James L. Davis
Vice President
Dimensional Fund Advisors

PURELY ACADEMIC

August 2015

Long-Term Predictability and Short-Term Asset Allocation

Stock market commentators frequently say that day-to-day fluctuations in share prices may be hard to predict, but long-term returns are more predictable. The implication seems to be that investors should reflect these long-term return predictions in their current asset allocations. While the degree of predictability of long-horizon returns is an interesting subject for debate and analysis, this article approaches the practical implications of predictability from a different angle. Suppose long-term returns really are predictable. How should that impact my allocation today? For example, assume that 10-year returns are highly predictable. An investor who revises her asset allocation once every 10 years would find that information very useful. Is that predictability just as useful for investors who adjust their allocations more frequently—e.g., annually? When the return predictability horizon and the allocation frequency are not aligned, the benefits of predictability may be limited.

This study provides evidence on that question by setting up an experiment in which the degree of long-term return predictability is known. The benefit of this predictability is then measured for investors who adjust their asset allocations once per year based on the prediction. The results indicate that when there is a large difference between the allocation frequency and the horizon of return predictability, even relatively strong predictability may not benefit investors much.

I simulate 5,000 investment lifetimes, each 40 years long, by bootstrapping annual returns from the historical sample of US equity and Treasury bill returns.[1] Each of my 5,000 hypothetical investors has a model for forecasting the average 10-year equity premium (i.e., the return difference between stocks and T-bills). I determine how good the forecasting model is by setting the R^2 of the model to a pre-specified level. The R^2 indicates how much of the behavior of the average premium is captured by the model. For example, an R^2 of 1.0 indicates that the model explains all the variation in the 10-year average equity premium, while an R^2 of 0.0 indicates that the model explains none of it. The Appendix describes how the desired R^2 is imposed on the bootstrapped samples of the equity premium. The 5,000 investors observe their forecasts of the average equity premium over the next 10 years at the beginning of the year and use that forecast to decide whether to invest in stocks for the year. If the forecast is positive, the portfolio is 100% stocks. If the forecast is negative, the portfolio is 100% Treasury bills.

1. Based on bootstrapped simulations using stock and Treasury bill historical returns for 5,000 lifetimes (each 40 years long). A bootstrap simulation is a method of analysis that can be used to approximate the probability of certain outcomes by running multiple trial runs, called bootstrapped samples, using historical returns. The projections or other information generated by bootstrapped samples regarding the likelihood of various investment outcomes are hypothetical in nature, do not reflect actual investment results, and are not guarantees of future results. Results will vary with each use and over time. See the Appendix for data sources.

Investors adjust their allocations every year based on their forecast, even though their models can only forecast the average premium for the next 10 years.

Investors need a benchmark to determine if their forecast-based allocation has been successful. The benchmark used in this study is the simple strategy of remaining invested in the stock market. If the long-horizon predictability is beneficial, the forecast-based allocation strategy should produce higher ending wealth for a high proportion of the 5,000 investor lifetimes. If the pre-specified R^2 of the forecasting model is high, investors have a lot of information about the average equity premium over the next 10 years. The unanswered question is: How much information does this long-horizon predictability provide for the equity premium in the very next year? For investors who change their allocations every year, that is the important question.

The main point of this exercise is to show what the results look like when the predictability horizon (10 years in this case) and the allocation frequency (one year) are not aligned. As the two become more closely aligned, intuition suggests that the predictability may become more valuable. To test this, I repeat the exercise described above for five-year predictability. In other words, each investor now has a model for the average equity premium over the next five years, and since the predictability horizon and allocation frequency are now closer to one another, the predictability is more likely to provide a benefit to investors.

Exhibit 1 shows success rates for both 10-year and five-year predictability. Success rates indicate the percentage of the 5,000 simulated lifetimes for which the prediction-based allocation produced a higher ending wealth than the constant all-equity allocation. As indicated on the horizontal axis, the R^2 of the forecast varies between 0

Exhibit 1: Success Rates for Annual Rebalancing
5-Year and 10-Year Predictability

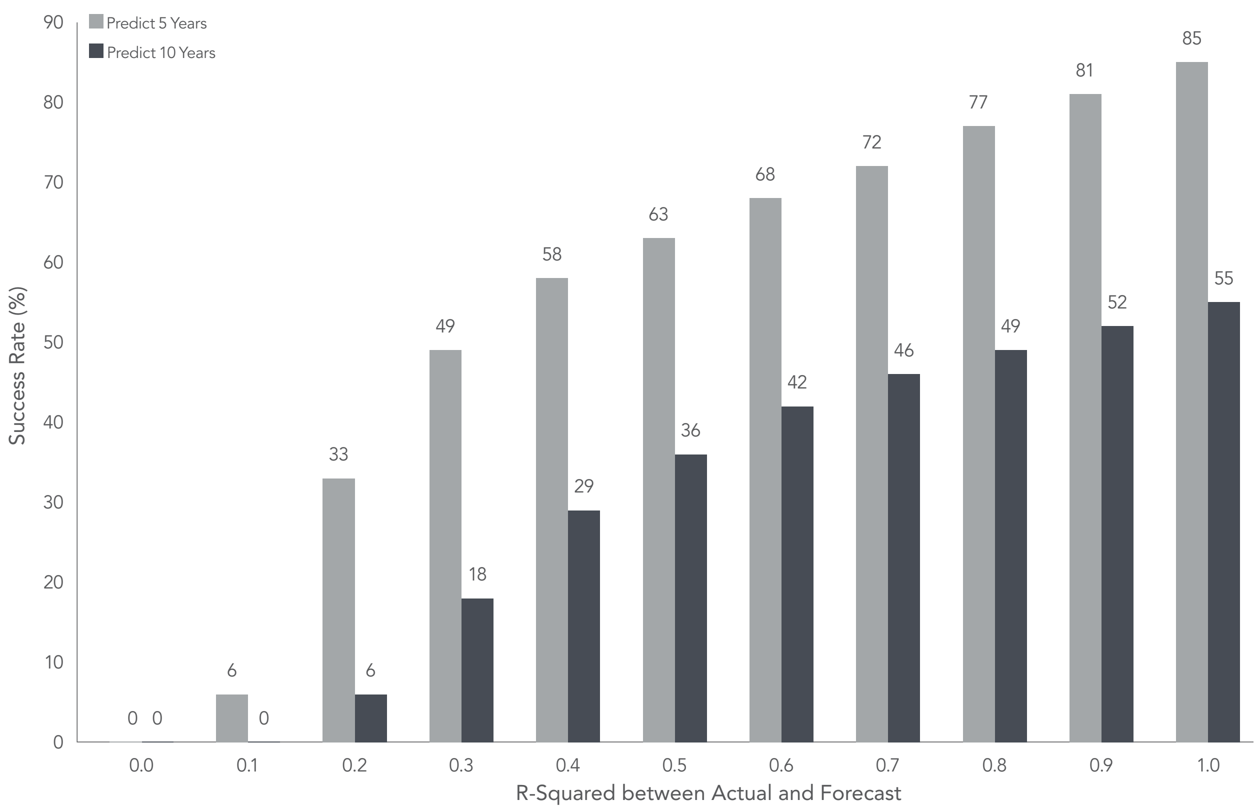

Success rates are the percent of simulations in which the dynamic allocation strategy produced a higher ending portfolio value than the simple strategy of remaining invested in stocks. Light and dark bars indicate that the five-year and 10-year average equity premiums, respectively, are predictable with the R^2 indicated along the horizontal axis. The Appendix contains data descriptions and the procedure used for setting the forecast R^2 to the desired level. See footnote 1 for important information regarding bootstrap simulations and limitations of the same.

and 1 in increments of 0.10. Since the success rates are calculated across the 5,000 simulated investor lifetimes, they estimate the probability that investors can look back at the end of their lifetimes and conclude that they were better off basing their asset allocations on the forecasts.[2]

Exhibit 1 shows that a high degree of long-term predictability is necessary to improve likely results for investors who rebalance annually. For 10-year predictability, it takes an R^2 in the 80%–90% range to produce a chance of success similar to a coin flip. As expected, a given level of predictability is more valuable for the five-year premium, since the forecast horizon is closer to the rebalance frequency. However, even in this case, it takes an R^2 in excess of 30% to make the odds better than a fair coin flip.

The success rates in **Exhibit 1** are probably overstated. This analysis assumes that the parameters of the forecasting model are fixed and known. In reality, the parameters of just about any forecasting model are unknown, and they likely vary over time. This means that the parameters must be estimated, and the individual developing the model is trying to estimate a moving target. This parameter uncertainty makes return forecasts less precise than the ideal situation studied here.

The simulations in this study assume that investors base their asset allocations on the most recent forecast. If the R^2 of a forecast of the 10-year average premium really were close to 1, the most recent forecast would not be the best one for setting the asset allocation for the next year. Investors could obtain a very precise estimate of next year's premium by comparing their 10-year forecast from nine years ago with the actual premiums from the intervening nine years. I am not aware of anyone who has so much faith in their model that they are willing to base their current allocation on a forecast from nine years ago.

Is the level of explanatory power necessary for a reasonable chance of success available to investors? While the answer to that question is well beyond the scope of this study, some perspective may be gained by looking at a single example. The cyclically adjusted price/earnings ratio (CAPE) of Campbell and Shiller (1998) has been advocated as a useful variable for predicting future stock returns.[3] Using S&P data for 1871–2013 from Robert Shiller's website to run regressions of real average returns on CAPE ratios, I find that the R^2 values for average 10-year and average five-year real returns are 32% and 18%, respectively. Based on the results in **Exhibit 1**, these levels of explanatory power are well below those required to give investors a reasonable probability of success. For example, consider the 32% R^2 for the 10-year average return. The dark shaded bar in **Exhibit 1** that is closest to 32% (the one for a 30% R^2) indicates a success rate of only 18%.

The results in **Exhibit 1** are specific to the US equity premium. The success rates for various R^2 levels will likely be different for other premiums. In general, the success rates for these types of simulations are driven by three variables—the accuracy of the forecast (R^2), the average premium, and the volatility of the premium. All else constant, the higher the average premium, the more explanatory power is required for success. Similarly, all else constant, the lower the volatility of the premium, the more explanatory power is required for success. A high average premium makes the cost of being wrong higher, and a low volatility makes the likelihood of a negative premium smaller.

The results of this study do not imply that investors should rebalance less frequently. The main purpose of rebalancing should be to manage the risk and enhance the diversification of a portfolio. This is likely best accomplished by rebalancing more frequently than once every 10 years.

Investors seeking to earn a 10-year return must hold their investments for 10 years. While this statement is painfully obvious, its implications for dynamic asset allocation are apparently not so obvious. Adjusting portfolio weights on an annual (or more frequent) basis in response to changes in long-run return forecasts is common advice in the financial press. The results of this study show that, unless those long-run forecasts are very precise, this is not the best approach for setting a portfolio's asset allocation. What seems to make more sense is the consistent practice of rebalancing to an allocation that is based on investor risk tolerance, goals, and needs, rather than a forecast of returns for the next several years.

2. The conclusions are not sensitive to the choice of weights for the asset allocation. Using 60/40, 70/30, 80/20, 90/10, and 100/0 for the asset weights all produce similar success frequencies as long as the benchmark weights are adjusted accordingly.
3. Campbell, John Y. and Robert J. Shiller. 1998. "Valuation Ratios and the Long-Run Stock Market Outlook." *Journal of Portfolio Management* 24: 11-26.

APPENDIX

Data Description

The equity premium used in this study is the return on US stocks in excess of the one-month US Treasury bill return. The monthly return on US stocks is the market cap-weighted return of all NYSE, NYSE MKT, and NASDAQ firms that are listed in the CRSP database with a share code of 10 or 11 (common stock) at the beginning of the month. Data for stocks are from the Center for Research in Security Prices (CRSP) at the University of Chicago Booth School of Business. The one-month Treasury bill rate is from Ibbotson Associates.

Data for regressions of real returns on CAPE ratios are from Robert Shiller's website: www.econ.yale.edu/~shiller/data.htm.

Specifying the Desired R^2

Define the following variables:

A is the actual average equity premium over the next 10 years (or five years).

F is the investor's forecast of this future average premium.

ε is a randomly generated error that is normally distributed with a mean of zero and a variance of 1.

σ^2 is the variance of the average premium A.

$\bar{A}$ is the overall average equity premium across all the bootstrapped samples.

The objective is to produce an unbiased forecast of the future average premium that has the desired level of explanatory power. Specifically, let R_D^2 be the desired level of explanatory power—how much of the variation in the average premium that the forecast can explain. Without showing all the details, a bit of algebra shows that the desired R^2 is achieved by the following forecast:

$$F = \bar{A}(1 - R_D^2) + R_D^2 A + \varepsilon[\sigma^2 R_D^2 (1 - R_D^2)]^{1/2}$$

Note that when R_D^2 is 1.0, $F = A$, so that the forecast of the long-run average is perfectly accurate. When R_D^2 is 0.0, the forecast is equal to $\bar{A}$, the average equity premium across all samples. Thus the forecast is unbiased, even when there is no explanatory power. In the case where R_D^2 is zero, the forecast remains constant, as does the investor's allocation between stocks and bills.