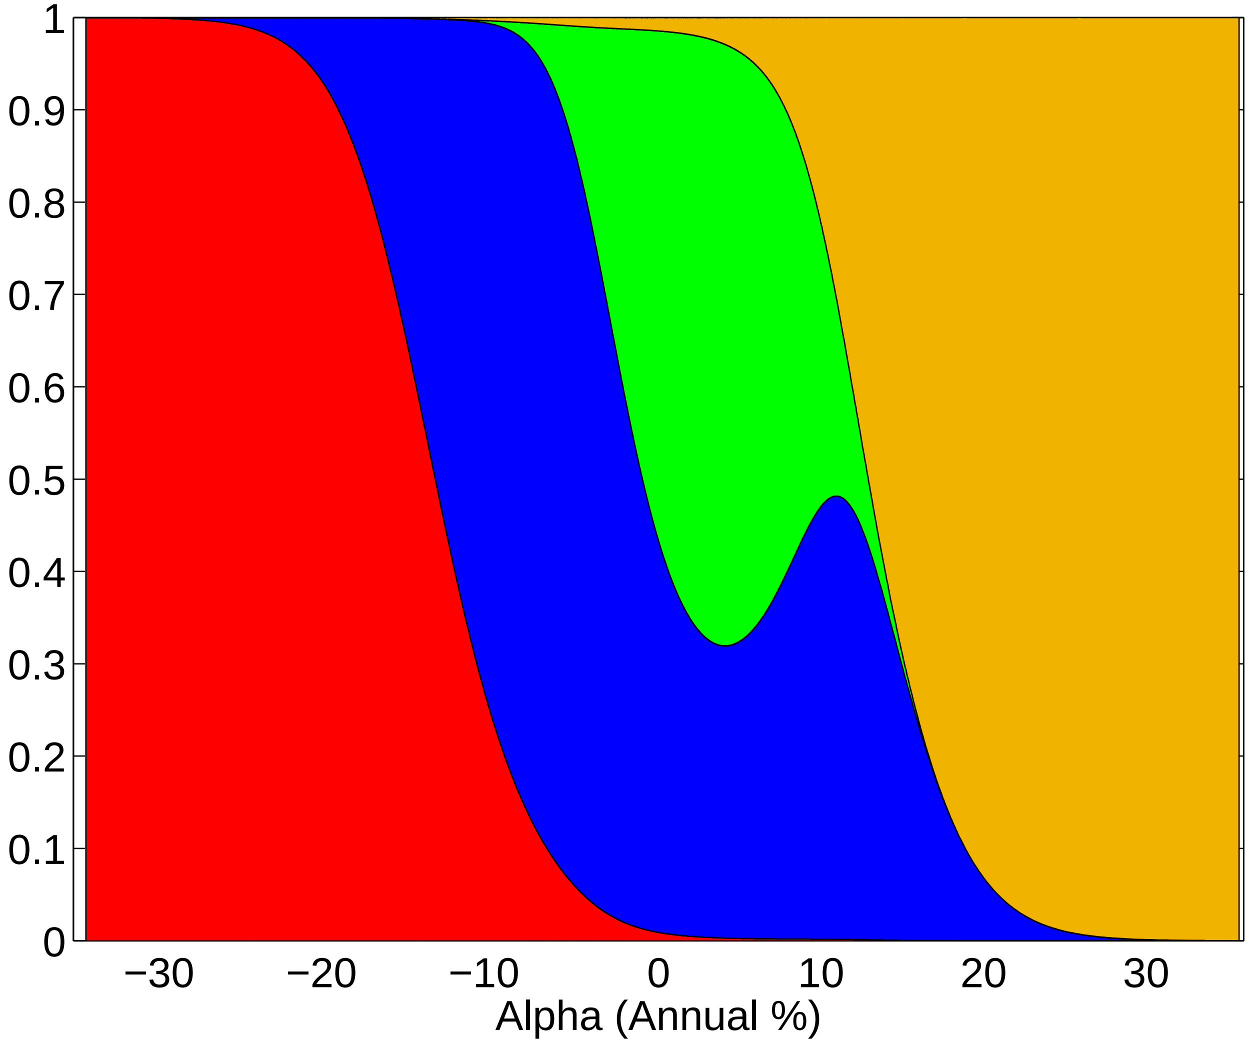

Figure 10: Alpha and Return History

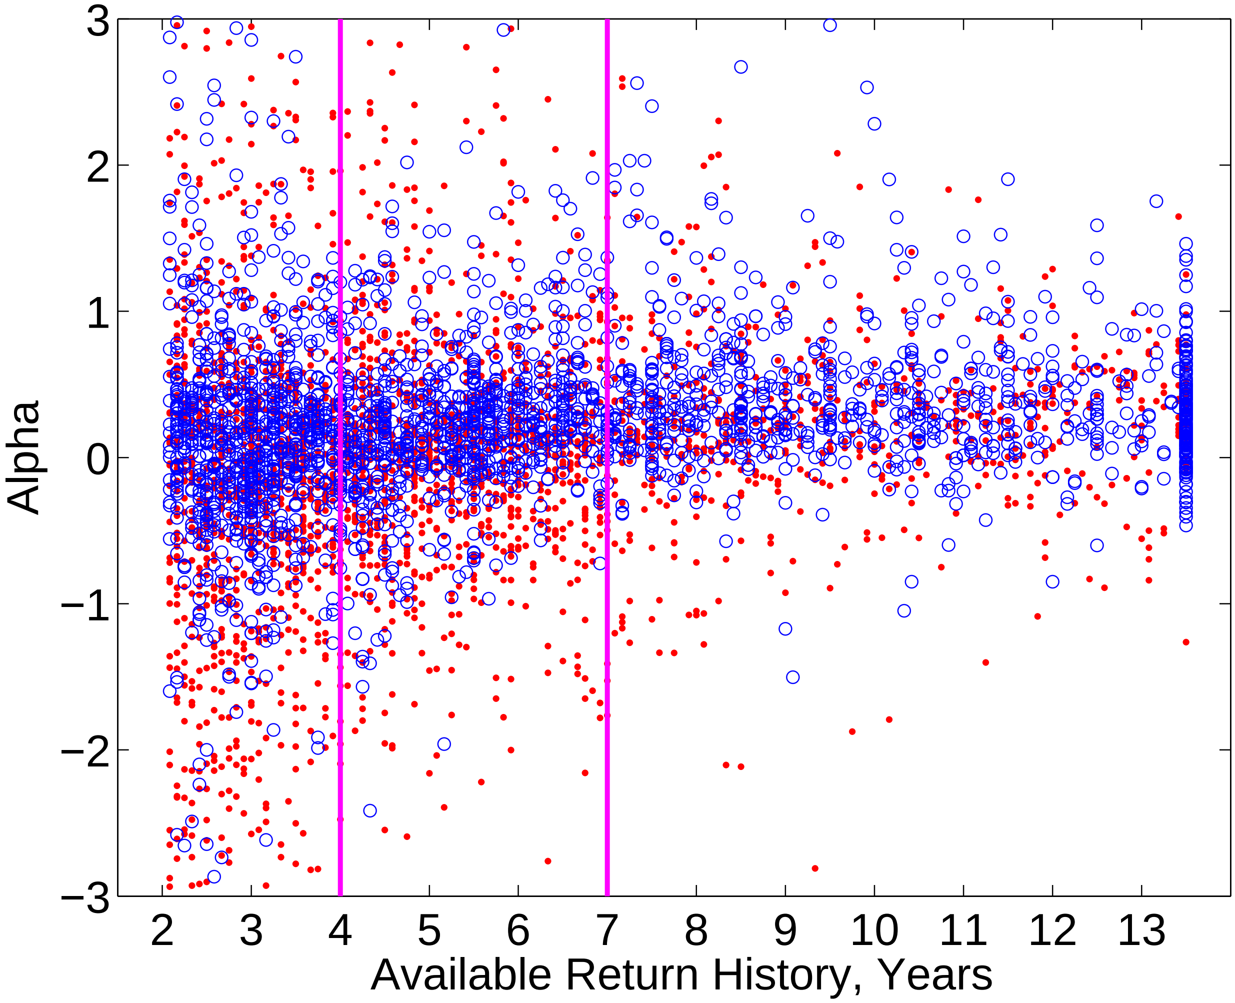

Figure 11: Distribution of Alpha and Return History