Evaluating Retirement Strategies: A Utility-Based Approach *

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. *Retirees need to make two critical financial decisions, namely, the rate at which they withdraw funds to support their retirement and the asset allocation of their portfolios. We propose a methodology that retirees, and particularly advisors, could use to make these decisions in an optimal way. We introduce a new variable, the coverage ratio, and an evaluation framework based on utility. Our approach can be used to make optimal decisions during both the accumulation and the retirement period, but we illustrate it by focusing on the latter, and particularly on the choice of an optimal asset allocation.*

October, 2018

Individuals planning for retirement need to determine the fraction of income they must save as they approach retirement, the fraction of their savings they need to withdraw throughout retirement to support their consumption, and how to allocate their savings between risky and safe assets. In this article we focus on the third decision, how to allocate retirement savings between risky and safe assets, but our ultimate goal is broader.

Academics and practitioners have invested considerable time and effort toward deriving measures to evaluate investment strategies. We contribute to this discussion by proposing both a new variable, the coverage ratio, and an evaluation framework based on utility, that aim to help investors and advisors determine an optimal asset allocation strategy.

We illustrate our approach by focusing on the retirement period, and particularly on the optimal choice among competing asset allocations. But importantly, our approach can be used just as well to choose an optimal withdrawal rate during the retirement period, or a saving rate and asset allocation during the accumulation period.

Of the many variables that have been proposed to evaluate investment strategies during retirement, the most widely used is the failure rate, which aims to capture how often a strategy failed to sustain a withdrawal plan through the end of the retirement period. We propose an

* We would like to thank Cel Kulasekaran and Jack Rader for their comments. Fran García 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.

alternative variable, the coverage ratio, which captures the fraction of the retirement period a strategy was able to sustain a withdrawal plan. For reasons we discuss below, we believe that the coverage ratio is a better variable than the failure rate to evaluate retirement strategies.

But we do not evaluate competing asset allocation strategies simply by comparing their coverage ratios. We consider the *utility* of a strategy's coverage ratio. Specifically, we propose a kinked utility function in which utility increases at a decreasing rate for returns that exceed the return required to sustain a withdrawal plan and decreases linearly with returns that fall short of sustaining a withdrawal plan. This utility function recognizes that investors are more displeased with failures than they are pleased with surpluses.

We test our approach empirically by applying it to a sample of 21 countries and the world market over a 115-year period. We find that the asset allocation strategies selected by our approach are, in general, more aggressive than we suspect most investors would be willing to tolerate. We attribute this outcome to the high equity risk premiums that prevailed throughout our sample. Although we are confident in the logic of our approach, we advise investors who choose to apply it to do so with simulations that reflect their expectations for future return distributions.

The Failure Rate

The failure rate is the most commonly used metric for evaluating investment strategies for retirement. It can be traced back to the groundbreaking article by Bengen (1994), in which he aims to determine a safe withdrawal rate. He showed that a retiree withdrawing 4% of a portfolio equally split between stocks and bonds at the beginning of retirement, and adjusting all subsequent annual withdrawals by inflation, would have never depleted a portfolio (failed) in less than 30 years, which he considered the minimum requirement for portfolio longevity. From that point on, the failure rate became the standard metric used to evaluate retirement strategies.

That said, the failure rate has a critical shortcomings, see Milevsky (2016). One of them is that it does not distinguish between a failure that occurs early in retirement from one that occurs near the end of retirement. Estrada (2017) addresses this issue by introducing the variable *shortfall years*, which aims to complement the failure rate by measuring the average number of years a strategy failed to support withdrawals over all the retirement periods in which it failed. Thus, the failure rate measures the frequency of failure and shortfall years complements it by measuring the magnitude of failure.

Estrada (2018*a*, 2018*b*) further proposes to change the perspective from failure to success and introduces the variable *years sustained*, which measures the average number of years a strategy sustained withdrawals both when it failed and when it succeeded. He also introduces the variables *risk-adjusted success* (RAS) and *downside risk-adjusted success* (D-RAS), which relate a strategy's years of withdrawals sustained to two different measures of risk.

However, we believe there is a simpler and more comprehensive way to evaluate retirement strategies, which accounts for both success and failure, as well as both their frequency and magnitude. We therefore introduce the coverage ratio.

The Coverage Ratio

The coverage ratio equals the fraction of years an investment strategy supported withdrawals, including surpluses that could have supported withdrawals beyond the end of retirement. Let f be a variable that takes a value of 1 in a retirement period in which a strategy failed and 0 otherwise. Then, the *failure rate* (F) is formally defined as

$$F = \left(\frac{1}{T}\right) \cdot \sum_{t=1}^{T} f_t \quad (1)$$

where T is the number of (historical or simulated) retirement periods considered and t indexes retirement periods, both typically measured in years.

Let Y_t be the number of years of inflation-adjusted (real) 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

$$C_t = Y_t/L \quad (2)$$

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 behind; and C=1 indicates that the strategy sustained withdrawals exactly through the end of the retirement period with no bequest left behind.

To illustrate the intuition behind our coverage ratio, consider (as we do later in our empirical section) a \$1,000 portfolio at the beginning of retirement, a 4% initial withdrawal rate, subsequent annual withdrawals adjusted by inflation, and a 30-year retirement period. This set-up yields an initial withdrawal of \$40, followed by 29 withdrawals of \$40 in inflation-adjusted dollars. Consider also three scenarios: In the first, the portfolio is depleted after 24 years, six years short of the end of the retirement period; in the second, the portfolio is depleted after 30 years, having sustained 30 years of withdrawals but leaving no bequest behind; and in the third, the portfolio sustained 30 years of withdrawals and left behind a bequest of \$240 in real dollars.

Then, by definition, in the first, second, and third scenarios our Y_t variable would be 24, 30, and 36; and our coverage ratios would be C=0.8, C=1.0, and C=1.2. Note that in the third scenario the strategy sustained 30 years of withdrawals during the entire retirement period and left behind a bequest of $240 in real dollars, which amounts to six additional years of $40 withdrawals in real dollars.

The appeal of our coverage ratio is that it consolidates in one variable both the frequency and the magnitude of success and failure. For this reason, it provides more information than the failure rate (which measures only the frequency of failure) and shortfall years (which measures only the magnitude of failure). Furthermore, although the failure rate does not distinguish between two strategies that succeeded but left behind very different bequests, the coverage ratio (which would increase with the size of the bequest) does.

That said, our coverage ratio and the failure rate are clearly related. If F=0, then it must be the case that $C_t \geq 1$ for all t, implying that the strategy sustained withdrawals for at least 30 years in every retirement period considered; if F>0, then it must be the case that $C_t < 1$ for at least some t, implying that at least in one retirement period the strategy fell short; and if F=1, then it must be the case that $C_t < 1$ for all t, implying that the strategy failed in every period.

A Utility-Based Approach

We believe that our coverage ratio provides a more comprehensive assessment of retirement strategies than the failure rate. However, we do not propose to calculate the average coverage ratio across all the retirement periods considered for a strategy, and compare it to that of a competing strategy. Although doing so would account for the success or failure of different strategies, it would fall short on at least two counts. First, it would fail to account for risk; and second, it would fail to account for the fact that retirees are much more averse to outcomes that fail to sustain withdrawals than they are attracted to outcomes that produce surpluses.

To that purpose, we propose a utility function in which as the coverage ratio increases above 1, a retiree's utility increases at a decreasing rate, thus implying risk aversion. This notion of utility was first proposed by Daniel Bernoulli in 1738 and is a widely accepted description of preferences throughout the finance literature.

We believe that this is a proper description of a retiree's utility as long as the coverage ratio is higher than 1. However, when the coverage ratio falls below 1, thus implying a strategy's failure to sustain withdrawals through the entire retirement period, it makes sense to assume that a retiree would experience a steep decline in utility. Therefore, we propose a kinked utility function given by the expression

$$U(C) = \frac{C^{1-\gamma} - 1}{1-\gamma} \qquad \text{for } C \geq 1$$

$$U(C) = \frac{1^{1-\gamma} - 1}{1-\gamma} - \lambda(1 - C) \qquad \text{for } C < 1 \qquad (3)$$

where U denotes utility; 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 $C < 1$. Exhibit 1 depicts the kind of utility function we propose, which is similar to that used by Adler and Kritzman (2007) and Czasonis et al (2018).

Exhibit 1: A Kinked Utility Function

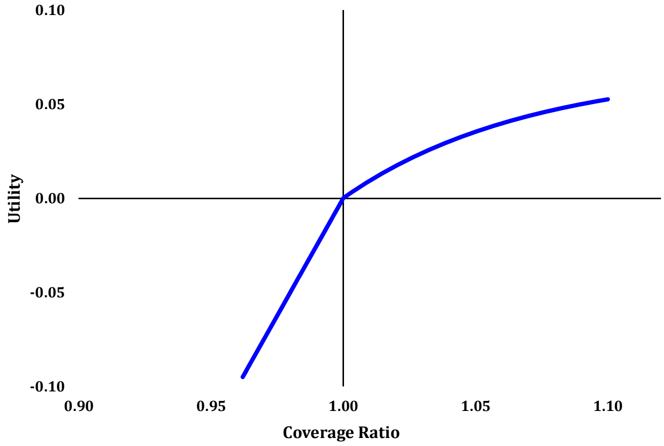

As both the expression and the exhibit show, this utility function has a very appealing property, namely, that the kink is located at a coverage ratio of 1. Therefore, when $C = 1$, implying that a strategy sustained withdrawals during the entire retirement period but left no bequest behind, $U(C) = 0$. When $C > 1$, implying that a strategy sustained withdrawals during the entire retirement period and in addition left a bequest behind, utility increases at a decreasing rate with the size of the bequest. And when $C < 1$, implying that a strategy failed to sustain withdrawals during the entire retirement period, utility falls steeply and linearly with the size of the shortfall.

Note that locating the kink below $C = 1$ would imply negative utility even when a strategy did not fail. This would imply an unintended penalty, and one that is difficult to interpret. We believe that locating the kink exactly at $C = 1$ is a very clean solution. Furthermore, we believe it is both plausible and intuitive to assume that when the coverage ratio is higher than 1, utility increases at a decreasing rate with the size of the bequest; and when it is lower than 1, utility falls linearly and steeply with the size of the shortfall. In other words, a retiree is less happy with a surplus than he is unhappy with a shortfall of the same amount.

It should be clear from (3) that γ determines the curvature of the function when $C > 1$, with higher values indicating that utility increases more slowly as the coverage ratio increases. Furthermore, λ determines the steepness of the function when $C < 1$, with higher values indicating that utility decreases more rapidly as the coverage ratio decreases. In our base case scenario we consider $\gamma = 0.9999$ (essentially logarithmic utility for positive coverage ratios) and $\lambda = 10$. While we are confident on the form of the utility function we propose, we understand that there are no universally accepted values for γ and λ. We believe that our choice of parameters provides sensible results, but we also perform a sensitivity analysis on these two parameters to provide a broader perspective.

Data and Methodology

The sample we consider is 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 real (adjusted by each country's inflation rate), in local currency (except for the world market, in dollars), and account for both capital gains/losses and cash flows (dividends or coupons). Exhibit A1 in the appendix summarizes some characteristics of all the series of stock and bond returns in the sample.

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.[1] 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 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.

It is important to note that the optimal allocations we report in the next section for each country do not presage what each country's future optimal allocation will be, unless its historical market conditions prevail into the future. Our purpose in showing results across many countries is 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. Having said that, we can only work with the available data.

[1] Changing the initial withdrawal rate would change all our results but none of the conclusions we draw from the analysis. As already mentioned, our main goal is to introduce a methodology that can be used to evaluate retirement strategies, and to that purpose we just focus on the choice among different asset allocations.

Results

In order to implement the utility-based framework we propose, we take the following steps for each of the 11 strategies and 22 markets we consider. First, we calculate the coverage ratio for each of the 86 rolling 30-year retirement periods in our sample. Then we calculate the utility a retiree 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 a figure that summarizes the average utility a retiree perceives from a given strategy in a given market.

We implement the process above first for the utility function in (3), assuming γ = 0.9999 and λ = 10 in our base case. As we already discussed, this utility function features a penalty for failing to sustain withdrawals during the entire retirement period, which increases with the size of the shortfall. In order to evaluate the impact of this penalty, we also explored the optimal choices that follow from the expression

$$U(C) = \frac{C^{1-\gamma} - 1}{1-\gamma} \qquad \text{for all } C \qquad (4)$$

that is, a power utility function over the entire range of coverage ratios, regardless of whether this variable is above or below 1. This utility function indicates that a retiree's utility decreases as the coverage ratio decreases, but there is no additional penalty when a strategy fails (that is, when C<1).

Exhibit 2 summarizes the main results of our analysis; panel 1 is based on the utility function in (3) and panel 2 on the modified utility function in (4). In both panels, we report the optimal asset allocation for each market, based on the highest average utility.

Panel 1, based on the utility function in (3), 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. The most conservative strategy selected, in only two countries (Portugal and Sweden), is a portfolio with 60% in stocks. In two other countries (Spain and Switzerland) the optimal choice consists of 70% in stocks; in one country (Austria) the highest utility corresponds to an allocation of 80% to stocks; and in the rest of the countries the strategy selected consists of 90% in stocks.

Panel 2 reveals that the strategies selected by a retiree with the utility function given by (4) would be somewhat more aggressive than those selected by another retiree with a utility function given by (3). In other words, the absence of an additional penalty for failure results in allocations with a higher proportion of stocks. On average across all markets, the strategies selected based on (3) allocate 91% to stocks; those selected based on (4), on the other hand, allocate 98% to stocks.

Exhibit 2: Coverage Ratio – Utility-Based Approach

This exhibit shows the asset allocation selected by two utility functions, from the 11 allocations considered, ranging from 100 (all stocks, no bonds) to 0 (no stocks, all bonds), with nine allocations (90, 80, …, 20, 10) in between, the rest being allocated to bonds. The results in Panel 1 are based on expression (3) in the text for γ=0.9999 and λ=10; those in Panel 2 are based on expression (4) for γ=0.9999. 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. The data is described in Exhibit A1 in the appendix.

Panel 1: With Kink	**Country**	**Allocation**	**Country**	**Allocation**
	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
Panel 2: No Kink	**Country**	**Allocation**	**Country**	**Allocation**
	Australia	100	Netherlands	100
	Austria	70	New Zealand	100
	Belgium	100	Norway	90
	Canada	100	Portugal	100
	Denmark	100	South Africa	100
	Finland	100	Spain	90
	France	100	Sweden	100
	Germany	100	Switzerland	100
	Ireland	100	UK	100
	Italy	100	USA	100
	Japan	100	World	100

As we already mentioned, we are confident on the shape of the utility function we propose, and we think our choice of parameters for the base case (γ = 0.9999 and λ = 10) is sensible. That said, we explored how sensitive the results of our base case are to changes in the value of these two parameters.

Although λ is a penalty coefficient specific to our utility function, there is a vast literature that discusses plausible values for the risk aversion coefficient (γ). Gandelman and Hernández-Murillo (2014) estimate this coefficient at the aggregate level for 75 countries and obtain a range between 0 and 3; a value in the vicinity of 1 for the vast majority of countries; and a cross-sectional average of 0.98. They conclude that their overall results support the use of a logarithmic utility function, which is essentially what we do for coverage ratios higher than 1. Furthermore, Thomas (2016) argues that the UK Treasury recommends the use of a risk-aversion coefficient equal to 1.

Exhibits A2 through A7 in the appendix report the results of a sensitivity analysis on the γ and λ coefficients of our utility function. The first three exhibits (A2 through A4) consider changes in the value of the risk aversion coefficient from 0.9999 in the base case to 0, 0.5, and 2. Predictably, higher values of γ (greater risk aversion), lead to more conservative strategies. Still, we find that the most conservative strategies selected consist of allocations with at least 40% in stocks.

Exhibits A5 through A7 report the results of a sensitivity analysis on the penalty coefficient that applies when the underlying strategy fails. In this case, we consider changes in the penalty from the original 10 to 1, 5, and 20. Again predictably, higher values of λ lead to more conservative strategies. Still, as in the previous case, we observe that none of the strategies selected consist of allocations with less than 40% in stocks.

We believe that our choices for the base case scenario are sensible, as is the fairly broad range of values we consider in the sensitivity analysis we just discussed. Still, we performed an even broader sensitivity analysis with coefficients of risk aversion and penalty coefficients much higher than what would be considered plausible, to explore whether we observed substantial changes in the strategies selected. Exhibit 3 displays our results for the U.S. market.

Exhibit 3: Utility-Based Approach – Broader Sensitivity Analysis – USA

This exhibit shows the proportion of stocks in the asset allocation (AA) selected, the rest of the portfolio invested in bonds, for different values of the risk aversion (γ) and penalty (λ) coefficients. The utility function is given by expression (3). The 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. The data is described in Exhibit A1 in the appendix.

$\boldsymbol{\gamma}$	**0.5**	**0.9999**	**1.5**	**2.0**	**2.5**	**3.0**	**3.5**	**4.0**	**4.5**	**5.0**	**10.0**	**15.0**	**20.0**
AA	100	100	100	90	90	80	80	80	80	80	80	80	80
$\boldsymbol{\lambda}$	**1**	**5**	**10**	**15**	**20**	**25**	**30**	**35**	**40**	**45**	**50**	**75**	**100**
AA	100	100	100	100	100	100	90	90	90	90	90	80	80

As the exhibit shows, the asset allocation selected does become more conservative as either γ or λ increase. However, even for very high (and perhaps rather implausible) values of these coefficients, the strategy selected in the U.S. market never allocates less than 80% to stocks.

We suspect that the generally aggressive strategies we find to be optimal in the U.S. and in most other countries are closely related to the length of the retirement period we consider, which is rather standard in the literature. Over long periods such as 30 years, stocks are very likely to outperform bonds by a substantial margin, which pushes the choice toward aggressive strategies.

That said, it may be wise for a retiree to revisit his asset allocation periodically, which would imply doing it on the basis of an ever shorter retirement period. Would a retiree reconsidering his asset allocation every five years always lean toward the rather aggressive strategies we find for a 30-year retirement period? We doubt it. Would a retiree's degree of risk aversion have the same impact on the optimal choice when he expects to live five more years than when he expects to live 30 more years? We doubt it.

These are important questions somewhat beyond the scope of this article, but we address them in separate research already in progress. In fact, our preliminary results for the U.S. show that the shorter is the length of the retirement period we consider, the more the optimal strategy varies as we change a retiree's level of risk aversion. Put differently, for long retirement periods the optimal strategy is very aggressive largely regardless of a retiree's risk aversion, but for short retirement periods the optimal choice varies substantially depending on a retiree's level of risk aversion.

To conclude, note that in our analysis we consider the asset allocation between only two asset classes, stocks and bonds. For only two asset classes, it is straightforward to calculate outcomes for all possible combinations and determine the optimal portfolio. Individuals, and particularly advisors, that want to consider more granular allocations must resort to full-scale optimization in order to identify the utility-maximizing allocation. This methodology applies a numerical algorithm based on evolutionary biology to identify the optimal mix of asset classes; see Cremers et al (2005) for a description and illustration of this approach.

Summary

One of the most important investment decision retirees and those planning for retirement face is the allocation to risky and safe assets. Researchers have put forth a variety of solutions for addressing this question, but in our view, none that adequately assesses a potential asset allocation strategy within the context of an investor's goals. We believe that our framework resolves this issue in two ways.

First, we introduce a new metric of a strategy's performance called the coverage ratio. It equals the fraction of years a strategy sustained an investor's withdrawals during retirement. For an investor who is retired for 30 years, if the strategy supported withdrawals for only 20 years, the coverage ratio would equal 0.67; if the strategy supported withdrawals for 30 years, it would equal 1.00; if it left a bequest that could have supported withdrawals for an additional 10 years, it would equal 1.33. Unlike other metrics used to evaluate asset allocation strategies for retirement, our coverage ratio captures not only failure and success but also the frequency and magnitude of both.

Second, we evaluate the coverage ratio in terms of an investor's preferences, recognizing that an investor is much more displeased with outcomes that fail to fully support retirement than pleased with outcomes in which the strategy leaves a bequest. Specifically, we assume that a retiree's utility increases at a decreasing rate as the coverage ratio rises above 1 and that it decreases linearly as the coverage ratio falls below 1.

We illustrate our utility-based approach by applying it to the choice of stocks and bonds for a wide range of countries over a large number of historical 30-year retirement periods. Our results suggest that retirees in most countries would be well served by selecting fairly aggressive asset allocations at the beginning of retirement. This may seem scary to many investors, but over long holding periods in most countries, compounded stock returns overwhelm compounded bond returns as well as other considerations, such as risk aversion.

Our methodology for evaluating asset allocation strategies has wide-ranging practical applications; to name but a few:

- It could be used to determine the sensitivity of the optimal solution to a variety of considerations such as the relative performance of risky and safe assets, the rate of inflation, the withdrawal rate, the risk aversion of the investor, and the length of the retirement period.
- Although we illustrate our methodology by focusing on the retirement period, it could be used to determine the optimal asset allocation strategy for the pre-retirement accumulation period.
- In our illustration, we hold constant the withdrawal rate and find the optimal asset allocation strategy. Instead, we could hold fixed various asset allocations and find the optimal withdrawal rate for each of them.
- If we focus on the pre-retirement accumulation period, we could apply our methodology to determine the required savings rate for a given asset allocation and target value for a retirement portfolio.
- We apply our methodology to a fixed 30-year retirement period, but it could be applied to progressively shorter retirement periods for the purpose of determining an optimal glidepath.

In short, we believe that our methodology offers a sound and efficient framework for addressing most of the key financial decisions associated with retirement planning. For this reason, the framework we propose should help advisors provide better recommendations to their clients than they are able to provide with other existing approaches.

Appendix

Exhibit A1: Summary Statistics

This exhibit shows, for the series of annual returns over the 1900-2014 period, the arithmetic (AM) and geometric (GM) mean return, standard deviation (SD), semideviation for a 0% benchmark (SSD), lowest return (Min), and highest return (Max). All returns are real (adjusted by each country's inflation rate), in local currency (except for the world market, in dollars), and account for capital gains/losses and cash flows (dividends or coupons). All figures in %.

	AM	GM	SD	SSD	Min	Max
A: Stocks						
Australia	8.9	7.3	17.9	9.2	−42.5	51.5
Austria	4.6	0.6	30.0	15.6	−60.1	127.1
Belgium	5.4	2.7	23.7	13.0	−48.9	105.1
Canada	7.2	5.8	16.9	8.4	−33.8	55.2
Denmark	7.2	5.3	20.7	8.9	−49.2	107.8
Finland	9.3	5.3	30.0	13.9	−60.8	161.7
France	5.7	3.2	23.1	12.3	−41.5	66.1
Germany	8.2	3.2	31.7	14.7	−90.8	154.6
Ireland	6.8	4.2	22.9	11.9	−65.4	68.4
Italy	5.9	1.9	28.5	15.6	−72.9	120.7
Japan	8.8	4.1	29.6	15.2	−85.5	121.1
Netherlands	7.1	5.0	21.4	10.3	−50.4	101.6
New Zealand	7.8	6.1	19.4	9.0	−54.7	105.3
Norway	7.2	4.2	26.9	11.7	−53.6	166.9
Portugal	8.4	3.4	34.4	15.3	−76.6	151.8
South Africa	9.5	7.4	22.1	9.0	−52.2	102.9
Spain	5.9	3.7	21.9	11.0	−43.3	99.4
Sweden	8.0	5.8	21.2	10.8	−42.5	67.5
Switzerland	6.3	4.5	19.5	10.1	−37.8	59.4
UK	7.1	5.3	19.6	9.7	−57.1	96.7
USA	8.5	6.5	20.0	10.4	−37.6	56.3
World	6.6	5.2	17.4	9.4	−41.0	68.2
B: Bonds						
Australia	2.5	1.7	13.2	7.6	−26.6	62.2
Austria	4.9	−3.8	51.2	20.1	−94.4	441.6
Belgium	1.6	0.4	15.0	9.9	−45.6	62.3
Canada	2.8	2.2	10.4	5.4	−25.9	41.7
Denmark	3.9	3.3	11.9	5.1	−18.2	50.1
Finland	1.5	0.2	13.7	10.9	−69.5	30.2
France	1.1	0.2	13.0	9.5	−43.5	35.9
Germany	1.3	−1.4	15.8	12.4	−95.0	62.5
Ireland	2.7	1.6	15.1	8.0	−34.1	61.2
Italy	0.2	−1.2	14.7	11.8	−64.3	35.5
Japan	1.7	−0.9	19.7	14.7	−77.5	69.8
Netherlands	2.2	1.7	9.8	5.2	−18.1	32.8
New Zealand	2.5	2.1	9.0	4.8	−23.7	34.1
Norway	2.6	1.9	12.0	6.8	−48.0	62.1
Portugal	2.5	0.8	18.7	11.2	−49.7	82.4
South Africa	2.4	1.9	10.4	5.9	−32.6	37.1
Spain	2.5	1.8	12.6	7.1	−30.2	53.2
Sweden	3.5	2.8	12.7	5.9	−37.0	68.2
Switzerland	2.7	2.3	9.4	4.3	−21.4	56.1
UK	2.4	1.6	13.7	7.1	−30.7	59.4
USA	2.5	2.0	10.4	5.3	−18.4	35.1
World	2.5	1.9	11.3	6.0	−32.0	46.7

Exhibit A2: Coverage Ratio – Utility-Based Approach – Sensitivity Analysis on γ

This exhibit shows the allocation selected by expression (3), for γ=0 and λ=10, from the 11 allocations considered, ranging from 100 (all stocks) to 0 (no stocks), with nine allocations in between, the rest allocated to bonds. 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	Allocation	Country	Allocation
Australia	100	Netherlands	100
Austria	80	New Zealand	100
Belgium	100	Norway	90
Canada	100	Portugal	70
Denmark	90	South Africa	100
Finland	100	Spain	90
France	100	Sweden	100
Germany	100	Switzerland	100
Ireland	100	UK	100
Italy	100	USA	100
Japan	100	World	100

Exhibit A3: Coverage Ratio – Utility-Based Approach – Sensitivity Analysis on γ

This exhibit shows the allocation selected by expression (3), for γ=0.5 and λ=10, from the 11 allocations considered, ranging from 100 (all stocks) to 0 (no stocks), with nine allocations in between, the rest allocated to bonds. 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	Allocation	Country	Allocation
Australia	100	Netherlands	100
Austria	80	New Zealand	100
Belgium	100	Norway	90
Canada	100	Portugal	60
Denmark	90	South Africa	100
Finland	100	Spain	90
France	100	Sweden	100
Germany	100	Switzerland	80
Ireland	100	UK	100
Italy	100	USA	100
Japan	100	World	100

Exhibit A4: Coverage Ratio – Utility-Based Approach – Sensitivity Analysis on γ

This exhibit shows the allocation selected by expression (3), for γ=2 and λ=10, from the 11 allocations considered, ranging from 100 (all stocks) to 0 (no stocks), with nine allocations in between, the rest allocated to bonds. 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	Allocation	Country	Allocation
Australia	100	Netherlands	70
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	40
Germany	100	Switzerland	40
Ireland	100	UK	100
Italy	100	USA	90
Japan	70	World	100

Exhibit A5: Coverage Ratio – Utility-Based Approach – Sensitivity Analysis on λ

This exhibit shows the allocation selected by expression (3), for γ=0.9999 and λ=1, from the 11 allocations considered, ranging from 100 (all stocks) to 0 (no stocks), with nine allocations in between, the rest allocated to bonds. 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	Allocation	Country	Allocation
Australia	100	Netherlands	100
Austria	70	New Zealand	100
Belgium	100	Norway	90
Canada	100	Portugal	100
Denmark	100	South Africa	100
Finland	100	Spain	100
France	100	Sweden	100
Germany	100	Switzerland	100
Ireland	100	UK	100
Italy	100	USA	100
Japan	100	World	100

Exhibit A6: Coverage Ratio – Utility-Based Approach – Sensitivity Analysis on λ

This exhibit shows the allocation selected by expression (3), for γ=0.9999 and λ=5, from the 11 allocations considered, ranging from 100 (all stocks) to 0 (no stocks), with nine allocations in between, the rest allocated to bonds. 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	Allocation	Country	Allocation
Australia	100	Netherlands	100
Austria	80	New Zealand	100
Belgium	100	Norway	90
Canada	100	Portugal	60
Denmark	100	South Africa	100
Finland	100	Spain	90
France	100	Sweden	100
Germany	100	Switzerland	80
Ireland	100	UK	100
Italy	100	USA	100
Japan	100	World	100

Exhibit A7: Coverage Ratio – Utility-Based Approach – Sensitivity Analysis on λ

This exhibit shows the allocation selected by expression (3), for γ=0.9999 and λ=20, from the 11 allocations considered, ranging from 100 (all stocks) to 0 (no stocks), with nine allocations in between, the rest allocated to bonds. 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	Allocation	Country	Allocation
Australia	100	Netherlands	70
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	40
Germany	100	Switzerland	40
Ireland	100	UK	100
Italy	100	USA	100
Japan	80	World	100

References

Adler, Timothy, and Mark Kritzman (2007). "Mean Variance versus Full Scale Optimization: In and Out of Sample." Journal of Asset Management, 7, 5, 302-311.

Bengen, William (1994). "Determining Withdrawal Rates Using Historical Data." Journal of Financial Planning, 7, 4, 171-180.

Czasonis, Megan, Mark Kritzman, Baykan Pamir, and David Turkington (2018). "A Comparative Analysis of Performance Fees." Journal of Portfolio Management, forthcoming.

Cremers, Jan-Hein, Mark Kritzman, and Sébastien Page (2005). "Optimal Hedge Fund Allocations." Journal of Portfolio Management, 31, 3, 70-81.

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 (2018*a*). "From Failure to Success: Replacing the Failure Rate." Journal of Wealth Management, Journal of Wealth Management, 20, 4, 9-21.

Estrada, Javier (2018*b*). "Replacing the Failure Rate: A Downside Risk Perspective." Journal of Retirement, 5, 4, 46-56.

Milevsky, Moshe (2016). It's Time to Retire Ruin (Probabilities)." Financial Analysts Journal, 72, 2, 8-12.

Gandelman, Néstor, and Rubén Hernández-Murillo (2014). "Risk Aversion at the Country Level." Federal Reserve Bank of St. Louis, Working Paper Series.

Thomas, Philip (2016). "Measuring Risk Aversion: The Challenge." Measurement, 79, February, 285-301.