Institutional Investor Journals

The Voices of Influence | iijournals.com

THE JOURNAL OF
RETIREMENT

SPRING 2014 ■ Volume 1 ■ Number 4 ■ www.iijor.com

Can Collars Reduce Retirement Sequencing Risk? *Analysis of Portfolio Longevity Extension Overlays (LEO)*

MOSHE A. MILEVSKY AND STEVEN E. POSNER

Can Collars Reduce Retirement Sequencing Risk? *Analysis of Portfolio Longevity Extension Overlays (LEO)*

MOSHE A. MILEVSKY AND STEVEN E. POSNER

MOSHE A. MILEVSKY
is an associate professor of finance at the Schulich School of Business, York University and a managing director of the π Longevity Extension (PiLE) Corporation in Toronto, Canada.
milevsky@yorku.ca

STEVEN E. POSNER
is an independent consultant operating out of Hong Kong and New York.
seposner@alumni.princeton.edu

It is by now a well-established fact that poor equity returns early on during the drawdown phase—especially when the withdrawal rate is too high—can greatly reduce the sustainability as well as longevity of a portfolio.[1] This simple fact, which is known as the "sequence of investment returns" risk, has convinced many practitioners to recommend modern insurance products, such as GLWBs and GMIBs,[2] as well as traditional income annuities, in various combinations, to generate income streams that are sustainable over an entire lifetime. The mixture and combination of these types of instruments is known as the product allocation approach to retirement portfolio construction.[3] Other practitioners, motivated by the same challenge, have instead advocated a dynamic asset-liability management strategy using real-return bonds (TIPS) and potentially leveraged ETFs and real annuities.[4] What these different approaches have in common is the recognition that generating a sustainable retirement income is a *different problem* than building the asset base to support withdrawals—i.e., the accumulation problem—and requires more than simply tinkering with stock and bond mixes.

In this article we take yet another, different but in some ways older, approach: we examine whether retirees can protect against a negative sequence of returns early in the withdrawal phase by using traded equity put and call options while continuing to hold or invest in a traditional (accumulation-type) asset allocation. We call this strategy a *longevity extension overlay* (LEO) and discuss how it can be implemented and when it might work. This strategy and our analysis are geared specifically toward wealth and money managers who, for whatever reasons, are unable, unwilling, or uninterested in using insurance-based solutions (i.e., product allocation) to address the problems induced by a poor sequence of returns and longer-than-average life spans.

It must be noted at the outset that the strategy of synthesizing or creating insurance for an investment portfolio using puts and calls, or "collaring," is not new; it has been available to sophisticated investors and used for decades. However, the novelty of this (exploratory) research is to i) discuss the conditions under which a collar can actually extend the *longevity* of an investment portfolio in the *withdrawal* phase, and ii) provide *evidence* on the efficacy of this strategy during the period 2007–2013 using real-world options data that we have managed to obtain.

Another topic that is obviously related to our overall agenda—and is almost impossible to bypass in the retirement income literature—is the appropriate (sustainable) withdrawal or spending rate from a balanced portfolio in retirement. One tends to see numbers in the 3% (safe) to 5% (aggressive) range recommended in the literature, with the understanding that these percentages are expressed in terms of the initial portfolio value and adjusted for inflation each year thereafter. Our only comment on this admittedly controversial topic is that advocating any fixed-in-advance, deterministic, and unyielding rule is irresponsible, and incompatible with rational economic behavior. Alas, stock markets move up and down, so spending rates should adjust to ongoing economic conditions. That said, given the inordinate amount of practitioner writing on safe withdrawal rules, we adopt the "*we assume you spend x% until you go broke*" approach in our simulations and analysis, without necessarily endorsing that philosophy.

On the same note, we leave for future research whether equity options are more or less effective than *tactically* shifting from stocks to bonds (and back again) at various milestones in retirement *vis-à-vis* a given spending rate. This too is a subject of much controversy. For example, a recent paper by Pfau and Kitces [2013] argues, counter-intuitively, that a rising equity glide path might be preferred at retirement. In this article we do not take a position on whether that is a better way to mitigate the sequence of returns. Rather, to put it bluntly, our main point is that *options are an option*.

MEASURING A PORTFOLIO'S LONGEVITY

In a world with no investment volatility or market uncertainty, there is a very simple mathematical relationship or expression for what we label the longevity of a portfolio, or the number of years the income will last. In particular, a portfolio growing at a real rate of g (for example 3%), under a real withdrawal rate of w (for example, \$50,000 per year) and an initial nest egg of M (for example \$1,000,000), will last for exactly L years, where "longevity" L satisfies the following equation:

$$L = \frac{1}{g}\ln\left[\frac{w/M}{w/\mathrm{M} - g}\right] \tag{1}$$

Mathematically speaking, the present value of consumption from time zero (retirement) to time L must exactly equal the "money in the nest egg" M, for L to be the longevity of the portfolio. Inverting that mathematical relationship (in continuous time), and then solving[5] for L leads to Equation (1).

Now let's plug some numbers into Equation (1) to get some intuition for a portfolio's longevity. Assume that you have \$300,000 in your retirement account (IRA, 401(k), and so on) and would like to spend a real \$30,000 per year, which is a very high spending rate of 10% per year, unsustainable in the real world. Next, assume that real inflation-adjusted growth rates are a fixed and constant 3% per year. (Recall that in the world of our example, there is no volatility or uncertainty.)

The ratio inside the brackets in Equation (1) is 1.42857, written to five digits. Taking the natural logarithm of 1.42857 gives us 0.35667, written to five digits; then dividing this number by the interest rate 0.03 leads to the longevity of the portfolio of **L = 11.9** years. *Voilà*: by using Equation (1) we can see that our portfolio will not last very long, consistent with intuition that: $\frac{w}{M} = 10\%$ is an abnormally high retirement spending rate. Notice that we have not used any simulation techniques. This is but a simple time-value-of-money relationship.

Let's now explore the results if the retirement withdrawal rate is reduced. Holding all other factors constant, if consumption is \$25,000 per year (instead of \$30,000 per year), but the growth rate is unchanged at 3% per year, the quantity in the brackets is 1.5625, and its natural logarithm is 0.44629. When you divide by 0.03, you arrive at **L = 14.9** years, which represents a gain of almost three years. This portfolio longevity gain is not very reassuring, but nevertheless it is intuitive: Cut down on your planned spending by \$5,000 per year, and the money will last three years longer. No mystery here.

But what if we take the case in which the consumption plan is fixed and the portfolio behavior (growth rate g) is random and changes from year to year? In this case there is no deterministic fixed, known, expression for L, as there was in Equation (1). Rather, we can now only talk about *anticipated* longevity. From a mathematical point of view, our main research question thus becomes whether the expected median or modal value can be extended using put and call options that are overlaid directly on the portfolio.

To revisit: There is nothing innovative in claiming that purchasing out-of-the-money put options can help protect or insure a retiree's income portfolio against large declines.[6] In fact, one might plausibly argue that this protection is the exact *raison d'être* of buying equity put options. On the other hand, and to the detriment of the strategy we are exploring, both the ongoing put option premiums for these options and the truncated upside for the calls that are sold will act as a substantial drag on the portfolio's growth rate. This is why many investors with long-term (saving for retirement) horizons often choose to forgo the protection provided by puts and calls—especially if there is a natural premium gained by those who sell puts. The question then becomes, *Can we identify **when** in retirement this trade-off might be "worthwhile" if the objective is to draw as much money for as long as possible (or, put another way, if the objective is to increase the longevity of the portfolio)?* Answering this question is the impetus of the research outlined in this article, and our initial results are encouraging.

Note that this article is something more than just an advocacy piece for buying puts funded by selling calls. And, in truth, our "simple" strategy is more difficult to implement than it sounds. We note that there are many ways to construct a (relatively) low-cost longevity overlay on a retirement income portfolio—defined as one in drawdown mode—that create greater sustainability, by judiciously selling the "right" calls and buying the "right" puts. And, while one outcome of using this protection strategy is that portfolio growth will be curtailed (due to the short calls, which often expire in the money), we nonetheless show that the life or longevity of the portfolio can often be extended using such techniques. In the language of stochastic processes: The conditional expected ruin time of the portfolio can be extended by reducing the *diffusion* coefficient, even if it is at the expense of the *drift*.[7]

Note that this longevity extension isn't guaranteed for all sample paths or eventualities, nor is it expected under all combinations of withdrawal and option parameter values. We openly admit that we do not have definitive answers and our work is intended to be exploratory, focusing initially on exploring, *When does a "longevity extension overlay" actually work*? We start to answer this question by examining how some very simple buy/write strategies would have performed historically, during a period covering the 2008–2009 financial crisis.

This article is a proof-of-concept that demonstrates that *even* such a very simple strategy can work in some circumstances. We conjecture there are numerous other dynamic option-based strategies that could further extend portfolio longevity, beyond what is available from a naïve strategy or what is described in this article. For example, preliminary research seems to indicate that using path-dependent (Asian) options can further extend the longevity of a portfolio and may be more effective than basic puts and calls. We leave that avenue of exploration for future research.

Of course, we completely bypass the topic of whether portfolio longevity and the amount of time a portfolio is anticipated to last is something worth targeting. Some might argue that the goal of portfolio longevity suffers from the same weakness as lifetime ruin probability (LRP) minimization, in that it ignores any magnitude of loss. Nevertheless, we proceed under the assumption that "longevity extension overlays" are something worth pursuing because they are a *natural generalization* of the deterministic calculation, underlying Equation (1), that buttresses all retirement income calculations.

The remainder of this article is organized as follows: in the next section we describe the limited historical results, in the section that follows we report on simulation results, and the last section concludes the article.

LIMITED HISTORICAL RESULTS

Let's work through some scenarios using historic data. In these scenarios, we are selling a call with a strike price K_c and then using the proceeds to purchase a put with a lower strike price K_p. Accordingly, if the market price of the asset falls below K_p, you are guaranteed a minimum return, as you have the right to sell it at a price of K_p. However, if the asset's value increases above a value of K_c, you will have to sell it to the holder of the call at a price of K_c, therefore limiting the gains you could have otherwise earned on the portfolio. We examine how this (K_p, K_c) combination would have worked out over the past seven years.

Our raw data—and the core of this analysis—consists of daily "implied volatility" (IV) values on the S&P 500 index, across a range of strike prices, during the period from January 2007 to December 2013, which obviously covers the period of the financial crisis. In particular, we have daily values of the 20-day, three-month and one-year IVs for options that were at the money (ATM), 5% in the money (ITM), and 5% out of the money (OTM). (Recall that the only free parameter in the Black–Scholes–Merton equation is the implied volatility, and expressing prices in terms of IV is common practice among traders.)

Over these 1,700 trading days, almost seven years, the highest IV was 66.75%, on November 20, 2008, and the lowest IV was 12.92%, on February 4, 2011. This, of course, implies that the purchase of options (and portfolio protection) was most expensive in late November 2008 and cheapest in early February 2011.

Methodologically, the entire set of IV numbers was then inverted and expressed as market prices for put and call options using the Black–Scholes–Merton formula, assuming a risk-free valuation rate of 1% and a dividend yield of 1.5% over the entire time period. We openly acknowledge that these two parameters fluctuated over the seven years and more precision could have been achieved.[8] That said, we do have precise IV values, which are the most important determinant of option prices, and thus we are confident that the resulting option prices for equity puts and calls are indicative of the cost of the options on each of those 1,700 trading days. Exhibit 1 displays the results from our strategy and the main (historical) take-away from this study.

Interpreting the Data

Here is how to interpret the results displayed in Exhibit 1. Let's start with Panel B. Looking here, we can see that a retiree who invested $100 on January 1, 2007, and allocated the entire $100 to the S&P 500 Index, did not withdraw anything from the portfolio, and did not protect the portfolio, would have ended up with $122.13 almost seven years later in October 2013 (after 27 quarters). The return for the S&P 500 during the period was approximately 22%. However, if the same retiree had withdrawn $10 each year, or $2.5 each quarter, for a total of $67.50, the portfolio would have been worth only slightly more than $27 in October 2013. This poor result is, of course, due to the very high withdrawal rate as well as the poor sequence of returns. Recall that the bear market came early in this seven-year period (during 2007–2009), and the bull market came later—presenting us (and our unfortunate sample investor) with a classic sequence of returns disaster. Now, looking in between the two extremes, had the retiree withdrawn $6 per year, or $1.5 each quarter, his portfolio would have been worth $65.26 by October 2013.

EXHIBIT 1

Impact of Combination of Written Calls and Purchased Puts on Risk of "Sequence of Returns" Performance of Various Strategies During the Period Jan/2007 to Oct/2013

Panel A: Longevity Extension Overlay: LOWEST Portfolio Value

Spending Rate	No Overlay LOWEST Portfolio Value	C = 106% P = 90%	C = 106% P = 91%	C = 105% P = 92%	C = 105% P = 93%	C = 104% P = 94%	C = 104% P = 95%	C = 103% P = 96%	C = 102% P = 97%
0%	$60.69	$70.76	$71.39	$72.31	$73.40	$75.12	$77.68	$81.75	$86.29
2%	$57.42	$67.12	$67.72	$68.72	$69.81	$71.29	$73.89	$77.89	$82.20
4%	$54.15	$63.49	$64.16	$65.05	$66.09	$67.50	$68.67	$71.28	$73.79
6%	$50.88	$59.81	$59.16	$58.38	$57.44	$57.39	$56.90	$59.24	$61.44
8%	$42.13	$47.83	$43.36	$45.30	$44.50	$43.74	$43.09	$44.91	$46.65
10%	$27.35	$31.46	$30.08	$28.82	$28.27	$27.75	$27.66	$29.75	$31.65

Panel B: Longevity Extension Overlay: FINAL Portfolio Value

Spending Rate	No Overlay FINAL Portfolio Value	C = 106% P = 90%	C = 106% P = 91%	C = 105% P = 92%	C = 105% P = 93%	C = 104% P = 94%	C = 104% P = 95%	C = 103% P = 96%	C = 102% P = 97%
0%	$122.13	$117.40	$113.86	$111.21	$109.05	$107.54	$104.91	$105.40	$105.97
2%	$103.17	$100.17	$96.99	$94.76	$92.89	$91.26	$89.28	$90.45	$91.19
4%	$84.22	$82.72	$80.33	$78.48	$76.50	$75.43	$73.88	$75.23	$76.30
6%	$65.26	$65.87	$63.73	$61.87	$60.34	$59.76	$58.25	$59.91	$61.44
8%	$46.30	$49.01	$46.89	$45.31	$44.50	$43.74	$43.09	$44.91	$46.65
10%	$27.35	$31.46	$30.08	$28.82	$28.27	$27.75	$27.66	$29.75	$31.65

Notice what happens once the "longevity extension overlay" is acquired (i.e., the portfolio is collared). The eight columns in Exhibit 1 represent different combinations of put purchases and written calls—which is an almost costless transaction. The value next to the C and P (in the header row on our columns) denotes the average strike price transacted during the 27 quarterly rebalancing dates. For example, in the first column we have displayed results for someone who, on average, sells three-month calls that are 6% OTM and uses the funds to purchase three-month puts that are 10% OTM, giving us the C = 106% and P = 90%. The paired transactions are not free, but as close to it as possible with a collar.

Now let's review the results of our strategy starting from the very bottom on the very left corner of our columns. Under an abnormally high, 10% withdrawal or retirement spending rate, the final portfolio value is $31.46 under a C = 106% and P = 90% (average) strategy compared to the $27.35 had the portfolio not been protected with an overlay. The same protective effect can be observed when the call is struck at C = 106% and the put is struck at P = 91%, which is slightly more expensive compared to P = 90%: The portfolio with the overlay is worth $30.08 versus the unprotected $27.35 value. Again, the overlay is effective after seven years of observations. In fact, every single combination of puts and calls we examined (i.e., the entire bottom row) results in higher values than the unprotected portfolio. The intuition here comes in two parts: first, *giving away upside participation in exchange for the (expensive) downside is worth the price*—and second, *this strategy is more effective when spending rates are very high*. These pieces of intuition will be confirmed with independent simulations in the next section.

In contrast, once the withdrawal rate is reduced to more reasonable levels, such as the traditional $4 to $6 range, this qualitative result no longer holds true in all cases. For example, at $6 spending the "looser" collar with parameters C = 106% and P = 90% results in a final portfolio value that is marginally higher: $65.87 versus $65.26. But, with "tighter" collars that are more expensive, the result is the opposite: The portfolio with the overlay provides a lower final value compared to the unprotected portfolio. Intuitively, when the spending rate goes to zero—during a period in which markets increased by over 22%—the portfolio with an overlay results in less wealth. Stated bluntly, in every quarter during which the S&P 500 increased by more than 2% to 6%, the gains belonged to someone else (namely the person to whom you sold the option).

But—and this is another important point—even if the final value of the portfolio was worse off in the presence of the longevity extension overlay, the strategy still has its own merits. In other words, even under lower spending rates, when the impact of the sequence of returns is not as pronounced, the overlay is worth considering. How so? The key lies in Panel 1A.

Panel 1A focuses attention on the lowest value of the portfolio during the almost seven-year period from 2007 to 2013. During this period, the S&P 500's lowest quarterly value would have been 60% of the January 2007 value, and a portfolio consisting of 100% equity would have declined in value to $60.69 on the quarter ending April 30, 2009. (Note: This assumes no portfolio withdrawals at all.) If $6 had been withdrawn every year, or $1.5 every quarter, the lowest portfolio value (maximum drawdown) would have been $50.88, also in April 2009. And, under a $10 withdrawal, the lowest portfolio value would have been $27.35, which is also identical to the final value. Finally, in the case of $8 withdrawals, the lowest value of $42.13 occurs in October 2011. The final value is $46.30 and the value in April 2009 is $47.61.

Now we get to the crux of the matter: The lowest portfolio value over a period is highly correlated with the sustainability of a withdrawal strategy. Whereas in the period we are examining, the S&P 500 ended with a gain of 22%, there was no guarantee *ex ante* that this would have been the final outcome. There was obviously the chance—invoking a roulette wheel analogy—that the last few spins could have been red instead of black. So, given the hindsight bias in any such analysis, it is worth examining the *lowest* portfolio value during this period even if the *final* value shows a full recovery.

This examination reveals the point at which even a naïve overlay strategy will look good. Notice how the lowest portfolio values are uniformly higher under all combinations of call and put options, compared to those without the options. In other words, the client will experience less-severe drawdowns and higher "bear market" account values *at all times* and *during all periods*. This is precisely the value this strategy may offer, and it has two effects. One effect is pure economics, which extends the expected life of the portfolio. The second, and perhaps more important, effect is psychological. Skittish investors might actually stay invested if they are more comfortable with the downside protection!

Alas, the Achilles heel of a historical study limited to seven years such as this one is that we cannot determine whether the portfolio longevity was actually extended. We would need, perhaps, a historical record of some 20 to 30 years, encompassing some periods of poor markets, to observe longevity values with and without extension overlays. But we don't have that, which brings us to Monte Carlo Simulations.

SIMULATION ANALYSIS

The previous section examined one realization of history over a relatively brief period of time. To extend our view, in this section we use Monte Carlo simulation to examine many runs over very long horizons. Although the methodology is quite different, the results are similar and for the most part, encouraging.

For this part of the analysis, we used hypothetical (simulated) option prices using the Black–Scholes–Merton formula, and generated 100,000 sample paths for the entire retirement horizon. Exhibit 2 provides a long-term perspective on the impact of collaring the investment portfolio. The table assumes that nominal investment returns follow a geometric Brownian motion with drift $\times \mu = 6\%$ and volatility $\sigma = 20\%$. (This has been calibrated to long-run historical S&P 500 index returns.) An inflation rate of $\pi = 2\%$ is added to the investment returns as well as to the quarterly withdrawals, which range from \$4 to \$6 per year. Dividends on the S&P 500 are 2% and the risk-free valuation rate (for option pricing purposes) is 3%. We simulated three-month returns and three-month option prices that were settled every three months. In practice, longer maturities could be used, partially to correspond with withdrawal (spending) dates.

Since we used the Black–Scholes–Merton formula to price options, and a fixed volatility of 20% (or in some later cases 25%) per year, the relationship between the collar's put and call was implicitly fixed for the entire horizon. This is in contrast to previous section, in which the volatility smile (skew) affected the relative strike prices over time. In our case, the strike prices of the puts and calls are at the top of the corresponding column. The nice thing about our Monte Carlo simulation is that it allows us to build a perfect zero-cost portfolio. That is, we can find the precise pairs of put/call strikes that have a net cost of zero. Once again, this can only be done in a theoretical simulation, since in practice it will be impossible to locate a call/put combination with exactly offsetting prices. We stress that in reality there would have to be some net cash inflow or outflow at the time of the option purchase/sale.

To be very specific, we again assumed a rather mechanical and naïve strategy in which the portfolio is collared anew every quarter by purchasing OTM put options (strike price K_p) with funds obtained from selling an OTM call option (with strike price K_c). The strike price of the call option was dictated by the premium received for the written put, which made the overlay (truly) zero cost. For the most part, though, the implied call option (with identical price) was approximately 4% to 6% above the market level, which is consistent with actual prices observed in Exhibit 1.

Note that for the simulation, the stochastic process generating returns was assumed lognormal and calibrated to the (historical) first two moments of the S&P 500 index. Finally, to partially increase realism, we assumed \$3, \$4, and \$5 consumption strategies (or cases) adjusted by 0.5% for inflation each quarter (so the nominal value of spending increased), which is approximately an inflation rate of 2% for retirees.

Following this logic, in the case where \$5 is withdrawn in addition to the option cash flow settlements,

EXHIBIT 2A

Conditional Longevity of Portfolio (20% Vol.)

Withdrawal	No Collar	P95/C106.16	P90/C112.29
\$3.0	19.5	25.4	22.0
\$4.0	17.8	22.5	19.9
\$5.0	16.3	19.9	18.1

Notes: Generated via Monte Carlo simulation. Conditional on ruin during lifetime.

EXHIBIT 2B

Conditional Longevity of Portfolio (25% Vol.)

Withdrawal	No Collar	P95/C106.28	P90/C112.52
\$3.0	17.5	24.8	21.1
\$4.0	15.9	22.1	19.1
\$5.0	14.6	19.6	17.4

Notes: Generated via Monte Carlo simulation. Conditional on ruin during lifetime.

we extracted \$1.25 from the portfolio at the beginning of each quarter in real terms, adjusted nominally for inflation by 0.5%. So, for example, at the end of Q1 in the simulation we withdrew $1.25(1.005)^1 = 1.25625$ from the portfolio. Then, at the end of Q2 we withdrew $1.25(1.005)^2 = 1.256253$, and so on. Again, this is somewhat of a fiction, since inflation does not progress in this perfect deterministic manner nor does consumption take place quarterly, and so on. The point here is to test the concept and see whether these overlays can extend the life of the portfolio over a very long time horizon. With that preamble, on to the results, which, as we have noted, are quite encouraging.

Notice that in Exhibit 2 the expected (i.e., weighted) longevity of the portfolio is higher with collars at all withdrawal rates, and the lifetime ruin probability values are lower as well. In other words, the collars work. These values are qualitatively consistent with the results from previous section. Relatively speaking, the longevity extension is more effective when the withdrawal rate is higher. Note also the impact or effect of higher volatility assumptions: The more volatile your asset mix, the lower the longevity, all else being equal. But the collar helps.

Stochastic Present Value (SPV)

In this section we offer a slightly more advanced analysis of the efficacy of collars on portfolio longevity, which involves the stochastic present value (SPV) of consumption. This technique can be used to compute lifetime ruin probabilities (LRP) analytically, and can also shed light on the *random* amount of capital required at retirement to sustain a particular lifestyle. The more capital you have, the greater the certainty you will be able to achieve the (desired) lifestyle. So, while the SPV doesn't tell us anything about longevity of the portfolio *per se*, it does give us a sense of how LEOs can make retirement "cheaper," providing the desired lifestyle, defined as withdrawal amounts, at lower cost. The stochastic present value is defined as:

$$SPV = \sum_{i=1}^{\infty} \frac{({}_i p_x)}{\prod_{j=1}^{i}(1+R_j)} \tag{2}$$

where R_j denotes the random portfolio return in the j'th period (year, quarter, month), the withdrawal is one dollar but can be scaled up, and $({}_i p_x)$ is the conditional survival probability. The random variable *SPV* can be simulated by generating sample paths for the entire vector R_j, and then addingup the relevant products.

Of course, if the portfolio is collared, the R_j will be restricted to the range dictated by the call and put options. The mortality probabilities that we used were based on the Gompertz-Makeham law of mortality. Equation (2) collapses to the actuarial present value of a life annuity when the return (denominator) is fixed and constant. And, Equation (2) collapses to the (MBA, textbook) present value of a stream of cash flows, in the absence of mortality. Recall that the right-hand side of Equation (2) is random. So, the SPV can be very high (*not good, you need a lot of money to finance the lifestyle*) or the SPV can be very low (*good, you don't need as much to finance the lifestyle*)—or it can be somewhere in between.

Exhibit 3 is a graphical idealization of our results. Again, the SPV is random regardless of whether you use a longevity extension overlay (e.g., collars) or not. There are no guarantees, and there is always a chance of prolonged poor markets, which can make any retirement plan that is contingent on markets unsustainable. However, notice that in Exhibit 3 the right tail of the SPV graph—the scenarios in which retirement spending is very expensive—is reduced because of the collars. This is yet another indication that portfolio longevity can be extended with the proper use of derivatives, and provides another way to think about the impact of collars, which shift the mass (or density) of the SPV away from the right tail, closer to the middle. And, while each set or blend of option parameters, asset allocation, and spending rate will generate its own unique SPV picture, the results we have examined are consistent with the notion that the right-tail results (bad returns, early on, with a long life) are truncated.

Exhibit 4 shows the stochastic present value explicitly. If you want \$4 for life, then to generate a 95% sustainability (which is less than 5% ruin probability) you need \$144 at retirement if you have an un-collared portfolio. But, if you put a tight collar on at retirement, all you need for the same level of statistical security is \$90 in your nest egg. And if you put on a looser collar, you need \$106, which is still less than the \$144. Ergo, collars help extend portfolio longevity. The expected SPV tells the same story: If you try to finance a \$4 real income stream with an unprotected portfolio, an initial capital of

EXHIBIT 3
Stochastic Present Value (SPV) of Retirement Consumption

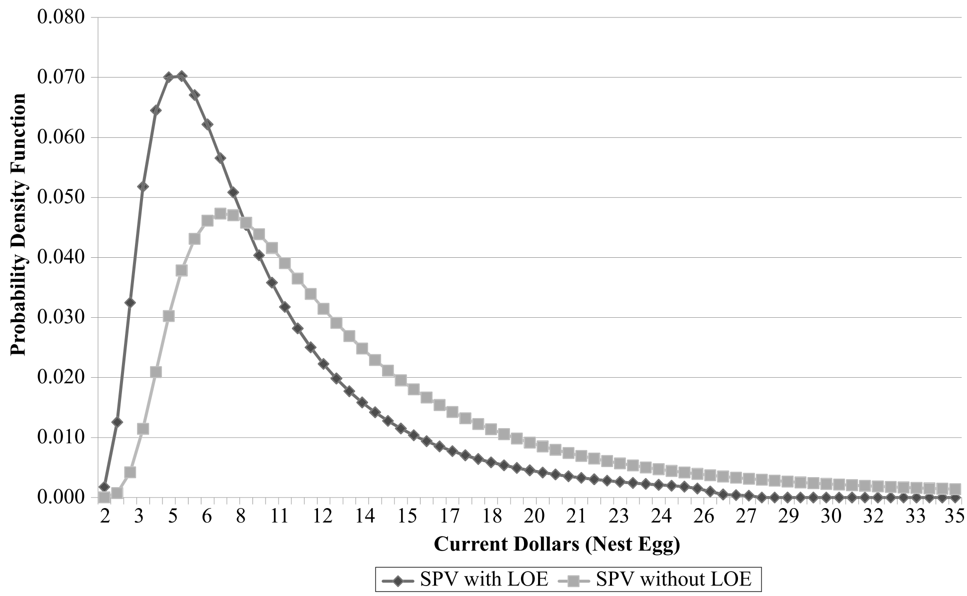

EXHIBIT 4
Stochastic Present Value (SPV) of $4 Withdrawals at 65

Percentile	No Collar	P95/C106.16	P90/C112.29
10%	$33.45	$45.68	$36.08
20%	$40.40	$50.28	$41.70
30%	$46.45	$53.94	$46.54
40%	$52.60	$57.33	$51.16
50%	$59.47	$60.72	$56.12
60%	$67.40	$64.35	$61.55
70%	$77.36	$68.69	$68.18
80%	$91.59	$74.12	$77.15
90%	$116.87	$82.62	$91.99
95%	$144.40	$90.43	$106.94
mean	**$69.69**	**$62.83**	**$61.08**
stdev	$40.61	$14.98	$24.32

Note: GoMa Mortality = [0.0014 85.3901 8.843], 20% Vol.

$69.69 will be enough on average only. These are horrible odds and one should never counsel a retiree to try to finance a $4 income with a $69.69 nest egg. Notice, however, that the tightly collared portfolio induces an expected SPV of $62.83, which is obviously lower than $69.69. Thus you need less capital on average to finance the same withdrawal flow.

The simulated SPV for other withdrawal rates exhibit the same pattern: A smaller nest egg can generate the same consumption stream with the same probability, provided you collar the portfolio with options. As a side note, we believe that the SPV should be the metric by which different investment strategies are compared at retirement. The ruin probability (or its inverse, the sustainability rate) is only the tail of the SPV distribution and might provide a distorted view of any strategy.

Analyzing a Particular Sample Path

Exhibit 5 provides an additional microscopic perspective on the evolution of retirement wealth under a collared strategy, i.e., longevity extension overlay. It is a snapshot of one of the 100,000 individual sample paths, and one in which the collared portfolio actually lasted longer. This outcome—lasting longer—is not a sure thing and is not guaranteed for all scenarios. Notice how, early on, the collared portfolio had a lower value (recall, this is during the most sensitive part of the sequence of return effect) and later on, the paths reversed so that the collar portfolio lasted longer and produced more income.

In addition to the average or median longevity, another way of measuring the success of a particular "longevity extension overlay" is to compare the number of individual simulation scenarios in which the portfolio lasted longer than the non-collared sample paths,

Exhibit 5

One Possible Path of Future Retirement Wealth

which can be considered an overlay success ratio. And, while the two quantities, a high overlay success ratio and increased longevity, are correlated, they are measuring slightly different things.

In the particular simulation run displayed, the short call option turned out to be quite detrimental in the first few years (the option writer gave up the upside). However, when the market declined substantially around retirement year three, the collared portfolio was insulated from the decline as a result of the puts.

Of course, there is an element of *ex post* storytelling in all of this. But the main takeaway is that longevity can be extended even though, or when, its growth rate is curtailed. This phenomenon obviously depends on the spending rate and the cost of the options, so it doesn't work for *any* overlay in *any* market situation.

CONCLUSION

In a recent article, Blanchet and Kaplan [2013] introduced a metric they call "gamma," which attempts to measure the benefit of proper financial advice and guidance. They claim this "gamma" can be more important than alpha and beta in determining the relative success of a given retirement plan. Another way to describe what we have done in this article is that we attempted to derive another source of "gamma" around the usage of equity options during retirement. We conclude with the following observations:

1. There are many different ways to structure an LEO using traded put and call options, but not all of them will extend the life of the portfolio.
2. In general we find that when withdrawal rates are low, i.e., less than \$3 per \$100 of initial nest egg, the LEO is not very effective and can sometimes reduce the longevity of the portfolio. However, with higher withdrawal rates, the expected longevity may be extended. In other words, giving up upside in exchange for downside protection is not a guaranteed proposition for retirees.
3. During the 2007 to 2013 period, the "longevity extension overlay" was more effective at higher spending rates, especially when the sequence of returns was more pronounced. The terminal values were not always higher in the collared portfolio, but the maximal drawdown, which is indicative of potential problems, was much lower.

We believe that while these are preliminary results, they are interesting in and of themselves. The next step is to investigate whether option-based strategies are more or less effective in increasing portfolio longevity relative to asset allocation-based strategies to help protect a given level of spending. Another avenue worth pursuing is whether path-dependent (versus puts and call) options are more effective as LEOs, and whether collaring only a fraction or percentage of the portfolio, perhaps to protect a given level of spending, is an even better option.

In sum, despite the complexity of traded options and the persistent questions regarding their suitability for elderly and retired investors, a strong case can be made that they can effectively help manage retirement risks. Indeed, we believe that the options embedded inside modern variable annuities (VA) with withdrawal benefit guarantees are no less complicated than the strategies described in this article. Buying them directly should be considered no less suitable than buying them as part of a more expensive package.

ENDNOTES

The authors would like to thank the editor (Sandy Mackenzie) and an anonymous reviewer for helpful comments on an earlier version of the manuscript, as well as Branislav Nikolic (QWeMA Group) for research assistance and Alexandra MacQueen for editorial assistance. The authors also acknowledge helpful conversations with Tom Salisbury, Huaxiong Huang, Faisal Habib, and Lowell Aronoff as this research was conducted.

[1]See, for example, the recent paper by Laster et al. [2013] on the pitfalls in retirement portfolio construction and the impact of spending rates on sustainability.

[2]Guaranteed Living (or Lifetime) Withdrawal Benefit (GLWB) and Guaranteed Minimum Income Benefits (GMIB) are variable annuities (VAs) fused with downside protection provided in the form of life annuities.

[3]See, for example, the work by Pfau [2013], and the references cited therein.

[4]This includes the stream of literature described in papers by Sexauer et al. [2012]; Scott and Watson [2013]; and Fan et al. [2013].

[5]See the book by Milevsky [2012, Chapter 1] for a derivation, explanation, and history of this equation. In particular, the relevant integral equation is: $M = \int_0^L we^{-gt}\,dt$.

[6]An early paper in this strand of literature was written by Dybvig [1999], advocating a dynamic portfolio strategy to protect spending for endowments, which is similar to the retirement income dilemma. Using options to collar retirement income portfolios was also proposed in Milevsky and Abaimova [2006] and was further analyzed in the PhD dissertation of Wang [2006].

[7]The key references on "expected ruin time" are Browne [1997] and Heath et al. [1987]. There is a technical problem in that the expected ruin time might not be defined, mathematically, because the non-zero probability the underlying diffusion process escapes to infinity. So, in our simulations, the definition of portfolio longevity is the expected amount of time before the portfolio is ruined, conditional on those sample paths in which ruin occurs prior to death. An alternative definition involving the median and/or modal value is also plausible.

[8]More to the point, if we had an endless supply of enslaved PhD students (which we do not), we would have forced them to use the daily interest rate and volatility for each of the 1,700 trading days.

REFERENCES

Blanchet, D., and P. Kaplan. "Alpha, Beta and Now… Gamma." *The Journal of Retirement*, Vol. 1, No. 2 (2013), pp. 29-45.

Browne, S. "Survival and Growth with a Liability: Optimal Portfolio Strategies in Continuous Time." *Mathematics of Operations Research*, Vol. 22, No. 2 (1997), pp. 468-493.

Dybvig, P.H. "Using Asset Allocation to Protect Spending." *Financial Analysts Journal*, Vol. 55, No. 1 (1999), pp. 287-313.

Fan, Y., S. Murray, and S. Pittman. "Optimizing Retirement Income: An Adaptive Approach Based on Assets and Liabilities." *The Journal of Retirement*, Vol. 1, No. 1 (2013), pp. 124-135.

Heath, D., S. Orey, V. Pestien, and W. Sudderth. "Minimizing or Maximizing the Expected Time to Reach Zero." *SIAM Journal of Control and Optimization*, Vol. 25, No. 1 (1987), pp. 195-205.

Laster, D., A. Suri, and N. Vrdoljak. "Pitfalls in Retirement," *The Journal of Retirement*, Vol. 1, No. 1 (2013), pp. 91-99.

Milevsky, M.A. *The 7 Most Important Equations for Your Retirement: The Fascinating People and Ideas Behind Planning Your Retirement Income*. Toronto: Wiley, 2012.

Milevsky, M.A., and A. Abaimova. "Risk Management During Retirement." In *Retirement Income Redesigned: Master Plans for Distribution*, edited by H. Evensky and D.B. Katz, New York: Bloomberg Press, (2006), pp. 163-184.

Pfau, W.D. "Analyzing an Income Guarantee Rider in a Retirement Portfolio," *The Journal of Retirement*, 1 (2013), pp. 100-109.

Pfau, W.D., and Kitces, M.E. "Reducing Retirement Risk with a Rising Equity Glide-Path." 2013. Available at SSRN: http://ssrn.com/abstract=2324930 or http://dx.doi.org/10.2139/ssrn.2324930.

Scott, J.S., and J.G. Watson. "The Floor-Leverage Rule for Retirement." *Financial Analysts Journal*, Vol. 69, No. 5 (2013), pp. 45-60.

Sexauer, S.C., M.W. Peskin, and D. Cassidy. "Making Retirement Income Last a Lifetime." *Financial Analysts Journal*, Vol. 68, No. 1 (2012), pp. 74-84.

Wang, J. "Numerical PDE Techniques for Personal Finance and Insurance Problems," 2006. Ph.D. Thesis on file at York University, Toronto, Canada.