



American Finance Association

Are Investors Rational? Choices among Index Funds

Author(s): Edwin J. Elton, Martin J. Gruber, Jeffrey A. Busse

Source: *The Journal of Finance*, Vol. 59, No. 1 (Feb., 2004), pp. 261-288

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/3694896>

Accessed: 23/11/2009 13:09

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>

Are Investors Rational? Choices among Index Funds

EDWIN J. ELTON, MARTIN J. GRUBER, and JEFFREY A. BUSSE*

ABSTRACT

S&P 500 index funds represent one of the simplest vehicles for examining rational behavior. They hold virtually the same securities, yet their returns differ by more than 2 percent per year. Although the relative returns of alternative S&P 500 funds are easily predictable, the relationship between cash flows and performance is weaker than rational behavior would lead us to expect. We show that selecting funds based on low expenses or high past returns outperforms the portfolio of index funds selected by investors. Our results exemplify the fact that, in a market where arbitrage is not possible, dominated products can prosper.

THE CLASSIC PARADIGM OF FINANCIAL THEORY assumes that investors operating in frictionless markets make rational decisions. Under this paradigm, rational investors set prices, and their actions lead to the elimination of dominated financial investments. Since S&P 500 index funds are a commodity that differ from each other principally in price, they represent one of the simplest vehicles for examining investor rationality under the assumption of frictionless markets. Although all of these funds hold the same securities in virtually identical percentages, they have substantial differences in fees and returns that investors should find economically significant. For example, the difference in annual return between the best performing and worst performing S&P 500 index fund is 2.09 percent per year.

In this paper, we show that the differences in future returns of S&P 500 index funds, unlike those for actively managed funds, are predictable to a very high degree of accuracy. Other aspects of performance that might concern investors, such as risk and tax efficiency, are also easily predictable. We show that although cash flows are related to performance, the relationship, while statistically significant, is much weaker than we would expect based on rational behavior in frictionless markets, the classic paradigm of financial economics. We investigate the ability of other influences to explain cash flows, and we find only modest success. Finally, we examine the actual return earned by investor cash flows into S&P 500 index funds, and we find that naïve rules for selecting index funds significantly outperform the returns that investors actually earn. We would expect the investors who buy index funds to be among the

*Elton and Gruber are at the Stern School of Business, New York University. Busse is at the Goizueta Business School, Emory University. We thank Rick Green (the editor) and two anonymous referees for helpful comments and LiTing Cheng for research assistance.

most knowledgeable of all investors and to make their allocation among index funds to maximize their economic payoff.¹ As we show, this is not the case. Possible explanations include substantial information-gathering costs, barriers to the flow of capital, certain funds providing investors with major benefits beyond returns, or that we as a profession have overestimated the rationality of investors.²

Other research questions the classic paradigm of financial economics. Some studies show that securities that are close substitutes can sometimes sell at different prices (Rosenthal and Young (1990) and Froot and Dabora (1999)). Other studies show that the behavior of an individual investor might not fit the classical paradigm. Some investors fail to exercise in-the-money options and exploit arbitrage opportunities (Rietz (1998) and Longstaff, Santa-Clara, and Schwartz (2001)). Individual investors trade too much, maintain undiversified portfolios, hold losing positions too long, require a risk premium for idiosyncratic risk, and overinvest in their own companies' stock (Blume and Friend (1975), Green and Rydqvist (1997), Ferris, Haugen, and Makhija (1998), Odean (1998), Barber and Odean (2000), Cohen, Gompers, and Vuolteenaho (2001), Grinblatt and Keloharju (2001), Hirshleifer et al. (2001), and Huberman (2001)). This paper continues this line of research using a widely held investment vehicle in an extremely liquid market.

The paper is organized as follows. Section I describes the data. Section II examines the predictability of performance, risk, and tax efficiency. Section III estimates the relationship between cash flows and performance as well as other fund characteristics. Section IV compares the actual return investors earn with the return earned by following simple strategies. Section V concludes.

I. Data

In this section we describe the sample we use to test hypotheses about the predictability of payoffs and future cash flows.

We take the initial list of index fund names from the January 1997 edition of Morningstar's *Principia Plus*. After we eliminate any enhanced return index funds, our sample consists of 52 open-end S&P 500 index funds. We track the 52 funds through name changes (31 funds incur 36 name changes) and mergers (three funds) in subsequent editions of Morningstar *Principia Plus*.³ Since our

¹ About one-seventh of the flows into mutual funds comes from 401K and 403B plans. Retirement plans typically constrain cash flows to go into funds offered by the plan. However, since the plan sponsor has a fiduciary responsibility to select the best performing funds for the plan, these flows should be allocated as rationally as direct investments.

² Christoffersen and Musto (2002) study the response of investors to fees and performance for money market funds. Although their analysis differs from ours, it represents the only other attempt to examine investor actions across a sample of funds where performance is determined principally by fees charged to the investor.

³ All mergers occur between index funds. We leave out the disappearing funds in the year of the merger, which could introduce some survivorship bias. To examine this possibility, we use the follow-the-money approach of Elton, Gruber, and Blake (1996a) that eliminates survivorship bias. The results are unaffected.

sample includes all the S&P 500 index funds that exist at the beginning of the sample period, the sample does not suffer from survivorship bias of the sort identified in Brown et al. (1992) and Elton, Gruber, and Blake (1996a), which occurs when a sample includes only funds that exist at the end of the sample period.

We use the January 1997 through 2002 editions of *Principia Plus* and *Principia Pro Plus* for fund data from January 1996 through December 2001. The data include monthly returns and annual data consisting of the net asset value, expense ratio, 12b-1 fee, actual management fee, load, capital gains in dollars, dividend income in dollars, and whether the fund is only available to institutional investors. We use the Center for Research in Security Prices (CRSP) Survivorship-free Mutual Funds Database for monthly total net assets. To improve accuracy, we compare monthly returns and annual expense ratios reported by CRSP and Morningstar, and we use fund prospectuses and data provided by the fund itself to reconcile differences. For each fund we also gather fund family data from Morningstar, including number of funds, number of different fund objective categories, total size, and whether the family includes a fund with a five-star Morningstar rating.

We use the S&P 500 return index including dividends as our benchmark. To compute excess returns on the funds and on the S&P 500, we use the CRSP monthly *T30RET* 30-day T-bill return.

II. Characteristics of Index Funds and Their Predictability

What should an investor who is purchasing an index fund care about? First and foremost, the investor should care about the fund return relative to the index. An unsophisticated investor might simply want a fund whose return in the next period is high compared to the index or to other index funds. A more sophisticated investor might wish to adjust any return above or below the index for differences in risk from the index risk. Second, the investor might be concerned explicitly with the fund's risk. How closely does the fund track the index, and is the fund's risk stationary over time? Finally, the investor might be concerned with the tax efficiency of the fund.

We show that all of the important characteristics of an index fund are highly predictable. We examine the predictability of return, risk, and tax efficiency each in turn. We use two different time horizons to examine predictability—1 year and 3 years—to see if our conclusions are robust across different potential holding periods for investors.

A. Predictability of Average Index Fund Return

Investors should first examine whether they can predict the payoff they will receive from holding a fund. The investor who holds an index fund intended to replicate the S&P 500 index has a clear benchmark: the total return on the

S&P 500 index.⁴ How will an investor measure the payoff? Naive investors will simply observe how much return they receive above or below the S&P 500 index (i.e., the differential return). A more sophisticated investor will define the payoff as risk-adjusted return.⁵ This investor will adjust the return for the fund's beta. The risk-adjusted return is alpha, defined in

$$R_i - R_f = \alpha_i + \beta_i(R_m - R_f) + e_i, \quad (1)$$

where R_i is the monthly return on index fund i , R_f is the return on a 30-day Treasury bill, R_m is the monthly return on the S&P 500 index, β_i is the sensitivity of fund i to the return on the S&P 500 index, and α_i is the risk-adjusted return on fund i .

Investors who care about future differential return are likely to examine past differential return to select a fund. If they do not believe that management skill varies across funds, they might also forecast differential return by examining expenses or by examining both past differential return and expenses together. A more sophisticated investor will adjust return for systematic risk and will likely examine past values of alpha to predict future alpha. This investor might also use expenses, or alpha and expenses together, to predict future alpha.

In this section of the paper we examine the relationship between differential return or alpha and a set of variables that might be useful in predicting them. We do so in three stages. First, we examine alpha and differential return to see if their dispersion across funds is sufficiently large to make prediction a worthwhile exercise. Then, we measure the association between both of these measures and past data to see if a statistically significant relationship exists. Finally, we test directly whether ranking funds on the basis of a set of past variables leads to superior or inferior performance in the future. We look at the latter two sets of tests over both 1-year and 3-year holding periods.

A.1. Size and Dispersion of Return Variables

Table I shows that the average differential return across S&P 500 index funds is -0.485 percent per year with a range of -1.857 percent to 0.232 percent. Alpha averages -0.410 percent per year with a range of -1.530 percent to 0.228 percent. Clearly, the index funds on average underperform the index by a significant amount, and economically significant differences exist in the performance of alternative funds. Expense ratios average 0.444 percent per year within a range of 0.060 percent to 1.350 percent. The size of expense ratios suggests that expenses may account for a large portion of the differential performance between index funds.

⁴ Unlike other work on performance measurement, we need not be concerned with identifying an appropriate benchmark or using a multi-index model.

⁵ If deviations of beta from one are random, the investor should be concerned with differential return rather than alpha.

Table I
Summary Statistics

The table shows summary statistics of the S&P 500 index fund sample. Differential return is the return on the fund minus the return on the S&P 500 index. Alpha is the intercept in the regression of the excess fund return (fund return minus risk-free rate) against the excess return on the S&P 500 index. Returns, alpha, expenses, and management fees are annualized. The sample includes 52 funds. The sample period is from January 1996 to December 2001.

	Performance			
	Differential Return	Alpha	Expenses	Management Fee
Mean (%)	-0.485	-0.410	0.444	0.177
Minimum	-1.857	-1.530	0.060	0.000
Maximum	0.232	0.228	1.350	0.700
25 th percentile	-0.683	-0.580	0.250	0.070
75 th percentile	-0.224	-0.211	0.590	0.250

	Managerial Skill		Risk	
	Alpha + Expenses	Diff Return + Expenses	Beta	R ²
Mean	0.034%	-0.041%	0.998	0.9999
Minimum	-0.415	-0.846	0.979	0.9991
Maximum	0.539	0.632	1.005	1.0000
25 th percentile	-0.077	-0.184	0.996	0.9998
75 th percentile	0.119	0.104	1.000	1.0000

A.2. Association of Return with Past Variables

We next examine the association between return measures and lagged values of several variables that might predict return.⁶ We examine the association over a 3-year horizon and then for a 1-year horizon. For a 3-year horizon, we estimate the relationship in the cross section using 3 years to formulate the predictors and 3 years to calculate return measures.

Panel A of Table II indicates that differential return has a very high R^2 with past expenses (0.768). The relationship is significant at the 1 percent level. Furthermore, expenses on average lower differential return by the amount of the expenses, since differential return goes down by 0.999 percent for every 1 percent increase in expenses. If higher expenses motivate management to greater effort that leads to better future performance, we would expect a coefficient of much less than 1. Past expenses affect future performance very strongly, because past expenses are almost perfect predictors of future expenses. We demonstrate the stability of expenses by noting that the coefficient of determination between past and future expenses is 0.931 with a slope of 0.997.

When we examine the association between past differential return and future differential return, we get even stronger results. Panel A of Table II shows

⁶ Persistence in active mutual fund performance has been studied by Elton, Gruber, and Blake (1996b), Gruber (1996), Carhart (1997), and Bollen and Busse (2002), among others.

Table II
Association over 3-Year Periods

The table shows statistics from regressions of 3-year measures of the indicated regressand variables at time $t + 1$ on 3-year measures of the indicated regressor variables at time t . Returns are annualized. The sample includes 52 funds. The sample period is January 1996 through December 2001.

Variables	Intercept	$t(\text{int})$	Slope 1	$t(\text{slope } 1)$	Slope 2	$t(\text{slope } 2)$	R^2
Panel A. Performance							
Differential return vs. differential return	-0.019%	-0.548	0.777	14.773			0.845
Differential return vs. expenses	0.015%	0.323	-0.999	-11.514			0.768
Differential return vs. expenses, differential return + expenses	-0.018%	-0.453	-0.782	-8.983	0.764	4.401	0.845
Alpha vs. alpha	-0.030%	-0.985	1.048	16.849			0.877
Alpha vs. expenses	0.010%	0.219	-1.011	-11.987			0.778
Alpha vs. expenses, Alpha + expenses	-0.036%	-1.002	-1.043	-16.121	1.108	5.528	0.877
Panel B. Managerial Skill							
Differential return + expenses vs. differential return + expenses	0.025%	1.318	0.265	2.328			0.119
Alpha + expenses vs. Alpha + expenses	-0.044%	-2.798	0.665	4.558			0.342
Panel C. Risk							
R^2 vs. R^2	0.498	3.840	0.502	3.870			0.272
Abs ($\beta - 1$) vs. Abs ($\beta - 1$)	0.001	2.432	0.487	2.727			0.157
Panel D. Tax Efficiency							
(Div yld + exp)/fund avg vs. (div yld + exp)/fund avg	0.471	5.072	0.520	5.667			0.445
Capital gain yld/fund avg vs. Capital gain yld/fund avg	-0.068	-0.148	1.070	2.752			0.159

that the coefficient of determination increases to 0.845 when we substitute past differential return for past expenses. Investors interested in differential return can apparently choose an index fund simply by looking at the past expense ratio, but they can do even better by looking at the past differential return.⁷

Panel A of Table II also shows that when we examine alpha as the dependent variable, the results closely parallel those for differential return. The only difference is an association between future alpha and past alpha that is even stronger than the association between future differential return and its past value.

Although we have studied the association over a 3-year period, the question remains whether similar results hold over a shorter 1-year time horizon. Since Gruber (1996) and Zheng (1999) show that investors react to short-term performance, we wish to establish whether the same variables that hold over longer horizons also hold in the short term. Panel A of Table III shows that the results for the 1-year horizon closely parallel the 3-year results.⁸ As we move from 3 years to 1 year, the R^2 across all relationships decreases slightly. Interestingly, lagged expenses show a higher association with future differential return and alpha than lagged values of either of these variables, unlike the results for the 3-year horizon.⁹ The reason for this is that the short-term return contains more random noise as a predictor of the future than long-term return, whereas expenses are much more constant over time. The R^2 of the regression of expenses in one year on expenses in the following year is 0.97.

A.3. Prediction of Return

Since we will be examining the impact of past differential return and alpha on subsequent cash flow, it is worthwhile to examine the subsequent performance of funds chosen on the basis of past predictive variables. Tables IV and V show this for 3- and 1-year periods. For each holding period, we rank funds according to some predictive variable and then divide them into deciles.

First, consider the 3-year results shown in Table IV. Note that buying the 10 percent of funds with the highest past return rather than the 10 percent with the lowest past return produces an extra return in the next period of

⁷ Panel A of Table II shows that using both expenses and preexpense differential return does not increase the R^2 above its value when we include only postexpense differential return. The coefficient of determination does not change, and if we examine the coefficient on expenses and the coefficient on differential return plus expenses, we see that the net weight on expenses is approximately zero.

⁸ We use the Fama and MacBeth (1973) techniques for calculating t -statistics because they conveniently summarize numbers. However, in every case where we discuss average significance, each of the coefficients in the yearly regressions has the same sign each year, and the coefficient in question is statistically significant at the 1 percent level each year.

⁹ In the 1-year case, using both preexpense differential return (or alpha) and fund expenses together leads to higher association with future performance than using past performance or expenses separately.

Table III
Association over 1-Year Periods

The table shows statistics from regressions of annual measures of the indicated regressand variables at time $t + 1$ on annual measures of the indicated regressor variables at time t . The intercept, slope coefficients, and R^2 represent averages of five annual regressions. We compute t -statistics in the spirit of Fama and MacBeth (1973). Returns are annualized. The sample includes 52 funds. The sample period is January 1996 through December 2001.

Variables	Intercept	$t(\text{int})$	Slope 1	$t(\text{slope 1})$	Slope 2	$t(\text{slope 2})$	R^2
Panel A. Performance							
Differential return vs. differential return	-0.079%	-0.871	0.767	34.353			0.700
Differential return vs. expenses	0.021%	0.434	-1.094	-13.954			0.772
Differential return vs. expenses, differential return + expenses	0.027%	0.489	-1.034	-12.240	0.284	3.603	0.798
Alpha vs. alpha	-0.083%	-1.113	0.806	14.119			0.663
Alpha vs. expenses	0.010%	0.261	-0.949	-20.400			0.766
Alpha vs. expenses, alpha + expenses	0.016%	0.350	-0.961	-21.843	0.240	2.540	0.781
Panel B. Managerial Skill							
Differential return + expenses vs. differential return + expenses	-0.007%	-0.089	0.221	5.170			0.073
Alpha + expenses vs. alpha + expenses	0.032%	0.882	0.126	1.649			0.025
Panel C. Risk							
R^2 vs. R^2	0.405	0.873	0.595	1.284			0.190
Abs $(\beta - 1)$ vs. Abs $(\beta - 1)$	0.002	3.818	0.138	0.865			0.098
Panel D. Tax Efficiency							
(Div yld + exp)/Fund avg vs. (Div yld + exp)/Fund avg	0.489	4.931	0.514	5.261			0.313
Capital gain yld/Fund avg vs. Capital gain yld/Fund avg	0.466	1.738	0.542	2.169			0.278

Table IV
3-Year Prediction Deciles

The table shows mean 3-year statistics at time $t + 1$ for funds sorted into deciles based on 3-year measures of the indicated variables at time t . Returns are annualized. The sample includes 52 funds. The sample period is January 1996 through December 2001.

Decile, t	Performance			Managerial Skill		
	Differential Return, $t + 1$ Sorted by Differential Return, t	Differential Return, $t + 1$ Sorted by Expenses, t	Alpha, $t + 1$ Sorted by Alpha, t	Alpha, $t + 1$ Sorted by Expenses, t	Differential Return + Expenses, $t + 1$ Sorted by Differential Returns + Expenses, t	Alpha + Expenses, $t + 1$ Sorted by Alpha + Expenses, t
1	-0.118%	-1.082%	-0.031%	-1.102%	0.128%	0.098%
2	-0.147	-0.677	-0.270	-0.689	0.042	0.035
3	-0.226	-0.698	-0.211	-0.709	0.027	0.118
4	-0.300	-0.410	-0.330	-0.399	-0.022	-0.030
5	-0.318	-0.345	-0.307	-0.394	-0.079	0.028
6	-0.375	-0.349	-0.295	-0.332	0.077	0.000
7	-0.448	-0.348	-0.557	-0.363	0.002	-0.087
8	-0.657	-0.268	-0.664	-0.262	-0.051	-0.136
9	-0.782	-0.119	-0.798	-0.133	-0.035	-0.033
10	-1.083	-0.157	-1.098	-0.159	-0.067	-0.064
Top	0.078	-1.442	0.081	-1.463	0.138	0.164
Bottom	-1.318	0.078	-1.346	0.081	0.032	-0.049
Spearman	1.000	-0.939	0.939	-0.964	0.661	0.830

(continued)

Table IV—Continued

Decile, t	Risk		Tax Efficiency	
	$R^2, t + 1$ Sorted by R^2, t	Abs $(\beta - 1), t + 1$ Sorted by Abs $(\beta - 1), t$	(Div yld + exp)/avg, $t + 1$ Sorted by (Div yld + exp)/avg, t	Capital Gain yld/avg, $t + 1$ Sorted by Capital Gain yld/avg, t
1	0.999930	0.0052	1.092	2.018
2	0.999932	0.0028	1.054	0.629
3	0.999964	0.0019	0.981	1.701
4	0.999919	0.0035	0.974	1.538
5	0.999917	0.0017	0.968	2.197
6	0.999943	0.0030	1.030	1.002
7	0.999951	0.0012	0.971	0.501
8	0.999919	0.0016	1.005	0.169
9	0.999872	0.0012	0.970	0.377
10	0.999792	0.0016	0.895	0.139
Top	0.999979	0.0046	1.237	1.188
Bottom	0.999707	0.0004	0.854	0.141
Spearman	0.515	0.782	0.697	0.782

Table V
Annual Prediction Deciles

The table shows mean annual statistics at time $t + 1$ for funds sorted into deciles based on annual measures of the indicated variables at time t . Returns are annualized. The sample includes 52 funds. The sample period is January 1996 through December 2001.

Decile, t	Performance			Managerial Skill		
	Differential Return, $t + 1$ Sorted by Differential Return, t	Differential Return, $t + 1$ Sorted by Expenses, t	Alpha, $t + 1$ Sorted by Alpha, t	Alpha, $t + 1$ Sorted by Expenses, t	Differential Return + Expenses, $t + 1$ Sorted by Differential Returns + Expenses, t	Alpha + Expenses, $t + 1$ Sorted by Alpha + Expenses, t
1	-0.125%	-1.155%	-0.138%	-1.005%	0.057%	0.040%
2	-0.229	-0.729	-0.218	-0.661	0.029	0.088
3	-0.267	-0.655	-0.249	-0.552	0.022	0.037
4	-0.298	-0.488	-0.288	-0.425	-0.009	0.094
5	-0.374	-0.377	-0.278	-0.342	-0.045	0.007
6	-0.456	-0.403	-0.396	-0.341	-0.051	-0.026
7	-0.520	-0.329	-0.451	-0.273	-0.066	-0.010
8	-0.609	-0.264	-0.530	-0.264	-0.093	-0.034
9	-0.726	-0.153	-0.664	-0.129	-0.067	0.046
10	-1.128	-0.137	-0.963	-0.134	-0.056	0.046
Top	0.041	-1.506	-0.045	-1.348	-0.016	0.055
Bottom	-1.423	0.092	-1.235	0.058	-0.039	-0.040
Spearman	1.000	-0.988	0.988	-0.988	0.915	0.236

(continued)

Table V—Continued

Decile, t	Risk		Tax Efficiency	
	$R^2, t + 1$ Sorted by R^2, t	Abs ($\beta - 1$), $t + 1$ Sorted by Abs ($\beta - 1$), t	(Div yld + exp)/avg, $t + 1$ Sorted by (Div yld + exp)/avg, t	Capital Gain yld/avg, $t + 1$ Sorted by Capital Gain yld/avg, t
1	0.999931	0.0047	1.156	2.340
2	0.999939	0.0024	0.996	1.730
3	0.999936	0.0034	0.999	1.450
4	0.999922	0.0029	1.033	0.959
5	0.999938	0.0025	0.997	1.070
6	0.999922	0.0029	0.982	0.962
7	0.999901	0.0021	0.974	0.442
8	0.999897	0.0026	0.969	0.412
9	0.999875	0.0023	0.984	0.293
10	0.999819	0.0025	0.937	0.380
Top	0.999981	0.0059	1.202	2.550
Bottom	0.999675	0.0025	0.914	0.940
Spearman	0.855	0.491	0.842	0.952

0.97 percent per year. Buying the single fund with the highest past return rather than the single fund with the worst past return leads to a difference of 1.40 percent per year in future return. Furthermore, performance decreases as we move down the deciles, with a highly significant rank correlation between past differential return and future differential return.

For alpha, the results are analogous: Funds ranked in the top 10 percent according to alpha have an alpha that is 1.07 percent per year higher than funds in the bottom 10 percent, and the best fund outperforms the worst fund by 1.43 percent. Once again, the rank correlation between the deciles formed on past alpha and the subsequent realized alpha is highly significant. Clearly, these results are economically significant and indicate that investors interested in good future return should buy funds with higher past returns.

What if we buy funds based on expenses? Buying the 10 percent of funds with the lowest expenses rather than 10 percent with the highest expenses results in an extra differential return of 0.92 percent per year, while the lowest expense fund outperforms the highest expense fund by 1.52 percent. Similarly, the difference in next year's alpha from buying the lower expense decile of funds rather than the higher is 0.94 percent per year, while the difference between funds at the extreme of expenses is 1.54 percent. Once again, all rank correlations are highly significant.

When we examine Table V for 1-year results, we find that they are very similar to the 3-year results. All of the measures work well and show even more predictive power.

B. Predictability of Management Skill

In this section, we add expenses back to return to examine how well managers do, as opposed to investors, relative to holding a passive portfolio at their chosen risk levels. If mutual funds did not charge expenses, their performance would still likely differ from the return on an index for a number of reasons. First, by purchasing and selling securities, funds incur transaction costs that reduce their return below that of an index. Second, index funds need cash management policies to handle inflows and outflows from investors and policies regarding the timing of the reinvestment of dividends. Funds can freely choose their policies, while index returns are calculated based on a mechanical rule for reinvesting dividends and assuming no inflows or outflows. A fund's cash management practices can enhance or decrease performance relative to an index. Third, funds can freely choose how they handle sales and purchases caused by changes in the companies contained in the S&P index. Again, these changes are handled mechanically when calculating a return on an index. Fourth, index funds have policies for handling tender offers and mergers, while these are handled mechanically in index construction. Finally, index funds can lend securities and earn a return on the lent securities; the index return cannot do so.

Do managers add value before expenses, and if so, is this value predictable? We look at this in two ways. First, we examine alpha plus expenses, which

represents the fund's risk-adjusted return if there were no expenses. Table I shows that, on average, management adds 3.4 basis points (t -statistic 3.69) in value per year. A more naïve way to examine management skill is to compare a performance measure that is not risk-adjusted to a benchmark. Examining differential return plus expenses, management underperforms by 4.1 basis points per year. Management skill in the future is related to past skill over long periods. Panel B of Table II shows that, for a 3-year holding period, alpha plus expenses in the last 3 years regressed on alpha plus expenses in the first 3 years has a slope of 0.665 with a t -statistic of 4.558. Similarly, differential return plus expenses in the second 3 years regressed on the same variable in the first 3 years has a slope of 0.265 with a t -statistic of 2.328. In both cases there is regression toward the mean, so that the best estimate of future management skill is a fraction of its prior value. However, in both cases a statistically significant association exists between past and future management performance. Ranking funds in the first 3 years on the basis of alpha plus expenses and selecting the top decile rather than the bottom decile leads to differential management performance of 16.2 basis points per year. Similarly, ranking funds by differential return plus expenses in the first 3 years and selecting the top 10 percent rather than the bottom 10 percent leads to differential management performance of 19.5 basis points per year. Management performance is not only predictable in the tails, but the rank correlations across deciles are also statistically significant for both measures of performance.

Less predictability exists over the 1-year horizon. We examine the 1-year association by regressing the value of the variable during 1 year on its value during the prior year. Panel B of Table III shows the regression results. Although past differential return plus expenses is associated with its future value, past alpha plus expenses is not. Since the difference between return and alpha is the beta adjustment, this suggests that the association of differential return with its past value comes from many funds maintaining a stable beta that is different from one over time.

When measured by differential return, management skill across funds is highly predictable. Table V shows the average results based on 1-year horizon deciles. For differential return plus expenses, the rank correlation between past and future values is 0.915. The 10 percent of funds with the highest management skill in period t outperforms the index by 5.7 basis points per year in the subsequent period. Likewise, the 10 percent of funds with the lowest management skill in period t underperforms the index by 5.6 basis points in the next period. When we examine alpha plus expenses, the rank correlation is insignificant, with no difference in the subsequent performance of funds in the top and bottom deciles.¹⁰ All of these results together suggest that management skill is predictable in the long run, but only weakly predictable in the short run.

¹⁰ When we examine funds that are only open for institutional sales separately from funds that are open to individuals, we observe the same relationship in the subsamples as in the combined sample.

C. Predictability of Risk

The basic risk in purchasing an index fund is that the fund's return pattern does not match the index. Two reasons explain why the fund return can deviate from the index. First, the fund's beta could systematically differ from one. This could occur if, for example, the fund held cash to service inflows and outflows, and did not adjust the beta by using futures. We measure this type of tracking error with the absolute value of $\beta - 1$. Second, less-than-perfect replication could cause a fund's return to randomly deviate from the index even if the fund's beta is one on average. We measure this type of tracking error with R^2 .

Panel C of Table II shows that R^2 and the absolute value of the difference between beta and one computed over a 3-year period are each associated with their values computed over the previous 3 years. In each case, the slope of the regression of the value in the second 3 years on the value in the first 3 years is significant at the 1 percent level. However, Table IV shows that the Spearman rank correlation between past and future levels of one of the two risk measures, R^2 , is not significant at the 5 percent level. Even less association emerges when we examine a 1-year horizon. The regressions in Table III show that neither measure is significantly related to its prior year value. In contrast to the 3-year decile analysis, when we examine the 1-year deciles in Table V, the absolute value of $\beta - 1$ is not predictable, but R^2 is predictable.

If we look across funds at the variations in R^2 (from 0.9991 to 1.0000) and beta (from 0.979 to 1.005), we see very little difference. Thus, risk differences are unlikely to explain a large amount of the differences in cash flows among funds.

D. Predictability of Tax Efficiency

Index funds also differ in their tax efficiency. Dividends and realized capital gains affect tax payments. To examine the predictability of dividends and capital gains as a percentage of fund NAV, we divide each variable by its mean value across funds for each 1- and 3-year period.¹¹ Doing so allows us to pool the results. Both of these tax variables are highly predictable.

First, examine the 3-year results. Panel D of Table II shows that the slope coefficient in the regressions of the latter 3 years on the first 3 years is highly significant for both dividends and capital gains. Table IV shows that the 3-year rank correlation is also significant. Tables III and V indicate that a similar pattern arises with a 1-year horizon. In particular, the rank correlation indicates that past values predict future values with a high level of significance. However, selecting funds based on past dividends leads to only a small variation in future

¹¹ Dividends paid to shareholders are dividends paid to the fund less expenses. The predictability of expenses has already been examined. Thus, in examining the predictability of dividends paid to shareholders, we concentrate on the component that we have not examined: dividends paid to the fund. An interesting implication is that for taxable investors a higher expense ratio results in lower taxes as well as lower dividends. Thus, the government pays for part of any increase in expenses.

dividends, with the top decile 16 percent above average and the bottom decile 6 percent below average. Given an average dividend yield of 1.25 percent per year implies that dividend differences have very little effect on tax efficiency. However, an investor selecting funds based on capital gains can obtain a great deal of future variation. Capital gains in the top decile are 234 percent of the average, and capital gains in the lowest decile are 38 percent of the average. Capital gains average 2.24 percent per year over this period. Thus, the difference in tax efficiency caused by selecting funds on the basis of past realized capital gains is substantial across index funds.

E. Other Considerations

Investors might concern themselves with two other characteristics of index funds: 12b-1 fees and loads. The expense ratio includes actual 12b-1 fees, and so actual 12b-1 fees directly affect returns. The investor could care about maximum 12b-1 fees because they might predict future actual 12b-1 fees and future expenses. Investors could also care about maximum 12b-1 fees because high fees may indicate a lack of management concern with performance. Ten of the funds in our sample have 12b-1 fees during at least one year. Of these, seven have them during all six years. Thus, the presence of 12b-1 fees is highly predictable. Furthermore, although four funds adjust the magnitude of their 12b-1 fees at some point during the sample period, these adjustments are small; once adjusted they tend to remain consistent. Thus, past 12b-1 fees predict future 12b-1 fees.

Nine of the funds have loads during the sample period, mostly of the same magnitude; six funds have loads in every year, two drop loads, and one adds a load. Two funds change the size of the load during the sample period, but the change is small enough not to affect its relative rank. Thus, both the presence of loads and their magnitude are highly predictable on the basis of their presence and size during the prior year.

Since both 12b-1 fees and loads primarily support marketing and distribution, an investor might examine them separately from other expenses in deciding which fund to buy.

III. Cash Flows and Fund Characteristics

In this section we discuss a set of variables, including performance characteristics, that might be associated with future fund cash flows, and we then examine their impact. As the prior section shows, the important performance characteristics of index funds are easily understood, and for most of these characteristics we can easily predict future differences across funds. We group the performance characteristics into four categories:

- A. Predictors of overall performance
 - 1. Differential return
 - 2. Alpha
 - 3. Expenses

- B. Predictors of management performance
 - 1. Differential return plus expenses
 - 2. Alpha plus expenses
- C. Predictors of risk
 - 1. Absolute value of $(\beta - 1)$
 - 2. R^2 /average R^2
- D. Predictors of tax efficiency
 - 1. Percentage dividends/average percentage dividends
 - 2. Percentage capital gains/average percentage capital gains

We measure each of these variables at the start of the period in which we measure cash flow. Thus, all of them could be known to the investor before the investor decides on an investment. Several researchers have suggested that, in addition to performance characteristics, a number of other variables might affect cash flow.¹² We can categorize these variables into two groups: marketing and spillover. Marketing includes two types of influences. The first is any action management takes to increase investor awareness of the fund, including advertising, promotional brochures, website development, etc. The second involves compensating brokers and investment advisors for selling the fund (rather than competing funds). Although we do not have direct measures for either of these influences, some good proxies exist. Funds primarily use a large portion of 12b-1 fees and the majority of loads to reward salespersons for selling the fund, with the remainder paying for the general marketing effort. Consequently, the size of 12b-1 fees and loads should be related to the marketing effort. In addition, the management fee is the part of the expense ratio that is less likely to be used for the marketing effort.¹³ Thus we use the difference between the expense ratio and the management fee as another measure of marketing effort. Investor awareness of a fund can also be a function of the fund's frequency of inclusion in financial publications, as well as its direct marketing effort. To proxy for general investor awareness we use both the size of the fund and the age of the fund.

Spillover refers to the idea that membership in a fund family increases inflows because investors prefer to concentrate their investments within one family.¹⁴ For example, an investor who is attracted to a fund family because of a desirable small stock fund might invest in the fund's index fund for diversification purposes and a desire to hold all mutual fund investments in one family. We use the following variables to capture spillover effects: (1) The natural log of the total net assets of the fund family (to measure family size); (2) the number of different Morningstar categories held by the family (to measure the investor's chance to diversify within the family); (3) fund flows into a family excluding

¹² See Sirri and Tufano (1998) for a discussion of marketing variables.

¹³ It might seem surprising that the part of expenses labeled "management fee" does not simply compensate managers. Conversations with fund managers indicate they can and do sometimes pay brokers part of the management fee.

¹⁴ Spillover becomes less important with the advent of systems that allow an investor to transfer funds among many families, such as the systems marketed by Fidelity and Schwab.

the flows into the index fund (to measure the popularity of the family); and (4) the presence of a fund within the family that received a Morningstar five-star rating. We include this last variable because Khorana and Servaes (1999) show that a major determinant of flows into a family is the presence of a highly rated fund.

Another variable might impact cash flow: The impression that investors and financial professionals have of the service provided by the fund family. Dalbar, Inc. constructs the leading measure of service rankings, surveying more than 2600 investors and professionals on their impression of fund family service quality. Dalbar has generously provided us with data over our sample period for the families it follows. However, Dalbar gathers data for less than half of the families represented in our sample (25 out of 52 in the year with the greatest coverage). Consequently, we examine the service variable separately rather than including it directly in our regression.

Like others, we use a percentage flow variable to measure the flow into a fund (see, e.g., Sirri and Tufano (1998) and Warther (1995)). However, our measure differs slightly from measures used in previous research. The index funds in our sample differ dramatically in size. An extremely large fund has a much harder time attracting the same percentage of new capital as a small fund. Thus, we first compute expected flow based on a regression of cash flow on size. We then define percentage surprise by taking the difference between the actual flow (the net growth rate in fund assets beyond that due to capital gains and reinvested dividends) and expected flow and dividing the difference by the fund's total net assets.

Table VI shows the results of our analysis.¹⁵ The odd-numbered columns show the coefficients of regressions that include one variation of the variables discussed above. The even-numbered columns show the coefficients of the regressions that include only those variables that have a significant relationship with the surprise in cash flow. Columns also differ by the performance variable used: one and two use differential return, three and four use alpha, and five and six use the expense ratio. The results are very similar regardless of the predictive measure of performance.

Although columns one, three, and five use one iteration of the full range of variables that the literature has suggested, we examine variations of similar variables separately. Thus, we use the log of total net assets and age, which are highly correlated, in separate regressions. Table VI reports only the most relevant alternative. However, in the text we discuss variations of variables not shown in the table.

We first examine the regressions shown in columns two, four, and six. First, note that the overall explanatory power is at least as high as in other studies

¹⁵ We also run each of the regressions on panel data incorporating an intercept for each fund in the sample. An *F*-test indicates that we cannot reject the hypothesis that the firm intercepts are jointly equal to zero. In addition, we also use the criterion proposed by Akaike (1974) to test for the inclusion of additional variables. This test indicates that we should not include the firm intercepts. Hence, we report the simpler results in Table VI.

Table VI
Cash Flow Regressions

The table shows regression coefficients from pooled regressions of annual measures of (actual cash flow – expected cash flow)/TNA at time $t + 1$ on annual measures of the indicated regressor variables at time t . The t -statistics are in parenthesis. The sample includes 52 funds. The sample period is January 1996 through December 2001.

	1	2	3	4	5	6
Intercept	0.078 (0.10)	0.041 (0.23)	0.101 (0.13)	0.150 (0.82)	–0.049 (–0.06)	0.269 (1.36)
Differential return	100.804 (2.92)	66.255 (3.33)				
Differential return + exp	–74.550 (–1.44)					
Alpha			1269.311 (3.07)	968.961 (3.34)		
Alpha + expense ratio			–826.258 (–1.33)			
Expense ratio					–103.225 (–3.00)	–98.076 (–3.28)
$R^2 - 1$	1307.616 (2.21)	1054.962 (2.09)	1299.104 (2.20)	1175.370 (2.34)	1332.432 (2.26)	1145.748 (2.27)
Abs ($\beta - 1$)	1.928 (0.09)		–3.861 (–0.17)		2.247 (0.11)	
(Div yld + exp)/Fd avg	0.601 (0.98)		0.617 (1.02)		0.729 (1.23)	
Capital gain yld/Fd avg	–0.018 (–0.33)		–0.017 (–0.32)		–0.023 (–0.44)	
12b – 1 dummy	–0.243 (–1.14)		–0.242 (–1.14)		–0.273 (–1.31)	
Load dummy	0.587 (2.33)	0.353 (1.68)	0.594 (2.36)	0.365 (1.73)	0.596 (2.36)	0.425 (1.90)
ln(TNA)	–0.012 (–0.21)		–0.013 (–0.23)		–0.004 (–0.07)	
Number of family fds	0.003 (2.49)	0.002 (2.14)	0.003 (2.50)	0.002 (2.09)	0.003 (2.44)	0.002 (2.10)
Fam 5-star fd dummy	–0.350 (–1.52)		–0.342 (–1.49)		–0.341 (–1.49)	
Institutional dummy	–0.286 (–2.10)	–0.213 (–1.70)	–0.281 (–2.06)	–0.209 (–1.68)	–0.302 (–2.25)	–0.260 (–2.07)
Year 1 dummy	–0.582 (–2.74)	–0.561 (–2.70)	–0.628 (–3.02)	–0.686 (–3.36)	–0.615 (–2.96)	–0.639 (–3.12)
Year 2 dummy	–0.329 (–1.58)	–0.179 (–0.95)	–0.391 (–2.16)	–0.362 (–2.08)	–0.405 (–2.25)	–0.379 (–2.18)
Year 3 dummy	0.043 (0.25)	0.084 (0.49)	–0.007 (–0.04)	–0.027 (–0.16)	0.029 (0.167)	0.045 (0.26)
Year 4 dummy	0.054 (0.28)	0.171 (0.98)	0.018 (0.10)	0.066 (0.39)	–0.020 (–0.12)	–0.010 (–0.06)
R^2	0.264	0.242	0.265	0.242	0.262	0.240

of this type based on actively managed funds.¹⁶ The high explanatory power is understandable, given the ease with which we can predict characteristics for index funds. Of the variables measuring fund characteristics, overall performance and one risk variable are significant. Whether we measure overall performance by differential return, alpha, or expense ratio, fund flow shows a highly significant coefficient on the performance measure. The coefficient is also economically significant. Taking the difference between the 75th and 25th percentiles of any performance measure leads to a difference in flows of about 30 percent of total net assets.

Although the regression coefficient on $R^2 - 1$ is very large and statistically significant, the differences in $R^2 - 1$ across funds are very small, so that the net impact is small. The average difference in cash inflow associated with the 25th percentile of $R^2 - 1$ rather than 75th percentile is about 0.2 percent. The other risk measure, the absolute value of beta $- 1$, is insignificant. The insignificance of $\beta - 1$ is not surprising given the small variation in the variable.¹⁷ The lack of significance for the tax efficiency variables is somewhat more surprising. Although dividends vary little, capital gains vary substantially, and capital gains efficiency is predictable. Investors are either unconcerned with capital gains or unaware of its predictability.

Load is the only significant marketing variable, and it enters with a positive sign.¹⁸ Load has two impacts. First, load constitutes a major cost for investors and substantially reduces their return. Second, load rewards brokers and financial planners for including the fund in their clients' portfolios. On net, the incentive for brokers and financial planners to push the fund matters more for new flows than the effect of reduced flows associated with investors avoiding high cost and poorly performing funds. None of the other marketing variables—12b-1 fees, fund size, age (results not shown), or expenses less management fee (results not shown)—are significant.

The number of types of funds in the fund family is the only important spillover variable. An index fund that is part of a family offering a variety of other types of funds attracts more cash flow. This is logical, because many investors wish to stay with one fund family to facilitate inter-fund transfers and record keeping. Neither family size, flow of funds into family (results not shown), nor the presence of a star fund explains cash flows.

We separately run the analysis for institutional funds and funds for individuals. The results are essentially the same, with the following caveats. First, we anticipated that capital gains would be important in the individual sample.

¹⁶ See Sirri and Tufano (1993, 1998) and Chevalier and Ellison (1997), for example. In fact, we find higher explanatory power since before running this regression, we remove the impact of size on cash flow. The R^2 between cash flow and size is 0.21. See Bergstresser and Poterba (2002) for a study of cash flows and taxes. See Del Guercio and Tkac (2001) for a study of cash flows and Morningstar ratings. See Jain and Wu (2000) for a study of advertising and fund flows.

¹⁷ The extremely small variation in beta indicates that S&P index funds have developed a cash management policy (possibly involving futures) that allows minimization of the impact of cash flows on their performance.

¹⁸ Excluding load decreases the explanatory power to 19 percent.

Most institutions are tax-exempt, but individuals are not. Furthermore, capital gains payouts vary substantially across funds. However, the relationship is not significant, which indicates that individuals either do not pay attention to taxes or are unaware of the differences across funds and their predictability. Second, we had anticipated that performance would produce a greater effect on cash flows for institutional funds. This does not occur. The principal reasons for this are likely to be the small size of the institutional sample and the lack of a substantial variation in the performance of funds primarily sold to institutions.

As a final step in this part of the analysis, we examine the relationship between unexplained cash flow (the residuals from each of the regressions described above) and service rankings. Dalbar ranks fund families from best to worst according to perceived service quality. We compute the Spearman rank correlation between the perceived service quality of quartiles formed according to unexplained cash flow. The rank correlation is insignificant each year and overall. In fact, it is very close to zero and often has the wrong sign. Dalbar carefully constructs its survey of perceived service. The results indicate that, at least for index funds, either the variables we incorporate in the regression already capture the important aspects of service, or the quality of service that the fund family offers does not matter to index fund investors.¹⁹

IV. How Well Do Investors Do?

In this section we examine the performance of investors who purchase S&P 500 index funds. At any point in time, the dollar flows into index funds purchase a portfolio of holdings for which we can easily measure future returns. We compare the actual returns earned by investors to returns of naïve portfolios and returns earned by investors making informed choices.

We first describe our methodology in more detail and then present our results. Each month we form a portfolio where the weight on each fund is the net cash flow into that index fund divided by the aggregate cash flow into all index funds. This converts the dollar cash flow into each fund into a fraction of the total cash flow invested in all S&P index funds that went into each fund. We then measure the return on this portfolio over 1- and 3-year holding periods. More specifically,

$$R_{A,t,t+j} = \sum_{i=1}^N \frac{CF_{i,t}}{\sum_{i=1}^N CF_{i,t}} R_{i,t,t+j}, \quad (2)$$

where $R_{A,t,t+j}$ is the actual annualized return earned on the portfolio of index funds bought by investors at time t assuming they hold the portfolio until time $t + j$, j is 12 or 36 months, N is the number of funds in our sample, $CF_{i,t}$ is the

¹⁹ A third explanation exists. Given the small number of funds for which we have service measurements, the sample may be inadequate to detect a relationship.

net cash inflow to fund i at time t ,²⁰ and $R_{i,t,t+j}$ is the actual annualized return of fund i between time t and $t + j$.

Equation (2) measures the actual return earned by new purchases of index funds, assuming either a 1-year or a 3-year holding period. We construct a new portfolio each month from January 1996 to January 2001 for the 1-year holding period and from January 1996 to January 1999 for the 3-year holding period.

For both the 1- and 3-year holding periods, we construct reference portfolios each month to determine how well investors do. The first two reference portfolios examine the return earned when we use naïve rules to form portfolios. The first portfolio is equally weighted. Each month we allocate all net new cash inflow equally to each fund in the sample. The second portfolio is a market-weighted reference portfolio. Each month we allocate all net new cash inflow across index funds in proportion to the total net assets of each index fund at that point in time.

We also compare actual performance to the performance an investor could earn by selecting index funds using past data. If an investor recognizes that index funds are a commodity, the investor should buy funds based on variables that predict return. We thus assume that the investor buys one of the following three portfolios: the 10 percent of funds with the lowest past expense ratios, the highest past returns, or the highest past alphas. This results in three additional reference portfolios that represent how well smart investors do.

To gauge how poorly investors acting perversely on past information would do, we form three portfolios based on the 10 percent of index funds with the highest past expenses, lowest returns, and lowest alphas. Finally, we examine the return earned by the 10 percent of the net cash flow that goes into the index funds with the worst past performance.

We begin with 1-year holding periods. Before examining investor performance, we examine the return from buying funds based on past returns. If an investor buys the top 10 percent of funds ranked on past return rather than the bottom 10 percent, he earns an extra 92 basis points per year. This is not only statistically significant at the 1 percent level, but is also clearly economically significant. How well do investors do? They do considerably better than spreading their money across all index funds in equal amounts, but they do four basis points worse than simply investing their money in index funds proportional to their market value. Furthermore, they do 15 basis points per year worse than simply buying the top 10 percent of index funds ranked on return in the previous period. Investors are making decisions that produce returns that are statistically significantly worse than buying based on past returns, or, even more naively, buying funds proportional to their size. Some investors are clearly making bad decisions in choosing index funds. If we examine the return earned by the 10 percent of the money that flows into the funds with the lowest past return, we see that these dollars earn on average 53 basis points less per

²⁰ We exclude funds with net outflows in the month the net outflow occurs. We intend to see how new cash invested in mutual funds performs.

year than investments in the top 10 percent of funds. Similar results hold for the 3-year case shown in Table VII.

If we examine the results for buying based on past expenses or past alpha, also shown in Table VII, we find very similar results. Investors earn 20 basis points per year fewer than buying the 10 percent of funds ranked lowest on past expenses and 14 basis points per year fewer than buying the 10 percent of funds ranked highest on past alpha. When we examine 3-year holding periods, the results are almost identical to the 1-year case.

We should mention that in performing the return calculations, we compute return as if loads do not exist. Although only a small percentage of index funds have loads, including the impact of loads on returns would make investors seem even more irrational. The return that investors receive on their actual portfolios would be 8 basis points lower for a 1-year holding period and 3 basis points lower for a 3-year holding period.²¹ Since none of the top 10 percent of fund portfolios has a load, those loads represent an additional loss to investors who do not buy funds based on past performance.

When they purchase index funds, many investors seem to be making decisions that violate rationality within the classic paradigm of financial economics. At the very least, many are ignoring every one of a set of indicators that predict future return. They could be relying in part on other factors, such as taxes or services.

If the portfolios investors hold have lower capital gains than the funds selected by using past performance, this might provide justification for investor actions. However, the capital gains earned by the funds investors hold are higher than the capital gains for a portfolio constructed on the basis of past return, alpha, or expenses, and the differences are statistically significant at the 1 percent level.

The last defense we can imagine for holding funds that offer a lower return is that they might provide better services. We do not have a direct measure of services, but two proxies seem appropriate. Many mutual fund investors want the ability to switch among different types of mutual funds. Consequently, investors might prefer to hold funds that belong to fund families that offer funds of more than one type. We use the number of Morningstar categories offered by a fund family as a measure of the number of types of funds offered. We find that investors hold index funds with slightly fewer Morningstar categories than portfolios selected on the basis of past return, expenses, or alpha. The differences are significant at the 5 percent level. Diversification across fund types and the ability to switch money around cannot account for the differences in investors' choice.

Another variable that might proxy for services is the total amount of dollars under management by the fund's family. Large families are in a position to provide more services than smaller families. However, this cannot serve as a defense for investor actions, because the families of the funds held by investors

²¹ It is interesting to note that funds with front end loads almost always have higher expense ratios, even though we do not include the loads in computing expense ratios.

are smaller than the families of funds in the top decile based on any of our three predictive variables.

What, then, can account for so many investors behaving inconsistently with the assumption of rationality in frictionless markets, the classical paradigm of financial economics? We believe that they are affected more by salesmanship and marketing activities than they are by financial analysis. Although we do not have direct proof of this, we have indirect proof. The portfolios investors hold have higher 12b-1 fees, higher loads, and higher expenses than the best portfolios. Funds use almost all 12b – 1 fees and loads and part of the expense ratio to reward salespersons and to market funds. Apparently, this marketing effort produces the desired effect.

V. Conclusion

In this paper we show that one can easily forecast all of the characteristics of an S&P 500 index fund that an investor might care about. An investor might be concerned with differential return or alpha. Future values of these variables are easily forecast by their past values or expense ratios. Similarly, evidence indicates that a fund's risk, whether measured by R^2 or by deviations of beta from one, can be forecast from prior values. Finally, differential tax efficiency comes about primarily from differences in capital gains distributions, and these differences can also be forecast by past values.

Given that one can forecast the characteristics of index funds that an investor should be concerned with, one would expect that these characteristics would primarily determine investor cash flows to different index funds. This is not the case. Adding the standard marketing and spillover variables improves the ability to account for cash flows, but a lot of cash flows remain unexplained. One possible reason for this is that our analysis does not include all of the relevant marketing variables. For example, funds sign agreements with financial advisors and brokers to pay for fund flow. Although some of this is captured by 12b-1 fees, loads, and expenses less management fees, much of it is not. The terms of these agreements are not observable. Thus, financial advisors or brokers who advise clients to purchase funds to improve their own profits could cause much of the flow into index funds. Although one could argue that the investor is simply paying for advice on which funds to buy, the individual supplying the advice is often motivated by a compensation system that is incompatible with the best interests of investors. Furthermore, this is an inefficient way of compensating for allocation advice.

Finally, we examine the actual return earned by investors on their purchases of index funds and compare it with the return earned by simply purchasing the 10 percent of funds with the highest differential return, alpha, or expenses in the prior period. Investment based on any of the predictors of future performance results in substantial extra return. Furthermore, these rules lead to selections that (1) have zero loads (where actual cash flows sometimes go to funds with loads); (2) are associated with families with a greater number of options (which facilitates diversification within a fund family); (3) are more

tax-efficient; and (4) correlate more highly with the index. On any dimension, selecting funds on the basis of past performance leads to owning a superior set of funds. Instead, investors buy funds with higher marketing costs than the best-performing funds.

S&P 500 index funds are a commodity. The difference in risk across funds is very small, and differences in returns are easily forecast. Thus, the classic paradigm of financial economics would imply that a rational investor should choose based on expected return, including tax efficiency. If investors acted in this way, one would expect cash flows to go to the fund or funds that expect to offer the highest return. Yet a large amount of new cash flow goes to the poorest-performing funds. Furthermore, in a frictionless rational market, one would expect the low-return funds, some of which have expense ratios over 1 percent, to disappear. This is not the case. In fact, index funds added to the market since the beginning of our sample period cost more, with expense ratios averaging 0.77 (and a maximum over 2 percent) compared to an average of 0.44 percent for the funds in our sample. A higher percentage of the funds that have recently been introduced have loads and 12b-1 fees. Loads on the new funds average 1.48 percent compared to 0.93 percent for our sample, and 12b-1 fees average 0.29 percent per year compared to 0.13 percent. Finally, in our sample period, the 10 percent of funds with the highest expenses grows at an average annual growth rate of 20.5 percent, compared to 11.8 percent for the low-cost funds.

How can we explain the results? Any market consists of a set of informed rational investors and a set of uninformed investors. The arbitrage activities of informed investors make markets efficient. But the only thing an informed investor can do in the market for index funds is to buy the good-performing funds—no arbitrage is possible. In such a market, all that is necessary for inferior funds to exist and grow is a set of uninformed investors and a set of distributors who have an economic incentive to sell inferior products. In a market where arbitrage is impossible, we may be disappointed, but we should not be surprised when inferior products exist and even prosper.

Much of the financial and economic literature assumes that the law of one price holds. As we show in this paper, in markets where arbitrage is not available, the law of one price need not hold. It will hold only if all investors are rational.²²

REFERENCES

- Akaike, H., 1974, A new look at the statistical model identification, *IEEE Transactions on Automatic Control* 19, 716–723.
- Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773–806.

²² Elsewhere, Elton, Gruber, and Rentzler (1989) find that public commodity funds grow and prosper despite a return below the T-bill rate. This is another market where the lack of arbitrage allows inferior investments to exist and prosper.

- Bergstresser, Daniel, and James Poterba, 2002, Do after-tax returns affect mutual fund inflows? *Journal of Financial Economics* 63, 381–414.
- Blume, Marshall E., and Irwin Friend, 1975, The asset structure of individual portfolios and some implications for utility functions, *Journal of Finance* 30, 585–603.
- Bollen, Nicolas P. B., and Jeffrey A. Busse, 2002, Short-term persistence in mutual fund performance, Working paper, Vanderbilt University.
- Brown, Stephen J., William N. Goetzmann, Roger G. Ibbotson, and Stephen A. Ross, 1992, Survivorship bias in performance studies, *Review of Financial Studies* 5, 553–580.
- Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.
- Chevalier, Judith, and Glenn Ellison, 1997, Risk taking by mutual funds as a response to incentives, *Journal of Political Economy* 105, 1167–1200.
- Christoffersen, Susan E. K., and David K. Musto, 2002, Demand curves and the pricing of money management, *Review of Financial Studies* 15, 1499–1524.
- Cohen, Randolph B., Paul A. Gompers, and Tuomo Vuolteenaho, 2002, Who underreacts to cash-flow news? Evidence from trading between individuals and institutions, *Journal of Financial Economics* 66, 409–462.
- Del Guercio Diane, and Paula Tkac, 2001, Star power: The effect of Morningstar ratings on mutual fund flows, Working paper, University of Oregon.
- Elton, Edwin J., Martin J. Gruber, and Christopher R. Blake, 1996a, Survivorship bias and mutual fund performance, *Review of Financial Studies* 9, 1097–1120.
- Elton, Edwin J., Martin J. Gruber, and Christopher R. Blake, 1996b, The persistence of risk-adjusted mutual fund performance, *Journal of Business* 69, 133–157.
- Elton, Edwin J., Martin J. Gruber, and Joel C. Rentzler, 1989, New public offerings, information, and investor rationality: The case of publicly offered commodity funds, *Journal of Business* 62, 1–15.
- Fama, Eugene, and J. MacBeth, 1973, Tests of the multiperiod two-parameter model, *Journal of Financial Economics* 1, 43–66.
- Ferris, Stephen P., Robert A. Haugen, and Anil K. Makhija, 1988, Predicting contemporary volume with historic volume at differential price levels: Evidence supporting the disposition effect, *Journal of Finance* 43, 677–697.
- Froot, Kenneth A., and Emil M. Dabora, 1999, How are stock prices affected by the location of trade? *Journal of Financial Economics* 53, 189–216.
- Green, Richard C., and Kristian Rydqvist, 1997, The valuation of nonsystematic risks and the pricing of Swedish lottery bonds, *Review of Financial Studies* 10, 447–480.
- Grinblatt, Mark, and Matti Keloharju, 2001, What makes investors trade? *Journal of Finance* 56, 589–616.
- Gruber, Martin J., 1996, Another puzzle: The growth in actively managed mutual funds, *Journal of Finance* 51, 783–810.
- Hirshleifer, David A., James N. Myers, Linda A. Myers, and Siew Hong Teoh, 2001, Do individual investors drive post-earnings announcement drift? Working paper, The Ohio State University.
- Huberman Gur, 2001, Familiarity breeds investment, *Review of Financial Studies* 14, 659–680.
- Jain, Prem C., and Joanna Shuang Wu, 2000, Truth in mutual fund advertising: Evidence on future performance and fund flows, *Journal of Finance* 55, 937–958.
- Khorana, Ajay, and Henri Servaes, 1999, The determinants of mutual fund starts, *Review of Financial Studies* 12, 1043–1074.
- Longstaff, Francis A., Pedro Santa-Clara, and Eduardo S. Schwartz, 2001, Throwing away a billion dollars: The cost of suboptimal exercise strategies in the swaptions market, *Journal of Financial Economics* 62, 39–66.
- Odean, Terrance, 1998, Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775–1798.
- Rietz, Thomas A., 1998, Enforcing arbitrage restrictions in experimental asset markets, Working paper, University of Iowa.
- Rosenthal, Leonard, and Colin Young, 1990, The seemingly anomalous price behavior of Royal Dutch/Shell and Unilever N.V./PLC, *Journal of Financial Economics* 26, 123–141.

- Sirri, Erik R., and Peter Tufano, 1993, Buying and selling mutual funds: Flows, performance, fees and service, Working paper, Harvard Business School.
- Sirri, Erik R., and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance* 53, 1589–1622.
- Warther, Vincent A., 1995, Aggregate mutual fund flows and security returns, *Journal of Financial Economics* 39, 209–235.
- Zheng, Lu, 1999, Is money smart? A study of mutual fund investors' fund selection ability, *Journal of Finance* 54, 901–933.