



Survivorship Bias and Mutual Fund Performance

Edwin J. Elton; Martin J. Gruber; Christopher R. Blake

The Review of Financial Studies, Volume 9, Issue 4 (Winter, 1996), 1097-1120.

Stable URL:

<http://links.jstor.org/sici?sici=0893-9454%28199624%299%3A4%3C1097%3ASBAMFP%3E2.0.CO%3B2-5>

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. 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.

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.

The Review of Financial Studies is published by Oxford University Press. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/oup.html>.

The Review of Financial Studies
©1996 Oxford University Press

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2002 JSTOR

Survivorship Bias and Mutual Fund Performance

Edwin J. Elton
Martin J. Gruber
New York University

Christopher R. Blake
Fordham University

Mutual fund attrition can create problems for a researcher because funds that disappear tend to do so due to poor performance. In this article we estimate the size of the bias by tracking all funds that existed at the end of 1976. When a fund merges we calculate the return, taking into account the merger terms. This allows a precise estimate of survivorship bias. In addition, we examine characteristics of both mutual funds that merge and their partner funds. Estimates of survivorship bias over different horizons and using different models to evaluate performance are provided.

The subject of mutual fund attrition and the effect of survivorship bias on performance has only recently begun to receive attention in the academic literature. Early studies of mutual fund performance were concerned with illustrating new methodologies for measuring performance and were less concerned with biases in the data. Later studies tended to neglect fund attrition: the most commonly used databases do not allow the user to either study it or correct for it.

We thank Robert Korajczyk (the editor), an anonymous referee, and participants at the 1995 European Finance Association meetings (Milan) for helpful comments. We are also grateful to Stephen Brown, Bill Crawford at Micropal (ICDD), Judy Doliner at Ivessco, Betsy Tompkins at Salomon Brothers, Chuck Webster at Lehman Brothers, and Interactive Data Corporation (IDC) for supplying data. Address correspondence to Edwin J. Elton and Martin J. Gruber, Leonard N. Stern School of Business, New York University, 44 West 4th Street, 9th Floor, New York, NY 10012.

Mutual fund attrition can create problems for a researcher because the funds that disappear tend to do so either because their performance is very poor over a period of time or because their total market value is sufficiently small that management judges that it no longer pays to maintain the fund. The latter reason for closing a fund is usually associated with the former reason: poor performance. Thus, to study only funds that survive overstates the measured performance. In the vast majority of cases, a fund that disappears is not dissolved but is merged into another fund, often within the same family of funds (sponsoring organization). The effect and perhaps intent of this is that the sponsoring organization continues to earn fees on the investors' capital while the record of the fund's poor performance is deleted from most hard copy and computerized sources of data. For example, drawing a sample from any of Wiesenberger's publications would present a history of return for mutual funds that exist at the time the issue was prepared but would give no details on funds that ceased to exist during the historic time period studied. Most of the classic studies of mutual fund performance ignore attrition and are subject to survivorship bias.¹

Correction for attrition is important for several reasons. First, samples that do not correct for attrition will overstate the return that mutual funds earn for their investors. Second, ignoring attrition may differentially impact the return reported for mutual funds with different objectives, because funds with different objectives may have different rates of attrition. Finally, some of the other variables studied may also be correlated with attrition and, thus, studying a sample with survivorship bias may introduce spurious correlation between these variables and performance.

The purpose of this study is to examine the impact of survivorship bias. We will examine both the frequency of mutual fund disappearance and the impact of this on investor return. In studying performance we examine raw returns as well as risk-adjusted returns from both a single- and a multi-index model.

This article differs from other studies of mutual fund attrition in that we trace the subsequent performance of all funds that existed at a prior point in time. If a fund disappeared from standard databases,

¹ This may have occurred for any one of several reasons: (1) The methodology of examining funds was new and the authors of these studies were interested at least as much in illustrating the methodology as in studying mutual funds; (2) The author desired to demonstrate new techniques over a time span with no missing observations, therefore making the application of the techniques straightforward; (3) Computerized databases were in their infancy and the work required to get any data was immense. To correct the data for missing funds would have been a Herculean task; even today it is extremely difficult. Refinement in correcting data should have proceeded step by step with refinement in the methodology of evaluating mutual fund performance. Yet only recently have researchers become concerned with correcting data for survivorship bias.

we trace what happened to it by contacting the master trustee or the management group associated with the original fund. Once we know what happened to a fund, we track the return an investor in that fund would earn over time. For example, if a fund merged with a second fund with the same objective, we compute risk-adjusted return by examining return for the original fund prior to the merger, incorporating actual merger terms to compute return in the month of the merger, and computing risk-adjusted return for the combined fund after the merger.² Since we track the performance of an investment in any fund that existed at the start of our sample period, our sample is free of survivorship bias. We calculate performance from this sample and from a sample with survivorship bias. We then compute the effect of survivorship bias on performance. In addition, we compute the amount of bias over different horizons to allow the reader to estimate the size of the bias present in the classic studies of performance.

This article contains eight sections. In Section 1, we present a brief review of how articles on mutual fund performance dealt with survivorship bias. In Section 2 we examine estimates of survivorship bias from other authors. In Section 3 we discuss our sample. In Section 4 we discuss our analysis and estimates of survivorship bias. In Section 5 we discuss the characteristics of funds that disappear by merging into other funds as well as the characteristics of the funds with which they merge. In Section 6 we examine the magnitude of survivorship bias over different time horizons and thus the impact of survivorship bias on the results of previous studies. In Section 7 we examine the effect of survivorship on estimates of the relationship between fund characteristics and performance. The final section contains our conclusions.

1. Review of the Literature

The majority of mutual fund performance studies selected funds that existed continuously over a period of time and had some stated investment policy or policies (for example, classification as a common stock fund) at the beginning or end of that period. Such a sampling scheme produces survivorship bias, since only funds that existed over the full period are selected and those that failed are omitted. This selection procedure was used by Connor and Korajczyk (1991), Grinblatt and Titman (1988, 1992)³, Henriksson (1984), Jensen (1968), Lehmann and Modest (1987), Sharpe (1966), and Treynor (1965).

² We call this approach "follow the money." In using it, we are careful to correct for differences in risk that might occur before and after the merger.

³ The Lehmann and Modest (1987) sample, also used by Connor and Korajczyk (1991), requires survival over most of, but not all of, the full period.

Some studies of common stock mutual funds which estimate or attempt to correct for survivorship bias are Brown and Goetzmann (1994), Carhart (1994), Elton et al. (1993),⁴ Grinblatt and Titman (1989), and Malkiel (1994). Grinblatt and Titman (1989) estimate survivorship bias via a simulation study. Brown and Goetzmann (1994) and Malkiel (1994) track all funds and record their returns up to the year they disappear. Elton et al. (1993) track the yearly returns for all funds that existed at the beginning of a sample period, including the year of merger, and the subsequent performance of the funds they merged into. Four of these latter studies provide estimates of survivorship bias for common stock funds, and they will be discussed in the next section of this article.

2. Prior Estimates of Survivorship Bias

Five studies have provided some estimate of survivorship bias: Blake, Elton, and Gruber (1993) for bond funds and Brown and Goetzmann (1994), Carhart (1994), Grinblatt and Titman (1989), and Malkiel (1994) for common stock funds. Blake, Elton, and Gruber (1993) estimate that survivorship bias raises return by 27 basis points per annum for bond funds. This estimate is obtained by taking the difference in excess risk-adjusted return (α from a multi-index model) between those funds that survive and those that don't survive. We might expect this number to be larger for stock funds, given the higher variance of the underlying securities.

Grinblatt and Titman (1989) use quarterly equity holdings to try to estimate the effect of survivorship bias. For each fund, they simulate quarterly returns by calculating the return as if the fund held the equity shares shown at the beginning of each quarter to the end of that quarter. Annual return is computed from the quarterly returns. They calculate the return on two equally weighted portfolios of the individual funds in their sample: one with survivorship bias and one they state is without bias.⁵ Their estimate is the difference in α between these two portfolios. They produce several estimates of bias ranging between 10 and 30 basis points.

⁴ Ippolito (1989) selects all funds that were classified by Wiesenberger as common stock funds at the end of his sample period. He traces these back to the beginning, including funds that merged into each of the funds he selected. However, at the beginning of his sample period the investment policies of many of the funds are not common stock. Thus, no investor interested only in stock funds would have selected the sample.

⁵ Their samples were affected by an inability to track funds because of name changes. Name changes are highly correlated with mergers and policy changes. Thus it is unclear if the sample is free of survivorship bias.

Brown and Goetzmann (1994) present annual returns for the years 1976 to 1988 for two samples. The first sample is all funds that exist as of 1988 and that didn't merge or disappear in the period 1976 to 1988. The second is all funds that existed in Wiesenberger any year for the period 1976 to 1988. Brown and Goetzmann do not track funds that disappear from Wiesenberger where Wiesenberger does not record what happened to them. Some of these disappearances are mergers and some are name changes that Wiesenberger did not record. This problem is recognized by the authors, and they refer to their estimates as coming from an almost survivorship-bias-free sample. Because of the way Brown and Goetzmann select their sample, funds that are included in their sample could have existed for 1 year, 2 years, or up to 12 years. Thus, unlike earlier researchers, they have not used the dual objective of survival and a minimum history. It is therefore difficult to use their results to understand the size of the bias in other studies of mutual fund performance. Brown and Goetzmann's estimates of the bias involved by not including merged funds vary between 20 and 80 basis points per year, depending on the weighting scheme used. Brown and Goetzmann's estimates are based on differences in annual raw returns.

Malkiel (1994), like Brown and Goetzmann, examines the performance of all funds that exist for any time over a period of years. He too reports unadjusted raw returns. He finds that survivorship bias increases the return on the surviving mutual funds by 150 basis points. Finally, Carhart (1994) measures survivorship as the difference in α using one- to four-factor models on equally weighted portfolios of surviving and nonsurviving funds. He estimates the difference in α between these portfolios at 3% to 5% per year. He does not provide a direct estimate of survivorship bias. The studies by Brown and Goetzmann (1994), Carhart (1994), and Malkiel (1994) have gone further to correct for survivorship bias than any previous studies.

Our study differs from those discussed above in one or more ways. First, unlike most studies, we use risk-adjusted returns in addition to raw returns. Risk adjustment is performed using a three-index model, that has been successfully employed in the past, as well as the standard single-index model. Second, we explicitly track every individual fund that existed at the beginning of our sample period to the end of our sample period. We call this technique "follow the money." Most prior studies that looked at bias had funds that disappeared from the sample, with no subsequent tracking of performance. Either funds disappeared from a data set for unknown reasons or, if they merged, they disappeared without accounting for performance in the final period before the merger or accounting for performance after the merger.

3. Sample

We initially started with the 361 funds categorized as having a "common stock" investment policy in the 1977 edition of Wiesenberger's *Investment Companies*.⁶ This directory lists data for the year 1976. We divided the sample into those funds with \$15 million or more in total net assets (207) and those with under \$15 million in total net assets (154) as of year-end 1976. It is extremely difficult to track returns for funds that are no longer reported in Wiesenberger. For very small funds it is often impossible.⁷ Thus our return calculations will be made using returns on funds that had capital of \$15 million or more at the end of 1976. However, when calculating the effect of survivorship bias we will use data on the incidence of survival for the full sample. Having divided our sample by size, we tracked each fund that initially had total net assets of \$15 million or more to December 1993. In this process (see Table 1) we found that 42 of the 207 funds merged, 146 survived to the end of the sample period, and 19 were restricted funds.⁸

There were two types of restricted funds. The first type, called variable annuities (like CREF), were listed in earlier editions of Wiesenberger's *Investment Companies* but were primarily available through insurance plans. The second type were restricted as to purchaser (e.g., a fund that could only be held by Lutherans). We eliminated variable annuity funds, because their sale was usually tied to an insurance product resulting in higher fees to cover insurance, because they are taxed differently, and because their objectives may well be different from the objectives of other funds in the sample.⁹ We excluded funds that could not be bought by the general public because, once again, they may have special objectives.

Of the remaining 188 funds that make up our sample of funds with initial assets of \$15 million or more, 13 disappeared from Wiesen-

⁶ In 1977, Wiesenberger was considered the standard directory for information on mutual funds.

⁷ Prior to 1985, funds had to have a minimum of 1,000 shareholders to be listed by the NASDAQ. Furthermore, the SEC supplies an annual list to the NASD of mutual funds that must be recertified (after their initial listing) as having a minimum of either 750 shareholders or \$15 million in total net assets in order to be included in the NASDAQ listing supplied to newspapers. Although the NASDAQ now provides a supplemental listing to so-called level-one quote vendors (such as IDC) that includes funds with under \$15 million in net assets, for return data in earlier years the NASDAQ newspaper listing is the basic source of return data for all newspapers and firms, and hence investors, that collect data on mutual funds.

⁸ In the process we recorded all name changes, mergers, investment policy changes (e.g., common stock to bond), and restrictions (including changes in variable annuity status). The source of these data was Wiesenberger's *Investment Companies*, supplemented by information from the funds themselves.

⁹ In 1982, the IRS issued a series of so-called wrap-around annuity rulings, designed to prevent tax avoidance. The results is that, since 1982, a fund cannot be both an open-end mutual funds available through direct purchase and a variable annuity fund medium.

Table 1
Distribution (number) of firms by 1976 year-end total net assets

	\$15 million or greater	Under \$15 million	All ^a
Merged	42 ^b	30	106
Survived	146 ^c	67	216
Disappeared		37	
Restricted	19	20	39
Total	207	154	361

^abased on allocation of disappearing firms.

^b12 of the merged funds also changed policy.

^c15 of the funds that survived to the end of the period changed policy during the period.

Surviving funds are defined as funds that existed from the end of 1976 to the end of 1993; merged funds are those funds that existed at the end of 1976 and subsequently merged into another fund; disappeared funds are funds that were listed in Wiesenberger at the end of 1976 but which disappeared from subsequent wiesenberger listings with no indication of what happened to them; restricted funds are funds that were not available to the general public for investment at the end of 1976.

berger with no indication in Wiesenberger of what happened to them.¹⁰ Each was traced: 12 of these merged and 1 had a name change. Applying the same ratio to the funds with net assets under \$15 million to allocate the disappearing funds as either merged or survived results in 106 mergers and 216 survivors for the full sample.

We explored the relationship between fund investment objective and survivorship. There was not a monotonic relationship between fund objectives (maximum capital gain, growth, income) and survivorship. Nor was there a meaningful difference in survival probabilities among categories. Thus, within common stock funds, survivorship does not seem to be related to investment objective.¹¹

Having determined our sample, we then collected returns for the funds. We calculate return for each fund on a monthly basis. In calculating return, dividends are assumed to be used to purchase additional shares in the fund at the reinvestment price (net asset value) that was available to shareholders of the fund. This is the assumption made by Morningstar and Investment Company Data, Inc. (ICDI) in constructing their databases.

For funds that existed over the entire period, returns were supplied by ICDI. For funds that ceased to exist, returns were calculated from data supplied by Interactive Data Corporation (IDC), supplemented by information from the fund management companies themselves.

¹⁰ A substantial number of others "disappear" in that they are not listed in the next issue of Wiesenberger. However, they reappear in subsequent issues of Wiesenberger or have a name change that can be traced through other data recorded in Wiesenberger.

¹¹ In general, survival seems to be a function of risk. Commodity funds have much greater failure rates than common stock funds, and bond funds have lower failure rates than either. Studies that included balanced funds may well have survival related to objective.

Merger terms (e.g., merge ratios) were obtained from the fund management companies themselves.

We checked our data by comparing our returns with those calculated from data reported by Wiesenberger on an annual basis and by comparing monthly returns derived from IDC data with monthly returns from ICDI for time periods over which these data were simultaneously available. Although the pattern of returns between Wiesenberger and ICDI was usually similar across funds, differences of several percent in annual returns were not uncommon. The return differences come about because, in calculating returns with Wiesenberger annual data, dividends are assumed to occur at the end of the year, and ICDI assumes reinvestment in shares of the fund at the time the dividends were paid. Differences between returns calculated from IDC data and ICDI returns rarely occurred. In each case where we found a difference, we went to original sources to resolve the differences. This process gave us great confidence in the accuracy of ICDI return data.¹²

4. Analysis

As explained above, to calculate return we start with all funds that were listed in Wiesenberger as of 1977 and that had \$15 million or more in total net assets. Each fund was tracked to the end of 1993, recording all name changes, policy changes, and mergers. We could track every fund in our sample, and none dissolved. We measure performance using excess return (α) as described below.

Alpha was calculated in two ways, based on a three-index model and one-index model. Our principal results utilize a three-index model. Alpha is defined by the following equation:

$$R_{it} = \alpha_i + \beta_{iL}R_{Lt} + \beta_{iS}R_{St} + \beta_{iB}R_{Bt} + \epsilon_{it} \quad (1)$$

where R_{it} is the excess return on fund i in month t (the return on the fund minus the 30-day T-bill rate); R_{Lt} is the excess return on the S&P 500 Index in month t ; R_{St} is the excess return on small stocks in month t , measured by the return on an equally weighted average of the smallest two deciles of CRSP NYSE stocks; R_{Bt} is the excess return on a bond index in month t , measured by a par-weighted combination

¹² One must be very careful when calculating returns involving dividends and reinvestment net asset values (NAVs) shown in IDC. While IDC is an excellent source of data, we occasionally found discrepancies in dividend reinvestment dates, dividends, and other dividend-related data. These tended to happen around the time of a merger or share change. We checked all data around the time of these events with the funds and examined NAVs surrounding dividend reinvestment dates to ensure accuracy.

of the Lehman Brothers Aggregate Bond Index and the Blume/Keim High-Yield Bond Index;¹³ β_{ik} is the sensitivity of return on fund i to return on index k ($k = L, S, B$); and ϵ_{it} is the random error in period t .

The three-index model compares the performance of the fund to a passive portfolio of large stocks, small stocks, bonds, and T-bills which in combination have similar risk.¹⁴ To help in comparing our results to other studies, and because it has been used historically to evaluate funds, we also present α 's calculated from a single-index model where the single index is the S&P 500 Index. Finally, because other authors have done so, we present results using raw (unadjusted) returns in addition to α 's.

Table 2 presents our estimates of survivorship bias. Note that in Table 2 and all subsequent tables our monthly results have been annualized. All estimates of bias are the difference in performance (α or return) between funds that survive and the full sample that existed at the beginning of the period (labeled all). What varies in the table is the definition of survival and how the performance on the nonsurviving funds is calculated. Both survival and performance are calculated in two ways, resulting in four estimates of bias for the three-index, one-index, and raw-return models.

Panel A of Table 2 shows results when performance on funds that do not survive is calculated up to and including the month when the fund ceased to survive. The upper part of panel A presents the results when survival is defined as "did not merge" over the period. The lower part of panel A defines a surviving fund as a fund that meets the dual criteria: no merger and maintaining a common stock investment policy throughout the sample period. The reason for this division is that some researchers draw their sample from funds that list common stock as an investment policy at the beginning of their sample period and survive throughout, while others insist on a common stock policy throughout the sample period. We wanted to measure bias for both sampling procedures.¹⁵ In calculating return for funds that "disappear in merger" (DIM), we calculate return through the month of the merger. If the fund merged midmonth, this includes any

¹³ See Elton, Gruber, and Blake (1994) for details on the construction of the bond index.

¹⁴ In Blake, Elton, and Gruber (1993) we showed that estimated β 's closely matched the weights of the asset groups in the portfolio. This type of evaluation was first discussed in Sharpe (1988). For most funds in our sample the bond β is close to 0. However, some funds with common stock as an investment policy still hold some bonds and hence have a positive β with respect to the bond index.

¹⁵ More funds disappear from the sample because they merge than because of policy changes. The only other reason we found for funds disappearing is because they are variable annuities (like CREF). Variable annuities were no longer available for direct purchase after 1982; thus they would disappear from all data sets that encompass this period. This affects three funds in our sample.

Table 2
Alternative estimates of survivorship bias (annualized)

	Surviving	Merged	Policy changed	Policy changed and merged	All ^a	Bias
Panel A: Assuming investment to earlier of merger and/or policy change						
Common stock investment policy at beginning of sample period						
1-Index α	(67.034%) ^a	(32.966%) ^a				
3-Index α	0.6712%	-0.2991%			0.3513%	0.3199%
Raw return	-0.1269%	-2.8779%			-1.0388%	0.9069%
S&P 500	14.6832%	8.9977%			12.8089%	1.8743%
Raw return - S & P 500	14.0373%	9.5373%				
	0.6459%	-0.5396%			0.2551%	0.3908%
Common stock investment policy throughout sample period						
1-Index α	(60.147%) ^a	(23.547%) ^a	(6.887%) ^a	(9.419%) ^a		
3-Index α	0.6355%	0.0180%	1.9178%	1.0412%	0.6081%	0.0274%
Raw return	-0.1352%	-2.9052%	-0.0051%	-1.0942%	-0.8688%	0.7356%
S&P 500	14.7504%	9.2448%	9.0589%	8.4516%	12.4687%	2.2817%
Raw return - S&P 500	14.0373%	9.0764%	6.5405%	7.2516%		
	0.7131%	0.1684%	2.5184%	1.2000%	0.7550%	-0.0419%
Panel B: Assuming reinvestment at earlier of merger and/or policy change						
Common stock investment policy at beginning of sample period						
1-Index α	(67.034%) ^a	(32.966%) ^a				
3-Index α	0.6712%	-1.5566%			-0.8985%	0.7716%
Raw return	-0.1269%	-2.4575%			-0.8985%	0.7716%
S&P 500	14.6832%	12.4684%			13.9531%	0.7301%
Raw return - S&P 500	0.6459%	-1.5689%			-0.0842%	0.7301%
	(60.147%) ^a	(23.547%) ^a	(6.887%) ^a	(9.419%) ^a		
1-Index α	0.6355%	-1.5165%	0.4138%	-1.0441%	-0.0447%	0.6802%
Raw return	-0.1352%	-2.3978%	-0.3549%	-1.8085%	-0.8407%	0.7055%
S&P 500	14.7504%	13.0762%	14.5171%	12.7555%	14.1522%	0.5982%
Raw return - S&P 500	14.0373%	14.0373%	14.0373%	14.0373%		
	0.7131%	-0.9611%	0.4798%	-1.2818%	0.1149%	0.5982%

^aAdjusted to reflect entire population of available common stock funds at end of 1976.

For surviving funds, raw returns and excess returns (α 's) are calculated over the full sample period. In panel A, for funds that merged or changed policy, returns are calculated through the month of merger of policy change. In panel B, a "follow the money" approach is used. After a merger, we assume the investor invests in the "partner" fund that the original fund merged into. After a policy change, we assume the investor invests in the average of the remaining funds in the sample. Excess or raw returns are then calculated as a weighted average of the returns in the pre- and post event periods. The category percentages are adjusted to reflect the entire population of available common stock funds at the end of 1976. The returns in the category "All" are weighted averages by percentage in each category. "Bias" is the difference between the value shown for "Surviving" and the value shown for "All."

dividends plus capital changes where capital changes are calculated using merger terms and end-of-month NAV for the partner fund.

Table 2 shows results for both α and raw return. We will first discuss the results using α , then using raw return. As shown in the top half of panel A, the average α using the three-index model for the survivors is -0.1269% , and the average α for those that merge calculated through the month of the merger is -2.8779% per year. Since 32.97% merge and 67.03% survive, the α for the combined sample is -1.0338% ; the estimate of bias is equal to the α in the surviving sample minus the α on the full sample, or 0.9069% per year.¹⁶

The bottom part of panel A in Table 2 calculates bias when the researchers used the dual criteria that the fund had data in each year and was a common stock fund throughout. To determine bias, the funds were divided into four groups. For those funds that neither merged nor had a policy change, α was calculated over the 17-year sample period. For those funds that merged or had a policy change, or did both, α was calculated up to the policy change or merger, whichever came first. The average α for funds that survive and are common stock throughout is -0.135% , for those that have a policy change and do not merge it is -0.0051% , for those that merge it is -2.9052% , and for those that both merge and have a policy change it is -1.094% . The fractions in each category are, respectively, 0.6015 , 0.0689 , 0.2355 , and 0.0942 . Multiplying the fraction in each group times the α for each group and summing across the four groups, we see that the α for the combined sample is -0.869% . Our estimate of bias is the difference between the α for the full sample and that for the surviving sample (-0.135%), or 0.734% per year.

The disadvantage of this method of determining survivorship bias is that disappearing funds and surviving funds exist over different periods of time.¹⁷ An alternative procedure that does not have this

¹⁶ Measuring the statistical significance of the difference in mean α 's between the groups is a more complex issue when the data for all funds are not over the same time period (as in panel A of Table 2) or when α 's over the entire sample period are calculated as weighted averages of α 's from two or more funds over different time periods (as for the nonsurviving funds in panel B). To test for significance, for each group we calculated the variance-covariance matrix of the residuals, utilizing only the data over the common time frame for each pairwise residual covariance for nonsurviving funds in panel A and utilizing the complete set of "spliced" residuals for nonsurviving funds in panel B, and assumed stationarity of the $(X'X)^{-1}$ matrix of the regressors to obtain an estimate of standard error of mean α . The advantage of this technique is that it utilizes actual correlation patterns, both cross-sectionally and over time, to adjust for both heteroskedasticity and cross correlation. Using the results from this procedure, the difference in mean α between surviving funds and the merged funds is statistically significant for the three-index model, where the t values are 2.69 , 2.57 , 2.14 , and 1.90 , proceeding from top to bottom in Table 2. (For the one-index model, no difference are significant at or below the 10% level.)

¹⁷ While differences in length of life is a problem in itself, it can bias the results even more if there is a time pattern to α . The bias occurs because funds that don't exist for the entire sample period have more of their observations in early years. We examined the return of surviving funds over

disadvantage is to make an explicit reinvestment assumption that investors stay invested over the entire period and follow the fund with their money. For funds that merge into another common stock fund, we assume the investor continues to hold the new fund. The α is computed for the original and partner funds separately and then weighted by the time the investor held each. For example, if the fund merged in 5 years we would calculate the α for the first 5 years for the original fund and the last 12 years for the partner fund, and then take 5/17 of the α on the original fund and 12/17 of the α on the partner fund to get an average α . If the fund changes policy and becomes a noncommon stock fund before a merger, after a merger, or without a merger, and if policy change is being used to determine the sample, we assume that at the time the policy change is recorded the investor sells the fund and purchases a new fund with a return equal to the average return on the remaining common stock funds. The overall α , following parallel methodology to that described above, is again a weighted average of the two α 's before and after the event. All of the results under these reinvestment assumptions are shown in panel B of Table 2. Let's first examine the situation where the researcher ignores policy changes and includes funds that changed policy in the surviving category.

If we apply the same methodology to panel B as we applied to panel A, we see that for the three-index model the difference in α between the survivor sample and the full sample is 0.77% per year. If the researcher forming the sample required that the mutual funds both had data over the full period and were listed as common stock in each year, the difference in α between the researcher's sample and the full sample is 0.71% per year.

All of the results discussed to this point measure α from the three-index model. We believe this is appropriate.¹⁸ However, for comparison with earlier studies we also calculated the effect of survivorship bias for a one-index model. These results are also shown in Table 2. The estimates of α for the one-index model are much higher than

different periods and found a slight tendency for higher α 's in early years. Thus, if there is a bias because of differential performance across time it will result in an overstatement rather than understatement of the performance of funds that don't exist for the entire period.

¹⁸ To examine the robustness of our estimate of bias, we also used a four-index model employing a value versus growth index. Researchers have shown that value-growth and market-to-book are potentially interesting additional indexes. We used a value versus growth index based on growth and value indexes produced by Prudential Bache. This index was extremely highly correlated with an index measuring the difference between high and low market-to-book ratios [similar to the index employed by Fama and French (1992)], and thus it may in part capture market-to-book influences. The estimates of bias using the four-index model were for panel A in Table 2: 1.4789% for merger only and 1.3181% for merger and policy change; and for panel B were 0.8388% and 0.7326%, respectively. These estimates are very close to the estimates using three indexes.

for the three-index model. This was a period where small stocks performed relatively well. We know from Elton et al. (1993) that this will cause funds with significant exposure to small stocks to appear to have superior performance when a one-index model is used. The estimates of bias for panel A in Table 2 are much lower for the one-index model than they are for the three-index model. Examining the regression results used to estimate α 's revealed that β 's (sensitivity) for merged funds with the small-stock index were generally higher than those for nonmerged funds. This causes the single-index model to confound the small-fund effect with α and thus to underestimate the bias for the one-index model. We feel the estimates shown in panel A for the one-index model are inaccurate and should be ignored. When we consider reinvestment, the same influences are present, but they are dampened by the averaging with the data from the partner fund.

While we believe the three-index model is the appropriate way to measure excess returns and bias, we also present the results in terms of raw (unadjusted) returns and the difference between return and the naive strategy of holding the S&P 500 Index.¹⁹

Examining panel A, we see that looking at raw returns would lead to a very high estimate of bias (1.875). The reason is easy to see. The S&P 500 Index performed worse in the months when the merged funds were around than it did in the months when the merged funds were not around. Measuring bias from raw returns confounds the effects of merger performance and market performance. Simply measuring each fund's performance net of the return of the S&P 500 Index eliminates the bulk of this problem.²⁰ Notice that the estimates from these excess returns are close to the estimates arrived at using the one-index model. The reason is that this procedure is the same as assuming a one-index model holds and the β of all funds with the S&P 500 Index is equal to 1. We have already explained why the results using the one-index model are inferior to the results using the three-index model. Forcing the β to be equal to 1 for all funds compounds the error and leads to inferior estimates of bias. It is important to point

¹⁹ In panel A, since our non-surviving funds exist for different numbers of months, each fund was paired with an investment in the S&P 500 Index over the months it existed. The overall return for the S&P 500 Index is a weighted average of these paired returns. The effect of this is to weight the monthly S&P 500 return by the number of funds in our sample in any month. Since panel B represents a "follow the money" approach and the investor is always fully invested, the appropriate value for the S&P 500 is its average return over the full sample period.

²⁰ We use the S&P 500 Index for comparison with, and adjustment of, raw returns. Comparison with an index fund is also interesting and can be done easily. Index funds available to individuals exist with a correlation of .99 with the S&P 500 Index and an annualized α of -0.30%, which is just about equal to the expense ratio for these funds. The comparison with an index fund simply lowers the appropriate index returns by about 30 basis points and leaves the estimate of bias unchanged.

this out, for using raw returns or excess returns over some index is a procedure that has been using extensively in the literature.

When we move to panel B, we see that the estimate of bias remains unchanged whether we measure returns in raw form or as excess return relative to the market. This is one of the advantages of the "follow the money" approach; it is much more robust to alternative metrics for performance measurement. However, while the alternative measures of bias are closer together, for the reasons given above we believe the three-index model gives the best results.

Before leaving this section, we should address a potential problem with our methodology. It is possible that funds systematically change their behavior prior to a merger. This could show up either as a change in β 's (sensitivities) or a change in residual risk. For example, one could hypothesize that a manager observing that a fund is having difficulties may increase risk dramatically by loading up on high-risk securities, reducing diversification, or betting on a type of stock (e.g., small stocks). If a manager were to do any of these, then the variance of the residuals around the three-index (or one-index) model should be higher right before the merger than in earlier proofs. Because we are using a regression over the full sample period (up to the time of the merger), changes in either the regression coefficients or residual risk should be reflected in an increase in the variance of the residuals. To examine this we calculated for all merged funds the variance of the residuals for the 6 months prior to the merger and for the 30 months that preceded this 6-month period. The average number and .0595 and .0591, respectively. The fact that the variances are virtually the same over both periods gives us confidence that our three-index model adequately describes the risk process prior to the merger.²¹

5. Characteristics of Funds that Merge into Partner Funds and Their Partner Funds

In this section we will examine the characteristics of funds that disappear in merger (DIM) and the partner funds into which they merge. We examined size (total net assets), expense ratio, load, investment objective, and performance. All analysis is done for the sample of funds over \$15 million in size. While characteristics of mutual funds that merge into another fund have been previously studied, no one heretofore has simultaneously analyzed both the characteristics of funds that merge into other funds and the partner funds into which they merge.

²¹ When using the one-index model, the average variance for the 6 months prior to the merger is .0671, and the average variance for the 30 months preceding that 6-month period is .0722. These numbers are also fairly close and indicate, if anything, a reduction in risk prior to the merger.

Funds that merge into partner funds are smaller than the average of all funds and smaller than the funds into which they merge. The average size of DIM funds in our sample was about \$52 million, while the average size of all funds weighted by the year of merger was about \$305 million.²² Twenty-eight of the 42 funds were under \$50 million, with 15 under \$25 million.

Both the average for the entire sample and for the DIM funds was strongly affected by a few large funds. To control for this we also computed their average rank, where funds were ranked from small to large. The average rank for the DIM funds was the 25% percentile.

The partner funds also tended to be smaller than average. The average size of the partner funds was \$155 million, about one-half the size of all funds. On average, the partner funds ranked in the 43% percentile of all funds (43% were smaller).

Before we examine expenses we must consider the interrelationship between size and expenses per dollar invested. Size and expenses are highly related. To examine this we ranked funds in our sample by size decile and computed the average expense ratio in each decile. To examine if our sample was in any way idiosyncratic, we also did the same analysis for the funds listed in Morningstar that had the same objectives as funds in our sample. Our sample contained a maximum of 207 funds with over \$15 million in assets at the end of 1976, and Morningstar included 773 funds for the year-end 1992 that we analyzed. Table 3 shows the results for the beginning of our sample period and for the 1992 Morningstar data. Examining Table 3 shows almost a monotonic decrease in expense ratios as size increased. This relationship is significant at the 1% level using a rank correlation test.

Thus, in examining difference in expenses, we need to hold the effect of size constant. To hold size constant we matched each DIM fund as well as each partner fund with the two funds closest in size: one larger and one smaller. The results are shown in Table 4. DIM funds have higher expenses than the average fund (1.17% compared to 1.03%). However, the expenses for DIM funds are lower than the matched pairs of similar size, although not significantly so. Thus the expenses of DIM funds are no different than funds of the same size that don't merge. A different pattern emerges for partner funds. Partner funds have significantly *higher* expenses than the matched funds of the same size, 1.24% compared to 0.97%. Furthermore, the partner funds are *higher* in expense than the DIM funds. We also examined what happened to the expenses of partner funds after the merger.

²² The average for all funds in our sample was computed by taking the average asset size in each year and multiplying by the fraction of funds that merged in that year. This averaging was done to correct for an increase in size of the average fund over time.

Table 3
Total year-end net asset size and annual expenses by decile (each decile constructed by ranking on asset size)

Decile	1976		1992	
	Assets	Expenses	Assets	Expenses
1	\$ 16.85	1.25%	\$ 5.43	2.38%
2	\$ 21.50	1.12%	\$ 15.51	1.51%
3	\$ 29.93	1.13%	\$ 29.27	1.55%
4	\$ 38.01	1.10%	\$ 45.35	1.55%
5	\$ 49.96	0.94%	\$ 76.17	1.34%
6	\$ 71.23	0.93%	\$ 120.91	1.28%
7	\$113.09	0.79%	\$ 182.47	1.33%
8	\$189.39	0.67%	\$ 312.55	1.16%
9	\$328.47	0.66%	\$ 621.88	1.05%
10	\$831.70	0.54%	\$2,469.14	0.90%

The 1976 data are from the 207 funds with more than \$15 million in total net assets as shown in the 1977 edition of Wiesenberger and described in Table 1. The 1992 sample is all common stock funds listed in Morningstar at the end of 1992; that sample size is 773 funds. Total net assets are in millions of dollars; expenses are expressed as a percentage of net assets.

Table 4
Average annual expense ratios of merged and partner funds

	Sample	Matched pair	Difference	Std. dev. of difference	t value of difference
Merged	1.170%	1.210%	-0.040%	0.065%	0.70
Partner	1.240%	0.970%	0.270%	0.075%	4.45
Average	1.025%				

The column labeled "Matched pair" contains the average expense ratio of the matched pair for each of the corresponding sample funds; a sample fund's matched pair consists of the funds with the next smaller and next larger asset sizes surrounding the asset size of the sample fund.

The columns labeled "Difference," "Std. dev. of difference," and "t value of difference" contain values computed using pairwise differences of the corresponding sample funds' expense ratios and the average expense ratios of the sample funds' matched pairs.

The row labeled "Average" contains the simple average of the expense ratios for all funds in the sample.

We could find no discernible pattern. The higher expenses of partner funds were not due to expenses incurred with the merger, nor did the merger lead to lower expenses in subsequent years.²³

While we have discussed expense ratios in general, we have not as yet examined one element of expense, the existence of loads. It is possible that funds merge into partners which, while having a higher annual expense, do not have a load. Table 5 shows that this is not the case. There is no tendency for funds that merge to select partner funds without loads. There is some tendency for funds to merge into partners with the same load policy as the merged fund.

²³ For a few funds the merger year seemed to be a local maximum. This pattern was not strong enough to show up for the average of all DIM funds.

Table 5
Change in load policy between merged funds and partner funds

	Actual occurrence	Expected occurrence
No change	29	20.3
Load to no load	4	9.1
No load to load	4	7.6

Expected occurrence is computed by assuming that the probability of a fund ending up in any load category is equal to the fraction of funds in that category in the year of the merger.

Actual occurrences do not total 42 (the number of merged funds in our sample) due to unavailability of load policy data for some partner funds.

Table 6
Change in investment objective between merged funds and partner funds

	Actual occurrence	Expected occurrence
No change	20	13.5
Riskier	12	11.1
Less risky	8	19.2

Investment objectives were ranked from most risky to least risky as follows: maximum capital gains, growth, growth and income, and income.

Expected occurrence was computed by assuming that the probability of any fund merging into a partner fund with a particular objective was equal to the percentage of funds that followed that particular objective in the merger year.

Actual occurrences do not total 42 (the number of merged funds in our sample) due to unavailability of investment objective data for some partner funds.

The next issue to be examined is the investment objectives of the DIM fund and the partner fund. Do funds tend to merge into funds with less risky objective to help out investors who have recently had a bad experience (bad return), or do funds pay no attention to risk in picking partners? Table 6 presents some evidence on this. The column labeled "actual occurrence" shows the number of funds that held investment objective constant, became riskier, or became less risky. The entries under "expected occurrence" show the number that would have been in each category if the decision as to merger partner were made at random. This number was calculated by assuming that the chance of picking a partner fund was equal to the percentage of funds that followed each investment objective. The table shows that there is no tendency for funds to pick merger partners that are less risky. One-half of the sample pick partners that have the same risk. Of the remaining group, more than half merge into funds that are more risky. When comparing these results to the expected occurrence, we see that the number of funds that became riskier is about what we would expect, while many fewer funds become less risky than expected, and a larger than expected number merge into funds with the same policy. There is no evidence that funds become less risky; in fact, on average they tend to become more risky after they merge.

The last aspect of DIM funds and partner funds we examine is the excess return (α) from the three-index model. We examined the α for the DIM fund for a 3-year period before the merger and the α for the partner fund both 3 years before and 3 years after the merger. The DIM funds had an average annualized α of -3.6% before merger, while for this same period of time the partner funds had an average annualized α of +1.7%.²⁴ It appears that funds that are doing poorly are merged into partner funds that have had very good performance relative to the population of funds from which partners can be selected. The next obvious question to ask is whether this excellent performance continues after the merger. The answer is that in the 3 years following the merger the α on the partner fund is actually slightly worse than the α for the average fund in our sample.²⁵

6. Survivorship Bias and Length of Sample Period

Researchers who have examined mutual fund performance have used sample periods ranging from 10 to 20 years. Clearly, the longer the sample period the greater the survivorship bias. In order to interpret the results contained in the literature, we need to estimate the survivorship bias for different holding periods. Merger activity is clearly a function of economic conditions. Thus to look at survivorship over the first 10 years of our sample is to make the estimate both a function of the shorter time interval and a function of the particular economic conditions that exist over that period. An alternative is to compute an average merger rate and use that average rate as the merger rate in each year. The argument against using the average rate is that each year in our sample we know that funds have existed one more year, and this could affect merger probability. Also, it is possible that mergers are related to market movements and tend to cluster in one year. To test this, we regressed merger rates as a function of time. We found no relationship (a negative adjusted R^2). Figure 1 shows the plot of the actual proportion that survived and a plot under the assumption of a constant survival rate when survival is defined as not merging.²⁶

²⁴ The α for merged funds is slightly different from that reported in Table 2. The reason is that funds that merged did slightly worse in a 3-year span prior to the merger than they did in earlier periods.

²⁵ We got similar results when performance was computed over a 1-year period (before and after) the merger.

²⁶ Brown and Goetzmann (1994) explore the relationship between funds disappearing from Wiesenberger and fund characteristics. They find that fund disappearance from Wiesenberger is a function of size, expenses, age, and relative performance. We do not find the same relationship with age. Malkiel (1994) has a sample where new funds enter as they are formed. He was concerned that the different life of his funds might affect his results. The survivorship pattern in our data indicates that this is not likely to be a problem in his sample.




Figure 1
Survival as a function of time: common stock at beginning




Figure 2
Survival as a function of time: common stock throughout

Figure 2 shows the same data when survival is defined as not merging and with no change in investment policy. The percentage of funds that merge in each year is 2.3%, while when we consider either a merge or a policy change, the percentage becomes 2.9%. As can be seen from the diagrams, the assumption of a constant survival rate closely matches the data for the actual survival rate. These same figures can be used to analyze another issue. Potentially, survivorship can be a function of market conditions. The figures suggest survivorship is not a function of market conditions. A regression of survival rates on the S&P 500 Index also showed no relationship.

Table 7
Survivorship bias (α) and length of sample period (three-index model)

Number of years in study	Assuming no reinvestment		Assuming reinvestment	
	Common stock at beginning	Common stock throughout	Common stock at beginning	Common stock throughout
10	0.489%	0.386%	0.396%	0.348%
11	0.545%	0.431%	0.443%	0.392%
12	0.602%	0.478%	0.493%	0.439%
13	0.660%	0.526%	0.544%	0.487%
14	0.719%	0.576%	0.598%	0.538%
15	0.780%	0.627%	0.654%	0.591%
16	0.843%	0.679%	0.712%	0.647%
17	0.907%	0.734%	0.772%	0.705%
18	0.972%	0.789%	0.834%	0.766%
19	1.039%	0.847%	0.899%	0.830%
20	1.108%	0.906%	0.966%	0.896%

Bias is defined as average α for surviving funds minus average α for all funds when a three-index model is employed.

Table 7 for the three-index model and Table 8 for the one-index model show the effect of survivorship bias over horizons from 10 to 20 years using each of the measurement techniques discussed earlier.

In Table 7 we calculate all results using a three-index model. Bias estimates are presented for 10 to 20 years because this encompasses the time frame of all prior studies. In the columns headed "Common Stock at Beginning" we are presenting bias estimates that are relevant for all studies that insisted that a fund exist for the entire period but only looked at investment policy at the beginning of the period. The failure rate used here is the 2.3% per year that we previously estimated from our database. In the two columns headed "Common Stock Throughout" we present estimates of bias to use in evaluating results from studies that insisted on a fund existing for the entire period and that eliminated funds that changed policy at any point in the study.

Table 7 shows us that the amount of bias present in other studies does not vary a great deal whether we measure α up to the month of a merger or policy change or measure α by making specific reinvestment assumptions. In addition, the estimate does not vary much whether researchers examine samples where funds that changed investment policy were included or excluded. For example, consider bias estimates for a 20-year sample. The largest estimate of bias is 1.108% and the smallest is 0.896%. All of the numbers are clearly of economic significance. The largest differences occur between the cases where the researcher rejects funds with investment policy changes and the cases where they are not rejected. The table shows that the bias is smaller when the investor rejects funds (throws

Table 8
Survivorship bias (α) and length of sample period (one-index model)

Number years in study	Assuming no reinvestment		Assuming reinvestment	
	Common stock at beginning	Common stock throughout	Common stock at beginning	Common stock throughout
10	0.173%	0.014%	0.378%	0.336%
11	0.192%	0.016%	0.423%	0.378%
12	0.212%	0.018%	0.470%	0.423%
13	0.233%	0.020%	0.519%	0.470%
14	0.254%	0.022%	0.570%	0.519%
15	0.275%	0.023%	0.624%	0.570%
16	0.297%	0.025%	0.679%	0.624%
17	0.320%	0.027%	0.737%	0.681%
18	0.343%	0.029%	0.796%	0.739%
19	0.367%	0.032%	0.858%	0.801%
20	0.391%	0.034%	0.923%	0.865%

Bias is defined as average α for surviving funds minus average α for all funds when a one-index model is employed.

out funds) that had a change in policy during a period. Why does this occur? The answer is that a large number of the funds that had policy changes subsequently merged. In fact, 58% of the funds that had policy changes subsequently merged. These funds had very poor performance between the time of the policy change and the merger (that is why they subsequently merged). By assuming the investor stops investing in these funds at the time of the policy change rather than at the time of the merger, we increase the performance of the unbiased sample (labeled "All" in Table 2). Therefore the "All" entries are much larger for the case where funds that have policy changes are excluded, and this results in a smaller bias.

Note that while the differences at a point in time across ways of measuring bias are not large, the estimate of bias varies a great deal with the length of the study regardless of the methodology adopted. For example, when ignoring policy change while assuming reinvestment, the estimate of bias is 0.396% per year in a 10-year study and 0.966% per year in a 20-year study.

The results for the one-index model are shown in Table 8. When we assume reinvestment, the bias produced by the one-index model is very similar to the results using the three-index model. A large difference does arise in the case where we assume no reinvestment. This case shows much lower values for bias. Again, the data that leads to this can be clearly seen from Table 2, where the estimates of bias without reinvestment for the one-index model are shown to be much smaller than for the three-index model. The reasons for this were discussed earlier. They account for the results in Table 8 and explain why we do not consider these cases in Table 8 to be appropriate.

Table 9
Size and performance

Size	α (from three-index model)	
	Biased	Unbiased
Smallest decile	-0.342%	-0.718%
Largest decile	-0.394%	-0.308%

The values in this table are the average decile α 's, where the α 's are calculated from our three-index model over our sample period; funds are assigned to deciles by ranking on size (total net assets) at the start of our sample period (year-end 1976).

Table 10
Investment objective and performance

	α (from three-index model)	
	Biased	Unbiased
Max. cap. gain	0.720%	-0.341%
Growth	0.680%	-0.655%
Growth & Income	-0.649%	-1.049%

The values in this table are the average α 's within investment objective categories, where the α 's are calculated from our three-index model over our sample period; funds are classified according to their investment objectives at the start of our sample period (year-end 1976).

7. Survivorship and Fund Characteristics

There is one more aspect of survivorship that deserves discussion. Failure to correct for bias can lead to incorrect inferences about the impact of fund characteristics. Two examples of this are presented in Tables 9 and 10.

When we examine the performance of the largest 10% of our sample and the smallest 10% of our sample ranked by assets under management, we find very different inferences about the impact of size (at the beginning of the sample) on performance. If we look at the biased sample (see Table 9), we find virtually no difference in performance between large funds and small funds. (Actually, small funds appear to perform just slightly better.) These results are consistent with previous findings of the impact of size on fund performance [e.g., Sharpe (1966)]. However, when α 's on the largest and smallest funds are examined for the unbiased sample, we find the smaller funds perform much worse, with a negative α more than twice the negative α of the large funds. The results are clearly consistent with the fact that a larger percentage of small funds relative to large funds fail to survive, and funds that fail to survive have poorer performance than funds that do survive.

As a second example of the effect of survivorship bias on the relationship of fund characteristics and α consider Table 10. Table 10 is a standard table for comparing α across fund investment objective. For

the biased sample, the maximum capital gain and growth funds appear to have positive α 's. This is consistent with the results reported by Grinblatt and Titman (1989) and Connor and Korajczyk (1991). However, when looking at the results accounting for survivorship bias, a different picture emerges. All fund categories have negative α 's. Further, the category "growth," which some prior researchers found had the best performance, has the biggest adjustment in α and no longer has the best performance.

These simple cases illustrate the fact that erroneous conclusions can be reached about the impact of a variable on performance when a sample of mutual funds that contains survivorship bias is examined.

8. Conclusions

Almost all prior mutual fund studies suffer from survivorship bias. Since funds that merge have worse performance than those that don't, those studies contain estimates of performance that overstate true performance. This study provides estimates of survivorship bias that can be used as benchmarks to determine the amount of bias in studies that do not take survivorship into account. This study also examines the characteristics of funds that disappear through merger and the funds into which they merge. Finally, we have presented two examples of the fact that a failure to eliminate survivorship bias can lead a researcher to spurious conclusions about the effect of fund characteristics on return.

References

- Blake, C. R., E. J. Elton, and M. J. Gruber, 1993, "The Performance of Bond Mutual Funds," *Journal of Business*, 66, 371-403.
- Brown, S., and W. Goetzmann, 1994, "Attrition and Mutual Fund Performance," working paper, New York University, forthcoming in *Journal of Finance*.
- Carhart, M. M., 1994, "On Persistence in Mutual Fund Performance," working paper, University of Chicago.
- Connor, G., and R. Korajczyk, 1991, "The Attributes, Behavior and Performance of U.S. Mutual Funds," *Review of Quantitative Finance and Accounting*, 1, 5-26.
- Elton, E. J., and M. J. Gruber, 1995, *Modern Portfolio Theory and Investment Analysis* (5th ed.), John Wiley, New York.
- Elton, E. J., M. J. Gruber, and C. R. Blake, 1994, "Fundamental Economic Variables, Expected Returns, and Bond Fund Performance," working paper, New York University, forthcoming in *Journal of Finance*.
- Elton, E. J., M. J. Gruber, S. Das, and M. Hlavka, 1993, "Efficiency with Costly Information: A Reinterpretation of Evidence for Managed Portfolios," *Review of Financial Studies*, 5, 1-22.

- Fama, E. F., and K. R. French, 1992, "The Cross-Section of Expected Stock Returns," *Journal of Finance*, 47, 427-465.
- Ferson, W., and R. Schadt, 1993, "Measuring Fund Strategy and Performance in Changing Economic Conditions," working paper, University of Washington School of Business Administration.
- Grinblatt, M., and S. Titman, 1988, "The Evaluation of Mutual Fund Performance: An Analysis of Monthly Returns," working paper, John E. Anderson Graduate School of Management, UCLA.
- Grinblatt, M., and S. Titman, 1989, "Mutual Fund Performance: An Analysis of Quarterly Portfolio Holdings," *Journal of Business*, 62, 393-416.
- Grinblatt, M., and S. Titman, 1992, "Performance Persistence in Mutual Funds," *Journal of Finance*, 47, 1977-1984.
- Henriksson, R. D., 1984, "Market Timing and Mutual Fund Performance," *Journal of Business*, 57, 73-96.
- Henriksson, R. D., and R. C. Merton, 1981, "On Market Timing and Investment Performance II: Statistical Procedures for Evaluating Forecast Skills," *Journal of Business*, 54, 513-553.
- Ippolito, R. A., 1989, "Efficiency with Costly Information: A Study of Mutual Fund Performance," *Quarterly Journal of Economics*, 104, 1-23.
- Jensen, M. C., 1968, "The Performance of Mutual Funds in the Period 1945-1964," *Journal of Finance*, 23, 389-416.
- Lehmann, B. N., and D. Modest, 1987, "Mutual Fund Performance Evaluation: A Comparison of Benchmarks and Benchmark Comparisons," *Journal of Finance*, 42, 233-265.
- Malkiel, B., 1994, "Returns from Investing in Equity Funds," working paper, Princeton University; forthcoming in *Journal of Finance*.
- Sharpe, W. F., 1966, "Mutual Fund Performance," *Journal of Business*, 39, 119-138.
- Sharpe, W. F., 1988, "Determining a Fund's Effective Asset Mix," *Investment Management Review*, June, 5-15.
- Treynor, J., 1965, "How to Rate Management of Investment Funds," *Harvard Business Review*, 44, 63-75.
- Wiesenberger Financial Service, 1977-1992, *Investment Companies*, Warren, Gorham & Lamont, Inc., Boston.