THE JOURNAL OF
RETIREMENT

VOLUME 4, NUMBER 3 WINTER 2017

There is general agreement that managing the drawdown or decumulation phase of the retirement finance cycle poses a challenge for most households. The typical 401(k) plan does not offer its members annuities or other instruments that might help prevent too rapid or too slow a drawdown of the retirement nest egg. In "Retirement Income Programs: *The Next Step in the Transition from DB to DC Retirement Plans*," Wade Pfau, Joseph Tomlinson, and Steve Vernon apply modern portfolio theory to show how an informed choice may be made regarding the allocation of the nest egg among these instruments, which they term "retirement income generators," or RIGs. In addition to single-purchase immediate annuities (SPIAs), RIGs can include systematic withdrawal plans (SWPs), deferred annuities (including qualified longevity annuity contracts, or QLACs) and GLWBs. A sound retirement income strategy can also include delaying the claiming of Social Security retirement benefits.

Modern portfolio theory introduced the concept of the efficient frontier, with its trade-off between risk and return to determine the appropriate allocation of an investor's portfolio between risky and less risky assets. The authors apply a version of the theory to decisions at the decumulation phase to shed light on the role of different combinations of RIGs, which they term retirement solutions, in providing retirement income. In their version, the principal trade-off is between the *expected* average income in retirement, and a measure of accessible savings. It is important to recognize that one particular combination of RIGs could have a higher expected income than another with the same accessible savings, but also have greater variability of income.

The authors test different combinations of RIGs using Monte Carlo simulation and assuming rates of return and standard deviations that reflect the current low-interest-rate environment. The simulations are based on the portfolios of several hypothetical retirees. The authors find that a solution taking the form of a 100% investment in a SPIA yields the highest income of any solution but also the lowest (zero) amount of accessible savings. A 100% investment in a SWP with a withdrawal rate of 7% and assets invested 100% in equities generates less income than the SPIA, but substantial available savings on average. Mixed solutions (e.g., 30% in a SPIA, 70% in a SWP) produce income and available savings combinations that lie somewhere in between.

One notable finding of this work is that, provided a retiree has income from an annuity or Social Security, a high allocation to equities of her SWP may be optimal. The authors point out that having pension and Social Security income is like an investment in bonds, which makes

holdings of equities less risky. One particular strategy the authors analyze is investing part of a retiree's assets in a deferred annuity or QLAC and then relying entirely on the income from that investment after age 85, and financing income in the intervening years with a SWP. This strategy requires dealing with the possibility that income falls when the QLAC kicks in.

The article includes a useful discussion of fiduciary considerations for plan sponsors. Regulatory guidance on retirement income solutions has not yet been issued, and the authors recommend that in the meantime plan sponsors offer a number of different RIGs and provide online support that enables their participants to make informed decisions. This final section also offers useful advice for financial institutions and advisors, stressing the importance of understanding each client's circumstances and her goals for retirement. This article has valuable insights for plan sponsors and financial intermediaries marketing instruments for the decumulation phase of retirement.

In "The Outsourced Chief Investment Officer Model of Management and the Principal-Agent Problem," Gordon Clark and Roger Urwin describe and appraise a relatively new form of pension plan management known as the Outsourced Chief Investment Officer (OCIO) model. They contrast it with the hub-and-spokes model, which was the conventional way of managing a plan's investments in the two decades before the financial crash. The hub-and-spokes model maintains governance and the formulation of investment strategy with the plan sponsor, but outsources the implementation of investment decisions to a group of asset managers. The OCIO model outsources both the design of investment strategy and its implementation to an outsourced CIO. OCIO services may be provided by specialized OCIO firms or by large financial institutions.

The authors argue that the OCIO has certain advantages over the hub-and-spokes model, especially for small and midsize plans: In particular, it adapts more flexibly to changes in financial markets and is able to internalize many of the transaction costs that the hub-and-spokes model incurs. More generally, the hub-and-spokes model of management suffers from weak control and coordination. Nonetheless, to realize its potential the OCIO requires effective governance.

The authors set out six principles of OCIO best practice. One principal that seems to be especially important is that, although the responsibility for the formulation of investment strategy is formally delegated to the OCIO, the pension sponsor must be able to provide effective oversight to challenge the OCIO's investment strategy when that is appropriate. This requires that management of the plan sponsor have adequate understanding of the OCIO's work, or be able to draw on the relevant expertise.

Any model of pension fund management has to deal with the principal agent problem: the problem that the agent (in this case the OCIO) can act in its own interest and not in the interest of the plan sponsor. For the OCIO model to be effective, the company providing the services has to bear the costs of its decisions. The principal–agent problem will be least troublesome when the plan sponsor has access to relevant expertise and information. Internalizing investment strategy might mitigate the principal agent problem, but might also make it harder to replace the CIO.

When it comes to retirement planning, long-term care (LTC) is the elephant in the room. It definitely needs to be taken account of in any analysis of the sustainability of a retiree's finances, but often is not. Part of the reason for this may be that only a minority of retired households incur really large LTC expenditure. In "Will Long-Term Care Ruin Retirement Plans?" Michael Crook and Ronald Sutedja make up for this common failure of analysis with a model-based analysis of the impact of LTC on the probability of a retiree's running out of money. They work with a dataset of the probabilities that older people will require different kinds of care (home health care, assisted living, nursing home), each with a different price tag.

The authors rely on the automatically calculated virtual annuity of Waring and Siegel to determine regular expenditure. This approach ensures that, given that LTC expenditure does not rear its head, retirees never run out of money regardless of the performance of their portfolio. In the authors' model, LTC expenditure

and longevity are both random variables. They run simulations assuming initial wealth of $1 million to $10 million, and find that the probability of ruin for a couple with $1 million is almost 1 in 3; it drops to 1 in 20 for a couple with $10 million. Extra wealth shields households from running out of money because at any given time what might prove to be a catastrophic expenditure on LTC is less likely to exceed household wealth. (For example, an elderly couple whose starting wealth was $10 million might have $3 million left after 25 years, while a couple with starting wealth of $1 million might have only $300,000. The second couple would be more vulnerable to large unexpected LTC expenditures late in life.

The authors note that their analysis has not incorporated Medicare benefits, which is provided only for a limited time, nor long-term care provided by Medicaid, and they recognize some additional qualifications. Their main conclusion is one that retirement planners and their clients should heed: Excluding LTC from a retirement plan will mean that estimates of the plan's sustainability can be substantially overestimated.

Retirees need to be concerned about both sustainable income and available saving. Traditional defined benefit (DB) pension plans are excellent at providing the first but not the second; defined contribution (DC) plans undoubtedly facilitate the accumulation of savings, but are not designed to provide a stream of guaranteed income. In "Balancing Income and Bequest Goals in a DB/DC Hybrid Pension Plan," Yelei Gu, David Kausch, Kristen Moore, and Virginia Young explore the possibilities of achieving both goals at least in part by combining DB and DC plans. The authors establish parameters for models of a DB and a DC plan so that the expected replacement rate of the DC plan and the DB plan are equal. The model assumes random asset prices and mortality post retirement. Because wage increases are assumed to be predetermined, DB plan members confront no uncertainty about the value of their future pension, although they do not know how long they will receive it. Returns to DB investments are higher than those of DC investments, so the contribution rate to the DC plan must exceed that of the DB plan to ensure the two plans have the same replacement rate on average. The authors provide a succinct justification for the higher rate of return they assume for the DB plan.

The authors then address the question of how splitting a member's participation between the two plans will affect the probability of achievement of a specified bequest target and of asset exhaustion, and the variability of post-retirement income. With contribution rates of 11.34% to the DB plan and 14.66% to the DC plan (assuming 100% participation in one or the other plan), investment of 50/50 in the two plans would entail a contribution of 0.5*11.34% + 0.5*14.66%, or 13%.[1] With a 25% investment in the DB plan, the contribution rate is 0.25*11.34% + 0.75*14.66%, or 13.83%.

The percentage allocation to the DB plan has little effect on the probability that a retiree depletes her DC holdings, which is about 50% regardless of the allocation mainly because the median rate of return on DC plan assets was chosen to be high enough to fund an annuity equal to that of the DB plan. The authors also model the probabilities of achieving a specific bequest goal of $100,000; they find that the probability, which is of course maximized when the share of the DC plan is 100%, does not decline appreciably as the share of the DB plan is raised to 50%. An experiment with a higher bequest goal ($500,000) results in a substantially large decline with a DB plan share of 50%. Another finding is that lifetime payouts with the DB plan are significantly higher than with the DC plan, because DB plan members are not affected by asset decumulation (assuming the plan sponsor can make good on any losses). The authors suggest that a 50/50 split could make a lot of sense for many participants, even those with a strong bequest motive, and might keep income variability at a reasonable level.

As any reader of this journal will appreciate, a particular, and even an overwhelming concern of a retiree is the sustainability of her income in retirement. No one wants to run out of money. Investment strategies (assuming that the nest egg is not devoted to an annuity) are commonly appraised by the probability that assets will be depleted while the retiree is alive. The very well-known article by William Bengen posed the question

of the likelihood of asset exhaustion given various asset strategies and assuming some given withdrawal rate.[2] In "Refining the Failure Rate," Javier Estrada proposes two refinements to the standard calculation. He relies on a dataset of stock and bond returns for 21 countries. He begins by calculating the failure rate assuming a 60/40 stock bond allocation over 86 rolling 30-year retirement periods, beginning in 1900. The failure rate is the share of these periods during which assets would be depleted for a given withdrawal rate (the author experiments with rates of 3% and 4%). It turns out that Canada has the lowest failure rate, and the United States has one of the lower rates under either withdrawal assumption.

The two refinements the author proposes are: (1) how long before the end of the retirement period depletion occurred; and (2) the proportion of the retirement period for which the given strategy was effective in sustaining the assumed rate of withdrawal. Javier is being a little modest in terming his proposed supplemental measures refinements. It makes a great deal of difference when a particular strategy fails. The consequences of failure after 15 years are considerably different from those of failure after 29 years, for example. Note that if the retirement period is the same in all countries then the two measures give essentially the same result. Interestingly, the two supplemental measures the paper proposes rank the countries more or less as the failure rate does. Nonetheless, they provide important additional information on the consequences of a particular investment strategy for the sustainability of retirement income.

Bootstrapping, of which there are a number of variants, is a statistical technique that can be applied to project the prices of financial assets and their variability. It is often used in exercises that project the sustainability of retirement income. In its most basic form, bootstrapping begins with a sequence of annual investment returns, usually over a period of 40 years. A random draw repeated 1,000 or more times of 40 of these returns is made, with each random draw representing a simulated outcome over a 40-year period. The results of this simulation exercise are then used to calculate a distribution of returns.

In "The Trimmed Bootstrap: *An Empirical Simulation Technique for Pension Finance Researchers*," Brett Doran and Graham Bornholt take issue with the bootstrap approach. They argue that the method, by treating each annual return as an independent variable, cannot capture situations when a return in a given year is partially dependent on returns in an earlier year or years, or when the stock market is prone to mean reversion. They compare historical outturns for U.S. stocks in rolling 40-year periods drawn over the period 1900–2013 with the results of a bootstrap simulation using the same dataset. The maximum and minimum values of simulated returns substantially exceed those of the historical returns.

The authors conclude that the bootstrap technique produces extreme and unrealistic return estimates. They propose an alternative to the simple bootstrap and its variants, which they dub the "trimmed" bootstrap. The trimmed bootstrap continues to rely on random draws from the historical returns data, but sequences of returns that exceed historical maximums or minimums are rejected. In addition, simulations continue until the distribution of returns is broadly in line with that of the historical data.

The authors test the trimmed bootstrap by comparing its performance against that of a number of variants of the bootstrap method. They find that their method generally outperforms the conventional bootstrap methods, and that its relative performance is not affected by changing the period of the simulation (e.g., from 40 years to 30 years) and the length of the period used to gauge simulated against actual performance. This is a technical but well-written article, which I would invite the interested reader to explore fully.

The life cycle hypothesis of household consumption is one of the workhorses of analytical economics. Its basic thesis is that a household bases its personal expenditure decisions on its wealth, or its expected income, rather than simply on current income. At the same time, the household tries to avoid drastic swings in consumption from one year to the next. Its major implication for any analysis of retirement security is that households will try to save during their working

lives so that their standard of living in retirement does not suffer. In practice, retirement planning relies on the concept of the replacement rate, the ratio of income in retirement to income during a period of working life. Retirement planners often recommend that households target a replacement rate of about 70%, given the expected decline in saving, work-related expenditures, and taxes in retirement.

Although some economists are critical of this application of the replacement rate, it is commonly used to assess the adequacy of income in retirement. In "The Life Cycle Model, Replacement Rates, and Retirement Income Adequacy," Andrew George Biggs presents a critical analysis of the way the replacement rate is calculated by the Social Security Administration (SSA). His critique focuses on two aspects of the calculation: the way wages are indexed and the treatment of years when a worker has no earnings. Biggs argues that the calculation of an average wage over a long period, like 35 years, should index nominal wages to consumer prices. This means that wages are averaged on the basis of their purchasing power. The SSA, however, indexes wages to the economy's average wage index. Because wages grow faster than consumer prices on average, the average wage this calculation produces is greater than the average calculated with the consumer price index. As a result, the replacement rate of Social Security benefits is lower with the SSA's method.

Biggs argues that the SSA's approach is inconsistent with the life cycle theory's assumption that households should be concerned with smoothing their real purchasing power over time. The SSA's approach assumes that a household's income, as calculated for the purpose of determining the replacement rate should reflect gains in economywide productivity. In effect, a household's income *n* years ago is adjusted to reflect the gains in productivity since that date, and its average income thus calculated may be very close to its income at the end of their working life. The resulting estimate will exceed its average career income in purchasing power terms, and is an unnecessary adjustment.

Another thorny issue arises with the treatment of years of no earnings, or "zero years." Counting only the years of positive earnings, as the SSA does in its calculation of the final earnings replacement rate lowers the calculated replacement rate. Andrew argues that workers experience zero years in periods other than pre-retirement. It can be argued that if households smooth consumption earlier in their working life taking into account any zero years, there is no reason to assume they do not do so in the run-up to retirement. This is less of an issue when the denominator of the replacement rate is an average calculated over a long period. These and related issues are carefully addressed in this topical and engaging article.

Although the poverty rate among older Americans as a group has declined substantially over time, the rates for Americans aged 80 and over are more than 40% above the rate of Americans aged 65 to 69. As John Turner, Gerard Hughes, Agnieszka Chłoń-Domińczak, and David Rajnes explain in "Improving Pension Income and Reducing Poverty at Advanced Older Ages: *Longevity Insurance Benefits in Ireland and Poland as Models for the United States*," various life events, including the death of a spouse, higher-than-expected LTC expenses, and greater-than-expected longevity, are behind the higher poverty rate.

The experience of two other countries—Ireland and Poland—suggests that the increase in poverty at older ages is not inevitable. Poverty rates for both men and women aged 75 and above are lower than the rates for the population aged 65 and older. Each country has a flat-rate social benefit that kicks in at age 80 (in Ireland's case) or at age 75 (in Poland's case). The authors attribute at least part of the lower incidence of poverty among the older aged to these benefits. In neither case is the allowance for age particularly high. In 2011, the cost of the Irish benefit was less than 0.1% of GDP.

The authors describe a benefit intended for America's older retirees proposed in the Save Our Social Security Act of 2016 that would pay a flat-rate benefit that would be phased in from age 82 to age 86. Taking account of population aging, the Social Security actuaries estimate that by 2089 the cost of the benefit would be 0.21% of taxable payroll. The authors maintain that this low cost may make the provision of the benefit feasible.

Finally, I am delighted to announce that the article written by Syl Schieber, "U.S. Retirement Policy Considerations for the 21st Century," which appeared in the Fall 2015 issue of *The Journal of Retirement*, has won this year's Peter L. Bernstein Award for the best article. Each year, the editors of the 11 Institutional Investor Journals are asked to nominate the best article published in the preceding year. The 11 articles are then judged by a panel of three outside experts on the basis of the originality or novelty of approach to its subject, the novelty and/or insightfulness of its results or implications, and its practical and academic relevance.

ENDNOTES

[1]In this case, the DB benefit would be 50% of the benefit of full participation in the DB plan. The benefit of the DC plan is set equal to the DB benefit until asset depletion occurs, if it does.

[2]See Bengen reference in the article by Pfau, Tomlinson, and Vernon.

George A. (Sandy) Mackenzie
Editor

Publisher's Note

Institutional Investor, the publisher of *The Journal of Retirement*, wants to extend a special thank you to Bank of America Merrill Lynch for its support of *The Journal of Retirement*. Please note that no sponsor has influence on the editorial content found in *The Journal of Retirement*. Representatives from any firm are encouraged to submit an article to our independent Editor, Sandy Mackenzie, for review and prospective acceptance into the publication. All editorial submissions, acceptance, and revisions are the sole decision of Mr. Mackenzie. The editorial submission guidelines are found on the last page of the publication and online at www.iijor.com. Thank you, and I hope you enjoy this and future issues of *The Journal of Retirement*.

Dave Blide
Publisher, Institutional Investor Journals, dblide@iijournals.com