How Variable Withdrawals Improve Retirement Outcomes

October 17, 2016

by Joe Tomlinson

Variable withdrawal strategies for retirement spending are receiving more attention from researchers and practitioners. Unlike fixed strategies, where the withdrawal plan is set at the beginning of retirement, variable strategies adjust spending over the course of retirement to reflect investment experience.

Optimal asset allocations for variable withdrawal strategies are quite different from the research findings and rules of thumb based on fixed strategies. Indeed, the implications go beyond asset allocation and show, for example, that equity glide paths in retirement are relatively unimportant.

Asset allocation in practice

Perhaps the best known rule of thumb in financial planning is the "100 minus age" stock allocation recommendation where a 65-year-old would hold a 35/65 stock/bond mix and the stock allocation would be reduced over the course of retirement. Given longer lifespans, some have argued that the 100 should be replaced by 110 or 120, which would raise the initial stock allocation as high as 55%. Charles Schwab recommends a 60% stock allocation from ages 60 to 69, 40% for ages 70 to 79 and 20% for ages over 80. My test of a Vanguard calculator based on a moderate-risk appetite showed that it recommends a 40% stock allocation at the beginning of retirement, and the Vanguard Target Retirement Income Fund (VTINX) uses a 30% stock allocation.

Research findings have favored somewhat higher stock allocations than those. Bill Bengen's classic 1994article that introduced the 4% rule also contained an asset allocation recommendation. He didn't recommend a specific stock/bond mix, but demonstrated that stock allocations in the 50% to 75% range worked best to eliminate plan failures with a 4% inflation-adjusted withdrawal rate. There have been lots of articles since, some by Bengen, that have refined his original research, and, in recent years, we have seen articles challenging the safety of the 4% rule in a lower return environment. There has also been much research, discussion and debate about the allocation to equities over the course of retirement – the equity-glide path – and whether it should decrease, increase or remain level. Another topic that has received attention is the extent to which asset allocations should reflect market valuation measures such as the CAPE ratio. All told, the range of recommended retirement stock allocations has been 30% to 75%.

With this discussion as a backdrop, I'll introduce variable withdrawal strategies and then develop asset allocation modeling to demonstrate how the optimal strategies with variable withdrawals differ significantly from the 30% to 75% range.

Introducing variable withdrawals

Variable withdrawal methods share the common characteristic that withdrawals adjust each year during retirement to reflect underlying investment performance, whereas, with the 4% rule and other fixed methods, the withdrawals are mapped out at the start of retirement. With fixed withdrawals, bad investment experience or longer-than-expected retirements can result in savings being depleted (plan failure), while better-than-expected investment experience will flow through to larger bequests. Although the 4% rule has been popular with researchers, practitioners recognize that good planning cannot ignore what is happening with the investment portfolio; spending adjustments will need to be made during retirement. One might determine these adjustments informally or apply a more disciplined variable withdrawal strategy.

I'll develop a hypothetical client example to demonstrate how outcomes using the 4% rule differ from those based on a variable withdrawal strategy. Then I'll turn to optimizing the asset allocation to support the variable strategy. This modeling will be based on a 65-year-old retired female with a remaining life expectancy of 25 years and $1 million in savings that she can dedicate to generating retirement income. Her basic living expenses are $50,000 per year, increasing with inflation, and she will receive an inflation-adjusted $30,000 annually from Social Security. She will utilize withdrawals from savings to cover the gap between basic living expenses and guaranteed lifetime income and take additional withdrawals for discretionary spending. Her main goal is to generate sustainable retirement income; leaving a bequest is of secondary importance. The analysis will be pre-tax.

The particular variable withdrawal approach I will use is described in detail in the Appendix, but it basically takes the portfolio at the start of each year during retirement and determines the withdrawal amount for that year that produces level real withdrawals over the remainder of retirement. Because of investment volatility, the annual resetting of withdrawals will cause year-to-year changes in retirement consumption.

The chart below shows the 4% rule versus variable withdrawals based on various retirement outcome measures, where detailed definitions can be found in the Appendix. I've assumed a 60/40 stock/bond allocation for both the 4% rule and variable withdrawals.

Projected outcomes -- Variable withdrawals versus the 4% rule

Performance measure	4% rule	Variable withdrawals
Average consumption	$67,500	$75,200
Average bequest	$647,000	$346,000
Failure percentage	26.3%	14.0%
Consumption change %	0.57%	6.09%

Source: Author's calculations

Compared to the 4% rule, the variable approach lowers bequests and increases consumption but is able to do so while actually reducing plan failures. This is a double improvement over the 4% rule for a client focusing mainly on the level of spending and wishing to avoid plan failure. The downside of the variable approach is more volatile consumption; this is shown by the consumption change percent, which measures the average absolute annual change in spending. The main takeaway from this illustration is that variable withdrawal strategies offer strong advantages that make them worth considering.

For those interested in learning more about different variable strategies, I recommend this 2015 paper by Wade Pfau. Pfau recently wrote this introductory post to begin a series he is writing about variable strategies. I also recommend this *Advisor Perspectives* article by Larry Siegel, which provides an in-depth presentation of a particular variable strategy he developed with Barton Waring. I've written two previous *Advisor Perspectives* articles on the advantages of variable withdrawals here and here.

Optimal asset allocation

The next task will be to determine which asset allocation produces the best retirement outcomes. To do this, I'll introduce two additional performance measures: average shortfall and certainty equivalent (CE) consumption, which are described in detail in the Appendix.

The average shortfall is a measure that combines both the probability and magnitude of plan failure. Magnitude is important because a retirement plan that fails 10 years before death is clearly a worse result than one that fails a year before death. However, the commonly used failure percentage doesn't recognize this difference. In the displays, I'll continue to show failure percentages because of the popularity of this measure but will focus attention on average shortfall.

The certainty equivalent (CE) consumption combines the measures of average consumption and consumption change percent into a single measure. I will show that increasing stock allocations with variable strategies typically increases average consumption but also increases the consumption change percentage. So we need a measure to compare high consumption that bounces around a lot from year to year with lower, steadier consumption. For this we'll apply economic utility analysis where the CE is level amount of consumption an individual would be willing to accept in exchange for variable consumption. The CE determination will depend on aversion to variable consumption (referred to by economists as risk aversion), which will vary by individual. This particular analysis is based on an individual with moderate risk aversion who would accept a CE of $66,600 as a substitute for consumption that bounced around randomly between $60,000 and $80,000.

The chart below shows outcomes for different stock allocations.

Projected outcomes--Variable withdrawals with different stock allocations

Performance measure	50% stocks	70% stocks	90% stocks	110% stocks
Average consumption	$72,500	$77,700	$84,200	$91,800
Average bequest	$322,000	$378,000	$450,000	$535,000
Failure percentage	14.5%	14.2%	14.8%	15.8%
Average shortfall	**-$9,300**	-$10,300	-$13,663	-$18,000
Consumption change %	5.11%	7.10%	9.35%	11.89%
Average CE consumption	$68,300	$70,800	$72,300	**$73,200**

Source: Author's calculations

The key measures here are the average shortfall and the average CE consumption, and we can see that there is a tradeoff: as the stock allocation increases, the CE improves but the shortfall gets worse. One might argue that the CE should be the primary measure because it takes bad outcomes (low consumption) into account; the counterargument is that not being able to pay for basic living expenses is an especially bad outcome not fully captured in the CE measure. The chart limits the range of stock allocations from 50% to 110% because separate testing showed that both shortfalls and CEs got worse below 50% and above 110%. The 110% is a leveraged stock allocation assuming borrowing at the bond return, but could also be thought of as an unleveraged allocation to an equity asset class with higher expected returns and higher risk.

The bottom line result is that, with variable withdrawals, we move from 30% to 75% stock allocation recommendation to a range of 50% to 110% – a significant increase in the upper end of the range. My modeling assumes a real arithmetic average stock return of 5% versus a 1926-2015 average of 8.92% and my assumed equity return premium of 5% is also lower than the historical 6.53%. Switching to variable withdrawals increases the upper end of the stock allocation range even with my assumption of lower-than-historical absolute and relative equity returns.

Including a SPIA

We can improve results further by using a portion of savings to purchase a SPIA and fill in the income gap between the $30,000 of Social Security and $50,000 of essential expenses. The chart below compares results with a SPIA to the best results from the prior chart.

Impact of SPIA purchase	without SPIA		With SPIA	
Performance measure	50% stocks	110% stocks	100% stocks	140% stocks
Average consumption	$72,500	$91,800	$79,200	$88,800
Average bequest	$322,000	$535,000	$256,000	$375,000
Failure percentage	14.5%	15.8%	0.0%	0.0%
Average shortfall	**-$9,300**	-$18,000	$0	**$0**
Consumption change %	5.11%	11.89%	6.02%	10.03%
Average CE consumption	$68,300	**$73,200**	$74,200	**$75,700**

Source: Author's calculations

The first two columns are from the 50% stock allocation that produced the lowest average shortfall, and the 110% stock allocation that produced the highest CE. Columns 3 and 4 are based on spending $490,000 to purchase an inflation-adjusted single-premium immediate annuity (SPIA) paying $20,000 per year initially with payments increasing at the inflation rate each year thereafter. (Many insurers do not offer inflation-adjusted SPIAs, but quotes can be obtained from Vanguard.) This amount of SPIA purchase completely eliminates the shortfall risk, so we can concentrate on the CE as the primary performance measure. Testing reveals that a 140% stock allocation produces the highest CE, and I also show the results for a 100% stock allocation. It is clear that adding a SPIA improves things, both in terms of eliminating shortfall risk and increasing the CE.

Additional considerations

Sensitivity testing: It's worth examining the extent to which the high stock allocations being recommended is a function of the particular modeling assumptions. I have run tests where I have lowered the real average stock return from 5% to 3% and raised the risk aversion coefficient from 5 to 20. This latter change goes from an individual who would accept a level $66,600 of consumption as a substitute for random bouncing between $60,000 and $80,000 (CE = 5) to one who would accept a level $62,200 (CE = 20). With those changes, our optimal case result that includes SPIA purchase goes from a 140% stock allocation to 60%, but there is very little difference in the CEs for stock allocations between 50% and 90%. So even with extremely conservative assumptions, results are still acceptable for high stock allocations.

Glide paths: There have been debates going back more than 50 years about how stock/bond mixes should change over the course of retirement. In the past few years, researchers Michael Kitces, Wade Pfau, Javier Estrada, Luke Delorme and David Blanchett have each produced papers or articles on this subject, reaching varying conclusions. Level, decreasing and increasing equity glide paths have been shown to be optimal depending on the particular methods used in the studies. Glide paths are an interesting topic when optimal stock allocations are in the 30% to 75% range as has been the case for research based on fixed withdrawal strategies. But when financially optimal results are produced with stock allocations above 100%, as I have demonstrated for variable withdrawals, glide-path analysis has little practical significance. Conveniently, with variable strategies, we don't really need to sort out the conflicting evidence about the best choice of a retirement glide path.

Final thoughts

My findings that stock allocations of 100% or more make sense for retirement may understandably provoke the reaction that there is something wrong. However, in the examples I've shown, sources other than the investment portfolio are generating a high percentage of consumption. For example, with a SPIA purchase and 100% stocks for the remaining investments shown in the final chart, average consumption is $79,200, and $50,000 (or 63%) of that is coming from Social Security and the SPIA. In terms of total retirement income, the stock allocation is only 37%. Once a retiree secures funding for essential spending needs, the remaining assets are "liberated" and can be invested more aggressively.

This view supports my recent *Advisor Perspectives* article that pointed out how a narrow focus on the investment portfolio distorts retirement planning. Taking a broader view combined with a greater focus on variable withdrawal strategies and their asset allocation implications moves retirement planning in a positive direction.

Joe Tomlinson, an actuary and financial planner, is managing director of Tomlinson Financial Planning, LLC in Greenville, Maine. Most of his current work involves research and writing on financial planning and investment topics.

Appendix

Investment assumptions: For the Monte Carlo analysis, stocks are assumed to earn an arithmetic average real return of 5% with a 20% standard deviation, and bonds (TIPS) are assumed to earn 0% with a 5.5% standard deviation. These returns are significantly lower than historical averages, reflecting current interest rates and a lower-than-historical equity risk premium. Allocations are rebalanced annually to maintain the initial allocation.

Variable withdrawal approach: This approach can be described based on the Excel PMT function. The maximum allowable withdrawal is recalculated each year as a function of the current portfolio balance, estimated remaining longevity and expected investment returns. The estimated remaining longevity is updated each year based on a Gompertz mortality function calibrated to a life expectancy of age 90 for a 65-year-old. For a return assumption, I use my estimated TIPS yield of 0%, to be conservative, even though my assumed portfolios contain both stocks and TIPS. The key difference between this approach and traditional approaches like the 4% rule is this variable method adjusts withdrawals each year to reflect emerging investment experience, whereas traditional approaches determine the pattern of future withdrawals at the start of retirement.

Modeling methodology: For each retirement income approach that I model, I generate 10,000 simulated retirements. Withdrawals each year are determined based on the particular retirement income approach being modeled. Investment returns for each year are generated randomly based on average return and standard deviation characteristics. The date of death for each of the 10,000 simulations is randomly determined based on the Gompertz mortality function.

Although methods that determine withdrawals as a percentage of the current portfolio balance will never completely deplete savings, I use a slightly different approach. If calculated consumption is not sufficient to cover basic living expenses but there are remaining savings, I take from the savings to cover the expense gap until savings are depleted. This way I can determine the percentage of simulations where savings are depleted.

Performance measures

Average consumption: Consumption equals Social Security income ($30,000) plus annual withdrawals. For each of the 10,000 simulations, I compute the lifetime average of the annual consumption amounts and then take the average of these averages.

Average bequest: Each of the 10,000 simulations produces a remaining savings at time of death, which can be zero or a positive amount. I calculate the average of the 10,000 bequest amounts.

Failure percentage: This represents the percentage of the 10,000 simulations where savings are insufficient to fully pay for the assumed $50,000 of basic living expenses.

Average shortfall: For each simulation that "fails," I calculate the amount of additional funds that would have been needed to pay for basic living expenses until the end of life. The sum of these amounts is divided by 10,000 to determine an average shortfall for all the simulations (including those with zero shortfall). This is a more useful failure measure than failure percentage because it incorporates both frequency and magnitude, but I also show failure percentage because it is a more commonly used measure.

Consumption change percent: This measures the absolute value of the change in consumption from one year to the next, averaging these for each of the 10,000 simulations, and then averaging these averages.

Average CE consumption: This certainty equivalent (CE) measure is based on an economic utility calculation that converts variable year-by-year consumption into a level amount that the recipient would view as equivalent. The CE amount depends on what economists refer to as the recipient's level of risk aversion. For example, if annual consumption bounced around randomly between $60,000 and $80,000, an individual with low risk aversion would demand close to $70,000 if offered a trade to level consumption. A highly risk averse individual would be willing to accept an amount closer to $60,000.

For this analysis, I have assumed a medium aversion to variable consumption – an individual would be willing to accept annual consumption of a level $66,600 in trade for consumption that bounced around randomly between $60,000 and $80,000. This translates to a risk aversion coefficient of 5 based on a CRRA utility function of the form U = (1/(1-RA))*C^(1-RA) to convert consumption into utility. For each of the 10,000 Monte Carlo iterations, I convert each year's consumption into utility, average the utilities based on the number of years in each iteration and convert to a CE using the inverse of the utility function.

