Journal Of Investment Management, Vol. 14, No. 2, (2016), pp. 5–32

www.joim.com

MARKET RISK, MORTALITY RISK, AND SUSTAINABLE RETIREMENT ASSET ALLOCATION: A DOWNSIDE RISK PERSPECTIVE

W. V. Harlow[a] *and Keith C. Brown*[b]

Despite its clear importance, there is no consensus on the optimal asset allocation strategy for retirement investors of varying age, gender, and risk tolerance. This study analyzes the allocation question by focusing on the downside risks that result from the joint uncertainty over investment returns and life expectancy. Using a new analytical approach, we show that concentrating on the severity of retirement funding shortfalls, rather than just the probability of ruin, markedly increases the sustainability of a retirement portfolio. We demonstrate that for retirement investors attempting to minimize downside risk while sustaining future withdrawals, appropriate equity allocations range between five and 25 percent, levels that are strikingly low compared to those typically found in life-cycle funds. Further, these optimal portfolio constructions appear to vary little with alternative capital market assumptions. We also show that more aggressive investors having substantial bequest motives should still be relatively conservative in their stock allocations. We conclude that the higher equity allocations commonly employed in practice significantly underestimate the risks that these higher-volatility portfolios pose to the sustainability of retirement savings and incomes.

1 Introduction

Few problems that investors face in their lifetimes are likely to be more significant or challenging than that of how to properly allocate the assets in their retirement portfolios. Conventional wisdom holds that an individual's exposure to higher-risk securities, like stocks, should decline as his or her retirement date nears. As Jagannathan and Kocherlakota (1996) note, the rationale for this tenet usually rests on three assumptions: (i) stocks are less risky over longer time horizons than they are over shorter ones, (ii) the higher expected returns associated with equity are necessary to meet the higher obligations typically faced by younger investors, and (iii) younger people have more years in which to produce labor income with which to offset any investment losses. Under these circumstances, a less volatile, increasingly conservative asset allocation pattern makes intuitive sense because a substantial equity market decline late in a person's career could

---

[a] Empower Retirement, One Post Office Square, Boston, MA 02109, USA. E-mail: van.harlow@empower-retirement.com

[b] University of Texas, Department of Finance B6600, Austin, TX 78712, USA. E-mail: keith.brown@mccombs.utexas.edu

affect the ability to fund his or her retirement, or even to retire at all.

Although the earliest studies on the topic (e.g., Samuelson (1969), Merton (1969)) established the conditions under which the optimal allocation between riskless and risky assets *would not* depend on an individual's investment horizon, there is by now considerable empirical support for the concept that allocating retirement assets is a function of an investor's age. In a survey of the allocation practices of various demographic groups, Riley and Chow (1992) find that risk tolerance decreases directly with age for investors over 65, a fact inferred from the increasing reluctance these investors show toward taking risky positions in their portfolios. Schooley and Worden (1999) document that the percentage of equity held by an investor generally increases with age, but then at some point—presumably at retirement—that allocation declines dramatically. Waggle and Englis (2000) show that, controlling for factors such as home ownership, education, marital status, and income level, investors exhibit a significant tendency to reduce the equity allocations in their retirement accounts as they grow older. Poterba and Samwick (2001) also demonstrate that younger investors tend to follow risky allocation strategies, including a greater use of leverage and investment in illiquid assets, than do older investors.

Despite the fact that reducing the risk level of a portfolio near retirement appears to be a widely accepted concept, there does not seem to be any consensus as to what the exact *level* of the equity allocation weight should be, either at the moment of an individual's retirement or throughout the retirement phase. (For the purpose of clarity, we define "retirement" here as the moment when a person begins net draw-downs from his or her accumulated savings to meet living expenses; see Kingston (2000) for more on this point.) While some simple rules of thumb have been expressed—for instance, Bodie and Crane (1997) and Booth (2004) note that investors often set the stock allocation in their portfolios equal 100 minus their age—a wide variety of optimal allocations can result from changing the conditions of the investor's level of wealth (Wachter and Yogo (2010)), risk tolerance (Hariharan *et al.*, 2000) or labor income flexibility (Bodie *et al.*, 1992; Cocco *et al.*, 2005; Gomes *et al.*, 2008).

This diversity of opinion among academics and financial services professionals alike has undoubtedly contributed to surprising differences in the portfolio allocation strategies faced by investors with ostensibly similar problems. Alestalo and Puttonen (2006) report that equity investments across the Finnish defined benefit pension plan industry range from zero to 70 percent of the fund. This disparity also appears clearly in the widely varied asset allocations offered by popular target date, or life-cycle, funds. For example, Figure 1 depicts the dramatic range of equity allocations in the "glide paths"

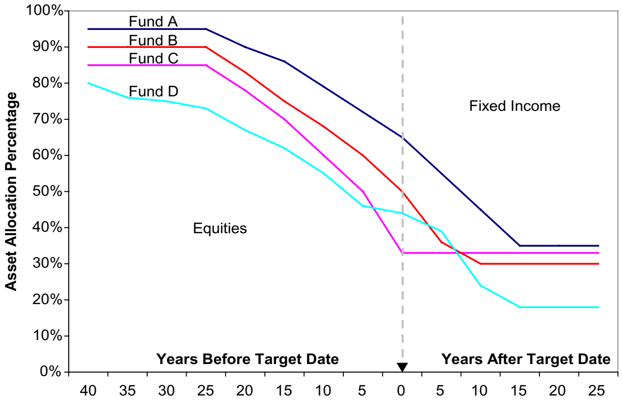

Figure 1 Target date glide paths for selected life-cycle mutual funds.

Source: United States Government accountability office, report GAO-11-118, "Defined contribution plans: Key information on target date funds as default investments should be provided to plan sponsors and participants," January 2011.

(i.e., how the fund's asset allocation scheme is scheduled to be adjusted over time) for four representative target date funds. Specifically, at the target date (presumably age 65), the commitments to the stock market vary from a high of 65 percent to a low of 33 percent, with an average of 48 percent. Clearly, the risk exposure for potential retirees would be significantly different depending on which of these investment products they chose for their retirement savings. Further, in their survey of target date fund investment practices, Elton *et al.* (2015) also note that professional portfolio managers are "very active" in changing their asset allocations and that these adjustments do not ultimately benefit investors.

Of course, the problem that these discrepancies create for individuals forced to make their own retirement investment decisions is that they may actually be harmed financially by selecting the wrong target date fund. Using simulated data, Schleef and Eisinger (2007) show that the dynamic nature of the equity allocation adjustments made by some life-cycle funds fail to increase the likelihood of reaching a targeted portfolio goal compared to a static allocation approach. Additionally, in their empirical analysis of 68 target date funds, Balduzzi and Reuter (2012) find a tremendous degree of heterogeneity in the performance of portfolios with the same target dates, which they attribute to pronounced differences in both equity allocation proportions and the systematic risk exposures of those equity positions. Finally, Basu and Drew (2009) demonstrate that it is possible that "contrarian" portfolios, which shift equity allocations opportunistically, can outperform life-cycle funds in certain market environments. These possibilities are at least somewhat troubling, given Merton's (2012) argument that intelligent financial product design can be an effective way of combatting deficiencies in consumer financial education.[1]

The purpose of this study is to take a closer look at this important decision for an investor who has already reached the point of retirement. To do so, we employ a unique set of analyses that summarizes the risk-return tradeoffs which go hand in hand with the asset allocation choice. An important feature of our approach is that we consider both financial market risk *and* investor mortality risk to be stochastic elements of the problem to be solved. For different amounts of withdrawal from retirement savings, we optimize the retirement portfolio's asset allocation mix so as to minimize three alternative measures of the risk of plan failure: probability of ruin and two measures of downside risk (i.e., expected shortfall and semi-deviation). A key finding of this analysis is that downside risk-based portfolios survive longer in worst-case scenarios than those minimizing the simpler probability of failure measure.[2]

Using downside risk-based metrics to minimize the severity of retirement outcomes, we then investigate how these allocations might change with varying sets of assets, market conditions, and assumptions about investor goals. Our analysis suggests that when limiting one's downside risk, the optimal asset allocation across a wide range of settings is strikingly conservative in terms of exposure to equities—far more conservative than those typically deployed in practice. In particular, we show that the optimal amount of equity in the portfolio of an investor between the ages of 65 and 85 can be as low as 5 percent and seldom needs to be greater 25 percent, considerably less than the almost 50 percent stock weight maintained by the average life-cycle fund in the marketplace. Even in situations where individuals want to take more risk in order to increase the potential value of the remaining assets to be left to their heirs, the range of equity allocations is still surprisingly conservative vis-à-vis conventional wisdom.

The remainder of the study is structured as follows. In the next section, we describe the *retirement present value* methodology central to our simulation analysis. Importantly, we demonstrate how this approach can combine the stochastic nature of both investment returns and the individual's life expectancy when developing a decision rule about how assets should be invested in order to minimize the shortfall risk that the investor will deplete his or her resources too soon. The following section then extends this discussion to consider minimum risk allocations under a wide variety of alternative scenarios, while the two subsequent sections provide sensitivity analysis on the simulation's range of assumed capital market conditions and investor bequest objectives, respectively. The discussion in the final section concludes the paper.

2 Optimizing retirement asset allocations: A new approach

2.1 A descriptive overview of the retirement present value process

Any individual investor faces multiple unknowns when planning for a retirement that could last for 20 to 30 years, or more. The most significant unknown variables are the future investment returns on retirement savings as well as the length of a person's remaining life span. When considering the retirement portfolio investment decision, more aggressive asset allocations have the potential to deliver higher average returns over time and thereby support longer retirement periods. Conversely, the higher risk and volatility inherent in more aggressive approaches to investing also raise the risk of depleting the investor's assets prematurely, thereby causing the retirement plan to fail.

An attractive way to reduce this uncertainty and more accurately evaluate the financial tradeoffs and overall health of a retirement plan is through the use of a method known as retirement present value (RPV) analysis. This technique considers the retirement plan to consist of both current and future assets and liabilities. Under this approach, savings contributions, for example, are treated as both assets and flows into the portfolio. The value of these assets fluctuates with uncertain and volatile investment returns over time. Retirement expenses, conversely, are both current and future liabilities reflected in the form of outflows from the portfolio.

Of course, the duration of any specific retirement plan will vary because of the uncertainty of how long one will live, which is the mortality risk problem faced by the investor. However, RPV analysis can integrate all of these dynamic components—flows, returns, and investor longevity—and then discount them into a positive or negative value expressed in current dollars. So, rather than simulating asset returns to project the value of a retirement portfolio out to some arbitrary point in the future (e.g., to age 85) when mortality is assumed, the simulated returns are used as discount factors in the RPV framework to compute the present value of future retirement cash flows. In this framework, mortality risk is then captured by weighting these projected cash flows by the probability of the investor being alive at each respective point in the future. In this context, a positive RPV statistic indicates the likelihood of having some assets remaining at the end of life, with higher levels clearly indicating better outcomes. Conversely, a negative RPV implies the possible (or even probable) depletion of all retirement assets well before death. Thus, the more negative the RPV statistic, the worse the outcome for the investor.[3]

One important caveat to note is that, for any given retirement plan, there is no single RPV amount, but rather a distribution of those present values. This is because of the “synergy” created by the uncertainty of future investment returns compounded by the uncertainty of how long the individual will live. If the distribution of RPV outcomes is completely positive (or nearly so)—which is to say that virtually all potential RPV states are greater than zero—then we would expect a successful retirement outcome with a high degree of confidence. Conversely, a highly negative RPV distribution suggests a situation in which an individual is highly likely to outlive his or her retirement resources for the assumed withdrawal rate.

Leaving aside the issue of providing a bequest to heirs, the theoretically perfect retirement plan would be one in which the RPV would be exactly zero. In that unique case, a person would have in place precisely the right amount of retirement funds to spend before expiring. In reality, of course, individuals’ retirement plans have a range of possible outcomes, from outliving their resources to dying early and leaving a sizeable unspent inheritance. For planning purposes, one reasonable goal would be to reduce the possibility of a negative RPV (i.e., the probability of failure, or ruin). An even more relevant goal might be to minimize the range or magnitude of potential portfolio shortfalls. In other words, beyond considering the possibility of failure, an investor might also want to minimize the magnitude and severity of the possible negative RPV levels (i.e., downside risks) to which a portfolio is subject.

2.2 *The mechanics of the RPV process*

As summarized by the preceding discussion, the RPV method is simply an expression of the financial value of a retirement plan in today’s dollars. It captures both mortality risk and the uncertainty around investment returns by discounting the cash inflows and outflows associated with the retirement plan in an appropriate manner. The calculation of RPV is straightforward and merely an adaptation of the familiar method of determining the discounted present value of a series of future cash flows. Mathematically, the equation for the probability-weighted discounted cash flows is:

$$\mathrm{RPV} = \sum_{t=0}^{\infty} \frac{p_t \mathrm{CF}_t}{(1 + R_t)^t} \tag{1}$$

where:

t = years into the future,

p_t = probability of being alive at time t,

CF_t = cash flow at time t, and

R_t = the period-specific risk-adjusted discount rate.

The forecasted cash flows of the retirement plan, CF_t in Equation (1), represent savings inflows into the portfolio prior to retirement age and the outflows from living expenses deducted after retirement. CF_0 in the RPV analysis represents the individual’s current savings at any desired time $t = 0$. The probability of being alive at future time t, p_t, can be obtained in a variety of ways: case histories for individuals with similar health statuses, directly from aggregated actuarial tables, or through standard mathematical models specified to approximate the actual probability values.[4] Finally, it is worth noting that the RPV expression in Equation (1) can be seen as a generalized version of the continuous-time, closed-form model of Milevsky and Robinson (2005) in that it allows for a user-specific set of mortality assumptions in discrete time, which makes it adaptable to any investor’s particular circumstances (e.g., current health status, spending needs).

In order to determine the discount rate applicable at each time t (i.e., R_t), the returns on the retirement fund investment portfolio available in each year are used. These returns, denoted by r_t, can be obtained from historical time series or through Monte Carlo simulation. Assuming for convenience (but without loss of generality) that the asset classes included in the investment portfolio are stocks, bonds, and cash equivalents, the discount rate can be expressed:

$$(1 + R_t)^t = (1 + r_1)(1 + r_2) \times (1 + r_3) \cdots (1 + r_t) \qquad (2)$$

where:

$$r_t = (w_{\mathrm{S}} \times r_{\mathrm{St}}) + (w_{\mathrm{B}} \times r_{\mathrm{Bt}}) + (w_{\mathrm{C}} \times r_{\mathrm{Ct}}). \qquad (3)$$

In Equations (2) and (3), we have the following definitions:

w_{S}, w_{B}, and w_{C} are the portfolio weights in stocks, bonds, and cash, respectively;

and:

r_{St}, r_{Bt}, and r_{Ct} are the real returns on stocks, bonds, and cash at time t, respectively.

In this study, we only consider the situation in which the investment portfolio is in a state of net withdrawals. In other words, we assume that the investor is already in retirement and so all future portfolio cash flows, CF_t, are negative. Said differently, we do not include any periods of savings accumulation subsequent to the date of retirement. It is also important to note that these spending amounts are assumed to be exogenous, pre-specified values represented in constant, inflation-adjusted dollars, which is typical of how someone setting his or her asset allocation policy at time $t = 0$ in light of projected future spending needs is likely to look at the problem. The Appendix provides an example of the RPV calculation for one realization of a sequence of future asset returns as well as for a specific mortality risk projection.

2.3 *The risk metrics of retirement asset allocation*

To show how this analysis might work in practice, Figure 2 provides the RPV distribution of a representative investor's retirement plan. In this case, a 65-year-old male has $100 in current retirement savings. He has just retired and plans to spend $7 per year in the future, expressed in real terms. Throughout the retirement period, we assume that his savings are invested in a constant mix of stock, bonds, and short-term instruments (to simplify things henceforth, we will continue to refer to these short-term instruments as "cash"). In this particular example, the allocation to stocks is 11 percent; bonds, 24 percent; and cash, 65 percent. (This is, in fact, the allocation that minimizes retirement downside risk for this retiree, as we will formally define shortly.) We also make the base case assumptions that stock, bonds, and cash have expected real returns of 6.0, 3.0, and 1.0 percent, respectively, as well as having respective volatilities of 16.0, 7.0, and 2.5 percent.[5]

The RPV analysis for this example shows a wide distribution of possible outcomes. Conditional on a spending level of $7, the retirement plan has, on average, a value of $10.21 (median value of $10.75). Thus, in current dollars, this is the net value of the plan—the present value of assets minus the present value of future liabilities. An alternative interpretation of this mean RPV is that it represents the amount that our retiree can expect to leave to his heirs, quoted in today's dollars. Also evident in Figure 2, however, is the fact that there is a range of negative RPVs that represent potential unsuccessful retirement outcomes, that is, total asset depletion. In fact, 9.96 percent of the possible outcomes have a negative present value, which thus can be considered the probability of ruin, or plan failure. This is equivalent to stating that there is a roughly one-in-ten chance of exhausting the portfolio's assets at some point

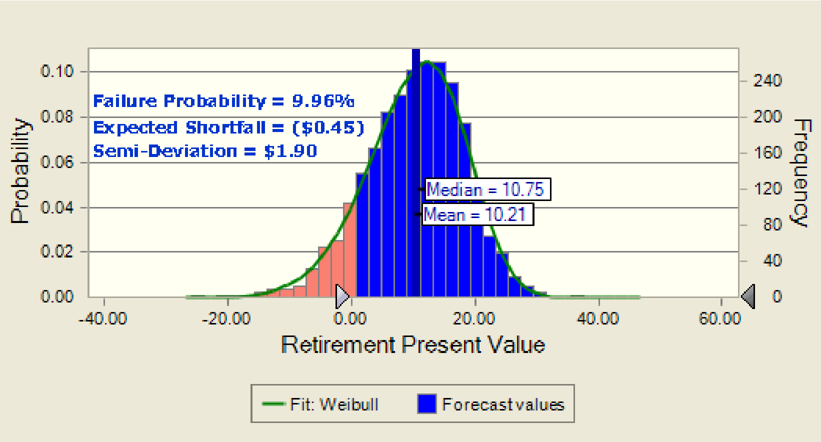

Figure 2 Example RPV distribution for a retirement plan minimizing risk for $7 real annual spending per $100 in savings for a 65-year-old male.

Note: The analysis assumes that a 65-year-old individual has $100 in retirement savings and plans to spend $7 per year, adjusted for inflation. The analysis ignores taxes and transaction costs. Mortality is modeled using the United States Social Security Administration's period life tables. The distribution displayed represents a histogram of 10,000 sequences of annual potential portfolio returns, using the following capital market assumptions: (i) Expected real return: 6.0 percent (stock), 3.0 percent (bonds), 1.0 percent (cash); (ii) Volatility: 16.0 percent (stock), 7.0 percent (bonds), 2.5 percent (cash); and (iii) Correlation: 0.20 (stock/bonds), 0.15 (stock/cash), 0.35 (bonds/cash). The discount factors used in the RPV computation for this particular distribution assume the following asset class allocation weights: 11 percent (stock), 24 percent (bonds), 65 percent (cash).

before the retiree's death. Negative RPV outcomes can be thought of as situations in which the retiree would have to borrow money from his heirs in order to support the desired spending level. Empirically, the probability of ruin is represented as a zero-order, lower-partial moment of the distribution of RPV outcomes about a target value of zero, and it is expressed as:

Probability of ruin:

$$\mathrm{LPM}_0 = \sum_{\mathrm{RPV}_j < \tau} \frac{(\tau - \mathrm{RPV}_j)^0}{n - 1} \tag{4}$$

where $\mathrm{RPV}_j =$ the j-th RPV outcome from the set of n observations using Equations (1)–(3), as generated from return simulations, and $\tau =$ target value of zero (i.e., $0.00).

Much of the extant literature on retirement asset allocation and spending decisions focuses on minimizing the probability of ruin expression in Equation (4) as the primary goal to be achieved. For instance, Rook (2014) formulates a dynamic programming problem that explicitly incorporates the probability of ruin into a multi-period objective function organized to establish an optimal annual withdrawal rate from the retirement portfolio. Stout (2008) also considers the problem of "optimal withdrawal management" from a retirement fund in a stochastic optimization setting with the aim of limiting the investor's probability of prematurely exhausting portfolio assets. Lastly, Milevsky and Robinson (2005) adapt a traditional discounting approach to include continuous-time equations for investment and longevity risk while also defining a failure to be a situation in which the present value of future spending needs exceeds the initial value of the investment portfolio (i.e., ruin, expressed in current dollars).

The popularity of using Equation (4) to describe what is meant by retirement failure makes it a useful starting point for considering how an individual might frame his or her goal in forming a sustainable retirement allocation. Indeed, for an investor who has no potential safety net in retirement (e.g., borrowing capacity, family members), the possibility of any negative RPV realization truly might represent failure. However, it is also possible that a simple probability of ruin measure is too basic in other circumstances, given that it misses some important dimensions of the problem that are crucial to making an optimal decision. In particular, by focusing on just the probability of a realizing a negative RPV, the LPM_0 statistic fails to account for the magnitude of any such shortfall. Thus, under Equation (4), a RPV distribution with a 10 percent probability of ruin would always be viewed as riskier than one with a 5 percent failure chance, even if the smallest possible value in the former is −\$0.01 whereas the latter might have potential RPV outcomes ranging as low as, say, −\$20.00. Clearly, many investors would not consider the slightly higher probability of falling short by as little as −\$0.01 to actually be a riskier outcome than the possibility of a significant negative tail-risk event (i.e., −\$20.00) occurring with a non-trivial, albeit slightly lower, probability. Thus, a statistical measure such as Equation (4) that ignores this magnitude differential might not be an accurate characterization of the objectives that a substantial group of investors hopes to accomplish in retirement.

In fact, probability of ruin is just one example of wider collection of statistics that can be used to describe retirement failure. Downside risk measures expand the class of objective functions by capturing both the probability and the magnitude of a potential shortfall (relative to τ, the investor's "pain threshold"). In our analysis, we consider two downside risk statistics, both of which can be defined as higher-order LPM statistics. First, the expected shortfall of the retirement portfolio can be written as a first-order LPM function as follows:

Expected Shortfall:

$$\mathrm{LPM}_1 = \sum_{\mathrm{RPV}_j < \tau} \frac{(\tau - \mathrm{RPV}_j)^1}{n-1} \tag{5}$$

where RPV_j and $\tau = \$0.00$ are defined as before. Notice that the statistic in Equation (5) calculates the average of just the negative RPV values in the projected distribution; this is the sense in which it is called an expected shortfall. Given the construction of the formula, this value is expressed in dollar terms—rather than in percentage form, as with Equation (4)—which provides a more intuitive interpretation. For example, for the distribution shown in Figure 2, the LPM_1 statistic is −\$0.45. This suggests that although the probability of a shortfall is around 10 percent, the average level of such a failure is quite small relative to the \$100 initial portfolio value.

Another downside risk measure can then be defined relative to the second-order lower-partial moment. This approach to measuring an investor's downside portfolio exposure is actually more popular in practice than Equation (5) because of its analytical correspondence to the variance (or standard deviation) statistic for the RPV distribution in question. Specifically, the semi-deviation is the square root of the average of the squared negative RPV outcomes:

Semi − Deviation:

$$\mathrm{LPM}_2 = \left[\sum_{\mathrm{RPV}_j < \tau} \frac{(\tau - \mathrm{RPV}_j)^2}{n-1} \right]^{1/2} \tag{6}$$

where the investor's target value is once again set at \$0.00. Like the expected shortfall metric, the semi-deviation is a more valuable assessment of the magnitude of the depletion risk problem than the probability of portfolio ruin because it captures the severity of the unsuccessful outcomes, some of which could be devastating. For example, Figure 2 indicates potential adverse outcomes that could range as high as $-\$20.00$, suggesting that there are combinations of market and mortality events that would have actually required 20 percent *more* in initial savings (i.e., \$20 plus the original \$100) at age 65 to completely fund a successful retirement involving the real spending goal of \$7 per year. In this example, the LPM_2 value is shown to be \$1.90.

Based on the distribution of RPV values illustrated in Figure 2, the metrics just discussed provide a convenient way to summarize the financial characteristics and overall sustainability of an investor's retirement plan. Retirement risk is captured by either the simple probability of a plan failure or by its downside exposure, which in turn can be measured as the average potential shortfall or as the standard deviation of that shortfall. The overall health and net value of the plan are represented by the average RPV.[6]

One final methodological issue merits consideration. The optimal asset allocations generated throughout this study are selected so as to minimize retirement downside risk for any given scenario. Given the complex nature of the problem we are examining, we utilize a stochastic optimization process to establish the best asset allocation mix for any set of capital market and mortality risk assumptions. This approach is different than conventional optimization procedures in that thousands of simulations are made with each step of the algorithm in its search for the best solution.[7] Mathematically, this is expressed as:

$$\text{Select } \{w_k\}(= \{w_S, w_B, w_C\})$$
$$\text{so as to minimize } \text{LPM}_X \qquad (7)$$

subject to:

(i) $\sum_k w_k = 1$; and
(ii) All $w_k > 0$.

In this formulation, RPV_j is once again defined by Equations (1)–(3) and $X = 0$, 1, or 2 in order to represent the probability of ruin or the two measures of downside risk, respectively.[8]

3 Sustainable retirement asset allocation: Probability of ruin versus downside risk

The RPV distribution shown in Figure 2 reflects the range of potential retirement net present values that arise due to mortality and investment uncertainty, given a pre-specified spending policy. A natural question to now ask is: How can we alter the shape of this distribution to best conform to an investor's preferences? In other words, what is the best objective (or loss) function to use in optimizing this distributional shape in order to improve the funding sustainability of the retirement portfolio?

There are two inputs into this decision that affect the shape and location of the resulting RPV distribution. The first is the investor's spending policy. As the amount of the annual withdrawals from retirement savings increases, the distribution will shift to the left with more of the potential outcomes being negative. The individual is therefore faced with the decision of trading off an increase in retirement income with an increase in the likelihood of negative outcomes (i.e., exhausting resources during retirement). The second input is the asset allocation strategy: What is the stock, bond, and cash mix that best shapes the RPV distribution, conditional on the selection of a given spending policy? In this context, "best" can mean any number of specifications of the RPV distribution that consider its mean, dispersion, skewness, etc.

As noted earlier, much of the financial literature on the sustainability of retirement income focuses on the basic probability of ruin notion. In other words, the probability of a negative outcome (the area of the red bars in Figure 2) becomes the operative metric to use in assessing sustainability and evaluating the tradeoffs between spending and risk. However, the literature on utility theory and the distribution of wealth have shown that simple measures of downside probability fail to adequately capture many salient dimensions of an investor's level of risk tolerance. For instance, as Bawa (1975, 1978), Bawa and Lindenburg (1977), Harlow and Rao (1989), and Harlow (1991) note, portfolio formation decisions based

Table 1 Asset allocations that minimize the probability of ruin and downside risk.

	65-year-old male						
Spending rate	Stocks	Bonds	Cash	LPM_0	LPM_1	LPM_2	Mean RPV
A. Minimize LPM_0 (*Probability of Ruin*)							
$6	8%	12%	80%	0.10%	($0.00)	$0.11	$20.15
$7	24%	56%	20%	7.39%	($0.55)	$2.75	$18.40
$8	44%	56%	0%	23.32%	($3.19)	$8.96	$11.64
B. Minimize LPM_1 (*Expected Shortfall*)							
$6	6%	15%	79%	0.10%	($0.00)	$0.10	$20.16
$7	14%	33%	53%	8.62%	($0.43)	$2.00	$12.40
$8	30%	70%	0%	24.59%	($2.86)	$7.83	$9.86
C. Minimize LPM_2 (*Semi-Deviation*)							
$6	5%	14%	81%	0.12%	($0.00)	$0.10	$19.59
$7	11%	24%	65%	9.96%	($0.45)	$1.90	$10.21
$8	20%	46%	34%	34.64%	($3.23)	$7.19	$4.01
	65-year-old female						
A. Minimize LPM_0 (*Probability of Ruin*)							
$5	7%	22%	71%	0.03%	($0.00)	$0.05	$24.42
$6	20%	54%	26%	6.03%	($0.43)	$2.36	$18.96
$7	41%	59%	0%	22.46%	($3.11)	$8.91	$12.23
B. Minimize LPM_1 (*Expected Shortfall*)							
$5	4%	16%	80%	0.03%	($0.00)	$0.04	$22.38
$6	14%	33%	53%	6.97%	($0.36)	$1.87	$14.01
$7	30%	70%	0%	23.49%	($2.92)	$8.15	$10.72
C. Minimize LPM_2 (*Semi-Deviation*)							
$5	3%	13%	84%	0.06%	($0.00)	$0.03	$21.68
$6	11%	25%	64%	8.29%	($0.38)	$1.80	$11.51
$7	21%	49%	30%	32.82%	($3.30)	$7.62	$4.96

Note: Spending rates represent the annual, inflation-adjusted withdrawal rates per $100 in retirement savings. Probability of ruin is measured by the zero-order lower-partial moment (LPM_0) of negative RPV outcomes while downside risk is measured by the first-order (LPM_1) and second-order (LPM_2) lower-partial moments, respectively.

on statistics such as the probability of ruin (i.e., LPM_0) are consistent only with judgments made by an individual possessing an upward sloping utility function ($U(w)$) for wealth (w), or $\mathrm{d}U(w)/\mathrm{d}w > 0$. On the other hand, risk aversion is only addressed through the use of the first- and second-order lower-partial moments of the distribution, or the expected shortfall (LPM_1) and semi-deviation (LPM_2) measures, respectively. These higher-order moments have the advantage of capturing the magnitude of shortfalls as well as the simple probability of achieving some sort of negative outcome.

One of our objectives in this study is to contrast the use of measures like the probability of ruin with more general measures of downside risk when used in conjunction with the RPV approach to retirement decision making. Accordingly, Table 1 provides a comparison of the asset allocations and risk statistics for various scenarios when the probability of ruin is minimized versus when controlling downside risk defines the investor's objective function. For this illustration, we list separate findings for both of the ways described earlier to think about downside risk: expected shortfall (LPM_1) and semi-deviation (LPM_2).

The upper panel of the exhibit illustrates the results for a male investor retiring at age 65. As a comparison, consider the RPV distribution metrics for the cases where the retirement spending rate is \$7 annually per \$100 saved. For plan failure probability (i.e., LPM_0) minimization, the probability of ruin is 7.39% and the downside risk measures are $-\$0.55$ for expected shortfall and \$2.75 for semi-deviation. By contrast, for LPM_1 minimization, these values are 8.62%, $-\$0.43$, and \$2.00, respectively, whereas when minimizing LPM_2 is the retiree's objective, the respective statistics are 9.96%, $-\$0.45$, and \$1.90, respectively. From these reported values, it is readily apparent that the minimum levels of every metric are indeed obtained when minimizing that specific outcome is the intended objective.

Perhaps the most interesting finding indicated in Table 1 involves the range of asset allocation strategies that result from the stochastic optimization process using each of the three objectives functions (i.e., LPM_0, LPM_1, and LPM_2). In particular, we see that the asset allocations associated with all of these measures of retirement failure are far more conservative in terms of the percentage of the portfolio invested in equity than the 50 percent level observed in practice for the typical target date fund. Focusing again on the 65-year-old male retiree with a \$7 annual spending need, the optimal levels of stock associated with the three objective functions are 24 percent (probability of ruin), 14 percent (expected shortfall), and 11 percent (semi-deviation). Furthermore, notice that the downside risk minimizations themselves lead to far more conservative allocations that do the simple probability of ruin. Thus, when limiting downside risk is the retiree's goal, the equity weight in the portfolio is at least ten percentage points *lower* than that associated with a more basic plan failure objective. However, it should also be noted that these downside risk minimization benefits (i.e., lower downside risk statistics, more conservative portfolio) do come with lower average RPVs: \$12.40 for LPM_1 and \$10.21 for LPM_2 versus \$18.40 for LPM_0.[9]

Returning to our original question, which of these strategies provides the best funding sustainability in this example? In order to provide a meaningful basis for addressing this issue, Figure 3 displays the frequency distributions of the ages at which the portfolios become depleted. In this comparison, we concentrate on the two most divergent objective strategies shown in Table 1: probability of ruin (LPM_0) versus semi-deviation downside risk (LPM_2). The blue bars for the downside

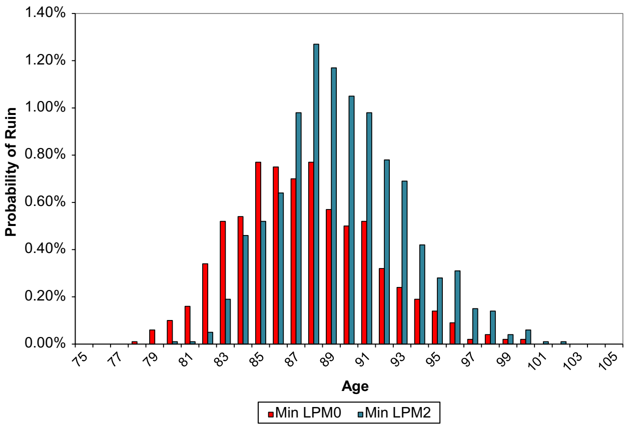

Figure 3 Comparing age at which retirement resources are depleted when asset allocation strategy minimizes probability of ruin (LPM_0) versus downside risk (LPM_2).

Note: The analysis assumes that a 65-year-old individual has $100 in retirement savings and plans to spend $7 per year, adjusted for inflation. The analysis ignores taxes and transaction costs. Mortality is modeled using the United States Social Security Administration's period life tables. Probability of ruin is measured by the zero-order lower-partial moment (LPM_0) of negative RPV outcomes while downside risk is measured by the second-order (LPM_2) lower-partial moment.

risk-based strategy and the red bars are for the probability of ruin-based strategy. It is clear that the downside risk-based portfolio survives longer in the worst-case scenarios than does the portfolio based on the probability of ruin objective. The downside risk portfolio has a non-zero probability of ruin beginning at age 80 versus at an age of 78 for the LPM_0-based portfolio. On average, the downside risk-based allocation, when it does become fully depleted, exhausts retirement assets at age 89.65. The probability of ruin-focused portfolio runs out of assets at age 87.41, more than two years sooner.

We should once again mention that while the downside risk-based allocation lasts longer in worst-case scenarios, it also has a somewhat higher overall probability of ruin: 9.96% versus 7.39%. However, as we discussed in the preceding section, many investors may well prefer this tradeoff, particularly those who are downside risk-averse. Only for individuals with explicit bequest motives—which may not be common in practice—would riskier asset allocations having higher average RPVs be the preferred solution. From a sustainability perspective, focusing on minimizing the severity of negative RPV outcomes rather than just on plan failure generates portfolio solutions that last longer in worst-case scenarios, a point that is missed by the literature relying exclusively on LPM_0 as the investor's objective function in retirement.

4 Minimum downside risk allocations

Having motivated the use of downside risk minimization in determining a sustainable retirement portfolio, we now turn our attention to establishing optimal asset allocation strategies for investors of different ages, different genders, and with different spending objectives. We consider the allocation problem for an individual who is primarily concerned with achieving a successful retirement, defined as minimizing both the probability and magnitude of any retirement failure. As noted, this goal is arguably the main concern for many individuals. However, notice that it differs sharply from the goal of individuals who also wish to leave a bequest of assets to their heirs; we also examine this case in a subsequent section.[10]

Table 2 provides the minimum downside risk allocations and retirement plan summary statistics for a wide range of scenarios. For both male and female investors, the display lists the optimal risk-minimizing asset allocations for retirees aged 65, 75, and 85, where the investment weights are assumed to be constant from each age going forward. A different set of three spending rates are also shown for each gender and age. These withdrawal levels were chosen to reflect a low,

Table 2 Asset allocations that minimize retirement downside risk in the three-asset class case: Stocks, bonds, and cash.

Spending rate	Stocks	Bonds	Cash	LPM$_0$	LPM$_2$	Mean RPV
			Male			
A. 65-year-old						
\$6	5%	14%	81%	0.12%	\$0.10	\$19.59
\$7	11%	24%	65%	9.96%	\$1.90	\$10.21
\$8	20%	46%	34%	34.64%	\$7.19	\$4.01
B. 75-year-old						
\$11	9%	17%	74%	4.84%	\$0.75	\$9.81
\$11.5	11%	21%	68%	14.00%	\$1.82	\$6.71
\$12	14%	27%	59%	27.80%	\$3.39	\$4.13
C. 85-year-old						
\$21	2%	6%	92%	0.60%	\$0.13	\$9.49
\$22	6%	13%	81%	7.68%	\$0.71	\$6.38
\$23	9%	19%	72%	24.60%	\$2.07	\$3.14
			Female			
A. 65-year-old						
\$5	3%	13%	84%	0.06%	\$0.03	\$21.68
\$6	11%	25%	64%	8.29%	\$1.80	\$11.51
\$7	21%	49%	30%	32.82%	\$7.62	\$4.96
B. 75-year-old						
\$9	8%	17%	75%	2.28%	\$0.46	\$12.25
\$9.5	11%	21%	68%	9.84%	\$1.49	\$8.58
\$10	14%	27%	59%	23.00%	\$3.18	\$5.47
C. 85-year-old						
\$17	2%	5%	93%	0.60%	\$0.11	\$10.31
\$18	6%	14%	80%	8.08%	\$0.82	\$6.71
\$19	10%	21%	69%	28.36%	\$2.57	\$2.95

Note: Spending rates represent the annual, inflation-adjusted withdrawal rates per \$100 in retirement savings. Probability of ruin is measured by the zero-order lower-partial moment (LPM$_0$) of negative RPV outcomes while downside risk is measured by the second-order (LPM$_2$) lower-partial moment.

moderate, and high retirement expenses relative to a starting pool of \$100 in retirement savings. The moderate annual spending rate was selected so that the probability of failure would be around 10 percent (i.e., a level often used in practice to define a reasonable withdrawal amount). The low withdrawal case reflects a probability of failure less than 5 percent; the high withdrawal case is associated with failure probabilities ranging from 20 to 30 percent. For example, as shown in the upper panel of the exhibit, the moderate spending rate for a 65-year-old male is \$7 per \$100 in savings. For 75- and 85-year-old males, the moderate spending rates are \$11.5 and \$22, respectively.

We can draw several important conclusions from the findings summarized in Table 2. Notice that each of the risk-minimizing asset allocation mixes is quite conservative, with virtually all equity weights being less than 20 percent. For sustainable and low spending rates where the probability of failure is 10 percent or less, the equity allocations tend to be in the 5 to 10 percent range. These equity exposures are significantly lower than those we saw in Figure 1 for typical retirement products, as well as those reported by Balduzzi and Reuter (2012) and Elton *et al.* (2015) in their examinations of target date fund usage. In addition, notice that for the same level of risk, the spending rate for females is lower than that for males. For example, a 65-year-old male spending \$7 has roughly the same risk and RPV profile as a 65-year-old female spending \$6. The same result holds at the \$6 withdrawal amount for males and the \$5 level for females. As stated earlier, this simply reflects the fact that females have a longer life expectancy and consequently need their investable savings to support a lengthier retirement period. Alternatively, for the same spending rate, the equity allocation for females would need to be larger than that for males to support the longer retirement. For example, at the \$6 spending rate, the optimal equity level is 11 percent for females versus only 5 percent for males. With a \$7 spending rate, the downside risk-minimizing equity allocation is 21 percent for women versus 11 percent for men.

It should be noted that the overall range of these reported spending rates are higher than those normally indicated for retirees by financial advisors. Often at age 65, for instance, a 3 or 5 percent spending rate is used as a rule-of-thumb guideline that should sustain an individual's retirement; see, for example, Scott *et al.* (2009) for an extensive discussion on this topic. However, most financial planning tools do not incorporate the effects of mortality risk on expected spending levels. Here, with mortality risk included in the optimization process, a sustainable annual spending rate for males of \$7 would be appropriate, while the comparable policy for females would be \$6. Conversely, for an individual who expects to live up to, say, age 95, the lower spending levels would be recommended.[11]

One final observation from Table 2 worth mentioning is that the optimal equity exposure level does not change much throughout an individual's retirement period. For instance, the equity allocation in the moderate spending case for a male investor is 11, 11 and 6 percent at ages 65, 75, and 85, respectively. For the female investor, the comparable investment proportions are also 11, 11 and 6 percent. Thus, at least for the first part of an investor's retirement period, the downside risk-minimizing stock allocations are fairly constant as well as being relative low. By contrast, the allocations to bonds and cash for both genders indicate a somewhat more conservative profile as age increases.

While the asset allocations in Table 2 represent those that minimize LPM_2-based downside risk in an RPV context, there are more practical implications that merit consideration as well. Specifically, we examine whether these risk-minimizing asset class allocations provide outcomes that would actually be preferred by a risk-averse individual. To address this issue, we simulate 10,000 retirement scenarios for a 65-year-old male who has \$100 in initial savings and plans to withdraw \$7 per year to pay for his inflation-adjusted retirement expenses. We compare the distribution of outcomes using the appropriate allocation in Table 2 against those that would be associated with the four popular target date funds illustrated in Figure 1. For each of these five simulated asset allocations, we capture the worst outcome, which we define here as the one generating the fewest number of funded retirement years. We also capture the 95 percent confidence level (i.e., 95 percent of the other allocations in the distribution produce a larger number of funded retirement years).

Figure 4 displays these results. In terms of the worst-case outcomes, the minimum downside risk allocation described in this study provides funding for 17 years, or out to age 82 for the retiree. The target date funds, on the other hand, can only be expected to remain solvent for between 11 and 13 years (i.e., to ages 76–78). The risk-minimizing asset allocation, therefore, is able to fund four to six more years of retirement expenses. Similarly, using the 95 percent confidence level criterion, the minimum downside risk allocation results in two to five years more funding compared with the four target date funds. In percentage terms, these are meaningful improvements in the number of years of funded retirement, ranging from an increase of 31 percent to more than 55 percent on a worst-case basis, and from nine to 25 percent based on the 95 percent confidence level rule. Taken together, these results suggest that the minimum risk allocations offer a significant increase in the sustainability of retirement compared to that provided by target date funds typically available in the marketplace.

From an asset allocation perspective, most existing retirement products and recommendations include only a modest amount of cash-equivalent instruments. Certainly, most such portfolios employ far less than the optimal cash allocations documented in Table 2. Often, the cash allocation is limited to 5 to 15 percent, if not excluded from consideration altogether. Thus, the typical allocation decision in practice is really one that comes down a choice between stocks and bonds. As we see from Table 2, however, a substantial commitment to short-term interest instruments is needed to minimize retirement downside risk across a wide variety of ages and spending rate policies. We now turn our attention to how the risk-minimizing allocations change when the availability of these cash-oriented investment vehicles is eliminated.

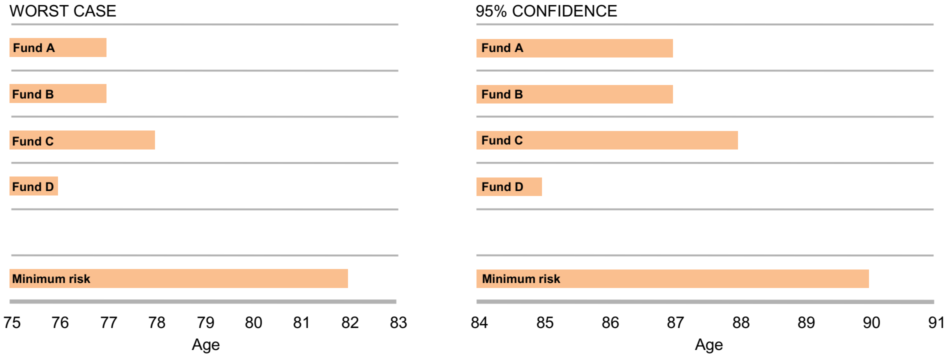

Figure 4 Comparison of the number of retirement funding years for a 65-year-old male spending $7 real annually with $100 in savings: Downside risk-minimizing portfolio versus existing target date funds.

Note: The analysis assumes that a 65-year-old individual has $100 in retirement savings and plans to spend $7 per year, adjusted for inflation. The analysis ignores taxes and transaction costs. Mortality is modeled using the United States Social Security Administration's period life tables. The asset allocation weights for the downside risk-minimizing portfolio are 11 percent (stock), 24 percent (bonds), 65 percent (cash); the actual portfolio weights are used for the target date funds (i.e., Funds A–D).

Table 3 Asset allocations that minimize retirement downside risk in the two-asset class case: Stocks and bonds only.

Spending rate	Stocks	Bonds	Cash	LPM_0	LPM_2	Mean RPV
Male						
A. 65-year-old						
\$6	21%	79%	0%	1.16%	\$0.94	\$31.53
\$7	23%	77%	0%	8.20%	\$3.21	\$20.54
\$8	26%	74%	0%	24.36%	\$7.73	\$9.73
B. 75-year-old						
\$11	25%	75%	0%	6.32%	\$2.17	\$18.33
\$11.5	25%	75%	0%	11.84%	\$3.27	\$14.60
\$12	25%	75%	0%	18.68%	\$4.68	\$10.99
C. 85-year-old						
\$21	22%	78%	0%	4.08%	\$1.42	\$16.76
\$22	22%	78%	0%	9.84%	\$2.39	\$12.76
\$23	22%	78%	0%	18.20%	\$3.76	\$8.80
Female						
A. 65-year-old						
\$5	22%	78%	0%	0.56%	\$0.73	\$35.38
\$6	24%	76%	0%	6.44%	\$3.03	\$22.91
\$7	27%	73%	0%	23.08%	\$8.02	\$10.68
B. 75-year-old						
\$9	26%	74%	0%	4.36%	\$1.77	\$21.76
\$9.5	25%	75%	0%	9.20%	\$2.92	\$17.31
\$10	26%	74%	0%	15.76%	\$4.44	\$13.00
C. 85-year-old						
\$17	23%	77%	0%	3.68%	\$1.37	\$18.61
\$18	22%	78%	0%	10.00%	\$2.53	\$13.75
\$19	22%	78%	0%	19.40%	\$4.22	\$9.01

Note: Spending rates represent the annual, inflation-adjusted withdrawal rates per \$100 in retirement savings. Probability of ruin is measured by the zero-order lower-partial moment (LPM_0) of negative RPV outcomes while downside risk is measured by the second-order (LPM_2) lower-partial moment.

Specifically, in Table 3 we replicate the analyses underlying Table 2 but without the possibility of investors making any cash allocation (i.e., restrict them to two-asset portfolios). In this setting, there are some interesting observations that can be made. First, notice that the optimal equity exposures rise to roughly 25 percent across all age, gender, and spending cases. Without cash to provide downside protection, the allocations to stock have increased since bonds don't provide as much protection against market price volatility. Second, notice also that without cash in the mix, the overall level of retirement downside risk will necessarily increase; the optimal outcome for a two-asset portfolio cannot be superior to that for a portfolio that includes the possibility of an additional asset class. For our base case of a 65-year-old male with a \$7 annual spending rate, the downside risk metric from Table 3 is \$3.21 compared to \$1.90 in the scenario when cash is included in the solution—a 69 percent increase. (The comparable increase for the 65-year-old female investor with a \$6 spending goal is from \$1.80 to \$3.03.) Finally, when cash is excluded, the average RPV also increases, which is an artifact of the higher expected returns of stocks and bonds relative to that for cash-equivalent instruments.

5 Sensitivity to capital market assumptions

Clearly, the ultimate success or failure of a retirement plan is closely tied to the returns and volatility of the assets in which individuals choose to invest their retirement savings. While the downside risk-minimizing allocations we saw in Table 2 used very reasonable capital market expectations based on long-term historical averages, it is nevertheless useful to test the findings of the RPV model with alternative sets of investment assumptions.

Returning to our base case of a 65-year-old male investor spending an inflation-adjusted \$7 per year in retirement, Table 4 provides a comparison of the optimal asset allocation weights, the probability of ruin percentage (LPM_0) and retirement

Table 4 Minimum risk allocations under different capital market assumptions for a 65-year-old male with a \$7 annual spending rate.

		Portfolio allocation:					
	Scenario description	Stocks	Bonds	Cash	LPM_0	LPM_2	Mean RPV
	Base case	11%	24%	65%	9.96%	\$1.90	\$10.21
1	Stock return: 6% → 7%	14%	24%	62%	7.16%	\$1.62	\$12.61
2	Stock return: 6% → 5%	7%	25%	68%	13.08%	\$2.13	\$8.45
3	Bond return: 3% → 3.5%	10%	32%	57%	6.84%	\$1.58	\$12.71
4	Bond return: 3% → 2.5%	11%	17%	72%	13.44%	\$2.17	\$8.49
5	Cash return: 1% → 0.1%	18%	56%	26%	10.35%	\$3.07	\$14.88
6	Cash return: 1% → 0%	19%	60%	21%	10.03%	\$3.13	\$15.77
7	Stock volatility: 16% → 18%	7%	25%	68%	11.76%	\$2.06	\$9.08
8	Stock volatility: 16% → 14%	15%	24%	61%	7.64%	\$1.65	\$12.00
9	Bond volatility: 7% → 8%	11%	16%	73%	12.44%	\$2.11	\$9.06
10	Bond volatility: 7% → 6%	10%	37%	53%	7.28%	\$1.58	\$12.04
11	Stock/bond correlation: 0.2 → 0.3%	10%	19%	70%	11.16%	\$1.98	\$9.26
12	Stock/bond correlation: 0.2 → 0.1%	12%	27%	61%	8.84%	\$1.75	\$11.48
13	Scenarios (2) + (7) + (10)	5%	26%	69%	14.20%	\$2.24	\$7.89
14	Scenarios (1) + (4) + (8) + (9) + (12)	22%	12%	66%	6.28%	\$1.43	\$14.04

Note: Spending rates represent the annual, inflation-adjusted withdrawal rates per \$100 in retirement savings. Probability of ruin is measured by the zero-order lower-partial moment (LPM_0) of negative RPV outcomes while downside risk is measured by the second-order (LPM_2) lower-partial moment.

downside risk level (LPM_2), and the average RPV profile as capital market assumptions are changed within the model. For example, in Scenario 1, the expected real return on stocks is increased from 6 to 7 percent. This adjustment results in an increase in the risk-minimizing stock allocation to 14 percent compared to the base case allocation of 11 percent. Furthermore, retirement risk decreases (to \$1.62 from \$1.90) and the mean RPV increases (to \$12.61 from \$10.21), reflecting the more attractive return expectations in the equity market. Other scenarios then look at the impact of changing the initial return expectations for each of the three asset classes, as well as considering adjustments to the volatilities of stocks and bonds and the correlation coefficient between them. Finally, Scenario 13 tests the combined effect of three scenarios—2, 7, and 10—that result in a decrease in the optimal equity allocation while Scenario 14 combines five scenarios—1, 4, 8, 9, and 12—in which the optimal equity allocation increases.

The key conclusion suggested by all of these assumption-changing scenarios is that the downside risk-minimizing portfolios remain relatively conservative in terms of their stock investment proportions. To this point, notice that setting the expected real return on cash-equivalent investments to essentially zero (i.e., Scenarios 5 and 6)—a condition prevalent in the capital markets in the post-2008 environment—never causes the optimal stock allocation to exceed 20 percent.[12] Indeed, even a combination of assumptions deliberately chosen to increase the aggressiveness of the equity allocation (i.e., Scenario 14) leads to an optimal equity level of just 22 percent, which is still significantly below what is seen in retirement products typically offered to investors.

At this point, some intuition is useful as to why, in general, the risk-minimizing portfolios have low equity allocations to begin with and remain low even in the various alternative scenarios we have examined. The answer is linked to the primary cause of retirement shortfall, namely sequence-of-returns risk, which suggests that a retiree who is already taking withdrawals from his or her savings account will be affected more adversely by an initial time series of negative returns followed by a series of positive returns than if that return pattern had been reversed (see, for instance, Pfau and Kitces (2014)). So, if an investor is unfortunate enough to be exposed to a sequence of negative returns early in retirement, the likelihood of an early depletion of savings rises dramatically. Such would have been the case, for example, for individuals retiring in 1973, 1999, or 2007.

Any large exposure to equities carries with it an added chance of increasing this sequence-of-returns risk. While the higher expected returns for stocks relative to bonds and cash equivalents are certainly advantageous for sustaining retirement savings, a key result of our investigation is that this benefit is outweighed by the potential for equity investments to induce downside return shocks that substantially increase the risk of financial ruin.

6 Optimal retirement asset allocation with bequest motives

As an alternative to the risk minimization objective in the previous sections, let us now consider the subset of retired individuals who still have a concern for retirement downside risk but who also want to leave assets to their heirs. As we discussed earlier, the average RPV of a retirement plan can be thought of as an estimate its net value in today's dollars. Someone who wants to bequeath money to his or her beneficiaries might be willing to take on some additional retirement risk in exchange for increasing the potential value of the assets remaining at the time of death. In this context, there is, in fact, a continuous set of tradeoffs between retirement downside risk and the average RPV level. So, just as there is a mean-variance optimal efficient frontier for investment allocations that minimize portfolio volatility for a given level of expected return, there is an analogous efficient "retirement frontier" that illustrates the optimal tradeoffs between retirement risk and the value of the potential bequest.

Figure 5 illustrates these efficient frontiers by depicting the tradeoffs facing an investor who has both risk-control and bequest goals. The retirement frontiers are shown for a 65-year-old male with annual spending rates of \$6, \$7, and \$8. The minimum downside risk portfolios are identified as the points falling at the bottom left of each iso-spending curve and are labeled as A, B, and C, respectively. (These portfolios and their characteristics were shown in Table 2 and have equity allocations of 5, 11 and 20 percent, respectively.) Once again, these portfolios are relevant for those individuals who are most concerned about the risk of outliving their retirement assets without having explicit bequest motives.

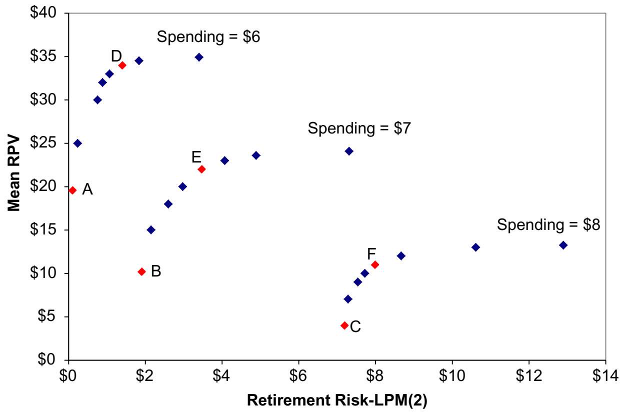

Figure 5 Retirement downside risk-RPV efficient Frontiers for a 65-year-old male.

Note: The analysis assumes that a 65-year-old individual has \$100 in retirement savings and plans to spend the indicated amount per year, adjusted for inflation. The analysis ignores taxes and transaction costs. Mortality is modeled using the United States Social Security Administration's period life tables.

For each of the spending policy frontiers, as we move upward and to the right along the curves, retirement risk increases. However, with this added risk there is also an increase in the average RPV of the plan. At first, the curves are very steep, meaning that relatively small increases in retirement downside risk are accompanied by relatively large increases in the average RPV. In other words, in this region of the curve, the "cost" to increase the possibility of higher RPV is relatively low in terms of the incremental downside risk the investor must assume. At approximately their midway points, though, the iso-spending frontiers become almost flat. Once this transition point is reached, any additional increase in the desired RPV levels comes at the expense of very large increases in retirement risk. The marginal cost of increasing potential bequests, therefore, becomes very high.

In Figure 5, we have indicated three portfolios—labeled D, E, and F—that would seem to reflect the upper limit of the RPV-Retirement Risk tradeoff that would be attractive to most investors. While these are highly subjective selections, they do allow us to investigate the change in the asset allocation policy that occurs as the retiree's objective moves beyond just being concerned about his or her own risk of ruin exposure to one that also includes the desire to leave money to others. (The specific values of the asset allocations and risk characteristics for all of the portfolios represented in Figure 5 are listed in Table 5.)

As a starting point for this comparison, Figure 6 displays the RPV distribution for Portfolio E; recall that Portfolio B from the same frontier is depicted in Figure 2. Portfolio E has a stock, bond, and cash mix of 34, 66 and 0 percent, respectively. While still a somewhat conservatively balanced portfolio, its more aggressive positioning relative to the downside risk-minimizing Portfolio B more than doubles the mean RPV, from \$10.21 to \$22.02. Along with this increase in plan value, however, is a marked rise in retirement downside risk from \$1.90 to \$3.47, an increase of 83 percent. Interestingly, while downside risk increases, the probability of failure actually decreases slightly from 9.96 to 7.96 percent. Therefore, the likelihood of failure occurring decreases by two percentage points, but the LPM_2 statistic indicates that, when failure does occur, it

Table 5 Retirement downside risk-RPV efficient frontiers: Optimal asset allocations.

				Male - age 65			
Highlighted portfolio (Figure 5)	Spending Rate	Stocks	Bonds	Cash	LPM_0	LPM_2	Mean RPV
A	\$6	5%	14%	86%	0.12%	\$0.10	\$19.59
	\$6	14%	30%	56%	0.28%	\$0.24	\$25.00
	\$6	16%	74%	10%	1.04%	\$0.76	\$30.00
	\$6	29%	64%	7%	0.92%	\$0.89	\$32.00
	\$6	34%	64%	2%	1.04%	\$1.07	\$33.00
D	\$6	44%	56%	0%	1.80%	\$1.40	\$34.00
	\$6	52%	48%	0%	2.40%	\$1.84	\$34.51
	\$6	69%	31%	0%	4.40%	\$3.40	\$34.92
B	\$7	11%	24%	65%	9.96%	\$1.90	\$10.21
	\$7	18%	41%	41%	8.20%	\$2.15	\$15.01
	\$7	21%	59%	20%	7.96%	\$2.60	\$18.00
	\$7	25%	66%	9%	7.76%	\$2.97	\$20.01
E	\$7	34%	66%	0%	7.96%	\$3.47	\$22.02
	\$7	44%	56%	0%	8.60%	\$4.07	\$23.00
	\$7	52%	48%	0%	9.28%	\$4.89	\$23.60
	\$7	69%	31%	0%	12.20%	\$7.31	\$24.09
C	\$8	20%	46%	34%	34.64%	\$7.19	\$4.01
	\$8	24%	55%	21%	28.24%	\$7.28	\$7.04
	\$8	28%	63%	9%	25.56%	\$7.54	\$9.01
	\$8	31%	66%	3%	24.40%	\$7.72	\$10.01
F	\$8	35%	65%	0%	22.92%	\$7.98	\$11.01
	\$8	44%	56%	0%	22.80%	\$8.67	\$12.02
	\$8	58%	42%	0%	23.84%	\$10.61	\$13.00
	\$8	70%	30%	0%	25.20%	\$12.89	\$13.25

Note: Spending rates represent the annual, inflation-adjusted withdrawal rates per \$100 in retirement savings. Probability of ruin is measured by the zero-order lower-partial moment (LPM_0) of negative RPV outcomes while downside risk is measured by the second-order (LPM_2) lower-partial moment.

Table 5 (*Continued*)

				Female - age 65			
Highlighted portfolio (Figure 5)	Spending Rate	Stocks	Bonds	Cash	LPM_0	LPM_2	Mean RPV
	$5	3%	13%	84%	0.06%	$0.03	$21.68
	$5	18%	34%	48%	0.12%	$0.21	$30.01
	$5	26%	62%	12%	0.44%	$0.59	$35.01
	$5	30%	65%	5%	0.52%	$0.72	$36.08
	$5	35%	65%	0%	0.64%	$0.88	$37.08
	$5	40%	60%	0%	0.72%	$0.99	$37.53
	$5	46%	54%	0%	1.24%	$1.18	$38.00
	$5	69%	31%	0%	3.20%	$2.85	$38.82
	$6	11%	25%	64%	8.29%	$1.80	$11.51
	$6	15%	34%	51%	7.00%	$1.87	$15.04
	$6	19%	66%	15%	6.68%	$2.57	$20.21
	$6	29%	64%	7%	6.32%	$2.93	$23.02
	$6	32%	67%	1%	6.36%	$3.16	$24.00
	$6	39%	61%	0%	6.72%	$3.54	$25.00
	$6	52%	48%	0%	7.92%	$4.57	$26.02
	$6	69%	31%	0%	10.60%	$7.07	$26.56
	$7	21%	47%	30%	32.82%	$7.62	$4.96
	$7	27%	53%	20%	26.96%	$7.66	$8.01
	$7	28%	65%	7%	24.36%	$7.88	$10.01
	$7	31%	67%	2%	23.04%	$8.04	$11.01
	$7	36%	64%	0%	22.24%	$8.29	$12.01
	$7	44%	56%	0%	22.00%	$8.96	$13.03
	$7	58%	42%	0%	22.80%	$10.83	$14.05
	$7	69%	31%	0%	24.24%	$13.15	$14.32

will be significantly more severe (i.e., the RPV values more highly negative).

Table 6 provides the asset class weights and RPV statistics for portfolios represented by points D, E, and F as well as for the comparable portfolios chosen using different age, gender, and spending rate assumptions. It should be noted that the equity allocations for all of these portfolios are approximately twice as large as those in the examples we have seen earlier when minimizing retirement downside risk was the investor's only goal. Given the bequest goal for these portfolios, their equity shares are roughly in the 35 to 45 percent range.

While these findings are explicitly computed using three asset classes—stocks, bonds, and cash—they also hold for the two-asset class analysis since the cash allocation in Table 6 is zero in all cases. The intriguing aspect of all of these results is that, even after we extend the risk positioning of the retirement portfolios to increase the potential for bequests, the proscribed

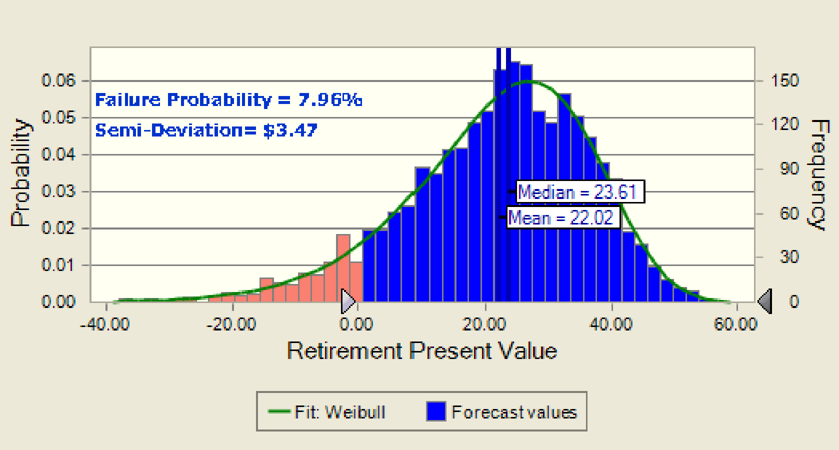

Figure 6 RPV distribution for an asset allocation that considers both risk and bequest for a 65-year-old male with $7 annual spending rate.

Note: The analysis assumes that a 65-year-old individual has $100 in retirement savings and plans to spend $7 per year, adjusted for inflation. The analysis ignores taxes and transaction costs. The analysis ignores taxes and transaction costs. Mortality is modeled using the United States Social Security Administration's period life tables. The distribution displayed represents a histogram of 10,000 sequences of annual potential portfolio returns, using the following capital market assumptions: (i) Expected real return: 6.0 percent (stock), 3.0 percent (bonds), 1.0 percent (cash); (ii) Volatility: 16.0 percent (stock), 7.0 percent (bonds), 2.5 percent (cash); and (iii) Correlation: 0.20 (stock/bonds), 0.15 (stock/cash), 0.35 (bonds/cash). The discount factors used in the RPV computation for this particular distribution assume the following asset class allocation weights: 34 percent (stock), 66 percent (bonds), 0 percent (cash).

Table 6 Asset allocations that consider both retirement downside risk and bequest motives.

	Male					
Spending rate	Stocks	Bonds	Cash	LPM_0	LPM_2	Mean RPV
A. 65-year-old						
$6	44%	56%	0%	1.80%	$1.40	$34.00
$7	34%	66%	0%	7.96%	$3.47	$22.02
$8	35%	65%	0%	22.92%	$7.98	$11.01
B. 75-year-old						
$11	42%	58%	0%	8.20%	$2.67	$20.01
$11.5	38%	62%	0%	11.64%	$3.62	$16.00
$12	34%	66%	0%	17.92%	$4.88	$12.00

Note: Spending rates represent the annual, inflation-adjusted withdrawal rates per $100 in retirement savings. Probability of ruin is measured by the zero-order lower-partial moment (LPM_0) of negative RPV outcomes while downside risk is measured by the second-order (LPM_2) lower-partial moment.

Table 6 (*Continued*)

Spending rate	Stocks	Bonds	Cash	LPM_0	LPM_2	Mean RPV
			Male			
C. 85-year-old						
$21	43%	57%	0%	6.20%	$1.89	$18.00
$22	41%	59%	0%	11.52%	$2.91	$14.01
$23	39%	61%	0%	18.96%	$4.27	$10.00
			Female			
A. 65-year-old						
$5	40%	60%	0%	0.72%	$0.99	$37.53
$6	39%	61%	0%	6.72%	$3.54	$25.00
$7	36%	64%	0%	22.24%	$8.29	$12.01
B. 75-year-old						
$9	38%	62%	0%	4.88%	$2.01	$23.01
$9.5	41%	59%	0%	9.84%	$3.40	$19.00
$10	34%	66%	0%	14.96%	$4.61	$14.00
C. 85-year-old						
$17	45%	55%	0%	5.48%	$1.93	$20.00
$18	39%	61%	0%	10.88%	$2.97	$15.01
$19	43%	57%	0%	19.88%	$5.01	$10.51

allocation schemes are still more conservative than the typical allocations seen in practice.

7 Conclusions

There are many ways to think about the risks of an individual investor's retirement plan and how the asset allocation decision can influence those risks. The retirement present value (RPV) method provides a useful starting point by modeling the retirement plan as the net present value of assets minus liabilities weighted by the probability of the investor's survival throughout his or her post-retirement life. Because there is a distribution of RPVs based upon the realization of future investment returns and mortality events, risk can be thought of as the potential for negative outcomes in net value of the plan.

While much of the financial economic literature on retirement investing has focused on the use of probability of ruin as a risk measure, we find that a more general metric that captures both the possibility *and* the severity of funding shortfalls offers more attractive outcomes. In particular, we show that a downside risk measure based on the semi-deviation of retirement present values produces portfolios that survive longer in worst-case scenarios. Therefore, this downside risk approach appears more useful in constructing sustainable retirement strategies.

We also find that when minimizing the risk of retirement plan shortfalls, the optimal asset allocation mix conditional on sustainable spending rates is surprisingly conservative, with equity commitments for 65 to 85-year-olds falling in the

5 to 10 percent range. With cash excluded from the asset mix, optimal equity allocations for the minimum risk portfolios are still only around 25 percent of the overall portfolio. In addition, these stock allocations remain little changed even when we make substantial adjustments to the underlying investment risk and return assumptions. The conservative nature of these results differs significantly from most of the investment products offered to investors in today's marketplace, which typically have an average equity allocation of around 50 percent.

Of course, not all retired investors are focused solely on minimizing the downside risk of their retirement plan. For some, taking on additional risk with a more aggressive asset allocation would be acceptable in exchange for the possibility of leaving their heirs with a larger estate upon their demise. However, even when we consider these tradeoffs, it is still the case that the optimal equity allocations are relatively conservative, falling in the 35–45 percent range.

Taken as a whole, the findings in this study should give any investor a considerable amount to ponder before setting his or her asset allocation path in retirement. If mitigating the risk of outliving one's retirement resources is the cornerstone of the asset allocation decision, it is critical to limit equity exposure and recognize the impact that investment volatility and mortality risk can have on the sustainability of the retirement plan.

Appendix. Calculating retirement present value: An example

For this example of an RPV calculation, as expressed by Equations (1)–(3), we assume the individual in question is a man who is 65 years old and has just retired. The investor currently has $1,000,000 in accumulated savings and intends to spend $70,000 per year in retirement, adjusted for inflation. (Notice that this is just a scaled version of the example used throughout the study, which assumes a $100 initial portfolio value and a $7 annual real spending goal.) Table A.1 depicts the elements of the RPV calculation for one specific sequence of potential portfolio return values using our base case capital market assumptions: (i) *Expected real return*: 6.0 percent (stock), 3.0 percent (bonds), 1.0 percent (cash); (ii) *Volatility*: 16.0 percent (stock), 7.0 percent (bonds), 2.5 percent (cash); and (iii) *Correlation*: 0.20 (stock/bonds), 0.15 (stock/cash), 0.35 (bonds/cash).

The retirement cash flows for the investor are projected for 45 years, or out to age 110; the probability of living beyond this age is effectively zero. The cash flows in each year (Column B) are multiplied by the probability of survival to that given age (Column A) and then divided by the applicable discount rate. In this example—as well as throughout the entire study—the mortality risk statistics are provided by the United States Social Security Administration's period life tables. Each period-specific discount rate is calculated by Equation (2) and indicates the compounded return on the portfolio up to a particular point in time *t*. The yearly values shown in the "Portfolio Return" column of Table A.1 represent the weighted combination of one particular sequence of single annual draws from each of the return distributions for the three asset classes using the formula in Equation (3).[13] For convenience, this sequence of discount rates for the projected 45-year retirement period is shown in their reciprocal form, which are the Present Value Factors listed in Column C.

Using these mortality risk and capital market forecasts, the probability-weighted discounted cash flow for a given year in the future can be calculated as the simple product of Column A, Column B, and Column C. For this retiree's 45-year potential remaining life span, the sequence of these

Table A.1 Example of an RPV calculation for a mortality risk forecast and one sequence of capital market return values.

Year (t)	Age	Probability of survival (A)	Cash flow (B)	Portfolio return	Present value factor: $1/(1+R)^t$ (C)	Probability-weighted discounted cash flow =A*B*C
0	65	1.0000	$1,000,000		1.0000	$1,000,000
1	66	0.9919	($70,000)	9.31%	0.9149	($63,520)
2	67	0.9830	($70,000)	−2.18%	0.9352	($64,355)
3	68	0.9735	($70,000)	14.76%	0.8149	($55,534)
4	69	0.9633	($70,000)	3.24%	0.7893	($53,228)
5	70	0.9524	($70,000)	9.68%	0.7197	($47,980)
6	71	0.9406	($70,000)	−3.43%	0.7453	($49,068)
7	72	0.9278	($70,000)	7.91%	0.6907	($44,856)
8	73	0.9140	($70,000)	8.15%	0.6386	($40,859)
9	74	0.8990	($70,000)	−4.00%	0.6652	($41,862)
10	75	0.8826	($70,000)	11.21%	0.5982	($36,959)
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
40	105	0.0148	($70,000)	3.97%	0.2091	($217)
41	106	0.0097	($70,000)	2.91%	0.2032	($138)
42	107	0.0062	($70,000)	2.80%	0.1977	($85)
43	108	0.0038	($70,000)	1.20%	0.1953	($52)
44	109	0.0022	($70,000)	8.49%	0.1800	($28)
45	110	0.0013	($70,000)	5.58%	0.1705	($15)
						$39,964

Note: This example of a single RPV calculation assumes an investor who is 65 years old, has just retired, and will not live past the age of 110. Mortality risk statistics are given by the U.S. Social Security Administration's period life tables. The sequence of annual investment returns is hypothetical and represents one path of future portfolio performance, using the following capital market assumptions: (i) Expected real return: 6.0 percent (stock), 3.0 percent (bonds), 1.0 (cash); (ii) Volatility: 16.0 percent (stock), 7.0 percent (bonds), 2.5 (cash); and (iii) Correlation: 0.20 (stock/bonds), 0.15 (stock/cash), 0.35 (bonds/cash). The discount factors used in the RPV computation for this particular distribution assume the following asset class allocation weights: 11 percent (stock), 24 percent (bonds), 65 percent (cash).

values is shown in the final column of the display. The RPV statistic is then computed as the sum of these probability-weighted discounted cash flows. In this example, the RPV outcome is equal to $39,964 in real dollars (which, when scaled to the $100 initial portfolio value used throughout the study, is equivalent to $4.00). Of course, this value represents only one of many possible outcomes due to the uncertainty around future investment returns. By sampling repeatedly from a collection of asset class returns based on historical statistics—generating, say, 10,000 potential sequences, as in a typical Monte Carlo-based simulation—or by forecasting returns in some other manner, an entire distribution of potential RPV values can be generated for a given investor's retirement plan. It is this sort of RPV distribution—as well as the various statistical measures of it—that are detailed and reported in this study (e.g., Figure 2).

Notes

[1] For a more thorough discussion on the role that investor education plays in the life-cycle investment problem, see Bodie *et al.* (2012).

[2] The issue of outliving one's assets in retirement is, of course, a primary concern of investors and one that is critical to the way we define what is optimal in our approach to asset allocation. Other studies that have considered the longevity risk problem in this context include Ho *et al.* (1994) and Horneff *et al.* (2006).

[3] The notion of mortality risk in retirement planning has also been considered in Milevsky *et al.* (1997) and Chen *et al.* (2006).

[4] In the analyses presented in this paper, retirement plan cash flows and simulated returns are estimated from the individual's current age out to age 110. Mortality effects are based on the United States Social Security Administration's period life tables.

[5] Our other base case assumptions are that real stock returns have a correlation with those of bonds and cash of 0.20 and 0.15, respectively, and that the correlation of real bond returns with cash returns is 0.35. These assumptions, as well as the expected returns and volatilities, are consistent with historical trends in the United States capital markets since 1946.

[6] The use of multiple measures to quantify the risk of possible outcomes is consistent with Das *et al.* (2010) whereby the LPM_2 measure captures information on the downside dispersion of RPV outcomes and the LPM_0 measure embodies information regarding risk preferences.

[7] An alternative to controlling retirement downside risk uses annuities to convert a portion of the investment portfolio into a guaranteed income stream for the duration of the investor's life, subject only to the financial solvency of the guarantor. Interestingly, Albrecht and Maurer (2002) find that a "self-annuitization" approach in which the retiree maintains control of the asset allocation strategy is often optimal; see also Kaplan (2006).

[8] The goal of the optimization problem in Equation (7) is to select the asset allocation strategy that minimizes the retiree's downside risk exposure. From an operational perspective, this is indeed consistent with choosing $\{w_k\}$ so as to *minimize* LPM_0 or LPM_2 when either of those two equations defines the investor's objective function. With LPM_1, however, minimizing downside risk actually entails *maximizing* the expected shortfall (i.e., making a negative value as small as possible). Subject to this clarification, throughout the paper we will continue to refer to the optimization problem as one in which the retiree minimizes LPM_X.

[9] The lower panel in Table 1 lists a comparable set of optimization results for a 65-year-old female investor. Generally speaking, these findings lead to comparable conclusions as for the 65-year-old male retiree just discussed. Specifically, for any given desired spending level in retirement, the optimal equity allocation is substantially lower when minimizing downside risk is the objective compared to minimizing the probability of ruin (e.g., 14 percent (expected shortfall) and 11 percent (semi-deviation) versus 20 percent (probability of ruin) for a \$6 spending goal). Notice also, though, that for any comparable spending goal—say \$7—the female retiree requires a *riskier* allocation than her male counterpart (i.e., equity exposure of 21 percent versus 11 percent using the LPM_2 minimization). As we discuss in Section 4, this is because of her longer expected life span from the same starting retirement age.

[10] Throughout the remainder of the study, we streamline the comparison of failure probability outcomes versus downside risk outcomes by presenting findings for just the two most contrasting statistics: LPM_0 (probability of ruin) and LPM_2 (semi-deviation). For each of the subsequent comparisons, we have also produced a complete set of results for the LPM_1 (expected shortfall) statistic as well, but these are suppressed in the presentation herein in the interest of space and comprehension.

[11] Using our downside risk framework, a retiree, age 65 (male or female) who will live to age 95 with certainty, has a sustainable spending rate (i.e., the rate associated with a 10 percent probability of ruin) of \$3.90 per \$100 of savings and a risk-minimizing allocation to stocks, bonds, and cash of 12, 31, and 57 percent, respectively.

[12] It is interesting to note the difference in the optimal stock allocation when the cash return is set at zero (19 percent in Table 4, Scenario 6) versus when cash is removed from the retirement portfolio altogether (23 percent for the comparable age and spending policy configuration in Table 3). Thus, relative to a two-asset portfolio that includes just stocks and bonds, it is still beneficial to include zero-return cash in a downside risk-minimizing portfolio for its ability to diversify the portfolio by virtue of being less-than-perfectly correlated with the other more volatile asset classes.

[13] For the purpose of this illustration, we use the following assumed asset class allocation weights to generate these yearly portfolio returns: 11 percent (stock), 24 percent (bonds), and 65 percent (cash). Note, however, that the primary purpose of the study is to strategically select these allocation weights so as to optimize the investor's objective function (e.g., minimize downside retirement risk), subject to some regularity constraints. We explain the mechanics of this optimization process in Section 2.3.

Acknowledgments

We are grateful for the comments of Jeffrey Brown, Dan Cassidy, Jeremy Crean, Javier Estrada, and Leonard Glynn, as well as participants in the University of Texas seminar series and the 2012 FMA European Conference in Istanbul, Turkey. The opinions and analyses presented herein are those of the authors and do not necessarily represent the views of Putnam Investments or Empower Retirement.

References

Albrecht, P. and Maurer, R. (2002). "Self-Annuitization, Consumption Shortfall in Retirement and Asset Allocation: The Annuity Benchmark," *Journal of Pension Economics and Finance* **1**, 269–288.

Alestalo, N. and Puttonen, V. (2006). "Asset Allocation in Finnish Pension Funds," *Journal of Pension Economics and Finance* **5**, 27–44.

Balduzzi, P. and Reuter, J. (2012). "Heterogeneity in Target-Date Funds and the Pension Protection Act of 2006," NBER Working Paper.

Basu, A. K. and Drew, M. E. (2009). "Portfolio Size Effect in Retirement Accounts: What Does it Imply for Lifecycle Asset Allocation Funds?" *Journal of Portfolio Management* **35**, 61–72.

Bawa, V. S. (1975). "Optimal Rules for Ordering Uncertain Prospects," *Journal of Financial Economics* **2**, 95–121.

Bawa, V. S. (1978). "Safety First, Stochastic Dominance, and Optimal Portfolio Choice," *Journal of Finance and Quantitative Analysis* **13**, 255–271.

Bawa, V. S. and Lindenburg, E. (1977). "Capital Market Equilibrium in a Mean, Lower Partial Moment Framework," *Journal of Financial Economics* **5**, 189–200.

Bodie, Z., Merton, R. C. and Samuelson, W. F. (1992). "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model," *Journal of Economics Dynamics and Control* **16**, 427–449.

Bodie, Z. and Crane, D. B. (1997). "Personal Investing: Advice, Theory and Evidence," *Financial Analysts Journal* **53**, 13–23.

Bodie, Z., Siegel, L. B., and Stanton, L. (2012). *Life-Cycle Investing: Financial Education and Consumer Protection.* Charlottesville, VA: Research Foundation of the CFA Institute.

Booth, L. (2004). "Formulating Retirement Targets and the Impact of Time Horizon on Asset Allocation," *Financial Services Review* **13**, 1–17.

Chen, P., Ibbotson, R. G., Milevsky, M. A., and Zhu, K. X. (2006). "Human Capital, Asset Allocation and Life Insurance," *Financial Analysts Journal* **62**, 97–109.

Cocco, J. F., Gomes, F. J., and Maenhout, P. J. (2005). "Consumption and Portfolio Choice Over the Life Cycle," *Review of Financial Studies* **18**, 491–533.

Das, S., Markowitz, H., Scheid, J., and Statman, M. (2010). "Portfolio Optimization with Mental Accounts," *Journal of Financial and Quantitative Analysis* **45**, 311–334.

Elton, E. J., Gruber, M. J., de Souza, A., and Blake, C. R. (2015). "Target Date Funds: Characteristics and Performance," New York University Working Paper.

Gomes, F. J., Kotlikoff, L. J., and Viceira, L. A. (2008). "Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds," *American Economic Review* **98**, 297–303.

Harlow, W. V. (1991). "Asset Allocation in a Downside-Risk Framework," *Financial Analysts Journal* **47**, 28–40.

Harlow, W. V. and Rao, R. K. S. (1989). "Asset Pricing in a Generalized Mean-Lower Partial Moment Framework," *Journal of Financial and Quantitative Analysis* **24**, 285–311.

Hariharan, G., Chapman, K. S., and Domian, D. L. (2000). “Risk Tolerance and Asset Allocation for Investors Nearing Retirement,” *Financial Services Review* **9**, 159–170.

Ho, K., Milevsky, M. A., and Robinson, C. (1994). “Asset Allocation, Life Expectancy and Shortfall,” *Financial Services Review* **3**, 109–126.

Horneff, W. J., Maurer, R., Mitchell, O. S., and Dus, I. (2006). “Optimizing the Retirement Portfolio: Asset Allocation, Annuitization and Risk Aversion,” NBER Working Paper.

Jagannathan, R. and Kocherlakota, N. R. (1996). “Why Should Older People Invest Less in Stocks than Younger People?” *Federal Reserve Bank of Minneapolis Quarterly Review* **20**, 11–23.

Kaplan, P. D. (2006). “Asset Allocation with Annuities for Retirement Income Management,” *Journal of Wealth Management* **9**, 27–40.

Kingston, G. H. (2000). “Efficient Timing of Retirement,” *Review of Economic Dynamics* **3**, 831–840.

Merton, R. C. (1969). “Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case,” *Review of Economic and Statistics* **51**, 247–257.

Merton, R. C. (2012). “Observations on Financial Education and Consumer Financial Protection,” in *Life-Cycle Investing: Financial Education and Consumer Protection*, edited by Z. Bodie, L. Siegel, and L. Stanton (Charlottesville, VA: Research Foundation of the CFA Institute).

Milevsky, M. A. and Robinson, C. (2005). “A Sustainable Spending Rate without Simulation,” *Financial Analysts Journal* **61**, 89–100.

Milevsky, M. A., Ho, K., and Robinson, C. (1997). “Asset Allocation Via the Conditional First Exit Time or How to Avoid Outliving Your Money,” *Review of Quantitative Finance and Accounting* **9**, 53–70.

Pfau, W. D. and Kitces, M. E. (2014). “Reducing Retirement Risk with a Rising Equity Glide-Path,” *Journal of Financial Planning* **27**, 38–45.

Poterba, J. M. and Samwick, A. (2001). “Household Portfolio Allocation Over the Life Cycle,” in *Aging Issues in the United States and Japan*, edited by S. Ogura, T. Tachibanaki and D. Wise (Chicago, IL: University of Chicago Press).

Riley, W. B. and Chow, K. V. (1992). “Asset Allocation and Individual Risk Aversion,” *Financial Analysts Journal* **48**, 32–37.

Rook, C. J. (2014). “Minimizing the Probability of Ruin in Retirement,” Stevens Institute of Technology Working Paper.

Samuelson, P. A. (1969). “Lifetime Portfolio Selection by Dynamic Stochastic Programming,” *Review of Economic and Statistics* **51**, 239–246.

Schleef, H. J. and Eisinger, R. M. (2007). “Hitting or Missing the Retirement Target: Comparing Contribution and Asset Allocation Schemes of Simulated Portfolios,” *Financial Services Review* **16**, 229–243.

Schooley, D. K. and Worden, D. D. (1999). “Investors’ Asset Allocations Versus Life-Cycle Funds,” *Financial Analysts Journal* **55**, 37–43.

Scott, J. S., Sharpe, W. F., and Watson, J. G. (2009). “The 4% Rule—At What Price?” *Journal of Investment Management* **7**, 31–48.

Stout, R. G. (2008). “Stochastic Optimization of Retirement Portfolio Asset Allocations and Withdrawals,” *Financial Services Review* **17**, 1–15.

Wachter, J. A. and Yogo, M. (2010). “Why Do Household Portfolio Shares Rise in Wealth?” *Review of Financial Studies* **23**, 3929–3965.

Waggle, D. and Basil Englis, B. (2000). “Asset Allocation Decisions in Retirement Accounts: An All-or-Nothing Proposition?” *Financial Services Review* **9**, 79–92.

Keywords: Retirement asset allocation; market risk; mortality risk; downside risk; stochastic optimization

JEL Classification: G11, G17