The Best Asset Allocation for Retirees

April 21, 2015

by David Blanchett

There has been growing interest recently in the optimal glide path shape for retirees. Previous research byPfau and Kitces (2014) noted that rising glide paths were optimal. My ownrecent research (2015), which incorporated today's low bond yields, noted declining glide paths were best. Even more recent research by Kitces and Pfau (2015), which incorporated market valuations (using the Shiller P/E metric), again confirmed that increasing glide paths are best.

All these smart people are reaching different conclusions. What gives?

It turns out that the optimal glide path depends significantly on the return assumptions, especially early in retirement. Using a model based entirely on long-term averages yields results that are very different than a model based on available returns today.

In this article, I determine the optimal glide path shape for retirees using varying initial bond yields and stock market valuations. Increasing glide paths perform best in moderate- and higher-return environments, while decreasing glide paths perform better in lower-return environments (especially when bond yields are low).

With retirees facing both low bond yields and a high market valuation, decreasing glide paths are slightly more optimal; however, the results do not vary that significantly across the glide paths (decreasing, increasing or constant).

Market conditions today

Retirees have to consider both the bond yields and stock market conditions at the onset of retirement. I use two key indicators to forecast the future returns for bonds and stocks: 10-year U.S. government bond yields and the Shiller CAPE ratio. Bond yields reliably predict the future returns of bonds because a majority of the return is the coupon, which is known at the time of purchase.

The CAPE ratio was introduced by Campbell and Shiller (1998) and is calculated by dividing the price of the S&P 500 by the average real earnings over the previous 10 years. Those researchers noted that this metric can explain about 30% of subsequent real stock returns over a 10-year horizon. Research by Davis, Aliaga-Díaz and Thomas (2012) also noted the historical predictive benefit of the CAPE ratio versus other common forecasting metrics. The historical relation between each of these metrics is included in Figure 1 for bonds and stocks, in Panels A and B, respectively.

Figure 1: Future Returns Based on Initial Conditions

Panel A: Bond Yields and Future Bond Returns   Panel B: The CAPE Ratio and Future Stock Returns

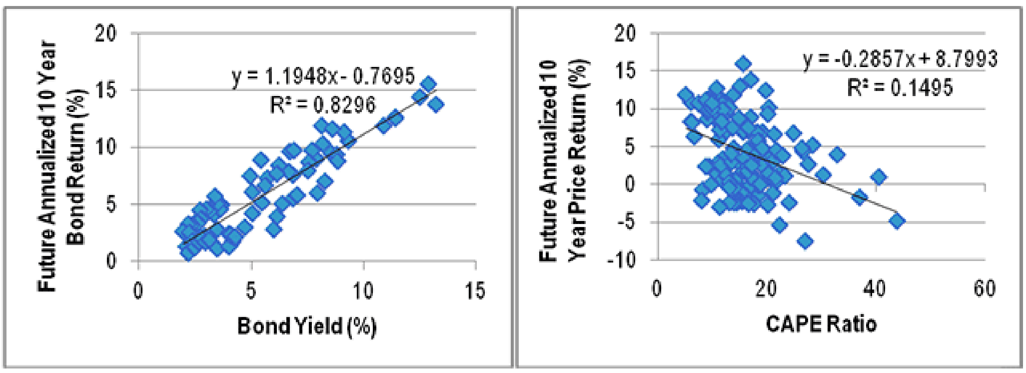

Bond yields and future bond returns are more highly correlated than the CAPE ratio and future stock returns. However, neither metric is favorable for retirees today, especially when compared to their long-term averages. Kitces (2009) noted the relation between safe-withdrawal rates (SWRs) based on various market valuation levels; individuals who retire in an overvalued market should withdraw less initially than retirees who retire in an undervalued market. Figure 2 provides some perspective on how CAPE ratios and bond yields have changed over time.

Figure 2: Historical Bond Yields and CAPE Ratios

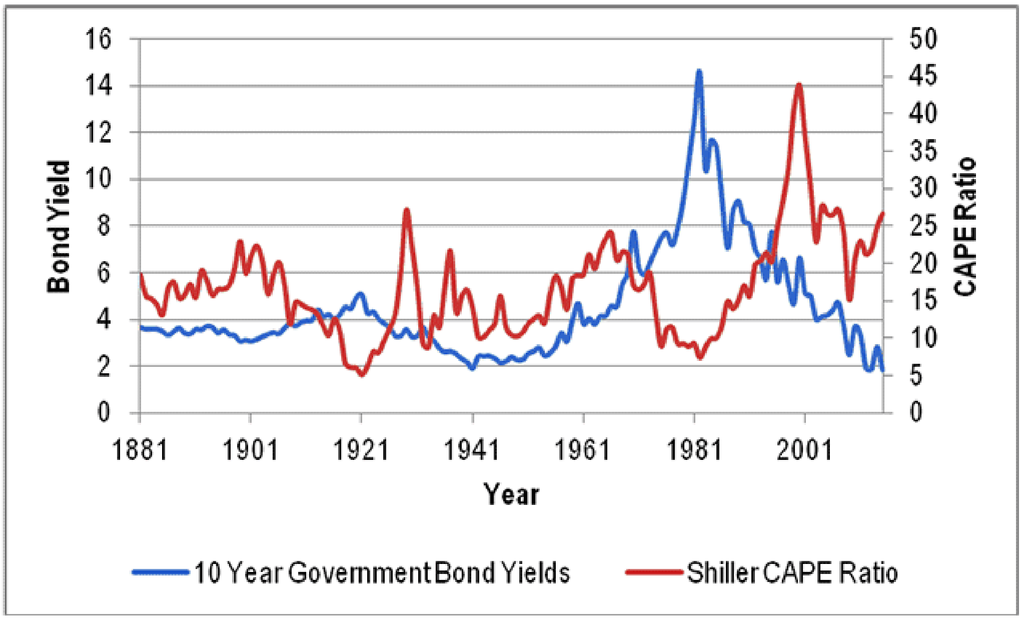

The correlation between bond yields and CAPE ratio has been low (-0.13) historically, and not significant even at the 10% level. Lower bond yields have generally been associated with higher CAPE ratios, but that the relation is incredibly weak.

As of February 1, 2015, the yield on 10-year Treasury bonds was 2.0% and the CAPE ratio was 27.85. Let me provide some perspective on how extreme the current values are. Only in four of the last 135 years (as of January 1) were yields lower than 2.0%, and only in six of the last 135 years was the CAPE ratio above 27. There has never been a time when yields have been this low and the CAPE ratio has been this high. Understanding where things are today is very important since today's conditions will have a significant impact on retiree outcomes.

Retirement researchers commonly use long-term historical return averages for SWR analysis. The most popular data is the Ibbotson time series data on stocks, bonds, bills and inflation (SBBI) data going back to 1926. While this historical series is extensive, its utility in the current market environment is limited. For example, the average historical yield on 10-year Treasury bonds since 1881 has been 4.5%. However, the yield on 10-year Treasury bonds was only 2.0% as of February 2015, which is less than half the long-term average. Therefore, an analysis based on long-term averages makes the implicit assumption that an investor can purchase a 10-year government bond yielding 4.5% despite the fact that this is impossible today.

The safety of an initial withdrawal rate, as well as conclusions about the optimal glide path, are heavily dependent on the return assumptions used in the model. For example, in a variety of past studies Blanchett, Finke and Pfau (*Journal of Financial Planning*, the *Journal of Wealth Management*[1] and the *Retirement Management Journal*) have explored how the safety of Bengen's (1994) initial "4% Rule" changes considerably when you use return expectations better calibrated to current market conditions. This is especially true because the returns experienced early in retirement have a disproportionate weight on the outcome of the retirement scenario, a concept commonly referred to as sequence-of-returns risk.

The returns model and its assumptions

In this section, I will briefly summarize the returns model. Specific equations are omitted, yet can be obtained here if you'd like more information on the exact model.

In a world of abundant and accurate data, it would not be necessary to create a model. However, there is relatively limited data that matches today's unique environment. For example, this is the only time since 1881 that the yield on 10-year government bonds is below 3.0% and the CAPE ratio is greater than 25. Therefore, the past is unlikely to provide us a large enough sample of return series for forecasts today.

Instead, I used a returns model that is an updated version of the approach I introduced with Finke and Pfau (2014) that we also updated recently for an *Advisor Perspectives* article. The returns model is based on the general relations noted in

Figure 1, whereby lower bond yields correlate to lower bond returns and higher CAPE ratios correlate to lower stock returns. The model begins with two separate assumptions: the interest rate and the CAPE ratio at retirement.

My base assumption is that the 10-year Treasury yield reverts to its long-term average (5%) over time. Therefore, if retirement begins in a low-yield environment (like today) the yield will increase each year. The low starting yield has a dual effect on the bond returns. The income return is lower and the price return (i.e., change in value of the portfolio) may be negative since the return on bonds is inversely correlated to changes in bond yields.

The average CAPE is assumed to be 17, which is close to the long-term average. With a CAPE value of 17 the arithmetic average return on stocks in the model is approximately 9% with an annualized standard deviation of 18%. The expected return on stocks is lower than the actual historical long-term return on stocks to be more consistent with equity returns in international markets. The long-term compounded average annual return of U.S. equities has been about 2.0 percentage points higher than the other 20 countries in the Dimson, Marsh and Staunton dataset (which goes back to 1900), and an adjustment is included to reflect that.

For the first 15 years of retirement, I assumed an inverse relationship between CAPE values and subsequent returns. For example, if the CAPE ratio is higher than average (i.e., above 17) the returns on equities will be lower than average, and vice versa. I assume there is no relation between the initial CAPE ratio and future equity returns after 15 years, based on historical evidence (i.e., the relation has almost entirely dissipated after 15 years).

Additionally, I included an investment fee of 50 basis points in the simulations to reflect the fact that it is impossible to invest for free. While it is possible to build a portfolio of index mutual funds or ETFs at a relatively low cost (e.g., 20 basis points or less) the average investor pays far more than this, either through actively managed investments, advisory fees or both.

Figure 3 provides some perspective on how key assumptions affect the probability of a portfolio being able to successfully fund a 4% initial withdrawal rate over a 30-year retirement period for varying equity allocations (from 0% to 100% equities in 10% increments). The initial withdrawal amount, which is 4% of the initial balance of the portfolio (e.g., $40,000 from a $1 million portfolio), is assumed to increase annually by inflation. The results in Figure 3 are based on a 10,000-run Monte Carlo simulation.

Figure 3: Success Rates and Model Adjustments for a 4% Initial Withdrawal Rate Over 30 Years

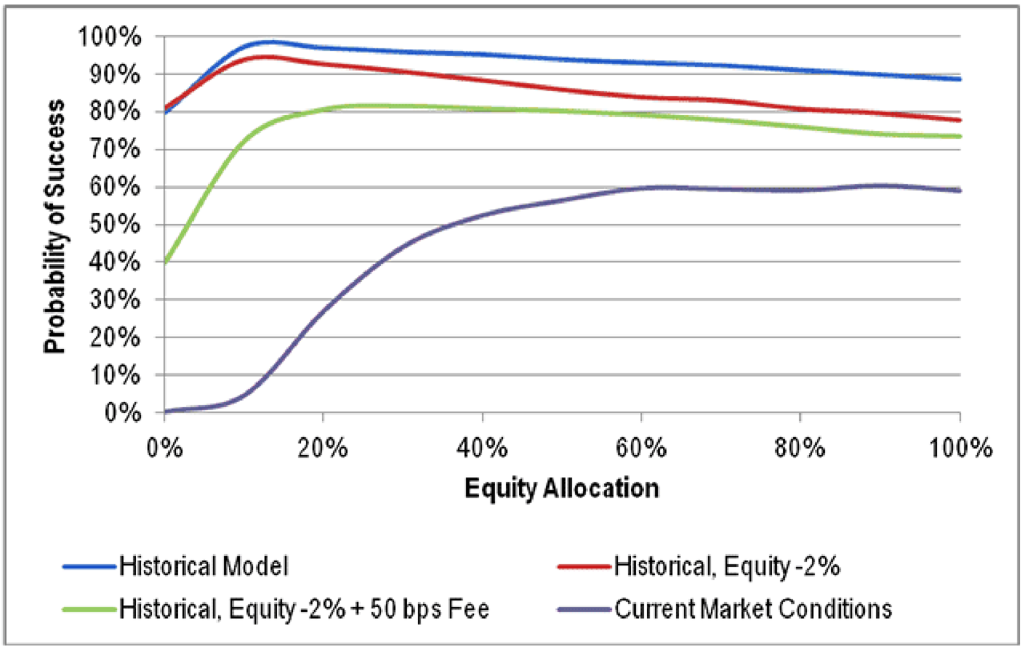

I made three key adjustments to the historical returns model represented by the top blue line in Figure 3: First, I reduced the rate of return for equities by 2% to better reflect the average historical return in equities internationally. Second, I included a 50 basis point fee for the portfolio. Third, I assumed that an environment similar to today, where the initial yield on bonds is 2.5% and the CAPE ratio is 27.

Analysis

I considered three primary glide paths: decreasing, constant and increasing. The decreasing glide path has an initial equity allocation of 60% that decreases by 1% each year during retirement. The constant glide path has a constant 45% equity allocation throughout retirement. The increasing glide path has an initial equity allocation of 30% that increases 1% per year in retirement.

The fixed-income portion of the portfolio is assumed to be 75% bonds and 25% bills. The results were not sensitive to this assumption and were effectively the same if these weights are reversed.

I tested three initial CAPE ratios (10, 17 and 27) and three bond yields (2.5%, 5.0% and 7.5%). The initial withdrawal rate is assumed to be 4.0%, and subsequent portfolio withdrawals are assumed to be increased annually by inflation (i.e., the "4% Rule"). The retirement period is assumed to be 30 years. Taxes are ignored.

Table 1 includes information about the optimal respective glide path shape for the nine different scenarios, as well as the respective probabilities of success for each.

T*able 1: The Optimal Glide Path Shape Assuming Various Initial Retirement Conditions*

Rank Order of Optimal Glide Path

		Yld Seeds		
		2.5%	5.0%	7.5%
CAPE Ratio	10	Dec/Cons/Inc	Inc/Cons/Dec	Inc/Cons/Dec
	17	Dec/Cons/Inc	Inc/Cons/Dec	Inc/Cons/Dec
	27	Dec/Cons/Inc	Inc/Cons/Dec	Inc/Cons/Dec

Probability of Success: Decreasing/Constant/Increasing

		Yld Seeds		
		2.5%	5.0%	7.5%
CAPE Ratio	10	84%/82%/79%	87%/89%/89%	90%/92%/95%
	17	74%/72%/69%	79%/81%/82%	83%/87%/90%
	27	56%/55%/52%	62%/65%/67%	67%/74%/81%

The results in the table demonstrate a few things. The increasing glide path was optimal for moderate- or high-initial yields, while the decreasing glide path was optimal for the low-yield scenarios. If you focus on both "moderate" scenarios (yield of 5.0% and initial CAPE ratio of 17) the increasing glide path is best (with an 82% success rate), while the decreasing glide path is worst (79%). The return assumptions really matter. Depending on whether you assume initial bond yields are low or high significantly impacts which glide path is optimal. Finally, the initial yield impacts the results more than the initial CAPE ratio.

Regardless of the scenario, however, the results weren't that different. For example, among the nine scenarios, the one that is most similar to conditions today had an initial yield of 2.5% and a CAPE ratio of 27. The best glide path is decreasing, followed by constant, and increasing is the worst, with probabilities of success of 56%, 55%, and 52%, respectively. The difference between decreasing and increasing glide path may seem material (an increase of 4%), but it really isn't. Although the differences will be statistically significant if one were to run enough simulations, they are not economically significant, i.e., there is not a material change to the outcome for a retiree.

Additional considerations

This article raises a number of important implications for planners. As previously found in a number of research projects by Finke, Pfau and myself, using a model based on expected returns results in a very different estimate of the probability of success for a 4% initial withdrawal rate for a 30-year retirement period compared with a model based on historical returns. This doesn't mean that a 4% or higher initial withdrawal rate can't still work given the nuances associated with implementation. For example, most retirees will not need to spend the same amount each year (in real terms) during retirement; retirees can adjust consumption based on portfolio performance and decrease spending as they experience physical and mental limitations. Also, probability of success is only one way to measure an outcome for a retiree – it doesn't incorporate the magnitude of failure.

The best glide path is heavily dependent on model assumptions, such as how long the retiree is going to live as well as the returns during retirement. Using long-term average bond and equity valuations suggests that an increasing glide path is best, but a model that is better calibrated to today's environment suggests a decreasing glide path is best for investors living at least 30 years past retirement. Overall, though, the differences among glide paths with similar lifetime risk levels aren't that different and do not significantly change the economic outcomes for a retiree.

One impact variable that is not considered in this analysis is risk tolerance (or risk preference). A retiree's risk tolerance tends to decrease as he or she ages (i.e., individuals prefer more conservative portfolios). It is difficult to incorporate risk tolerance into common retirement income planning models because they are largely focused on accomplishing a goal. However, if risk preference were directly incorporated into the models it would likely provide an additional "nudge," making decreasing glide paths more optimal than their more aggressive alternatives. This would be especially true if taking additional risk can't be shown to produce more secure outcomes.

David M. Blanchett, CFA, CFP® is the head of retirement research for Morningstar Investment Management in Chicago, IL.

---

1. See the working paper version here

