

Do Vanguard ETF Investors Make Good Decisions? – Testing the Bogle Hypothesis

Meng Xie

Professor Edward Tower, Faculty Advisor

*Honors Thesis submitted in partial fulfillment of the requirements for Graduation with Distinction in
Economics in Trinity College of Duke University*

Duke University
Durham, North Carolina
2012

Acknowledgements

I would like to thank Professor Tower for his support and expertise. This thesis is the result of our collaborative research project this semester. We are grateful to Alecia Waite for writing the VBA program we used and for her comments. We are also grateful to Bill Bernstein, John Bogle, Mel Lindauer, Larry Swedroe, and Allan Roth for their discussions, without implying their approval.

Abstract

John Bogle, the founder of Vanguard, is a notable opponent of frequent ETF trading. We test his hypothesis that Vanguard investors are not trading ETFs intelligently. A comparison of dollar-weighted and time-weighted returns is the typical method used for assessing investor timing. We instead employ Sharpe's style analysis techniques to compare the returns of a portfolio of ETFs to a basket of standard Vanguard funds that mimics the ETF portfolio's pattern of returns. We find that the ETF portfolio underperforms the standard Vanguard funds, providing empirical evidence supporting Bogle's view that ETFs are misused.

JEL classification: G11

Keywords: Exchange-traded funds; Mutual funds; Vanguard; Investment

Section I: Introduction

Exchange traded funds (ETFs) were launched in the United States in 1993 and have seen rapid growth since their inception. According to Ferri (2009), “ETFs available for investment rose more than ten-fold between December 2003 and December 2008, from 71 to 747, including a 221-fund increase in 2008.” ETFs are funds that can be traded throughout the trading day like stocks. Unlike a company stock, the number of shares can change throughout the day due to the existence of a process for the creation and redemption of shares that allows the ETF price to match that of its underlying securities. Benefits of ETFs include the fact that they can be used as a tool for fast portfolio readjustment and that they allow investors to gain exposure to certain sectors of the market almost instantly. Furthermore, ETFs typically have lower costs than managed or index mutual funds because clients interact with brokerages rather than the ETF companies. Finally, ETFs can be more tax-efficient because shareholders do not have to pay capital gains taxes on any excess capital gains earned and distributed by the mutual fund. (Ferri, 2009)

There has long been debate over whether investors trade any asset intelligently, and this is particularly true for ETFs, given their rapid surge in popularity. It is unclear whether ETFs enhance the abilities of investors to increase returns or whether they simply tempt investors to make poor decisions through over-frequent trading and poorly timed performance chasing. Research into the ability of ETF investors to time their trades correctly can contribute to the discussion of whether this increased liquidity is worth the potential costs. This would help investors to make more informed decisions. It is also of interest from a regulatory standpoint, addressing the question of whether regulators should instate rules constraining ETF trading activity to protect investors from themselves.

John Bogle, the founder of Vanguard, the world’s largest mutual fund company, is a notable opponent of ETF trading. He is known for believing in a “staying-the-course” strategy, in which investors pick a portfolio mix and stay with it through rebalancing. This means that investors decide on an asset allocation that suits their personal level of risk and trades assets with the goal of maintaining that balance. A portfolio can become unbalanced due to particular market conditions. For example, consider a portfolio of stocks and bonds. If the stock market is booming, then the allocation to stocks will increase and the relative allocation to bonds will decrease correspondingly, increasing the risk of the portfolio. To rebalance, the investor now needs to sell stocks and buy bonds. Bogle believes that ETFs tempt investors to trade excessively, resulting in lower long-term investment returns than staying the course. Given his long-standing experience in and reputation for providing investors with low-cost efficient investment options, we aim to implement an empirical test of his hypothesis. Vanguard launched its first ETF in 2001. We will assess specifically whether investors in Vanguard ETFs successfully use this asset to their advantage.

The issue of investor timing has been long-researched. Traditionally, the returns of a mutual fund were assessed with the time-weighted returns of the asset, calculated as the $(\text{final price} - \text{initial price}) / \text{time period}$. Currently, the widespread strategy is to use dollar-weighted returns instead. Dollar-weighted returns weight the returns of a mutual fund over a certain period by the amount of money invested in the fund at that time. The dollar weighted return will be higher if the firm makes high returns when there are large investments in that fund and thus reflects the earnings that the average investor experiences. Dollar-weighted returns, also called investor returns, take into account the fact that returns earned by investors can differ from the

return of the fund itself. If the dollar-weighted returns for a certain mutual fund are lower than the corresponding time-weighted returns, we say that there is evidence of poor investor timing.

In this paper, we further expand the past research in investor timing through examining returns in a portfolio context. While dollar-weighted returns are a better indicator of investor performance than time weighted returns, they do not take into account changes in the number of funds over time. For example, an investor who sells a high-performing fund early may be reinvesting his money into an even higher-returning fund. Alternatively, investors may just be spreading their money more thinly as more funds come into existence. This issue is particularly pertinent to ETFs given their huge growth. We instead consider an ETF portfolio to take into account the change in number of funds over time. We compare the performance of this portfolio to an index and near-index fund benchmark of similar style to see whether an investor would fare better investing in the ETF portfolio or the standard portfolio. Specifically, we calculate the returns of the ETF portfolio weighted by the size of each ETF at each time point and compare to the corresponding return of a similarly-returning portfolio composed of standard funds. Thus, we take into account the idea of dollar-weighted returns by weighting the returns by investor holdings at each time point and also making a comparison with an intuitive benchmark.

Section II: Background

Previous research using dollar-weighted returns have suggested that investors do not move in and out of mutual funds optimally. Morningstar examined investor performance in 17 mutual fund categories and found that the returns earned by investors were below the returns of the funds themselves in every category. For large-, mid-, and small-cap growth funds, the 10-year annualized dollar-weighted return was 3.4, 2.5, and 3.0 percent lower than the time-weighted return, respectively. Large- and small-cap value investors underperformed their funds by 0.4 and 2.0 percent per year, respectively. Health care sector and financial sector investors underperformed by 4.0 and 1.6 percent, respectively. However, the most drastic result was that technology sector investors underperformed the funds they invested in by 14 percent per year. This supports the study's further observation that investors in more volatile funds see greater differences in the time-weighted and dollar-weighted returns because these investors unsuccessfully attempt to chase unpredictable returns. In a separate study, Bogle Financial Markets Center researchers compared time-weighted and dollar-weighted returns from 1996 through 2000 for the 200 mutual funds with the largest cash inflows. The average time-weighted return of all the funds was 8.9 percent per year, while the average dollar-weighted return was only 2.4 percent per year. The dollar-weighted return was higher in only two of these funds. For 76 funds, the dollar-weighted returns fell short of the time-weighted returns by 50 to 95 percent. Both studies provide evidence of poor investor timing. (Swedroe, 2011)

Bogle (2007) uses the period from 1997 to 2002, in which the economy experienced a bull market followed by a bear market, to study examine investors' performance chasing tendencies. Looking at the top ten funds during the bull market, he found that while the net return over the entire period ended up positive even after the economic downswing, the fluctuation in prices causes investors to take chances and mistime the market. The top ten funds had an average net gain of 13 percent, while these funds' investors lost an average of 57 percent in that same period. In another example, Bogle examines the case study of Vanguard's Growth Index Fund and Value Index Fund, both created in 1992. He explains that the former was

targeted toward young investors willing to take on more risk to accumulate wealth, while the latter was aimed toward older investors wanting to preserve wealth with low-risk assets. Again, however, once the market started moving drastically, people were tempted to chase upward price trends rather than holding onto their respective funds. In the bull market from 1997 through 2000, growth stocks boomed. From 1992 to March 2000, the Growth Index total return was 364 percent, as compared to the 229 percent return of the Value Index. Correspondingly, people invested \$11 billion into the Growth Index Fund, but only \$3 billion in the Value Index Fund. In the switch to a bear market in 2000 through 2002, the growth stocks dropped. From 2001 to 2006, there were net redemptions of \$850 million in the Growth Index Fund and net purchases of almost \$2 billion in the Value Index Fund. Between 1993 and 2006, the time-weighted returns of the Growth fund and Value fund were 9.1 and 11.2 percent per year, respectively. Due to investors' "counterproductive timing and selection", the dollar-weighted return of investors in the Growth and Value funds were only 0.9 and 7.6 percent per year, respectively.

These previous studies using dollar-weighted returns show consistently that investors time the market poorly even with respect to regular mutual funds. Our study will use a new methodology to look at ETF investors, who have an even greater opportunity to chase returns given the potential frequency of their trades.

Section III: Methodology

Overview

Our strategy is to ask whether the aggregate of Vanguard investors have experienced higher risk-adjusted returns from investing in ETFs than they would have from investing in comparable Vanguard investor class mutual funds. To do this, we first calculate the series of ETF portfolio returns weighted by the percentage of total investments in each ETF. Then, we follow the Style Analysis technique of Nobel laureate William Sharpe to infer the "style" of investment in the ETF portfolio based on its pattern of returns (without needing to look at the composition of the portfolio). The style of investment is then replicated by calculating the corresponding weights of investment in each asset of a benchmark portfolio. We use a basket of Vanguard index and near-index funds as our benchmark portfolio, since we are interested in the differences between the returns of Vanguard ETFs and what we will henceforth refer to as Vanguard's standard mutual funds. An example of an index fund would be the fund tracking the S&P500, and examples of near-index funds include the international growth and international value funds. This set of standard funds is chosen to reflect the broad range of investments also available through ETFs.

To determine whether Vanguard investors make wise style adjustments over time, we perform two comparisons of the Vanguard ETF portfolio to the basket of standard Vanguard funds. First, we compare the performance when the weights of the standard Vanguard funds are chosen so that the pattern of monthly returns of the benchmark portfolio matches that of the ETF portfolio. We will refer to this as the stay-the-course clone of the ETF portfolio. This represents a clone portfolio in which the investor chooses a portfolio mix and then maintains it through rebalancing. Second, we compare performance when we allow the weights in the benchmark to vary each period thereby reflecting the changing style of the ETF portfolio. We refer to this portfolio as the market-timing clone. Taken together, if the ETF portfolio outperforms the stay-

the-course clone by more than the market-timing clone, we conclude that part of the ETF outperformance is due to style adjustments. On the contrary, if the ETF portfolio underperforms the stay-the-course clone portfolio by more than the market-timing clone, we conclude that poor market timing contributed further to the underperformance of the ETFs.

ETF Categories and Calculation of the ETF Portfolio Return Series

We downloaded monthly data for returns, prices, and number of shares outstanding for all of Vanguard's ETFs from the Center for Research in Security Prices (CRSP) website. These returns have been adjusted for expense ratios. We obtained the list of all Vanguard ETFs from Morningstar. We calculated the market capitalization (by multiplying the price by the number of shares outstanding) of each ETF for each month and then found its percentage of the total market capitalization of all ETFs in that month. We then weighted each ETF's return by this percentage and added all the weighted monthly returns at each period to calculate the series of monthly returns for the ETF portfolio. In addition, we downloaded the monthly returns for each of the standard Vanguard funds in the benchmark basket. We use these return series to compare the ETF portfolio with a portfolio of standard funds. The number of ETFs in existence increased substantially from the beginning of our series to the end, confirming the necessity of a method of assessing investor competence that takes into account this increase in investment options.

We use the calculation above for a portfolio of all existing ETFs. In addition, we perform the same calculations to find the returns series of portfolios containing only certain categories of ETFs. The categories considered are stock index ETFs, bond ETFs, and sector ETFs, and international ETFs. Also, the combination of the stock index ETFs, bond ETFs, international ETFs, and two specialty equity ETFs make up a portfolio of stock ETFs. The Appendix lists all the funds used along with their assigned categories. We also specify the span of the returns series for each category of funds investigated. We collected monthly data from February 2005 through December 2011, but because not all types of ETFs existed for that entire period, some calculations were done over a shorter period. For example, Vanguard did not launch its first bond ETF until April 2007. See the appendix for specific dates used.

Use of the Geometric α

While the data we collect and returns series we calculate are arithmetic returns, we convert these values to geometric returns when implementing the analysis. Thus, the criterion used for outperformance of the ETF portfolio is the geometric α , the continuously compounded geometric return of the ETF portfolio minus that of the clone. Geometric average returns are typically calculated as $\ln(\text{value at end of period} / \text{value at beginning of period}) / \text{length of period}$, and they reflect continuously compounding returns. We prefer the geometric α because it reflects how rapidly a portfolio grows from the beginning of a time span to the end. In addition, geometric returns are conveniently additive across adjacent time periods, while arithmetic returns are not. This property allows us to simply multiply by 12 when finding annual α 's using monthly data. The following arguments regarding the benefits of using the geometric α over the arithmetic α are discussed in Tower (2009b) and developed with a numerical example in Tower (2009a). The following are a few arguments presented in these two papers. First, two funds

experiencing the same cumulative return over a particularly time period will differ in their arithmetic α values, with the more volatile fund having higher arithmetic α , but the two funds will have the same geometric α . Second, the geometric α is a better predictor of future returns when the distribution of future returns is drawn from a distribution of past returns without replacement. Third, geometric returns, which reflect continuous compounding, are likely to generate a more stable α . When using arithmetic returns to compare two portfolios that differ only in expense ratios, α can be lower in falling markets than in rising markets. For example, if the value of a fund falls to zero during the year, its return rate will be close to 100 percent regardless of the expense ratio. In a rising market, though, differences in expense ratios can lead to larger differences in cumulative return over the year. However, under continuous compounding, constant expense ratios will lead to constant α .

Constructing the Stay-the-Course Clone Portfolio

We use Sharpe's Style Analysis to construct the stay-the-course clone portfolio. Following the Style Analysis presented in Sharpe (1992), the return of the ETF portfolio is described as the return of a basket of standard Vanguard funds plus a constant term, plus an error term with mean zero. That is,

$$R_{ETF} = \alpha + \beta R_{Standard} + \varepsilon,$$

where $\beta = \text{Cov}(R_{ETF}, R_{Standard}) / \text{Var}(R_{Standard})$.

In our case, β is constrained to be zero to risk-adjust the portfolios, as explained in the next paragraph. The size and sign of the constant term α reflects how well investors perform when investing in ETFs relative to investing in the basket of standard Vanguard funds. Thus, α is the average excess return of the ETF portfolio compared with the benchmark portfolio over all the months included in the analysis. The portfolio of the standard funds that best tracks the ETF portfolio is chosen to match the coverage of the ETFs examined, and to minimize the mean squared sum of the error terms. If the ETF portfolio is largely composed of international stocks, then the composition of the clone will largely consist of funds with international stocks, since the returns of that combination of funds will best mimic the returns of the ETF portfolio. The size of the constant term α depends positively on the expenses and turnover rates of the Vanguard standard funds relative to the ETF portfolio. The expense ratio of a fund is the fund's operating expenses divided by the average value of the assets in the fund. The return of the fund is lowered by this amount. The turnover ratio is the percentage of a mutual fund's holdings that are replaced within a year. Sharpe's methodology is appealing in that it can capture individual nonlinear relationships between returns and various fund characteristics like size and value.

Furthermore, we want to compare the ETF portfolio with the standard Vanguard mutual funds portfolio with the same level of risk, as measured by standard deviation. Modigliani and Modigliani (1997) risk adjust by adding a risk-free asset to the standard portfolio and allow the weight on this asset to fluctuate. We use a variant of this risk-adjustment technique by choosing the weights on the standard funds while constraining the standard deviation of the risk-free-asset augmented standard portfolio and the ETF portfolio to be the same. In solving for the weights of the clone portfolio including this extra constraint, weight put on low-risk bond assets can adjust the risk of the two portfolios so that they are equal.

We execute this methodology using Microsoft Excel’s Solver utility. Solver is instructed to find the weighted sum of the returns of the Vanguard indices and the constant term that minimizes the mean squared error of prediction subject to the constraints that all the coefficients are positive and add to one. That is, Solver is instructed to find the weights ω to

$$\min\{\varepsilon\} \text{ such that } \omega_i > 0, \text{sum}(\omega_i)=1, \text{ and } SD_{\text{ETF}} = SD_{\text{Standard}}.$$

Qualitatively, this means that there is no short selling and that the investor invests all the money allocated for this portfolio (without borrowing) at each time period. We use continuously compounded monthly returns for both the ETF portfolio and the standard basket of funds in this calculation.

Constructing the Market-Timing Clone Portfolio

To assess the month-to-month ETF portfolio adjustments, we calculate the weights of the clone portfolio at each time period. Following the methodology of Blanchett (2010), we determined the benchmark composition in each month by requiring similar performance of the ETF portfolio and the benchmark over a window of months before and after the month in question. We follow this approach using a 15-month window, with seven months before and after the month in question. We calculate for the portfolio weights at each month by performing the same minimization with the same constraints as for the stay-the-course clone. We then record the excess return between the benchmark portfolio and the ETF portfolio for the month in question. We calculate the clone portfolio weights that mimic the ETF return for each month starting at the eighth month in our dataset. The α reported is then the average of the recorded excess return for each month. Using Blanchett’s terminology, we can say that we derived the stay-the-course α from a “full-period regression”, while we calculated the market-timing α from a series of “rolling regressions”. To have comparable coverage in the market-timing and stay-the-course clones, we adopt the same start and end dates for both analyses. Again, the start and end dates used for each ETF portfolio are specified in the appendix. We differentiate the α ’s for the two clones by calling them the stay-the-course α and the market-timing α .

Section IV: Results

We first analyzed the ETF portfolio by including all the existing ETFs. The stay-the-course α for this portfolio is -1.25 percent per year, meaning the ETF portfolio underperforms its corresponding basket of standard funds by 1.25 percent per year. The market-timing alpha of this ETF portfolio is -1.05 percent per year. This means that investors would have performed 1.25 percent per year better if they had invested in the stay-the-course clone, but only 1.05 percent per year better if they had invested in the market-timing clone. This suggests that not only were ETFs a worse investment given investor behavior, the style adjustments used in the ETF portfolio further hurt the ETF investors’ returns. ETF investors would have done better on average if they had invested in standard mutual funds and maintained a constant style through rebalancing. We perform two-tailed significance tests of on the α values calculated using Microsoft Excel’s “t-test: paired two sample for means” functionality on the two series of

monthly geometric returns. Both of these α values are statistically significant at the five percent level. These results are presented in Exhibits 1 through 3.

Second, we consider a portfolio of only equity ETFs, i.e. all except bond ETFs. The benchmark funds used are now the equity subset of the entire basket of standard Vanguard funds. In this case, the stay-the-course α and the market-timing α are -0.74 percent per year and -0.51 percent per year. Again, the ETF portfolio underperforms the portfolios of standard Vanguard funds and style adjustments further contribute to the underperformance. The two-tailed significance levels on these α values are 10.75 percent and 4.11 percent, respectively. These results are reported in Exhibits 4 through 6. Overall, ETF investors make poor decisions in choosing the mix of equity ETFs and even worse decisions in choosing the mix of equity and bond ETFs.

Third, we consider four ETF portfolios, each representing a particular category of ETFs: equity index ETFs, bond ETFs, international equity ETFs, and industry sector stock ETFs. For each type of ETF portfolio, the benchmark funds used are those corresponding to the composition of the ETF portfolio. For example, the ETF portfolio comprised of bond ETFs will have the standard bond funds as its benchmark, excluding equity funds. The α results for each of these four ETF types are presented in Exhibit 7. None of the calculations were statistically significant, but we present the numbers to emphasize further that there is no evidence of outperformance of the ETF portfolio compared to the benchmark of standard Vanguard funds. For both the bond ETF portfolio and the international equity ETF portfolio, both the stay-the-course and market-timing α 's are negative, suggesting poor investor decisions in making style adjustments. The insignificance of these values prevents us from making a stronger statement. For the sector and equity index ETF portfolios, the stay-the-course α 's were positive, while the market-timing α 's were negative. Because the positive α 's were not statistically significant, we cannot say that there is evidence of outperformance by the ETF portfolio.

Discussion of Admiral Shares

Vanguard offers its standard funds in different share classes. Share classes are differentiated by varying expense ratios that reflect the cost to Vanguard of providing services to the investors in each share class. Of particular interest to us are admiral shares, a class of shares with lower expense ratios available only to Vanguard investors whose accounts exceed certain lower limits. According to the Vanguard website, this share class was created to “recognize and encourage the cost savings stemming from large investment accounts”. The calculations in this paper use the returns of Vanguard’s investor shares, Vanguard’s traditional shares aimed at individual investors and requiring low initial investments. We chose this share class because some of the corresponding standard Vanguard index and near-index funds do not have an Admiral class. To examine the difference in results if we use admiral shares rather than investor shares, we could redo the analysis, replacing the expense ratio-adjusted returns for the admiral shares for investor shares whenever possible. A related future study could also perform our calculations across all share classes, weighted for the percentage of investors in each share class.

Whereas previously, admiral shares either did not exist for most funds or required a very large initial investment, admiral shares have recently become more prevalent and accessible. Thus, an extrapolation of this study for future use might consider an investor trading admiral shares instead of investor shares. To begin exploring this issue, we compare the weighted

average expense ratios of the investor shares used in our calculations (from December 2007 through May 2011) with the expense ratios of potential investments in Admiral shares in 2012. To calculate the former, we use the average yearly expense ratio based on values found on Morningstar. We then multiply these yearly averages by the benchmark portfolio weights calculated for December 2011 to obtain a weighted average expense ratio. For the 2012 expense ratio for the benchmark funds, we use the expense ratios reported by Vanguard in March 2012. The expense ratios for the Admiral shares were used whenever possible, but if the fund still does not offer Admiral shares, expense ratios for the investor shares were used. Again, these expense ratios were weighted by the December 2011 portfolio weights for the benchmark mimicking the ETF portfolio consisting of all existing ETFs. The portfolio-share-weighted average change in expense ratio is -0.0927 percent per year.

For the ETFs, we consider the expense ratios of the eleven largest ETFs in December 2011. For these shares, the difference in current expense ratio and the average expense ratio over our period, weighted by portfolio share, is 0.0136 percent per year. The current expense ratio is lower than in the past for ETFs. The lower expense ratio in ETFs, would cause the α to be less negative, while the lower expense ratios in the standard Vanguard Admiral shares in 2012 would cause the α value calculated to be even lower. In total, these differences would cause α to be lower by 0.0791 percent per year, about eight basis points per year. This not only further strengthens the conclusion that ETF investors make poor timing decisions on average, it would increase the statistical significance of the negative α values shown in Exhibits 1 through 3. This exploratory calculation suggests that a full study including share classes could be worthwhile.

Section V: Conclusion

We conclude from this study that investors get better returns when investing with standard index and near-index mutual funds than when trading ETFs. We find that Vanguard ETF portfolios underperform when compared with standard Vanguard portfolios of the same style. That is, ETF investors receive lower returns than if they invest in baskets of standard Vanguard funds that mimic their investing style. We find this effect to be significant when using a portfolio consisting of all the Vanguard ETFs, and there is no evidence of outperformance using portfolios consisting only of particular ETF categories. Moreover, the underperformance of the market-timing clone portfolio is larger on average than for the constant weight, stay-the-course clone portfolio. This leads us to conclude that ETF investors make poor decisions in both the short and long run, and that style drift leads to even more detriment in returns. Our calculations support Bogle's hypothesis that ETF investors are not trading intelligently and that their desire to time the market hurts rather than helps their returns. Thus, ETF investors should reconsider their trading strategies, potentially opting instead for Bogle's recommended stay-the-course strategy.

Several further considerations and ideas for future research have emerged from this study. One potential shortcoming in our research is that we only consider a portfolio comprising entirely of Vanguard ETFs. In reality, Vanguard ETFs are often a supplement to other portfolio holdings. For example, someone may hold standard Vanguard mutual funds in a retirement account that does not permit ETFs, but supplement that account with a personal account comprised of ETFs. Thus, further research could look into whether ETFs help or hurt portfolio returns when they are a supplement to standard mutual funds. We expect that the conclusions

would be similar because our methodology adjusts for the average style of ETF investing across investors, and this style could transfer to ETF investments even when the portfolio is comprised of more assets than just ETFs. Furthermore, Bogle is particularly opposed to frequent ETF trading, but he also states that frequent trading in Vanguard index funds is unwise as well. Thus, further research could explore whether investors also make poor decisions when trading Vanguard index mutual funds. We could also study whether investors switch successfully between ETFs and index mutual funds. In addition, future research could compare our research on Vanguard funds with data from other fund families, like the Dimensional Fund Advisors (DFA) or Fidelity. Finally, the poor decisions made in ETF investment could also be because they are relatively new. Another study could look at the quality of ETF investment decisions over time from the time of their inception. This study would likely require examining ETFs from more than one fund family, as we have done here. The researchers could then suggest whether they believe ETF investment decisions will improve as investors learn more about these assets. Further research will provide more insight into the behavior of investors responding to the different incentives of various asset types and enrich the empirical grounding for individual investor decisions.

Appendix: Results

Stay-the-Course Results for the Full ETF Portfolio					
alpha (% per year)	MSE %	Correlation %	Stock share %	International stock share %	Two-tailed significance %
-1.25	0.24	0.999	88	31	2.63

Exhibit 1: Stay-the-course alpha for the ETF portfolio containing all existing ETFs, including both stocks and bonds. The ETF portfolio underperforms the stay-the-course clone.

Standard Vanguard Funds	proportion
500 Index	0.30
Emerging Markets Stock Index	0.12
European Stock Index	0.02
Extended Market Index	0.00
Growth Index	0.09
High-Yield Corporate	0.00
Intermediate-Term Bond Index	0.00
International Explorer	0.00
International Growth	0.15
International Value	0.02
Long-Term Investment-Grade	0.05
Mid-Capitalization Index	0.08
Pacific Stock Index	0.00
Prime Money Market	0.07
REIT Index	0.06
Short-Term Bond Index	0.00
Small Cap Growth Index	0.00
Small Cap Index	0.00
Small Cap Value Index	0.03
Vanguard Total Bond Market Index	0.00
Vanguard Total International Stock Index	0.00
Vanguard Value Index	0.00

Exhibit 2: Vanguard index and near-index funds share coefficients from the stay-the-course simulation with the full ETF portfolio. The standard Vanguard funds used are all Investor class shares. This also displays all the standard Vanguard funds used in our analysis.

Market-Timing Results for the Full ETF Portfolio								
alpha (% per year)	MSE %	Average stock share %	Average international stock share %	Correlation between predicted and actual return	Standard deviation of stock share %	Standard deviation of fund return (% per mo)	Excess standard deviation of fund return (% per mo)	Two-tailed significance %
-1.05	0.20	90	33	99.9	4	5	-0.0058	0.73

Exhibit 3: The full ETF portfolio loses significantly to the market-timing portfolio.

Stay-the-Course Results for the Stock ETF Portfolio					
alpha %	MSE %	Correlation %	Stock share %	International stock share %	Two-tailed significance %
-0.74	0.30	0.998	100	31	10.57

Exhibit 4: The ETF portfolio of all stock ETFs loses to the stay-the-course clone.

Standard Vanguard Funds	proportion
500 Index	0.13
Emerging Markets Stock Index	0.10
European Stock Index	0.09
Extended Market Index	0.11
Growth Index	0.25
International Explorer	0.00
International Growth	0.04
International Value	0.00
Mid-Capitalization Index	0.00
Pacific Stock Index	0.00
REIT Index	0.07
Small Cap Growth Index	0.00
Small Cap Index	0.00
Small Cap Value Index	0.00
Vanguard Total International Stock Index	0.08
Vanguard Value Index	0.12

Exhibit 5: Vanguard index and near-index funds share coefficients from the stay-the-course simulation with the stock ETF portfolio. The standard Vanguard funds used are all Investor class shares.

Market-Timing Results for the Stock ETF Portfolio							
alpha (% per year)	MSE %	Average stock share %	Average international stock share %	Correlation between predicted and actual return	Standard deviation of international stock share %	Excess standard deviation of fund return (% per mo)	Two-tailed significance %
-0.51	0.16	100	32	99.96	0	-0.0085	4.11

Exhibit 6: The ETF portfolio of all stock ETFs loses to the market-timing clone.

Alpha's for Different Types of ETF Portfolios				
Portfolio	Stay-the-course alpha (% per yr)	Market-timing alpha (% per yr)	Stay-the-course significance %	Market-timing significance %
Bond ETF Portfolio	-0.51	-0.20	44.5	75.7
International Equity ETF Portfolio	-0.74	-0.69	40.9	42.7
Industry Sector ETF Portfolio	0.10	-0.74	90.7	48.1
Equity Index ETF Portfolio	0.10	-1.22	68.6	61.9

Exhibit 7: None of the alpha values for specific types of ETFs were statistically significant, but they do not provide any evidence against our conclusions of the ETF portfolios' underperformance. We report the p-values for two-tailed significance tests.

Appendix: Data

Beginning and End Dates for Analyses		
	Begin	End
Full ETF Portfolio	December 2007	May 2011
Stock ETF Portfolio	November 2005	May 2011
Bond ETF Portfolio	December 2007	May 2011
International Equity ETF Portfolio	November 2005	May 2011
Industry Sector ETF Portfolio	September 2005	May 2011
Equity Index ETF Portfolio	September 2005	May 2011

Exhibit 8: Beginning and end dates for the analyses on each of the ETF portfolios. We gathered data starting in February 2005 whenever possible. However, due to the later introduction of international and bond funds, we had to start some series later. As explained in the section of the creation of the market-timing clone, we start both the stay-the-course and market-timing analyses 7 months after the first month of portfolio data and end the analyses 7 months before December 2011, the last period of data collected.

Number of ETFs During Our Study		
	February 2005 (number in first month of existence if none existed in 2/05)	December 2011
Equity Index ETFs	10	30
Bond ETFs	0 (4)	12
International ETFs	0 (3)	9
Industry Sector ETFs	11	11
Specialty ETFs	2	2

Exhibit 9: ETF fund growth over the period of our study.

Fund Name	Ticker	Type
Vanguard Extended Dur Trs Idx ETF	EDV	bond
Vanguard Interm-Tm Corp Bd Idx ETF	VCIT	bond
Vanguard Interm-Tm Govt Bd Idx ETF	VGIT	bond
Vanguard Intermediate-Term Bond ETF	BIV	bond
Vanguard Long-Term Bond Index ETF	BLV	bond
Vanguard Long-Term Corp Bond Idx ETF	VCLT	bond
Vanguard Long-Term Govt Bd Idx ETF	VGLT	bond
Vanguard Mortgage-Backed Sec Idx ETF	VMBS	bond
Vanguard Short-Term Bond ETF	BSV	bond
Vanguard Short-Term Corp Bd Idx ETF	VCSH	bond
Vanguard Short-Term Govt Bd Idx ETF	VGSH	bond
Vanguard Total Bond Market ETF	BND	bond
Vanguard FTSE All-Wld ex-US SmCp Idx ETF	VSS	intl
Vanguard FTSE All-World ex-US ETF	VEU	intl
Vanguard Global ex-US Real Estate ETF	VNQI	intl
Vanguard MSCI EAFE ETF	VEA	intl
Vanguard MSCI Emerging Markets ETF	VWO	intl
Vanguard MSCI Europe ETF	VGK	intl
Vanguard MSCI Pacific ETF	VPL	intl
Vanguard Total Intl Stock Idx ETF	VXUS	intl
Vanguard Total World Stock Index ETF	VT	intl
Vanguard Consumer Discretionary ETF	VCR	sector
Vanguard Consumer Staples ETF	VDC	sector
Vanguard Energy ETF	VDE	sector
Vanguard Financials ETF	VFH	sector
Vanguard Health Care ETF	VHT	sector
Vanguard Industrials ETF	VIS	sector
Vanguard Information Technology ETF	VGT	sector
Vanguard Materials ETF	VAW	sector
Vanguard REIT Index ETF	VNQ	sector
Vanguard Telecom Services ETF	VOX	sector
Vanguard Utilities ETF	VPU	sector
Vanguard Dividend Appreciation ETF	VIG	specialty (stock)
Vanguard High Dividend Yield Indx ETF	VYM	specialty (stock)
Vanguard Extended Market Index ETF	VXF	stock
Vanguard Growth ETF	VUG	stock
Vanguard Large Cap ETF	VV	stock
Vanguard Mega Cap 300 Gr Index ETF	MGK	stock
Vanguard Mega Cap 300 Index ETF	MGC	stock
Vanguard Mega Cap 300 Value Index ETF	MGV	stock

Vanguard Mid-Cap ETF	VO	stock
Vanguard Mid-Cap Growth ETF	VOT	stock
Vanguard Mid-Cap Value ETF	VOE	stock
Vanguard Russell 1000 Growth Index ETF	VONG	stock
Vanguard Russell 1000 Index ETF	VONE	stock
Vanguard Russell 1000 Value Index ETF	VONV	stock
Vanguard Russell 2000 Growth Index ETF	VTWG	stock
Vanguard Russell 2000 Index ETF	VTWO	stock
Vanguard Russell 2000 Value Index ETF	VTWV	stock
Vanguard Russell 3000 Index ETF	VTHR	stock
Vanguard S&P 500 ETF	VOO	stock
Vanguard S&P 500 Growth Index ETF	VOOG	stock
Vanguard S&P 500 Value Index ETF	VOOV	stock
Vanguard S&P Mid-Cap 400 Growth Idx ETF	IVOG	stock
Vanguard S&P Mid-Cap 400 Index ETF	IVOO	stock
Vanguard S&P Mid-Cap 400 Value Index ETF	IVOV	stock
Vanguard S&P Small-Cap 600 Gr Idx ETF	VIOG	stock
Vanguard S&P Small-Cap 600 Index ETF	VIOO	stock
Vanguard S&P Small-Cap 600 Value Idx ETF	VIOV	stock
Vanguard Small Cap ETF	VB	stock
Vanguard Small Cap Growth ETF	VBK	stock
Vanguard Small Cap Value ETF	VBR	stock
Vanguard Total Stock Market ETF	VTI	stock
Vanguard Value ETF	VTV	stock

Exhibit 10: Fund names, tickers, and categories of the ETFs included in the analysis.

References

- Blanchett, D. M. (2010). Exploring the “good guys”: An empirical study of Vanguard’s actively managed domestic equity mutual funds. *The Journal of Investing*, 19(2), 39-48.
- Bogle, J. C. (2007). *The little book of common sense investing: The only way to guarantee your fair share of stock market returns*. (pp. 50-51, 92-93, 133-135). Hoboken, New Jersey: John Wiley and Sons, Inc.
- Bogle, J. C. (2002). The Telltale Chart. Keynote Speech before the Morningstar Investment Forum, Chicago, IL, June 26, 2002; available online from the Bogle Financial Markets Research Center.
- Fama, E. F. & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56.
- Ferri, R. A. (2009). *The ETF book: All you need to know about exchange-traded funds*. (pp. 3-4, 23-39, 57-78). Hoboken, New Jersey: John Wiley and Sons, Inc.
- Modigliani, F., & Modigliani, L. (1997). Risk-adjusted performance. *The Journal of Portfolio Management*, 23(2), 45-54.
- Rodriguez, A., & Tower, E. (2008). Do Vanguard’s managed funds beat its index funds?: Looking for prescient stock and style selection. *The Journal of Indexes*, 11(2), 26-35.
- Roth, A. (2012, Jan 19). *Have DFA advisors helped clients avoid mistakes?* Retrieved from http://www.cbsnews.com/8301-505123_162-57359825/have-dfa-advisors-helped-clients-avoid-mistakes/
- Sharpe, W. (1992). Asset allocation: Management style and performance measurement. *Journal of Portfolio Management*, 18(2), 7-19.
- Swedroe, L. (2011, Jan 18). *Mutual fund returns: Worse than you thought*. Retrieved from http://www.cbsnews.com/8301-505123_162-51499520/mutual-fund-returns-worse-than-you-thought/ (Excerpted from *The quest for alpha: The holy grail of investing*. (2011). Hoboken, New Jersey: John Wiley and Sons, Inc.)
- Tower, Edward. (2009a). Performance of actively managed versus index funds: the Vanguard case. Chapter 12 of *Mutual Funds: Portfolio Structures, Analysis, Management, and Stewardship* (pp. 211-236), edited by John A. Haslem. Hoboken: Wiley.
- Tower, Edward. (2009b). “Classic and enhanced index funds: Performance and issues.” Chapter 13 of *Mutual Funds: Portfolio Structures, Analysis, Management, and Stewardship* (pp. 237-264), edited by John A. Haslem. Hoboken: Wiley.
- The Vanguard Group, Inc. (2012). *Share class*. Retrieved from <https://personal.vanguard.com/us/content/Funds/FundsVGFundsShareClassDefJSP.jsp>