



American Finance Association

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

Author(s): Klaas P. Baks, Andrew Metrick, Jessica Wachter

Source: *The Journal of Finance*, Vol. 56, No. 1 (Feb., 2001), pp. 45-85

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/222463>

Accessed: 24/11/2009 12:08

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

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

KLAAS P. BAKS, ANDREW METRICK, and JESSICA WACHTER*

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.

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

* Baks and Metrick are from the Department of Finance, The Wharton School, University of Pennsylvania. Wachter is from the Department of Finance, The Stern School, New York University. We thank Nick Barberis, Gary Chamberlain, Ken French, Will Goetzmann, Karsten Hansen, Chris Jones, Tom Knox, Tony Lancaster, Ľuboš Pástor, André Perold, Steve Ross, Andrei Shleifer, Rob Stambaugh, René Stulz, Sheridan Titman, an anonymous referee, and seminar participants at Columbia, Wharton, the NBER, the 1999 NBER Summer Institute, and the 2000 AFA meetings for helpful comments. Wachter thanks Lehman Brothers for financial support.

¹ Recently, Carhart (1995), Malkiel (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), Lehman and Modest (1987), Grinblatt and Titman (1988, 1992), Hendricks, Patel, and Zechhauser (1993), Goetzmann and Ibbotson (1994), Brown and Goetzmann (1995), Elton, Gruber, and Blake (1996), and Carhart (1997).

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, it is not possible to distinguish between these two possibilities. Nevertheless, such small differences may have large consequences for investors.

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 nonbenchmark assets (actively managed mutual funds). We propose and employ a Bayesian method of performance evaluation; a key 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 and Stambaugh (2000), and Pástor (2000)). Like the latter two papers, our techniques build upon the work of Pástor and Stambaugh (1999); 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 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. Although 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

³ Gruber (1996) and Zheng (1999).

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

⁵ There is a related literature that employs Bayesian methods to explore the role of estimation risk on portfolio choice. See Klein and Bawa (1976), Bawa, Brown, and Klein (1979), Brown (1979), Jobson, Korkie, and Ratti (1979), Jobson and Korkie (1980), Jorion (1985, 1986, 1991), Frost and Savarino (1986), and Barberis (2000). In these applications, however, prior beliefs about parameters are typically noninformative or come from empirical Bayes procedures. Another related line of research focuses on the role of prior beliefs in model testing. See Shanken (1987), McCulloch and Rossi (1990), and Kandel, McCulloch, and Stambaugh (1995).

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 basis points?”, 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 1,437 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 quantify the economic importance of these results by estimating the portfolio share in active managers and the certainty-equivalent loss if this share is set to zero. We then discuss how to reconcile 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 nonbenchmark asset (an active investment manager). Under what conditions does the investor place any of her portfolio in the active manager? In this section, we derive the necessary condition for positive investment and find a counter-intuitive outcome for an investor who relies only on the data. This motivates a Bayesian performance-evaluation approach with the use of informed prior beliefs about managerial 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

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

where ε is normally distributed with mean zero and variance σ^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 the assumptions used in this paper, one can adopt a Bayesian procedure for estimating α and solve for w^* as proportional to the posterior mean of α , where this (positive) proportion would be a function of the inves-

tor's risk aversion and the posterior variance of manager returns. Then, the decision rule would be to invest in the active manager if and only if the posterior mean of α is positive. In this case, the decision rule is reduced to a Bayesian inference problem on α .

The implications of this simple decision rule are best illustrated through two polar examples. At one extreme, if the investor has diffuse (noninformative) prior beliefs for the parameters in equation (1), then posterior beliefs would be completely determined by the data and the posterior mean of α would be equal to its OLS estimate $\hat{\alpha}$. Thus, an investor with diffuse prior beliefs would invest with the active manager as long as $\hat{\alpha}$ is positive.

This result seems counterintuitive—shouldn't managers need more than just a positive point estimate before they merit an investment? Unease with this result may be related to the assumption of diffuse prior beliefs for α . In fact, some academics might tend towards the opposite extreme: a perfectly informed, or "dogmatic," belief in the impossibility of persistent managerial skill. Dogmatic prior beliefs do not allow the data to play any role and imply that investors should avoid all active managers, no matter how strong their record might be. These two extreme cases—diffusion and dogma—are easy to solve. The more interesting cases are in between, particularly when prior beliefs are so close to dogmatic beliefs that they cannot be distinguished by frequentist tests.

In this paper, we explicitly model the prior beliefs for α and solve the investor's problem. 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. 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 a dogmatic belief in the impossibility of skill.

The point mass in Figure 1 occurs at $\underline{\alpha} = \alpha - fee - cost < 0$. Here, we set α 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: α , his losses due to transactions with skilled managers, fee , his total fees, and $cost$, his transactions costs. While α is the same for all managers, the other components, fee and $cost$, generally differ across managers.

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




Figure 1. Prior distribution of α . 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 $\alpha > \underline{\alpha}$.

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 equation (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 $\underline{\alpha}$. Although many managers have return realizations that imply an $\hat{\alpha}$ below $\underline{\alpha}$, one would need to be systematically trading on nonpublic “misinformation” to have an *expected* α below $\underline{\alpha}$. If, instead, the market is not semistrong 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 behav-

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 judgment or the ability to exploit the behavioral anomalies of other investors, there is another manager who systematically does the opposite.

ioral biases are evenly distributed among all managers. Although this assumption does 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 . We solve this Bayesian inference problem in the next section.

II. Bayesian Performance Evaluation and Portfolio Choice

This section, along with Appendix A, 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 equation (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

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

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 equation (2). 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 . One can think of q as the probability of drawing a skilled manager from the population of managers. 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

$$p(\alpha, \beta, \sigma^2) = [p(\alpha|Z=0)P(Z=0) + p(\alpha|Z=1)p(Z=1)]p(\beta, \sigma^2). \quad (3)$$

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

$$p(\beta, \sigma^2) \propto \frac{1}{\sigma^2}. \quad (4)$$

The diffuse prior on β is necessary 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 50 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

$$P(Z=1) = q, \quad (5)$$

$$P(Z=0) = 1 - q, \quad (6)$$

$$p(\alpha|Z=0, \sigma^2) = \delta_\alpha, \quad (7)$$

$$p(\alpha|Z=1, \sigma^2) = 2N\left(\alpha, \sigma_\alpha^2 \left[\frac{\sigma^2}{s^2}\right]\right) \mathbf{1}_{\alpha > \alpha}, \quad (8)$$

where δ_x is the Dirac delta function with mass point at x , $\mathbf{1}_X$ is the indicator function for the set X , α 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.

The prior link between α and σ is first suggested by MacKinlay (1995) and is implemented in Pástor and Stambaugh (1999, 2000) and Pástor (2000). 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 nonbenchmark 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 equation (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 equation (8). 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 basis points? (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 transaction costs for the manager? (Call this answer *cost*.)

Note that all quantities are measured at the monthly frequency, so the qualifier “per month” is assumed every time a quantity is measured in basis points (bp). In answering these questions, the researcher should not consider any return-based information about the manager that coincides with the sample period under study. For example, 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 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.

⁸ The same argument can be used to motivate a link between β and σ . Because 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 α (and, by extension, $\underline{\alpha}$) to the ratio σ^2/s^2 . Unfortunately, this link is not tractable, so we model α as a constant. As long as α is small, however, the restriction is not quantitatively important.

Note that Question 2 is not asking about the probability of different realizations of $\hat{\alpha}$, but about “true” values of α . Realizations of $\hat{\alpha}$ depend on sampling variability. The true α , on the other hand, is the $\hat{\alpha}$ we would observe 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 zero as well. The use of 25 bp in this question is arbitrary, and any other point in the support of α could be substituted.

Question 2 is conditioned on a specific level of residual variance, s^2 . This conditioning is crucial, and 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:

$$q(25) = P(\alpha > 25 | \sigma^2 = s^2) = 2q \left(1 - \Phi \left(\frac{25 - \underline{\alpha}}{\sigma_\alpha} \right) \right), \quad (9)$$

$$\alpha = -q\sigma_\alpha \sqrt{\frac{2}{\pi}}, \quad (10)$$

and

$$\underline{\alpha} = \alpha - fee - cost, \quad (11)$$

which we can solve for the three unknowns, α , $\underline{\alpha}$, and σ_α . Equation (9) relates $q(25)$ to σ_α (given the other parameters), equation (10) imposes the constraint¹⁰ that the expectation of α , conditional on $\sigma^2 = s^2$, is equal to $(-fee - cost)$, and equation (11) 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 , α , and $\underline{\alpha}$. In most of the examples in Table I, α 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 set fee equal to zero on the right-hand-side of equation (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 equation (10) 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)).

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 $\underline{\alpha} = a - \text{fee} - \text{cost}$ 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 transaction costs, respectively. σ_α , a , and $\underline{\alpha}$ are expressed in basis points per month; $q(25)$ and q are expressed as probabilities.

$q(25)$	q	σ_α	a	$\underline{\alpha}$
Panel A: <i>fee</i> = 8 bp per month and <i>cost</i> = 6 bp per month				
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: <i>fee</i> = 8 bp per month and <i>cost</i> = 9 bp per month				
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

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 equations (12)–(20). The details are given beginning with equation (21) 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

$$\tilde{\alpha} = \tilde{q}E[\alpha|Z = 1, r, F] + (1 - \tilde{q})\underline{\alpha}. \quad (12)$$

The first term on the right-hand-side of equation (12) 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 equation (12) reflects the contribution coming from the possibility that the manager is unskilled; the posterior probability of no skill ($1 - \tilde{q}$) is multiplied by $\underline{\alpha}$.




Figure 2. Prior and posterior distribution of 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 $\underline{\alpha} = a - \text{fee} - \text{cost}$ 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. 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 $\alpha > \underline{\alpha}$.

To calculate $\tilde{\alpha}$, we need to solve for the two unknown elements on the right-hand side of equation (12): $E[\alpha | Z = 1, r, F]$ 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:

$$X \equiv (\iota_T \quad F), \quad (13)$$

$$\begin{pmatrix} \hat{\alpha} \\ \hat{\beta} \end{pmatrix} \equiv (X'X)^{-1}X'r, \quad (14)$$

$$m \equiv \text{top left element of } (X'X)^{-1}, \quad (15)$$

$$\text{var}(\hat{\alpha}) \equiv m\sigma^2. \quad (16)$$

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

$$p(\alpha|Z = 1, r, F, \sigma^2) \propto N(\alpha', \sigma'^2) \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (17)$$

where

$$\alpha' = \lambda \hat{\alpha} + (1 - \lambda) \underline{\alpha}, \quad (18)$$

$$\sigma'^2 = \left(\frac{1}{\text{var}(\hat{\alpha})} + \frac{1}{\sigma_\alpha^2 \left(\frac{\sigma^2}{s^2} \right)} \right)^{-1}, \quad (19)$$

$$\lambda = \frac{\sigma'^2}{\text{var}(\hat{\alpha})}. \quad (20)$$

Equations (17)–(20) are illustrated graphically in Figure 2. α' is the mode of the skilled posterior distribution given in equation (17); it would also be the mean, and σ'^2 the variance, of the untruncated version of this distribution. α' is written in equation (18) as a weighted average of the maximum likelihood estimate ($\hat{\alpha}$) and the prior mode ($\underline{\alpha}$), with weights given by λ and $1 - \lambda$, respectively. In equation (19), $\text{var}(\hat{\alpha})$ 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 equation (20) 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:

$$p(\alpha|Z = 1, r, F) \propto t_\nu \left(\alpha', \frac{\lambda m h}{\nu} \right) \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (21)$$

¹¹ See Appendix A for the details of these calculations.

where

$$h = S + \frac{1 - \lambda}{m} (\hat{\alpha} - \underline{\alpha})^2, \quad (22)$$

$$\hat{\theta} = \begin{pmatrix} \hat{\alpha} \\ \hat{\beta} \end{pmatrix}, \quad (23)$$

$$S \equiv (r - X\hat{\theta})'(r - X\hat{\theta}), \quad (24)$$

$$\nu = T - K. \quad (25)$$

The parameter λ has the same interpretation here as in equation (20); 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

$$E(\alpha|Z = 1, r, F) = \alpha' + \frac{\lambda mh}{\nu - 2} t_{\nu-2} \left(\underline{\alpha}; \alpha', \frac{\lambda mh}{\nu - 2} \right) \frac{1}{\int_{\underline{\alpha}}^{\infty} t_{\nu} \left(\alpha; \alpha', \frac{\lambda mh}{\nu} \right) d\alpha}. \quad (26)$$

The first term on the right-hand-side of equation (26) is just the mode of the skilled posterior distribution (the mean of the untruncated t -distribution), and 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 equation (12), $\tilde{q}$, the posterior probability that the manager is skilled. From Bayes' formula for binomial variables, it follows that

$$\tilde{q} \equiv P(Z = 1|r, F) = \frac{p(r|Z = 1, F)q}{p(r|Z = 1, F)q + p(r|Z = 0, F)(1 - q)}. \quad (27)$$

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

$$\tilde{q} = \frac{q}{q + \frac{1 - q}{B}}, \quad (28)$$

where B is given by

$$B \equiv \frac{p(r|Z = 1, F)}{p(r|Z = 0, F)}. \quad (29)$$

If observing the realized data is equally likely whether the manager is skilled or unskilled, then $B = 1$, and the 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,

$$B = \frac{t_{\nu-1}\left(\alpha; \hat{\alpha}, \frac{mS}{(1-\lambda)(\nu-1)}\right)}{t_{\nu-1}\left(\alpha; \hat{\alpha}, \frac{S}{\nu-1}m\right)} \left(2 \int_{\alpha}^{\infty} t_{\nu}\left(\alpha; \alpha', \frac{\lambda mh}{\nu}\right) d\alpha\right). \quad (30)$$

This is the ratio of two t -distributions, multiplied by a term to correct for the truncation at α .

Finally, with values for $E(\alpha|Z = 1, r, F)$ (equation (26)) and $\tilde{q}$ (equations (28) and (30)), we can substitute into equation (12) 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

$$p(r_j|\alpha_j, \beta_j, \sigma_j^2, F) = N(\alpha_j \iota_T + F\beta_j, \sigma_j^2 I_T). \quad (31)$$

To shorten notation, let $\psi_j = (\alpha_j, \beta_j, \sigma_j^2)$ and $\Psi = (\psi_1, \dots, \psi_N)$. For each manager, the prior on the parameters, $p(\psi_j)$, is given in Section II.B.

The key condition that allows us to extend our result to multiple managers is that “no manager conveys information on any other manager.” To obtain this condition, we make two assumptions:

ASSUMPTION 1: *The likelihoods are independent across managers:*

$$\begin{aligned} p(r|\Psi, F) &= \prod_j p(r_j|\Psi, F) \\ &= \prod_j p(r_j|\psi_j, F). \end{aligned} \quad (32)$$

¹² Note that by dividing the numerator and denominator of equation (28) by q , the posterior odds ratio $[(1-q)/q] \{ [p(r|Z = 0, F)] / [p(r|Z = 1, F)] \}$ 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 .

Assumption 1 does not mean that raw manager returns are independent, but rather that the factors capture all the dependencies. This is equivalent to stating that the residuals in equation (1) are independent across managers.

ASSUMPTION 2: *The priors are independent across managers:*

$$p(\Psi) = \prod_j p(\psi_j). \quad (33)$$

For some applications, Assumptions 1 and 2 would be problematic. Here, we feel that they are innocuous. Independence across managers would be a dangerous assumption if we intend to make strong statements about the total fraction of a portfolio invested in the full set of active managers. If we restrict ourselves to statements about portfolio shares when managers are considered one at a time, then inference about covariances is much less important. Of course, ignoring possible dependencies means that we lose some information, but there is no reason to believe that losing this information biases the results in either direction.

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

$$\begin{aligned} p(\Psi|r, F) &\propto p(r|\Psi, F)p(\Psi) \\ &= \prod_j p(r_j|\psi_j, F)p(\psi_j) \\ &\propto \prod_j p(\psi_j|r_j, F). \end{aligned} \quad (34)$$

Therefore, the calculation for the posterior of α_j when there are multiple managers is identical to the solution derived in Section 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.

Let $\tilde{E}$ and $\tilde{V}$ denote the expectation and variance-covariance matrix, respectively, of the predictive return distributions for the N active managers and the K index funds. Denote w and x as the vectors of weights, expressed as a share of invested wealth, on the active managers and index funds, respectively. Then, the weight on the risk-free asset is given by $1 - \sum_{i=1}^N w_i - \sum_{j=1}^K x_j$.¹³ Using these weights, we can calculate the mean and variance for the next-period return on any portfolio consisting of managers and index funds. We denote this mean and variance as $E[R_p]$ and $\text{Var}[R_p]$, where these expectations are taken with respect to the predictive return distribution of the managers and the factors.

¹³ If some of the index funds are zero-investment spread positions, then this equation would be modified. See Appendix C for a discussion.

The investor's problem is then to maximize over w and x

$$U = E[R_p] - \frac{1}{2} A \text{Var}[R_p] \quad (35)$$

where $A (> 0)$ is interpreted as the coefficient of relative risk aversion.

It is well known that the solution to this problem yields optimal weights w^* and x^* given by

$$\begin{pmatrix} w^* \\ x^* \end{pmatrix} = \frac{\tilde{V}^{-1} \tilde{E}}{A}, \quad (36)$$

In Appendix C, it is shown that the vector of weights on the managers is given by

$$w^* = \frac{\Omega^{-1}}{A} \tilde{\alpha}, \quad (37)$$

where Ω is a diagonal matrix with only positive elements. Thus, 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,” that is, 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

$$p(\psi_j | r, F) = p(\psi_j | r, F, \text{survival}_j). \quad (38)$$

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

ASSUMPTION 3: $p(\text{survival}_j | r, F, \psi_j) = p(\text{survival}_j | r, F)$.

Assumption 3 states that survival depends only on realized returns. Conditional on realized returns, the manager's skill (and the parameters β_j and σ_j) do not affect the probability of survival. Realized returns are, of course, observable, whereas ψ_j is unknown. It is quite plausible that survival depends on the manager's observed returns, 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 Section II.E.

Using Assumptions 1–3, it follows that survivor bias is not a problem for our analysis. In particular, by Bayes' rule:

$$p(\psi|r, F, \text{survival}) = \frac{p(\text{survival}|r, F, \psi)p(\psi|r, F)}{p(\text{survival}|r, F)} \quad (39)$$

$$= p(\psi|r, F). \quad (40)$$

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,

$$p(r|F, \psi, \text{survival}) \neq p(r|F, \psi). \quad (41)$$

That is, the likelihood conditional on survival is not the same as the likelihood without conditioning on survival. The prior on ψ conditional on survival also differs from the unconditional prior; in particular, values of ψ that increase the likelihood of survival would receive greater weight in the prior. Equations (39) and (40) 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.¹⁴ If, for example, we perform inference on the unconditional probability of skill in the population, $P(Z = 1)$, then the disappearance of poorly performing funds must affect the analysis. For this reason, we cannot provide guidance on the unconditional probability of skill for a population that includes both survivors and nonsurvivors. Alternatively, if we are missing data on a particular fund manager in years where that fund manager does particularly poorly, this would also bias our conclusions. But this is not the case in our sample.

III. Empirical Results

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

¹⁴ For studies of the implications for survivor bias on inference, see Brown et al. (1992), Brown, Goetzmann, and Ross (1995), and Goetzmann and Jorion (1997).

this sample. Section III.C contains the main analysis and answers the question posed in the title of the paper. Section III.D calibrates the economic significance of our findings by estimating the fraction of the portfolio allocated to active managers and computing the certainty-equivalent loss if investment in active managers is set to zero. Section III.E discusses the sensitivity of our analysis to alternative assumptions. Section III.F 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 1962 to 1996. We restrict ourselves to the subset of domestic diversified equity funds still operating at the end of 1996, and only include the monthly 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 this 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 1,437 funds with an average of 51 months of returns. This sample suffers from survivor bias and is not 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).¹⁵ The model is given by

$$r_{jt} = \alpha_j + \beta_{j1}RMRF_t + \beta_{j2}SMB_t + \beta_{j3}HML_t + \varepsilon_{jt}, \quad (42)$$

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, size, and value strategies.¹⁶ Although 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

¹⁵ The qualitative results do not change if we use the CAPM or the four-factor model of Carhart (1997).

¹⁶ See Fama and French (1993) for details on the construction of these portfolios. We are grateful to Ken French for providing the factor returns.

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. Although the benchmark returns in equation (42) are not themselves available as passive index funds, they are very similar to (combinations of) index products available in the late 1990s.¹⁷ 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 equation (42) for all 1,437 managers yields 705 managers with a positive $\hat{\alpha}$ and 732 with a negative $\hat{\alpha}$. Note that these estimates reflect performance after expenses, and thus show almost half of the managers succeeded in earning 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 meaningful 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, with 64 months of returns, has an $\hat{\alpha}$ of 92.1 bp, and a standard deviation for this estimate of 24.0 bp. 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 1,437 funds, the probability of observing an $\hat{\alpha}$ with a p -value less than or equal to 0.00014 is

$$1 - (1 - 0.00014)^{1,437} = 0.18, \quad (43)$$

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 equation (43), thus raising the probability of observing an extreme

¹⁷ 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, because momentum investing is, by nature, a high-turnover activity. This is the main reason we do not include a momentum benchmark in our analysis.

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.¹⁸ 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 in the database; we use a single value, 6 bp, as the *cost* for every manager.¹⁹ In Section III.E, we discuss the implications of changing this assumption.

For the denominator of the leverage term, given as s^2 in equation (8), we use the cross-sectional mean of the frequentist maximum-likelihood estimates of σ^2 for all funds that have at least 24 monthly observations. This empirical Bayes procedure yields $s^2 = 0.00029$. This level of s^2 is a useful normalization that makes it easier to interpret the results, because 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 bp, after transactions costs but before fees, and conditional on the average level of residual risk in the sample. We use this definition of $q(25)$ for interpreting our results.

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 24.1 bp. The standard error on this α estimate is 8.6 bp. If we perform a frequentist test of the null hypothesis that α is zero, we obtain a p -value of 0.003.

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 monthly transactions costs are 6 bp, the prior mean for Guardian's α is

¹⁸ Chevalier and Ellison (1999) provide evidence that many such characteristics are correlated with alphas.

¹⁹ 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.




Figure 3. Guardian Park Avenue Fund: $\tilde{\alpha}$ as a function of $q(25)$ when $q = 0.01$. α is the intercept in the Fama and French (1993) three-factor model (see equation (42)). $\tilde{\alpha}$ 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. $q(25) \equiv P(\alpha > 25 | \sigma^2 = s^2)$. The prior constant s^2 is 0.00029. The sample period for Guardian Park Avenue is July 1972 to December 1996.

equal to -12.8 bp.²⁰ (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 $q = 0.01$. Thus, the investor believes that one percent of all managers are expected to have some skill. Holding q constant at 0.01, 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. Recall that α adjusts for different levels of $q(25)$, so that the prior mean of α is always equal to -12.8 bp. The figure covers the range $q(25) \in [0, 0.001]$. For the very smallest levels of $q(25)$, $\tilde{\alpha}$ is weighted heavily towards the prior mean

²⁰ Guardian also had a maximum load fee of 450 bp in 1996. Many of the other top performing funds do not charge any load fees.




Figure 4. The highest (solid) and tenth highest (dashed) $\tilde{\alpha}$ over all managers when $q = 0.01$. This figure shows the highest and tenth highest $\tilde{\alpha}$ over all 1,437 managers for $q(25) \in [0, 0.001]$ and $q = 0.01$. α is the intercept in the Fama and French (1993) three-factor model (see equation (42)). $\tilde{\alpha}$ is the posterior expectation of α . $q(25) \equiv \Pr(\alpha > 25 | \sigma^2 = s^2)$. 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.

and is negative. This negative range is very narrow, however, and $\tilde{\alpha}$ is positive for all values of $q(25)$ greater than 0.00006. By the time we reach $q(25) = 0.001$, $\tilde{\alpha}$ is about 13 bp. Thus, if an investor believes that one manager in a 100 has skill, and one in 1,000 has sufficient skill so that their α is greater than 25 bp, then the posterior mean for Guardian would be 13 bp.

As impressive as Guardian's performance is, it does not provide the highest $\tilde{\alpha}$ in the sample for this range of prior beliefs. 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 $q(25) \in [0, 0.001]$, holding q constant at

0.01.²¹ As seen in the figure, the best performing manager's $\tilde{\alpha}$ becomes positive at about $q(25) = 0.00003$, and the tenth best manager becomes positive at about $q(25) = 0.00083$. At $q(25) = 0.001$, the best performing manager has an $\tilde{\alpha}$ of 17 bp, and the tenth best performing manager has an $\tilde{\alpha}$ of 2 bp.

In Figure 4, the use of a fixed $q = 0.01$ is done only to provide an illustrative example. We could also draw this figure for any other level of q . In every case, if $q > 0.0002$ and $q(25) > 0.0001$, there is at least one manager with $\tilde{\alpha} > 0$. Thus, as long as the investor believes that at least two in 10,000 managers has skill, and one in 10,000 managers has an α of at least 25 bp, then she invests in at least one manager. In a frequentist test, such low 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 is due to chance (in equation (43)), this test has very little power against the alternative that $q = 0.0002$ and $q(25) = 0.0001$. Thus, we conclude that zero investment in active managers cannot be justified solely on the basis of the available statistical evidence.

D. Economic Significance

In the previous section, we showed that an investor with weak prior beliefs in the possibility of skill would still choose to invest some of her portfolio in active managers. In this section, we calibrate the approximate size and economic significance of the investor's position in active managers. This calibration follows the procedures introduced by Pástor and Stambaugh (2000). In our analysis, we ignore load fees, taxes, margin requirements, and short sale constraints on the index funds, so our results should be interpreted with these caveats. Our intention is not to provide specific investment advice, but rather to estimate the economic significance of our findings in a way that can be compared with other studies.

We begin again with the investor's problem as discussed in Section II.E. To complete the problem, we specify the benchmark assets as the three factors used in equation (42): *RMRF*, *HML*, and *SMB*. Because each of the factors is composed of both a long and a short position, investments in them are all zero-cost spread positions. For simplicity, we assume that there are no margin requirements, so that the investor can take positions of any size. Weights on these factors are then expressed as a percentage of total invested wealth. For example, optimal weights of 150 on each factor would correspond to a spread position on each factor of \$1.50 for each dollar of invested

²¹ 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 equation (43).

wealth.²² For the means and variances of the factor returns, we use the predictive moments calculated from the monthly returns of July 1963 to December 1996 (See Appendix B for these procedures). Following the same normalization as Pástor and Stambaugh (2000), we set the risk-aversion parameter so that an investor choosing only an optimal level of *RMRF*—with the other factors set to zero—would choose to be “fully invested”; that is, *RMRF* equal to 100 with an implied zero position in the risk-free asset. This normalization is $A = 2.47$ over our sample period.

To keep our analysis conservative, we restrict the investor to choosing no more than one manager. With this constraint, our assumption of independent manager returns (Assumption 1) does not inflate the total weight placed on active managers.²³ The remainder of the analysis then exactly follows the solution derived in Section II.E and Appendix C. For every set of prior beliefs, we begin by computing $\tilde{\alpha}$ for all 1,437 managers and finding the set of managers that have $\tilde{\alpha} > 0$. Then, for each manager in this set, we separately compute the optimal weights in equation (36) for a choice problem among that manager, the risk-free asset, and the three factors. For example, if 10 managers have $\tilde{\alpha} > 0$ for some level of prior beliefs, then we solve 10 different portfolio-choice problems, each time considering one manager along with the other assets. We refer to the “optimal portfolio” as the maximum-utility portfolio among all these solutions; we refer to the manager held in this optimal portfolio as the “best manager.” This computation requires simulated draws from the predictive distributions; these procedures are discussed in Appendix B. Note that the best manager is not necessarily the manager with the highest $\tilde{\alpha}$ from Figure 4, as the variance of portfolio returns also affects the utility of the optimal portfolio.

The results are summarized in Figures 5 and 6. Figure 5 plots the optimal weight on the top manager and the corresponding weight on *RMRF* for the same prior parameters as used in Figures 3 and 4: $q = 0.01$, $q(25) \in [0, 0.001]$. We interpret the optimal weight on *RMRF* as the residual weight on a “market-index fund.” For very low levels of $q(25)$, there is no manager in the optimal portfolio; this range corresponds exactly to the $\tilde{\alpha} < 0$ range from Figure 4. Over this range, the weight on *RMRF* is 161; that is, a \$1.61 spread position for each dollar of invested wealth. The weights on *HML* and *SMB*, not shown in the figure, are 377 and 70, respectively. For levels of $q(25) > 0.00003$, there is at least one manager with $\tilde{\alpha} > 0$, and so there is positive investment in the best manager. This weight rises, and the weight on *RMRF* falls, with $q(25)$. At $q(25) = 0.00005$, there is a weight of 175 on the best manager and zero weight on *RMRF*. At this point, the

²² Because the factors are all zero-investment positions, the “residual” weight on the risk-free asset is calculated as 100 minus the weight on the active manager. Because *RMRF* includes a short position in the risk-free asset, the “actual” weight on the risk-free asset is the residual weight minus the weight on *RMRF*.

²³ Without this constraint, the investor would perceive investment in multiple managers as a diversification of independent risks, and every manager with positive $\tilde{\alpha}$ would have positive investment.




Figure 5. Weight on the best manager (solid) and the corresponding weight on *RMRF* (dashed) when $q = 0.01$. This figure shows the weight (expressed as a percentage of invested wealth) in the best manager (solid) and the corresponding weight on *RMRF* (dashed) as a function of $q(25)$ for $q(25) \in [0, 0.001]$ and $q = 0.01$. Weights are the solution to the portfolio-choice problem in equation (35), where the assets are a single manager, the risk-free asset, and the three factors from equation (42). For each level of $q(25)$, we solve equation (35) separately for each manager that has positive $\tilde{\alpha}$. The “best manager” is the manager held in the portfolio that yields the highest-utility solution to this problem. The investor’s coefficient of relative risk aversion is set to $A = 2.47$. α is the intercept and *RMRF* is the market factor in the Fama and French (1993) three-factor model (see equation (42)). 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. $q(25) \equiv P(\alpha > 25 | \sigma^2 = s^2)$. The prior constant s^2 is 0.00029.

investor would only take on market risk through her investment in the active manager—she would take no additional position in the broad market index (*RMRF*). At $q(25) = 0.001$, the weights on the best manager and on *RMRF* would be 249 and -66 respectively.²⁴

²⁴ Figure 5 plots x_{RMRF}^* , the optimal weight on *RMRF* and w^* , the weight on the best manager. Note that the “actual” weight on *RMRF* would be $x_{RMRF}^* + \beta_1^* w^*$: the optimal weight shown in Figure 5 plus a component due to the factor loading of the best manager on *RMRF*. This actual weight is constant at 161 for all values of $q(25)$.




Figure 6. Certainty equivalent loss for a manager-restricted investor when $q = 0.01$.

This figure shows the certainty equivalent loss expressed in basis points per month when the investor is precluded from investment in all managers. This loss is plotted as a function of $q(25)$ for $q(25) \in [0, 0.001]$ and $q = 0.01$. The loss is defined as the difference in certainty equivalents between the portfolio with the best manager (as plotted in Figure 5) and a constrained-optimum portfolio with all managers set to zero. The investor's coefficient of relative risk aversion is set to $A = 2.47$. α is the intercept and $RMRF$ is the market factor in the Fama and French (1993) three-factor model (see equation (42)). 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. $q(25) \equiv P(\alpha > 25 | \sigma^2 = s^2)$. The prior constant s^2 is 0.00029.

Another way to calibrate the economic significance of these results is to compute the certainty equivalent (CE) loss to an investor if she is precluded from any investment in active managers. To do this, we calculate the difference between the CE return for the optimal portfolio used for Figure 5 and the CE return for a “manager-restricted” investor who is allowed to take positions in the factors but not in any manager. This difference is plotted in Figure 6. For the zero-investment range of $q(25)$, the manager-restricted investor has no CE loss. At $q(25) = 0.0005$, a manager-restricted investor would suffer a CE loss of about 8 bp per month. At $q(25) = 0.001$, the CE loss rises to 16 bp per month. As a comparison, we calculate the CE loss if the

investor is also restricted from investments in *HML* and *SMB*. The loss from this additional restriction would be approximately 104 bp per month. Thus, although the returns to investing in the best manager seem economically significant, they are still much lower than the returns to investing in the size and value factors.

E. Sensitivity of Analysis to Modeling Assumptions

Our analysis assumes that *cost* is known and equal to 6 bp for all managers. How sensitive are our results to this assumption? Suppose for example, that we use 9 bp as our baseline level. Then, for any given level of q and $q(25)$, α 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 $\bar{\alpha}$ 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. Although we cannot obtain analytical solutions using a fatter-tailed likelihood, we can analyze a related question: Could the results of Section III be driven by incorrectly applying our methods to fat-tailed data? To answer this question, we simulate 10 years of returns for 1,000 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 has any skill. We then replicate Figure 4 for these data. The results show very little difference across the three return distributions.²⁵

The results of this section are based on a specific performance-evaluation model, but other popular models lead to the same qualitative conclusions.²⁶ 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, induces bias in

²⁵ Details of these simulations are available from the authors. Although the results are not directly comparable to those of the actual data set, it is interesting to note that for $q = 0.01$ (as in Figure 4) no simulated manager has a positive $\bar{\alpha}$ for any $q(25) < 0.0003$.

²⁶ Results for the CAPM and the four-factor model (Carhart (1997)) are available from the authors.

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 conclusions would change: The prior beliefs necessary to support investment in active managers are virtually indistinguishable from either "no skill" or "no persistence of skill."

F. Comparison of Frequentist and Bayesian Results

In Section III.C, we implement a Bayesian approach 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 does 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 equation (43), the null hypothesis is that α is zero for all managers. 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 then pays 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 to reject a null hypothesis, they may still be sufficiently extreme to justify their expenses.

IV. Conclusion

Should investors avoid all actively managed mutual funds? A natural frequentist approach to this question is in three steps: (step 1) My null hypothesis is that no manager has skill; (step 2) The data do not reject this null hypothesis; (step 3) I will not invest in active managers. Although this may seem like a reasonable approach, it does not have a sound decision-theoretic justification; the evidence necessary to reject the null hypothesis in step 2 is different from the evidence necessary to justify investment in active managers. Hence, step 2 does not imply step 3. Current data and methods have insufficient power to distinguish between the null hypothesis in step 2 and

close alternatives. The main contribution of our paper is to show that some of these close alternatives imply economically large investment in active managers. Thus, we conclude that the case against investing in actively managed funds cannot rest solely on the available 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-based intuition.

Appendix A. Posterior Distribution and Expectation of α

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

$$p(\beta) \propto 1, \quad (\text{A1})$$

$$p(\sigma^2) \propto \frac{1}{\sigma^{\nu_0+2}} \exp \left\{ -\frac{h_0}{2\sigma^2} \right\}. \quad (\text{A2})$$

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:

$$p(F_t | \mu_F, \Sigma_F) = N(\mu_F, \Sigma_F), \quad (\text{A3})$$

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

$$p(\mu_F, \Sigma_F) \propto |\Sigma_F|^{-(K+1)/2} \quad (\text{A4})$$

A.1. Posterior Distribution of α

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

$$\begin{aligned} p(\theta, \sigma^2, \mu_F, \Sigma_F | r, F) &\propto p(r | \theta, \sigma^2, F) p(F | \mu_F, \Sigma_F) p(\theta, \sigma^2) p(\mu_F, \Sigma_F) \\ &\propto p(\theta, \sigma^2 | r, F) p(\mu_F, \Sigma_F | F), \end{aligned} \quad (\text{A5})$$

where we use 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 , whereas 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

$$p(\theta, \sigma^2 | r, F) \propto p(r | \theta, \sigma^2, F) p(\theta, \sigma^2). \quad (\text{A6})$$

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 make use of equation (A6) throughout this appendix.

The likelihood for r conditional on factors is given by

$$\begin{aligned} p(r | \theta, \sigma^2, F) &\propto \frac{1}{\sigma^T} \exp \left\{ -\frac{1}{2\sigma^2} (r - X\theta)' (r - X\theta) \right\} \\ &\propto \frac{1}{\sigma^T} \exp \left\{ -\frac{1}{2\sigma^2} (S + (\theta - \hat{\theta})' X' X (\theta - \hat{\theta})) \right\}, \end{aligned} \quad (\text{A7})$$

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

Combining the likelihood and the prior yields

$$\begin{aligned} p(\theta, \sigma^2 | Z = 1, r, F) &\propto \frac{1}{\sigma^{\nu_0 + T + 2}} \frac{1}{\sigma} \\ &\times \exp \left\{ -\frac{1}{2\sigma^2} \left(h_0 + S + \frac{s^2}{\sigma_\alpha^2} (\alpha - \underline{\alpha})^2 + (\theta - \hat{\theta})' X' X (\theta - \hat{\theta}) \right) \right\} \mathbf{1}_{\alpha > \underline{\alpha}}. \end{aligned} \quad (\text{A8})$$

Integrating with respect to β yields

$$\begin{aligned} p(\alpha, \sigma^2 | Z = 1, r, F) &\propto \frac{1}{\sigma^{\nu_0 + T - K + 2}} \frac{1}{\sigma} \\ &\times \exp \left\{ -\frac{1}{2\sigma^2} \left(h_0 + S + \frac{s^2}{\sigma_\alpha^2} (\alpha - \underline{\alpha})^2 + \frac{(\alpha - \hat{\alpha})^2}{m} \right) \right\} \mathbf{1}_{\alpha > \underline{\alpha}}. \end{aligned} \quad (\text{A9})$$

Completing the square in α yields

$$p(\alpha, \sigma^2 | Z = 1, r, F) \propto \frac{1}{\sigma^{\nu+2}} \frac{1}{\sigma} \exp \left\{ -\frac{1}{2\sigma^2} \left(h + \frac{(\alpha - \alpha')^2}{\lambda m} \right) \right\} \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (\text{A10})$$

where

$$h = h_0 + S + \left(\frac{1 - \lambda}{m} \right) (\hat{\alpha} - \underline{\alpha})^2,$$

$$\nu = \nu_0 + T - K. \quad (\text{A11})$$

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,

$$p(\alpha | \sigma^2, Z = 1, r, F) \propto \exp \left\{ -\frac{1}{2\sigma'^2} (\alpha - \alpha')^2 \right\} \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (\text{A12})$$

and we have shown equation (17).

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 equation (A10) using the properties of the gamma distribution (see, e.g., Gelman et al. (1995)). Define a change of variables $u = \frac{1}{2} D \sigma'^{-2}$, where $D = h + (\lambda m)^{-1} (\alpha - \alpha')^2$. The resulting function of u is the pdf of a gamma distribution without the normalizing constant:

$$p(\alpha | Z = 1, r, F) \propto \left(\frac{D}{2} \right)^{-(\nu+1)/2} \left(\int u^{(\nu-1)/2} \exp\{-u\} du \right) \mathbf{1}_{\alpha > \underline{\alpha}}. \quad (\text{A13})$$

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

$$p(\alpha | Z = 1, r, F) \propto \left(\frac{h}{2} \right)^{-(\nu+1)/2} \Gamma \left(\frac{\nu+1}{2} \right) \left(1 + \frac{(\alpha - \alpha')^2}{\lambda m h} \right)^{-(\nu+1)/2} \mathbf{1}_{\alpha > \underline{\alpha}}, \quad (\text{A14})$$

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

$$p(\alpha | Z = 1, r, F) \propto t_\nu \left(\alpha', \frac{\lambda m h}{\nu} \right) \mathbf{1}_{\alpha > \underline{\alpha}}. \quad (\text{A15})$$

A.2. Posterior Expectation of α Conditional on Skill

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

$$p(\alpha | Z = 1, r, F) = \frac{1}{\int_{\underline{\alpha}}^{\infty} t_\nu \left(\alpha; \alpha', \frac{\lambda m h}{\nu} \right) d\alpha} t_\nu \left(\alpha; \alpha', \frac{\lambda m h}{\nu} \right) \mathbf{1}_{\alpha > \underline{\alpha}}. \quad (\text{A16})$$

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

$$\begin{aligned}
 E[\alpha | Z = 1, r, F] &= \frac{1}{\tilde{t}} \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma(\nu/2)\sqrt{\lambda m h \pi}} \int_{\underline{\alpha}}^{\infty} \alpha \left(1 + \frac{(\alpha - \alpha')^2}{\lambda m h}\right)^{-(\nu+1)/2} d\alpha \\
 &= \alpha' + \frac{1}{\tilde{t}} \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{(\nu-1)\Gamma(\nu/2)} \sqrt{\frac{\lambda m h}{\pi}} \left(1 + \frac{(\underline{\alpha} - \alpha')^2}{\lambda m h}\right)^{-(\nu-1)/2} \\
 &= \alpha' + \frac{\lambda m h}{\nu-2} t_{\nu-2}\left(\underline{\alpha}; \alpha', \frac{\lambda m h}{\nu-2}\right) \frac{1}{\int_{\underline{\alpha}}^{\infty} t_{\nu}\left(\alpha; \alpha', \frac{\lambda m h}{\nu}\right) d\alpha}. \quad (A17)
 \end{aligned}$$

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

A.3. Derivation of B

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

$$p(r | Z = 1, F) = \int p(r | Z = 1, F, \theta, \sigma^2) p(\theta, \sigma^2 | Z = 1) d\theta d\sigma^2. \quad (A18)$$

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

$$\begin{aligned}
 p(r | Z = 1, F) &= C \sqrt{\frac{2}{\pi}} \frac{s}{\sigma_{\alpha}} \int_0^{\infty} \int_{\underline{\alpha}}^{\infty} \frac{1}{\sigma^{\nu+2}} \frac{1}{\sigma} \\
 &\quad \times \exp\left\{-\frac{1}{2\sigma^2} \left(h_0 + S + \frac{s^2}{\sigma_{\alpha}^2} (\alpha - \underline{\alpha})^2 + \frac{(\alpha - \hat{\alpha})^2}{m}\right)\right\} d\alpha d\sigma^2, \quad (A19)
 \end{aligned}$$

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

$$\begin{aligned}
 p(r | Z = 1, F) &= C \sqrt{\frac{2}{\pi}} \frac{s}{\sigma_{\alpha}} \int_0^{\infty} \int_{\underline{\alpha}}^{\infty} \frac{1}{\sigma^{\nu+2}} \frac{1}{\sigma} \\
 &\quad \times \exp\left\{-\frac{1}{2\sigma^2} \left(h + \frac{(\alpha - \alpha')^2}{\lambda m}\right)\right\} d\alpha d\sigma^2. \quad (A20)
 \end{aligned}$$

The form of the equation is the same as in equation (10). As above, the properties of the gamma distribution are used to integrate out σ^2 :

$$p(r|Z = 1, F) = C \sqrt{\frac{2}{\pi}} \frac{s}{\sigma_\alpha} \left(\frac{h}{2}\right)^{-(\nu+1)/2} \Gamma\left(\frac{\nu+1}{2}\right) \int_{\underline{\alpha}}^{\infty} \left(1 + \frac{(\alpha - \alpha')^2}{hm\lambda}\right)^{-(\nu+1)/2} d\alpha. \quad (\text{A21})$$

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

$$p(r|Z = 1, F) = 2C\sqrt{1-\lambda} \left(\frac{h}{2}\right)^{-\nu/2} \Gamma\left(\frac{\nu}{2}\right) \int_{\underline{\alpha}}^{\infty} t_\nu\left(\alpha; \alpha', \frac{\lambda mh}{\nu}\right) d\alpha. \quad (\text{A22})$$

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

$$p(r|Z = 0, F) = C \int_0^\infty \frac{1}{\sigma^{\nu+2}} \exp\left\{-\frac{1}{2\sigma^2} \left(S + h_0 + \frac{(\hat{\alpha} - \underline{\alpha})^2}{m}\right)\right\} d\sigma^2. \quad (\text{A23})$$

Integrating with respect to σ^2 yields

$$p(r|Z = 0, F) = C \left(\frac{S + h_0 + (\hat{\alpha} - \underline{\alpha})^2/m}{2}\right)^{-\nu/2} \Gamma\left(\frac{\nu}{2}\right). \quad (\text{A24})$$

Therefore,

$$B = \sqrt{1-\lambda} \left(\frac{S + h_0 + \frac{1-\lambda}{m} (\hat{\alpha} - \underline{\alpha})^2}{S + h_0 + m^{-1}(\hat{\alpha} - \underline{\alpha})^2}\right)^{-\nu/2} 2 \int_{\underline{\alpha}}^{\infty} t_\nu\left(\alpha; \alpha', \frac{\lambda mh}{\nu}\right) d\alpha. \quad (\text{A25})$$

Both the numerator and the denominator are proportional to a t -distribution with $\nu - 1$ degrees of freedom (in one case, α is known, whereas in the other case we integrate with respect to α). Multiplying and dividing by a constant yields expression (30) in the text.

Appendix B. Predictive Return and Factor Distribution

This section describes how to draw from the predictive distribution of r and F and derives an expression for their first and second moments.

B.1. Drawing from the Posterior Distribution of $(\alpha, \beta, \sigma^2)$

We first show how to draw from the posterior distribution of $(\alpha, \beta, \sigma^2)$. We describe how to draw from (1) the skilled distribution, (2) the unskilled distribution, and (3) the full posterior using (1) and (2).

(1) Drawing from the distribution of $(\alpha, \beta, \sigma^2)$ conditional on $Z = 1$ ("skilled posterior"): Consider random variables $\check{\sigma}^2$ and $\check{\alpha}$ such that

$$\check{\sigma}^2 | r, F, Z = 1 \sim IG\left(\frac{\nu}{2}, \frac{h}{2}\right), \quad (\text{B1})$$

$$\check{\alpha} | \sigma^2 = \check{\sigma}^2, r, F, Z = 1 \sim N(\alpha', \sigma'^2), \quad (\text{B2})$$

where IG denotes an inverse gamma density (see, e.g., Gelman et al. (1995)). Then

$$p(\check{\alpha}, \check{\sigma}^2 | r, F, Z = 1) \propto \frac{1}{\check{\sigma}^{\nu+2}} \frac{1}{\check{\sigma}} \exp\left\{-\frac{1}{2\check{\sigma}^2} (h + (\lambda m)^{-1}(\check{\alpha} - \alpha')^2)\right\}. \quad (\text{B3})$$

Using the procedure above, but discarding the draw whenever $\check{\alpha} \leq \underline{\alpha}$ produces a distribution that has zero mass when $\check{\alpha} \leq \underline{\alpha}$, but where the relative densities of any other points are the same as in equation (B3). By equation (A10), this is exactly the joint distribution of α and σ^2 . Therefore, drawing $\check{\sigma}^2$ from $IG(\nu/2, h/2)$ and $\check{\alpha} | \sigma^2 = \check{\sigma}^2$ from $N(\alpha', \sigma'^2)$ and discarding the draws whenever $\check{\alpha} \leq \underline{\alpha}$ produces a draw from the joint posterior of α and σ .

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

$$\beta | \alpha, \sigma^2, r, F, Z = 1 \sim N(\hat{\beta} + m^{-1}y(\alpha - \hat{\alpha}), \sigma^2(F'F)^{-1}), \quad (\text{B4})$$

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

$$\begin{pmatrix} m & y' \\ y & Q \end{pmatrix}. \quad (\text{B5})$$

(2) Drawing from the posterior distribution of $(\alpha, \beta, \sigma^2)$ conditional on $Z = 0$ ("unskilled posterior"): It follows from the likelihood and the prior, that conditional on $Z = 0$, we have

$$\sigma^2 | Z = 0, r, F \sim IG\left(\frac{\nu}{2}, \frac{h_0 + S + m^{-1}(\hat{\alpha} - \underline{\alpha})^2}{2}\right). \quad (\text{B6})$$

Conditional on σ^2 and on $\alpha = \underline{\alpha}$, β is drawn from equation (B4).

(3) Drawing from the full posterior of $(\alpha, \beta, \sigma^2)$: For any given draw $(\alpha, \beta, \sigma^2)$, 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 be a draw from the distribution with uniform mass on $[0, 1]$. If $u < \tilde{q}$, then $(\alpha, \beta, \sigma^2)$ is drawn from the skilled posterior. Otherwise, a draw is made from the unskilled posterior.

B.2. Drawing from the Posterior Distribution of (μ_F, Σ_F)

Let L be the number of periods for which factor data is available and

$$\hat{\mu}_F = \frac{1}{L} \sum_{t=1}^L F_t, \quad (\text{B7})$$

$$\hat{\Sigma}_F = \frac{1}{L} \sum_{t=1}^L (F_t - \hat{\mu}_F)(F_t - \hat{\mu}_F)'. \quad (\text{B8})$$

Combining the likelihood (A3) and prior (A4) for the factors, it follows (see, e.g., Gelman et al. (1995)) that the posterior distribution of Σ_F^{-1} is a Wishart distribution with parameter matrix $(L\hat{\Sigma}_F)^{-1}$ and $L - 1$ degrees of freedom. Conditional on Σ_F , μ_F is normally distributed with mean $\hat{\mu}_F$ and variance Σ_F/L . Thus to draw from the posterior distribution of μ_F , Σ_F , first draw Σ_F from its inverted Wishart posterior distribution, and then draw μ_F from its normal posterior distribution conditional on Σ_F .

B.3. Drawing from the Predictive Distribution of r and F

From the likelihood:

$$p(r_{T+1} | \alpha, \beta, \sigma^2, F_{T+1}) = N(\alpha + F_{T+1}\beta, \sigma^2), \quad (\text{B9})$$

and

$$p(F_{T+1} | \mu_F, \Sigma_F) = N(\mu_F, \Sigma_F). \quad (\text{B10})$$

Because

$$\begin{aligned} p(r_{T+1} | r, F) &= \int_{F_{T+1}} \int_{(\alpha, \beta, \sigma^2)} p(r_{T+1} | \alpha, \beta, \sigma^2, F_{T+1}) p(\alpha, \beta, \sigma^2 | r, F) d\alpha d\beta d\sigma^2 \\ &\times \int_{(\mu_F, \Sigma_F)} p(F_{T+1} | \mu_F, \Sigma_F) p(\mu_F, \Sigma_F | F) d\mu_F d\Sigma_F dF_{T+1}, \end{aligned} \quad (\text{B11})$$

we can obtain a draw from predictive distribution of r using the following procedure. First, draw (μ_F, Σ_F) from their posterior distribution, then draw F_{T+1} from equation (B10), conditional on those values of μ_F and Σ_F . Next, draw $(\alpha, \beta, \sigma^2)$ from the posterior in Appendix B.1. Finally, draw r_{T+1} from equation (B9), conditional on the values of the parameters, and on F_{T+1} .

B.4. Moments of the Predictive Distribution of r and F

This section shows how to efficiently calculate the predictive mean and joint variance–covariance matrix of the manager returns and the factors. This is a necessary step for computing the portfolio weights as described in Appendix C. The first moment of the joint predictive distribution of r and F is obtainable analytically—the second moment is only partially obtainable analytically. By equation (B10) and the law of iterated expectations

$$\begin{aligned} E[F_{T+1}|F] &= E[E[F_{T+1}|\mu_F, \Sigma_F, F]|F] \\ &= E[\mu_F|F] \equiv \tilde{\mu}_F = \hat{\mu}_F. \end{aligned} \quad (\text{B12})$$

Again, using the law of iterated expectations and using equation (B9) and the posterior independence of (μ_F, Σ_F) and $(\alpha, \beta, \sigma^2)$, it follows that

$$\begin{aligned} E[r_{T+1}|r, F] &= E[E[r_{T+1}|\alpha, \beta, \sigma^2, \mu_F, \Sigma_F, r, F]|r, F] \\ &= E[\alpha + \mu_F \beta|r, F] \\ &= \tilde{\alpha} + \tilde{\mu}_F \tilde{\beta} = \tilde{\alpha} + \hat{\mu}_F \tilde{\beta}, \end{aligned} \quad (\text{B13})$$

where $\tilde{\beta} \equiv E[\beta|r, F] = \hat{\beta} + (\tilde{\alpha} - \hat{\alpha})y/m$. By the covariance decomposition for conditional distributions,

$$\text{Cov}[x, y] = E[\text{Cov}(x, y|z)] + \text{Cov}[E(x|z), E(y|z)], \quad (\text{B14})$$

and the properties of the inverted Wishart distribution, it follows that

$$\begin{aligned} \text{Var}[F_{T+1}|F] &= \text{Var}[E(F_{T+1}|\mu_F, \Sigma_F, F)|F] + E[\text{Var}(F_{T+1}|\mu_F, \Sigma_F, F)|F] \\ &= \text{Var}[\mu_F|F] + E[\Sigma_F|F] \\ &= \left(\frac{1}{L} + 1\right)E[\Sigma_F|F] \\ &= \frac{L+1}{L-K-2} \hat{\Sigma}_F, \end{aligned} \quad (\text{B15})$$

Again, using the covariance decomposition for conditional distributions and using equations (B12) and (B13), it follows that

$$\begin{aligned}
 \text{Cov}[r_{T+1}, F_{T+1}|r, F] &= \text{Cov}\{E[r_{T+1}|\alpha, \beta, \sigma^2, \mu_F, \Sigma_F, r, F], \\
 &\quad E[F_{T+1}|\alpha, \beta, \sigma^2, \mu_F, \Sigma_F, r, F]|r, F\} \\
 &\quad + E[\text{Cov}[r_{T+1}, F_{T+1}|\alpha, \beta, \sigma^2, \mu_F, \Sigma_F, r, F]|r, F] \\
 &= \text{Cov}[\alpha + \mu_F \beta, \mu_F \beta|r, F] + E[\Sigma_F \beta|r, F] \\
 &= (\text{Var}[\mu_F|F] + E[\Sigma_F|F])\tilde{\beta} = \text{Var}[F_{T+1}|F]\tilde{\beta} \\
 &= \frac{L+1}{L-K-2} \hat{\Sigma}_F \tilde{\beta}.
 \end{aligned} \tag{B16}$$

The only remaining part of the joint predictive variance–covariance matrix is the variance of r_{T+1} , which is obtained by simulation as outlined in Appendix B.3.

Appendix C. The Portfolio-Choice Problem and the Positive-investment Condition

In this section we show that if an asset has a positive posterior expectation of α (i.e., $\tilde{\alpha} > 0$), then this asset is held in positive quantities in the mean-variance tangency portfolio. In addition, we derive the results stated in Section II.E.

The next-period rate of return on a portfolio consisting of the factors and the managers is

$$R_p = \sum_{i=1}^N w_i r_{i,T+1} + \sum_{j=1}^K x_j F_{j,T+1} + r_f, \tag{C1}$$

where w_i is the share of period T wealth invested in manager i , x_j is the share of period T wealth invested in factor j , and r_f is the return on the risk-free asset in period $T+1$. The weight on the risk-free asset is given by $1 - \sum_{i=1}^N w_i - \sum_{j=1}^K x_j$. If index fund j is a zero-investment spread position (as in Section III), then it receives a weight of x_j on its long component and $-x_j$ on its short component, for a net contribution of zero towards the risk-free asset.

The mean-variance investor's problem is

$$\max_{\omega} E[R_p] - \frac{A}{2} \text{Var}[R_p], \tag{C2}$$

where $\omega \equiv (w \ x)' \equiv (w_1 \cdots w_N \ x_1 \cdots x_K)'$ are the portfolio weights on the managers and index funds. This optimization problem can be rewritten as

$$\max_{\omega} \omega' \tilde{E} - \frac{A}{2} \omega' \tilde{V} \omega, \quad (\text{C3})$$

with the well-known solution given by

$$\omega^* \equiv \begin{pmatrix} w^* \\ x^* \end{pmatrix} = \frac{\tilde{V}^{-1} \tilde{E}}{A}, \quad (\text{C4})$$

where w^* and x^* are, respectively, the optimal weights on the N managers and the K index funds and $\tilde{E}$ and $\tilde{V}$ are, respectively, the mean and variance–covariance matrix of the predictive distribution of all the $K + N$ assets.

Next we show that if an asset has a positive posterior expectation of α (i.e., $\tilde{\alpha} > 0$), then this asset is held in positive quantities in the mean–variance tangency portfolio defined by equation (C4). First, consider a portfolio that is

- long a manager and
- short the index funds with weights equal to $\tilde{\beta} \equiv E[\beta|r, F] = \hat{\beta} + (\tilde{\alpha} - \hat{\alpha})y/m$.

We call this asset the “alpha portfolio” and denote its return by r_{T+1}^α . Moreover, we replace the managers with their corresponding alpha portfolios and consider the new problem of determining the weights on the N alpha portfolios and the K index funds. Because any alpha portfolio is a linear combination of a manager and the index funds, it follows that in this new problem, the weight on alpha portfolio i equals the weight on the corresponding manager i defined in equation (C4). Thus if we can show that the alpha portfolio has a positive weight it follows that the manager would have a positive weight as well.

For simplicity, we keep the same notation as above; $\tilde{E}$ and $\tilde{V}$ are the predictive variance–covariance matrix and the predictive expectation of the N alpha portfolios and K index funds. If we can show that $\tilde{E}$ is a vector of length N with elements $\tilde{\alpha}_i$, that $\tilde{V}$ is block diagonal with elements $\tilde{V}_{11}$ ($N \times N$) and $\tilde{V}_{22}$ ($K \times K$), that is,

$$\tilde{V} = \begin{pmatrix} \tilde{V}_{11} & 0 \\ 0 & \tilde{V}_{22} \end{pmatrix}, \quad (\text{C5})$$

and that $\tilde{V}_{11}$ is diagonal, then by equation (C4) it follows that if and only if $\tilde{\alpha}_i > 0$, then there is positive investment in alpha portfolio i , and thus in manager i .

Conditional on the parameters and past data, next period's total return on an alpha portfolio, r_{T+1}^α , is given by

$$r_{T+1}^\alpha \equiv r_{T+1} - F_{T+1}\tilde{\beta} = \alpha + F_{T+1}(\beta - \tilde{\beta}) + \varepsilon_{T+1}. \quad (C6)$$

Using equations (B12) and (B13), it follows that the predictive expectation of an alpha portfolio equals $\tilde{\alpha}$:

$$\begin{aligned} E[r_{T+1}^\alpha | r, F] &= E[r_{T+1} | r, F] - E[F_{T+1} | r, F]\tilde{\beta} \\ &= \tilde{\alpha}. \end{aligned} \quad (C7)$$

Moreover, using equation (B16), it follows that the predictive covariance between the returns to index funds and to an alpha portfolio is zero:

$$\begin{aligned} Cov[r_{T+1}^\alpha, F_{T+1} | r, F] &= Cov[r_{T+1} - F_{T+1}\tilde{\beta}, F_{T+1} | r, F] \\ &= Cov[r_{T+1}, F_{T+1} | r, F] - Var[F_{T+1} | r, F]\tilde{\beta} \\ &= 0. \end{aligned} \quad (C8)$$

Therefore, $\tilde{V}$ is block diagonal as in equation (C5).

Finally, assumptions 1 and 2 (from Section II.E) imply that the posterior distributions for each manager are independent. Using the covariance decomposition for conditional distributions (see equation (B14)), it follows that the covariance between $r_{i,T+1}^\alpha$ and $r_{j,T+1}^\alpha$ equals zero for any $i \neq j$:

$$\begin{aligned} Cov[r_{i,T+1}^\alpha, r_{j,T+1}^\alpha | r, F] &= Cov[E(r_{i,T+1}^\alpha | F_{T+1}, r, F), E(r_{j,T+1}^\alpha | F_{T+1}, r, F) | r, F] \\ &\quad + E[Cov(r_{i,T+1}^\alpha, r_{j,T+1}^\alpha | F_{T+1}, r, F) | r, F] \\ &= Cov[\tilde{\alpha}_i, \tilde{\alpha}_j | r, F] + 0 = 0. \end{aligned} \quad (C9)$$

Therefore $\tilde{V}_{11}$ is diagonal. The expression in Section II.E, equation (37) is obtained by substituting $\Omega = \tilde{V}_{11}$.

We have shown that the weight on alpha portfolio i equals $w_i^* = (1/A)(Var(r_{i,T+1}^\alpha | r, F))^{-1}\tilde{\alpha}_i$. Because the weight on the alpha portfolio must equal the weight on the corresponding manager, it follows that there is positive investment in the manager i if and only if $\tilde{\alpha}_i > 0$.

REFERENCES

- Baks, Klaas P., Andrew Metrick, and Jessica Wachter, 1999, Bayesian performance evaluation, NBER Working paper 7069.
- Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225–264.
- 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., William N. Goetzmann, Roger Ibbotson, and Stephen Ross, 1992, Survivorship bias in performance studies, *Review of Financial Studies* 5, 553–580.
- Brown, Stephen J., William N. Goetzmann, and Stephen A. Ross, 1995, Survival, *Journal of Finance* 50, 853–873.
- Carhart, Mark M., 1995, Survivorship Bias and Mutual Fund Performance, Ph.D. thesis. Graduate School of Business, University of Chicago.
- Carhart, Mark M., 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.
- 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.
- CRSP, 1999, *CRSP survivor bias free US mutual fund database file guide*. Center for Research in Security Prices, Graduate School of Business, University of Chicago.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic based benchmarks, *Journal of Finance* 52, 1035–1058.
- 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.
- Frost, Peter A., and James 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, London).
- Goetzmann, William N., and Roger G. Ibbotson, 1994, Do winners repeat? Patterns in mutual fund performance, *Journal of Portfolio Management* 20, 9–18.
- Goetzmann, William N., and Philippe Jorion, 1997, Re-emerging markets, Working paper, University of California at Irvine.
- 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, University of California at Los Angeles.
- Grinblatt, Mark, and Sheridan Titman, 1992, The persistence of mutual fund performance, *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, *Journal of Finance* 23, 389–416.
- Jobson, J.D., and Bob Korkie, 1980, Estimation of 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, in *Proceedings of the Business and Economic Statistics Section, The American Statistical Association* 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, Transaction costs and investment style: An inter-exchange analysis of institutional equity trades, *Journal of Financial Economics* 46, 265–292.
- Klein, Roger W., and Vijay S. Bawa, 1976, The effect of estimation risk on optimal portfolio choice, *Journal of Financial Economics* 3, 215–231.
- Lehman, Bruce N., and David Modest, 1987, Mutual fund performance evaluation: A comparison of benchmarks and a benchmark of comparisons, *Journal of Finance* 42, 233–265.
- MacKinlay, A. Craig, 1995, Multifactor models do not explain the CAPM, *Journal of Financial Economics* 38, 3–28.
- Malkiel, Burton, 1995, Returns from investing in mutual funds 1971 to 1991, *Journal of Finance* 50, 549–572.
- 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.
- Pástor, Ľuboš, 2000, Portfolio selection and asset pricing models, *Journal of Finance* 55, 179–223.
- Pástor, Ľuboš, and Robert F. Stambaugh, 1999, Costs of equity capital and model mispricing, *Journal of Finance* 54, 67–121.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2000, Comparing asset-pricing models: An investment perspective, *Journal of Financial Economics*, 56, 335–381.
- Perold, André F., 1988, The implementation shortfall: Paper versus reality, *Journal of Portfolio Management* 14, 4–9.
- Shanken, Jay, 1987, A Bayesian approach to testing portfolio efficiency, *Journal of Financial Economics* 19, 195–215.
- Zheng, Lu, 1999, Is money smart? A study of mutual fund investors' fund selection ability, *Journal of Finance* 54, 901–933.