The Journal of Financial Research • Vol. XX, No. 2 • Pages 145–158 • Summer 1997

COGNITIVE DISSONANCE AND MUTUAL FUND INVESTORS

William N. Goetzmann
Yale School of Management

Nadav Peles
J. P. Morgan Investment Management

Abstract

We present evidence from questionnaire responses of mutual fund investors about recollections of past fund performance. We find that investor memories exhibit a positive bias, consistent with current psychological models. We find that the degree of bias is conditional upon previous investor choice, a phenomenon related to the well-known theory of cognitive dissonance. Psychological and economic frictions in the mutual fund industry are examined via a cross-sectional study of equity mutual funds. We find an unusually high frequency of poorly performing funds, consistent with investor "inertia." We also examine the differential responses of investment dollars to past performance, controlling for survivorship. These show that the effect is confined to the top quartile. We find little evidence that the response to poor performance is unusual.

I. Introduction

One of the greatest mysteries in the mutual fund industry is why some investors stay with funds that consistently perform poorly. Several researchers note that investor dollars flow into winning funds more rapidly than they flow out of losing funds. This differential is taken as evidence of irrationality (Ippolito (1992)), differential management services and high transactions costs (Sirri and Tufano (1992)), and a failure of investor probability heuristics (Harliss and Peterson (1994)). In this paper we provide evidence that investor psychology may affect the fund-switching decision. Questionnaires taken from two groups of mutual fund investors suggest that investor aversion to switching from poor performers may be explained by overly optimistic perceptions of past mutual fund

The authors thank Stephen Brown, Ken French, Bruce Greenwald, Richard Herrnstein, David Ikenberry, Narasimhan Jegadeesh, Ivo Welch, an anonymous referee, and the participants in the NEBER Behavioral Finance Program for helpful discussion. We thank Stephen Livingston for data collection. Gur Huberman contributed substantially to the preliminary draft of this paper. The first author thanks Columbia Business School and the Yale Finance Institute for summer research support.

performance. Samples of both educated and casual mutual fund investors show that investor recollections of past performance are consistently biased above actual past performance. This recollection bias may be why investors justify remaining in funds that consistently perform poorly. Although investor inertia might actually be due to high economic switching costs, our evidence suggests that investors nonetheless adjust their beliefs to support past decisions.

The tendency to adjust beliefs to justify past actions is an example of the psychological phenomenon termed by Festinger (1957) as cognitive dissonance. Festinger's theory asserts that individuals are distressed by conflicting cognitive elements, such as a discrepancy between empirical evidence and past choice, and that they alter their beliefs to reduce this discomfort. The key feature of dissonance is that individuals alter their beliefs to conform to their past actions. In the context of investment decision making, cognitive dissonance can be considered a psychological cost that investors seek to reduce by adjusting their beliefs about the efficacy of past investment choices.

The theory of cognitive dissonance has long been useful in studies of consumer behavior.[1] Erlich, Guttman, Schonbach, and Mills (1957) examine consumer response to advertising after a major purchasimg decision: choosing a new car. They find that new car owners selectively noticed advertisements that re-enforced the efficacy of their recent decision. The advertisements reduced the uncertainty they felt about the wisdom of their choice. Akerlof and Dickens (1982) consider cognitive dissonance in the labor market. They show how a rational individual may adjust beliefs about job risk to reduce dissonance. The choice of a mutual fund is no less anxiety producing than the choice of a new automobile or a new job; indeed, it is arguably more so. Thus, dissonance may be a factor in the fund-evaluation process.

Several researchers investigate the possible psychological basis for investor behavior. De Bondt and Thaler (1985) argue that mean reversion in stock prices is evidence of investor overreaction where investors overemphasize recent firm performance in forming future expectations. Shiller (1988) interprets evidence of excessive volatility in asset returns as suggestive of investing "fads." DeLong, Shleifer, Summers, and Waldmann (1990) incorporate irrational traders with erroneous stochastic beliefs into a model of asset markets to explain the equity premium puzzle. These and related studies typically focus on price and return behavior to infer something about investor attitudes, beliefs, and behavior.[2] We find experimental evidence on the role of investor psychology in decision making. Thaler, Kahneman, and Knetsch (1992), for instance, find an "endowment" effect among subjects endowed with even a relatively low-cost gift. That is, people are more likely to believe something they own is better than something

[1]See Britt (1978) for a survey of the application of the theory of cognitive dissonance to marketing research.
[2]Shiller's (1988) survey of institutional investor beliefs about the 1987 crash is an exception.

they do not own. Samuelson and Zeckhauser (1988) find an endowment effect among experimental subjects asked to consider different investment strategies.

In this study we offer direct evidence of investor beliefs, based on actual rather than hypothetical choices about investments. Our findings are consistent with an endowment effect. In addition, we have a small but unique sample of responses that allows us to differentiate between an endowment effect and beliefs conditional upon past choice.[3] The experimental evidence in the cognitive psychology literature suggests both mechanisms should influence investor beliefs and actions. In fact, in this smaller sample, we find evidence of both a dissonance effect and a strong endowment effect.

II. Mutual Fund Investor Questionnaire

While information about mutual funds is broadly accessible, information about investors is not. Because of the desire for privacy, no single source of information about mutual fund investors identifies personal investments. Consequently, we used a questionnaire to gather information from several mutual fund investors. The questionnaire requested information about exactly what mutual funds they use and how they believe these funds performed in the past. We collected samples from two groups of mutual fund investors. The first sample was collected from members of a state chapter of the American Association of Individual Investors (AAII). This group is assumed to be well informed about the past performance of their investments, relative to the general population. The second sample was collected from a group of professional architects who have a defined contribution profit sharing plan and who also invest in mutual funds. Their profit sharing plan does not allow the beneficiaries to choose among plans and managers. Thus, the profit sharing plan may be considered an endowment for vested employees, while the subjects' personal mutual fund investments require choice.[4]

The questionnaire for both groups asked for a list of the mutual funds owned by the respondent, an estimate of the difficulty in choosing the fund, and whether it was a load fund. In addition, the questionnaire (circulated in 1993) asked for an estimate of the 1992 return of the fund and an estimate of the

[3]Early researchers into cognitive dissonance recognize the difficulty of rejecting the endowment effect alternative. Brehm (1956) controls for an endowment effect in his experimental study of post-decision changes in the desirability of alternatives. He finds some evidence of an endowment effect, but concludes that changes in beliefs and information evaluation by his subjects were, nonetheless, strongly conditioned upon the active choice decision.

[4]The profit sharing plan retained a money manager who maintains a 50/50 balance between stocks and bonds. The investment results are reported annually to the beneficiaries. Additional information about the securities held by the plan, as well as the management fees, is available upon request.

TABLE 1. Results of Questionnaire on Perceptions of Investment Manager Performance.

	Sample I Funds	Sample II Funds	Sample II Profit Share
Estimated return - Actual return	3.40%	8.58%	2.36%
Standard errors	2.32%	3.38%	1.22%
Estimated relative return - Actual relative return	5.11%	4.17%	3.42%
Standard errors	2.27%	2.49%	1.36%
Satisfaction	2.41	4.17	5.64
Standard errors	0.29	0.62	0.72
Difficulty	7.08	6.22	
Standard errors	3.7	2.25	
Years before switch	2.15	2.39	
Average three-year alpha ranking	NA	0.34	
Counts	57	18	11

Notes: These results are based on responses to questionnaires described in the text. Actual returns were obtained from Morningstar for individual funds. Benchmark performance figures were obtained from Ibbotson Associates. For 1992, they were 7.67 percent for the S&P 500 and 7.68 percent for the Merrill Lynch bond index. The benchmark for the profit sharing account is assumed to be based on a 50/50 mix between stocks and bonds; however, during 1992 the results were not sensitive to manager deviations from this allocation. The probability of rejection is based on a two-tailed *t*-test with unknown variances. Although reported, dependency across observations and deviations of variables from normality violate necessary assumptions for the test to be correctly specified. Average three-year alpha ranking is taken from estimates by Morningstar. It is based on rankings of the capital asset pricing model alpha among funds of similar style from July 1991 to July 1993.

percentage by which the fund led or lagged a benchmark.[5] Subjects were asked how many years of poor performance it would take before they would switch to another fund and how satisfied they were with the manager's performance. Subjects in the second group were asked to estimate the absolute and relative performance of the profit sharing plan and to estimate their satisfaction with the manager.

We analyzed the results of the two groups separately. The perceptions of mutual fund performance were compared with actual performance reported by Morningstar, which provides summary performance statistics for all U.S. mutual funds, on both a raw and a risk-adjusted basis. This allows us to calculate the spread between what investors thought they earned and what they actually earned.

Sample I contains twenty-nine questionnaire responses from members of the AAII. These are tabulated in Table 1. Responses include information about fifty-seven mutual fund investments. The results indicate that even relatively sophisticated investors display a positive bias in their recollection of past fund performance. Sample I investors overestimated the actual fund return by 3.40 percentage points. When asked to estimate whether their fund beat the appropriate

[5]The benchmark was the S&P 500 for equity funds, the Merrill Lynch bond index for bond funds, and an appropriate percentage mixture between the two for the profit sharing plan.

benchmark, investors provided even higher misestimates. The average relative return exceeded the actual relative return by 5.11 percentage points, which is significant at the 95 percent level. Sample I investors indicated they were highly satisfied with their choice. On a scale of one to ten, with one being very satisfied, the average score was 2.41. This chapter of the AAII, or at least those members that reported their funds, beat the benchmarks by an average of 6.57 percent. Although the informed investor group did well, it did not do as well as its members thought.[6]

Investors did not consider themselves slow to respond to a drop in performance. The questionnaire asked how long they would continue to hold a fund if it consistently performed poorly. The average response was 2.15 years. This is a relatively short interval, given the statistical difficulties in identifying a fund as a winner or a loser. This last response is interesting in light of the hypothesis that an investor's unwillingness to remain in poorly performing funds is due to transactions costs. If investors interpret transactions costs as high, this should be reflected in how long they wait to withdraw their investment. Instead, respondents claim to have a relatively high level of impatience with poor performance. The analysis of Sample I is consistent with psychological motives for remaining in underperforming mutual funds. It suggests that the cognitive processes possibly used to justify investor inaction are based on biased beliefs about past performance. These biased beliefs are, in turn, consistent with the endowment effect as well as the cognitive dissonance effect.

Sample II is much smaller, but it allows us to distinguish between an endowment effect and a cognitive dissonance effect. Sixty questionnaires were circulated to professionals in an architecture office. Those who invested in mutual funds and/or were vested in the company profit sharing plan were invited to respond. Twelve questionnaires were returned, and these yielded response information on twenty-nine investments: eighteen mutual funds and eleven profit sharing plan shares. The analysis of these responses are reported in the second column of Table 1. Although the sample is small, the results reveal that, on average, respondents overestimated the annual return of the fund or plan by 6.22 percentage points. Respondents overestimated the amount by which the fund or plan exceeded the benchmark by 4.62 percentage points. When we separate these results according to whether the investor had a choice of manager, the optimism in total returns was higher for the mutual fund sample: 8.58 percent compared with 2.36 percent. The difference was smaller, but positive, for relative returns. In other words, when we condition upon investor choice, investors have a higher opinion of their personal choice.

---

[6]We assume they disclosed both their poor choices and their good choices. While censorship is a possibility, it does not obviously bias the results of the questionnaire.

The relative satisfaction with the manager differed according to whether the subject had control over selecting the manager. Sample II subjects were asked to rank satisfaction with the manager on a scale of one to ten, with one indicating extreme satisfaction. When subjects had a choice of manager, they expressed relative satisfaction: 4.17. When subjects had no choice, they expressed relative dissatisfaction: 5.64. Subjects appear not to have based satisfaction on performance relative to a benchmark. Since the questionnaire invited direct comparison, subjects may have been expressing dissatisfaction relative to the mutual fund alternative. Similar to Sample I, Sample II respondents considered the choice of fund relatively easy, scoring the task over the middle value of five. Sample I respondents considered the task somewhat easier, perhaps, because they enjoy investing. Sample II respondents apparently were no better at picking funds than the more sophisticated investors. Their average relative return was near zero.

III. Effects of Psychological and Economic Switching Costs

The mutual funds themselves are another source of information about investor behavior. Morningstar surveys virtually all publicly available mutual funds in the United States and reports past pricing statistics such as returns, standard deviations, and net asset values. This provides cross-sectional information that may be used to estimate how many consistently poorly performing funds exist, as well as to investigate the investor dollars in these poorly performing funds.

We begin by examining the relation between fund investment and the probability that a fund mean will exceed the riskless rate. Unfortunately, given that the Morningstar sample includes only existing funds, we have a survivorship-biased sample. Thus, correctly estimating the unconditional probability of a positive excess fund return is difficult. An approximation conditional upon survival is the *t*-statistic, or the Sharpe ratio scaled by the square root of the number of observation periods.[7] Figure I plots the fund scores against the cumulative distribution of funds, shareholders, and investor dollars for all "Growth" mutual funds with at least three years of data.

---

[7]Under the assumption that the time series of returns to the fund are independent and identically distributed and approximately normal, and that the riskless rate is constant, we form a *Z*-statistic for the probability that the fund return is greater than the riskless rate as:

$$Z = \frac{R_{\text{fund}} - R_{\text{T-bill}}}{\frac{\sigma_{\text{fund}}}{\sqrt{N}}}$$

where R_{fund} is the historical arithmetic average of the fund return over N periods, $R_{\text{T-bill}}$ is the T-bill return over the same period, and σ_{fund} is the sample standard deviation of the fund return.

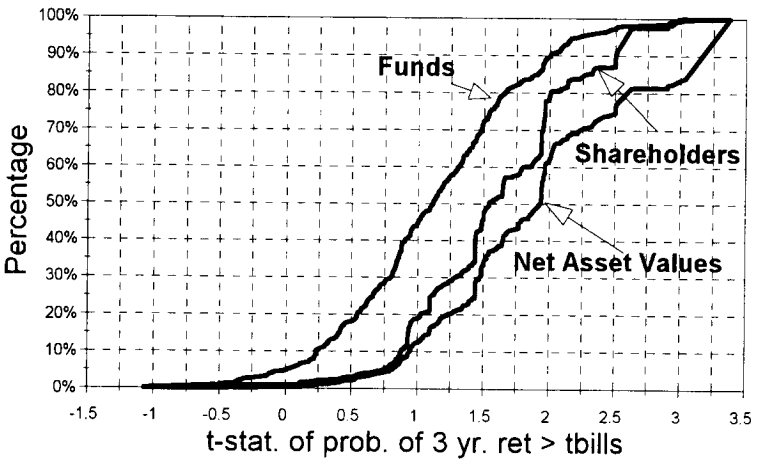

Figure I. Growth Fund Measures.

Notice that the median fund has a *t*-statistic of about 1, indicating the typical fund does not have a very high probability of exceeding Treasury bill rates. In fact, as many as 5 percent of the funds have average returns below the T-bill rate over three years. Even with a three-year history, it is hard to understand how such a fund can survive. At the other end of the scale, only about 10 percent of the funds have *t*-statistics over 2. In other words, only a small proportion of the funds have a high probability of exceeding Treasury bills. We cannot attribute this distribution pattern to irrational investors, however. The decision to shut down or merge a mutual fund is not made by the investors, but is instead made by the investment company. It is based upon the current or expected future profitability of the fund, not on any performance measure.

Figure I also shows the distribution of shareholders. Virtually no investor holds funds with negative *t*-statistics. About 20 percent of the investors hold shares in funds with *t*-statistics over 2. The median investor holds shares in a fund with a *t*-statistic of about 1.5.

The third cumulant in the figure represents the fund dollars. Nearly half of the invested dollars are in funds with *t*-statistics over 2. The dollar cumulant always lies below the shareholder cumulant. This is as expected, since we have conditioned upon past performance: investors who selected the winning funds are now wealthier, on average.

Although one question may be why most wealth is not invested in the funds with the highest probability of exceeding Treasury bills, our simplistic probability measure cannot identify it as such. The funds we observe have been subjected to performance hurdles throughout their existence. Brown, Goetzmann, Ibbotson, and Ross (1992), for instance, show that not only is there an average

bias in the fund mean induced by survivorship, but that the bias for an individual fund is positively related to its variance. A rational investor using past returns will condition upon the probability of having survived. Even when all funds have the same mean, the higher variance funds that survive appear to have a higher expected return. Thus, we cannot conclude the average fund investor is irrational for not selecting a fund with a *t*-statistic over 2. Indeed, the heavy dollar weight on funds with high *t*-statistics may be ignoring the possibility of regression toward the mean.

The cross-section of mutual fund investments suggests that while cognitive dissonance may be a factor in investment decision making, it is not a major problem in the mutual fund industry. No large pools of wealth are invested in obviously underperforming funds, even though investors appear to be overly optimistic about their past investment choices and/or investment endowments. This has implications for regulatory policy. Some policymakers suggest requiring mutual fund managers to state the risk and return characteristics of their funds in common terms so that investors can make educated choices. The existence of a large sector of the investing public holding shares of consistent underperformers should warrant a closer regulatory look. Apparently, this is not the case.

One possibility suggested by the distribution of funds versus the distribution of fund dollars is that investors in poor funds are not really investors, but are instead principals in the mutual fund companies themselves. Companies may be "propping up" losing funds with small amounts of capital to keep them alive. The motivation for such support is clear. If the ugly duckling ever turns into a swan over the long term, it will have a long track record that can be promoted. The fund can be revived at some future date, not as a start-up, but as a long-term performer. Thus, perhaps investment companies as "investors" in poor performers are motivated by the long-term option value of the fund. If this is the case, we would expect to find a large number of funds with only a few investor accounts. Figure II plots the cumulative distribution of growth fund size measured by number of investors.

Nearly 10 percent of the growth funds in the sample have 100 or fewer investors, despite economies of scale in the industry. Consequently, even the small number of funds we find with negative *t*-statistics may actually be maintained by the mutual fund company, rather than by disinterested investors. In fact, the cross-sectional purchase and redemption of mutual fund shares shed some light on this agency issue.

IV. Differential Response to Past Performance

Several researchers point out that investor dollars flow into funds that have superior track records (e.g., Ippolito, (1992), Sirri and Tufano (1992),

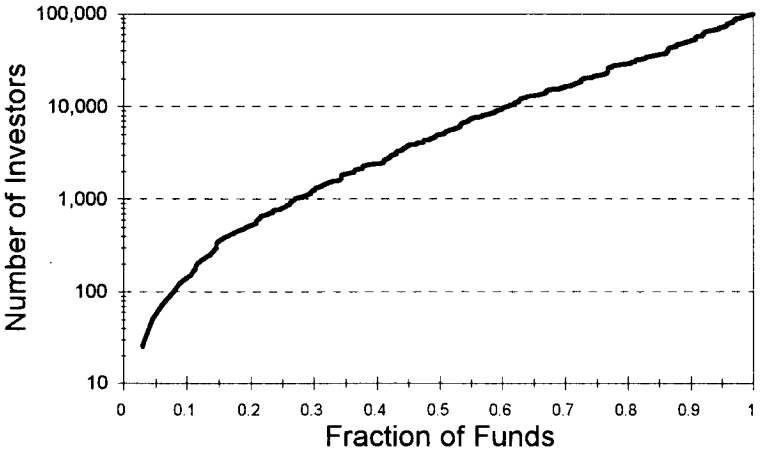

Figure II. Distribution of Fund Size: Growth Funds with Three Years of Data.

Patel, Hendricks, and Zeckhauser (1990)). This suggests rational probability assessment: investors appear to act on the possibly rational expectation that past performance predicts future performance. This rational response apparently differs depending on whether the fund did well or poorly. Sirri and Tufano (1992), for instance, find that money flows slowly out of losers and quickly into winners. This differential response suggests economic or psychological switching costs. Once investors select a fund, they tend not to chase recent winners. Chevalier and Elison (1995) suggest that their performance flow response is nonlinear and that the functional form may have implications for incentives of managers to take risks. Roston (1995) shows that differential performance flow response appears to be conditioned upon the age of the fund. When the find is older, differential investor cash flows are less sensitive.

One problem with identifying differential response is the survivorship factor. Good performers tend to survive, while poor performers are often merged into other mutual funds. Brown, Goetzmann, Ibbotson, and Ross (1992) show that survivorship bias is potentially severe in repeat winner studies. Hendricks, Patel, and Zeckhauser (1997) show how a pattern of censorship can induce a J-shaped response. Survivorship is even more of a problem in cross-sectional analysis that predicts future cash flows based on current performance. This is because survivorship is directly related to the value of the dependent variable. We expect a poor performer that survives to be a fund that had relatively low cash outflows. Thus, absent a correction for survivorship bias, we expect to observe a differential response of cash flows to performance.

To address this survivorship bias possibility, we use a database that contains information about defunct funds. The database is from the Weisenberger

Investment Companies Service *Investment Companies* annual reports. It provides information about virtually all publicly offered open-end mutual funds on an annual basis. Data were collected by hand from the *Mutual Funds Panorama* section for 1976 through 1988, for all firms listed as common stock funds or for specialty funds that invested in common stock (typically sector funds). We selected funds denoted as CS, or common stock, and funds denoted as SPEC, or specialty funds, if their primary objective was investment in equities. We included specialty funds because this category includes sector funds.

We study 1976–88 for comparison with other studies. For each fund we recorded the name as it appeared that year, the year of origin, the fund objective, the net asset value at the end of the year, the net asset value per share at the beginning of the period, the twelve-month percentage change in net asset value per share adjusted for capital gains distributions, the income return, the capital gains distributions, and the expense ratio. We calculated the total return including capital appreciation, income, and capital gains distributions. In some cases, one or more of these data were not reported, and this prevented total return calculations. Occasionally, this occurred because the fund had begun the same year, so the net asset value per share for the beginning of the year was not available. We calculated net fund inflows, "new money," as the total fund value at time t less the total fund value at time $t-1$ times the total fund return over the period, adjusted for capital gains. This makes the conservative assumption that fund investors automatically reinvest dividends. To track funds through time, we assigned a unique number to each. Footnotes at the end of the *Panorama* section indicate merged funds and name changes of funds. When one fund was merged into another, the "acquired" fund was deemed to have disappeared, while the "acquiring" fund was deemed to have continued.[8] This allows us to identify funds that disappeared and to make the conservative assumption that a defunct fund had a −100 percent outflow in the year of its disappearance. Thus, all missing funds are accounted for in the cross-section, and our assumption of a −100 percent outflow biases the coefficient of the regression of new money inflows on last year's returns toward a positive slope.

Our test of the differential response follows Sirri and Tufano (1992). We regress this year's new money on last year's fund return, including year dummy variables to capture annual differences in mean returns and flows. We test for differential response by dividing fund returns into quartiles and then performing a Chow test of equality across the four coefficients. Table 2 reports the results. Each year except for 1978 is assigned a dummy variable. The response to the other explanatory variables is constrained to be the same across years. The

[8]For additional details about the database, see Brown and Goetzmann (1995).

TABLE 2. Regression of Dollar-valued New Money on Past Year's New Money and Past Year's Return.

Variable	Coefficient	Standard Error
1979 dummy[a]	-16.07	6.54
1980 dummy	-14.37	7.92
1981 dummy	-9.93	8.58
1982 dummy	-1.72	6.19
1983 dummy	14.87	7.39
1984 dummy	-8.02	7.01
1985 dummy	1.25	5.77
1986 dummy	3.92	6.87
1987 dummy	-4.26	5.41
1988 dummy	-36.42	4.50
1st Q Ret (t-1)	-23.11	31.62
2nd Q Ret (t-1)	-49.78	31.26
3rd Q Ret (t-1)	32.79	24.00
4th Q Ret (t-1)	54.87	11.91
Constrained Ret (t-1)	55.63	11.79
F-statistic (23, 3948) = 4.97	p-value = .0001	

Notes: In the unconstrained regression, the past year's return is divided into quartiles

$$NM_t = \alpha + \beta_1 NM_{t-1} + \beta_2 R_t + \beta_3 R_{t-1} I_1 + \beta_4 R_{t-1} I_2 + \beta_5 R_{t-1} I_3 + \beta_6 R_{t-1} I_4 + \varepsilon_t$$

where I_k equals one if R_t is in the k^{th} quartile, and zero otherwise. The constrained regression requires the response to be the same for each quartile of the past year's returns. The F-statistic for the Chow test is provided in the final row. The p-value indicates the probability that the quartile coefficients are equal to each other.

[a]All dummies are defined as the difference between the current year and 1978.

combined regression suggests the relation between new money and last year's return is significantly positive.

Table 2 also reports an unconstrained regression for which the returns for each year are divided into quartiles. The response is significant only for the best quartile. The Chow test of the equality of these coefficients rejects the null hypothesis with a probability of .99. The top quartile of funds exhibits a response pattern different from the other quartiles. More interesting, the bottom quartile is not significantly different from the second and third quartiles.[9] We cannot reject the hypothesis that the response is the same for all but the top performing funds. In other words, even controlling for survivorship, we document a differential response conditional upon past performance. We see nothing special about the coefficient on poor performance, however. In fact, although the Weisenberger

[9]This is true regardless of whether we make the assumption of -100 percent cash outflows in the following year for defunct funds. Omitting defunct funds from the sample changes the results very little.

dataset contains defunct funds as well as surviving funds, the J-shape identified in Table 2 corresponds closely to the pattern identified by Hendricks, Patel, and Zeckhauser (1997).[10]

These results suggest that, although the market rewards the top performers each year, it does little to discipline poor performers. This contradicts the evidence in Brown and Goetzmann (1995) that fund disappearance is highly correlated with past performance. Their probit analysis on the same Weisenberger dataset indicates that funds disappear (typically by merger into another fund) when the previous year's performance is relatively low. The contradiction is more apparent than real, however. Brown and Goetzmann find that new money flows do not explain fund disappearance when past returns are used as an explanatory variable. Only when past returns are omitted from the specification do fund flows become significant, probably because they are an instrument for earlier differential returns. Apparently, the fund flow is not the principal factor driving fund merger; instead, the return history is the principal factor. This supports the hypothesis that the manager or investment company, not the customer, decides whether to maintain a fund. In general, lower past returns reduce the option value of maintaining a fund; however, redemptions by fund investors have little effect on the ex-ante probability the fund will look like a winner. One caveat to this rule is that conditioning upon past very low returns implicitly selects on high variance, as does conditioning on very high returns. As the variance increases, the option value in keeping the fund open increases. Small flows into very poor performers to keep the funds open is consistent with a strategy of "propping up" poorly performing funds to maintain their option value.

V. Conclusion

Festinger's (1957) cognitive dissonance hypothesis asserts that people tend to revise their beliefs to reduce apparent logical contradictions. Market research indicates the anxiety associated with a major purchase decision, such as choosing an automobile, tends to induce selective consumer perception of information about the efficacy of their choice. We find some evidence that the same is true for the mutual fund purchase decision. Our questionnaire responses from two groups of investors about their personal holdings and mutual fund choice suggest that even well-informed investors tend to bias their perceptions about past performance. This positive bias does not preclude the possibility that investors confront genuine economic costs that lead to their inertia. However, where they are slow to respond

[10]The pattern of standard errors is also consistent with the hypothesis that the J-shape is an artifact of survivorship. Brown, Goetzmann, Ibbotson, and Ross (1997) find that the standard errors increase substantially for the right-hand tail of the conditional performance test.

to past poor performance, they may justify their behavior through biased beliefs about performance.

Our study of the cross-sectional distribution of fund size, using a simple probability measure of performance that exceeds Treasury bill returns, suggests the number of investors in poor funds, as well as the total wealth invested in poor funds, is small. The existence of many irrational investors in the mutual fund industry is contradicted by our cross-sectional evidence. Apparently, P. T. Barnum's wisdom does not extend to the mutual fund industry. The proportion of uninformed mutual fund managers seems to be greater than the proportion of naive mutual fund investors.

Unlike previous researchers, such as Ippolito (1992) and Sirri and Tufano (1992), we are not concerned that the market fails to discipline poor performers. In fact, our cross-sectional study indicates that few investors hold shares in lagging funds. The question remains whether the positive flow of funds into the biggest losers represents outside money or inside support. If it is outside money, we need models to explain such odd behavior, if only by a small minority. The finance literature lacks useful conceptual models that describe bounded rationality and provide testable implications. The natural place to find such models is the social psychology literature, where behavioral patterns, rational or not, are developed and empirically tested.

The importance of identifying cognitive dissonance in mutual fund investor behavior extends beyond the desire for models of bounded investor rationality to reporting requirements and mutual fund industry regulation. If new investors focus on past performance rankings, the optimal mutual fund company strategy is to increase the number of funds under management, increase the volatility of individual funds, and decrease the cross-fund correlation. To the extent that the principal benefit of equity mutual funds is to provide low-cost diversification, this strategy does little to benefit mutual fund investors. The inertia caused by psychological or economic factors tempts mutual fund companies to slowly raise fees on poor performers. Fortunately, there is little evidence that this particular strategy is pursued.

The cognitive dissonance of mutual fund investors provides some positive information for mutual fund companies on strategic use of information. Given a high level of cognitive dissonance surrounding the choice of a mutual fund, the principal value of advertising by the fund is in confirming that its current investors made a wise investment choice. Although advertisements may not be able to influence new investor decisions, they may be able to help funds retain their current customers.

References

Akerlof, G. and W. Dickens, 1982, The economic consequences of cognitive dissonance, *American Economic Review* 81, 1–13.

Brehm, J., 1956, Post-decision changes in the desirability of alternatives, *Journal of Abnormal and Social Psychology* 52, 384–89.

Britt, S. H., 1978, *Psychological Principles of Marketing and Consumer Behavior* (Lexington Books, Lexington, MA).

Brown, S. and W. N. Goetzmann, 1995, Performance persistence, *Journal of Finance* 50, 679–98.

Brown, S., W. N. Goetzmann, R. Ibbotson, and S. Ross, 1992, Survivorship bias in performance studies, *Review of Financial Studies* 5, 553–80.

———, 1997, Rejoinder: The J-shape of performance given survivorship bias, *Review of Economics and Statistics*, Forthcoming.

Chevalier, J. and G. Elison, 1995, Risk-taking by mutual funds as a response to incentives, Working paper, University of Chicago, Graduate School of Business.

De Bondt, W. F. M. and R. Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–805.

DeLong, J., A. Shleifer, L. Summers, and R. Waldmann, 1990, Noise trade risk in financial markets, *Journal of Political Economy* 98, 703–38.

Erlich, D., P. Guttman, V. Schonbach, and J. Mills, 1957, Postdecision exposure to relevant information, *Journal of Abnormal and Social Psychology* 54, 98–102.

Festinger, L., 1957, *A Theory of Cognitive Dissonance* (Stanford University Press, Stanford, CA).

Harliss, D. and S. Peterson, 1994, The persistence of poorly performing mutual funds, Working paper, Virginia Commonwealth University, Department of Economics.

Hendricks, D., J. Patel, and R. Zeckhauser, 1997, The J-shaped pattern of performance persistence given survivorship bias, *Review of Economics and Statistics*, Forthcoming.

Ippolito, R., 1992, Consumer reaction to measures of poor quality: Evidence from the mutual fund industry, *Journal of Law and Economics* 35, 45–70.

Patel, J., D. Hendricks, and R. Zeckhauser, 1990, Investment flows and performance: Evidence from mutual funds, cross border investments and new issues, Working paper, Harvard University.

Roston, M., 1995, Mutual fund managers and life cycle risk: An empirical investigation, Working paper, University of Chicago, Economics Department.

Samuelson, W. and R. Zeckhauser, 1988, Status quo bias in decision-making, *Journal of Risk and Uncertainty* 1, 7–59.

Shiller, R., 1988, Stock prices and social dynamics, *Brookings Papers on Economic Activity*, 457–98.

Sirri, E. and P. Tufano, 1992, The demand for mutual fund services by individual investors, Working paper, Harvard Business School.

Thaler, R., D. Kahneman, and J. Knetsch, 1992, The endowment effect, loss aversion and status quo bias, in R. Thaler, ed., *The Winner's Curse: Paradoxes and Anomalies of Economic Life* (The Free Press, New York).