

Toward Determining the Optimal Investment Strategy for Retirement*

Javier Estrada

IESE Business School, Department of Finance, Av. Pearson 21, 08034 Barcelona, Spain
Tel: +34 93 253 4200, Fax: +34 93 253 4343, Email: jestrada@iese.edu

Mark Kritzman

Windham Capital Management, 800 Boylston Street, 30th Floor, Boston, MA 02199, USA
Tel: +1 617 419-3900, Fax: +1 617 236-5034, Email: mkritzman@windhamcapital.com

Abstract. *Investors who are about to retire are first and foremost concerned with supporting their spending needs throughout retirement. But they also derive satisfaction from growing their wealth beyond what is needed to support consumption in order to leave a bequest to their heirs or chosen charities. The predominant metric for evaluating retirement investment strategies is the failure rate. However, it fails to distinguish between strategies that fail early in retirement from those that fail near the end of retirement, and it fails to account for potential bequests. To overcome these shortcomings we propose a new metric, the coverage ratio, which is more comprehensive and informative than the failure rate. In addition, we propose a utility function to evaluate the coverage ratio, which penalizes shortfalls more than it rewards surpluses. Finally, we use the framework we propose to determine the optimal allocation to stocks and bonds using both historical and simulated returns.*

March, 2019

1. Introduction

The most important attribute of a retirement investment strategy is that it provides a return that is sufficiently large and reliable to support spending throughout retirement. Beyond this primary concern, investors also derive satisfaction from generating surplus wealth in order to leave a bequest to their heirs or chosen charities.

A critical issue for retirees and advisors is to determine whether a retirement investment strategy is better than another, and ultimately which is the best one. The predominant metric for such evaluation is the failure rate, which measures how frequently a strategy failed to sustain a withdrawal plan over all the (historical or simulated) retirement periods considered. This failure rate has at least two critical shortcomings. First, it fails to distinguish failures that occur near the beginning of retirement from those that occur near the end. Second, it fails to account for surpluses that could be left as a bequest.

* We would like to thank Jack Rader for his comments. Fran García and Cel Kulasekaran provided valuable research assistance. IESE's Center for International Finance (CIF) kindly provided partial support for this research. The views expressed below and any errors that may remain are entirely our own.

To overcome these limitations of the failure rate, we propose a new metric for evaluating retirement investment strategies called the coverage ratio, which measures the number of years of withdrawals (during and after retirement) supported by a strategy, relative to the length of the retirement period considered. In addition, we recognize that investors are more averse to failures than they are attracted to surpluses. Therefore, we propose a particular utility function for evaluating the coverage ratio which penalizes failures more than it rewards surpluses. Finally, we apply our approach to evaluate static investment strategies based on both historical and simulated returns.

For the sake of concreteness, we illustrate our approach by focusing on the retirement period, and particularly on the choice among competing asset allocations. However, our approach is more general and could be used to choose an optimal withdrawal rate during the retirement period; a saving rate or asset allocation during the accumulation period; or, more generally, to evaluate the success or failure of other financial goals.

The remainder of the paper is organized as follows. In Section 2 we introduce the coverage ratio, and we show how to evaluate it based on a kinked utility function, which we believe is a plausible description of investor preferences. In Section 3 we apply our framework to 11 static strategies based on historical returns across 21 countries and the world market. In Section 4 we apply our framework using simulation to evaluate static strategies given different assumptions for the length of the retirement period and the distribution of stock returns. We conclude in section 5.

2. The Coverage Ratio

Both academics and practitioners have invested considerable effort toward deriving measures to help retirees and advisors determine the appropriate blend of risky and safe assets for managing wealth throughout retirement. The most common measure is the failure rate, which is the fraction of all the (historical or simulated) retirement periods considered in which an investment strategy failed to support spending through the end of retirement. An important shortcoming of the failure rate is that it assumes investors are indifferent between running out of money early in retirement or near the end of retirement. In addition, it fails to account for potential surpluses that could be applied toward a bequest.¹

¹ To overcome the first limitation, Estrada (2017) introduced the concept of shortfall years to complement the failure rate; the latter measures how often a strategy failed, and the former measures by how much it failed. To overcome the second limitation, and therefore to account for funds left as bequest, Estrada (2018a and 2018b) introduced the risk-adjusted success ratio (RAS) and the downside risk-adjusted success ratio (D-RAS), both of which are linked to the variable ‘years sustained’ (i.e., the average number of years a strategy sustained withdrawals both when failing and when succeeding). The coverage ratio we introduced in this article is a simpler and more intuitive version of the years sustained variable, and the

To overcome both limitations, and ultimately to enable a more comprehensive evaluation of retirement strategies, we suggest that investors and advisors focus on the metric we introduce here, the coverage ratio, which aims to capture the number of years of withdrawals (during and after retirement) supported by a strategy, relative to the length of the retirement period considered. Importantly, our coverage ratio accounts for the years a strategy sustained withdrawals both when failing (not supporting withdrawals through the entire retirement period) and when succeeding (leaving a bequest).

Formally, let Y_t be the number of years of inflation-adjusted withdrawals sustained by a strategy, both during and after the retirement period, and L be the length of the retirement period considered. Then we define the *coverage ratio* in retirement period t (C_t) as

$$C_t = Y_t/L \tag{1}$$

By definition, $C < 1$ indicates that the strategy depleted the portfolio before the end of the retirement period; $C > 1$ indicates that the strategy sustained withdrawals through the entire retirement period and left a bequest; and $C = 1$ indicates that the strategy sustained withdrawals exactly through the end of the retirement period and left no bequest.

To illustrate, consider a 30-year retirement period, a \$1,000 retirement portfolio, annual inflation-adjusted withdrawals of \$40, and three strategies. The first strategy depletes a portfolio in 24 years, the second does so in exactly 30 years, and the third sustains withdrawals for 30 years and leaves a bequest of \$240 (which can support another six years of \$40 withdrawals). Then, Y_t would be 24, 30, and 36, for the first, second, and third strategies, and C_t would respectively be 0.8, 1.0, and 1.2.

However, the coverage ratio by itself is incomplete for evaluating the suitability of alternative investment strategies, because investors are much more likely to be displeased with outcomes that fall short of fully supporting retirement spending than pleased with outcomes that generate surpluses. In order to account for this asymmetry in preferences, we evaluate a strategy's coverage ratio within a utility-based framework. We propose a kinked utility function that assumes that as the coverage ratio increases above 1, the investor derives increasing satisfaction but at a decreasing rate; and as the coverage ratio falls below 1, the investor experiences a steep and linear decline in utility.² Formally, our utility function is given by

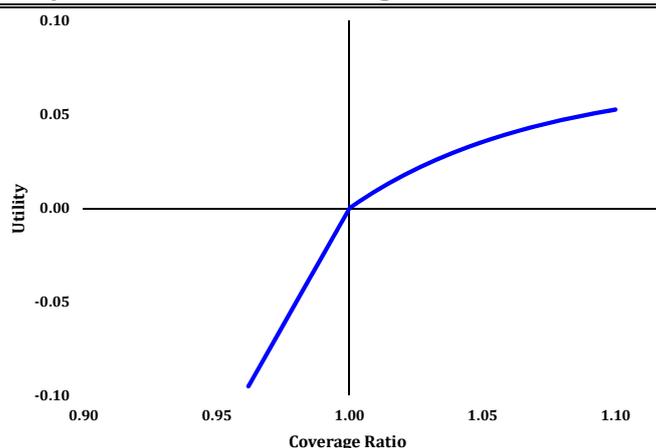
utility function we consider here accounts for risk as both RAS and D-RAS aim to do, but further takes into account an investor's attitude toward risk

² The notion of a kinked utility function predates our application here. Kahneman and Tversky (1979) proposed a utility function that assigns value to gains and losses rather than wealth levels. They argued that investors have concave utility in the domain of gains and convex utility in the domain of losses, which gives an S-shaped utility function. Thaler et al (1997) also assumed kinked utility in their experimental analysis of the effect of myopia and loss aversion on risk taking. And Yamashita (2014) showed how to employ option strategies to maximize utility for investors with kinked utility. Our application of kinked

$$\begin{aligned}
 U(C) &= \frac{C^{1-\gamma} - 1}{1-\gamma} && \text{for } C \geq 1 \\
 U(C) &= \frac{1^{1-\gamma} - 1}{1-\gamma} - \lambda(1 - C) && \text{for } C < 1
 \end{aligned} \tag{2}$$

where U denotes utility; C denotes the coverage ratio; γ is the coefficient of risk aversion, which determines the curvature of the slope when $C > 1$; and λ is a linear penalty coefficient when $C < 1$. This utility function is depicted in Exhibit 1.

Exhibit 1: Kinked Utility Function of the Coverage Ratio



This utility function has a very appealing property when we locate the kink at a coverage ratio of 1. When a strategy exactly funds a retirement period ($C=1$), utility equals zero. When a strategy leaves a bequest ($C > 1$), utility is positive and increases as the coverage ratio increases above 100% but at a diminishing rate. And when a strategy fails before the end of the retirement period ($C < 1$), utility declines steeply and linearly. If we were to set the kink below 1, the coverage ratio would convey negative utility before the penalty function takes effect. We would thus be adding an unintended penalty, and one that is difficult to interpret; it is therefore very clean to locate the kink at 1.

While we are confident of the general form of our proposed utility function, there are no universally appropriate values for the curvature of the utility function above the kink or the slope of the penalty function below the kink. Higher values of γ , which defines the degree of curvature, cause utility to increase more slowly as the coverage ratio increases; and higher values for λ penalize spending shortfalls more severely than lower values.

utility, however, is uniquely calibrated for the problem we address. Specifically, we locate the kink at a coverage ratio equal to 1, which is required for our particular application.

3. Historical Results

We first apply our approach using returns from the Dimson-Marsh-Staunton database, described in detail in Dimson, Marsh, and Staunton (2002, 2016). It contains annual returns for stocks and long-term government bonds over the 1900-2014 period for 21 countries and the world market. Returns are adjusted by each country's periodic inflation rate, and they are expressed in local currency (except for the world market, in dollars). All returns account for capital gains and losses as well as cash flows such as dividends and coupons.

The analysis is based on a \$1,000 portfolio at the beginning of retirement, a 4% initial withdrawal rate, annual inflation-adjusted withdrawals, and a 30-year retirement period.³ At the beginning of each year the annual withdrawal is made, the portfolio is then rebalanced to the target asset allocation for the year, and then it compounds at the observed return of stocks and bonds for that year. This process is repeated at the beginning of each year during the 30-year retirement period, at the end of which the portfolio has a terminal wealth or bequest that may be positive or 0. The first 30-year retirement period considered is 1900-1929 and the last one is 1985-2014, for a total of 86 rolling (overlapping) periods.

The analysis considers 11 static (annually-rebalanced) stock-bond allocations ranging from 100% (all stocks, no bonds) to 0% (no stocks, all bonds), with nine allocations (90%, 80%, ..., 20%, 10%) in between, all indicated by the proportion of stocks in the portfolio, with the balance allocated to bonds.

For each of the 11 strategies and 22 markets we consider we proceed as follows. We first calculate the coverage ratio for each of the 86 rolling 30-year retirement periods in our sample. Then we calculate the utility an investor derives from each coverage ratio. Finally, we calculate expected utility by averaging the utilities in the previous step across the 86 retirement periods. This process yields the average utility an investor obtains from a given strategy in a given market. We implement the process using the kinked utility function described earlier, assuming that the risk aversion coefficient γ equals 0.9999, and the linear penalty coefficient λ equals 10.⁴ Exhibit 2 summarizes the results of our empirical analysis.

Exhibit 2 reveals that our approach results in the selection of relatively aggressive strategies, with an average allocation of 91% to stocks and 9% to bonds. In over half of the markets, including the U.S. and the world market, the strategy selected is the most aggressive of those considered, namely, 100% stocks. It is important to note that the optimal allocations we

³ Needless to say, the length of the retirement period neither has to be 30 years nor is certain. Individuals evaluating a strategy well into their retirement may want to use shorter periods, conservative individuals at the beginning of retirement may want to use longer periods. Longevity risk has been widely explored in the literature, often within the framework of annuities; see, for example, Maurer et al (2013).

⁴ As γ approaches 1, utility equals the natural logarithm of the coverage ratio; hence, we are effectively using a log-wealth utility function for the coverage ratio above 1.

report for each country do not suggest what each country's future optimal allocation will be, unless its historical market conditions prevail into the future. We report results across many countries in order to provide a broad perspective and to show how results differ based on different market conditions. Moreover, we recognize that our results are based on overlapping retirement periods and are, therefore, less reliable than they would be if we had data for many independent retirement periods.

Exhibit 2: Optimal Stock Allocation Based on Historical Returns

This exhibit shows the allocation to stocks (S) selected by the utility function in (2) for $\gamma=0.9999$ and $\lambda=10$, with the rest allocated to bonds. The 11 asset allocations considered range between 100 (all stocks) and 0 (no stocks), with nine allocations (90, 80, ..., 20, 10) in between. All strategies are evaluated over 86 rolling 30-year retirement periods between 1900-1929 and 1985-2014; a starting capital of \$1,000; a 4% initial withdrawal rate; subsequent annual withdrawals adjusted by inflation; and annual rebalancing.

| Country | S (%) | Country | S (%) |
|-----------|---------|--------------|---------|
| Australia | 100 | Netherlands | 90 |
| Austria | 80 | New Zealand | 100 |
| Belgium | 100 | Norway | 90 |
| Canada | 100 | Portugal | 60 |
| Denmark | 90 | South Africa | 100 |
| Finland | 100 | Spain | 70 |
| France | 100 | Sweden | 60 |
| Germany | 100 | Switzerland | 70 |
| Ireland | 100 | UK | 100 |
| Italy | 100 | USA | 100 |
| Japan | 90 | World | 100 |

As we already mentioned, we are confident about the shape of the utility function we propose, and we think our choice of parameters for the base case ($\gamma=0.9999$ and $\lambda=10$) is sensible. That said, in Exhibit 3 we explore the sensitivity of the results in our base case to changes in the value of these two parameters, just for the U.S. market. As the exhibit shows, and as expected, the optimal asset allocation does become more conservative as either γ or λ increase. However, even for high values of both coefficients, the optimal strategy never allocates less than 80% to stocks.

Exhibit 3: Sensitivity Analysis (U.S. Market)

This exhibit shows the allocation to stocks (in %) selected by the utility function in (2) for different values of the risk aversion (γ) and penalty (λ) coefficients, with the rest allocated to bonds. The 11 asset allocations considered range between 100 (all stocks) and 0 (no stocks), with nine allocations (90, 80, ..., 20, 10) in between. All strategies are evaluated over 86 rolling 30-year retirement periods between 1900-1929 and 1985-2014; a starting capital of \$1,000; a 4% initial withdrawal rate; subsequent annual withdrawals adjusted by inflation; and annual rebalancing.

| λ (Penalty) | γ (Degree of Curvature) | | | | | |
|------------------------|--------------------------------|--------|-----|-----|-----|-----|
| | 0.5 | 0.9999 | 1.5 | 2.0 | 3.0 | 5.0 |
| 10 | 100 | 100 | 100 | 90 | 80 | 80 |
| 30 | 100 | 90 | 90 | 80 | 80 | 80 |
| 50 | 100 | 90 | 80 | 80 | 80 | 80 |
| 70 | 90 | 80 | 80 | 80 | 80 | 80 |
| 90 | 90 | 80 | 80 | 80 | 80 | 80 |

We suspect that our approach yields relatively aggressive strategies for two reasons. First, the equity risk premium was relatively high in most markets in our sample; second, we use a 30-year retirement period, and for such long investment horizons stocks have far more often than not outperformed bonds.⁵

4. Simulated Results

We now present results based on simulated returns, which have several advantages relative to results that are based on actual historical returns. First, simulation allows us to apply our approach using different assumptions for the distribution of stock returns. Moreover, simulated retirement periods are independent of each other unlike historical retirement periods which overlap; hence, the simulated results are statistically more reliable. Finally, simulation allows us to eliminate any biases that may have arisen from the particular sequence of returns that occurred historically.

The sequence of returns could have a drastic impact on the progression of wealth throughout retirement. In the early years of retirement, there is greater wealth because less of it has been spent. Thus, relatively high (low) returns in the early (later) years will favor (hinder) the compounding of capital. Whereas historical returns allow only for the return sequences that actually occurred, simulation captures a much broader range of potential return sequences.

We employ Monte Carlo simulation to generate 25,000 annual real returns for stocks and bonds drawing from two uncorrelated normal distributions, one with a particular mean and standard deviation to characterize stocks and another with different mean and standard deviation to characterize bonds. We always assume that the distribution of bond returns has a mean (real) return of 2% and a standard deviation of 3%. However, we allow the mean (real) return and standard deviation for stock returns to vary.

For each set of the 25,000 simulated returns we compute for each 30-year retirement period the coverage ratio for stock allocations ranging between 0% and 100%, in 10% increments, with the balance allocated to bonds. As in the historical analysis, we assume a \$1,000 retirement portfolio, an initial withdrawal rate of 4%, and subsequent annual withdrawals adjusted by inflation.

In our base case we compute the utility of each strategy's coverage ratio for each of the 25,000 trials assuming, as in the historical analysis, $\gamma=0.9999$ (essentially log-wealth utility for coverage ratios higher than 1) and $\lambda=10$ (penalizing coverage ratios lower than 1 more than we reward coverage ratios higher than 1). We store these utilities for each of the 25,000 trials and

⁵ Across the 21 markets in our sample, the average arithmetic mean return, geometric mean return, and standard deviation were 7.3%, 4.6%, and 23.9% for stocks, and 2.4%, 1.0%, and 14.9% for bonds.

compute the average utility for each of the mixes of stocks and bonds. We then identify the asset allocation with the highest average utility. Finally, we repeat this exercise for stock and bond return distributions that have different combinations of means and standard deviations for stocks. Exhibit 4 presents these results.

Exhibit 4:

Optimal Stock Allocation Based on Simulated Returns for 30-Year Retirement Periods

This exhibit shows the allocation to stocks (in %) selected by the utility function in (2) for $\gamma=0.9999$ and $\lambda=10$, with the rest allocated to bonds, given different values for the mean and standard deviation of stock returns. The 11 asset allocations considered range between 100 (all stocks) and 0 (no stocks), with nine allocations (90, 80, ..., 20, 10) in between. These results assume a starting capital of \$1,000, a 4% initial withdrawal rate with subsequent annual withdrawals adjusted by inflation, and annual rebalancing. The mean and standard deviation for bond returns are assumed to be 2% and 3%.

| Stocks Mean | Standard Deviation of Stocks | | | |
|----------------|------------------------------|-----|-----|-----|
| | 15% | 20% | 25% | 30% |
| 3% | 50 | 30 | 20 | 10 |
| 4% | 100 | 60 | 40 | 30 |
| 5% | 100 | 80 | 60 | 40 |
| 6% | 100 | 100 | 70 | 50 |

Exhibit 4 reveals that the optimal allocation to stocks is significantly lower than the stock allocations suggested by historical returns when we assume lower mean returns and standard deviation for stocks than those that prevailed historically. These simulated results may provide better guidance for investors who believe that the historical record represents a relatively favorable pass through history that is unlikely to recur.

We next explore whether the aggressive stock allocations suggested by historical returns may result, in part, from our choice of a 30-year retirement period. We recognize that many investors may be well into retirement and hence face a shorter horizon than 30 years. We, therefore, produce results for horizons of 20, 10, and 5 years. We follow the same steps in our simulation as we described earlier with one important exception. When we simulate each of the 25,000 paths for a particular stock-bond allocation for these shorter horizons, the wealth at the beginning of these shorter horizons is determined by the strategy returns and withdrawals that led up to that point in time for each of the allocations, assuming the strategy was invested in the same stock-bond allocation in the years leading up to the beginning of the shorter horizon.

Thus when we evaluate a 60-40 stock-bond allocation looking ahead 20 years, we assume the wealth at that point in time was generated by a 60-40 stock-bond allocation throughout the prior 10 years. We, therefore, make no presumption that the investor chose the optimal strategy before the inception of these shorter horizons. We simply assume, that for each asset mix, the investor invested in the same asset mix leading up to these shorter horizons. Obviously, the level of wealth 10 years into an individual's retirement period will be specific to the returns and spending that actually occurred. Our results are based on 25,000 paths; hence they represent what one might expect on average.

Exhibit 5:**Optimal Stock Allocation Based on Simulated Returns for Different Retirement Periods**

This exhibit shows the allocation to stocks (in %), for different lengths of the retirement period, selected by the utility function in (2) for $\gamma=0.9999$ and $\lambda=10$, with the rest allocated to bonds, given different values for the mean and standard deviation of stock returns. The 11 asset allocations considered range between 100 (all stocks) and 0 (no stocks), with nine allocations (90, 80, ..., 20, 10) in between. These results assume a starting capital of \$1,000, a 4% initial withdrawal rate with subsequent annual withdrawals adjusted by inflation, and annual rebalancing. The mean and standard deviation for bond returns are assumed to be 2% and 3%.

| 20-Year Horizon (Years 11-30) | | | | |
|--------------------------------------|-------------------------------------|------------|------------|------------|
| Stocks | Standard Deviation of Stocks | | | |
| Mean | 15% | 20% | 25% | 30% |
| 3% | 40 | 20 | 10 | 10 |
| 4% | 70 | 40 | 30 | 20 |
| 5% | 100 | 60 | 40 | 30 |
| 6% | 100 | 70 | 50 | 40 |
| 10-Year Horizon (Years 21-30) | | | | |
| Stocks | Standard Deviation of Stocks | | | |
| Mean | 15% | 20% | 25% | 30% |
| 3% | 20 | 10 | 10 | 0 |
| 4% | 30 | 20 | 10 | 10 |
| 5% | 40 | 20 | 20 | 10 |
| 6% | 60 | 30 | 20 | 10 |
| 5-Year Horizon (Years 26-30) | | | | |
| Stocks | Standard Deviation of Stocks | | | |
| Mean | 15% | 20% | 25% | 30% |
| 3% | 10 | 10 | 0 | 0 |
| 4% | 10 | 10 | 10 | 0 |
| 5% | 20 | 10 | 10 | 10 |
| 6% | 40 | 20 | 10 | 10 |

Exhibit 5, together with Exhibit 4, reveal that the optimal allocation to stocks differs in accordance with conventional wisdom. If we hold constant standard deviation and horizon, the optimal stock allocation rises with the expected return of stocks. If we hold constant the expected return of stocks and horizon, the optimal stock allocation falls as the standard deviation of stock returns increases. And if we hold constant both the expected return and standard deviation, the optimal allocation to stocks falls as the retirement period becomes shorter. In some cases optimal allocations do not change from one cell to another; this is because we only allow the asset allocation to change in 10% increments. If we were to allow for more granular changes in the stock-bond mix, we would see differences across all combinations of expected return, standard deviation, and retirement period.

It is worth noting that although the sensitivity of the optimal allocation to stocks to the length of the retirement period is consistent with conventional wisdom, it appears to run counter to Samuelson's famous contradiction to the notion of time diversification (Samuelson, 1963). In this classic paper Samuelson argued that the optimal mix of risky and safe assets should be invariant to the length of the investment horizon. He showed that under certain conditions (essentially that investors have a power utility function and that returns are serially

independent), the likelihood of loss and the potential magnitude of loss exactly offset each other as the investment horizon increases.

In our analysis we assume that returns are serially independent, but we do not assume power utility across the entire range of returns. We assume that investors have power utility for returns that generate a coverage greater than 1, but that they have linear utility for returns that produce coverage ratios below 1. Stated more prosaically, we assume that investors face a critical threshold (a coverage ratio equal to 1) at which their utility changes abruptly. Thus our framework is not at all inconsistent with Samuelson's observation about the connection between time and risk.

5. Conclusion

We propose a new framework for evaluating retirement investment strategies. We begin with the assumption that investors are primarily concerned with generating returns that are sufficiently large and stable to sustain spending throughout retirement, and that they derive additional satisfaction from leaving a bequest.

We then introduce a new metric, the coverage ratio, which measures the number of years of withdrawals (during and after retirement) supported by a strategy, relative to the length of the retirement period considered. Importantly, this metric accounts for everything the failure rate accounts for. In addition, it accounts for when a strategy fails if it does, and by how much it succeeds when it does, neither of which is taken into account by the failure rate.

Moreover, we propose a particular kinked utility function for evaluating the coverage ratio. We assume that utility rises at a diminishing rate as the coverage ratio rises above 1; that is, with increases in the potential bequest, and that it falls sharply and linearly as the coverage ratio falls below 1, which corresponds to increasing failure to support spending. This particular utility function is consistent with our initial assumption that investors are primarily concerned with supporting spending and secondarily concerned with leaving a bequest.

We then applied our framework to determine the optimal allocation to stocks and bonds based on historical returns across 21 countries as well as the world market over a 110-year period. Our analysis revealed that in most countries investors should have implemented very aggressive strategies, with an average allocation to stocks of 91%. We argued that this strong preference for stocks was caused by the favorable equity risk premium that prevailed across these markets as well as by the (30-year) length of the retirement period we considered.

We next applied our framework using simulated returns. Simulation allowed us to experiment with different assumptions for the expected return and risk of stocks and with different retirement horizons. Our simulations revealed that the optimal allocation to stocks

varied in accordance with conventional wisdom. All else equal, the optimal allocation to stocks was higher for higher expected returns for stocks, lower volatility, and longer retirement periods.

We believe that our framework for determining the optimal retirement investment strategy is conceptually appealing and empirically reasonable. We also believe that it is superior to a simplistic framework based solely on the limited information provided by the failure rate. We thus encourage investors and particularly advisors to take it into account when making financial decisions for retirement.

References

- Dimson, Elroy, Paul Marsh, and Mike Staunton (2002). *Triumph of the Optimists – 101 Years of Investment Returns*. Princeton University Press.
- Dimson, Elroy, Paul Marsh, and Mike Staunton (2016). *Credit Suisse Global Investment Returns Sourcebook 2016*. Credit Suisse Research Institute, Zurich.
- Estrada, Javier (2017). “Refining the Failure Rate.” *Journal of Retirement*, 4, 3, 63-76.
- Estrada, Javier (2018a). “From Failure to Success: Replacing the Failure Rate.” *Journal of Wealth Management*, *Journal of Wealth Management*, 20, 4, 9-21.
- Estrada, Javier (2018b). “Replacing the Failure Rate: A Downside Risk Perspective.” *Journal of Retirement*, 5, 4, 46-56.
- Kahneman, Daniel, and Amos Tversky (1979). “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica*, 47, 2, 263-292.
- Maurer, Raimond, Olivia Mitchell, Ralph Rogalla, and Vasily Kartashov (2013). “Lifecycle Portfolio Choice with Systematic Longevity Risk and Variable Investment-Linked Deferred Annuities.” *Journal of Risk and Insurance*, 80, 3, 649-676.
- Samuelson, Paul (1963). “Risk and Uncertainty: A Fallacy of Large Numbers.” *Scientia*, 98, 4, 108-113.
- Thaler, Richard, Amos Tversky, Daniel Kahneman, and Alan Schwartz (1997). “The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test.” *Quarterly Journal of Economics*, 112, 2, 647-661.
- Yamashita, Miwaka (2014). “Optimal Investment Strategy for Kinked Utility Maximization; Covered Call Option Strategy.” *Journal of Mathematical Finance*, 4, 2, 55-74.