Reverse Survivorship Bias

JUHANI T. LINNAINMAA[*]

February 2010

ABSTRACT

Mutual funds often disappear following poor performance. When this poor performance is partly attributable to negative idiosyncratic shocks, the fund's estimated alpha understates its true alpha. This paper develops and estimates a structural model to evaluate this bias. I find that the bias in the mean of the observed alpha distribution is approximately 1 percent per year. When I correct for this bias using historical data, I find that the majority of fund managers still have negative net alphas but the average is not nearly as low as what the fund-level estimates suggest. This reverse survivorship bias affects all studies that run fund-level regressions to draw inferences about fund managers' abilities.

JEL classifications: G11, G12, C11

Keywords: Performance evaluation, mutual funds, reverse survivorship bias

---

The typical survivorship-bias argument starts from the observation that mutual funds often disappear following poor performance. Thus a study that conditions on fund survival overstates mutual fund performance.[1] In this paper I show that the correlation between poor performance and fund disappearance induces another pattern with the opposite sign: mutual fund alphas estimated from a survivorship-bias-free data set are biased downwards relative to the true distribution of alphas.

The mechanics of this bias are transparent in a setting in which investors learn about fund alphas. (As I discuss below, the bias arises even in the absence of learning. The only requirement is that funds tend to disappear following poor performance.) Suppose, for the sake of an argument, first, that investors have a prior belief that a fund's alpha is zero; second, that the true alpha is fixed; and third, that investors abandon a mutual fund (and the fund shuts down) when their posterior belief is that the fund's alpha is less than $-\bar{\alpha}$. Each month, investors update their beliefs about the alpha based on the risk-adjusted return. If this return is positive, investors infer skill. If the return is negative, investors infer less skill. The posterior mean always lies between the prior mean and the average realized risk-adjusted return. As a consequence, if the posterior mean ever falls below $-\bar{\alpha}$ and the fund disappears then the *observed* alpha must have been strictly lower than $-\bar{\alpha}$. If not, the posterior mean, which is an unbiased estimate of the fund's true alpha, could not have crossed this threshold. I call the resulting positive gap between the true alpha and the alpha estimated from the data the *reverse survivorship bias*.[2]

The reverse survivorship bias arises from the correlation between risk-adjusted returns and fund survival probabilities. It does not matter whether this correlation is the product of an underlying learning process about fund alphas. To appreciate this bias as a statistical result, suppose a mutual fund has a fixed alpha α and that its risk-adjusted returns are $\tilde{R}_t^e \equiv \alpha + \tilde{\varepsilon}_t$, where $\tilde{\varepsilon}_t$'s have zero means and are independently distributed. A fund survives for a random $\tilde{T}$ number of months. The

---

[1]See, for example, Brown, Goetzmann, Ibbotson, and Ross (1992), Elton, Gruber, and Blake (1996), Carpenter and Lynch (1999), and Carhart, Carpenter, Lynch, and Musto (2002).

[2]Pástor, Taylor, and Veronesi (2009) present a closely related argument in the IPO literature. They note that if a firm has an IPO after the posterior mean about the firm's profitability exceeds some threshold then the observed pre-IPO profitability must have been strictly higher than this threshold. As a consequence, the market rationally expects the firm to experience a post-IPO drop in profitability equal in size to the gap between the average pre-IPO profitability and the posterior mean at the time of the IPO.

optional stopping-time theorem[3] states that

$$E\left[\sum_{t=1}^{\tilde{T}} \tilde{\varepsilon}_t\right] = 0. \tag{1}$$

It follows immediately from expression (1) that the expected *average* risk-adjusted return is

$$E\left[\frac{1}{\tilde{T}}\sum_{t=1}^{\tilde{T}} \tilde{R}_t^e\right] = \alpha + \text{cov}\left(\frac{1}{\tilde{T}}, \sum_{t=1}^{\tilde{T}} \tilde{\varepsilon}_t\right). \tag{2}$$

This expression shows that if the survival probability is increasing in the risk-adjusted return, that is, $\text{corr}(\tilde{T}, \tilde{\varepsilon}_t) > 0$, the average risk-adjusted return is a downwards-biased measure of the fund's true alpha, $E\left[\frac{1}{\tilde{T}}\sum_{t=1}^{\tilde{T}} \tilde{R}_t^e\right] < \alpha$.

Intuitively, this result arises from the ambiguity about whether a realized return is low because the true alpha is low or because the fund-specific shock is low. (This same noise leaves a Bayesian investor's posterior mean between the prior mean and the signal.) A fund manager may have a small positive alpha, but a negative idiosyncratic shock results in a low return. If such a fund dies, it leaves behind an in-sample (i.e., frequentist) alpha estimate that is too low. No mechanism eliminates just-lucky mutual funds to offset this bias. I note that one cannot resolve this bias by studying mutual funds in isolation of each other. One can only address the covariance term in expression (2), which represents the bias, by studying the cross section of mutual funds.

I note that this bias is related to, but not an example of, the "baby-boy fallacy": if parents stop having children after their first son, there will not be more boys than girls in the population. The difference between this fallacy and the reverse survivorship bias is the same as the difference between the sums and averages in expressions (1) and (2): do we count the total number of boys and girls in the population, or do we compute the average fraction of boys in each family? If parents were to follow a stop-at-a-boy rule, the boy and girl counts would be the same in the population but the fraction of girls in the average family would not equal the unconditional probability, $\frac{1}{2}$. If each family always stops having children after the first boy, the fraction of girls will be zero in $\frac{1}{2}$ of families, $\frac{1}{2}$ in $\frac{1}{4}$ of families (girl-boy), $\frac{1}{3}$ in $\frac{1}{8}$ of families (girl-girl-boy), and so forth. Continuing *ad infinitum,*

---

[3]See, for example, Williams (1991).

the expected fraction of girls in a family is then $\lim_{n \to \infty} \left\{ \left(\frac{1}{2} \right)^1 \frac{0}{1} + \left(\frac{1}{2} \right)^2 \frac{1}{2} + \cdots + \left(\frac{1}{2} \right)^n \frac{n-1}{n} \right\} = \sum_{n=1}^{\infty} \left(\frac{1}{2} \right)^n \frac{n-1}{n} = 0.3069$. Similarly, if mutual funds are more likely to disappear following "bad" outcomes than "good" outcomes, the observed fund data oversample "bad" outcomes.

This paper investigates the reverse survivorship bias by estimating a structural learning model from the data. The model assumes each mutual fund's alpha is fixed and drawn from a distribution that is known to investors. Each fund generates monthly risk-adjusted returns that exhibit two sources of uncertainty: idiosyncratic shocks to portfolio returns and the estimation uncertainty from the asset pricing model. Each period, investors use Bayes' rule to combine the prior distribution with the monthly risk-adjusted return to arrive at a posterior distribution for a fund's alpha. I remain agnostic about the mechanism that causes funds to disappear from the data and model the fund disappearance probability as a free function of the posterior distribution. This function nests, as a special case, the possibility that a fund disappears when the posterior mean falls below some critical threshold. However, the structural model also can undo the built-in learning process if the data suggest such a reversal is warranted.

This paper estimates the shape of the alpha distribution, the total variance of risk-adjusted returns, and the exit-probability function from the CRSP mutual fund data by using the Simulated Method of Moments. I match, between the data and simulations, the average alphas of both surviving and disappearing funds, the differences in these averages over fund age, mutual fund survival rates, and the mean and variance of the observed distribution of alphas. The model matches these salient characteristics of the mutual fund data. The estimated form of the exit-probability function supports the learning mechanism: a mutual fund typically disappears when the market's belief about its (CAPM) alpha decreases from the prior mean of .05 percent per year to a posterior mean of $-.21$ percent per year.

The structural model estimates indicate the reverse survivorship bias is economically important. Whereas the average true (CAPM) alpha at the estimated parameter values is .05 percent per year, the average alpha estimated from the data is just $-.71$ percent per year. The difference in these figures, 76 basis points, is the estimate of the magnitude of the reverse survivorship bias. When the structural model is estimated by using the three- and four-factor model alphas, the magnitude of the bias is 85 basis points and 83 basis points per year, respectively. These estimates of the size of

the reverse survivorship bias, about 1 percent per year, are similar in magnitude to the estimates of the direct survivorship bias that plagued early mutual fund databases.[4] These computations suggest that if we take a database that omits all dead funds and then start adding them back in, the (positive) direct survivorship bias decreases, but at the same time, the (negative) reverse survivorship bias begins to drag fund-specific alpha estimates down. When all dead funds have been added back in, as is done in the survivorship-bias-free databases, the mean alpha estimate is too low by approximately 1 percent per year.

The reverse survivorship bias affects the measurement of fund managers' true alphas but does not bias the estimates of the returns available to mutual fund *investors*. For example, the average return on a strategy that invests the same amount into each actively managed fund is unaffected because the profits of such a strategy do not depend on the counterfactual of how well a mutual fund would have performed had it not disappeared. The reverse survivorship bias is, however, of crucial importance if we are to draw inferences about fund managers' stock-picking abilities: How many fund managers have positive alphas? Or what is the average fund manager's alpha? The answers to these questions are important for understanding whether any active fund managers have access to valuable information. By contrast, the average return available to mutual fund investors, although an important statistic in its own right, does not measure heterogeneity in managers' access to information.

Several recent studies have focused on the question of whether any fund managers are skilled or whether all seemingly superior performance can be attributed to luck. Kosowski, Timmermann, Wermers, and White (2006) use a bootstrap technique and find that a "sizable minority" of managers pick stocks well enough to more than cover their costs. Barras, Scaillet, and Wermers (2010) examine the distribution of alphas (and the associated p-values) and find the number of negative-alpha funds (24%) far outweighs the number of positive-alpha funds (.6%). Fama and French (2009) extract demeaned risk-adjusted returns and run simulations to examine whether chance alone could generate the alpha distribution observed in the data. They find only weak evidence in net returns of some managers having enough skill to cover the costs they impose on investors. By contrast, they note

[4]Grinblatt and Titman (1989) estimate that the direct (upwards) survivorship bias is "relatively small" and between 10 and 30 basis points per year; Brown and Goetzmann (1995) get estimates between 20 and 80 basis points; Elton, Gruber, and Blake (1996) find estimates between 71 and 91 basis points based on the three-factor model alphas; and Carhart, Carpenter, Lynch, and Musto (2002) show the bias can be as large as 1 percent in samples longer than 15 years.

the left tail of the actual alpha distribution is far thicker than the tail of the simulated distribution. This left tail could indicate some funds have either negative stock-picking skills or high trading costs or both. I note that this finding also is consistent with the reverse survivorship bias. The alpha estimates are often too low for funds that disappear and so they thicken the left tail of the distribution.

This bias has significant implications beyond the evaluation of the abilities of mutual fund managers. First, individual investors, similar to mutual funds, often stop trading following poor performance.[5] The reverse survivorship bias argument suggests individual investors' true alphas also are not necessarily as low as what their estimated alphas suggest. Second, inferences about CEOs can also be problematic if CEOs' "survival" correlates with the firm's performance and policies. A pooled regression with CEO fixed effects[6] will then produce fixed effect estimates that are biased away from the true distribution of CEOs' abilities or from their contributions to firms' policies.

The paper is organized as follows. Section I describes the data and reports on the correlations between mutual fund survival and alternative alpha estimates. Section II formulates the learning-based structural model, estimates it by using the simulated method of moments, and uses the estimated model to draw inferences about the size of the reverse survivorship bias. Section III concludes.

I. Data and Performance Measurement

I use the mutual fund data from the CRSP (Center for Research in Security Prices) database. I follow French (2008) and Fama and French (2009) and, first, include only funds that invest in U.S. common stocks and, second, combine different share classes of the same fund into a single fund. I restrict the sample to mutual funds that start on January 1984 or after this date. Although the CRSP mutual fund data start in 1962, the pre-1984 part of the data is not reliable.[7] Fama and French (2009) also cut out the pre-1984 part of the data and note the average returns in this part of the data are significantly higher for mutual funds that report monthly returns than for those that

---

[5]See, for example, Seru, Shumway, and Stoffman (2010).

[6]See, for example, Bertrand and Schoar (2003).

[7]See, for example, Elton, Gruber, and Blake (2001) for a critical assessment of the accuracy of the pre-1984 segment of the CRSP Mutual Fund Database.

report annual returns.

I let a mutual fund enter the sample after its combined net asset value across all share classes exceeds $5 million in December 2006 dollars. Once a fund has exceeded this threshold, I keep the fund in the sample no matter what happens to avoid introducing a selection bias. This net-asset-value screen guards against the incubation bias of Evans (2009).[8] The CRSP mutual fund files contain monthly returns up to the end of September 2009. I use returns up to July 2009 to assess which funds are no longer active in September 2009. My sample contains data on 2,599 mutual funds that have return data for at least six months.

A. On the Correlation between Alphas and Fund Disappearance

In this section, I examine the correlation between fund returns and disappearance because of the importance of this correlation in determining the size of the reverse survivorship bias (see expression (2)). Although a number of studies[9] find a positive relation between past returns and fund disappearance, I report on these relations because, first, my sample period (and sample construction) differs from those used in earlier studies, and second, because these estimates serve as inputs for the structural model estimation.

I use three measures of performance throughout this study: the CAPM alpha, the alpha from the three-factor model of Fama and French (1993), and the alpha from the four-factor model of Carhart (1997). I measure these alphas by employing time-series regressions such as

$$R_{i,t} - r_{f,t} = \alpha_i + b_i \left(R_{m,t} - r_{f,t} \right) + s_i \, \text{SMB}_t + h_i \, \text{HML}_t + m_i \, \text{MOM}_t + e_{i,t}, \tag{3}$$

where $R_{i,t}$ is the return on fund i in month t, $r_{f,t}$ is the risk-free rate, $R_{m,t}$ is month-t return on the value-weighted CRSP index, and SMB_t, HML_t, and MOM_t are month-t returns on long-short portfolios for size, value, and momentum. I require a fund to have at least six months of return data

---

[8]Incubation bias arises from fund management companies providing seed money to new funds (to develop a return history) and then selectively opening the funds with best histories to the public. When these funds were allowed to backfill their return histories, such incubation drove a positive wedge between the fund's past performance, which was selected to be good, and its expected performance. Fama and French (2009) use the same $5 million threshold as their main specification and suggest this limit probably exceeds the fund's seed money and thus cuts out the pre-release period returns.

[9]See, for example, Brown and Goetzmann (1995), Khorana (1996, 2001), Chevalier and Ellison (1999), and Lynch and Musto (2003).

to estimate its alpha. I update alphas continuously as funds age to measure the covariance between realized alphas and fund survival.

Table I reports on average alphas (as percentage points per year) for mutual funds that either survive through the t^{th} year or disappear in year t. A mutual fund is included in year t analysis if the fund is still alive at the end of year $t-1$. If the fund disappears by the end of year t, I compute the fund alpha by using all available data from inception until the last month of returns. If the fund is still alive at the end of year t, I compute the alpha by using data up to the end of year t. Each column in Table I reports on average alphas for funds that either survive or disappear in year t. For example, the four-factor model alphas indicate the average alpha is -3.2 percent per year for a fund that disappears at some point during its fifth year. By contrast, the alpha estimate is .07 percent per year for the average fund that survives through its fifth year. Because the dead-fund regressions in the first column are based on at most 12 data points, and possibly as few as six, the three-factor and four-factor alpha estimates are very noisy.

The last row in the table, which reports the fraction of funds that survives at least t years, punctuates the economic importance of the disappearance of mutual funds. For example, 10 percent of mutual funds disappear before reaching the age of four and approximately two thirds of mutual funds survive through the tenth year. (If a mutual fund is T-years old at the end of the CRSP mutual fund data set, I count the fund as a survivor in year 1, ..., T computations but ignore it in year $T+1$, ..., 10 numbers.)

The alpha estimates suggest mutual funds that disappear perform considerably worse than surviving mutual funds. This conclusion is not sensitive to the year in which the comparison is made or to the choice of the asset pricing model used to estimate alphas. However, the size of the gap between the dead and surviving funds varies across the models. These differences in average alphas suggest the disappearance of poorly performing funds is potentially a significant factor in performance evaluation via the reverse survivorship bias channel. For example, the reverse survivorship bias argument suggests the four-factor model alpha estimate of -3.2 percent for those funds that die in their fifth year is probably too low relative to these fund managers' true abilities. The heart of the issue is that these realized alphas may be low not only because the true alpha is low but also because these funds experienced negative idiosyncratic shocks. Because of the latter possibility, the

8

expected alpha for these funds at the time of disappearance was probably higher than -3.2 percent per year. Put differently, had these funds remained in existence, their asymptotic alpha estimates would probably have been higher than their historical alphas.

Figure 1 plots cumulative risk-adjusted returns for surviving and disappearing funds to show how these returns vary from year to year leading up to fund disappearance (or survival). In this figure, I compute the risk-adjusted returns by using both the CAPM and the four-factor model. (I note that the three-factor model estimates are very similar to the four-factor model estimates and thus not reported.) The month-t risk-adjusted return from the four-factor model is

$$\hat{R}_{i,t}^e = R_{i,t} - r_{f,t} - \hat{b}_i \left(R_{m,t} - r_{f,t} \right) - \hat{s}_i \, \mathrm{SMB}_t - \hat{h}_i \, \mathrm{HML}_t - \hat{m}_i \, \mathrm{MOM}_t, \tag{4}$$

where $\hat{b}_i$, $\hat{s}_i$, $\hat{h}_i$, and $\hat{m}_i$ are fund i's loadings on the market, size, value, and momentum factors. I estimate these loadings by using a time-series regression. If a fund disappears in year t, I again estimate its factor loadings by using all data up to the month of disappearance. If the fund survives, I use all data up to the end of year t. Figure 1 reports on cumulative risk-adjusted returns for funds that survive for at least two, four, six, eight, or 10 years, as well as for funds that disappear in years two, four, six, eight, or ten. The estimates suggest mutual funds often disappear following a string of low returns. For example, mutual funds that disappear during their sixth year cumulatively lose approximately 15 percent on a risk-adjusted basis in both specifications.

Table II reports on a set of probit regressions to measure the strength of the correlation between alpha estimates and mutual fund survival. In each regression, the dependent variable takes the value of one if the fund disappears in year t and zero otherwise. I estimate these regression month by month, based on how long each fund has been in existence. For example, one regression uses data on funds that have existed for eight years and two months. I drop such months from the analysis in which no mutual funds disappear. I note that the dependent variables are independent of each other across these cross-sectional regressions because for a fund to appear in the month t regression, its dependent variable must have been zero in all previous regressions. The regressor in Panel A is the fund's alpha (the three leftmost columns) or the t-value associated with the fund's alpha (the three rightmost columns), estimated using monthly returns up to month of the cross-sectional regression. Panel B reports on otherwise identical regressions but lags alpha estimates by one year. I estimate

9

alphas using the CAPM, the three-factor model, and the four-factor model.[10] Table II reports on the average coefficient estimates across the monthly cross-sectional regressions.

The estimates in Table II support the same conclusion as the averages in Table I and the cumulative risk-adjusted returns in Figure 1. The probability that a mutual fund disappears from the data decreases significantly as the fund's alpha estimate increases. The estimates for the CAPM regression in Panel A indicate the annualized exit probability is 0.027 for a fund with $\hat{\alpha} = 0$. This probability decreases to almost zero (0.0003) if the alpha estimate is 2 percent instead. By contrast, if the alpha estimate is -2 percent, the exit probability jumps up to 0.498.[11] Thus, whereas the exit probability is close to zero for slightly negative alphas and for all positive alphas, it increases rapidly as the alpha estimate falls. The results are similar when alpha estimates are replaced with the t-values associated with these first-stage estimates. In the CAPM regression, for example, the annualized exit probability increases from 0.006 to 0.055 when the t-value for the fund's alpha estimate decreases from $+2$ to -2. Panel B shows the results are largely the same when we replace current alpha estimates by estimates lagged by one year.

The probit regression estimates are similar across the three asset pricing models. I note that two effects may make empirically distinguishingly between these models in exit regressions difficult. First, the more complicated the asset-pricing model, the more data are needed to estimate alphas with the same precision. If the amount of data is fixed then the addition of each new factor decreases the precision at which alphas are estimated. For example, even if the four-factor model alphas correlated perfectly with survival, the noisiness of these alphas in the data relative to the CAPM alphas could tilt the regressions to favor the CAPM. Second, if fund disappearance is related to investors' inferences about alphas, the relevant question is not what is the true or ex post best asset pricing model but what is the asset pricing model used by the investors. Before Carhart (1997),

---

[10]These regressions are not free of the errors-in-variables problem, because the regressor is a first-stage alpha estimate. However, if these first-stage residuals are uncorrelated with fund survival, this errors-in-variables problem does not bias the estimates but lowers the test's power by adding noise. Moreover, the average alphas reported in Table I show that the errors-in-variables problem should not materially influence the inferences. These Table I averages could be estimated from a similar survival regression but with its two sides reversed: the alpha estimate would be the dependent variable and the regressors would represent interactions between years and survival.

[11]I compute the annualized numbers as follows. If the alpha estimate is -2 percent, the monthly exit probability, computed from the average point estimates for CAPM, as reported in the first column of Panel A in Table II, is $\Phi\big((-2.84) + (-0.02) * (-62.64)\big) = 0.0558$. Thus the fund disappearance probability is 0.0558 in month 1, $(1 - 0.0558)(0.0558) = 0.0527$ in month 2, and so forth. The annualized exit probability is the sum of these projected exit probabilities. This computation assumes the alpha estimate remains unchanged over the year.

10

for example, mutual fund performance studies did not typically include the momentum portfolio as a risk factor (or as a passive benchmark). Cremers, Petäjistö, and Zitzewitz (2008) note that practitioners commonly evaluate fund managers by comparing their returns with benchmark indices, such as the S&P 500 for large-cap stocks and Russell 2000 for small-cap stocks.

The estimates in Tables I and II and in Figure 1 indicate fund survival correlates significantly with mutual fund performance when performance is measured by fund alphas. These findings, which are consistent with the prior literature on mutual funds, suggest the reverse survivorship bias will probably affect inferences about the distribution of fund managers' abilities.

II. Simulated Method of Moments Estimation

A. Structural Model

In this section, I construct a structural model of mutual fund survival that I then estimate using the CRSP fund data. The Appendix constructs and calibrates an alternative model to investigate the size of the reverse survivorship bias.

I assume each mutual fund's (unobserved) alpha is drawn from a normal distribution with a mean μ and variance σ^2. The market knows the parameters of this distribution. I assume investors cannot distinguish new funds from each other so this normal distribution is also the market's prior distribution about each fund's alpha. The prior distribution's mean and variance at date zero are thus $m_0 = \mu$ and $v_0 = \sigma^2$, respectively.

Each mutual fund generates monthly return observations. I assume the market uses some asset pricing model to obtain an estimate of the fund's monthly risk-adjusted return, $\tilde{R}_{i,t}^e \equiv \alpha_i + \tilde{\varepsilon}_{i,t}$, where $\tilde{\varepsilon}_{i,t}$ is normally distributed with a mean zero and variance σ_e^2. This risk-adjusted return variance reflects two sources of uncertainty: the idiosyncratic shocks to returns and the estimation uncertainty that arises from having to use a finite sample to estimate the asset pricing model for risk adjustment. I note that the market's asset pricing model is correct in that $E[\tilde{R}_{i,t}^e] = \alpha_i$.

The market uses each risk-adjusted monthly return realization to update from the prior distribution to the posterior distribution. It follows from the normality assumptions that the mean and

the variance of the posterior distribution evolve by

$$m_{i,t} = \frac{\frac{m_{i,t-1}}{v_{t-1}} + \frac{\tilde{R}^e_{i,t}}{\sigma^2_e}}{\frac{1}{v_{t-1}} + \frac{1}{\sigma^2_e}}, \text{ and} \tag{5}$$

$$v_t = \frac{1}{\frac{1}{v_{t-1}} + \frac{1}{\sigma^2_e}}, \tag{6}$$

where the posterior variance does not have the fund subscript i because its process is deterministic.

I assume the probability that a fund disappears after month t is a function of the fund's posterior distribution $N(m_{i,t}, v_t)$. I construct this exit function as follows. First, let $\bar{\alpha}$ denote some critical level of alpha. I then compute and record the probability that the fund's alpha is below this level, $p_{i,t} \equiv \int_{-\infty}^{\bar{\alpha}} \phi(\alpha; m_{i,t}, v_t) \, d\alpha$. I assume a fund disappears from the data between months t and $t+1$ with probability $\xi(p_{i,t}, t)$, where $\xi(\cdot)$ is a possibly non-linear and time-dependent function. With these assumptions, the probability that a fund disappears is a function of the amount of mass in the posterior distribution below some critical threshold $\bar{\alpha}$. I expand $\xi(\cdot)$ twice around $p_{i,t} = 0$ and allow it to depend linearly on the fund's age. The approximation of $\xi(\cdot)$ I estimate from the data is then

$$\xi(p_{i,t}, t) \approx \hat{\xi}(p_{i,t}, t) = \gamma_0 + \gamma_1 \, p_{i,t} + \gamma_2 \, p^2_{i,t} + \gamma_3 \, t, \tag{7}$$

where t is the fund's age in years. If the implied exit probability is below zero or greater than one, I set the probability to zero or one, respectively. I adopt this free-function approach to avoid the need to specify the economic mechanism that drives the disappearance of mutual funds. Moreover, because this approach does not force the exit probability to decrease in alpha, it also allows for the possibility that some successful funds may close as well if the manager, for example, moves to a hedge fund. When I estimate both the critical threshold and parameters of $\hat{\xi}(p_{i,t}, t)$ from the data, the structural model can "undo" the built-in learning process if the data suggest such a reversal is warranted. However, I note this five-parameter specification nests, as a special case, the possibility that a fund disappears when the posterior mean falls below some endogenous "learning-based" threshold.

The eight structural parameters of the problem I estimate from the data are μ (the mean of the alpha distribution), σ^2 (the variance of the alpha distribution), σ^2_e (the total variance of risk-adjusted returns), $\bar{\alpha}$ (the critical level in the posterior distribution computation), and γ_0, γ_1, γ_2,

12

and γ_3 (the parameters of the function that translate the amount of probability mass below $\bar{\alpha}$ into an exit probability).

I use 17 moment conditions to identify these eight structural parameters. The first five moment conditions represent the average alphas of the funds that survive through different points in time. These are the same numbers as the per-year averages reported in Table I except I aggregate the data to biennial frequency to reduce noise. The next five moment conditions represent the average alphas of funds that disappear in these same two-year periods. These 10 moment conditions instruct the structural model to match the levels of alphas for surviving and disappearing funds and to match the changes in these alphas as funds mature. The next five moment conditions consist of attrition rates for these two-year periods: How many new funds disappear during the first two years? How many of the funds alive after the first two years disappear during years three or four? And so forth. The inclusion of these attrition rates ensures the rate at which mutual funds fall out of the model matches the empirically observed rate. The average and the variance of the overall alpha distribution constitute the last two moment conditions. I compute each fund's alpha by using up to 10 years of data to construct this overall alpha distribution. These two moments instruct the model to push the observed alpha distribution in the model close to the alpha distribution in the data. I use these 17 moments to estimate the eight structural parameters $(\mu, \sigma^2, \sigma_e^2, \bar{\alpha}, \gamma_0, \gamma_1, \gamma_2, \gamma_3)$ to have an overidentified model.

Because the structural model does not yield closed-form estimation equations, I use the simulated method of moments (SMM) for indirect inference. I draw random funds from the true alpha distribution, generate monthly risk-adjusted returns, update the market's beliefs using expressions (5) and (6), and then compute the probability that a mutual fund disappears from the data. I create new funds until I have generated a large simulated sample of mutual funds from which to compute the simulated moments. The SMM-estimator $\hat{\theta}$ is then

$$\hat{\theta} = \arg\min_{\theta} \left(\hat{M} - M(\theta) \right)' W \left(\hat{M} - M(\theta) \right), \tag{8}$$

where θ is an eight-by-one vector of the structural parameters, $\hat{M}$ is the vector of estimated moments from the data, $M(\theta)$ is the vector of model-implied moments, and W is an arbitrary positive-definite weighting matrix. I use the optimal weighting matrix for estimation.

13

The downside of this simulation approach is that the simulated moments reflect randomness in the simulation process. Even if the number of simulated funds is large, the resultant simulated moments are too noisy to use derivative-based methods to minimize expression (8). Instead of applying the simulated-annealing method, I exploit the specific structure of the problem and compute the simulated moments for expression (8) via numerical integration. Before I begin the minimization routine, I generate a set of primitive $U(0, 1)$ random variables. If the simulations are for N funds, I generate N primitive draws of mutual fund alphas and $N \times 120$ (i.e., 10 years) primitive draws for monthly risk-adjusted returns. I then save these primitive random variables, and each time I re-evaluate expression (8), I recall this set of draws. For example, as the structural parameters μ and σ^2 change by a small amount, indicating a change in the shape of the alpha distribution, the same N primitive random variables now correspond to (slightly) different fund alphas because a fund's alpha α_i can be constructed from a primitive draw ϵ_i as $\alpha_i \equiv \mu + \Phi^{-1}(\epsilon_i)\,\sigma$.

I also allow each mutual fund to both disappear and stay alive (probabilistically) each month in these model-based computations. Each month, I compute the probability that a mutual fund disappears or stays alive. I then use each simulated fund to construct $2 * 120$ fund observations (with different probability weights) so each observation corresponds to a fund either disappearing or staying alive in month t. I estimate the model by setting N to 100,000 in this Monte Carlo integration procedure.

B. Estimation Results

Table III compares the moment conditions between the data and the structural model (Panel A) and reports on the estimates of the eight structural parameters (Panel B). The χ^2 test statistic in Panel B indicates the learning-based model has the best fit to the data when the alphas are estimated from the CAPM. This result may suggest the market uses the CAPM (and not the three- or four-factor models) to evaluate fund performance and to determine which funds should disappear, or that the three- and four-factor models produce noisy alpha estimates relative to the CAPM. However, although the model is rejected under the three- and four-factor model alphas, I note that the inferences from the model are similar irrespective of which asset pricing model one uses to construct the sample moments.

The structural model matches the salient features of the mutual fund data. First, the observed

14

alphas are low for funds that disappear early on but then increase monotonically as the disappearance date recedes. For example, the average CAPM alpha in the data is −4.1 percent for funds that disappear within the first two years. The corresponding simulated number at the estimated parameter values is −4.5 percent. For funds that disappear during years nine and 10, these alphas are −1.9 percent and −1.8 percent, respectively. The CAPM alphas of surviving funds also increase in fund age both in the data and simulations. These moment conditions, however, cause difficulties for the model with the three- and four-factor model alphas. Unlike the CAPM alphas, the three- and four-factor alphas in the data start positive but then turn negative after five years. This reversed U-shaped pattern in the surviving funds' risk-adjusted returns is shown in Figure 1. Figure 2 replicates this risk-adjusted return computation by using data simulated from the model. This figure shows that although the model is partially successful in matching the U-shape pattern, the remaining discrepancies are sufficiently large to reject the model in the three- and four-factor model specifications.

In contrast to the alphas of the surviving funds, the model matches each of the remaining moment conditions irrespective of which asset pricing model one uses to measure alphas. The attrition rates, for example, are nearly the same as in the data. Although the mean of the overall alpha distribution is different depending on the asset pricing model, the model accommodates these shifts and also matches the variance of the alpha distribution.

The first two parameters in Panel B describe the shape of the true alpha distribution. (This distribution also constitutes mutual fund investors' prior distribution for each fund alpha in the model.) While I report on the annualized parameters, I note that the fund alphas in the model are drawn from the underlying monthly distribution. Both the mean and the dispersion of alphas are significantly greater under the CAPM than they are under the three- and four-factor models. For example, whereas 1.2 percent of funds have (true) annual alphas greater than 2 percent per year under the CAPM, this proportion is effectively zero for both the three- and four-factor models.

The risk-adjusted returns, which comprise both the idiosyncratic fund-return component and the estimation uncertainty inherited from the asset-pricing model, have an annualized standard deviation of 8.5 percent under the CAPM. (This estimate is close to the median annualized standard deviation of the CAPM residual in fund-specific regressions that use data on all U.S. equity mutual funds.

This median is 7.6%.) This estimate is high in terms of what it implies about the speed at which the market can resolve uncertainty about fund alphas. I note that in the underlying normal-normal updating problem, the variance of the date T posterior, ignoring the effect of fund attrition, is

$$v_T^2 = \frac{1}{\frac{1}{v_0^2} + \frac{T}{\sigma_e^2}}.$$ (9)

This expression indicates the time T it takes to reach any given posterior variance v_T is $T = \left(\frac{1}{v_T^2} - \frac{1}{v_0^2}\right)\sigma_e^2$. The fact that the variance of the idiosyncratic component is orders of magnitude higher than the variance of the prior implies investors resolve uncertainty about alphas at a low pace. The variance of the posterior distribution, for example, decreases by approximately 10 percent in 10 years.

The critical alpha level, which investors in the model use to construct an input for the mutual fund exit rule, is -2.3 percent (per year) for the CAPM and approximately -1.8 percent for both the three- and four-factor models. (This critical alpha level, which is estimated quite precisely in each of the three models, gives the boundary that the market uses to construct the probability-mass input for the stochastic exit function.) Unlike for the critical boundary $\bar{\alpha}$, the model returns imprecise estimates of the four parameters $(\gamma_0, \gamma_1, \gamma_2, \gamma_3)$ of the exit-probability function. The point estimates of the parameters suggest this imprecision is due to over-parametrization of the model. Although the model permits the exit rule to be stochastic (i.e., the exit probability of a fund can lie between zero and one), the estimated function suggests the mutual fund disappearance pattern in the model follows a simpler rule. In the CAPM-based estimates, a fund stays alive almost always if the probability that the fund's alpha is below -2.3 percent per year is less than 0.02. However, when the amount of probability mass below this critical threshold increases above 0.02, the exit probability jumps to one. In simulations, the result of this rule is that the typical fund disappears when the market's belief about its (CAPM) alpha decreases from the prior mean of .05 percent per year to a posterior mean of -0.21 percent per year.

C. The Size of the Reverse Survivorship Bias

Panel C of Table III reports on the means and percentiles of three different alpha distributions. I report on these distributions separately for the CAPM and the three- and four-factor models. The

first column in each block, labeled "Data," reports the empirical distribution of alphas in the data. For example, the average alpha over all sample mutual funds, and with each alpha estimated using all available data, is −0.71 percent per year in the CAPM, −1.24 percent per year in the three-factor model, and −1.36 percent per year in the four-factor model. (These numbers are not identical to the numbers in Panel A because Panel A's numbers use at most 10 years of fund data.) The second column reports the observed alpha distributions that are computed by simulating from the model at the estimated parameter values. The means of the model-based (observed) alpha distributions are somewhat higher but close to the data-based distribution. The overall shapes of these distributions, however, are close to each other. The similarities in distributions are particularly striking in the CAPM-based estimates in which even tails of the distributions are similar between the data and the model. (I note that no moment conditions in the model force the estimated model distribution to match these percentiles.)

Given that the observed distributions are similar between the data and the model, using the model to examine how the true, unobserved alpha distribution differs from the observed distribution is now useful. What underlying true distribution of fund manager ability is responsible for the observed alpha distribution? The third column in each block, which reports on the true distribution, shows the differences between the true and observed distributions are economically significant. For example, starting from the CAPM-based estimates, the observed average alpha is −.71 percent per year over all funds, but the true mean alpha is far higher, .05 percent per year. The difference between these numbers, 76 basis points per year, is the mean effect of the reverse survivorship bias. The estimated magnitude of this bias is similar when one uses either the three- or four-factor model to estimate the input alphas. The gap between the true and observed mean alpha is 85 basis points per year in the three-factor model and 83 basis points per year in the four-factor model. These estimates suggest that although the average mutual fund manager's true alpha is negative, between −.5 percent and 0 percent per year *after* fees, these alphas are not nearly as low as what the empirical distribution of fund-specific alphas suggests.

The estimates of the true alpha distribution suggest a minority of mutual managers may pick stocks well enough to cover the fees they impose on their investors. In the CAPM-based estimates, 1 percent of managers have alphas greater than 2.06 percent per year. However, adjustments for the size, book-to-market, and momentum factors result in a less optimistic assessment. In the three-

factor model, the top 1 percent of managers have alphas greater than 1.02 percent per year, and in the four-factor model, these right-tail alphas are greater than 0.65 percent per year. Thus, although the reverse survivorship bias significantly distorts the observed alpha distribution relative to the true alpha distribution, the model's conclusion about fund managers' abilities is not unlike those we find in the literature: most mutual fund managers cannot pick stocks well enough to cover the costs they impose on their investors. However, a minority of managers appear to be able to do so.

The three- and four-factor model estimates of the size of the skilled-manager group lie between the estimates in Kosowski, Timmermann, Wermers, and White (2006) and Barras, Scaillet, and Wermers (2010). I note, however, that the benefit of the structural model approach is that it deals with the luck-versus-skill problem in an interesting way. Whereas the extant literature adjusts test statistics to account for the multiple-comparisons problem, a structural model can back out the true alpha distribution from a set of observables. The usual caveat, however, applies and the resultant skill estimates could be sensitive to the modeling assumptions. The structural model I consider in this section is fairly flexible, in particular, with respect to the mutual fund attrition mechanism. Moreover, the model I consider matches not only the explicit moment conditions but also the general shape of the observed alpha distribution. These two considerations increase my confidence in the validity of the inferences I draw about the true alpha distribution.

III. Conclusions

The reverse survivorship bias affects the measurement of mutual fund managers' stock-picking abilities. Because mutual funds often disappear following poor performance, some funds disappear because they experience negative idiosyncratic shocks and not because their true alpha is low. A fund that disappears because of a negative idiosyncratic shock leaves behind an alpha estimate that is too low. Because no offsetting mechanism exists to eliminate just-lucky mutual funds, the observed distribution of alphas is biased downwards relative to the true distribution of alphas. In this paper, I first solve a portfolio choice model for a representative fund investor that endogenizes mutual fund survival to theoretically demonstrate the economic significance of the reverse survivorship bias. I then estimate a variant of this model to measure the size of the bias in the CRSP mutual fund data. I find the mean effect of this bias is large, between 76 and 85 basis points per year depending on

the asset pricing model. Thus the average mutual fund manager's alpha is not nearly as low as the observed alpha distribution suggests.

The question of whether any active fund managers have valuable information is one of the central questions in the empirical asset pricing literature. Although the average actively managed dollar must lose money because of the market-clearing constraint and transaction costs (French 2008), the necessity of these aggregate losses do not rule out the possibility that some active investors profit at the expense of their peers. A small number of mutual fund managers could, for example, be privy to a stream of signals that allows them to profit at the expense of other mutual funds, household investors, or pension funds. The reverse survivorship is of crucial importance in a study that examines whether any mutual managers trade on valuable information. Although the bias described in this paper does not influence the returns available to mutual fund investors, this average return also does not measure heterogeneity in mutual fund managers' information. The correct answer (to the question of how many managers, if any, have enough skill to cover their costs) depends on the counterfactual of what would have happened had no fund actually disappeared following poor performance.

In this paper, I do not consider the equilibrium effects of money flows on alphas. In Berk and Green (2004), alphas get pushed to zero because investors can perfectly diversify all non-market risk and because the model's mutual fund can monopolistically set its fee (to choose its own size). In Pástor and Stambaugh (2009), by contrast, equilibrium alphas do not equate to zero because, first, investors cannot diversify away all risk; second, because the active management industry is competitive; and third, because there is a finite number of mutual fund investors. I note that even if expected alphas are sensitive to money flows, the equilibrium considerations should not alter the premise of the reverse survivorship bias. As long as poorly performing mutual funds still shut down, as they do in the data, the observed alphas will deviate from the true, unobserved alphas.

No simple resolution to the reverse survivorship bias exists, because the endogeneity between fund performance and survival–not the lack of data–drive it. One cannot, for example, resolve this bias by applying a rule that leaves out funds that survive for fewer than k years. The issue is that although the fund performance improves as k increases, we do not know *a priori* what k should be. Moreover, even if we could choose k so the mean of the observed distribution coincides with

the true mean, the overall shape of the observed alpha distribution would still differ from the true distribution. Any inferences drawn about the managers in the tails would not be reliable. A viable resolution to the bias would be to estimate fund alphas by using as little data as possible from the very beginning of funds' lives, before any funds shut down. (This period would need to be less than a year because 1% of funds already shut down during the first year.) However, even though this approach in principle circumvents the bias, the problem is that the alpha estimates based on a just handful of data points would be extremely noisy.

The ramifications of the reverse survivorship bias are not limited to the evaluation of mutual fund managers' abilities. This bias also undoubtedly influences inferences about individual investors' abilities. Similar to mutual funds, individual investors stop trading following poor performance. This sensitivity to poor performance, which also could arise because individual investors learn about their abilities[12], must drive a wedge between investors' observed returns and their true, unobserved abilities. The size of the reverse survivorship bias is the product of three factors: this bias increases in the volatility of returns, the dispersion in survival times, and the correlation between returns and survival. This bias is thus probably large for individual investors as well because, first, individual investors hold under-diversified portfolios, which increases return volatility, and, second, because individual investors appear to be at least as sensitive as mutual funds to poor performance.

---

[12]See, for example, Seru, Shumway, and Stoffman (2010) and Linnainmaa (2009).

Appendix: A Bayesian Portfolio Choice Model with Endogenous Fund Attrition

This appendix calibrates a portfolio choice model with endogenous fund attrition to measure the economic significance of the reverse survivorship bias. In this model, a single mutual fund investor can invest in a risk-free asset, the market portfolio, and an actively managed mutual fund. The key features of the model are that, first, the investor is uncertain about the fund's alpha and, second, that the fund's survival depends on the investor's continuing investment. Each period, the investor can abandon the current fund and, if the investor exercises this option, the old mutual fund disappears and the investor draws a new fund with an unknown alpha. This switch "resets" the investor's prior belief about the fund's alpha.[13]

I estimate a different model in the body of the paper for two reasons. First, the main model is more parsimonious because it directly specifies a mapping from posterior distributions to exit probabilities. The portfolio choice model described here generates qualitatively similar mapping, but it requires one to specify (or to estimate from the data) many parameters that are not of direct interest, such as the properties of market returns, the risk-free rate, subjective discount rates, and so forth. Second, the model described here makes restrictive assumptions about the mutual fund industry and equilibrium: there is a single investor and the fund's survival is contingent on the investor's continuing investment; this is a partial equilibrium model with exogenously specified return processes; and the investor can invest in, and learn about, only one mutual fund at a time. By contrast, the use of the exit function in the estimated model sidesteps such concerns because as long as the functional form of the exit function is general enough, it can match any underlying economic model.

---

[13]Dangl, Wu, and Zechner (2008) also study a model in which investors learn about mutual fund managers' abilities and the management company has the option to replace the manager. However, whereas they focus on the model's implications on the relations between fund size, portfolio risk, and the fund manager's tenure, I examine how the endogenous survival mechanism biases fund-specific performance estimates.

A. Assumptions

I assume an infinite-horizon investor maximizes log-utility over consumption:

$$E\left[\sum_{t=0}^{\infty} \beta^t \log c_t\right], \tag{10}$$

where β is the investor's discount rate. The wealth dynamics are given by

$$W_{t+1} = (W_t - c_t)\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f) + \theta_z(\tilde{z}_t - r_f)\big), \tag{11}$$

where r_f is the risk-free rate, θ_m is the proportion of wealth invested in the market portfolio, $\tilde{r}_{m,t}$ is the return on the market portfolio, θ_z is the proportion of wealth invested in the mutual fund, and $\tilde{z}_t$ is the return on the mutual fund. I assume $\theta_z \geq 0$ so the investor cannot short the fund. The investor knows all parameters of the problem except the fund's alpha. The date t return on the market portfolio is

$$\tilde{r}_{m,t} = \mu_m + \tilde{\varepsilon}_{m,t}, \tag{12}$$

where is μ_m is the (known) expected market return and $\tilde{\varepsilon}_{m,t}$'s are i.i.d. from period to period. I assume each $\tilde{\varepsilon}_{m,t}$ is drawn from a left-truncated normal distribution with truncation at $x = -1$. The underlying *untruncated* distribution has a mean zero and variance σ_m^2. I assume the fund's market-model beta is one so that its return is

$$\tilde{z}_t = \alpha + \tilde{r}_{m,t} + \tilde{\varepsilon}_{z,t}, \tag{13}$$

where α is not known to the investor. The model I have in mind is one where the fund's market exposure is unknown and the estimation of this exposure from the data increases the variance of signals about alpha. However, instead of formalizing this estimation uncertainty by modeling β as either unknown to the investor or an i.i.d. draw each period, I interpret expression (13) as delegating, for simplicity, this estimation uncertainty component to the residual $\tilde{\varepsilon}_{z,t}$. I assume each $\tilde{\varepsilon}_{z,t}$ is also drawn from a left-truncated normal distribution with truncation at $x = -1$. The underlying untruncated distribution has a mean zero and variance σ_z^2. I apply these truncation assumptions to give both the market portfolio and the mutual fund limited liability.

The investor updates his beliefs after each date as follows. First, after observing both fund and market returns, the investor backs out the signal $s_t \equiv \tilde{z}_t - \tilde{r}_{m,t}$ about the fund's alpha. The investor then reverses the truncation by computing the signal realization s_t' by mapping the truncated distribution to an untruncated distribution. The investor's date t prior belief about the mean of this untruncated distribution is normal with mean m_t and variance v_t. The investor knows the true population distribution of alphas, $N(\mu_\alpha, \sigma_\alpha^2)$, so $m_0 = \mu_\alpha$ and $v_0 = \sigma_\alpha^2$. The belief dynamics are then given by[14]

$$m_{t+1} = \frac{\frac{m_t}{v_t} + \frac{s_t'}{\sigma_z^2}}{\frac{1}{v_t} + \frac{1}{\sigma_z^2}}, \tag{14}$$

$$v_{t+1} = \frac{1}{\frac{1}{v_t} + \frac{1}{\sigma_z^2}}. \tag{15}$$

Each period, the investor can abandon the current fund (thus causing the fund to disappear) and, if the investor so chooses, draw a new fund with an unknown alpha. I assume that abandoning a fund and drawing a new fund costs the investor κ percent of wealth. This cost may represent components such as front-end loads, back-end loads, and the abnormal transaction costs incurred when assets are sold in the event of fund termination.

B. Solution

The investor's indirect utility is a function of three state variables: the current wealth W_t, the mean of the prior distribution m_t, and the variance of the prior distribution, v_t. Because the investor maximizes log-utility, the wealth and belief terms are additive in the indirect utility function. I conjecture that

$$V(W_t, m_t, v_t) = A + B \log W_t + g(m_t, v_t), \tag{16}$$

---

[14]The trick I apply here to keep the posterior distribution closed under updating is more transparent in an alternative setup with log-normally distributed returns. Instead of working with a log-normal likelihood function, an investor could be uncertain about the mean of the normal variate $\tilde{x}$ in $\tilde{r} \equiv e^{\tilde{x}} - 1$. The investor would then back out from each return observation the normal variate realization, $\log(1 + \tilde{r})$, and use this signal to update beliefs about the mean of $\tilde{x}$. I do not use this log-normal assumption because of its downside that the variance of $\tilde{r}$ also changes as the mean of $\tilde{x}$ changes. By contrast, the truncation of the normal distribution at $x = -1$ is an innocuous return-distribution twist because, for reasonable return volatilities, the amount of truncated mass is effectively zero. See Johnson, Kotz, and Balakrishnan (1994) for details on truncated normal distributions.

where A and B are constants and $g(\cdot)$ is a function of the investor's date t beliefs. The investor's optimization problem with this conjecture becomes

$$
\begin{aligned}
V(W_t, m_t, v_t) & \\
= & \max_{c_t, \theta_m, \theta_z \geq 0} \mathrm{E}\left\{ \log c_t + \beta V\left((W_t - c_t)\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f) + \theta_z(\tilde{z}_t - r_f)\big), \tilde{m}_{t+1}, v_{t+1} \right) \right\} \\
= & \ \beta A + \max_{c_t} \left\{ \log c_t + \beta B \log(w_t - c_t) \right\} \\
& + \max \left\{ \beta B \max_{\theta_m} \left\{ \mathrm{E}\big[\log\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f)\big) \big] \right\} + \beta g(m_t, v_t), \right. \\
& \ \ \beta B \max_{\theta_m, \theta_z \geq 0} \left\{ \mathrm{E}\big[\log\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f) + \theta_z(\tilde{z}_t - r_f)\big) \big] \right\} + \beta\mathrm{E}\big[g(\tilde{m}_{t+1}, v_{t+1})\big], \\
& \ \ \left. \beta B \max_{\theta_m, \theta_z \geq 0} \left\{ \mathrm{E}\big[\log\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f) + \theta_z(\tilde{z}_t - r_f) - \kappa\big) \big] \right\} + \beta\mathrm{E}\big[g(\tilde{m}_1, v_1)\big] \right\}.
\end{aligned}
$$

$$(17)$$

The last three lines of expression (17) account for the fact that each period, an investor can make a choice that disrupts the natural evolution of beliefs from (m_t, v_t) to $(\tilde{m}_{t+1}, v_{t+1})$. The first possibility, shown on the third to the last line, is that the investor invests in neither the current nor a new mutual fund. The investor withdraws the money from the current fund, the fund disappears, and the evolution of beliefs stops. I note that if κ is low and if the investor has ever invested in a mutual fund, the investor can never reach a belief state in which he would choose this option. If the investor has ever invested in a fund, the option to draw a new fund must dominate this quitting choice. This option is chosen only if the distribution of alphas is so unattractive that the investor never invests in a mutual fund.

The second possibility, shown on the second to the last line, is that the investor remains with the current mutual fund. The investor's belief about the mean mutual fund return $\tilde{z}_t$ is m_t, and the investor updates his beliefs to $(\tilde{m}_{t+1}, v_{t+1})$ after observing the risk-adjusted (and transformed) return. The third possibility, shown on the last line, is that the investor abandons the current mutual fund (the fund disappears) and draws a new fund. The investor's prior distribution about the new fund is the same as the original prior distribution, (m_0, v_0), and so the investor's beliefs "restart" if the investor exercises this abandonment option. If the investor draws a new fund, the

investor's prior mean about the next-period mutual fund return is m_0 and the investor updates to $(\tilde{m}_1, v_1)$ based on the return realization. The investor pays κ to exercise this option.

The optimal consumption from the first-order condition of expression (17) is $c_t^* = \frac{1}{1+\beta B} W_t$. Constants A and B can be solved by inserting the optimal consumption back into expression (17) and matching the coefficients against the conjecture in expression (16). The value function then simplifies to

$$V(W_t, m_t, v_t) = \frac{\beta \log \beta + (1-\beta) \log(1-\beta)}{(1-\beta)^2} + \frac{1}{1-\beta} \log W_t + g(m_t, v_t), \tag{18}$$

where $g(m_t, v_t)$ solves the following functional equation:

$g(m_t, v_t)$

$$
\begin{aligned}
= \quad \beta \max \Bigg\{ & \frac{1}{\beta(1-\beta)^2} \max_{\theta_m} \Big\{ \mathrm{E}\big[\log\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f)\big)\big] \Big\}, \\
& \frac{1}{1-\beta} \max_{\theta_m, \theta_z \geq 0} \Big\{ \mathrm{E}\big[\log\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f) + \theta_z(\tilde{z}_t - r_f)\big)\big] \Big\} + \mathrm{E}\big[g(\tilde{m}_{t+1}, v_{t+1})\big], \\
& \frac{1}{1-\beta} \max_{\theta_m, \theta_z \geq 0} \Big\{ \mathrm{E}\big[\log\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f) + \theta_z(\tilde{z}_t - r_f) - \kappa\big)\big] \Big\} + \mathrm{E}\big[g(\tilde{m}_1, v_1)\big] \Bigg\}.
\end{aligned}
\tag{19}
$$

Expression (18) verifies the conjecture about the form of the value function. The investor's optimal investment decisions are determined by $g(m_t, v_t)$ in expression (19). The first argument in the inner-maximization problem simplifies the corresponding line in expression (17) by noting that if an investor ever abandons a mutual fund without picking a new one, the investor's beliefs must forever remain stuck in the same belief state. The investor's problem then becomes a standard stochastic log-utility investment problem that has the solution shown on the first line of expression (19). I note that the function $g(m_t, v_t)$ also can be solved, up to a static portfolio choice problem, in those belief states where the investor has resolved all uncertainty about a fund's alpha. Analogously to the permanent-exit case, such an investor must remain in the same belief state, and the value of $g(m_t, v_t)$ at such a $v_t = 0$ boundary point is

$$g(m_t, 0) = \frac{\beta}{(1-\beta)^2} \max_{\theta_m, \theta_z \geq 0} \Big\{ \mathrm{E}\big[\log\big(1 + r_f + \theta_m(\tilde{r}_{m,t} - r_f) + \theta_z(\tilde{z}_t - r_f)\big)\big] \Big\}. \tag{20}$$

The form of function $g(m_t, v_t)$ in expression (19) shows that an investor's decision to either remain with the current fund or to switch to a new fund depends crucially on the evolution of beliefs. The path of beliefs, in turn, depends crucially on the posterior uncertainty about the fund's alpha. If the prior distribution about a fresh mutual fund's alpha "dominates" the posterior distribution about the current fund's alpha (net of the strike price κ), the investor abandons the current fund and restarts the problem with a new fund. The log-utility assumption, which shuts down the intertemporal hedging demand component, cleanly isolates the value of this abandonment option.

I solve the problem in three steps. First, I create a mean-variance belief grid with 1,000 points for the posterior mean m_t and 1,200 points for the posterior variance v_t.[15] Second, I solve the static portfolio choice problems in expression (19) for each grid point in the m_t dimension of the grid. Third, I generate initial guesses about the value of $g(m_t, v_t)$ for each grid point and then start iterating over the grid, sweeping recursively from v_T toward v_0 at each iteration. I compute the value of $g(m_t, v_t)$ from expression (19) at each grid point given the initial guesses of the value of this function at other grid points. I note that the expectation about next period's $g(\tilde{m}_{t+1}, v_{t+1})$ depends on the uncertain evolution of beliefs. I compute the transition probabilities from m_t to all different states $\tilde{m}_{t+1}$ using expression (14) together with the distributional assumptions about the signal $\tilde{s}_t$. I iterate over the mean-variance belief grid until the values of $g(m_t, v_t)$ have converged at each grid point. The solution to the investor's problem requires value-function iterations because of the abandonment option. An investor who abandons the fund transitions back to $g(m_0, v_0)$ and the function value in this grid point, in turn, depends on his optimal choices in every possible state that follows.

C. Calibration

I assume each period in the model represents one month and fix the parameters of the model to the following values. First, I set the mean and standard deviation of the true distribution of annual (net) alphas to $\mu_\alpha = -0.5\%$ and $\sigma_\alpha = 1.25\%$. This standard deviation is the middle estimate in

---

[15]I choose the grid for v_t to match the deterministic evolution of the posterior variance in expression (15). Thus the last node corresponds to the variance of beliefs after having invested in the same fund for 100 years. (I calibrate the model so that one period corresponds to one month.) I assume the variance of beliefs drops to zero after this date. I create the grid for m_t so it covers 99.99 percent of the true population distribution of α.

Fama and French (2009). I set the standard deviation of the idiosyncratic fund return component $\tilde{\varepsilon}_{z,t}$ to 8 percent per year, which is between the mean and median estimates for the annualized standard deviation of the CAPM residual in monthly fund-by-fund regressions. Second, I fix the non-mutual fund parameters of the model by setting the risk-free rate to $r_f = 5\%$ and the expected return and standard deviation of the market portfolio to $\mu_m = 10\%$ and $\sigma_m = 30\%$. I set the investor's discount rate (β) to .9. Finally, I choose the cost of switching a fund, κ, so that the 10-year survival rate in the model is close to the empirical survival rate in the mutual fund data. A value of $\kappa = 0.064$ gives a survival rate of 69.3 percent, which is close to the 70 percent survival rate reported in Table I. I note that survival rate increases monotonically in κ because it only decreases the one-period expected return in expression (19).

Panel A of Table IV reports the average alphas (in percentage points per year) for mutual funds that either survive through the t^{th} year or that disappear in year t. Both the underlying computations and the presentation of the results in this table are identical to Table I, which is based on the CRSP sample. The alpha estimates in the model-based simulations exhibit the same upward-sloping pattern observed in the actual sample. The average alpha is negative for both surviving and dead funds in year one (-0.5% and -35.5%, respectively), but both of these averages increase over time. For example, the average observed alpha is 1.98 percent per year for funds that survive through the tenth year. By contrast, the average alpha is -4.12 percent for funds that disappear during the tenth year.

Figure 3 plots the critical posterior mean value in the model that the investor uses to decide whether to abandon the existing fund and to draw a replacement fund from the true alpha distribution. The critical threshold is initially low, approximately -10 percent per year. This low threshold value suggests the investor abandons a fund early only if the realized return is low. This low threshold value thus explains why alphas are as low as they are in Panel A for funds that disappear early on. The critical alpha threshold increases smoothly over time. After 10 years, the critical value is approximately -3 percent per year, and after 25 years, the threshold is -1 percent per year. The increase in the critical threshold leads to the increasing pattern in average alphas seen in Panel B of Table IV.

Panel B of Table IV reports the size of the reverse survivorship bias within this portfolio choice

model. The first column ("biased") reports the distribution of observed alphas from the actual model. The second column ("unbiased") changes the model so that every time a mutual fund disappears in the real model, a randomly chosen fund disappears in this alternative model. This computation constitutes a benchmark for the "biased" column because the distributions of mutual fund sample lengths are identical. Finally, the third column reports the true (noiseless) alpha distribution. I simulate at most 10 years of monthly return data from the model for 100,000 mutual funds to construct these distributions. Although the true mean of the alpha distribution is -0.5 percent in the calibrated model, the observed distribution has a significantly lower mean of -1.13 percent because of the disappearance of poorly performing funds. Thus, in this rough calibration, the mean effect of the reverse survivorship bias is 63 basis points per year.

The percentiles for the observed and true alpha distributions show that the reverse survivorship bias does not evenly shift the distribution downwards. The funds that disappear must have low realized alphas, so the salient result of the reverse survivorship bias is the thickening of the left tail of the observed alpha distribution relative to the true distribution. The worst 5 percent of funds in the observed alpha distribution have alpha realizations that are -11.6 percent or lower. If funds were to disappear randomly, as they do in the second column, this 5^{th} percentile of the distribution would be just -9.1 percent. Although this shift is most pronounced in the left tail of the distribution, it also appears subtly in the right tail of the distribution. The biased and unbiased values for the 95^{th} percentile of the empirical alpha distribution are, for example, 7.8 percent and 8.1 percent, respectively. This calibration suggests the reverse survivorship bias can significantly distort the shape of the observed alpha distribution. This distorting effect of the reverse survivorship bias is important because the literature on alpha distributions draws inferences about the tails of this distribution and not about its mean.

REFERENCES

Barras, L., O. Scaillet, and R. Wermers (2010). False discoveries in mutual fund performance: Measuring luck in estimated alphas. *Journal of Finance 65*, 179–216.

Berk, J. B. and R. C. Green (2004). Mutual fund flows and performance in rational markets. *Journal of Political Economy 112*(6), 1269–1295.

Bertrand, M. and A. Schoar (2003). Managing with style: The effect of managers on firm policies. *Quarterly Journal of Economics 118*(4), 1169–1208.

Brown, S. J., W. Goetzmann, R. G. Ibbotson, and S. A. Ross (1992). Survivorship bias in per-

formance studies. *Review of Financial Studies 5*(4), 553–580.

Brown, S. J. and W. N. Goetzmann (1995). Performance persistence. *Journal of Finance 50*(2), 679–698.

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance 52*(1), 57–82.

Carhart, M. M., J. N. Carpenter, A. W. Lynch, and D. K. Musto (2002). Mutual fund survivorship. *Review of Financial Studies 15*(5), 1439–1463.

Carpenter, J. N. and A. W. Lynch (1999). Survivorship bias and attrition effects in measures of performance persistence. *Journal of Financial Economics 54*(3), 337–374.

Chevalier, J. and G. Ellison (1999). Career concerns of mutual fund managers. *Quarterly Journal of Economics 114*(2), 389–432.

Cremers, M. K. J., A. Petäjistö, and E. Zitzewitz (2008). Should benchmark indices have alpha? Revisiting performance evaluation. Working paper, Yale University.

Dangl, T., Y. Wu, and J. Zechner (2008). Market discipline and internal governance in the mutual fund industry. *Review of Financial Studies 21*(5), 2307–2343.

Elton, E. J., M. J. Gruber, and C. R. Blake (1996). Survivorship bias and mutual fund performance. *Review of Financial Studies 9*(4), 1097–1120.

Elton, E. J., M. J. Gruber, and C. R. Blake (2001). A first look at the accuracy of the CRSP mutual fund database and a comparison of the crsp and morningstar mutual fund databases. *Journal of Finance 56*(6), 2415–2430.

Evans, R. B. (2009). Mutual fund incubation. *Journal of Finance*. Forthcoming.

Fama, E. J. and K. French (1993). Common risk factors in the returns of stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Fama, E. J. and K. French (2009). Luck versus skill in the cross section of mutual fund returns. University of Chicago working paper.

French, K. R. (2008). Presidential address: The cost of active investing. *Journal of Finance 63*(4), 1537–1573.

Grinblatt, M. and S. Titman (1989). Mutual fund performance: An analysis of quarterly portfolio holdings. *Journal of Business 62*(3), 393–416.

Johnson, N. L., S. Kotz, and N. Balakrishnan (1994). *Continuous Univariate Distributions*, Volume 1. New York, NY: John Wiley and Sons.

Khorana, A. (1996). Top management turnover: An empirical investigation of mutual fund managers. *Journal of Financial Economics 40*(3), 403–426.

Khorana, A. (2001). Performance changes following top management turnover: Evidence from open-end mutual funds. *Journal of Financial and Quantitative Analysis 36*(3), 371–393.

Kosowski, R., A. Timmermann, R. Wermers, and H. White (2006). Can mutual fund "stars" reall pick stocks? New evidence from a bootstrap analysis. *Journal of Finance 61*(6), 2551–2595.

Linnainmaa, J. (2009). Why do (some) households trade so much? Working paper, University of Chicago.

Lynch, A. W. and D. K. Musto (2003). How investors interpret past fund returns. *Journal of Finance 58*(5), 2033–2058.

Pástor, Ľ. and R. F. Stambaugh (2009). On the size of the active management industry. Working paper, University of Chicago.

Pástor, Ľ., L. A. Taylor, and P. Veronesi (2009). Entrepreneurial learning, the IPO decision, and the post-IPO drop in firm profitability. *Review of Financial Studies 22*(8), 3005–3046.

Seru, A., T. Shumway, and N. Stoffman (2010). Learning by trading. *Review of Financial Studies 23*(2), 705–739.

Williams, D. (1991). *Probability with Martingales*. Cambridge, Massachusetts: Cambridge University Press.
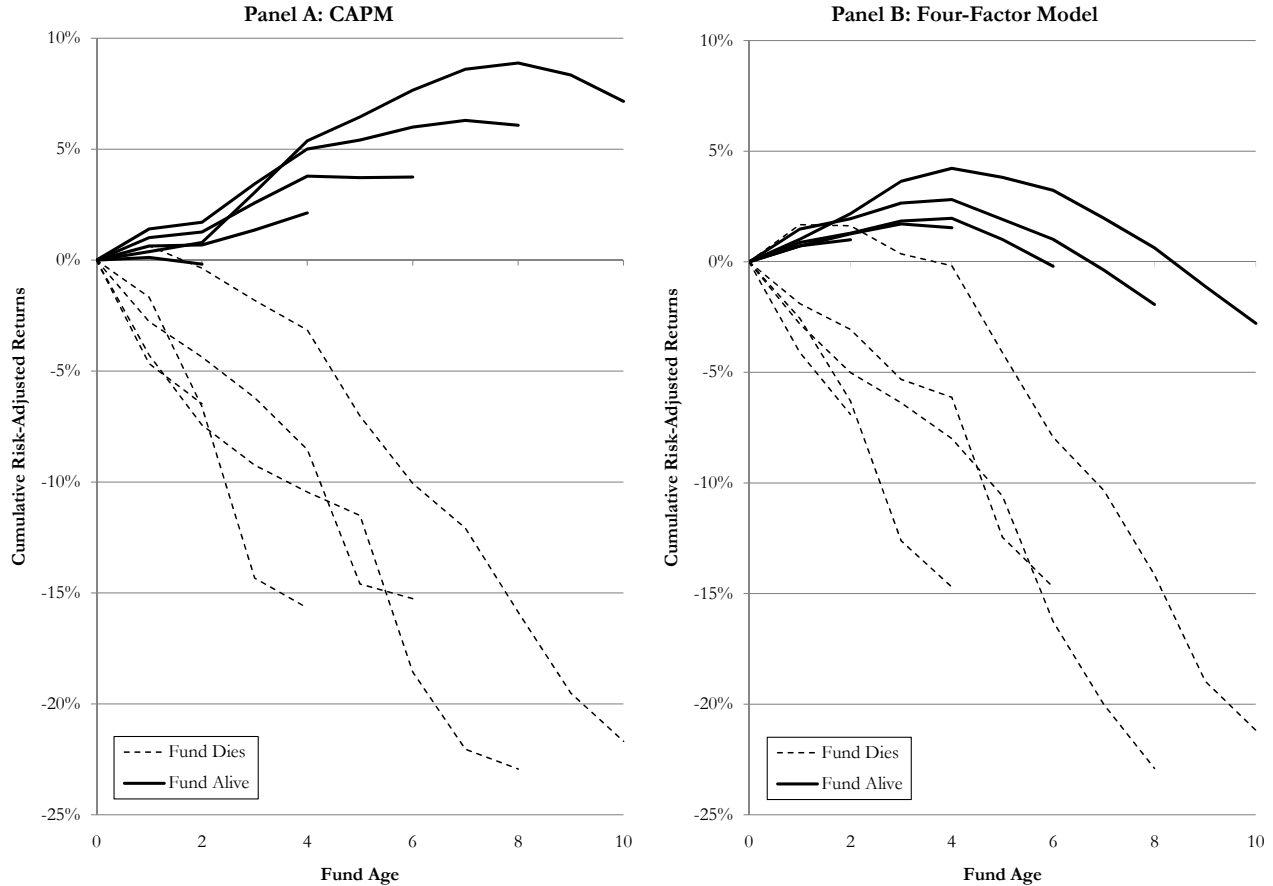
Figure 1. Cumulative risk-adjusted returns for mutual funds conditional on fund survival. I use time-series regressions to estimate the CAPM (Panel A) and Carhart's four-factor model (Panel B) loadings for all mutual funds that either survive through year t or that disappear in year t. I then compute from these estimates the cumulative risk-adjusted returns up to the end of year t or the month the fund disappears. The sample includes all 2,599 mutual funds that invest primarily in U.S. common stocks, that first appear in the CRSP data on or after January 1984, and that have at least six months of return data. The sample covers a 25-year period up to the end of September 2009. Different share classes of the same fund are combined into a single fund. Thick solid lines denote average cumulative risk-adjusted returns of the funds that survive through the year; thin dashed lines denote average cumulative risk-adjusted returns of the funds that disappear during the year. This figure plots cumulative risk-adjusted returns for even-numbered survival and disappearance years.
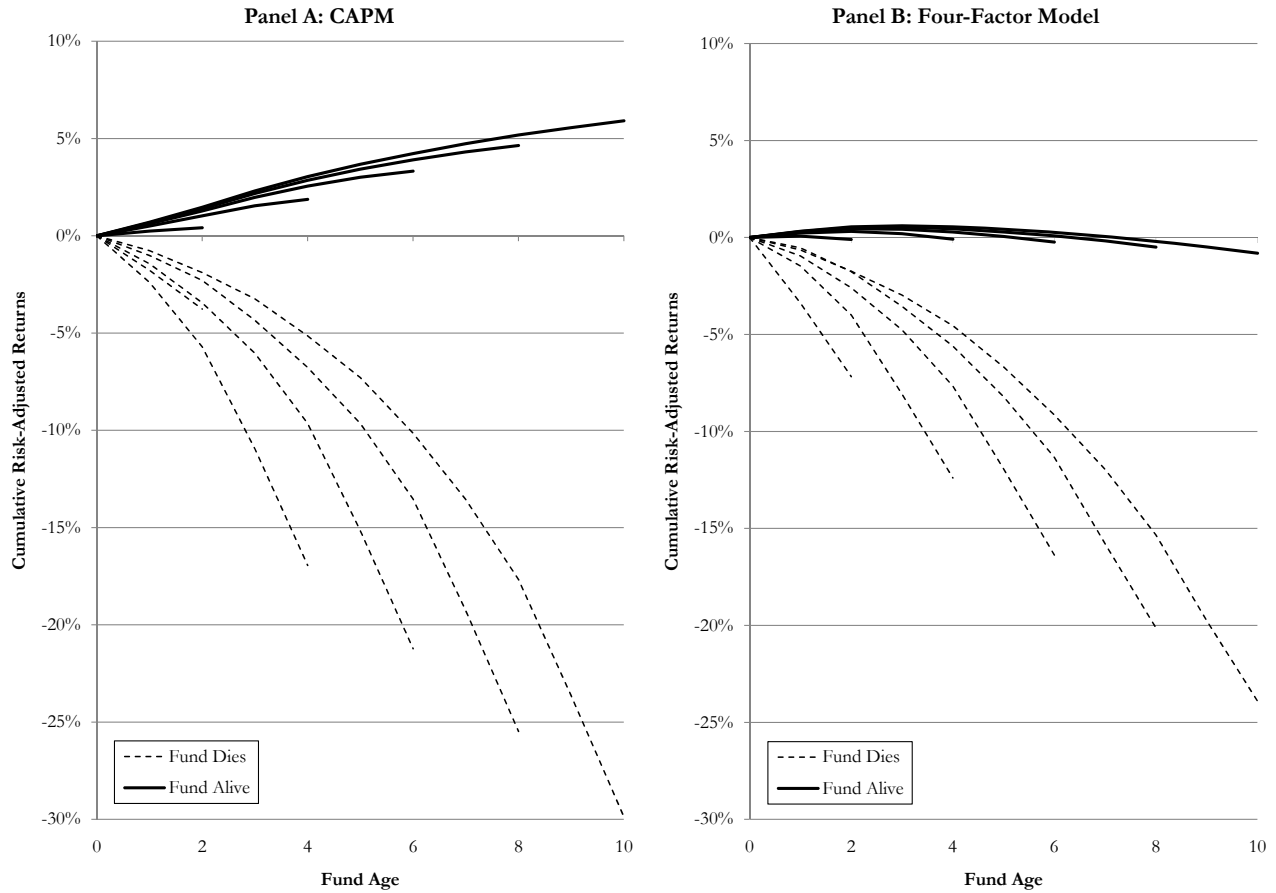
Figure 2. Cumulative risk-adjusted returns for mutual funds conditional on fund survival, simulated from a structural model. I simulate data from the structural model for 100,000 mutual funds and compute cumulative risk-adjusted returns up to the end of year t or the month the fund disappears. Thick solid lines denote average cumulative risk-adjusted returns of the funds that survive through the year; thin dashed lines denote average cumulative risk-adjusted returns of the funds that disappear during the year. This figure plots cumulative risk-adjusted returns for even-numbered survival and disappearance years. This figure is a model-based counterpart of Figure 1.
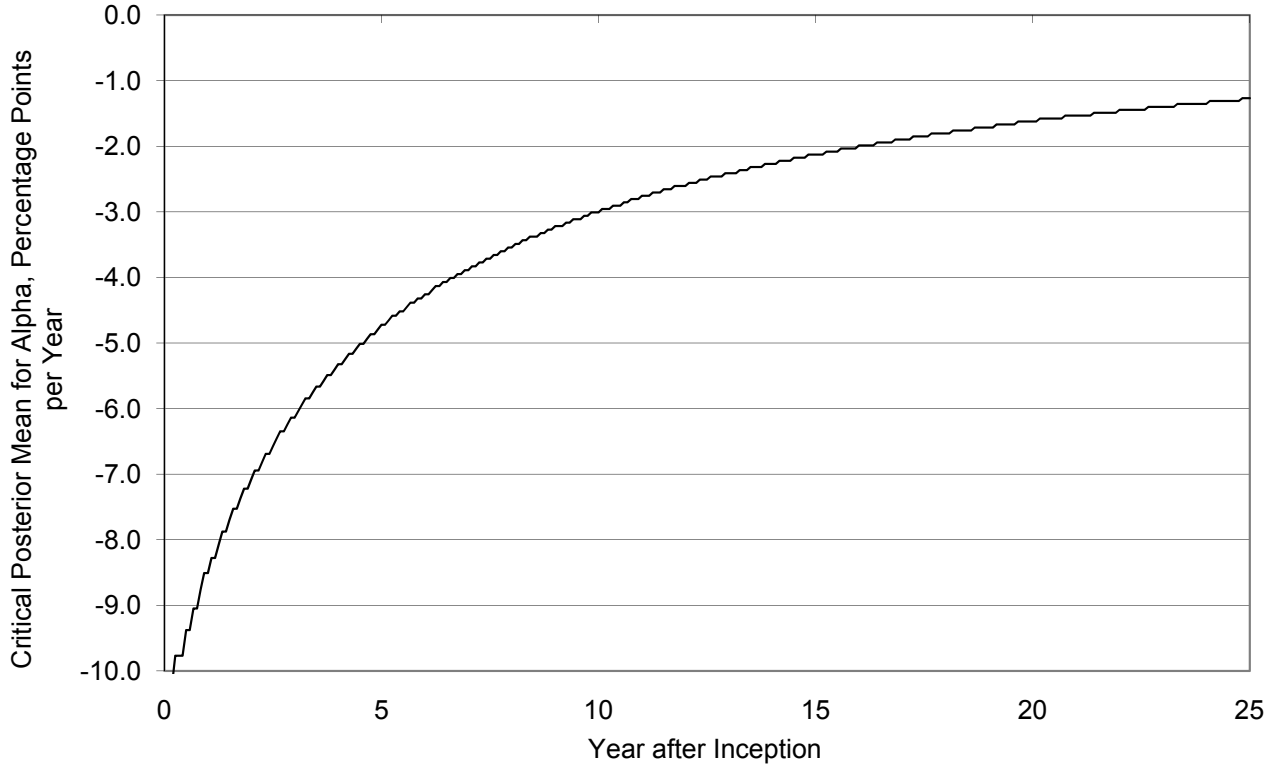
Figure 3. Critical posterior mean boundary for alpha (in percentage points per year) in a Bayesian portfolio choice model with endogenous fund attrition. This figure shows the critical level for the posterior mean for alpha in a Bayesian portfolio choice model below which a fund shuts down. In this model, an infinitely lived representative mutual fund investor with log-utility can invest in a risk-free asset, the market portfolio, and an actively managed mutual fund. The investor is uncertain about the fund's alpha but updates beliefs each month using Bayes' rule. Each month, the investor can abandon the existing mutual fund by paying a cost of κ. He then draws a new fund from the true alpha distribution and restarts the problem. The mean of the true alpha distribution (μ_α) is -0.5 percent and its standard deviation (σ_α) is 1.25 percent; the expected return of the market portfolio (μ_m) is 10 percent, and its standard deviation (σ_m) is 30 percent; the risk-free rate (r_f) is 5 percent; and the investor's discount rate (β) is .9. (These are annualized parameter values. Each period in the model is one month long.) The standard deviation of the idiosyncratic return component is 8 percent. The cost of switching a fund is $\kappa = 0.064$.

Table I
Mutual Fund Alpha Estimates Conditional on Fund Survival, Percentage Points per Year

This table reports the average alpha estimates for mutual funds based on whether the mutual fund survives or does not survive over t years. The sample includes all 2,599 mutual funds that invest primarily in U.S. common stocks, that first appear in the CRSP data on or after January 1984, and that have at least six months of return data. The sample covers a 25-year period up to the end of September 2009. Different share classes of the same fund are combined into a single fund. If a fund is still alive at the end of year $t - 1$, I estimate a time-series regression by using monthly returns up to the end of year t. If the fund disappears before the end of year t, I include returns up to the last available one. A fund that disappears by the end of year t is assigned to a pool of year-t dead funds; funds still alive at the end of the year are assigned to a pool of year-t alive funds. This table reports average α's for funds that either survive or do not survive year t. I use CAPM, a three-factor model, and a four-factor model to estimate alphas. The minimum time-series length requirement is six months for the CAPM, eight months for the three-factor model, and nine months for the four-factor model to equalize the minimum number of degrees of freedom in the regressions. I report t-values in parentheses. Row "Survival rate" reports the fraction of funds that survive for at least t years.

Model	Survive	Fund Age in Years									
		1	2	3	4	5	6	7	8	9	10
CAPM	Yes	0.25	−0.07	0.18	0.55	0.55	0.63	0.69	0.75	0.74	0.70
		(1.00)	(−0.38)	(1.12)	(3.99)	(4.25)	(5.33)	(6.12)	(6.96)	(7.24)	(6.88)
	No	−4.08	−4.08	−4.66	−4.51	−3.66	−2.87	−2.54	−3.04	−1.60	−2.23
		(−2.93)	(−4.94)	(−6.55)	(−5.92)	(−7.22)	(−2.92)	(−6.20)	(−6.50)	(−3.32)	(−6.75)
Three-Factor Model	Yes	0.95	0.65	0.60	0.57	0.30	0.14	−0.03	−0.11	−0.13	−0.13
		(4.01)	(4.01)	(4.16)	(4.44)	(2.53)	(1.36)	(−0.25)	(−1.13)	(−1.35)	(−1.31)
	No	−0.69	−3.91	−3.73	−4.24	−3.35	−2.60	−2.25	−2.83	−2.46	−2.09
		(−0.25)	(−5.41)	(−6.42)	(−7.21)	(−6.43)	(−2.88)	(−6.45)	(−7.25)	(−6.78)	(−6.33)
Four-Factor Model	Yes	0.95	0.51	0.43	0.38	0.07	−0.04	−0.19	−0.25	−0.29	−0.29
		(4.07)	(3.18)	(3.16)	(3.23)	(0.68)	(−0.40)	(−2.05)	(−2.77)	(−3.19)	(−3.19)
	No	−2.91	−4.45	−3.36	−4.22	−3.17	−2.71	−2.39	−2.90	−2.25	−2.19
		(−2.05)	(−5.28)	(−5.80)	(−8.13)	(−7.43)	(−3.12)	(−6.03)	(−6.76)	(−6.37)	(−6.20)
Survival Rate		0.990	0.966	0.929	0.900	0.861	0.830	0.801	0.767	0.738	0.700
N(Alive)		2,603	2,517	2,300	2,137	1,974	1,834	1,697	1,559	1,384	1,214
N(Dead)		18	64	90	72	79	70	59	65	65	72

Table II
Mutual Fund Survival and Alphas: Probit Analysis

This table reports Fama-MacBeth-style probit regressions of mutual fund disappearance on rolling estimates of fund alphas (three leftmost columns) or on rolling estimates of t-values associated with fund alphas (three rightmost columns). The sample includes all 2,599 mutual funds that invest primarily in U.S. common stocks, that first appear in the CRSP data on or after January 1984, and that have at least six months of return data. The sample covers a 25-year period up to the end of September 2009. Different share classes of the same fund are combined into a single fund. The dependent variable in this table's exit regressions takes the value of one if a fund disappears during year t and zero if it does not. I estimate the exit regression separately for each month (based on how long each fund has existed) and drop out months in which no mutual funds disappear. For example, one regression uses data on funds that have existed for eight years and two months. This table reports the averages and t-values of the resultant 125 cross-sectional estimates. I estimate fund alphas from monthly time-series regressions by using all data from the first month of fund returns up to the last available return (Panel A). Panel B lags the alpha estimates by a year. The underlying panel data contain 294,771 observations.

	Panel A: Current Alpha Estimates					
	Slope Estimates as Regressors			t-values as Regressors		
	CAPM	FF3	FF4	CAPM	FF3	FF4
Intercept	−2.84	−2.84	−2.84	−2.89	−2.89	−2.88
	(−160.01)	(−171.00)	(−181.20)	(−139.88)	(−138.33)	(−162.36)
Slope	−62.64	−55.39	−60.56	−0.19	−0.15	−0.15
	(−8.40)	(−10.48)	(−12.53)	(−12.49)	(−11.59)	(−13.43)

	Panel B: Alpha Estimates Lagged by a Year					
	Slope Estimates as Regressors			t-values as Regressors		
	CAPM	FF3	FF4	CAPM	FF3	FF4
Intercept	−2.81	−2.79	−2.80	−2.85	−2.84	−2.84
	(−163.35)	(−173.08)	(−174.92)	(−144.84)	(−157.35)	(−163.56)
Slope	−59.85	−46.24	−53.18	−0.20	−0.14	−0.15
	(−8.42)	(−9.08)	(−11.10)	(−13.62)	(−12.18)	(−13.13)

Table III
Estimates of the Structural Learning Model

This table reports the estimates of the structural learning-based model of mutual funds. The model is estimated using data on 2,599 mutual funds that invest primarily in U.S. common stocks, that first appear in the CRSP data on or after January 1984, and that have at least six months of return data. Each mutual fund's alpha is drawn from a normal distribution with a mean of μ and a variance of σ^2. Each mutual fund generates monthly return observations $\tilde{R}^e_{i,t} \equiv \alpha_i + \tilde{\varepsilon}_{i,t}$, where $\tilde{\varepsilon}_{i,t}$ is normally distributed with a mean of zero and a variance of σ^2_e. The market uses each risk-adjusted monthly return realization to update from the prior distribution to a posterior distribution. The probability that the fund disappears after month t is a function of the posterior distribution $N(m_{i,t}, v_t)$. Investors first compute the amount of mass below some critical level of alpha $\bar{\alpha}$. A fund disappears after period t with probability $\xi(p_t, t)$, where $\xi(\cdot)$ is approximated by $\xi(p_t, t) \approx \hat{\xi}(p_t, t) = \max(\min(\gamma_0 + \gamma_1 \, p_t + \gamma_2 \, p_t^2 + \gamma_3 \, t, 1), 0)$. I estimate the model using a simulated method of moments. The first five moment conditions represent the means of the alpha distributions for the funds that survive through years two, four, six, eight, and 10; the next five moment conditions represent the means of the alpha distributions for funds that disappear in the same two-year periods; and the next five moment conditions indicate the fraction of funds that disappear during these two-year periods. The last two moment conditions are the mean and the variance of the overall alpha distribution. I estimate alphas using at most 10 years of data. All parameters are annualized. I multiply the mean-alpha moment conditions by 100, and the variance moment condition by 10,000. I use the optimal weighing matrix to estimate the model. I estimate the model three times by using the CAPM, the three-factor model, and the four-factor model. Panel A compares simulated moments to sample moments, Panel B reports the (annualized) structural parameter estimates, and Panel C tabulates three alpha distributions for each asset pricing model (the actual distribution in the data, the observed distribution in the model, and the true [unobserved] distribution).

Panel A: Actual Moments versus Simulated Moments

Moment Condition	CAPM		Three-Factor		Four-Factor	
	Data	Model	Data	Model	Data	Model
Alpha Estimates for Disappearing Funds						
Years 1–2	−4.077	−4.531	−3.329	−4.909	−4.210	−5.674
Years 3–4	−4.593	−4.892	−3.956	−4.270	−3.738	−4.203
Years 5–6	−3.286	−3.436	−2.993	−2.995	−2.954	−2.806
Years 7–8	−2.803	−2.543	−2.552	−2.435	−2.660	−2.273
Years 9–10	−1.933	−1.845	−2.264	−1.893	−2.219	−1.806
Alpha Estimates for Surviving Funds						
Years 1–2	−0.073	0.183	0.054	−0.021	0.042	−0.032
Years 3–4	0.552	0.465	0.048	−0.003	0.032	−0.015
Years 5–6	0.627	0.576	0.012	0.003	−0.003	−0.010
Years 7–8	0.751	0.617	−0.009	0.006	−0.021	−0.008
Years 9–10	0.699	0.647	−0.011	0.008	−0.024	−0.006
Fraction of Funds Disappearing						
Years 1–2	0.031	0.030	0.031	0.033	0.031	0.033
Years 3–4	0.065	0.068	0.064	0.068	0.065	0.069
Years 5–6	0.065	0.066	0.062	0.062	0.065	0.064
Years 7–8	0.058	0.061	0.055	0.056	0.058	0.058
Years 9–10	0.065	0.057	0.064	0.053	0.065	0.054
Mean of the Observed Alpha Distribution	−0.562	−0.389	−1.084	−0.777	−1.205	−0.937
Variance of the Observed Alpha Distribution	0.171	0.178	0.140	0.144	0.125	0.129

Panel B: Structural Parameter Estimates

	CAPM		Three-Factor		Four-Factor	
	EST	SE	EST	SE	EST	SE
μ (%)	0.052	(0.052)	−0.383	(0.044)	−0.529	(0.044)
σ (%)	0.249	(0.003)	0.173	(0.001)	0.147	(0.001)
$\bar{\alpha}$ (%)	−2.254	(1.268)	−1.795	(0.162)	−1.779	(0.082)
σ_e (%)	8.496	(2.419)	7.350	(2.119)	6.961	(2.051)
γ_0	0.084	(0.081)	0.270	(0.117)	0.279	(0.116)
γ_1	−4.240	(14.145)	−5.667	(2.964)	−7.969	(3.242)
γ_2	49.393	(363.714)	28.313	(29.563)	55.022	(29.337)
γ_3	−0.012	(0.008)	−0.027	(0.010)	−0.025	(0.009)
J-test, χ^2 (p-value)	12.781	(0.173)	88.312	(0.000)	90.075	(0.000)

Panel C: Observed versus True Alpha Distributions

	CAPM			Three-Factor			Four-Factor		
		Model			Model			Model	
	Data	Obs.	True	Data	Obs.	True	Data	Obs.	True
Mean	−0.709	−0.427	0.052	−1.237	−0.830	−0.383	−1.358	−0.989	−0.529
Percentiles									
1%	−16.208	−16.396	−1.958	−14.198	−15.264	−1.781	−13.499	−14.904	−1.711
5%	−8.147	−9.065	−1.369	−7.623	−8.544	−1.371	−7.317	−8.366	−1.365
25%	−2.577	−2.131	−0.531	−2.988	−2.281	−0.788	−3.069	−2.316	−0.871
50%	−0.540	0.014	0.052	−1.099	−0.431	−0.383	−1.156	−0.569	−0.529
75%	1.618	2.005	0.635	0.624	1.275	0.023	0.526	1.019	−0.186
95%	5.783	5.328	1.474	4.674	4.169	0.606	4.043	3.677	0.307
99%	9.996	13.940	2.063	8.780	12.691	1.015	7.877	11.478	0.654

Table IV

Calibration Results for a Bayesian Portfolio Choice Model with Endogenous Fund Attrition

This table reports the calibration results for a Bayesian portfolio choice model. In this model, an infinitely lived representative mutual fund investor with log-utility can invest in a risk-free asset, the market portfolio, and an actively managed mutual fund. The investor is uncertain about the fund's alpha but updates beliefs each month using Bayes' rule. Each month, the investor can abandon the existing mutual fund by paying a cost of κ. He then draws a new fund from the true alpha distribution and restarts the problem. The mean of the true alpha distribution (μ_α) is -0.5 percent and its standard deviation (σ_α) is 1.25 percent; the expected return of the market portfolio (μ_m) is 10 percent and its standard deviation (σ_m) is 30 percent; the risk-free rate (r_f) is 5 percent; and the investor's discount rate (β) is .9. The standard deviation of the idiosyncratic fund return component ($\tilde{\varepsilon}_{z,t}$) is 8 percent per year. (These are annualized parameter values. Each period in the model is one month long.) Panel A reports average alphas conditional on fund survival. The cost of switching a fund, κ, is set to .064 to match, between the model and the data, the 10-year survival rate of 70 percent. Panel B simulates from the model and reports the observed distribution in the correct model ("biased"), the observed distribution in an alternative model in which funds disappear randomly ("unbiased"), and the actual alpha distribution. The observed distribution in the model is different from the unbiased distribution because of the disappearance of poorly performing funds. The results in this table are based on simulating 100,000 funds through the model. I simulate at most 10 years of monthly returns for each fund.

Panel A: Mutual Fund Alpha Estimates Conditional on Survival

Model Specification	Survive	Years after Inception									
		1	2	3	4	5	6	7	8	9	10
	Yes	-0.50	-0.33	0.05	0.46	0.84	1.14	1.41	1.63	1.82	1.98
	No	-35.54	-21.70	-14.68	-10.88	-8.61	-7.10	-6.03	-5.24	-4.61	-4.12

Panel B: Observed versus True Alpha Distributions

	Biased	Unbiased	True
Mean	-1.131	-0.502	0.000
Percentiles			
1%	-17.908	-12.694	-10.573
5%	-11.617	-9.087	-7.622
25%	-4.857	-3.993	-3.421
50%	-0.524	-0.505	-0.500
75%	2.862	2.986	2.421
95%	7.758	8.079	6.622
99%	11.174	11.784	9.573