

An Efficient Frontier for Retirement Income

by

Wade D. Pfau

Associate Professor

National Graduate Institute for Policy Studies (GRIPS)

7-22-1 Roppongi, Minato-ku, Tokyo 106-8677 Japan

Email: wadepfau@gmail.com

phone: 81-3-6439-6225

Abstract

This article outlines a different way to think about building a retirement income strategy, which moves dramatically away from the concepts of safe withdrawal rates and failure rates. The focus is how to best meet two competing financial objectives for retirement: satisfying spending goals and preserving financial assets. The process described here focuses on allocating assets between a portfolio of stocks and bonds, inflation-adjusted and fixed single-premium immediate annuities (SPIAs), and immediate variable annuities with guaranteed living benefit riders (VA/GLWBs). This process incorporates unique client circumstances, bases asset return assumptions on current market conditions, uses a consistent fee structure for a fair comparison between income tools, operationalizes the concept of diminishing returns from spending by incorporating a minimum needs threshold and a lifestyle spending goal, uses survival probabilities to calculate outcomes, and incorporates client preferences to balance the competing financial objectives for the final choice among the collection of allocations that define the efficient frontier for retirement income. Results are presented for a 65-year old couple whose lifestyle needs require a 4% inflation-adjusted withdrawal rate from retirement date assets. Their efficient frontier generally consists of combinations of stocks and fixed SPIAs. Perhaps surprisingly, bonds, inflation-adjusted SPIAs, and VA/GLWBs do not serve a useful role in the couple's optimal retirement income portfolio.

JEL Codes: C15, D14, G11, G17

Keywords: retirement planning, retirement income modeling, efficient frontier, annuities, systematic withdrawals

Acknowledgements: Many individuals provided me with useful feedback and discussions about the topics in this article. In particular, I wish to thank William Bernstein, Jason Branning, Michael Finke, Francois Gadenne, Michael Kitces, David Littell, Manish Malhotra, Dick Purcell, Bob Seawright, Joseph Tomlinson, and Michael Zwecher. I am grateful for financial support from the Japan Society for the Promotion of Science Grants-in-Aid for Young Scientists (B) #23730272.

Introduction

William Bengen's seminal 1994 article on sustainable withdrawal rates in the *Journal of Financial Planning* provided a much needed reality check on popular discourse by introducing how sequence of returns risk causes the sustainable withdrawal rate from a portfolio of volatile assets to fall below the average return to those assets. Bengen described the SAFEMAX, which he defined as the sustainable withdrawal rate from the worst-case scenario in history. It was closer to 4% than to numbers like 7% bandied about in the media at that time. Bengen's research answered an important question about sustainable spending rates. Several years later, Cooley, Hubbard, and Walz (1998) published a study popularly known as the Trinity study. It introduced a minor but significant modification to Bengen's work. Rather than reporting the historical worst-case scenario, the Trinity authors calculated success rates and the corresponding failure rates for different withdrawal rate and asset allocation strategies over differing retirement durations. Based on the U.S. historical data since 1926, success rates are the percentage of rolling historical periods in which some financial wealth remained at the end. Financial wealth depletion becomes synonymous with a failed retirement in this framework, as seen, for instance, when Terry (2003) wrote, "I believe that most investors would find even a one percent probability of failure to be excessively high when dealing with irreplaceable assets and considering the extreme costs of failure." Perhaps it was an unintended consequence, but the research on safe withdrawal rates (SWR) which has followed from these notions provides a rather incomplete picture of retirement income.

The idea that retirees should focus on finding a spending strategy that maintains a rather low failure rate using a diversified portfolio over a fixed retirement period, as is par for the course with safe withdrawal rate studies, is not adequate. Problems abound, including:

1. Failure rates do not consider the retiree's entire balance sheet of assets for income generation. With Social Security and other defined-benefit pensions, some clients may find that financial asset depletion is not disastrous.
2. Failure rates are not compatible with financial products designed to provide lifetime retirement income, such as single premium immediate annuities (SPIAs). Such annuities are assets and the present value of their remaining lifetime income payments can be included on a personal balance sheet, but this value is not a part of financial assets. Failure rates are not meaningful for clients seeking to understand the implications of partial annuitization.
3. Failure rates ignore the lost potential enjoyment from spending more early in retirement. They are an extreme outcome measure which puts weight only on financial wealth depletion. Client spending potential is irrelevant. Clients must find an appropriate personal balance between the aims of spending more and then having to make potentially larger subsequent cutbacks in the event of a long life and a sequence of poor market returns. According to Milevsky and Huang (2011), the decision about this tradeoff depends both on longevity risk aversion (the fear of outliving financial assets) and the amount of "pensionized" income from outside the financial portfolio. Working toward a similar end, Finke, Pfau, and Williams (2012) show how clients may potentially tolerate a higher failure rate in order to maximize their spending power and overall lifetime enjoyment in retirement.
4. Failure rates ignore the magnitude and severity of the "failure" condition. For how long and by how much will a client's income fall short of what is desired or needed? Failure is failure regardless of whether spending falls short by \$1 or by \$100,000.

5. The failure rate framework is based too heavily on the U.S. historical record since 1926 without proper consideration that things may change and current conditions may suggest a different starting point for today's retirees. Failure rates seemingly justify the use of historical averages for Monte Carlo simulations rather than current market conditions. Why else would it be relevant to a client today if a 5% withdrawal rate worked for someone retiring in 1947, for instance?

Planners and clients need to think more broadly beyond failure rates when developing their retirement income strategies. The objective here is to outline characteristics of a broader retirement income framework and to provide an example of its use with an efficient frontier for retirement income.

Literature Review

Tresidder (2012) provides an overview of the history of safe withdrawal rate research, in which he classifies three generations of past studies. The first generation simply assumes an average annual portfolio return, providing withdrawal rate guidance based on that. The second generation, developed by Bill Bengen, incorporates the sequence-of-returns risk associated with volatile portfolios and finds that historically a more conservative 4% withdrawal rate is much closer to being safe. Third generation researchers generally seek to incorporate more realism into second generation results by including factors such as market conditions at the retirement date, fees, longevity beyond 30 years, and an acknowledgment that the post-1926 US historical data is not sufficient to provide much confidence about the viability of the 4% rule. The third generation clarifies how 4% is only an educated guess based on limited historical data and some rather simplifying assumptions. There are other ways to view and interpret the historical data, and retirees only have one opportunity to get things right in retirement.

A growing body of research now exists which could be considered as the fourth generation of studies. This research broadens the safe withdrawal rate question to place it into the wider context of a complete retirement income strategy that can also include forms of annuitization. Ameriks, Veres, and Warshawsky (2001) was one of the first studies to investigate the beneficial impact of partial annuitization with a fixed SPIA on portfolio sustainability. Chen and Milevsky (2003) also examine the optimal allocation between immediate annuities and other financial assets. Milevsky (2009) further develops his product allocation framework. He builds a retirement income frontier in which users can aim to maximize spending power for an acceptable level of retirement sustainability and expected bequests by allocating between a total returns based portfolio for systematic withdrawals, a fixed SPIA, and a VA/GLWB. For guaranteed income sources, Milevsky creates the Retirement Sustainability Quotient (RSQ) as the portfolio success rate times the fraction of income taken from a volatile portfolio plus the fraction of annuitized income. This helps account for the availability of guaranteed income sources even after the portfolio is depleted. He defines "Financial Legacy Value" (FLV) as the average discounted bequest value at death for the strategy across the Monte Carlo simulations.

Huang, Grove, and Taylor (2012) develop their own version of Milevsky's product allocation with an efficient income frontier for a 65-year old male seeking to use financial assets to obtain a constant inflation-adjusted retirement income equal to 4% of retirement date assets. They also simulate various combinations of systematic withdrawals from stocks and bonds, a fixed SPIA, and a deferred VA with a guarantee rider. They plot "income risk" (the probability of financial wealth depletion at age 92 – which the 65 year old has a 25% chance to survive to) against "legacy potential" (the median amount of remaining wealth at age 86 – which is the life expectancy for the 65 year old male). Income risk and legacy potential provide the risk and return measures that allow

them to create a corresponding efficient frontier as found in Modern Portfolio Theory, but for the case of lifelong retirement income rather than single period returns. They emphasize that partial annuitization can potentially reduce income risk while also raising the median legacy value.

Several other studies from 2012 also investigate questions about product allocation by simulating the performance of different income tools using a variety of other outcome measures. For instance, Tomlinson (2012) uses a loss aversion utility function to compare spending shortfalls and bequest values for an inflation-adjusted SPIA, systematic withdrawals from portfolios of stocks and bonds, and partial annuitization approaches. He finds that depending on how a client views the tradeoff between bequests and spending shortfalls, optimal income strategies side toward either all stocks or all SPIAs.

Malhotra (2012) also seeks to develop general guidelines for a comprehensive framework to evaluate retirement income strategies. He defines two reward and three risk metrics, which are summary measures to analyze different income strategies. Reward measures include average income and average legacy, while risk measures include the probability of failure, the magnitude of failure, and the percentage of income from fixed sources. He incorporates several case studies to show how clients can evaluate the tradeoffs and seek solutions. His framework already incorporates the Social Security claiming decision, bond ladders, and time segmentation.

Finally, Pfau (2012) builds a framework to analyze eight different retirement income strategies, including constant and variable withdrawal rate strategies, SPIAs, and VA/GLWBs. Outcome measures emphasize downside risk, upside potential, and bequests, and the entire distribution of outcomes is shown. As each strategy is simulated in isolation, issues of product allocation are not directly explored.

Methodology

Retirement Income Planning Process

Retirement income planning is a process which should be agnostic with respect to the choice of products and withdrawal approaches. The performance of various combinations of income tools can be simulated and optimized to individual client circumstances.

Product Allocation

Though not comprehensive, retirement income tools considered here will follow the product allocation framework defined by Moshe Milevsky. Clients may rely on systematic withdrawals from a portfolio of stocks and/or bonds invested with a total returns perspective, or they may allocate some of their financial assets at retirement to purchase inflation-adjusted SPIAs, SPIAs with fixed nominal payments, or a VA with a GLWB rider.

Consistent Fees

It is important to maintain a consistent fee structure among their various retirement income tools in order to avoid biasing results, as the compounding effects of fees over time can be very dramatic and create large biases about optimal strategies.

The analysis here is based on realistic low-cost versions for the various income tools as are available in the marketplace. For systematic withdrawals, I assume that investors use low-cost index funds, and the stock and bond funds are assumed to have a 0.2% annual fee. SPIA prices are from Vernon (2012), who obtained them using the Income Solutions platform at the start of April 2012. The 65-year old couple can purchase a 100% joint-and-survivors SPIA with a payout rate of

3.875% for the inflation-adjusted version and 5.84% for the fixed version. Pricing aspects of the VA/GLWB include a payout rate of 4.5%, annual fees on the VA contract value of 0.6%, an annual guarantee rider fee of 0.95% on the high-watermark benefit base, and an annual step-up feature to increase payouts if the contract value reaches a new high watermark. The assumed asset allocation for the VA/GLWB is 70% stocks and 30% bonds. Using low-cost versions for each tool allows for more direct and meaningful comparisons. Further advisory fees could also be added if appropriate.

Current Market Conditions and Capital Market Expectations

A limitation which can be found in a surprising number of existing studies that analyze both annuity products and systematic withdrawals is that annuity prices are based on what are currently available in the market (and thus based on current market conditions), while the assumptions for systematic withdrawals from a portfolio of stocks and bonds are based on the much more optimistic historical average returns. If the historical averages could be expected to repeat in the future, then insurance companies would be able to offer higher payouts on their guaranteed products. But the reality is that interest rates are currently much lower than historical averages, suggesting lower long-term returns as well. The assumptions for each income tool/strategy must be comparable, or otherwise the simulated asset returns will lead to apples and oranges comparisons.

Table 1: Assumptions for Real Asset Returns

	Arithmetic Mean	Geometric Mean	Standard Deviation	Correlation Coefficients		
				Stocks	Bonds	Inflation
Stocks	5.1%	3.1%	20.0%	1	0.1	-0.2
Bonds	0.3%	0.1%	7.0%	0.1	1	-0.6
Inflation	2.1%	2.0%	4.2%	-0.2	-0.6	1
Equity Premium	4.8%					

Note: Standard deviations and correlation coefficients are based on *Stocks, Bonds, Bills, and Inflation* data provided by Morningstar and Ibbotson Associates, in which the U.S. S&P 500 index represents the stock market and intermediate-term U.S. government bonds represent the bond market. The arithmetic mean for bond returns is calibrated to recent TIPS yields. The arithmetic mean for inflation is based on the breakeven inflation rate implied by TIPS and Treasury yields. The arithmetic mean for stock returns is calibrated to allow an equity premium of 4.8% above the bond return, which is the equity premium for a GDP-weighted portfolio of 19 developed market countries between 1900 and 2010 from the Dimson, Marsh, and Staunton Global Returns Dataset provided by Morningstar and Ibbotson Associates.


Table 1 provides the assumptions for asset markets which guide the underlying 200 Monte Carlo simulations used for each product allocation strategy. While standard deviations and correlations are calibrated to the U.S. historical data since 1926, the return assumptions are connected to current market conditions rather than historical averages. The arithmetic mean for bond returns is calibrated to recent TIPS yields. The arithmetic mean for inflation is based on the breakeven inflation rate implied by TIPS and Treasury yields. The arithmetic mean for stock returns is calibrated to allow an equity premium of 4.8% above the bond return, which is the equity premium for a GDP-weighted portfolio of 19 developed market countries between 1900 and 2010 from the Dimson, Marsh, and Staunton Global Returns Dataset. The assumed distribution for asset returns is a multivariate

lognormal distribution which incorporates the means, standard deviations, and correlations for stocks, bonds, and inflation.

Retirement Financial Goals

As argued, retirees should not be narrowly focused on a singular goal to avoid financial wealth depletion. Financial goals for retirement can essentially be reduced to two competing objectives: to support minimum spending needs and lifestyle spending goals as best as possible, however high they may be, and to maintain a reserve of financial assets to support risk management objectives such as protecting from expensive health shocks, divorce, unexpected needs of other family members, severe economic downturns, etc., or to otherwise provide a legacy. To the extent that greater spending means using more financial resources, and to the extent that more aggressive investment strategies imply greater upside wealth as well as greater downside spending risks, these objectives must generally be balanced in a manner that best supports the clients preferences.

Utility measures translate spending into values which incorporate the concept of diminishing marginal enjoyment from spending increases. But defining an appropriate utility function for clients may be a very difficult task. As an alternative here to operationalize how well a strategy meets spending goals, I propose summing the amounts of spending shortfalls created by a strategy. With Monte Carlo simulations for financial market returns, each strategy will have a distribution of spending shortfalls. Figure 1 shows how the shortfall calculation can be made for a single hypothetical spending pattern in one simulation. The focus is on when the spending level is forced below minimum needs. This shortfall area can be summed and then divided by the total lifetime minimum needs. With this information, we can calculate, in turn, the percentage of lifetime minimum needs which are met by the strategy. This measure provides the magnitude of shortfalls relative to a spending level the client views as essential. For clients, the minimum needs line could be less than the lifestyle goal, as shown in Figure 1, or it could be at the same level.



Academics thinking in terms of lifecycle finance theory will aim to make a clear distinction between minimum needs and lifestyle goals. Clients may be unhappy if an event such as a bad sequence of markets returns makes their lifestyle goals unsustainable. But what could be truly disastrous is if their spending is forced to fall below a minimum level to meet basic needs.

Planners with practical client experience, though, may be more inclined to focus on spending below the lifestyle goal instead, as spending less than this may be viewed by the client as failure regardless of any hypothetical minimum needs. For examples of the planner's perspective, see Evensky, Horan, and Robinson (2011), Guyton (2011), and Kitces (2012). For the case study in the results, the minimum needs and lifestyle goals will be set equal out of respect for the planner's perspective.

Efficient Frontier

The efficient frontier for retirement income described here identifies the specific allocations among various combinations of asset classes and financial products which maximize the reserve of financial assets for a given percentage of spending goals reached, or which maximizes the percentage of spending goals reached for a given reserve of financial assets. Though scalable, each client requires a personalized efficient frontier matching household characteristics including age, gender, marital status, health, life expectancy, desired spending patterns, and constraints about asset and product allocation (such as a restriction that no more than 50% of assets will be annuitized). Customized expectations about the returns, volatilities, and correlations for asset classes and inflation, as well as mortality, can also be included, as can details about the current pricing of available retirement income tools. Different clients will desire different withdrawal rates for meeting basic expenses and for lifestyle spending goals, and they will also have varying amounts of financial assets and varying inflation-adjusted and fixed income from sources outside of the financial portfolio. Once an efficient frontier of product allocations is constructed with respect to all of these characteristics, clients can select one of the points on the frontier reflecting their personal preferences about the tradeoff between meeting spending goals and maintaining sufficient financial reserves. Results will be presented for a 65-year old heterosexual couple who are retiring and have already claimed Social Security (incorporating optimal claiming decisions into the framework will be left for subsequent research).

At retirement, assets are divided between stocks, bonds, inflation-adjusted SPIAs, fixed SPIAs, and VA/GLWBs. Each of these five components is simulated in 10 percentage point increments from zero to 100%, and there are 1,001 possible product allocations. For example, a product allocation of [20,30,10,30,10] means that at the retirement date, 20% of assets are invested in stocks, 30% in bonds, 10% into a real SPIA, 30% into a fixed SPIA, and 10% into a VA/GLWB. In this case, half of the assets remain in the financial portfolio with a fixed asset allocation of 40% stocks and 60% bonds, as these represent the relative shares of the total assets devoted to stocks and bonds. For the portfolio of financial assets, annual rebalancing restores the fixed asset allocation.

To obtain each year's spending amount (which is withdrawn at the start of each year), the maximum allowable amount is taken first from the annuitized financial products. If this sum exceeds what is needed to meet the lifestyle spending goal, then the excess funds are added to the financial portfolio. If the guaranteed income sources do not provide enough income to meet the lifestyle goal, then the remainder needed is withdrawn from the financial portfolio (stocks and bonds). Whenever fewer assets remain in the financial portfolio (stocks and bonds) than are needed to meet the spending goal, the client is unable to meet their full lifestyle spending goal for that year. They can only spend the income available from any annuities. The contract value of the VA/GLWB is included as part of the financial assets plotted on the y-axis, but clients do not withdraw more than the allowable amount each year so as to avoid breaking the contract terms for the guarantee. A SPIA allocation of 100% can still support a reserve of financial assets if the income received from the SPIA exceeds the lifestyle spending goal, as the excess is reinvested. In the case that the initial allocations to stocks and bonds are both zero, income returned to the investment portfolio in this manner is allocated to 100% stocks. Lifestyle goals and minimum needs are defined in real terms, while fixed SPIAs and

VA/GLWBs only provide nominal guarantees (GLWB step-ups cannot generally be expected to keep pace with inflation). As inflation erodes the real value of the nominal income, the percentage of lifestyle goals which these sources support declines.

Survival Probabilities

The outcome measures, remaining financial assets and percentage of lifetime spending needs which are satisfied, are calculated using survival probabilities from the Social Security Administration 2007 Period Life Table. The remaining financial assets measure is the sum of remaining financial assets in each year of retirement times the probability that the second member of the couple dies in that year. With this measure, the real value of remaining financial assets at death cannot be zero as some retirees will die shortly after retirement when assets will surely remain. This measure is an ex ante estimate of remaining financial assets based on survival probabilities for an unknown lifetime. Meanwhile, the percentage of lifetime spending needs which are satisfied is calculated as the sum of spending in each year as a percentage of the minimum spending need, times the probability that at least one member of the couple survives to that age.


Distribution of Outcomes

In presenting outcomes, the part of the distribution of outcomes which should be highlighted is not clear. For financial assets, should we focus on the mean amount, the median amount, or the remaining amount in a bad luck outcome such as the 10th percentile for values ranked from lowest to highest across the simulations? Likewise, the Monte Carlo simulations will produce a distribution of percentages of minimum spending needs which are met. Should we focus on the mean or median percentage of met needs across the simulations, or instead identify the percentage of needs met in a bad luck case like the 10th percentile? Different possibilities should be investigated and the issue becomes less important if the optimal strategies are fairly consistent between these different ways of measuring remaining financial assets and percentages of spending needs which are satisfied.

Results

Results are provided for a 65-year old couple. Their Social Security benefit is equal to 2% of their retirement date assets. Their lifestyle spending goal is 6% of retirement date assets, which requires them to use a 4% withdrawal rate above Social Security to meet their goal. Minimum needs are also 6%. Social Security is relevant because increasing the Social Security benefit will help support a higher percentage of spending needs being met, as it provides a further guaranteed income floor even when financial assets are depleted. Though higher Social Security benefits will not change the product allocations found on the efficient frontier, it will shift the frontier to the right, which could increase client comfort with comfortable choosing more aggressive allocations with a greater probability to increase financial asset holdings.

The efficient frontier for this case is shown in Figure 2, with the median value of remaining financial assets at death shown on the y-axis, and the 10th percentile bad luck outcome of spending needs which could be satisfied shown on the x-axis. The figure illustrates how 1,001 different product allocations perform with respect to meeting the two objectives. Moving toward the upper right hand corner of the figure is advantageous for retirees. A blue curve is added to show all of the product allocations consisting only of financial assets (stock and bonds). These outcomes represent some of the worst possible, demonstrating a clear role for partial annuitization to improve retiree outcomes, even in today's low interest rate environment. In particular, the red curve highlights most of the points on the efficient frontier. It shows the allocations between stocks and fixed SPIAs.




The results will vary when changing the personal characteristics and spending goals of the client. For instance, if spending needs require a withdrawal rate of only 1% of assets, then it is likely that spending needs will be met with any product allocation and the client can focus on the allocation which maximizes the median value of remaining financial assets: 100% stocks.

But in the particular case of Figure 2, the evidence suggests that optimal product allocations consist of stocks and fixed SPIAs, and clients need not bother with bonds, inflation-adjusted SPIAs, or VA/GLWBs. Though SPIAs do not offer liquidity, they provide mortality credits and generate bond-like income without any maturity date, and they support a higher stock allocation for remaining financial assets. Altogether, this allows a client to better meet both retirement financial objectives. Any of the product allocations on the efficient frontier represent a potentially optimal point, and clients can choose which one they think best balances their own objectives. Of course, any product allocations for which the client is not comfortable using (for instance, clients may fear too high of stock allocations or too much annuitization) could be removed and a constrained efficient frontier could be determined instead.

It is surprising to the author that fixed SPIAs dominate inflation-adjusted SPIAs in the retirement portfolio. In personal correspondence, Joseph Tomlinson explains why this may be. Essentially, either due to a lack of competition or because of the difficulties of hedging inflation risks, inflation-adjusted SPIAs are not priced competitively with fixed SPIAs. Fixed SPIAs supply more inflation-adjusted income than real SPIAs until cumulative inflation sufficiently reduces the real value of the fixed SPIA income stream. With payout rates of 5.84% and 3.875%, fixed SPIAs hold the advantage long enough when modeled with the breakeven inflation rate shown in Table 1. As using retirement date assets to purchase a fixed SPIA provides more income than needed until inflation sufficiently reduces the real value of the SPIA payments, and as excesses are invested in a portfolio of 100% stocks, a 100% allocation to a fixed SPIA can support more of the spending needs and result in the same median amount of financial reserves as a portfolio of 40% stocks and 60% bonds.

To provide a better idea about the range of potential spending outcomes, Figure 2 shows the distributions for the 11 allocations of stocks and fixed SPIAs, rather than only showing the 10th percentile of outcomes. This provides more information about the potential spending risks with the higher stock allocation portfolios.



Finally, Figure 4 provides the efficient frontiers for increasing lifestyle goals. Increasing spending to meet the lifestyle goal will result in less financial reserves in the median case and a smaller portion of spending needs which can be met in the 10th percentile. Evidence found in this sort of figure can help retirees determine a reasonable lifestyle goal with respect to their views about financial reserves and minimum needs.

