The Search for the Holy Grail

Remarks by
John C. Bogle, Founder and Senior Chairman
The Vanguard Group
Before
"The Q Group"
(The Institute for Quantitative Research in Finance)
Palm Beach, FL
March 29, 1998

In its hell-bent mission to transmogrify the finest investment for long-term investing ever created into a vehicle for intermediate-term—and even short-term—speculation, the mutual fund industry has implicitly conceded what knowledgeable observers have long known to be the ultimate reality of finance: the central task of investing is to realize the highest possible portion of the return provided by the class of financial assets in which you invest—*recognizing that that portion will be less than 100%.*

Simply put, then, the first reality is that investors, as a group, do not, cannot, and will not "beat the market." And the second reality is that the overwhelming odds are against any particular mutual fund doing so over an investment lifetime.

Stated another way, long-term investors must accept the fact that the fund or funds they select, irrespective of their past performance, will gradually regress to the mean over time.[*] But *not* to the market mean. To the market mean, *reduced by the costs the fund incurs*—advisory fees, operating expenses, marketing costs (in all, the "expense ratio"), plus the cost of buying and selling portfolio securities ("transaction costs").[1] And in the world of mutual funds, these costs are extremely large. The annual expense ratio of the median equity fund is now 1.6 percent, and rising. Transaction costs are difficult to quantify with precision, but at the high portfolio turnover rates of the past decade (80%-plus), an estimate of 0.5 percent to 1.0 percent hardly seems excessive. "All-in" costs, then, can be conservatively estimated at upwards of 2.0 percent per year.

So the search for the holy grail of market-beating long-term returns is every bit as frustrating to the fund investor of the 20th Century—and will surely be the same, and perhaps even more so, in the 21st Century—as the search for the Holy Grail of the Last Supper was to the Knights of King Arthur's Round Table in the 15th century. And it is for that reason that I use this search as my theme this evening.

The Equity Fund Record

Let me first speak of the records of equity mutual funds as a group in the recent past. And the recent past is particularly relevant, since it is in that era that equity fund assets have become the largest pool of assets (nearly 25%) in the stock market, and that mutual fund expense ratios and portfolio

[*] As I understand it, "regression to the mean" is a phrase of quantitative art, suggesting that above-market returns will be followed by below-market returns in a cyclical fashion. However, in my discussion of mutual fund performance, I define it by its precise words: that the before-cost returns of both above-average and below-average funds demonstrate a powerful tendency to regress to the mean of their style categories and to the market mean over time. See page 13 for source.

turnover activity have risen to the highest levels in history. Unless they are to decline materially, the future is almost certain to resemble the recent pattern. The fact is that in the past 15 years—coinciding almost precisely with the great bull market—equity fund annual returns (unadjusted for survivor bias) have averaged 14.3 percent, providing 86 percent of the return of the total stock market (Wilshire 5000 Equity Index) of 16.7 percent. (Note that that gap of 2.4 percent is somewhat larger than my earlier 2.0 percent estimate of fund costs, in part because of the drag of cash reserves).

These undeniable facts about fund returns and costs and the relevance of past records, of course, have led to the boom in index funds today. "If you can't beat the market—to say nothing of *meet* the market—why not join it?" And an all-market index fund, operated at a cost of 0.2 percent—one-tenth of the industry norm—would have provided an annual return of 16.5 percent—or nearly 99 percent of the market return during the same period. In fact, the comparison of 99 percent of *annual* market return for the index fund and 86 percent for the managed fund conceals a much larger gap than merely 13 points.

In fact, the *terminal* value of an initial investment of $10,000 in the index fund would have been worth $98,800 (97 percent of the market result), while the terminal value of the same investment in the traditionally-managed active fund would have been $74,200, only 73 percent of the market. The increase in the performance gap from 13 percentage points to 24 percentage points can best be described, not as the conventionally-rosy "magic of compounding," but as the unconventionally-ominous "tyranny of compounding."

I would also add that, in all probability, "you ain't seen nuthin' yet." For in a market that gives a more modest account of itself—(which, we are sure to see, assuming only that this bull market tree doesn't grow to the sky)—the performance gap gets larger. For example, let's hold costs constant and take stock market returns down to, say, 8 percent annually over fifteen years. The result would be a return of 7.8 percent for the index fund (97 percent of the market return) and but 6 percent for the managed equity fund (75 percent). And at the end of the period, the $10,000 investment would be valued at $30,800 in the index fund, a gain of $20,800 (96 percent of the market return), versus a gain of $14,000 for the costly mutual fund (65 percent). "Tyranny," as it were, has increased the performance gap from 22 percentage points to 31 percentage points. It is not a very happy prospect for investors, for it means simply that the mutual funds, accepted for a half-century as the holy grail of high performance, have failed the test of time.

Enter the Index Fund

It was looking at historical performance relationships similar to these (although less unfavorable to the funds) almost 25 years ago that encouraged me to recommended to the Board of Directors of our newly-created (1974) firm that its first action should be to embark on a newly-created concept: In 1975, we formed the first index mutual fund, modeled on the Standard and Poor's 500 Composite Stock Price Index. Sad to say, I didn't invent the concept—I'd give the lion's share of the credit for that to William E. Fouse and his colleagues at Wells Fargo Bank—but I was a believer, even a missionary, in the concept, and was confident that it could—against all odds—become a reality in the world of mutual funds.

In the early years, there was lots of joking about the first index mutual fund. "Index funds are un-American," "Bogle's Folly," and "seeking mediocrity," were hardly the worst things that were said. But, after a shaky start (mediocre relative returns during its first five years), and minuscule assets (just $90 million in 1980), our 500 index fund turned on its jets, and became both an artistic success (even better 15-year returns, given its large cap bias, than those of the total stock market) and a commercial success (now the second largest fund—and, with a companion institutional portfolio, the largest—fund in today's burgeoning mutual fund industry).

And, as the performance data indicated earlier show, it is no longer a competitive joke. Although uncopied—even shunned—for a full decade, it has now been joined by some 120 competitive index funds: few have been formed by missionaries (or converts, and "there's no one more religious than a convert"), but most by opportunistic no-load firms, "eating crow" and dragged—kicking and screaming, and with considerable misgiving—into the fray by the institutional 401-k savings plan market. While 42 have reasonable expense ratios (0.25 percent or less), most of those are the result of temporary fee waivers; 20 have unacceptable ratios of 0.75 percent or more; and an appalling 25 even charge sales loads or 12b-1 fees. (They don't seem to realize that minimal cost accounts for virtually *all* of the index advantage.) About two-thirds of the U.S. equity funds are targeted against the S&P 500 Index.

The Industry's Response to Indexing

The indexing concept is much broader than a single index fund, however. First, true indexing works best against the total stock market, and the all-market fund is only at the beginning of its acceptance. Second, for investors who, for one reason or another, seek to earn higher returns in specific broad market sectors, funds modeled on growth indexes and value indexes, as well as small cap and mid cap indexes, are now available. It is only a matter of time until someone has the good sense to offer index funds that match each of the nine Morningstar style/market-cap boxes. The record is crystal clear that index funds would have produced highly effective risk-adjusted returns in each. Index funds are in the incipient stage, too, in the bond market and the international market. As they must be, by definition, they are equally effective.

So, indexing is threatening to become the new holy grail for fund investors, and it should be. Yet, if investors increasingly see the merit of indexing strategies as a means to outpace the long term returns of individual funds—a point that the data abundantly demonstrate—how, finally, can sponsors of traditional funds compete? Material cuts in the fees they charge seem out of the question, for such actions would slash their profits unmercifully. Reducing portfolio transaction costs is also unlikely. It would result in unacceptably radical changes in today's silly, but chic, high turnover investment policies. Hence the answer: *create a new holy grail!*

The idea now seems to be to offer investors a new mantra: actively manage your own fund portfolio. Don't own a fund for the long-term. Dash to the transaction cost-free (or so it is incorrectly alleged) fund supermarkets. Exercise your freedom of choice. Treat funds as stocks. Switch and (or so it is suggested) get rich. Pay attention to the drumbeat of advertising for those handsome past performers (the only funds that are advertised) whose names you see on your television screen. This grotesque transfiguration of the long-term nature of fund investing is what is now coming to pass.

But does it work? Are there methods for selecting and swapping mutual funds that might result in superior returns? Are there strategies that have worked in the past? Well, on the first point, the ability of the computer to spit out endless performance comparisons, multiple regressions, and complex formulas has enabled academia to test, well, *everything*. There may be a lot of "data mining" involved in what is duly recorded on this subject in *The Journal of Finance*, *The Journal of Portfolio Management*, the *American Economic Review*, and the like. But persistent past factors that may (or may not) be valuable in selecting funds that will provide superior future returns is on the agenda of scores of respected academics.

Selecting Winning Funds—Academic Evidence

What have they found? Let's examine a few examples. We'll start with the mutual fund king of the academic profession, Professor William F. Sharpe of Stanford and his analysis of large seasoned U.S. mutual funds (1995).[2] He has carefully examined the 10-year records of the 100 largest funds (measured each year), accounting for more than 40% of the assets of all funds, and calculated their returns relative to the returns of comparably-weighted market sector indexes, including a U.S. Treasury bill component (therefore eliminating the negative effect of the persistent performance lag created by fund cash positions). He properly acknowledges that the cost-advantage ascribable to large funds probably provides superior returns to this group, but they nonetheless fell short of the multi-index return by 0.64% per year on average over the past decade. While the shortfall could not be deemed significantly different from zero, the data do undermine the belief that a typical active equity fund will outperform the passive alternative. The data would be even more damning if Sharpe had included sales charges.

Sharpe then examines the value that some of these managers have added by their apparent stock selection ability over prior periods, and relates that to their future success. He investigates common measures for judging funds—size, past performance, and Sharpe ratios. His best evidence of performance persistence showed that, using the results for the previous 12 months, if you held the top 25 funds in his study—the first quartile—over the subsequent five and ten year periods, you would have added an annual return of 0.8 percent relative to the index return; in the bottom quartile, you would have lost 0.5 percent per year over the five years and lost 1.3 percent annually over the ten years. Even disregarding the extra taxes incurred by switching funds regularly, this seems to form a pretty shaky basis for an investment strategy.

Do winners repeat, then? Sharpe summarizes his results with these conclusions: "If the past ten years are indicative of the next ten, one might answer in the affirmative." (Although, I would note, the positive margin is modest to a fault.) However, perhaps a neutral ("not proven") position is appropriate, for "the evidence is far from conclusive, statistically or economically."

Mark Carhart (1997) of the University of Southern California has also tackled the issue of persistence in fund performance, evaluating 1,892 diversified equity funds over 16,109 fund years (amazing!) from 1962-1993.[3] First, he finds that "common factors in stock returns (high beta vs. low beta, value vs. growth, large cap vs. small cap) and investment expenses almost completely explain persistence in equity fund returns." Properly adjusting for the customary failure to consider the effect of the sub-average returns of funds that have gone out of existence, he confirms Professor Burton Malkiel's conclusion that survivor bias has enhanced past annual returns reported for funds by about 1.5% per year. Malkiel found some evidence of persistence during the 1970s, but none during the 1980s.[4]

Looking at past one-year returns relative to those of the subsequent year, Carhart concludes, among other things, that relatively few funds stay in their initial decile ranking, although funds in the top and bottom deciles maintain their rankings more frequently than expected—17 percent in decile 1 and 46 percent in decile 10, doubtless because many low decile funds tend to be trapped there by high costs. (10 percent in each decile would be the random expectation.) In his conclusion, he warns: "while the popular press will no doubt continue to glamorize the best-performing mutual fund managers, the mundane explanations of strategy and investment cost account for almost all of the important predictability of mutual fund returns."

In a third study, William Goetzmann and Roger Ibbotson (1995) tested the repeat winner hypothesis over two-year, one-year and monthly intervals from 1975-1987.[5] For all periods, they ranked equity mutual funds in terms of both raw returns and risk-adjusted returns (alpha), and split them into two categories: "winners" (top 50%), and "losers" (bottom 50%). Their analysis indicated that investing in winners slightly increased the chance of outperforming the return of the all-fund average in the subsequent period.

By way of example, their study of growth mutual funds over two-year periods revealed that past top performers had a 60% chance of being winners over the subsequent two years. Therefore, one might conclude that the chance of a fund being better than average in all four subsequent two-year periods would have been about one-in-ten. Exceeding the average fund return in each two-year period, in short, was hardly an odds-on wager—and the odds would get far worse if sales charges and taxes had been taken into account.

To complicate matters, since funds as a group lag the market over the long term, even consistent winners might be losers relative to the market index. Indeed, Goetzman and Ibbotson concede that picking winners—even when defined as those funds in the top quartile in performance—may not be enough to beat the market. They conclude, "while the 'repeat-winner' pattern may not be a guide to beating the market, it does appear to be a guide to beating the pack over the long-term." In the face of index fund competition, then, to what avail is a strategy that relies on the tenuous persistence of fund performance?

Selecting Winning Funds—Real World Evidence

With that background from the theoretical world of academia, let's look at fund selection in the real world of industry practice. Let's consider: (1) the *actual* records of the funds that *did* beat the market over the past 15 years; (2) the *actual* records of investment advisers that recommend mutual fund portfolios; and (3) the *actual* records of "funds-of-funds" that invest solely in mutual funds.

Over the past 15 years, as I noted at the outset, the average managed mutual fund has provided only 86% of the annual return of an unmanaged total stock market index. But the fact is that of the 272 general equity funds that survived that period (this was a far smaller industry in 1982!), there were 41 that succeeded in outpacing the index return of 16.7%. The investor's odds of picking a winner, therefore, were about one in seven. But only nine of those 41 (3 percent of the total) did so by a margin of 1.5 percentage points. If we assume that the funds' annual tracking error relative to the index was a fairly modest 3%, then a return of 1.5% in excess of the index return would represent statistically significant outperformance. In fact, only one of the funds leaped the hurdle of statistical significance (although, in fairness, two others came close).

A bit of micro-analysis shows these nine funds to be a rather motley group. Three carved out their entire long-term margins in the early years, when their assets were small, and have been mediocre performers for years. That leaves six legitimate top performers. Interestingly, and importantly, all six had the same portfolio managers throughout all or most of the period (though their average age is now 60); three closed to new cash flow before assets reached $1 billion. In any event, suffice it to say that the nine winners could not have been easy to identify in advance, simply because the evidence is clear that they were not identified in advance by very many investors. (Their aggregate 1982 assets totaled but $1.4 billion, only 3% of total equity fund assets.) Today, would you select one of the three (what I might have the temerity to call) legitimate champions which remain open to investors? You can decide for yourselves.

Next, let's examine the public records of advisers who recommend mutual funds. For the past five years, *The New York Times* has published, regularly each quarter, the records of the equity fund portfolios selected and supervised by five respected expert advisers. During this period, not one of them has come close to matching the record of the market index fund chosen by the *Times* as the appropriate comparative standard. The average adviser's annual return of 13.7% provided 70%, the index fund 99% of the return of the market. In fairness, some of these advisers chose portfolios that were slightly less risky (less volatile) than the market. Still, providing 70% of the market's return during a period when even the *average* fund provided 80% suggests that selecting winning funds is hardly bereft of challenges.

Another, longer-run evaluation of the success of advisors in selecting fund portfolios is the *Hulbert Financial Digest*. It reports that, of the 59 advisory letters that have been tracked for a full decade, the portfolio of the average adviser provided a return of +9.5%, or the same 70% of the return earned by the *Times* group. Only eight letters outpaced the market with their recommendation, interestingly, the same one-in-seven chance for the funds themselves over the past 15 years. While many of these advisers recommended (for better or worse, during this bull market era) portfolios far more conservative than the stock market itself, on average they carried a market risk. Measured by their Sharpe Ratios, their risk-adjusted return amounted to just 58% of the markets. So, the accumulated evidence regarding the ability of the experts to select winning funds remains negative.

The third real world test we might consider is a review of the actual records of "funds-of-funds"—mutual funds that select other mutual funds for their portfolios. And this record, as it happens, is the most deplorable of the lot. For these funds not only lag the market—after all, six of every seven funds, as we know now, have done *that*—but they seriously lag even the style categories of the funds in which they invest, in part because of the extra layer of costs they almost universally add. For example, of the 11 funds-of-funds investing in large blend (value and growth) funds last year, five ranked in the 98-99-100th percentile (one was dead last), and five in about the 90th, with the champions, as it were, ranking in the 67th percentile. In all, the 74 funds with full year records in 1997 achieved about what might have been expected by random selection of funds and an added layer of costs: an average 75th percentile standing among their fund style peers.

Here, I have had to rely primarily on one year data, short as it is, because so few funds-of-funds have been around very long. (Only 27 have a three-year-return; they averaged a 72nd percentile ranking). A decade ago, there were but nine. There the record is a bit better: a 60th percentile average. But excluding the single fund that did *not* add a layer of extra expenses (23rd percentile), the average quickly drops down to the 66th percentile, a neighborhood, if hardly posh, that is surely familiar, based on my earlier numbers. Moreover, remember that this is the 66th percentile among a collection of funds most of which have themselves failed to outpace the market.

To make matters worse, I believe that it would be optimistic to expect that a 66th percentile rank to be sustained. Those funds-of-funds that bear an extra layer of fees add about 1.5 percent to the expenses of the underlying funds, which themselves have expense ratios of a comparable magnitude—bringing total annual costs to the three percent range. Such an extra deduction—assuming even their managers, on average, pick average funds—should produce about a 90th percentile rank. In any event, it would take naiveté to undreamed of heights to believe that such a heavily-loaded package of funds could outpace appropriate market indexes.

No Holy Grail Here—Academic or Pragmatic

So, whether we consider academic studies, which I can only presume tried many tests of predicting future returns that were found wanting and were never published, or the pragmatic and unforgiving actual results of the funds with the best long-term records, the funds selected by advisory news-letters, or the funds selected by funds-of-funds, the odds of selecting future top performers based on past statistics are poor. The chances, then, of successfully locating the holy grail—identifying *in advance* future superior performers—seem dismal to a fault.

Before this day of sophisticated return attribution on a factor basis and performance evaluation on a relative basis, equity mutual funds with the best sustained long-term records represented the holy grail of performance. In recent years, the acceptance of this thesis has been endangered by the considerable artistic success—the clear performance superiority—of the index fund—and, in very recent years, by its so-far modest commercial success. (Index funds now account for about 6% of the assets of all equity funds.) The industry has, at least implicitly, mounted a counterattack. If only the rare fund can hope to go toe-to-toe with the market on a long-term basis, let's abandon the conventional buy-and-hold fund strategy and switch opportunistically among funds to gain an edge. Leave aside what I hope is the self-evident fact that, while *some* investors may succeed in doing so, the eternal thesis under which investors *as a group* must underperform the market by the amount of their costs remains, as it must, firmly in place.

The Quantitative Question

That said, being well aware that you here tonight represent a whole new, non-traditional approach to investing, I owe you a comment about whether your quantitative strategies are likely to enable you to carve out an edge—a positive alpha, as it were—at the expense of money managers following traditional investment research and security analysis techniques. Here, I bring you both bad news and good news.

The bad news first: I do not believe that, over time, quantitative strategies—as a group, although it is a group composed of an incredibly wide range of multifaceted, independent, imaginative, unusually brilliant (I mean it!) practitioners—will outperform traditional strategies. In the United States, markets are becoming more efficient, almost daily it seems, an efficiency that now is well on its way to permeating the small cap segment as clearly as it has permeated the large cap segment. And international markets seem to make up with excessive transaction costs what they lack in relative efficiency. And, of course, the eternal bugaboo: to the extent that particular quantitative strategies *do* provide sustained superiority for a time, they will attract massive dollars that ultimately vitiate success, and make the pervasive likelihood of mean regression (a tendency that affords many of you the opportunity to earn a living) into a virtual certainty.

Now to the good news: even if your quantitative *strategies* are unlikely to defy the odds, your funds and private *clients* have a good opportunity to outpace traditional approaches. That, it seems to me, is because your investment costs are but a fraction of those that are encountered by traditional firms. You don't need offices populated with high-powered security analysts, research directors, strategy specialists, or policy committees, nor a compliance apparatus to assure the avoidance of insider information. But you must be wise, maintaining asset size that comports with your trading strategies, and holding your fees to reasonable levels (two disciplines almost totally shunned by the giant mutual fund industry). If you share with your clients some of the economies of scale you will enjoy as you grow in a controlled fashion, earning merely competitive gross returns will enable you to provide at least marginally superior net returns. Compounded over time, that edge will result in immensely superior capital accumulations for your clients. But if there is a holy grail here, it is apt to be found only at the margin.

So, I'm optimistic about "quants," and just as much a believer as I was twelve years ago, when we formed Vanguard Quantitative Portfolios, the first mutual fund managed on the basis of quantitative, disciplined, computer-directed investment techniques. In 1986, we were convinced that we could parlay our pioneering foray into the index fund arena to a higher level. Even though our original brand name has been sort of changed into a bland name—Vanguard Growth and Income Portfolio—that hardly suggests its still distinctive mission, I can report that "Quant" (as we still call it at the office) is alive and well.

The term "mission completed" is rarely applied to a mutual fund—nor should it ever be. But, so far, at least, "mission successful" is a fair enough description. The fund has succeeded in outpacing its target index, the Standard & Poor's 500. If the margin is skinny, it is nonetheless there: 0.06 percent per year, net of all transaction costs and an expense ratio averaging 0.50 percent. That may seem but a trivial accomplishment, and perhaps it is, although it has outpaced our Index 500 Fund by 0.3 percent. But, in healthy measure because of the hefty costs in this over-priced industry, Quant's annual net return of 17.5 percent in the 12 years since its inception ranks it in the 3rd percentile among its peers. Its fourth place rank among all 121 growth and income mutual funds (that's all that have existed for the full period) is a tribute to Franklin Portfolio Associates of Boston, the external manager we initially selected to run Quant, and that manages it to this day.

But our conviction—so far, rarely shared by others in this industry—about the value of quantitative strategies is also manifested in the fact, quietly and without fanfare, that our own Core Management Group has been managing portions of three of our existing funds, one of them for six years now, one for four, and a fourth fund (in its entirety) for nearly three years. If our success has been uneven, I think it is nonetheless fair also to describe our own foray into quantitative strategies at this early stage as "mission successful." And I look with enthusiasm for quantitative strategies to become an increasingly important weapon in our arsenal. Nonetheless, it would require more imagination than I can muster to envision it ever having the impact on our firm of our pioneering foray into index funds in 1974. Today we manage 26 index and index-oriented funds, constituting one-third of our $350 billion asset base.

Let me sum up briefly by returning to my initial theme. The traditional investor strategy of holding mutual funds for the long term has not provided the holy grail of market-superior returns—not by a long shot. The index strategy, by definition, must provide less-than-market returns—but only by a slight margin. That is what the real search for the holy grail is all about—achieving as close to 100% participation in market returns with a diversified investment portfolio as is possible. After all, mutual fund managers are mere mortals, operating in highly efficient markets. Maximum participation is about as good as it is likely to get for the long-term investor. The bogus fund-switching strategy in vogue today, implicitly designed to counter the index strategy by misleading investors into thinking that they can individually somehow out-fox the market is certain (I choose that word carefully) to be a loser's game. And funds of funds attempting to emulate the result of a long term buy-and-hold index fund strategy, to say nothing of a buy-and-hold mutual fund strategy (which fails largely in ratio to its 2 percent annual cost) by *adding* a cost of 1.5 percent or more, simply defies all reason. You've just seen the evidence of the past on this point.

So, that leaves the unconventional quantitative strategies you have espoused with reasonable success so far: If you do not learn from the mistakes committed by your traditional peers—most notably carrying excessive costs and growing fat with unbridled asset growth (and unbridled profit growth to the adviser), yours may well be "mission impossible," too. The holy grail—or at least *a* holy grail—is there for you to find. If only you take one small step for a manager—discipline—you will have taken one giant step for your investors. The golden rule—putting the client first—is the eternal holy grail we should all be seeking.

---

[1] John C. Bogle, "On Reversion to the Mean: Sir Isaac Newton's Revenge on Wall Street," unpublished speech, 1998.

[2] William F. Sharpe, "The Styles and Performance of Large Seasoned U.S. Mutual Funds, 1985-1994," published on the World Wide Web, 1995.

[3] Mark M. Carhart, "On Persistence in Mutual Fund Performance," *The Journal of Finance*, Volume LII, No. 1, March 1997, pages 57-82.

[4] Burton G. Malkiel, "Returns from Investing in Equity Mutual Funds 1971 to 1991," *The Journal of Finance*, Volume L, No. 2, June 1995, pages 549-571.

[5] Stephen J. Brown and William N. Goetzmann, "Performance Persistence," *The Journal of Finance*, Volume L, No. 2, June 1995, pages 679-698.

Note: Revised Copy, April 24, 1998