



2001

Should Investors Avoid All Actively Managed Mutual Funds? A Study in Bayesian Performance Evaluation

Klaas Pieter Baks
University of Pennsylvania

Andrew Metrick
University of Pennsylvania

Jessica A. Wachter
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/fnce_papers

 Part of the [Finance Commons](#), and the [Finance and Financial Management Commons](#)

Recommended Citation

Baks, K., Metrick, A., & Wachter, J. A. (2001). Should Investors Avoid All Actively Managed Mutual Funds? A Study in Bayesian Performance Evaluation. *The Journal of Finance*, 56 (1), 45-85. <http://dx.doi.org/10.1111/0022-1082.00319>

At the time of publication, author Jessica A Wachter was affiliated with New York University. Currently, she is a faculty member at the Wharton School at the University of Pennsylvania.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/fnce_papers/287
For more information, please contact repository@pobox.upenn.edu.

Should Investors Avoid All Actively Managed Mutual Funds? A Study in Bayesian Performance Evaluation

Abstract

This paper analyzes mutual-fund performance from an investor's perspective. We study the portfolio-choice problem for a mean-variance investor choosing among a risk-free asset, index funds, and actively managed mutual funds. To solve this problem, we employ a Bayesian method of performance evaluation; a key innovation in our approach is the development of a flexible set of prior beliefs about managerial skill. We then apply our methodology to a sample of 1,437 mutual funds. We find that some extremely skeptical prior beliefs nevertheless lead to economically significant allocations to active managers.

Disciplines

Finance | Finance and Financial Management

Comments

At the time of publication, author Jessica A Wachter was affiliated with New York University. Currently, she is a faculty member at the Wharton School at the University of Pennsylvania.



The Rodney L. White Center for Financial Research

*Should Investors Avoid All Actively Managed Mutual Funds?
A Study in Bayesian Performance Evaluation*

**Klaas Baks
Andrew Metrick
Jessica Wachter**

018-99

The Wharton School
University of Pennsylvania

The Rodney L. White Center for Financial Research

The Wharton School
University of Pennsylvania
3254 Steinberg Hall-Dietrich Hall
3620 Locust Walk
Philadelphia, PA 19104-6367

(215) 898-7616

(215) 573-8084 Fax

<http://finance.wharton.upenn.edu/~rlwctr>

The Rodney L. White Center for Financial Research is one of the oldest financial research centers in the country. It was founded in 1969 through a grant from Oppenheimer & Company in honor of its late partner, Rodney L. White. The Center receives support from its endowment and from annual contributions from its Members.

The Center sponsors a wide range of financial research. It publishes a working paper series and a reprint series. It holds an annual seminar, which for the last several years has focused on household financial decision making.

The Members of the Center gain the opportunity to participate in innovative research to break new ground in the field of finance. Through their membership, they also gain access to the Wharton School's faculty and enjoy other special benefits.

Members of the Center

1999 – 2000

Directing Members

**Ford Motor Company Fund
Geewax, Terker & Company
Miller, Anderson & Sherrerd
The New York Stock Exchange, Inc.
Twin Capital Management, Inc.**

Members

**Aronson + Partners
Credit Suisse Asset Management
EXXON
Goldman, Sachs & Co.
Merck & Co., Inc.
The Nasdaq Stock Market Educational Foundation, Inc.
Spear, Leeds & Kellogg**

Founding Members

**Ford Motor Company Fund
Merrill Lynch, Pierce, Fenner & Smith, Inc.
Oppenheimer & Company
Philadelphia National Bank
Salomon Brothers
Weiss, Peck and Greer**

**Should investors avoid all actively managed
mutual funds?
A study in Bayesian performance evaluation**

Klaas Baks

*Department of Economics, Brown University
Providence, RI 02912*

Andrew Metrick

*Department of Finance, University of Pennsylvania and NBER
2300 SH-DH
3620 Locust Walk
Philadelphia, PA 19104-6367*

Jessica Wachter

*Department of Economics, Harvard University
Littauer Center, North Yard
Cambridge, MA 02138*

July 1999

We thank Nick Barberis, Gary Chamberlain, Ken French, Karsten Hansen, Tom Knox, Tony Lancaster, Lubos Pástor, André Perold, Andrei Shleifer, Rob Stambaugh, René Stultz, an anonymous referee, and seminar participants at Columbia, Wharton, and the NBER for helpful comments.

Should investors avoid all actively managed mutual funds?
A study in Bayesian performance evaluation

Abstract

This paper analyzes mutual-fund performance from an investor's perspective. We study the one-period portfolio allocation problem for a mean-variance investor choosing from a risk-free asset, benchmark assets (passively managed index funds) and non-benchmark assets (actively managed mutual funds). To solve this problem, we propose and employ a Bayesian method of performance evaluation; the main innovation in our approach is the development of a flexible set of prior beliefs about managers' abnormal performance ("alphas"). We motivate this Bayesian approach by demonstrating unrealistic results for an investor who ignores prior beliefs and relies only on the data. We then apply our methodology to a sample of domestic diversified equity mutual funds and ask "what prior beliefs would imply zero investment in active managers?" In this sample, it is not possible to reject the null hypothesis that the best performance is due to chance. Nevertheless, we find that the policy of zero investment in active managers can only be supported by extremely skeptical prior beliefs about the probability of skill; such extreme skepticism could not possibly be "proved" using current methods and data.

Actively managed equity mutual funds have trillions of dollars in assets, collect tens of billions in management fees, and are the subject of enormous attention from investors, the press, and researchers. For years, many experts have been saying that investors would be better off in low-cost passively managed index funds. Notwithstanding the recent growth in index funds, active managers still control the vast majority of mutual-fund assets. Are any of these active managers worth their added expenses? Should investors avoid all actively managed mutual funds?

Since Jensen (1968), most studies have found that the universe of mutual funds does not outperform its benchmarks after expenses.¹ This evidence indicates that the average active mutual fund should be avoided. On the other hand, recent studies have found that future abnormal returns (“alphas”) can be forecast using past returns or alphas,² past fund inflows,³ and manager characteristics such as age, education, and SAT scores.⁴ Given this evidence, it is possible that alphas are persistent, and that some managers have positive “expected” alphas. Perhaps 0.1 percent of all managers have positive expected alphas. Perhaps none do. Using current data and methods, tests that distinguish between these two possibilities will have very low power. Nevertheless, such small differences may have large consequences for investors.

¹ Recently, Malkiel (1995), Carhart (1995) and Daniel et al. (1997) all find small or zero average abnormal returns by using modern performance-evaluation methods on samples that are relatively free of survivorship bias.

² Carlson (1970), Lehmann and Modest (1987), Grinblatt and Titman (1988, 1992), Hendricks, Patel, and Zeckhauser (1993), Goetzmann and Ibbotson (1994), Brown and Goetzmann (1995), Elton, Gruber and Blake (1996), Carhart (1997).

³ Gruber (1996), Zheng (1999).

⁴ Golec (1996), Chevalier and Ellison (1999).

In this paper, we explore these consequences by explicitly taking an investor’s perspective. We study the one-period portfolio allocation problem for a mean-variance investor choosing from a riskless asset, benchmark assets (passively managed index funds), and non-benchmark assets (actively managed mutual funds). We propose and employ a Bayesian method of performance evaluation; the main innovation in our approach is the development of a flexible set of prior beliefs about alphas that are consistent with intuition about managerial skill. In this framework, the prior probability of managerial skill can be made arbitrarily small (or zero), so investors can interpret the results filtered through their own beliefs.

Our approach is similar to several recent papers that take an investment perspective and use prior beliefs centered on an economic model. (Kandel and Stambaugh (1996), Pástor (1999), Pástor and Stambaugh (1999b), and MacKinlay and Pástor (1999)). Like the latter three papers, our techniques build upon the work of Pástor and Stambaugh (1999a); in particular, our definition of “skill” among managers plays the same mathematical role as “model mispricing” does in their analysis.⁵

In Section I, we formally pose the investor’s problem and discuss the conditions under which there is positive investment in an active manager. This exercise shows that an investor who relies only on the data would choose to invest in an active manager whenever the point estimate of alpha is greater than zero. This result seems contrary to most investment advice

⁵ There is a related literature that employs Bayesian methods to explore the role of estimation risk on portfolio choice. See Barberis (1999), Bawa, Brown and Klein (1979), Brown (1979), Frost and Savarino (1986), Jobson and Korkie (1980), Jobson, Korkie and Ratti (1979), Jorion (1985, 1986, and 1991) and Klein and Bawa (1976). In these applications, however, prior beliefs about parameters are typically noninformative or come from empirical Bayes procedures. Also, another related line of research focuses on the role of prior beliefs in model testing. See Kandel, McCulloch, and Stambaugh (1995), McCulloch and Rossi (1990), and Shanken (1987).

about active management and motivates the use of informed prior beliefs about the frequency and magnitude of manager skill. We then posit a flexible functional form for these beliefs.

Once prior beliefs have been specified, the investor's decision reduces to a Bayesian inference problem, which we solve in Section II. Using prior beliefs motivated in Section I, we derive an analytical solution for the posterior expectation of alpha. Our solution is expressed as a formula whose inputs are modified moments of well-known distributions. While our focus is on mutual-fund managers, this formula can also be applied to managers or portfolio strategies in other contexts. It can be applied to a single manager in isolation and does not require a comprehensive or bias-free database. In each context, the prior beliefs may be different, and rightly so. We also show how prior beliefs can be elicited by intuitive questions such as "what is the probability that a manager has an expected alpha greater than 25 bp per month?", and we map the answers to these questions into the parameters of the prior belief distribution.

Section III applies our methodology to an investor's choice over a large set of equity-mutual-fund managers. We use a sample of 1437 domestic diversified equity funds in existence at the end of 1996, and look at the full return history for the managers in place at that time. Using the three-factor model of Fama and French (1993), we calculate the posterior expectation of alpha for each manager over a wide range of prior beliefs. We then ask, "what prior beliefs would imply zero investment in active managers?" To justify such a zero-investment strategy, we find that a mean-variance investor would require extremely skeptical beliefs about the possibility of managerial skill. We then discuss how to recon-

cile the frequentist and Bayesian evidence for this sample. Section IV concludes with an interpretation of our results.

I. The investor’s problem and prior beliefs

Consider a mean-variance investor choosing from a risk-free asset, a set of K benchmark assets (passively managed index funds) and a single non-benchmark asset (an active investment manager). Assume that this manager has been randomly selected from a bias-free sample. Under what conditions will the investor place any of her portfolio in the active manager? In this section, we derive the necessary condition for positive investment and show that an investor who relies only on the data will arrive at a counterintuitive outcome. This motivates a Bayesian performance-evaluation approach with the use of informed prior beliefs about manager skill. We then propose a flexible and reasonable form for these prior beliefs.

Let F and r denote the excess returns on the index funds and active manager, respectively. Let w be the weight on the active manager, with its optimal level written as w^* . Next, define the “performance-evaluation equation” for the active manager as

$$r = \alpha + \beta F + \varepsilon, \tag{1}$$

where ε is distributed $N(0, \sigma^2)$. In our analysis, we treat α and β as fixed parameters that “belong” to a manager, and not to the mutual fund that he manages. In principle, the assumption of fixed parameters can be relaxed and they can be allowed to vary over time and with the characteristics of the manager’s portfolio.

Under some conditions (see Appendix C), one can adopt a Bayesian procedure for estimating α and solve for w^* as

$$w^* = c_1 \tilde{\alpha}, \tag{2}$$

where $\tilde{\alpha}$ is the posterior mean for α and $c_1 > 0$. Thus, the decision rule is to invest in the active manager if and only if the posterior mean of α is positive. In this case, we have reduced the decision rule to a Bayesian inference problem on α . For example, if the investor has diffuse (noninformative) prior beliefs for these parameters, then posterior beliefs would be completely determined by the data and we could write

$$w^* = c_2 \hat{\alpha}, \tag{3}$$

where $\hat{\alpha}$ is the frequentist OLS estimate (= posterior mean) for α and $c_2 > 0$. Thus, an investor with diffuse prior beliefs would invest with the active manager as long as $\hat{\alpha}$ is positive.

This result seems counterintuitive – with diffuse prior beliefs, any positive realization for $\hat{\alpha}$ would lead to positive investment in the active manager. While the “statistical significance” of $\hat{\alpha}$ would play a role in the calculation of c_2 and the level of w^* , it does not affect the sign of w^* . Moreover, it is easy to construct examples where $\hat{\alpha}$ is not statistically significant but w^* is economically large. We believe that the reason this result seems perverse is that few investors actually have diffuse prior beliefs about α . For example, a typical argument against active management stresses the expenses of active management and the rarity (or nonexistence) of skilled managers. This argument implies prior beliefs that most managers have *negative* expected alphas (after expenses) and that most positive realizations of $\hat{\alpha}$ are

due to luck. An investor that implicitly has such prior beliefs should explicitly use them, or she risks arriving at incorrect and suboptimal outcomes. This is the primary motivation for our Bayesian performance evaluation methods.

In this paper, we explicitly model such prior beliefs for α . Figure 1 gives a graphical representation. For now, we consider the case where the variance of ε is known, so that the manager's level of residual risk is held constant.⁶ The prior separates managers into two types: skilled (with probability q) and unskilled (with probability $1 - q$). The key features of the distribution are the lower bound and point mass of unskilled managers at $\underline{\alpha} < 0$, and the right tail of a normal distribution as the functional form for α among skilled managers. The normal distribution is chosen for analytic tractability, but is also a useful approximation for skill distributions in many other contexts. The parameters q and σ_α allow the investor great latitude in her beliefs about managerial skill. If she is very skeptical about the frequency and magnitude of skill, then she would set both parameters to be small. In the limit, either $q = 0$ or $\sigma_\alpha = 0$ implies no possibility of skill.

The point mass in Figure 1 occurs at $\underline{\alpha} = a - fee - cost < 0$. Here, we set a so that $E(\alpha) = -fee - cost$. This restriction forces the average α , before fees and costs, to be zero. With these prior beliefs, all abnormal returns earned by skilled managers must come at the expense of their unskilled counterparts. Thus, an unskilled manager is expected to earn a negative α that consists of three components: a , his losses due to transactions with skilled managers, fee , his total fees, and $cost$, his transactions costs. While a is the same for all

⁶ In Section II, we introduce a link between manager's residual risk and prior expectations for α .

managers, the other components, *fee* and *cost*, will generally differ across managers.

Why do we assume a lower bound at $\underline{\alpha}$? Under some interpretations of equation (1), this assumption is logical. For example, if the market is semi-strong efficient (Fama (1970)) with respect to (1), where the index funds are interpreted as a complete set of risk factors, then no manager should be expected to have an α below $\underline{\alpha}$. While many managers will have *realizations* below this level, one would need to be systematically trading on non-public “misinformation” in order to have an *expected* α below $\underline{\alpha}$. If, instead, the market is not semi-strong efficient and managerial skill is based upon the use of public information to exploit decision-making biases, then differential incidence of such biases would result in a some prior mass below $\underline{\alpha}$.⁷ By imposing a lower bound at $\underline{\alpha}$, we are assuming that any behavioral biases are evenly distributed among all managers. While this assumption will affect inference for the worst-performing managers, it should have little effect on posterior beliefs for the best-performing ones.

Once prior beliefs have been specified, the next step is to combine these beliefs with data and compute a posterior estimate for α . This posterior estimate can then be substituted into (2) to obtain the investor’s solution. We solve this Bayesian inference problem in the next section.

⁷ One way to model this possibility would be to make the prior distribution for α symmetric around $\underline{\alpha}$. In fact, this assumption would greatly simplify our analysis. We do not use a symmetric distribution because we find such prior beliefs to be implausible. A symmetric prior distribution for α would imply that for every skilled manager with superior judgement or the ability to exploit the behavioral anomalies of other investors, there is another manager who systematically does the opposite.

II. Bayesian performance evaluation and portfolio choice

This section, along with the Appendix, provides the details of our methodology. In Section II.A, we provide the likelihood function for a general (unconditional) factor representation of manager returns. Section II.B gives a complete mathematical representation for prior beliefs, and Section II.C poses four questions sufficient to elicit this representation. In Section II.D, we combine these prior beliefs with the likelihood function and derive an analytical solution for the posterior expectation of α . Taken together, Sections II.A through II.D solve the inference problem for a single manager studied in isolation. Section II.E extends the portfolio-choice analysis to multiple managers. Finally, Section II.F discusses the implications of survivor bias for our analysis.

A. Likelihood

Let r denote a $T \times 1$ vector of excess returns for a manager and F a $T \times K$ matrix of factor returns. The regression disturbance ε in (1) is assumed to be a serially uncorrelated, homoskedastic realization from a normal distribution, with zero mean and variance equal to σ^2 . Then, we write the likelihood for r conditional on F as

$$p(r \mid \alpha, \beta, \sigma^2, F) = N(\alpha \iota_T + F\beta, \sigma^2 I_T), \quad (4)$$

where ι_T is a T -vector of ones, and I_T is a $T \times T$ identity matrix. Thus, manager returns conditional on factor returns are normally distributed and have a standard factor structure. We assume that the factors F do not depend on α , β or σ , so the exact specification of the factor likelihood is not necessary for our analysis in this section.

B. Prior beliefs

The next step is to state the prior beliefs for the parameters in (4). As discussed in the previous section, managers are either skilled or unskilled. These two states of the world are indexed by the state variable Z , with $Z = 1$ denoting the skilled state and $Z = 0$ denoting the unskilled state. The probability of the skilled state is q . In the absence of data on the returns of a manager, the probability that the manager has skill is simply given by the probability that skill exists in the population, or $P(Z = 1)$. We assume that only α depends on whether or not the manager is skilled; the factor loadings, β , and residual risk, σ , do not. Under these assumptions, the prior distribution can be written as

$$p(\alpha, \beta, \sigma^2) = [p(\alpha | Z = 0) P(Z = 0) + p(\alpha | Z = 1) p(Z = 1)] p(\beta, \sigma^2). \quad (5)$$

We use a diffuse prior on β and σ^2 (Gelman et al. (1995)):

$$p(\beta, \sigma^2) \propto \frac{1}{\sigma^2}. \quad (6)$$

The diffuse prior on β is necessary in order to obtain analytical results. It is also a reasonable starting point when analyzing managed portfolios, where β can be estimated relatively precisely (as compared to the β of individual stocks). The diffuse prior on σ^2 is not necessary for analytical results, but it simplifies notation and allows us to focus our attention on the role played by α . Appendix A relaxes this second assumption and solves for the posterior when σ^2 has an informative prior; then, the diffuse prior used here becomes a limiting case.

We turn next to α , the main parameter of interest. Essentially, we want to write down a mathematical representation of Figure 1. The one additional twist is to recognize that prior

beliefs for α should be conditioned on some level of residual risk. For example, consider a fully-invested manager who has an α of α_j and is taking on s units of residual risk. Then, if this manager were to take on a fifty-percent cash position, his residual risk would decrease to $s/2$ and his α would fall to $\alpha_j/2$. The full specification of the prior for α recognizes this relationship and is written as

$$P(Z = 1) = q, \tag{7}$$

$$P(Z = 0) = 1 - q, \tag{8}$$

$$p(\alpha | Z = 0, \sigma^2) = \delta_{\underline{\alpha}}, \tag{9}$$

$$p(\alpha | Z = 1, \sigma^2) = 2N\left(\underline{\alpha}, \sigma_\alpha^2 \left[\frac{\sigma^2}{s^2}\right]\right) \mathbf{1}_{\alpha > \underline{\alpha}}, \tag{10}$$

where δ_x is the Dirac delta function with mass point at x , $\mathbf{1}_X$ is the indicator function for the set X , $\underline{\alpha}$ is a negative constant representing the expected α for an unskilled manager, and s^2 is an arbitrary constant specified by the researcher before priors are elicited. Finally, we assume that the parameters of the (unspecified) factor prior are independent of α , β , and σ .

The ratio σ^2/s^2 effectively links the posterior distributions of σ and α . As discussed above, this link allows us to adjust for the fact that a skilled manager can control his expected α through the strategic use of leverage.⁸ The importance of this relationship becomes clearer in the next section when we discuss the elicitation of priors.

⁸ The same argument can be used to motivate a link between β and σ . Since we use a diffuse prior for β , such a link is not applicable. To be completely consistent, we should also link our prior beliefs for a (and, by extension, $\underline{\alpha}$) to the ratio σ^2/s^2 . Unfortunately, this link is not tractable. As long as a is small, however, the omission will not be quantitatively important.

The prior link between α and σ is first suggested by MacKinlay (1995) and is implemented in Pástor (1999), Pástor and Stambaugh (1999a), Pástor and Stambaugh (1999b), and MacKinlay and Pástor (1999). Mathematically, our link is identical to theirs, although their motivation is somewhat different. In these papers, σ_α is an index of potential “mispricing”, and the motivation for the link is to reduce the ex-ante probability of very high Sharpe ratios among portfolios that combine benchmark and non-benchmark assets.

C. Elicitation of prior beliefs

It is possible to elicit prior beliefs using straightforward questions about performance, fees, and transaction costs. For example, consider any specific factor representation for (1). Then, given this factor representation, assume that the manager under study has a residual variance, σ^2 , equal to a specific value. We call this level s^2 , and it serves as the constant denominator term in (10). Then, conditional on $\sigma^2 = s^2$ and the chosen factor representation, the researcher should answer the following four questions:

- Question 1: What is the probability that the manager is skilled? (i.e., that he has an α greater than would be earned by randomly selecting stocks while incurring the same fees and costs.) [Call this answer q .]
- Question 2: What is the probability that the manager has an α greater than 25 bp per month? [Call this answer $q(25)$.]
- Question 3: What are the expected fees for the manager? [Call this answer fee .]
- Question 4: What are the expected costs for the manager? [Call this answer $cost$.]

In answering these questions, the researcher should not consider any return-based information about the manager that coincides with the sample period under study; this forces the exclusion of all information about the length of time the manager has survived or the level of assets that he has under management, both of which will tend to be correlated with past returns. Instead, the answers should reflect a thought experiment about a new manager before any return information has been observed.

Note that Question 2 is not asking about the probability of different realizations of $\hat{\alpha}$, but about “true” values of α . Realizations of $\hat{\alpha}$ will depend on sampling variability. The true α , on the other hand, is the $\hat{\alpha}$ we would expect as the number of time periods goes to infinity. Thus, if a researcher believes that no managers are skilled, then $q = 0$, and thus $q(25)$ should be 0 as well. The use of 25 bp in this question is arbitrary, and any other point in the distribution could be substituted.

Question 2 is conditioned on a specific level of residual variance, s^2 . This conditioning is crucial, and we believe that prior beliefs on α are not well-defined without it. The same reasoning discussed in Section II.B also applies here: if a manager has an expected α of α_j when his residual standard deviation is s , then his expected α would be $\alpha_j/2$ if he levered down his portfolio and took on only $s/2$ units of residual risk. By including the σ^2/s^2 term in the prior beliefs for α , we link our beliefs for α and σ in a way consistent with our elicitation procedures. In this respect, the prior beliefs elicited through these questions are not really about α , but are instead about Sharpe ratios for combinations of the manager

and the benchmarks.⁹ This returns us to the original motivation for the link as given by MacKinlay (1995).

Given the answers to these questions, we can solve for the remaining parameters of the prior belief distribution. Let $\Phi(x)$ denote the cdf of a standard normal distribution evaluated at x . Then, we have three equations,

$$q(25) = P(\alpha > 25 \mid \sigma^2 = s^2) = 2q \left(1 - \Phi \left(\frac{25 - \underline{\alpha}}{\sigma_\alpha} \right) \right), \quad (11)$$

$$a = -q\sigma_\alpha \sqrt{\frac{2}{\pi}}, \quad (12)$$

and

$$\underline{\alpha} = a - fee - cost, \quad (13)$$

which we can solve for the three unknowns, a , $\underline{\alpha}$ and σ_α . (11) relates $q(25)$ to σ_α (given the other parameters), (12) imposes the constraint¹⁰ that the expectation of α , conditional on $\sigma^2 = s^2$, is equal to $-fee - cost$, and (13) is just the definition of $\underline{\alpha}$. Table I illustrates some solutions to this system: given inputs of q , $q(25)$, fee , and $cost$, we provide the solutions for σ_α , a , and $\underline{\alpha}$. In most of the examples in the table, a tends to be very small, so $\underline{\alpha}$ is close to $-fee - cost$. Alternatively, one can elicit $q(25)$ before fees. To compute the prior parameters under this alternative method, we just replace $\underline{\alpha}$ with $a - cost$ on the right-hand-side of (11). This is the elicitation method used in Section III.

⁹ We are grateful to Rob Stambaugh for suggesting this interpretation.

¹⁰ In some applications, one may wish to relax the constraint in (12) and allow certain types of managers to have a positive expectation of α . This case is solved in a previous version of the paper (Baks, Metrick, and Wachter (1999)).

D. Posterior beliefs

Our goal in this section is to calculate the mean of the posterior distribution for α . Most of the intuition for this solution is contained in Figure 2 and its corresponding notation as developed in (14) - (22). The details are given beginning with (23) and in Appendix A.

We denote the mean of the posterior distribution for α , $E[\alpha | r, F]$, as $\tilde{\alpha}$. Similarly, we denote the posterior probability that a manager is skilled, $P(Z = 1 | r, F)$, as \tilde{q} . Then, it follows that we can write $\tilde{\alpha}$ as

$$\tilde{\alpha} = \tilde{q}E[\alpha | Z = 1, r, F] + (1 - \tilde{q})\underline{\alpha}. \quad (14)$$

The first term on the right-hand-side of (14) reflects the contribution to the posterior mean coming from the possibility that the manager has skill; the posterior probability of skill (\tilde{q}) is multiplied by the posterior expectation of α conditional on skill. The second term on the right-hand-side of (14) reflects the contribution coming from the possibility that manager is unskilled; the posterior probability $(1 - \tilde{q})$ is multiplied by $\underline{\alpha}$. Note that our procedure is not updating q for the “population”; q is treated as *known*. Rather, our inference problem only concerns the specific manager under study.

To calculate $\tilde{\alpha}$, we need to solve for the two unknown elements on the right-hand-side of (14): $E[\alpha | Z = 1, r, F]$ and \tilde{q} . The problem of computing $\tilde{\alpha}$ is thus considered in two parts. First, we calculate $E[\alpha | Z = 1, r, F]$, the expectation conditional on skill. Second, we calculate \tilde{q} , the posterior probability that the manager is skilled.

In expositing our solution, it is helpful to introduce some notation:

$$X \equiv \begin{pmatrix} \iota_T & F \end{pmatrix}, \quad (15)$$

$$\begin{pmatrix} \hat{\alpha} \\ \hat{\beta} \end{pmatrix} \equiv (X'X)^{-1} X'r, \quad (16)$$

$$m \equiv \text{top left element of } (X'X)^{-1}, \quad (17)$$

$$\text{var}(\hat{\alpha}) \equiv m\sigma^2. \quad (18)$$

Then, the posterior of α conditional on the variance σ and $Z = 1$, which we call the “skilled posterior distribution”, is given by a truncated normal distribution:¹¹

$$p(\alpha \mid Z = 1, r, F, \sigma^2) \propto N(\alpha', \sigma'^2) \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (19)$$

where

$$\alpha' = \lambda \hat{\alpha} + (1 - \lambda) \underline{\alpha}, \quad (20)$$

$$\sigma'^2 = \left(\frac{1}{\text{var}(\hat{\alpha})} + \frac{1}{\sigma_\alpha^2 \left(\frac{\sigma^2}{s^2} \right)} \right)^{-1}, \quad (21)$$

$$\lambda = \frac{\sigma'^2}{\text{var}(\hat{\alpha})}. \quad (22)$$

Equations (19) - (22) are illustrated graphically in Figure 2. α' is the mode of the skilled posterior distribution given in (19); it would also be the mean, and σ'^2 the variance, of the untruncated version of this distribution. α' is written in (20) as a weighted average of the maximum likelihood estimate ($\hat{\alpha}$) and the prior mode ($\underline{\alpha}$), with weights given by λ and

¹¹ See Appendix A for the details of these calculations.

$1 - \lambda$, respectively. In (21), $\text{var}(\hat{\alpha})$ represents the variance (in a frequentist sense) of the maximum likelihood estimate for α , conditional on a known residual variance of σ^2 . The posterior precision, $1/\sigma'^2$, is the sum of the precision of the prior and the precision of the data. Intuitively, this says that after having observed the data, there is greater certainty about the location of the posterior distribution of α than there was for the prior. Thus, the weight λ is determined in (22) by the relative precision of prior beliefs versus sample information. The greater the precision of $\hat{\alpha}$, the more the mode is shifted towards $\hat{\alpha}$ and away from the prior mode $\underline{\alpha}$.

The marginal posterior for α (conditional on skill) can be obtained in closed form by successively integrating out β and σ from the joint posterior. Once β is integrated out, the model resembles one where normal data is combined with a conjugate prior. Therefore, familiar techniques (see, e.g. Gelman et al. (1995)) can be used to integrate out σ , suitably adjusted to reflect the truncation at $\underline{\alpha}$. The marginal distribution is then given by:

$$p(\alpha | Z = 1, r, F) \propto t_\nu \left(\alpha', \frac{\lambda m h}{\nu} \right) \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (23)$$

where

$$h = S + \frac{1 - \lambda}{m} (\hat{\alpha} - \underline{\alpha})^2, \quad (24)$$

$$\hat{\theta} = \begin{pmatrix} \hat{\alpha} \\ \hat{\beta} \end{pmatrix}, \quad (25)$$

$$S \equiv (r - X\hat{\theta})' (r - X\hat{\theta}), \quad (26)$$

$$\nu = T - K. \quad (27)$$

The parameter λ has the same interpretation here as in (22); the greater the precision of the data relative to the prior precision, the more the posterior mode is shifted towards the maximum-likelihood estimate.

The posterior expectation of α in the skilled state can then be calculated as

$$E(\alpha \mid Z = 1, r, F) = \alpha' + \frac{\lambda mh}{\nu - 2} t_{\nu-2} \left(\underline{\alpha}; \alpha', \frac{\lambda mh}{\nu - 2} \right) \frac{1}{\int_{\underline{\alpha}}^{\infty} t_{\nu} \left(\alpha; \alpha', \frac{\lambda mh}{\nu-2} \right) d\alpha}. \quad (28)$$

The first term on the right-hand-side of (28) is just the mode of the skilled posterior distribution (the mean of the untruncated t -distribution), while the second term is an adjustment for the truncation at $\underline{\alpha}$. We use the notational convention that $t_{\nu}(\alpha; x, y)$ is a t -distribution evaluated at α with mean x , variance y and ν degrees of freedom.

We next solve for the second unknown element in (14), \tilde{q} , the posterior probability that the manager is skilled. From Bayes' formula for binomial variables, it follows that

$$\tilde{q} \equiv P(Z = 1 \mid r, F) = \frac{p(r \mid Z = 1, F) q}{p(r \mid Z = 1, F) q + p(r \mid Z = 0, F) (1 - q)}. \quad (29)$$

Dividing through by $p(r \mid Z = 1, F)$ yields

$$\tilde{q} = \frac{q}{q + \frac{1-q}{B}}, \quad (30)$$

where B is given by

$$B \equiv \frac{p(r \mid Z = 1, F)}{p(r \mid Z = 0, F)}. \quad (31)$$

If observing the realized data is equally likely whether the manager is skilled or unskilled, then $B = 1$, and our posterior probability of $Z = 1$ equals the prior probability: $\tilde{q} = q$. The

more likely the data are for a skilled manager relative to an unskilled manager, the higher is B and thus, the higher is \tilde{q} .¹²

As shown in Appendix A,

$$B = \frac{t_{\nu-1} \left(\underline{\alpha}; \hat{\alpha}, \frac{mS}{(1-\lambda)(\nu-1)} \right)}{t_{\nu-1} \left(\underline{\alpha}; \hat{\alpha}, \frac{S}{\nu-1} m \right)} \left(2 \int_{\underline{\alpha}}^{\infty} t_{\nu} \left(\alpha; \alpha', \frac{\lambda mh}{\nu} \right) d\alpha \right). \quad (32)$$

This is the ratio of two t -distributions, multiplied by a term to correct for the truncation at $\underline{\alpha}$.

Finally, with values for $E(\alpha | Z = 1, r, F)$ (equation (28)) and \tilde{q} (equations (30) and (32)), we can substitute into (14) and obtain a solution for $\tilde{\alpha}$.

E. Portfolio choice over multiple managers

The previous analysis applies to an investor with a choice of one manager and K index funds. In practice, investors can choose among many managers. This section gives the assumptions that allow our framework to extend to the case of multiple managers. This is necessary for the application in Section III.

Consider a mean-variance investor choosing among the K index funds, N manager portfolios, and a riskless asset. Let r_j denote the $T \times 1$ vector of returns on manager j , and let r denote the $T \times N$ matrix of returns on all the managers. As in Section II.A, the likelihood for returns is given by:

$$p(r_j | \alpha_j, \beta_j, \sigma_j^2, F) = N(\alpha_j \iota_T + F\beta_j, \sigma_j^2 I_T). \quad (33)$$

¹² Note that by dividing the numerator and denominator of (30) by q , the the posterior odds ratio $\frac{1-q}{q} \frac{p(r|Z=0,F)}{p(r|Z=1,F)}$ for testing the hypothesis $H_0 : q = 0$ versus $H_1 : q > 0$ appears. Thus, B is a ‘‘Bayes factor’’ associated with testing H_0 versus H_1 .

To shorten notation, let $\phi_j = (\alpha_j, \beta_j, \sigma_j^2)$, and $\Phi = (\phi_1, \dots, \phi_N)$. For each manager, the prior on the parameters, $p(\phi_j)$, is given in Section II.B.

The key assumption that allows us to extend our result is that no manager conveys information on any other manager. Specifically, we require that the parameters for the managers are independent in the prior, and that the likelihoods for each manager are independent:

Assumption 1 : The likelihoods are independent across managers:

$$p(r \mid \Phi, F) = \prod_j p(r_j \mid \Phi, F) \quad (34)$$

$$= \prod_j p(r_j \mid \phi_j, F). \quad (35)$$

Assumption 2 : The priors are independent across managers:

$$p(\Phi) = \prod_j p(\phi_j). \quad (36)$$

For some applications, the assumptions above might be problematic. Here, we feel that they are innocuous. Independence across managers would be dangerous assumption if we intended to make strong statements about the total fraction of a portfolio invested in active managers. If we restrict ourselves to statements about whether any of a portfolio is in active managers, then inference about covariances is much less important. There is no reason to believe that ignoring the information from other managers would bias the results in either direction.

Assumptions 1 and 2 imply that the posterior distributions across managers are independent:

$$p(\Phi \mid r, F) \propto p(r \mid \Phi, F)p(\Phi)$$

$$\begin{aligned}
&= \prod_j p(r_j | \phi_j, F) p(\phi_j) \\
&\propto \prod_j p(\phi_j | r_j, F).
\end{aligned} \tag{37}$$

Therefore, the calculation for the posterior of α_j when there are multiple managers is identical to that in II.D for a single manager.

More to the point, the zero-investment condition for multiple managers is analogous to the zero-investment condition for each manager. An outline of the proof is given here; details can be found in Appendix C. First, it simplifies the analysis to consider portfolios which offer a pure play on the manager’s alpha. That is, consider the portfolio that consists of going long one manager and short the index funds with weights equal to the posterior betas $\tilde{\beta}_j$. Call this portfolio an “alpha” portfolio. It is clear that the portfolio weight on an alpha portfolio must equal the weight on the corresponding manager. It is shown in Appendix C that the alpha portfolios have expected return $\tilde{\alpha}$.

Let \tilde{V} denote the posterior variance-covariance matrix of the N alpha portfolios and the K index funds, and let \tilde{E} denote the vector of means. Then the optimal weights are given by:

$$\begin{pmatrix} w^* \\ x^* \end{pmatrix} = \frac{\tilde{V}^{-1} \tilde{E}}{\text{normalizing constant}}, \tag{38}$$

where w^* is the vector of weights on the alpha portfolios and x^* is the vector of weights on the index funds. In Appendix C, it is shown that V is of the form:

$$\begin{pmatrix} V_{11} & 0 \\ 0 & V_{22} \end{pmatrix}. \tag{39}$$

and that V_{11} is diagonal. Therefore, the vector of weights on the managers is given by:

$$w^* = c_4 \tilde{\alpha}, \tag{40}$$

where $c_4 > 0$. That is, the investor puts positive weight on a manager, if and only if the posterior expectation of α is greater than zero.

F. Survivor bias

A possible objection to our framework is that it fails to recognize the possibility of survivor bias. The investor sees only the fund managers that “survive”, i.e. do not leave the sample. The question is, does this change the inference problem *for the managers that survive?*

Survivor bias can impact the analysis in two ways. First, the fact that poorly performing managers are not observed could, in principle, affect the posterior distribution of a manager that is observed. Under the independence assumptions of Section II.E, this first kind of survivor bias is not a problem. Second, knowing that the manager in question has survived might impact the posterior for that manager. This second type of survivor bias is also not a problem, as this section demonstrates.

We represent survival for manager j by a binary random variable $survival_j$. The question is whether

$$p(\phi_j | r_j, F) = p(\phi_j | r_j, F, survival_j) \quad ? \tag{41}$$

That is, does inference on ϕ_j change if conditioned on $survival_j$? The answer to the question is “no”, under the following reasonable assumption on conditions for survival:

Assumption 3 : $p(\text{survival}_j | r_j, F, \phi_j) = p(\text{survival}_j | r_j, F)$.

Assumption 3 states that survival depends only on realized returns. Conditional on realized returns, the manager's skill (and the parameters β and σ), do not affect the probability of survival. Realized returns are, of course, observable, while ϕ is unknown. It is quite plausible that survival depends on the manager's observed success, not on unobserved skill. In what follows, we suppress the j subscript. The discussion easily extends to the case of multiple managers, using the posterior independence shown in II.E.

Using Assumption 3, it follows that survivor bias is not a problem for our analysis. In particular, by Bayes' rule:

$$p(\phi | r, F, \text{survival}) = \frac{p(\text{survival} | r, F, \phi)p(\phi | r, F)}{p(\text{survival} | r, F)} \quad (42)$$

$$= p(\phi | r, F). \quad (43)$$

The intuition behind this result is that the returns are already observed, so there is no additional information in return-based survival. Note that in general:

$$p(r | F, \phi, \text{survival}) \neq p(r | F, \phi). \quad (44)$$

That is, the likelihood conditional on survival is not the same as the likelihood without conditioning on survival. The prior on ϕ conditional on survival will also differ from the unconditional prior. Equations (42) and (43) demonstrate that the effect on the prior and the likelihood must exactly cancel, and thus the posterior remains the same.

It is helpful to contrast our setting to those where survivor bias would be a problem.¹³ If, for example, we were performing inference on the prior parameter q , then knowing that poorly performing funds disappear would change our posterior. For this reason, we do not attempt to perform inference on q , and instead treat it as fixed in each implementation. Then, we can map different “known” levels of q into different posterior beliefs. Alternatively, if we were missing data on a particular fund manager in years where that fund manager did particularly poorly, this would also bias our conclusions. But this is not the case in our sample.

III. Should an investor avoid all actively-managed mutual funds?

In this section, we apply our methodology and ask, “given the evidence, what prior beliefs would induce positive investment in at least one active mutual-fund manager?” Section III.A discusses the data and performance-evaluation regression. Section III.B summarizes the frequentist evidence for this sample. Section III.C contains the main analysis and answers the question posed in the title of the paper. Section III.D discusses several criticisms and limitations of the analysis. Section III.E contrasts the frequentist and Bayesian results.

A. Setup

Our data is drawn from the Center for Research in Security Prices (CRSP) mutual-fund database (CRSP, 1999). This database includes information collected from several sources and is designed to be a comprehensive sample of all mutual funds from 1963 to 1996. We

¹³ For studies of the implications for survivor bias on inference, see Brown et al. (1992), Brown, Goetzmann, and Ross (1995), Goetzmann and Jorion (1997)

restrict ourselves to the subset of funds still operating at the end of 1996, and only include returns that have been earned by current (as of December 1996) managers. We include only the returns earned by current managers because we interpret α as a fixed parameter that is a characteristic of managers, not of funds. In the remainder of the section, we use the terms “fund” and “manager” interchangeably. We include team-managed funds only if a name is provided for at least one member of the team; returns for such funds are included for the tenure of the team’s longest-standing member. Furthermore, we restrict ourselves to funds with at least one complete year of return history. The resulting sample includes 1437 funds with an average of 51 months of returns. This sample suffers from survivor bias and will not be representative of mutual-fund performance as a whole. As discussed in Section II.F, the assumption that survival is based only on observed returns allows us to ignore survivor issues in our Bayesian analysis.

The next step is to choose a set of benchmarks for the evaluation. For conciseness, we restrict our presentation to a single well-known model – the three-factor model of Fama and French (1993).¹⁴ The model is given by

$$r_{jt} = \alpha_j + \beta_1 RMRF_t + \beta_2 SMB_t + \beta_3 HML_t + \varepsilon_{jt}, \quad (45)$$

where r_{jt} is the excess return to fund j in year t , α_j is the performance measure, and $RMRF_t$, SMB_t , and HML_t are the time t returns to benchmark portfolios constructed using market,

¹⁴ The qualitative results do not change if we use the CAPM or the four-factor model of Carhart (1997).

size, and value strategies.¹⁵ While there is an ongoing debate about whether these factors are proxies for risk, we take no position on this issue and simply view the three-factor model as a method of performance attribution. Thus, we interpret the estimated alphas as abnormal returns in excess of what could have been achieved by a matched investment in the benchmark portfolios. This model suits our purposes, as we wish to determine whether a manager can outperform an available set of passive index funds, and the main style categories for both indexation and active management are along size and value/growth dimensions. While the benchmark returns in (45) are not themselves available as passive index funds, they are very similar to (combinations of) index products available in the late 1990s.¹⁶ The fact that we ignore any transactions costs that would be incurred in constructing these benchmark portfolios is keeping with the conservative bias of our analysis; inclusion of such transaction costs would make managers look better.

B. Frequentist results

Before proceeding with the Bayesian performance evaluation, it is useful to summarize the frequentist evidence for this sample. The estimation of (45) for all 1437 managers yields 711 managers with a positive $\hat{\alpha}$, and 726 with a negative $\hat{\alpha}$. Note that these estimates reflect performance after expenses, and thus show almost half of the managers succeeded in earning

¹⁵ See Fama and French (1993) for details on the construction of these portfolios. We are grateful to Ken French for providing the factor returns.

¹⁶ Low-cost index funds are available in 1999 from the Vanguard mutual-fund family (among other places) in large-capitalization value, small-capitalization value, large-capitalization growth, and small-capitalization growth categories. Note that low-cost “momentum” index funds are not available, since momentum investing is by nature a high-turnover activity. This is the main reason we do not include a momentum benchmark in our analysis.

back their fees and transactions costs. A big reason for this high success rate is the survivor bias in the sample, and this bias prevents any useful inference about sample averages.

Despite the limitations of survivor bias, it is interesting to test the null hypothesis that the best performance in the sample is due to chance. To do this, we first test the null hypothesis that α is zero for each manager in the sample. On this test, the lowest p -value among all managers is achieved by Robert Sanborn of the Oakmark Fund, who has an $\hat{\alpha}$ of 91.7 bp per month, with a standard deviation for this estimate of 23.8. This yields a t -statistic of 3.8 and a p -value of 0.00014. Thus, under the null hypothesis that α is zero, we would expect to see such an extreme performance about 1.4 times in a sample of 10,000 managers. Even if we assume that our sample is randomly selected and consists of independent draws, such an extreme result is not that surprising. Under the null hypothesis that α is zero for all 1437 funds, the probability that the lowest p -value is less than or equal to 0.00014 can be calculated as

$$1 - (1 - 0.00014)^{1437} = 0.18, \tag{46}$$

or 18 percent. If survivor bias causes the sample to have a disproportionate number of good performers, then we would need to adjust upwards the exponent in (46), thus raising the probability of observing an extreme outcome. In any case, we *cannot* reject the null hypothesis that the best performer in this sample has an α equal to zero. At the end of this section, we explain the relationship between this result and the Bayesian inference described below.

C. Bayesian results

In the Bayesian analysis, we combine the return evidence with a range of possible prior beliefs and then map them into posterior beliefs. To simplify the analysis, we elicit priors before fees so that we can use the same interpretation of q , $q(25)$, and σ_α for all managers. In principle, we could use different prior parameters for each manager depending on their style, education, or other characteristics.¹⁷ Total fees are reported in the database and vary across managers and across time. Consistent with our elicitation, we analyze gross returns (by adding back fees for each year), and then subtract the current *fee* at the end. Transactions costs are not reported; we use a single value, 6 bp per month, as the *cost* for every manager.¹⁸ In Section III.D, we discuss the implications of changing this assumption.

For the denominator of the leverage term, given as s^2 in equation (10), we use the procedure of Pástor and Stambaugh (1999a) to obtain an empirical-Bayes estimate of $\hat{\sigma}^2$ for our sample (= 0.00029). This level of s^2 is a useful normalization that makes it easier to interpret the results, since prior beliefs can then be stated relative to an average level of residual risk in the sample. Thus, the elicitation uses the questions from Section II.C., and includes an answer for $q(25)$; here, the proper interpretation of $q(25)$ is the probability of α greater than 25 basis points per month, after transactions costs but before fees, and conditional on the average level of residual risk in the sample. We will use this definition of $q(25)$ for interpreting our results. Note that we cannot use our data to infer anything

¹⁷ Chevalier and Ellison (1999) provide evidence that many such characteristics are correlated with alphas.

¹⁸ This value roughly corresponds to the average monthly transactions costs for mutual funds and large institutions found in other studies; see Carhart (1997) for turnover rates and implied trading costs, Keim and Madhavan (1997) for per-trade costs, and Perold (1988) for the methodology behind these calculations.

about the “true” values of q , $q(25)$, σ_α , or $\underline{\alpha}$; this kind of analysis would require an unbiased sample (ours is not) and an additional adjustment for dependencies across managers. Thus, we are not attempting to say what prior beliefs “should be”. Rather, our goal is to better understand how these prior beliefs matter for inference and decision-making.

Given prior beliefs, the next step is to combine these beliefs with the data and calculate posterior beliefs. As an example, consider the Guardian Park Avenue fund (Class A shares). Charles Albers managed the fund from July 1972 through the end of our sample in December 1996. (He subsequently left to manage another fund.) Over this sample period, the fund earned an $\hat{\alpha}$ of 24.6 bp per month. The standard error on this α estimate is 8.5 bp. If we perform a frequentist test of the null hypothesis that α is zero, we obtain a p -value of 0.002.

How different are posterior beliefs when using an informed prior for α ? Combining Guardian’s 1996 monthly fees of 6.8 bp with our assumption that transactions costs are six bp, the prior mean for Guardian’s α is equal to -12.8 bp.¹⁹ (All statements about prior expectations are made conditional on $\sigma^2 = s^2 = 0.00029$.) Informed prior beliefs tend to shrink $\tilde{\alpha}$ towards its prior mean. For good performers, this shrinkage tends to be stronger the smaller are q and $q(25)$. As an illustration, consider the case where $q = 0.001$. Thus, the investor believes that only one in 1000 managers does any better than just throwing darts at the stock tables. Holding q constant at 0.001, Figure 3 plots Guardian’s $\tilde{\alpha}$ as a function of the free parameter, $q(25)$. The higher is $q(25)$, the more prior probability the investor is placing on $\alpha > 25$ bp per month. Recall that a adjusts for different levels of

¹⁹ Guardian also had a maximum load fee of 450 basis points in 1996, but this fee can often be waived for some customers. Many of the other top performing funds do not charge any load fees.

$q(25)$, so that the prior mean of α is always equal to -12.8 bp. The figure covers the range $q(25) \in [0, 0.0001]$. For the very smallest levels of $q(25)$, $\tilde{\alpha}$ is weighted heavily towards the prior mean and is negative. This negative range is very small, however, and $\tilde{\alpha}$ is positive for all values of $q(25)$ greater than 0.000007 . By the time we reach $q(25) = 0.0001$, $\tilde{\alpha}$ is 18 basis points. Thus, if an investor believes that one manager in a 1000 has skill, and one in 10,000 has sufficient skill so that their α is greater than 25 basis points per month, then the posterior mean for Guardian would be 18 bp.

Why does the posterior mean of Guardian converge so quickly to the frequentist estimate of 24.6 bp? The main reason is that the point mass for an unskilled manager occurs at -12.8 bp, and this is very far from the frequentist estimate of 24.6. (In frequentist terms, it is almost two standard deviations to the left of zero.) Thus, the posterior probability of skill for this manager, \tilde{q} , is updated considerably, even for very low levels of q and $q(25)$.

As impressive as Guardian’s performance is, it is not the best in the sample. In general, the best performing managers at low levels of $q(25)$ are those with a positive and “significant” $\hat{\alpha}$ and a long history of returns. Like Guardian, these managers tend to have low frequentist standard errors for their $\hat{\alpha}$ estimates, and large updates for their probability of skill. Figure 4 plots the highest and tenth highest $\tilde{\alpha}$ among all managers for $q(25) \in [0, 0.0001]$, holding q constant at 0.001.²⁰ As seen in the figure, the best performing manager’s $\tilde{\alpha}$ becomes positive at about $q(25) = 0.000003$, and the tenth-best manager becomes positive at about

²⁰ Our focus on the “best” managers does not run into the statistical difficulties that would occur in a frequentist analysis. Here, the assumptions discussed in Section II.E imply that information about manager i does not tell us anything about manager j . Thus, conditioning on the best or tenth-best manager does not affect inference about posterior means for those managers. Effectively, our prior beliefs serve the same role as the exponent used in the calculation of (46).

$q(25) = 0.000094$. At $q(25) = 0.0001$, the best performing manager has an $\tilde{\alpha}$ of 24 bp, and the tenth best performing manager has an $\tilde{\alpha}$ of 2 bp.

In Figure 4, the use of a fixed $q = 0.001$ is done only to provide a concise example. We could draw this figure for any other level of q . In every case, if $q > 0.00002$ and $q(25) > 0.00001$, there is at least one manager with $\tilde{\alpha} > 0$. Thus, as long as the investor believes that at least two in 100,000 managers has skill, and one in 100,000 managers has an α of at least 25 bp, then she will invest in at least one manager. In a frequentist test, such small values of q and $q(25)$ would be statistically indistinguishable from $q = q(25) = 0$, even in unbiased samples far larger than can currently be constructed. In other words, even though we could not reject that the best performance was due to chance (in (46)), this test has very little power against the alternative that $q = 0.00002$ and $q(25) = 0.00001$. Thus, we conclude that zero investment in active managers cannot be justified solely on the basis of the statistical evidence.

D. Discussion

A natural next step would be to quantify the expected losses to an investor who ignored active managers in favor of the benchmarks. To carry out this analysis, one would need a full predictive distribution for both benchmark and manager returns (as discussed in Appendix B), a utility function, and a complete description of the trading environment.²¹ While a complete treatment of this portfolio-choice problem is outside the scope of our paper, simple calculations suggest that even relatively low levels of $q(25)$ might lead to significant

²¹ See Pástor (1999) and Pástor and Stambaugh (1999b) for examples of this quantification.

investment in active managers for all but the most risk-averse investors. Even over the limited range of $q(25)$ considered in Figure 4, there are many managers with positive values of $\tilde{\alpha}$ that also have posterior standard deviations (for α) and residual standard deviations (σ) of similar magnitudes. Essentially, the trade-off comes down to a maximum expected loss of fees and costs if the best managers turn out to be unskilled versus the potential expected gain of several times this amount if the manager is skilled.

Our analysis assumed that *cost* was known and equal to six bp per month for all managers. How sensitive are our results to this assumption? Suppose for example, that we had used 9 bp as our baseline level. Then, for any given level of q and $q(25)$, $\underline{\alpha}$ shifts to the left, but σ_α increases (in order to maintain the same probability that α is greater than 25 bp). This effect can be seen by comparing the third column of Panels A and B in Table I. For the best performing funds, the second effect tends to dominate and $\tilde{\alpha}$ becomes higher over most of the range of Figure 4. In fact, the most conservative possible results for the best performing funds occur if we assume *cost* is zero. Even in this unrealistic case, however, the results are qualitatively similar to the 6 bp case. If *cost* is uncertain, then we cannot obtain analytical solutions but our intuition is that these effects would be second-order compared to shifting the baseline level. For example, uncertainty around 6 bp should not have a larger effect than the most conservative possible shift to 0 bp. Different assumptions about *cost* would, however, affect inference about poorly performing funds, but this is not the main subject of our analysis.

One possible criticism of our results is that they are driven by the restrictions of our

parametric structure. In particular, one might believe that returns have fatter tails than do a normal distribution, and that the best-performing managers would not look as good if we took this into account. While we cannot obtain analytical solutions using fatter-tailed return distributions, it is possible to gauge the effect of fat tails by looking at simulated data. We simulated ten years of returns for 1000 funds under three possible distributions for returns: normal, t -distributed with 10 degrees of freedom, and t -distributed with 3 degrees of freedom. In each case, we set $q = 0$, so that no fund had any skill. We then replicated Figure 4 for these data. The results show very little difference across the three return distributions.²²

The results of this section are based on a specific performance-evaluation model, but other popular models lead to the same qualitative conclusions.²³ A more serious concern is our reliance on a factor model with fixed parameters. It is well-known that successful timing ability, as manifested by changes in betas in response to informed forecasts of factor returns, will induce bias in the estimation of alphas. To deal with this concern, the methods developed here could be extended to conditional factor models. Our procedures also assume that alphas are constant over a manager’s career. One could argue that as markets grow more competitive, we should expect alphas to shrink for skilled managers. Also, as managers age and/or their portfolios grow, they may have different abilities, incentives, and opportunities, and their “true” alpha may change. Such possibilities add more dimensions to the space of prior beliefs, but we doubt that the main conclusion will change – the prior beliefs necessary

²² Details of these simulations are available from the authors. While the results are not directly comparable to those of the actual data set, it is interesting to note that for $q = 0.001$ (as in Figure 4) no simulated manager had a positive $\tilde{\alpha}$ for any $q(25) < 0.0004$.

²³ Results for the CAPM and the 4-factor model (Carhart (1997)) are available from the authors.

to support investment in active managers are virtually indistinguishable from either “no skill” or “no persistence of skill”.

E. Comparison of frequentist and Bayesian results

In Section I, we show how a naive data-driven approach leads to positive investment in an active manager whenever his $\hat{\alpha}$ is positive. This counterintuitive outcome motivates a Bayesian approach with informed priors. In this section, we implement a Bayesian approach (Section III.C) and find positive investment for all but the most skeptical prior beliefs. For this same sample, we show (Section III.B) that the best performance did not seem too extreme for the sample size. How can we reconcile the Bayesian and frequentist results?

In frequentist language, one can begin to reconcile the results by recognizing that the “null hypothesis” is different in the two analyses. In the frequentist test of (46), the null hypothesis is that α is zero for all managers. There are many managers with low p -values on this test, but none so low so that we can reject the null. In the Bayesian analysis, the effective null hypothesis of no skill occurs at a negative level of α ; for some managers, this level may be two standard deviations to the left of zero. In the Bayesian analysis, degrees of “rejection” of the null lead us to update the posterior probability of skill, and these rejections may be much stronger than if the null were at zero.

Another consideration in reconciling the frequentist and Bayesian results is the recognition that “insignificant” evidence may have large investment implications. This point is first made, in another context, by Kandel and Stambaugh (1996). The downside of choosing an active manager is that he may be unskilled. In expectation, the investor will then have

paid the expenses for nothing. This expected downside is limited. The upside of skilled management is potentially much larger. Even if the best performers in the sample are not significantly extreme, they may still be sufficiently extreme to justify their expenses.

IV. Conclusion

Should investors avoid all actively managed mutual funds? The average active fund underperforms index funds on a risk-adjusted basis. Skilled management, if it exists at all, is difficult to detect. When we analyze a sample of 1437 managers extant at the end of 1996, we cannot reject the null hypothesis that the best performance is due to chance. These facts by themselves might lead investors to shun actively managed funds. Our analysis shows that this conclusion is premature. Given our current methods of performance evaluation, the prior beliefs necessary to support some investment in active managers could not possibly be distinguished from “zero skill among managers” unless we could observe hundreds of thousands of managers over many decades. Thus, we conclude that the case against actively managed funds cannot rely solely on the statistical evidence.

Our analysis does not include elements of the investor’s decision such as load fees, taxes, and limitations on short sales. Furthermore, nobody knows the correct model of performance evaluation. Given these limitations, we do not claim to provide a definitive analysis of the portfolio-choice decision. Most investors may be best served by simple rules-of-thumb, especially if they do not possess the discipline or technology to implement sophisticated trading strategies. Nevertheless, we believe that the investor’s perspective motivates the importance of using informed prior beliefs in a Bayesian method of performance evaluation. This

method provides a new lens on the performance-evaluation evidence, with the final image in sharp contrast to frequentist results.

Appendix

A. Details for Section II.D.

Throughout this Appendix, we assume an informative prior on σ^2 , and, as in the text, a diffuse (improper) prior on β :

$$p(\beta) \propto 1, \quad (47)$$

$$p(\sigma^2) \propto \frac{1}{\sigma^{\nu_0+2}} \exp \left\{ -\frac{h_0}{2\sigma^2} \right\}. \quad (48)$$

Results in the text can be obtained by substituting $\nu_0 = 0$ and $h_0 = 0$ into the expressions below. Otherwise, the setup is the same as in Section II, and we make use of the same notation.

The likelihood for factors, unspecified in the text, is assumed to take the following form:

$$p(F_t | \mu_F, \Sigma_F) = N(\mu_F, \Sigma_F), \quad (49)$$

with realizations independent across t . The prior on μ_F and Σ_F is assumed to be diffuse:

$$p(\mu_F, \Sigma_F) \propto |\Sigma_F|^{-\frac{K+1}{2}}. \quad (50)$$

This assumption is made for notational convenience and, except for the discussion of the predictive distribution below, does not affect our results.

Derivation of the posterior distribution for α :

From Bayes' rule, the joint posterior for $(\theta, \sigma^2, \mu_F, \Sigma_F)$ is given by:

$$\begin{aligned} p(\theta, \sigma^2, \mu_F, \Sigma_F | r, F) &\propto p(r | \theta, \sigma^2, F)p(F | \mu_F, \Sigma_F)p(\theta, \sigma^2)p(\mu_F, \Sigma_F) \\ &\propto p(\theta, \sigma^2 | r, F)p(\mu_F, \Sigma_F | F), \end{aligned} \quad (51)$$

where we have used the prior independence of (θ, σ^2) and (μ_F, Σ_F) , and the fact that the likelihood for r conditional on factors depends only on θ and σ^2 , while the likelihood for the factors depends only on (μ_F, Σ_F) . Therefore, (θ, σ^2) , and (μ_F, Σ_F) are independent in the posterior.

The above arguments imply

$$p(\theta, \sigma^2 | r, F) \propto p(r | \theta, \sigma^2, F)p(\theta, \sigma^2). \quad (52)$$

Because Z is independent from (μ_F, Σ_F) in the prior, and because the likelihood for factors does not depend on Z , the equations above are also valid conditional on a value of Z . We will make use of (52) throughout this Appendix.

The likelihood for r conditional on factors is given by:

$$\begin{aligned} p(r | \theta, \sigma^2, F) &\propto \frac{1}{\sigma^T} \exp \left\{ -\frac{1}{2\sigma^2} (r - X\theta)'(r - X\theta) \right\} \\ &= \frac{1}{\sigma^T} \exp \left\{ -\frac{1}{2\sigma^2} \left(S + (\theta - \hat{\theta})'X'X(\theta - \hat{\theta}) \right) \right\}, \end{aligned} \quad (53)$$

where S , θ , $\hat{\theta}$, and X are defined as in Section II.D.

Combining the likelihood and the prior yields:

$$p(\theta, \sigma^2 | Z = 1, r, F) \propto \frac{1}{\sigma^{\nu_0+T+2}} \frac{1}{\sigma} \quad (54)$$

$$\exp \left\{ -\frac{1}{2\sigma^2} \left(h_0 + S + \frac{s^2}{\sigma_\alpha^2} (\alpha - \underline{\alpha})^2 + (\theta - \hat{\theta})'(X'X)(\theta - \hat{\theta}) \right) \right\} \mathbf{1}_{\alpha > \underline{\alpha}}.$$

Integrating with respect to β yields:

$$p(\alpha, \sigma^2 \mid Z = 1, r, F) \propto \frac{1}{\sigma^{\nu_0 + T - K + 2}} \frac{1}{\sigma} \exp \left\{ -\frac{1}{2\sigma^2} \left(h_0 + S + \frac{s^2}{\sigma_\alpha^2} (\alpha - \underline{\alpha})^2 + \frac{(\alpha - \hat{\alpha})^2}{m} \right) \right\} \mathbf{1}_{\alpha > \underline{\alpha}}. \quad (55)$$

Completing the square in α yields:

$$p(\alpha, \sigma^2 \mid Z = 1, r, F) \propto \frac{1}{\sigma^{\nu+2}} \frac{1}{\sigma} \exp \left\{ -\frac{1}{2\sigma^2} \left(h + (\lambda m)^{-1} (\alpha - \alpha')^2 \right) \right\} \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (56)$$

where

$$\begin{aligned} h &= h_0 + S + \left(\frac{1 - \lambda}{m} \right) (\hat{\alpha} - \underline{\alpha})^2, \\ \nu &= \nu_0 + T - K. \end{aligned} \quad (57)$$

Conditional on σ^2 , all terms involving σ^2 (as well as all terms involving only the data, such as $\hat{\alpha}$ and S), can be considered constants. Therefore,

$$p(\alpha \mid \sigma^2, Z = 1, r, F) \propto \exp \left\{ -\frac{1}{2\sigma^2} (\alpha - \alpha')^2 \right\} \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (58)$$

and we have shown (19).

In spite of the truncation, the functional form of the posterior is that of a conjugate prior distribution. Therefore, σ^2 can be integrated out of (56) using the properties of the gamma distribution (see, e.g., Gelman et al. (1995)). Define a change of variables $u = \frac{1}{2}A\sigma^{-2}$, where $A = h + (\lambda m)^{-1}(\alpha - \alpha')^2$. The resulting function of u is the pdf of a gamma distribution without the normalizing constant:

$$p(\alpha \mid Z = 1, r, F) \propto \left(\frac{A}{2} \right)^{-\frac{\nu+1}{2}} \left(\int u^{\frac{\nu-1}{2}} \exp\{-u\} du \right) \mathbf{1}_{\alpha > \underline{\alpha}}. \quad (59)$$

Using the proportionality constant for the gamma distribution, and dividing through by h yields:

$$p(\alpha | Z = 1, r, F) \propto \left(\frac{h}{2}\right)^{-\frac{\nu+1}{2}} \Gamma\left(\frac{\nu+1}{2}\right) \left(1 + \frac{(\lambda m)^{-1}(\alpha - \alpha')^2}{\nu(h/\nu)}\right)^{-\frac{\nu+1}{2}} \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (60)$$

which is proportional to the pdf of a t -distribution. Therefore,

$$p(\alpha | Z = 1, r, F) \propto t_\nu\left(\alpha', \frac{\lambda m h}{\nu}\right) \mathbf{1}_{\alpha > \underline{\alpha}}. \quad (61)$$

Derivation of the posterior expectation $\tilde{\alpha}$ conditional on skill:

Because $p(\alpha | Z = 1, r, F)$ must integrate to 1:

$$p(\alpha | Z = 1, r, F) = \frac{1}{\int_{\underline{\alpha}}^{\infty} t_\nu\left(\alpha; \alpha', \frac{\lambda m h}{\nu}\right) d\alpha} t_\nu\left(\alpha; \alpha', \frac{\lambda m h}{\nu}\right) \mathbf{1}_{\alpha > \underline{\alpha}}. \quad (62)$$

Let $\tilde{t} = \int_{\underline{\alpha}}^{\infty} t_\nu(\alpha, \alpha', \lambda m h/\nu) d\alpha$. Then

$$\begin{aligned} E[\alpha | Z = 1, r, F] &= \frac{1}{\tilde{t}} \frac{\Gamma(\nu+1)/2}{\Gamma(\nu/2)} \sqrt{\frac{\lambda m h}{2\pi}} \int_{\underline{\alpha}}^{\infty} \alpha \left(1 + \frac{(\alpha - \alpha')^2}{\lambda m h}\right)^{-\frac{\nu+1}{2}} d\alpha \\ &= \alpha' + \frac{1}{\tilde{t}} \frac{\Gamma(\nu+1)/2}{\Gamma(\nu/2)} \sqrt{\frac{\lambda m h}{2\pi}} \left(1 + \frac{(\underline{\alpha} - \alpha')^2}{\lambda m h}\right)^{-\frac{\nu-1}{2}} \\ &= \alpha' + \frac{\lambda m h}{\nu - 2} t_{\nu-2}\left(\underline{\alpha}; \alpha', \frac{\lambda m h}{\nu - 2}\right) \frac{1}{\int_{\underline{\alpha}}^{\infty} t_\nu(\alpha; \alpha', \frac{\lambda m h}{\nu-2}) d\alpha}. \end{aligned} \quad (63)$$

The first line follows from the pdf of the t distribution, and the last from multiplying and dividing by the necessary constants.

Derivation of B :

Under our assumptions, $p(\theta, \sigma^2 | Z = 1, F) = p(\theta, \sigma^2 | Z = 1)$. Therefore,

$$p(r | Z = 1, F) = \int p(r | Z = 1, F, \theta, \sigma^2) p(\theta, \sigma^2 | Z = 1) d\theta d\sigma^2. \quad (64)$$

Substituting in for the likelihood and the prior and integrating with respect to β yields:

$$p(r | Z = 1, F) = C \sqrt{\frac{2}{\pi}} \frac{s}{\sigma_\alpha} \int_0^\infty \int_{\underline{\alpha}}^\infty \frac{1}{\sigma^{\nu+2}} \frac{1}{\sigma} \exp \left\{ -\frac{1}{2\sigma^2} \left(h_0 + S + \frac{s^2}{\sigma_\alpha^2} (\alpha - \underline{\alpha})^2 + \frac{(\alpha - \hat{\alpha})^2}{m} \right) \right\} d\alpha d\sigma^2, \quad (65)$$

where C is a constant that is identical for $Z = 1$ and $Z = 0$. Completing the square in α yields:

$$p(r | Z = 1, F) = C \sqrt{\frac{2}{\pi}} \frac{s}{\sigma_\alpha} \int \frac{1}{\sigma^{\nu+2}} \frac{1}{\sigma} \exp \left\{ -\frac{1}{2\sigma^2} \left(h + (\lambda m)^{-1} (\alpha - \alpha')^2 \right) \right\} d\sigma^2. \quad (66)$$

The form of the equation is the same as in (56). As above, the properties of the gamma distribution are used to integrate out σ^2 :

$$p(r | Z = 1, F) = C \sqrt{\frac{2}{\pi}} \frac{s}{\sigma_\alpha} \left(\frac{h}{2} \right)^{-\frac{\nu+1}{2}} \Gamma \left(\frac{\nu+1}{2} \right) \int_{\underline{\alpha}}^\infty \left(1 + \frac{(\lambda m)^{-1} (\alpha - \alpha')^2}{\nu(h/\nu)} \right)^{-\frac{\nu+1}{2}} d\alpha. \quad (67)$$

The term inside the integral is proportional to a t -distribution. Therefore,

$$p(r | Z = 1, F) = C (1 - \lambda)^{\frac{1}{2}} \left(\frac{h}{2} \right)^{-\frac{\nu}{2}} \Gamma \left(\frac{\nu}{2} \right) 2 \int_{\underline{\alpha}}^\infty t_\nu(\alpha; \alpha', \lambda m h / \nu) d\alpha. \quad (68)$$

The calculation for $Z = 0$ follows along the same lines. As above, after integrating out with respect to β , we obtain:

$$p(r | Z = 0, F) = C \int \frac{1}{\sigma^{\nu+2}} \exp \left\{ -\frac{1}{2\sigma^2} \left(S + h_0 + \frac{(\hat{\alpha} - \underline{\alpha})^2}{m} \right) \right\} d\sigma^2. \quad (69)$$

Integrating with respect to σ^2 yields:

$$p(r | Z = 0, F) = C \left(\frac{S + h_0 + (\hat{\alpha} - \underline{\alpha})^2 / m}{2} \right)^{-\frac{\nu}{2}} \Gamma \left(\frac{\nu}{2} \right). \quad (70)$$

Therefore,

$$B = (1 - \lambda)^{\frac{1}{2}} \left(\frac{S + h_0 + \frac{1-\lambda}{m} (\hat{\alpha} - \underline{\alpha})^2}{S + h_0 + m^{-1} (\hat{\alpha} - \underline{\alpha})^2} \right)^{-\frac{\nu}{2}} 2 \int_{\underline{\alpha}}^\infty t_\nu \left(\alpha; \alpha', \frac{\lambda m h}{\nu} \right) d\alpha. \quad (71)$$

Both the numerator and the denominator are proportional to a t -distribution with $\nu - 1$ degrees of freedom (in one case, α was known, while in the other case we integrated with respect to α). Multiplying and dividing by a constant yields expression (32) in the text.

B. Drawing from the predictive distribution

This section describes how to draw from the joint predictive distribution for r and F . We first show how to draw from posterior distribution of $(\alpha, \beta, \sigma^2)$. This is done by first drawing from the skilled distribution, then the unskilled distribution, and then the full posterior. Finally, we draw from the predictive distribution of r and F conditional on all parameters.

Drawing from the distribution conditional on $Z = 1$ (“skilled posterior”):

Consider random variables $\check{\sigma}^2$ and $\check{\alpha}$ such that:

$$\check{\sigma}^2 \mid r, F, Z = 1 \sim IG\left(\frac{\nu}{2}, \frac{h}{2}\right), \quad (72)$$

$$\check{\alpha} \mid \sigma^2 = \check{\sigma}^2, r, F, Z = 1 \sim N(\alpha', \sigma'^2). \quad (73)$$

Then

$$p(\check{\alpha}, \check{\sigma}^2 \mid r, F, Z = 1) \propto \frac{1}{\check{\sigma}^{\nu+2}} \frac{1}{\check{\sigma}} \exp\left\{-\frac{1}{2\check{\sigma}^2}(h + (\lambda m)^{-1}(\check{\alpha} - \alpha')^2)\right\}. \quad (74)$$

Using the procedure above, but discarding the draw whenever $\check{\alpha} < \underline{\alpha}$ produces a distribution that has zero mass when $\check{\alpha} < \underline{\alpha}$, but where the relative densities of any other points are the same as in (74). By (56), this is exactly the joint distribution of α and σ^2 . Therefore, drawing $\check{\sigma}^2$ from $IG(\nu/2, h/2)$ and $\check{\alpha} \mid \sigma^2 = \check{\sigma}^2$ from $N(\alpha', \sigma'^2)$, and discarding the draws whenever $\check{\alpha} < \underline{\alpha}$, produces a draw from the joint posterior.

The posterior for β conditional on α and σ^2 follows from the properties of the multivariate normal:

$$\beta \mid \alpha, \sigma^2, r, F, Z = 1 \sim N(\hat{\beta} + m^{-1}y(\alpha - \hat{\alpha}), \sigma^2(Q - m^{-1}yy')), \quad (75)$$

where y and Q are submatrices of $(X'X)^{-1}$:

$$\begin{pmatrix} m & y' \\ y & Q \end{pmatrix}. \quad (76)$$

Drawing from the distribution conditional on $Z = 0$ (“unskilled posterior”):

It follows from the likelihood and the prior, that conditional on $Z = 0$ we have

$$\sigma^2 \mid Z = 0, r, F \sim IG\left(\frac{\nu}{2}, \frac{h_0 + S + m^{-1}(\hat{\alpha} - \underline{\alpha})^2}{2}\right). \quad (77)$$

Conditional on σ^2 and on $\alpha = \underline{\alpha}$, β is drawn from (75).

Drawing from the full posterior:

For any given draw $(\alpha^j, \beta^j, \sigma^j)$, there is a probability \tilde{q} that the draw comes from the skilled posterior, and probability $1 - \tilde{q}$ that the draw comes from the unskilled posterior. Let u^j be a draw from the distribution with uniform mass on $[0, 1]$. If $u^j < \tilde{q}$, then $(\alpha^j, \beta^j, \sigma^j)$ is drawn from the skilled posterior. Otherwise, a draw is made from the unskilled posterior.

Drawing from the predictive distribution

From the likelihood:

$$p(r_{T+1} \mid \theta, \sigma^2, F_{T+1}) = N(\alpha + F_{T+1}\beta, \sigma^2), \quad (78)$$

and

$$p(F_{T+1} \mid \mu_F, \Sigma_F) = N(\mu_F, \Sigma_F). \quad (79)$$

The following procedure produces draws from the predictive distribution of r_{T+1} . First draw (μ_F, Σ_F) from their posterior distribution (see Pástor (1999)), then draw F_{T+1} from (79), conditional on those values of μ_F and Σ_F . Next draw $(\alpha, \beta, \sigma^2)$ from the posterior, as described above. Finally, draw r_{T+1} from (78), conditional on those values of the parameters, and of F_{T+1} .

C. Derivation of the positive-investment condition for Section I and Section II.E.

In this Appendix we show that if an asset has a positive posterior expectation of α (i.e. $\tilde{\alpha} > 0$) then this asset will be held in positive quantities in the mean-variance tangency portfolio. The result is demonstrated first for a single asset and then extended to the multiple asset case of Section II.E.

First consider the single-manager case. In a Bayesian setting the asset weights of the tangency portfolio are defined by

$$\begin{pmatrix} w^* \\ x^* \end{pmatrix} = \frac{\tilde{V}^{-1} \tilde{E}}{\text{normalizing constant}}, \quad (80)$$

where w^* and x^* are, respectively, the optimal weights on the active manager and the index funds and \tilde{E} and \tilde{V} are, respectively, the mean and variance covariance matrix of the predictive distribution of all the assets (see Pástor (1999)). Suppose that the first asset consists of the following portfolio:

- long the manager
- short the index funds with weights equal to $\tilde{\beta} \equiv E[\beta | r, F]$.

We call this asset the “alpha portfolio” (as in Section II.E). Since this portfolio is a linear combination of the manager and an index funds, we drop the manager from the analysis and consider the remaining $K + 1$ portfolios: K index funds and one alpha portfolio. If we can show that the alpha portfolio has a positive weight in the tangency portfolio defined by (80), then it follows that the manager would have a positive weight as well.

Conditional on the parameters and past data, next period’s total return on the alpha portfolio is given by

$$R_{T+1} \equiv \alpha + (\beta - \tilde{\beta}) F_{T+1} + \varepsilon_{T+1}, \quad (81)$$

while next period’s return on the index funds is

$$F_{T+1} \equiv \mu_F + \eta_{T+1}, \quad (82)$$

where $\varepsilon_{T+1} \sim N(0, \sigma^2)$ and $\eta_{T+1} \sim N(0, \Sigma_F)$ are independent. It follows from the law of iterated expectations that the posterior expectation of index-fund returns is given by:

$$\begin{aligned} E[F_{T+1} | r, F] &= E \left[E \left[\mu_F + \eta_{T+1} \mid \mu_F, \Sigma_F, r, F \right] \mid r, F \right] \\ &= \tilde{\mu}_F, \end{aligned} \quad (83)$$

where $\tilde{\mu}_F \equiv E[\mu_F | r, F]$.

The predictive expectation for the return on the alpha portfolio is given by:

$$\begin{aligned} E[R_{T+1} | r, F] &= E \left[\alpha + (\beta - \tilde{\beta}) F_{T+1} + \varepsilon_{T+1} \mid r, F \right] \\ &= \tilde{\alpha} + E \left[(\beta - \tilde{\beta}) F_{T+1} \mid r, F \right], \end{aligned} \quad (84)$$

and, from the posterior independence of β and μ_F , we have

$$E[R_{T+1} | r, F] = \tilde{\alpha} + E \left[E \left[(\beta - \tilde{\beta}) F_{T+1} \mid \theta, \sigma, \mu_F, \Sigma_F, r, F \right] \mid r, F \right]$$

$$\begin{aligned}
&= \tilde{\alpha} + E\left[(\beta - \tilde{\beta})\mu_F \mid r, F\right] \\
&= \tilde{\alpha}.
\end{aligned} \tag{85}$$

We next show that the predictive covariance between the returns to index fund i and to the alpha portfolio is zero. The predictive expectation of the product of the returns is given by

$$\begin{aligned}
E[R_{T+1}F_{i,T+1} \mid r, F] &= E[F_{i,T+1}\alpha \mid r, F] + E\left[F_{i,T+1}(\beta - \tilde{\beta})F_{T+1} + F_{i,T+1}\varepsilon_{T+1} \mid r, F\right] \\
&= E[\alpha E[F_{i,T+1} \mid \mu_F, \Sigma_F, r, F] \mid r, F] + \\
&\quad E\left[(\beta - \tilde{\beta})E[F_{i,T+1}F_{T+1} \mid \mu_F, \Sigma_F] \mid r, F\right] \\
&= \tilde{\alpha}\tilde{\mu}_{i,F},
\end{aligned} \tag{86}$$

where $F_{i,T+1}$ is the next period return on index fund i and $\tilde{\mu}_{i,F}$ is the posterior expectation of index fund i . The second line follows from the law of iterated expectations, and the third line follows from posterior independence of θ and F (note that $E[F_{i,T+1}F_{T+1} \mid \mu_F, \Sigma_F, r, F]$ is a function of elements of μ_F and Σ_F). Therefore, the predictive covariance of F_{T+1} and the next period return on alpha portfolio i , denoted by $R_{i,T+1}$, is zero for all i . The structure of \tilde{V} is

$$\begin{pmatrix} v_{11} & 0 \\ 0 & V_{22} \end{pmatrix}, \tag{87}$$

where v_{11} is the predictive variance of the return on the alpha portfolio and V_{22} is the predictive variance of the index funds. Hence $w^* \propto \tilde{\alpha}/v_{11}$, which is positive if and only if $\tilde{\alpha} > 0$.

Next consider the multiple-managers case of Section II.E. Now there are N alpha portfolios. For simplicity, we keep the same notation as above; \tilde{V} is still the predictive variance-covariance matrix, and \tilde{E} is still the vector of means. Assumptions 1 and 2 (from Section II.E) imply that the posterior distributions for each manager are independent. Therefore, all the calculations above remain valid in the case of multiple managers. In particular, the covariance of the return on any alpha portfolio with any index fund is zero. Therefore, \tilde{V} takes the form:

$$\begin{pmatrix} V_{11} & 0 \\ 0 & V_{22} \end{pmatrix}. \quad (88)$$

Moreover, the posterior expectation on the i^{th} alpha portfolio is equal to $\tilde{\alpha}$. For the conclusion, it suffices to show that V_{11} is diagonal, i.e. that the covariance between $R_{i,T+1}$ and $R_{j,T+1}$ equals zero for any $i \neq j$. Conditioning on the values of the parameters and using the law of iterated expectations yields:

$$\begin{aligned} E[R_{i,T+1}R_{j,T+1} | r, F] &= E \left[E \left[(\alpha_i + (\beta_i - \tilde{\beta}_i)F_{T+1})(\alpha_j + (\beta_j - \tilde{\beta}_j)F_{T+1}) \mid \mu_F, \Sigma_F, r, F \right] \mid r, F \right] \\ &= E \left[E \left[\alpha_i\alpha_j + \alpha_i(\beta_j - \tilde{\beta}_j)F_{T+1} \right. \right. \\ &\quad \left. \left. + \alpha_j(\beta_i - \tilde{\beta}_i)F_{T+1} + (\beta_i - \tilde{\beta}_i)F_{T+1}F'_{T+1}(\beta_j - \tilde{\beta}_j)' \mid \mu_F, \Sigma_F, r, F \right] \mid r, F \right] \\ &= \tilde{\alpha}_i\tilde{\alpha}_j. \end{aligned} \quad (89)$$

The last equation follows because of the posterior independence of the manager parameters from the index fund parameters, and the parameters for manager i from the parameters from manager j . Therefore, the posterior covariance between $R_{i,T+1}$ and $R_{j,T+1}$ is zero, and the V_{11} is diagonal. Thus, there is positive investment in any manager with $\tilde{\alpha}_i > 0$.

References

- Baks, Klaas, Andrew Metrick, and Jessica Wachter, 1999, Bayesian performance evaluation, NBER Working Paper # 7069.
- Barberis, Nicholas, 1999, Investing for the long run when returns are predictable, *Journal of Finance*, forthcoming.
- Bawa, Vijay S., Stephen J. Brown, and Roger W. Klein, 1979, *Estimation Risk and Optimal Portfolio Choice* (North-Holland, Amsterdam).
- Brown, Stephen J., 1979, Optimal portfolio choice under uncertainty: A Bayesian approach, in Vijay S. Bawa, Stephen J. Brown, and Roger W. Klein, Eds: *Estimation Risk and Optimal Portfolio Choice* (North-Holland, Amsterdam).
- Brown, Stephen J., and William N. Goetzmann, 1995, Performance persistence, *Journal of Finance* 50, 679-698.
- Brown, Stephen J., and William N. Goetzmann, Roger Ibbotson and Stephen Ross, 1992, Survivorship bias in performance studies, *Review of Financial Studies* 5, 553-580.
- Carhart, Mark M., 1995, Survivorship bias and mutual fund performance, Unpublished Ph.D. dissertation Graduate School of Business, University of Chicago.
- Carhart, Mark, 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57-82.
- Carlson, Robert S., 1970, Aggregate performance in mutual funds, *Journal of Financial and Quantitative Analysis* 5, 1-32.
- Center for Research in Security Prices, CRSP survivor bias free US mutual fund data base file guide, 1999.
- Chamberlain, Gary, 1998, Econometrics and decision theory, Working paper, Harvard University.
- Chevalier, Judith, and Glenn Ellison, 1999, Are some mutual fund managers better than others? Cross-sectional patterns in behavior and performance, *Journal of Finance* 54, 875-900.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic based benchmarks, *Journal of Finance* 52, 1035-1058.
- Daniel, Kent, and Sheridan Titman, 1997, Evidence on the characteristics of cross-sectional variation in stock returns, *Journal of Finance* 52, 1-34.

- Elton, Edwin J., Martin J. Gruber, and Christopher R. Blake, 1996, The persistence of risk-adjusted mutual fund performance, *Journal of Business* 69, 133-157.
- Fama, Eugene, 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383-417.
- Fama, Eugene, and Kenneth French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3-56.
- Ferson, Wayne, and Rudi Schadt, 1996, Measuring fund strategy and performance in changing economic conditions, *Journal of Finance* 51, 425-461.
- Frost, Peter A., and James E. Savarino, 1986, An empirical Bayes approach to efficient portfolio selection, *Journal of Financial and Quantitative Analysis* 21, 293-305.
- Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin, 1995, *Bayesian Data Analysis*, Chapman and Hall.
- Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121-1152.
- Goetzmann, William N. and Phillippe Jorion, 1997, Re-emerging markets, Working paper, University of California at Irvine.
- Goetzmann, William N. and Roger G. Ibbotson, 1994, Do winners repeat? Patterns in mutual fund performance, *Journal of Portfolio Management* 20, 9-17.
- Golec, Joseph H, 1996. The effects of mutual fund managers' characteristics on their portfolio performance, risk, and fees, *Financial Services Review* 5, 133-148.
- Grinblatt, Mark, and Sheridan Titman, 1988, The evaluation of mutual fund performance: An analysis of monthly returns, Working paper, The John E. Anderson Graduate School of Management at UCLA.
- Grinblatt, Mark, and Sheridan Titman, 1992, Performance in mutual funds, *Journal of Finance* 47, 1977-1984.
- Gruber, Martin J., 1996, Another puzzle: The growth in actively managed mutual funds, *Journal of Finance* 51, 783-810.
- Hendricks, Darryl, Jayendu Patel and Richard Zeckhauser , 1993, Hot hands in mutual funds: The persistence of performance, 1974 - 1988, *Journal of Finance* 48, 93-130.

- Jensen, Michael C., 1968, The performance of mutual funds in the period 1945-1964, *The Journal of Finance* 23, 389-416.
- Jobson, J.D., and Bob Korkie, 1980, Estimation for Markowitz efficient portfolios, *Journal of the American Statistical Association* 75, 544-554.
- Jobson, J.D., Bob Korkie, and V. Ratti, 1979, Improved estimation for Markowitz efficient portfolios using James-Stein type estimators, *Proceedings of the American Statistical Association, Business and Economics Statistics Section* 41, 279-284.
- Jorion, Philippe, 1985, International portfolio diversification with estimation risk, *Journal of Business* 58, 259-278.
- Jorion, Philippe, 1986, Bayes-Stein estimation for portfolio analysis, *Journal of Financial and Quantitative Analysis* 21, 279-292.
- Jorion, Philippe, 1991, Bayesian and CAPM estimators of the means: Implications for portfolio selection, *Journal of Banking and Finance* 15, 717-728.
- Kandel, Shmuel, Robert McCulloch, and Robert F. Stambaugh, 1995, Bayesian inference and portfolio efficiency, *Review of Financial Studies* 8, 1-53.
- Kandel, Shmuel, and Robert F. Stambaugh, 1996, On the predictability of stock returns: An asset-allocation perspective, *Journal of Finance* 51, 385-424.
- Keim, Donald B., and Ananth Madhavan, 1997, Transactions costs and investment style: An inter-exchange analysis of institutional equity trades, *Journal of Financial Economics* 46, 265-92.
- Klein, Roger W., and Vijay S. Bawa, 1976, The effect of estimation risk on optimal portfolio choice, *Journal of Financial Economics* 3, 215-231.
- Lakonishok, Josef, Andrei Shleifer, and Robert Vishny, 1992, The structure and performance of the money management industry, *Brookings Papers: Microeconomics*, 339-391.
- Lakonishok, Josef, Andrei Shleifer, and Robert Vishny, 1994, Contrarian investment, extrapolation, and risk, *Journal of Finance* 49, 1541-1578.
- Lehman, Bruce N., and David Modest, 1987, Mutual fund performance evaluation: A comparison of benchmarks and a benchmark of comparisons, *Journal of Finance* 21, 233-265.

- McCulloch, Robert, and Peter E. Rossi, 1990, Posterior, predictive, and utility-based approaches to testing the arbitrage pricing theory, *Journal of Financial Economics* 28, 7-38.
- Malkiel, Burton, 1995, Returns from investing in mutual funds 1971 to 1991, *Journal of Finance* 50, 549-572.
- Pástor, Lubos, 1999, Portfolio selection and asset pricing models, *Journal of Finance*, Forthcoming.
- Pástor, Lubos, and Robert F. Stambaugh, 1999a, Costs of equity capital and model mispricing, *Journal of Finance* 54, 67-121.
- Pástor, Lubos, and Robert F. Stambaugh, 1999b, Comparing Asset-Pricing Models: An Investment Perspective, Working paper, The University of Pennsylvania.
- Perold, André F., 1988, The implementation shortfall: paper versus reality, *The Journal of Portfolio Management* 14, 4-9.
- Shanken, Jay, 1987, A Bayesian approach to testing portfolio efficiency, *Journal of Financial Economics* 19, 195-215.
- Shleifer, Andrei, 1999, *Clarendon Lectures: Inefficient Markets*, Oxford University Press, Forthcoming.
- Stambaugh, Robert F., 1997, Analyzing investments whose histories differ in length, *Journal of Financial Economics* 45, 285-331.
- Zheng, Lu, 1999, Is money smart? A study of mutual fund investors' fund selection ability, *Journal of Finance* 54, forthcoming.
- Zellner, Arnold, 1971, *An Introduction to Bayesian Inference in Econometrics*, John Wiley and Sons: New York.




Figure 1:

Prior distribution of α

Figure 1 illustrates the (conditional) prior distribution for α . Please see Section I for a detailed discussion of the parameters in this figure. α is the intercept in a factor model (see equation (1)). q is the probability that a manager is skilled; conditional on skill, we have $\alpha \sim N(\underline{\alpha}, \sigma_\alpha^2)$ with a left truncation at $\underline{\alpha}$, where $\underline{\alpha} = a - \text{fee} - \text{cost}$ is the expected abnormal return for an unskilled manager, a is the expected negative return from transactions with skilled managers, and fee and cost are the manager's fees and transactions costs, respectively. Note that this plot combines a point mass at $\alpha = \underline{\alpha}$ and a density for $\alpha > \underline{\alpha}$.

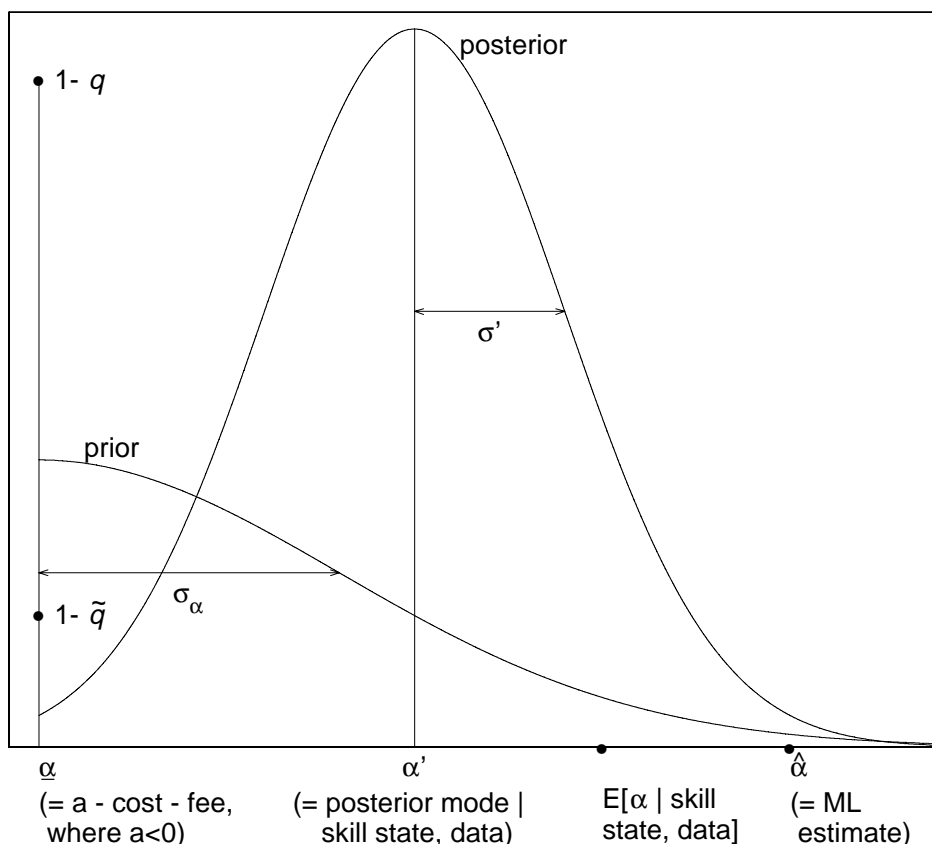


Figure 2:

Prior and posterior distribution of α

Figure 2 illustrates the prior and posterior distributions for α . α is the intercept in a factor model (see equation (1)). q is the probability that a manager is skilled; conditional on skill and $\sigma^2 = s^2$, we have $\alpha \sim N(\underline{\alpha}, \sigma_\alpha^2)$ with a left truncation at $\underline{\alpha}$, where $\underline{\alpha} = a - fee - cost$ is the expected abnormal return for an unskilled manager, a is the expected negative return from transactions with skilled managers, and fee and $cost$ are the manager's fees and transactions costs, respectively. Conditional on skill and $\sigma^2 = s^2$, the posterior distribution of α is $N(\alpha', \sigma'^2)$ with a left truncation at $\underline{\alpha}$. \tilde{q} is the posterior probability that $Z = 1$. $\hat{\alpha}$ is the maximum likelihood estimate of α . Note that this plot combines a point mass at $\alpha = \underline{\alpha}$ and a density for $\alpha > \underline{\alpha}$.




Figure 3:

Guardian Park Avenue Fund
 $\tilde{\alpha}$ as a function of $q(25)$ when $q = 0.001$

Figure 3 shows the relation between $\tilde{\alpha}$ and $q(25)$ for $q = 0.001$ for the Guardian Park Avenue Fund, Class A (Manager: Charles Albers). α is the intercept in the Fama-French (1993) three-factor model (see equation (45)). $\tilde{\alpha}$ is the posterior expectation of α . q is the probability that a manager is skilled; conditional on skill and $\sigma^2 = s^2$, we have $\alpha \sim N(\underline{\alpha}, \sigma_\alpha^2)$ with a left truncation at $\underline{\alpha}$, where $\underline{\alpha}$ is the expected abnormal return for an unskilled manager. $q(25) \equiv P(\alpha > 25 \mid \sigma^2 = s^2)$. The prior constant s^2 is 0.00029. The sample period for Guardian Park Avenue is 7/72 - 12/96.




Figure 4:

The highest (top line) and tenth highest (bottom line) $\tilde{\alpha}$ over all managers when $q = 0.001$

Figure 4 plots the highest and tenth highest $\tilde{\alpha}$ over all 1437 managers for $q(25) \in [0, 0.0001]$ and $q = 0.001$. α is the intercept in the Fama-French (1993) three-factor model (see equation (45)). $\tilde{\alpha}$ is the posterior expectation of α . $q(25) \equiv \Pr(\alpha > 25 \mid \sigma^2 = s^2)$. The plot shows the highest and tenth highest value of $\tilde{\alpha}$ among all managers for each level of $q(25)$. The prior constant s^2 is 0.00029.

Table I
Elicitation of Prior Beliefs

This table illustrates the mapping from $q(25)$, q , fee and $cost$ into the parameters σ_α , a and $\underline{\alpha}$. $q(25)$ is defined as $P(\alpha > 25 | \sigma^2 = s^2)$; q is the probability that a manager is skilled; conditional on skill and $\sigma^2 = s^2$, we have $\alpha \sim N(\underline{\alpha}, \sigma_\alpha^2)$ with a left truncation at $\underline{\alpha}$, where $\underline{\alpha} = a - fee - cost$ is the expected abnormal return for an unskilled manager, a is the expected negative return from transactions with skilled managers, and fee and $cost$ are the manager's fees and transactions costs, respectively. σ_α , a and $\underline{\alpha}$ are expressed in bp per month; $q(25)$ and q are expressed as probabilities.

Panel A: $fee = 8$ bp and $cost = 6$ bp				
$q(25)$	q	σ_α	a	$\underline{\alpha}$
0.0001	0.001	23.72	-0.019	-14.02
	0.01	15.19	-0.121	-14.12
	0.1	12.15	-0.969	-14.97
0.001	0.01	23.83	-0.190	-14.19
	0.1	15.62	-1.247	-15.25
0.01	0.1	24.92	-1.988	-15.99
Panel B: $fee = 8$ bp and $cost = 9$ bp				
$q(25)$	q	σ_α	a	$\underline{\alpha}$
0.0001	0.001	25.55	-0.020	-17.02
	0.01	16.36	-0.131	-17.13
	0.1	13.08	-1.044	-18.04
0.001	0.01	25.66	-0.205	-17.20
	0.1	16.83	-1.343	-18.34
0.01	0.1	26.84	-2.141	-19.14