Bolton

July 2024

Risk Transfer and Liability Measurement for Variable Annuity Plans with Caps and Floors

PREPARED BY

Ellen Kleinstuber, FSA, FCA, FSPA, EA, MAAA
Scott Steadman, FSA, EA, MAAA, FCA
Tom Lowman, FSA

Bolton
1 W. Pennsylvania Ave.
Suite 600
Towson, MD 21204

410.547.0500
www.boltonusa.com
solutions@boltonusa.com

CONTENTS

INTRODUCTION

Variable annuity plans have been growing in popularity among plan sponsors as a tool for improving the sustainability of defined benefit plans and their guarantee of a lifetime income stream. By sharing the investment risk with plan participants (like in a defined contribution plan), plan sponsors defray their risks while maintaining the benefit of longevity and other demographic experience pooling to provide a more cost-effective means of providing participants with annuity income than having to purchase an individual annuity benefit.

A "pure" variable annuity plan transfers all investment experience above/below the "hurdle" rate to participants by adjusting their benefit up or down in proportion to the actual return compared to the hurdle rate. To mitigate the risk to participants during times of extreme investment losses and provide the plan sponsor with a cushion after achieving significant investment gains, many of these plan designs include features to limit the annual adjustment made to participant benefits.

- A "floor" rate limits the downward adjustment to the participant's benefits, causing the plan to retain any investment losses below the floor.
- A "cap" rate limits the upward adjustment to the participant's benefit, allowing the plan to benefit from the excess returns as an offset to any retained or future losses from the operation of the floor rate.
- The combination of a cap and a floor is sometimes referred to as a "collar," which can be symmetric around the hurdle rate or skewed toward either the cap or floor.

This paper is written to assist actuaries and investment professionals with understanding the issues to consider when attempting to design or properly value benefit liabilities for variable annuity plans that include these features.[1]

The American Academy of Actuaries' December 2022 paper on Measuring Pension Obligations for Difficult-to-Value Plan Provisions (the "Difficult-to-Value Paper") provides a framework on how to value many designs including variable benefits with floors and caps. Such modified variable benefit designs straddle the line between a fixed benefit and a pure variable benefit.

This paper takes the concepts introduced in the Difficult-to-Value Paper a step further by discussing: (1) how to measure transferred and retained risk, and (2) how the use of median and average benefits and discount rates can be used to determine liabilities. We focus on median values but also discuss more conservative approaches. We show that as a design gets closer to a pure variable benefit design, it is best not to assume below mean benefit increases (e.g., median) with a mean discount rate or the funding result will probabilistically be under 100%. Similarly, using the higher mean expected benefit increase and the lower median discount rate will be conservative with an expected funding result of over 100%.

[1] The American Academy of Actuaries also published a paper on Variable Annuity Plans in 2019 that provides a broad overview of the administrative, funding, and investment issues for considerations with these programs. That paper focuses primarily on the case of the pure variable annuity plan.

The discussion herein focuses on issues related to these programs for private sector, multiemployer defined benefit plans. While much of this material is also relevant to the valuation of these programs for governmental employers, the regulatory requirements related to funding are different and may require additional considerations beyond the scope of this paper. Likewise, private sector single employer defined benefit plans have specific funding rules that make valuation of liabilities for variable annuity plans more complex, and those requirements are not addressed in this paper.

We have divided this paper into five sections.

- **Section 1** is about the transfer of investment risks between the plan and the participants for variable annuities with floors and/or caps.
- **Section 2** is about determining present values and discount rates that might be used for funding. These calculations are based on a multi-year, single payment forecast and follow the concepts discussed in the Difficult-to-Value Paper.[2]
- **Section 3** address discounting based on bond rates and looks at investments in bonds as it relates to these variable annuity plans with caps and floors. We discuss how this calculation is related to the determination of a Low-Default-Risk Obligation Measure ("LDROM") under ASOP No. 4.[3]
- **Section 4** is about the need for uniformity when selecting the discount rate and projected benefit changes and understanding whether asymmetries are biasing the results.
- **Section 5** discusses other issues such as how changes in capital market assumptions impact plans with floors or caps even though they do not impact the valuation of pure variable benefits, recovering from poor investment return years and vulnerabilities that variable annuity designs have to sequence of return risk, and harmonic averages.

The examples in this paper are generally based on having both a floor and a cap that are symmetrical around the hurdle rate (e.g., if the hurdle rate is 5%, the floor might be 0% and the cap 10%). However, the calculation concepts can be adapted to designs that are either asymmetrical or have only a floor or a cap.

---

[2] Table 2 shows what we mean by a single payment forecast. While annuities are a series of payments, the model is sufficient to develop expected benefit increases and net discount rates for many designs.
[3] See Section 3.11 of ASOP No.4, Measuring Pension Obligations and Determining Pension Plan Costs or Contributions (updated December 2021).

I. INVESTMENT RISK TRANSFER

Pure variable annuity plans transfer most of the investment risk to plan participants while still providing benefits in the form of an annuity. Plan Sponsors or Trustees also value designs that are expected to be essentially cost-neutral to a pure variable annuity design but transfer less investment risk to participants. One way to do this is to limit the losses that are applied to benefits in years where the investment returns are below the hurdle rate. The increase in minimum funding requirements for downside protection and other valuable plan provisions (e.g., benefit lock-in at retirement) can be offset through setting a cap on gains that are applied to benefits. For example, a variable annuity design might have a "hurdle" rate of 5% and exclude losses for investment returns below 0% and investment gains above 10% when adjusting benefits.[4] The cost can be much different from that of a design with no floors or caps (see Section 2) and the risk borne by the plan (and participants) is different.

Investment risk is often the biggest risk associated with traditional defined benefit plans. Future investment risk is usually expressed as the standard deviation of return, often called volatility. Values for annual standard deviations of returns for common investment mixes used by plans typically range from 10% to 14%. With a Variable annuity design with floors and caps of + or - 5%, most of the risk remains with the plan and not the participant and this paper seeks to quantify that for different designs and investment mixes.

For our analysis, we start with an investment mix that produces an expected arithmetic return of 7% and a standard deviation of 12% and a benefit design with a hurdle rate of 5%. We review three different caps and floors, each of which are symmetrical above and below the hurdle rate. The combination of a cap and a floor is often referred to as a "collar". For example, if the return is -2% on a plan with a 5% hurdle rate and a collar of 0% - 10% (i.e., 5% above and below the 5% hurdle rate), the benefit loss is based on a 5% investment loss (0% floor on return less 5% hurdle rate) and the plan's investment loss is 2% (return less floor). We distinguish the standard deviation of returns separately for both the returns used to determine benefits and net returns applied to the plan assets.

To evaluate the risk, return, and liabilities associated with various variable Annuity designs, we performed a Monte Carlo simulation using 10,000 trials. This model is used to develop the results presented in all the tables included in this paper. Details regarding the assumptions and Monte Carlo modeling technique used are included in Appendix B.

Table 1a below shows the standard deviation of returns retained by the plan and the participant using three different symmetrical collars as well as a pure variable Annuity design. The "Average net return above the hurdle rate credited" is the average (over the 10,000 trials) of the return not retained by the plan (i.e. provided to the members). The "Average return credited to plan" is the average of the return retained by the plan (i.e. not provided to the members). These

[4] Some plans define the cap as a maximum benefit adjustment of +5% vs. a maximum investment return applied to the adjustment (e.g., 1.1/1.05 – 1 = 4.76%). Both styles of adjustment have received favorable determination letters from the IRS. The adjustment should be discretely defined in the plan documents to satisfy ERISA's "definitely determinable" rules.

first two values should sum up to about 7%. The "Standard deviation of participant credit" is the standard deviation of the results used to calculate the first column of averages and the "Standard deviation retained by the plan" is the standard deviation tied to the second column of numbers. Since standard deviations measure changes around the average and neither the plan nor the participant have an average return of 7%, the sum of the last two columns will not add up to 12% (the standard deviation of the portfolio).

Table 1a – Reallocation of return and risk tied to design

COLLAR = HURDLE RATE +/-	AVERAGE NET RETURN ABOVE HURDLE RATE CREDITED	AVERAGE RETURN CREDITED TO PLAN	STANDARD DEVIATION OF PARTICIPANT CREDIT	STANDARD DEVIATION RETAINED BY PLAN
1%	0.09%	6.94%	0.97%	11.26%
5%	0.45%	6.58%	4.38%	8.53%
10%	0.87%	6.16%	7.54%	5.84%
Pure Variable	2.03%[5]	5.00% (all trials equal hurdle rate)	12.02%	0.00%

Alternatively, a plan could adopt a more conservative asset allocation. Table 1b below reflects a portfolio with an expected arithmetic return of 6% and a standard deviation of 10%.

Table 1b – Revised Table 1a with lower expected return

SYMMETRICAL CAP AND FLOOR	AVERAGE NET RETURN ABOVE HURDLE RATE CREDITED	AVERAGE RETURN CREDITED TO PLAN	STANDARD DEVIATION OF PARTICIPANT CREDIT	STANDARD DEVIATION RETAINED BY PLAN
+/- 1%	0.04%	5.98%	0.97%	9.24%
+/- 5%	0.22%	5.81%	4.28%	6.55%
+/- 10%	0.43%	5.59%	7.11%	4.03%
Pure Variable	1.02%	5.00%	10.01%	0.00%

[5] In theory this is 2% (7% - 5%), however it is common for Monte Carlo results to differ slightly.

Two items worth noting:

1. **Investment advisors typically target asset allocations with an expected average return greater than the hurdle rate.6 This implies two things: the variable benefit participant is more likely to have their benefits adjusted upwards than downwards in any single year, and the cost of the benefit for minimum funding purposes under ERISA is greater than a traditional benefit valued at the expected rate of return if the starting benefit is the same.**
2. **An asymmetric collar is needed in order to have costs similar to a pure VAP (assuming an expected portfolio return higher than the hurdle rate). While picking a symmetrical collar around the hurdle rate is an option, if the inclusion of these features is driven by finding a design that costs about the same as a pure variable annuity (given a common accrual rate), the collar would necessarily be asymmetrical when the expected return for the portfolio is greater than the hurdle rate.**

How do plans deal with retained risk?

The transfer of any of the investment risk to participants, particularly downside investment risk, will generally help the plan's projected solvency and lessen the potential liability of the plan sponsor. However, any retained risk should be viewed over multiple time frames (e.g., 5, 10, 15, 20+ years) and economic scenarios so that plan sponsors can appreciate the true risk profile.

For example, assume we have an assumed rate of return of 7.0%, a symmetrical 5% collar around a 5% hurdle rate and the plan experiences a return of -7%. The total loss retained by the plan is 7%. In the next year, returns of 10% or less will not offset any of the prior year's loss that was borne by the plan. A single year return of 18.28% is required to offset the entire prior-year loss (ignoring the effects of cash flows).[7] With a traditional design the plan's loss might have been 14% (7% assumed rate of return compared to -7% actual experience) or more. Is it easier to cover a 7% variable design loss or a 14% traditional design loss? We attempt to answer this in Section 5. In any case, it is important to understand whether contributions (exceeding normal cost) are sufficient to cover material losses both now and in the future. We expect that many plans will convert to variable designs for active or new members only. In this situation, when future retirees also have modified variable annuities the plan may retain greater degrees of investment risk due to floors and caps (but still less than for fixed benefits).

---

[6] The authors recognize that a subset of the investment advisor community regularly advises variable annuity plans to optimize the plan's asset allocation to target the hurdle rate as the expected return for the portfolio. Whether such approach is prudent is outside the scope of this paper.

[7] To offset a 7% loss in one year, an excess return of 7.53% is needed in another year (1.00 / 0.93 = 1.07527). This must be on top of the positively adjusted liabilities reacting to a 5%-capped return, meaning a total return of 18.28% (1.07527 * 1.10 = 1.182797) is needed to be completely back to liabilities and assets matching. For mathematical purposes here, we have assumed no mortality and no benefit payments removed from the liabilities.

II. DETERMINING A PRESENT VALUE AND FUNDING BASED ON EXPECTED RETURN

For a pure variable annuity plan the identical present value can be calculated by either:

- projecting out the expected benefit stream (with expected increases or decreases relative to the hurdle rate) and discounting at the expected rate of return, or
- using the current benefit (as if it does not change) and discounting at the hurdle rate.

For cash flow purposes, option (1) is preferred, while option (2) requires less computing power and is easiest to explain. The expected benefit increase using a 7% expected rate of return with a 5% hurdle rate is 1.07/1.05 -1 = 1.9%/year, which would be used to project the expected benefit stream in option (1). Alternatively, under option (2) you could use the median benefit discounted by the geometric rate and get the same answer in terms of present value even though the average benefit increase is different (as shown in Table 4 below). While Section 4 focuses on understanding the relationship between the discount rate and the expected increases it produces and why that is important.

When we implement caps and floors, the two options in the prior paragraph no longer provide the same result if:

- investment returns are projected stochastically by year, generating a sequence of returns displaying volatility above a cap or below a floor, or
- investment returns are projected at the expected return or median return which sits above the cap or falls below the floor.

In either of these situations, using the current benefit and discounting at the hurdle rate no longer produces the same value as projecting out the expected benefit stream and discounting at the expected rate of return. In either of these cases, a plan with a floor and/or cap will often be valued somewhere between a fixed benefit and a pure variable annuity (as shown in Table 2 below).[8] If there is only a floor or a cap, values will be more or less valuable (respectively) than a pure variable annuity.

[8] We know that fixed benefits have a different present value that depends on whether we discount at the expected arithmetic mean rate or the median rate. The present value for a pure variable annuity does not vary. Therefore, a benefit that falls between the two designs will still be sensitive to the discount rate selected. This makes for an interesting choice when a plan has both fixed benefits and variable benefits subject to a cap and floor.

When working with plans with floors and/or caps three possible valuation methods emerge:

1. **Projecting the expected benefit payment stream (reflecting variable annuity adjustments) and discounting at the rate used to determine the expected benefits.**
2. **Using the current benefit and discounting at a rate (different than the hurdle rate) determined relative to the expected annual change in benefits.**
3. **Valuing either a fixed or variable benefit and applying a load.**

Often 1. and 2. will be calculated using the same exercise.

Monte Carlo Projections of Benefits and Present Values

Using the concepts in the Difficult-to-Value Paper, we produced the following table showing expected benefits for a five-year period using a single investment return trial.[9] This table compares the present value methods one and two described above and provides an example of the range of results that might occur for a range of variable annuity plan designs.

Table 2 – Single Trial

	STOCHASTIC PRIOR YEAR RETURN	FIXED BENEFIT	PURE VARIABLE WITH 5% HURDLE RATE	VARIABLE WITH 0%/10% FLOOR/CAP	VARIABLE WITH -5%/15% FLOOR/CAP
Time 0	N/A	$1,000	$1,000.00	$1,000.00	$1,000.00
Time 1	18.56%	$1,000	$1,129.14	$1,047.62	$1,095.24
Time 2	16.48%	$1,000	$1,252.60	$1,097.51	$1,199.55
Time 3	-4.30%	$1,000	$1,141.65	$1,045.24	$1,093.30
Time 4	-5.00%	$1,000	$1,032.92	$995.47	$989.18
Time 5	-0.48%	$1,000	$979.01	$948.07	$937.55
Present Value at time 0 of benefit at time 5 discounting using actual returns		$800.32	$783.53	$758.76	$750.34
Present Value at time 0 of benefit at time 5 discounting using 5%		$783.53	$767.08	$742.84	$734.59

Although illustrative, focusing on the benefit after five years in one single trial does not provide a definitive answer with respect to determining the expected present value at time 0 of the projected benefit at time 5. However, the average and median present value of 10,000 trials

[9] For simplicity, we have dropped the +/- 1% scenario discussed in Section 1 and focus on the design options with a wider collar.

produces the following present values at time 0 of the projected benefit at time 5 and the equivalent discount rate:

Table 3 – Multiple Trials

PRESENT VALUE METRIC (TIME 5)	FIXED BENEFIT	PURE VARIABLE WITH 5% HURDLE RATE	VARIABLE WITH 0%/10% FLOOR/CAP	VARIABLE WITH -5%/15% FLOOR/CAP
Median	$734.47 (6.37% implied discount rate)	$783.53 (5.00% implied net discount rate)	$751.28 (5.89% implied net discount rate)	$767.35 (5.44% implied net discount rate)
Mean[10]	$758.13	$783.53	$759.11	$762.03[11]

What is the appropriate time horizon for reviewing trials?

The appropriate time horizon for projections depends on the purpose of the projection. If our goal is to project every benefit until death for every participant, extensive models will be required that go out 50+ years. If we focus on percentile results, the projection period matters particularly for the more extreme percentiles which should converge as the projection period lengthens. This might mean using a projection period of 30 or more years. If there are minimum benefits (e.g., initial $1,000), trials of at least 15–20 years will be needed. However, in many cases, mean and median results for expected increases and discount rates can be divined over short and intermediate term periods of 5–10 years to simplify valuations.

Can we estimate an average expected increase and value using the gross return rate?

In Table 3, the implied discount rate for a fixed benefit is 6.37% over the five-year projection period and 5.89% for a variable benefit with a 0%/10% cap/floor. Consider Table 1a above where the average return credited to the benefit under the 0%/10% floor/cap design was 0.45%. This compares to 6.37% - 5.89% = 0.48% in Table 3. Valuing an annual increase in this range and discounting at 6.37% would offer a reasonable present value and cash flow forecast.

[10] We have focused more on median results than average results because median results are consistent with what is produced by the geometric rate of 6.37%. Average results are harmonic averages and tend to have higher present values and lower discount rates than either arithmetic or geometric results, however harmonic averages are not commonly used in current practice. See Section 5 for further discussion.
[11] It is perhaps surprising that an arithmetic mean value would be below the median value for the same dataset. This perceived anomaly is discussed in Section 5.

As noted previously, the projection period matters when determining median and mean benefits and present values. In Tables 2 and 3, we used a relatively short five-year projection period so we could concisely illustrate the annual projected benefits under each plan design. Now, we extend our projection period to 10 years to reflect something more akin to a reasonable investment cycle. If, instead of focusing on the present value, we focus on the projected benefit we produce these values:

Table 4 – Projected Benefits

	FIXED BENEFIT	PURE VARIABLE WITH 5% HURDLE RATE	VARIABLE WITH 0%/10% FLOOR/CAP	VARIABLE WITH -5%/15% FLOOR/CAP
Median Benefit at Time 10	$1,000	$1,133.27 (1.26% average annual increase)	$1,034.27 (0.34% average annual increase)	$1,059.19 (0.58% average annual increase)
Mean Benefit at Time 10	$1,000	$1,206.75 (1.90% average annual increase)	$1,043.30 (0.42% average annual increase)	$1,087.57 (0.84% average annual increase)

Does the choice whether to discount at the arithmetic mean or median matter?

Yes, although decreasingly so as the design approaches a pure variable benefit, and ultimately it is more important that the measure used is handled consistently between the discount rate and expected benefit adjustments (see Section 4).

We expect that a fixed benefit will have a different present value depending on whether the discount rate is the arithmetic mean (no gain or loss) rate or the median (long-term geometric) rate. The median rate produces a higher present value of traditional benefit liabilities, but this may not always be the case with a variable annuity with a floor/cap structure.

For a pure variable benefit, the selected rate should not matter, and all present values should be identical and tied to the hurdle rate. As shown in Table 4 above, the median annual increase in our example for a pure variable annuity is 1.26%, and yet the mean rate is much higher at 1.90%.[12] This 0.64% difference between the 1.90% and 1.26% increases should approximate the difference between the assumed arithmetic return (7%) and the median return (6.33%) divided by 1+hurdle rate ((.07-.0633)/1.05 = 0.638%). Theoretically, it should not matter whether we use an assumed 1.90% increase and discount at an arithmetic rate, or an assumed 1.26% increase and discount at the median rate. In the examples illustrated, there may be some small variation since this is a Monte Carlo model and results can fluctuate slightly from their theoretical values.

[12] We expected 1.07/1.05 – 1 = 1.9%. Any individual trial in the Monte Carlo simulation produces a set of both returns and benefit increases for a variable annuity design. These increases are specific to that trial. Similarly, the median trial has a specific set of returns and benefit increases. Arguably, there may even be a different median trial at each time horizon of the projection. These are not the same as the set of returns and benefit increases based on the average of all trials. Present values should not be based on mixing the benefit projections on one basis (e.g., the average of all trials) with the investment returns from another basis (e.g., median returns or similar long term geometric returns).

As previously noted, variable annuity designs with symmetrical caps and floors will likely produce values that fall somewhere in between pure variable annuities and fixed benefits. We expect that the spread between the average and median expected benefit increase will be lower (less than 1.90% minus 1.26% in our Table 4 example) and the decision to determine present values using either arithmetic mean or median expected returns will still matter, even if not as much as for a fixed benefit. In our example above (Table 4) a variable annuity with a floor of 0% and a cap of 10% has an average benefit increase of 0.42% and a median benefit increase of 0.34%. If the plan is funding toward a median result, the 0.34% median benefit increase and median (long term geometric) discount rate should be used. If a 0.42% average annual benefit increase is assumed, then the arithmetic mean discount rate should be used.

As the percentage of returns falling between the floor and cap increases, the normal relationship of the median and mean benefits no longer applies. If more than half of the results are within the cap and floor range, the median benefit and present value will be the same as a pure variable annuity benefit. In this case, tail events will make the mean benefit higher than the median benefit, but the harmonic present value (see Section 5) could be lower than the median present value.[13]

As noted earlier, projecting the expected benefit (with expected increases or decreases) and discounting at the expected geometric rate of return is most appropriate with respect to plan cash flows. However, this is not always practical. Just as the pure variable annuity can be valued using the current benefit and using the hurdle rate as the discount rate, the two designs above could be valued using the current benefit and the implied discount rate. It is important to consider what time horizon to use to determine the implied discount rate. A five-year projection may not be enough. For example, longer projections might be necessary when the plan design uses investment returns averaged over several years (and recent history needs to be considered) or the initial benefit at retirement is a minimum/floor benefit.

At the top of this section, we mentioned that “loads” can be used to approximate the extra value that certain variable plans provide on a present value basis, without having to compute the present value directly. Sometimes loads might be appropriate particularly when the results approximate either a fixed benefit or a pure variable benefit. If a design is close to pure variable annuity design the load might be applied to the variable annuity present value. If a design is closer to fixed annuity design the load might be applied to the fixed benefit annuity present value. If we look at Table 3 we see that a variable design with a 0%/10% floor/cap has a mean present value of $759.11 which is close to the $758.13 present value for a fixed benefit. That might imply a load of $759.11/$758.13 – 1 = 0.12%. Table 3 was based on a 5-year projection so other longer periods might also need to be reviewed before setting a load.

[13] The incidence of tail events is an important consideration when selecting an economic scenario generator (“ESG”) for use in developing the investment returns used to generate projected investment returns. This is beyond the scope of this paper.

III. DETERMINING A PRESENT VALUE AND FUNDING BASED ON BOND RATES

As an alternative to measuring variable annuity benefit present values using expected return assumptions, we could measure the liability assuming plan investments are in high quality bonds.[14] Normally when we use bond rates to measure liabilities, we use the expected return for "securities whose cash flows are reasonably consistent with the pattern of benefits expected to be paid in the future."[15] Section 3.11 of ASOP No. 4 says:

For purposes of this obligation measure, the actuary should consider reflecting the impact, if any, of investing plan assets in low-default-risk fixed income securities on the pattern of benefits expected to be paid in the future, such as in a variable annuity plan.

Should the expected total return be used?

Note that the expected total return reflects price fluctuations of underlying fixed income securities. The standard deviation of investment returns for 5–10-year Treasuries is about 4% while for 10–20-year Treasuries it is about 8.0%. To the extent that this is partly driven by price changes due to changes in duration and the shape of the yield curve, significant complexity may be introduced to any valuation model.[16]

If a plan combines traditional and modified variable annuity designs (perhaps where only new entrants receive benefits subject to variable annuity adjustments), selecting the duration presents the question whether different durations are selected for the two aspects of the benefit design. If the duration of the variable benefit is determined separately, the duration of that portion of the total plan liability may start very high and decline materially over time as the plan matures.[17]

Moving in the direction described in the prior paragraphs will reflect the impact "of investing plan assets in low-default-risk fixed income securities on the pattern of benefits expected to be paid in the future."[18] It may provide "appropriate, useful information for the intended user regarding the funded status of a pension plan" and "a more complete assessment of a plan's funded status and provides additional information regarding the security of benefits that members have earned as of the measurement date."[19]

---

[14] This paper uses a Monte Carlo model to simulate future economic scenarios. The Difficult-to-Value Paper notes other types of models that can be used (e.g., option pricing). These alternative models produce market related values of liabilities which is not the funding basis for ERISA multiemployer plans. However, for purposes of the Low Default-Risk Obligation Measure ("LDROM") under ASOP 4, these other models (while more complex) may also be appropriate.

[15] ASOP No. 4 – Measuring Pension Obligations and Determining Pension Plan Costs or Contributions, Last Revised December 2021 ("ASOP 4")

[16] Some ESGs used in non-pension fields incorporate the shape of the yield curve into their models, such as the one promulgated by the National Association of Insurance Commissioners ("NAIC"). Such ESGs are not regularly used in the valuation of pension plans in the United States.

[17] These designs, where variable benefits are provided to new entrants beginning after a certain date, are increasingly popular.

[18] ASOP No. 4, Section 3.11.

[19] Ibid, ASOP 4 History section in Transmittal Memo.

One approach is to use current market rates (including yield curve/forward rates) and assume they do not change. However, this could be viewed as changing the plan design to allow investment returns to be calculated based on book value returns. We can test this simplified method to see if it is reasonable for a specific plan design.

Using our Monte Carlo model, we updated the analysis presented in Table 3 using a fixed income portfolio an expected return of 4% and a standard deviation of 4% and looking at the benefit after 10 years (rather than five): [20]

Table 5a – Present Value Measures Using a Fixed Income Portfolio (4%/4%)

PRESENT VALUE METRIC (TIME 10)	FIXED BENEFIT	PURE VARIABLE WITH 5% HURDLE RATE	VARIABLE WITH 0%/10% FLOOR/CAP	VARIABLE WITH -5%/15% FLOOR/CAP
Median	\$680.83 (3.92% implied discount rate)	\$613.91 (5.00% implied net discount rate)	\$623.28 (4.84% implied net discount rate)	\$613.91 (5.00% implied net discount rate)
Mean	\$685.63	\$613.91	\$626.77	\$614.36[21]

Replacing the 4%/4% assumptions for return and standard deviation with 4%/8% (to reflect longer term bonds) changes the median present value for plans with a collar, as illustrated in Table 5b.[22]

Table 5b – Present Value Measures Using a Fixed Income Portfolio (4%/8%)

PRESENT VALUE METRIC (TIME 10)	FIXED BENEFIT	PURE VARIABLE WITH 5% HURDLE RATE	VARIABLE WITH 0%/10% FLOOR/CAP	VARIABLE WITH -5%/15% FLOOR/CAP
Median	\$696.31 (3.69% implied discount rate)	\$613.91 (5.00% implied net discount rate)	\$650.77 (4.39% implied net discount rate)	\$619.21 (4.91% implied net discount rate)
Mean	\$716.49	\$613.91	\$658.85	\$625.93

[20] Expected returns below the hurdle rate implies benefits will decline over time.
[21] Does not match other values of \$613.91 since it captures tail returns. Table 5b has the same result, plus the median return is also impacted.
[22] Typically, longer-term bonds are also expected to have an expected return that exceeds that of shorter-term bonds. For illustration purposes, we assume the expected return does not change to isolate the effect of increased portfolio volatility. This could also be considered as illustrating a similar termed bond portfolio evaluated under circumstances leading to greater expected return volatility.

Reflecting a higher standard deviation for longer-term bonds does matter.

While this approach to estimating the consequence of investing in fixed income investments is straightforward, we have a couple of concerns.

One is whether an all-bond portfolio would have a lognormal distribution of returns like an open-ended portfolio (such as a diversified or equity-only portfolio). Arguably, individual fixed-income securities held to maturity could be valued at their yield-to-worst (with an allowance for default risk) at the time of purchase and held to maturity without caring about volatility at all. In times of relatively low interest rates, such cash-flow-matching may be perceived by plan sponsors as prohibitively expensive. At the time this paper is published, interest rates have increased significantly from their levels over most of the last decade, which has made adopting a cash-flow-matching approach more affordable given the higher expected returns that are attainable in the market.

The second concern relates to the concept of autocorrelation: the relationship between a return today compared to yesterday. Most scholarly research on autocorrelation focuses on equities and over short time intervals (daily or weekly). Here we are interested in the relationship between bond returns between years, particularly when driven by shifts in the yield curve. Our approach to considering autocorrelation is not focused on market observation but simply our Monte Carlo results as discussed in the next paragraph.

If our 4%/8% assumption scenario produces a first year (single trial) return of 12% (for example, based on a drop in yields, although there can be other reasons), the return in the second year shown in Tables 5a and 5b above are still tied to the long-term 4% assumption of the average return over the duration of the fixed income instrument. While autocorrelation might not materially matter for an equity-only or balanced portfolio, given that those autocorrelations seem to be weak and historically have varied between time periods, a bond return of 12% compared to a long-term yield-to-maturity of 4% means that the returns related to that bond instrument have been pulled forward in time and cannot be recognized again at a future period.[23] It would be logical to assume an expected future return in the second year that is adjusted to reflect the excess return already experienced in the first year. For example, assuming the bonds have an approximate duration of 15, the expected return in the second year is closer to 4.0% - (12%-4%)/15 = 3.47%. The 3.47% is the revised mean which is next adjusted by the percentile return for year two.

Let's apply this concept to our Monte Carlo model. Assume the random numbers for years one and two under the first of our 10,000 trials are 0.84 (84th percentile which is about an 11.92% return) and 0.40, respectively. The returns are shown in the chart below. For a given trial, autocorrelation could make the difference between a 1.70% return and a 1.17% return in year two.[24]

[23] This notion has been tested increasingly in the 21st century with the persistence of negative-yielding sovereign debt, most notably in Japan and Germany. Investment managers in the US would presumably face significant fiduciary pressure if they purchased securities that promised an absolute loss. As a simplification, our model has assumed that interest rates on fixed income securities will not drop below 0%.

[24] We have assumed for these purposes that the long-term volatility will remain constant for purposes of the distribution. If volatility were to compress around the lower expected return in this case, the difference because of autocorrelation would be even greater.

Table 6 – Impact of Autocorrelation on Year 2 Return

	YEAR 1 RETURN, 84TH PERCENTILE	YEAR 2 RETURN ASSUMING 4.00% AVERAGE AND 40TH PERCENTILE	YEAR 2 RETURN ASSUMING 3.47% AVERAGE AND 40TH PERCENTILE
Total Return	11.92%	1.70%	1.17%

When we revised our Table 5b model to reflect this autocorrelation adjustment using a factor of 15 to approximate the duration and a standard deviation of 8%, we have the following result:

Table 7 – Results Using Original and Revised Model with Autocorrelation Adjustment

PRESENT VALUE OF BENEFIT AT TIME 10	FIXED BENEFIT	PURE VARIABLE WITH 5% HURDLE RATE	VARIABLE WITH 0%/10% FLOOR/CAP	VARIABLE WITH 5%/15% FLOOR/CAP
Original – Median	$696.31	$613.91	$650.77	$619.21
Original – Mean	$716.49	$613.91	$658.85	$625.93
Revised – Median	$697.37	$613.91	$653.23	$620.54
Revised – Mean	$707.23	$613.91	$657.26	$626.29

While the difference between these values might not be characterized as material, they reflect an important underlying assumption of the models used to generate the economic scenarios and determining the discount rate. Considering ASOP No. 56, *Modeling*, requires an actuary to understand any known weaknesses in assumptions, methods, or other limitations of a model "that have material implications," actuaries could consider the role that autocorrelation may play in valuing a LDROM for a variable annuity plan.[25]

We also looked at partially tying the random numbers after the first year to the prior year's random numbers. However, this requires an assumption regarding what is a reasonable amount of correlation. Certainly, it is not 100% or the percentile return would be the same for all years under a given trial. However, when we use 20% we see the median and mean values of $653.23/$657.26 in Table 7 change to $652.00/$655.92 for a plan with 0% floor and 10% cap. [26]

[25] See Section 3.2b of ASOP No. 56, Modeling.

[26] In the base trial run, the random number (percentile of return) was 0.3764 for Trial 1 year 1 and 0.5127 for Trial 1 year 2. With a 20% correlation, the random number was 0.4854 for trial 1 year 2 = 20% x 0.3764 + 80% x 0.5127.

IV. VALUING BENEFITS BASED ON ASSUMED AVERAGE INCREASE WITHOUT MISMATCHING DISCOUNT RATE

For fixed benefits we use median (long-term geometric) discount rates when we want to value the long-term median expected value of the liabilities and the higher arithmetic mean return when we want to have no expected gains or losses. Given the competing ideals of precision vs lack of actuarial bias, neither method is unimpeachable. For pure variable annuities we can simply use the hurdle rate as the discount rate, which has the same mathematical consequence as using any interest rate due to the self-correcting nature of benefit adjustments relative to the hurdle rate.

Adding floors or caps to a variable annuity benefit adds complexity to the valuation process compared to a traditional fixed benefit or a pure variable annuity benefit. Not only are there arithmetic mean and geometric median returns that can be used as a discount rate, there are average and median expected benefit increases. Mismatching these two creates a tension in the valuation which could, in some cases (median discount rate and mean (average) benefit increases), be viewed as a margin of conservatism while in others (mean (average) discount rates and median benefit increases) cause systematic underfunding.

We can start by illustrating the baseline case with a pure variable annuity and attempting to value the liabilities using the gross expected return (with corresponding benefit adjustments) rather than the hurdle rate. For this purpose, we use a 7% arithmetic mean return assumption, 12% standard deviation, a $1,000 starting benefit, a 5% hurdle rate and a 10-year Monte Carlo model.[27]

After ten years, a pure variable annuity benefit will have a mean benefit value of $1,206.75 and a median value of $1,133.27. These imply annual average net increases of 1.897% and 1.259% respectively.[28] The median present value at time 0 of the benefit due at time 10 is $613.91 for all 10,000 trials when discounting at the rates produced in each trial.

---

[27] Throughout this Section, we use additional decimal places in presenting the interest rate results from our Monte Carlo analysis to better illustrate some of the mathematical results where the formulas would not tie as closely to actual results if we used the typical four decimal place standard as in other sections of the paper.

[28] Valuing a benefit at a 5% discount rate is the same as assuming a 1.9% COLA and using 7% as a discount rate (1/1.05 = 1.019/1.07 = 0.9523 since 1.9% is based on 1.07/1.05-1).

Now, think of $613.91 as the funding cost or liability at time zero. Table 8 below shows what $613.91 will accumulate to using an assumed net annual increase and either an interest rate of 7.000% (mean return) or 6.322% (the median 10-year return for the 10,000 trials in our original model) or the liability single payment:

Table 8 – Illustration of Two Different Projected Variable Benefits and Funding Mismatch

MONTE CARLO MEASURE	BENEFIT IN YEAR 10	GROWING PRESENT VALUE	MISMATCH BETWEEN RETURN ASSUMPTION USED FOR INCREASE AND DISCOUNT RATE
Median	$\$1{,}000 \times 1.01259^{10}$ = $1,133.27	$\$613.91 \times 1.06322^{10}$ = $1,133.27	$\$1{,}000 \times \left(\frac{1.01259}{1.07}\right)^{10}$ = $576.10 < $613.91
Mean	$\$1{,}000 \times 1.01897^{10}$ = $1,206.75	$\$613.91 \times 1.07^{10}$ = $1,206.75	$\$1{,}000 \times \left(\frac{1.01897}{1.06322}\right)^{10}$ = $653.71 > $613.91

Note that for both the median and the mean, the Growing Present Value matches the projected benefit in year 10. The time 0 liabilities in the last column are either overstated (if you use a mean increase of 1.897% and a discount rate of 6.322%) or understated (if you use a mean increase of 1.259% and a discount rate of 7.000%).

If this were a fixed benefit of $1,000 in year ten, we expect to simply have a choice as to which rate to discount at (7.000% or 6.322%) and unlike a pure variable annuity have two different present values. Since we have no changes in the fixed benefit, the issue of being consistent between the net benefit change and the discount rate does not arise, except to the extent that it might bias the actuarial valuation towards gains or losses.

If we have a variable annuity benefit with a floor and cap, we should avoid valuing the benefit using mismatched assumptions. Consider a variable annuity benefit with a floor and a cap of 0%/10%. If we use a mean average net increase assumption (0.425%) with a median expected return (6.322%) or a median net increase assumption (0.338%) with the mean expected return (7.000%) the result is a present value that is less than the respective median/mean present value from our series of 10,000 trials.

When we look at mean values, we calculate the harmonic average (e.g., the $576.66 value in Table 9) which we will discuss later in Section 5. While it may seem obvious that for a pure variable annuity plan we avoid this issue by valuing a fixed benefit at the hurdle rate, we use this illustration to set up a key conclusion regarding variable annuity benefits that have floor and cap features.

For variable annuity plans with floors and caps, it is important to avoid a mismatch when selecting the benefit increase and discount rate assumptions.

We should not assume a median net benefit increase (e.g., 0.338%) if we are using the arithmetic mean return (7.00%) as our discount rate. It may be easier to focus on the median values than the mean. A 0.338% median average net increase and a median discount rate of 6.287% work well.

Table 9 – Illustration of Mismatch Present Value Calculations (using 0%/10% Floor/Cap)

EXPECTED RETURN METRIC	BENEFIT AT TIME 10 (FROM 10,000 TRIALS)	PRESENT VALUE AT TIME 0 (FROM 10,000 TRIALS)	MISMATCH PRESENT VALUE
Median	\$1,034.27	\$562.11	$\$1{,}000 \times \left(\frac{1.00338}{1.07}\right)^{10}$ = \$525.80 < \$562.11
Arithmetic Mean	\$1,043.30	\$576.66	$\$1{,}000 \times \left(\frac{1.00425}{1.06322}\right)^{10}$ = \$565.18 < \$576.66

Like Table 8 above, we can check our math by calculating whether funding these benefits at these present value levels is expected to accumulate to cover the expected benefit amount:

Table 10 – Growing Present Value with Mismatched Assumptions

GROWING PRESENT VALUE	MISMATCH BETWEEN RETURN ASSUMPTION USED FOR INCREASE AND DISCOUNT RATE
$\$562.11 \times 1.06322^{10} = \$1{,}037.65$	$\$562.11 \times 1.07^{10} = \$1{,}105.76$

V. OTHER ISSUES AND CONSIDERATIONS

Changes in Capital Market Assumptions

We have established that funding (expected return) liabilities for pure variable annuities are not impacted by changes in expected return while traditional plans are impacted (i.e., the lower the discount rate, the higher the liability). As expected, a variable annuity with a cap/floor will likely be impacted in some reduced fashion compared to a traditional fixed annuity.

This chart shows the impact of changing our 7.00% arithmetic mean return assumption to 6.50% (keeping 12% standard deviation) on the present value of a benefit payable in ten years based on three different designs:

Table 11 - Effect of Changing Mean Return Assumption from 7.00% to 6.50%

	Fixed $1,000 Benefit	Pure Variable $1,000 Benefit	Variable $1,000 With 0%/10% Floor/Cap
Change in liability[29]	4.86% $\left(\frac{\$568.03}{\$541.72} - 1\right)$	0.00% $\left(\frac{\$613.91}{\$613.91} - 1\right)$	3.11% $\left(\frac{\$579.60}{\$562.11} - 1\right)$

Table 11 shows that expected returns still matter if the plan is not a pure variable annuity plan. If the benefit is valued instead using a liability load, that load may need to change as the return assumption changes.

Recovering from Losses Below the Floor

Because changing capital market assumptions impact the calculation of liabilities and funding of variable annuity benefits, there is some concern about the ability to recover from actuarial losses borne by the plan from years when returns are below the floor. Variable annuity plans are sometimes viewed as a panacea and the cure for underfunded plans. Floors on downward adjustments reintroduce the possibility of unfunded liabilities, with some downside protection to the plan compared to a fixed annuity benefit. Using a hurdle rate below the level of expected returns will reduce the losses absorbed by the plan. The level of the floor, and the spread between the floor and the hurdle rate, will also affect the likelihood of the floor being pierced. Standard deviations consistent with diversified investment portfolios are typically large enough that we expect that floors will be pierced in some years.

[29] Present values have been calculated here using the median geometric return as the discount rate, except for the pure variable benefit, which keeps the 5% hurdle rate as its discount rate.

When a floor is used, an actuary should consider the following concerns:

- **After a return below the floor occurs, it may take years before gains above the cap cover prior losses.[30]**
- **Modeling may be difficult if the plan is headed toward insolvency. Anytime a plan determines benefits on percentile returns and assets are less than liabilities the plan design could create a special risk issue. A plan moving toward insolvency has this issue and the degree of the mismatch (percent benefit increases without asset backed dollar returns) increases over time.**
- **Losses might not be offset by earlier gains if those gains were allocated for other purposes.**
- **Some designs might try to hold back excess gains which could lead to some generational transfers of risk.**

When valuing tail risks, it may be valuable to consider the likelihood of going below a floor or above a cap. If we use our prior assumptions (7% arithmetic mean return, 12% standard deviation) and assume a floor of 0% and a cap of 10% we get the following results:

Table 12 – Tail Risks

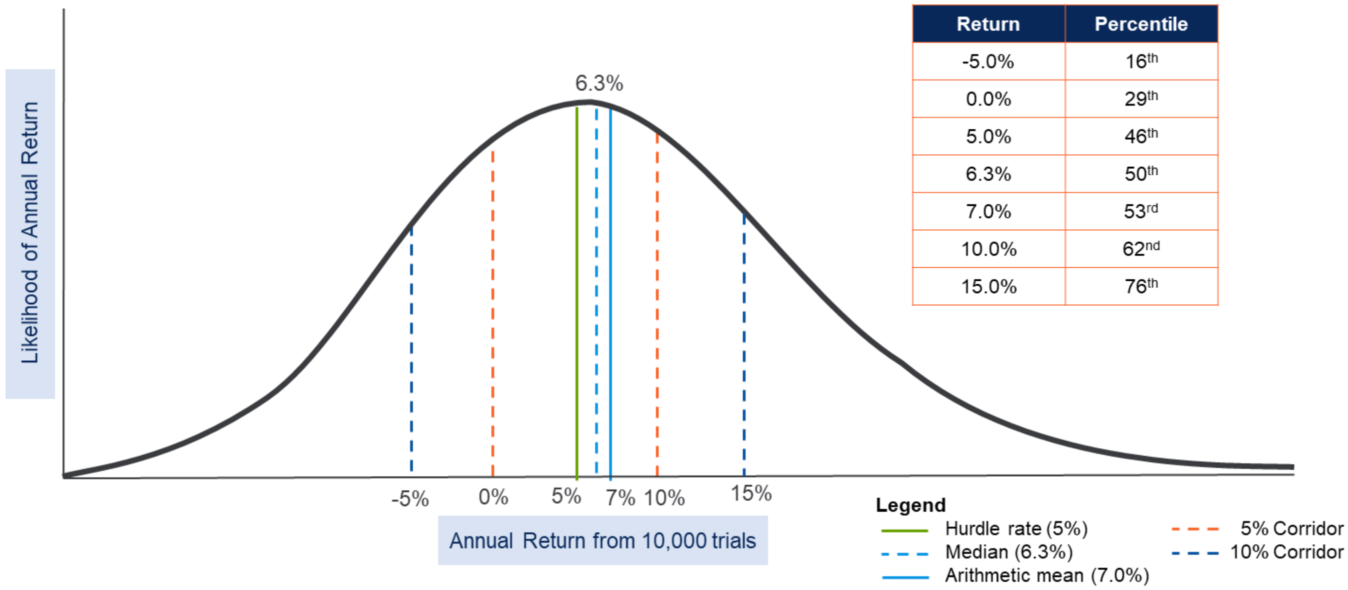

Return	Percentile
-5.0%	16th
0.0%	29th
5.0%	46th
6.3%	50th
7.0%	53rd
10.0%	62nd
15.0%	76th

Probability of going under 0%	29%
Probability of going over 10%	38%

[30] Here is another place where an actuary needs to be mindful of the tail returns to the downside implicit in their chosen ESG. A floor on adjustments is debilitating when the possibility of extreme negative returns, even spread over multiple years, is considered. As a thought exercise, consider a variable annuity plan with a 5% hurdle rate, a 0% floor and a 10% cap that earns the S&P 500 total return every year and started on January 1, 1929. This plan would have been underfunded every year until the end of 1936 and would have had a funding percentage nadir of 43% funded after 1932. Such plan would not begin to be consistently fully funded until after 1942. It can take a very long time to recover from bad enough returns.

In the above graph, the height represents relative frequency and the percentile shown in the box is the percent of the total area under the graph which is to the left of the point for the return. The median return is just to the right of the peak.

About a third of the time returns will be between 0% and 10%. Even if returns are below 0%, the losses to the plan might not be significant. To lose 5% or more below the floor has about a 16% (1/6) chance and to earn 5% or more above the cap (15% gross) has about a 24% chance in our model. While the odds are favorable (gains more likely than losses when the hurdle rate is less than expected return), it still could take several years (expected to be about four years if there are no additional intervening returns below the floor) to recover a one-out-of-six-year loss.

Ultimately, floors do create the possibility of asymmetrical downside risk for a plan's funding that might take many years to fully recover from, or even be unrecoverable in the event of a total breakdown in capital markets, as has historically been seen in non-US markets. [31] **Footnote 30 provides an illustration of a particularly brutal sequence of historical US returns and what effect a floor would have had on hypothetical plan funding.**

Harmonic Averages

We have focused this paper on arithmetic and geometric returns, rather than the harmonic average of returns. However, all three of these averages are available as outputs of our Monte Carlo model and were analyzed.

Arithmetic returns are basic. We used 7.00% as an input in our model, and when we do a simple average of 10,000 trials over 10 years, we get an average of these 10,000 numbers that is also 7.00%. Converting a sequence of arithmetic returns determines the "terminal wealth" and captures the tail returns in a skewed lognormal distribution.

Geometric returns are derived from the median benefits and present values by assuming a lognormal distribution of returns. When geometric returns are determined, they are based on median values and taking the product of 1+i for n years and to the 1/n power. We do not average values of v, or (1/(1+i)).

Harmonic mean calculations involve two inverse calculations. For example, v is the inverse of 1+i. In our model, we sometimes calculate a mean benefit and discount it back to the present value. Ignoring the benefit amount, we effectively have an average of 10,000 v^{10} values (first inversion). We invert this average to get $(1+i)^{10}$ values and can easily get an average value for i from this. However, this is a harmonic mean and is expected to be lower than either the arithmetic or geometric mean values.

While we sometimes found it useful to take the projected mean benefit and calculate an average (harmonic) present value, we have avoided taking the next step and calculating a harmonic mean interest rate. Generally, we would expect the mean benefit to be discounted back at the arithmetic rate (or optionally for a fixed benefit at the geometric rate).

[31] A common critique of using US historical returns is that there is an implicit survivorship bias to using equity markets that have remained open and thriving for more than a century, while other developed economies such as Germany, Japan, Brazil, and, most recently, Russia, have suffered catastrophic equity losses that have been unrecoverable. While this paper is focused on US pension plans, the authors are aware of this critique and have attempted to address it in our focus on asymmetrical downside risk here in Section 5. The context of this paper focuses generally on a range of normal economic conditions –extreme tail risk events of a political or non-economic nature are outside the scope of this paper.

We note that in Table 3 the harmonic present value for a design with a -5%/15% floor/cap was less than the median PV (harmonic $761.27 < median $768.85) which is not common. This occurs because extreme outliers do not impact the median but do impact the average. For example, if we make the caps +/- 20% the median will match the present value for a pure variable annuity however the mean will be different and factor in extreme tail events. The normal relationship is based on a lognormal distribution which we no longer have when looking at terminal wealth in a modified variable annuity plan design.

CONCLUSION

Our goal with this paper has been to highlight some of the considerations actuaries face as they value and design modified variable annuity plans. There are many other considerations that are worthy subjects for further study, ranging from the philosophical debate about the optimal use of plan reserves generated by a cap to the ideal asset allocation for investments in a plan with variable benefits. As these plans become more common, unique plan provisions will create new and different challenges for actuaries.

Ellen Kleinstuber is the Chief Actuary at Bolton, leading the firm's single employer sector and providing technical and consulting support to all sectors of our actuarial consulting practice.

Scott Steadman is a Consulting Actuary at Bolton. Scott focuses on ERISA multiemployer pension plans and has presented on issues facing variable annuity pension plans.

Tom Lowman retired from Bolton in 2023. Tom had 45 years of pension actuarial experience, the last 26 years of which were at Bolton.

APPENDIX A: GLOSSARY

The following definitions are used throughout this paper.

ASOP	Actuarial Standard of Practice
Mean	The average of possible values for a random variable weighted by the probability associated with each value.
Arithmetic Mean (Average)	The sum of a series of numbers divided by the number of items in the series, also referred to as the average value of the data set. The future-looking arithmetic mean return is typically the assumption that will produce no expected gains or losses.
Geometric Mean (Average)	The result of taking the Nth root of the product of N single-period values representing the amount that would be accumulated during the period from an investment of $1. The geometric mean will always be less than the arithmetic mean over multiple periods.
Harmonic Mean (Average)	Determined as the reciprocal of the arithmetic means of a given set of observations which are "reciprocal" (1/(1+i)) present values. Generally, the harmonic mean interest rate will be smaller than the arithmetic and geometric mean values.
Median	The value that separates the upper 50% from the lower 50% of the distribution of outcomes for a random variable. Over a long time period, the median and geometric mean will converge.
Terminal Wealth	The amount that accumulates from an initial investment of $1. For any value of terminal wealth at the end of N periods, the equivalent discount rate is determined by taking the Nth root of terminal wealth and subtracting 1.
Economic Scenario Generator (ESG)	A mathematical model that simulates possible future paths of economic and financial market variables.

APPENDIX B: MONTE CARLO MODELING ASSUMPTIONS AND METHODS

The model used for this paper projects benefits and present values using a Monte Carlo model created using Microsoft Excel. Projected investment returns are calculated using a log normal distribution of investment returns. The initial benefit is $1,000.

The focus of the calculation is based on a single payment in year ten. This is similar to the Difficult-to-Value paper which focused on a payment after five years. That paper discussed when there may need to be an analysis other than one focused on a single payment.

The model assumes that bond interest rates have a floor at 0%.

Initial Capital Market ("CAPM") assumptions used in the model for this paper are:

- Expected Return (arithmetic): 7%
- Standard Deviation: 12%

These are used for Tables 1a, 3 and 4.

The study notes when other rates are used, e.g. Table 1b replaces the assumptions of 7%/12% with 6%/10%.

Section 3 uses lower assumed values for expected return and standard deviation. Tables 5b and 7 replace the assumption of 7%/12% with 4%/8% in the initial year.

Two other assumptions are introduced in Section 3 related to autocorrelation:

1. Table 7 also adjusts the following year's return for each trial by an amount that equals (the prior year's expected return less the actual prior year's trial return)/15.
2. Near the end of section 3, we show a second separate type of autocorrelation based on adjusting the random number (a value between 0 and 1, representing percentile results). We adjusted the random number to weight a percent of the prior year's value. We used a 20% assumption, which changed the median and mean values of $653.23/$657.26 in Table 7 to $652.00/$655.92 for a plan with 0% floor and 10% cap. As noted in the report, there is little information to know what this assumption should be. However, it appears it may not be material.

Bolton
1 W. Pennsylvania Ave.
Suite 600
Towson, MD 21204

410.547.0500
www.boltonusa.com
solutions@boltonusa.com