



Aging and
Retirement

Benefit at Risk in Lifetime Pension Pools

FEBRUARY 2023



SOA
Research
INSTITUTE



**Canadian
Institute
of Actuaries**

**Institut
canadien
des actuaires**

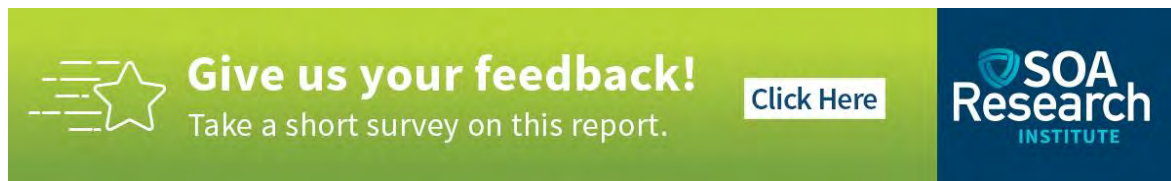
Benefit at Risk in Lifetime Pension Pools



Authors Jean-François Bégin, FSA, FCIA, PhD
Assistant Professor
Simon Fraser University

Barbara Sanders, FSA, FCIA, MSc
Associate Professor
Simon Fraser University

Sponsors Retirement Section Research
Committee of the Society of
Actuaries Research Institute

Canadian Institute of Actuaries



 **Give us your feedback!**
Take a short survey on this report. [Click Here](#) 

Caveat and Disclaimer

The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the Society of Actuaries Research Institute, the Society of Actuaries, the Canadian Institute of Actuaries, or their members. The Society of Actuaries Research Institute and the Canadian Institute of Actuaries make no representation or warranty as to the accuracy of the information.

CONTENTS

Executive Summary	4
Section 1: Introduction, Literature Review, and Scope	5
Section 2: An Introduction to Benefit at Risk	8
2.1 Value at Risk	8
2.2 Benefit at Risk	9
2.3 BaR as a Budgeting Tool	10
2.4 BaR as a Decision-Making Tool	12
2.5 Limitations of BaR	13
Section 3: Lifetime Pension Pools	15
3.1 Simple Framework for Varying Annuity Payouts	15
3.2 Example	17
3.3 Adjustment Factor for Systematic Mortality Improvements	18
3.4 Discussion on Mortality Pooling	19
Section 4: Benefit at Risk in Lifetime Pension Pools	20
4.1 Analytical Exploration	20
4.1.1 Minimum Benefit at Risk	21
4.1.2 Average Benefit at Risk	22
4.2 Sensitivity Testing and Robustness Analyses	23
4.2.1 Hurdle Rate	23
4.2.2 Asset Allocation	25
4.2.3 Combined Effects of the Hurdle Rate and the Asset Allocation	25
4.3 Nominal versus Real Terms	27
4.4 Bootstrapping Backtesting	29
4.5 Impact of Idiosyncratic Mortality	30
Section 5: Concluding Remarks and Further Developments	34
Acknowledgments	35
Appendix A: Proofs	36
A.1 Comparator Minus Minimum Benefit Distribution	36
A.2 Comparator Minus Average Benefit Distribution	36
Appendix B: Exogenous Comparators	38
B.1 Redefinition of Minimum and Average Benefit at Risk with Exogenous Comparators	38
B.1.1 Minimum Benefit at Risk with Exogenous Comparator	38
B.1.2 Average Benefit at Risk with Exogenous Comparator	38
B.2 Examples with Exogenous Comparators	38
B.2.1 Hurdle Rate	38
B.2.2 Asset Allocation	39
Appendix C: Block Bootstrapping	42
References	43
About the Society of Actuaries Research Institute	46
About the Canadian Institute of Actuaries	47

Benefit at Risk in Lifetime Pension Pools

Executive Summary

Several modern retirement arrangements, including lifetime pension pools, allow retirees to convert a single premium into income for life that varies with investment and mortality experience. To assess how much risk members bear in these arrangements, this report introduces a collection of new risk measures—called benefit at risk, or BaR for short—to be used in the context of varying benefits.

Following the rich literature on risk measures in risk management and actuarial science, we specifically consider quantile-related measures, which identify the level above which the benefit is likely to stay with a given confidence level and over a specified time horizon. These measures are very similar in nature to the value at risk widely used in banking and insurance.

The focal point of the present report is member communication and disclosure: we wish to create meaningful measures for members to understand the benefit risk borne in these pools. We focus on two different uses of such a measure: budgeting and decision making.

- Budgeting tool: For existing participants who must plan for (and possibly adjust) consumption based on the income to be provided by the pool.
- Decision-making tool: For prospective participants who must weigh the risks and rewards of voluntarily allocating retirement assets to the pool compared to other available decumulation options.

These two applications drive the creation of two original benefit at risk–type measures:

1. The minimum BaR: In the short term, we expect members to construct budgets based on the benefits they are currently receiving and focus on their ability to meet those budgets, even in the worst year. For this purpose, we propose to use a short-term benefit at risk measure based on the minimum benefit realized over this period. Using the minimum allows members to be confident that they will likely be able to pay for all their planned expenses, even when benefits are low.
2. The average BaR: From a decision-making perspective, members could be interested in comparing different arrangements before their retirement (i.e., before buying a retirement product or entering a pension arrangement). The benefit at risk could be useful from this perspective as this measure could inform members of the risk associated with each potential deal if the horizon selected is long enough.

The new measures are applied to a stylized lifetime pension pool in the spirit of Piggott et al. (2005) to understand how they work in this specific context.¹ We apply the measures to stylized pools and assess their behaviour when inputs used in their calculation are changed (e.g., horizon, hurdle rate, asset allocation strategy, risky asset return distribution, idiosyncratic mortality). We find that the measures are robust to changes in most assumptions, including more realistic risky asset return distributions and the inclusion of idiosyncratic mortality risk.

We conclude this report by discussing further developments and applications of the benefit at risk measure.

¹ A stylized pool refers to a pool represented in a way that simplifies details rather than trying to show reality. These are selected for simplicity's sake.

Section 1: Introduction, Literature Review, and Scope

As the prevalence of guaranteed pension arrangements decreases worldwide, flexible schemes such as lifetime pension pools are expected to become more popular. Lifetime pension pools allow retiring individuals to convert a lump sum into income for life. The pool does not guarantee a specific level of income; instead, the pension payable varies with the investment and mortality experience of the group.

Pooling mortality risk has become a priority in recent years in the defined contribution (DC) context. Accordingly, the Organization for Economic Cooperation and Development recently released recommendations for good design of DC plans; one of the recommendations is to ensure protection against longevity risk in retirement, and lifetime pension pools certainly address this suggestion.² By pooling individual participants' mortality risks, lifetime pension pools can generate higher income than an individual systematic withdrawal plan (see, e.g., Australian Government Actuary, 2014). Lifetime pension pools are also expected to outperform retail annuities as there is no requirement for risk capital to support expensive guarantees.

Various arrangements and products fit the broad description of lifetime pension pools in the literature: group self-annuitization plans (Piggott et al., 2005; Valdez et al., 2006; Qiao and Sherris, 2013; Hanewald et al., 2013), pooled annuity funds (Stamos, 2008; Sabin, 2010; Donnelly et al., 2013), annuity overlay funds (Donnelly et al., 2014; Donnelly, 2015), retirement tontines (Milevsky and Salisbury, 2015, 2016; Sabin and Forman, 2016; Fullmer and Sabin, 2019; Fullmer, 2019; Iwry et al., 2020; Chen et al., 2021), assurance funds (Fullmer and Forman, 2022), variable payout annuities (Horneff et al., 2010; Boyle et al., 2015), and variable payment life annuities (ACPM, 2017).³ Note that all these designs can be viewed as (implicit or explicit) tontines.⁴

Working examples of these lifetime pension pools include the College Retirement Equities Fund operated by the Teachers Insurance and Annuity Association of America (TIAA) in the US since 1952, the Variable Payment Life Annuities run by the Faculty Pension Plan of the University of British Columbia (UBC) since 1967, and the Lifetime Pension introduced to the Australian market by QSuper in 2021.⁵

We expect to see more of these lifetime pension pools in the years to come. In Canada, recent changes to income tax regulations that accompanied the Budget Implementation Act of 2021 will facilitate this trend. In the US, the bulk of these pension pools are currently not permitted, and Congress would need to craft legislation to allow them; there is some interest, nonetheless, among pension practitioners (see, e.g., Shemtob, 2021, 2022).

Given that the income provided through a lifetime pension pool is expected to vary as a function of the investment and mortality experience, it would be helpful to characterize the associated benefit risk. Following the rich literature on risk measures in risk management and actuarial science, this report considers a quantile-related measure, which identifies the level above which the benefit is likely to stay with a given confidence level and over a specified time horizon. This measure—very similar in nature to the value at risk, or VaR—is called benefit at risk (BaR). In lay terms, the BaR measure informs members of the extent of possible benefit losses with respect to a basis of comparison and over a specific time frame.

² "DC pension plans should provide some level of lifetime income as a default for the pay-out phase, unless other pension arrangements already provide for sufficient lifetime pension payments. Lifetime income can be provided by annuities with guaranteed payments or by non-guaranteed arrangements where longevity risk is pooled among participants." (OECD, 2022)

³ Academic research on lifetime pension pools and tontine-like designs has been very fruitful in recent years. Many new products have been introduced; for instance, combinations of annuities and tontines (Chen and Rach, 2019; Chen et al., 2020) such as tonuities (Chen et al., 2019), tontines with bequest (Bernhardt and Donnelly, 2019), and unit-linked tontines (Chen et al., 2022), among others.

⁴ Implicit tontines promise to pay the participants an income for life, but longevity credits are not explicitly allocated to the participants; explicit tontines explicitly allocate longevity credits to the individual accounts of participants (see Bernhardt and Donnelly, 2019, for more details).

⁵ See Forman and Sabin (2015), Milevsky (2015), and CREF (2022a,b) for more information on TIAA's lifetime pool.

Instead of investigating profits and losses like VaR, BaR focuses on benefits—or possible loss thereof. The notion of benefits is very broad; one could be interested in benefits received at a specific point in time, whereas other members could focus on the minimum benefit or the average benefit over a given period.

Another main difference of BaR when compared to the classic VaR measure is that we contrast the members' benefits to some other quantity or basis of comparison called the *comparator*; the comparator could be the current level of benefits, the future expected benefits, the total lifetime benefits, or the benefits from an alternative retirement option, for instance.

The benefit at risk is a valuable new metric for disclosure and communication. Specifically, it could help as a:

- Budgeting tool: For existing participants who must plan for (and possibly adjust) consumption based on the income to be provided by the pool.
- Decision-making tool: For prospective participants who must weigh the risks and rewards of voluntarily allocating retirement assets to the pool compared to other available decumulation options.

In the short term, we expect members to create budgets based on the worst year. For this purpose, we propose to use a benefit at risk measure based on the minimum benefit likely to be realized over a relatively short period; this specific measure is called the minimum BaR (mBaR) in this report. Using the minimum allows members to be confident that they will likely be able to pay for all their budgeted expenses, even when benefits are low.

From a decision-making perspective, members could be interested in comparing different arrangements before their retirement (i.e., before buying a retirement product or becoming a member of a pension arrangement). The benefit at risk could be useful from this perspective as this measure could inform members of the risk associated with each potential deal if the horizon selected is long enough. The short-term budgeting mentioned above focuses on the worst outcome over a given horizon; however, focusing on the worst year makes less sense over longer-term horizons as bad years might be offset by better ones. Over the long term, budgets will need to be adjusted to reflect emerging benefits and needs. Because of this, we use the average benefit at risk (aBaR) when using BaR as a decision-making tool.

These two new measures, mBaR and aBaR, are applied to a stylized lifetime pension pool to understand how they work in this specific context. We apply the measures and assess their behaviour when inputs used in their calculation are changed (e.g., horizon, hurdle rate, asset allocation strategy, risky asset return distribution, and idiosyncratic mortality).

The contribution of this report is the introduction of a collection of new measures to be used in the context of varying benefits—and more specifically, in the context of lifetime pension pools. As mentioned above, our focus is member communication and disclosure: we wish to create meaningful measures for members to understand the benefit risk borne in these arrangements. We believe the two specific measures introduced in this report fulfill this aim.

One could be interested in using BaR in other contexts than communication and disclosure (e.g., pool design). We warn users that solely relying on BaR might be unsatisfactory for such other applications as it does not tell the whole story. For instance, the benefit at risk does not provide information on the variability of the benefits from one year to the next. Indeed, BaR considers only the distribution of benefits across all possible scenarios at specific times and not how benefits might evolve as a function of time along a single scenario.

Despite the possible limitations of BaR, we believe that it is a great device when used appropriately—as a budgeting and decision-making tool for members, for instance. In addition, when combined with other metrics, BaR could be used by plan sponsors and providers as well as their actuaries when designing lifetime pension pools. In this report, however, we focus only on communication and discretionary disclosure applications of the measure.

This report is structured as follows. Section 2 presents the benefit at risk in all generality and proposes two specific applications of the new risk measure: the mBaR and the aBaR. Then, Section 3 introduces a stylized lifetime pension pool for illustrative purposes. Section 4 applies the risk measures to the lifetime pension pool introduced in Section 3 and assesses the robustness of mBaR and aBaR by changing various assumptions (i.e., horizon, hurdle rate, asset allocation strategy, risky asset returns, and mortality). Section 5 concludes and discusses avenues for future research.

Section 2: An Introduction to Benefit at Risk

This section introduces a new measure—the benefit at risk, or BaR—for effectively communicating risk to members. As mentioned in the introduction, many pension arrangements and products now exist without a hard guarantee, and it is paramount for members to understand the amount of risk they are bearing.

Inspired by the literature on risk measures, we first define the BaR measure in all generality; this definition is broader than the lifetime pension pool context we consider in the subsequent sections of this report. Indeed, the benefit at risk could be used by actuaries to communicate risk in virtually any arrangement for which benefits are not made certain.

The general BaR uses four inputs—the benefit statistic, the comparator, the horizon, and the probability level—and specific choices of these four inputs lead to different measures. Because the focus of this report is member communication, we propose two distinct BaRs: one for budgeting purposes and one for decision-making purposes. The former measure is used to help current members understand how risky their benefits are in the short term. The latter serves as a comparative tool for prospective members to contrast the pros and cons of different arrangements; it gives them information about benefit risk in the medium to long term.

2.1 VALUE AT RISK

Before introducing our new measure—the BaR—we first recall the rudiments of VaR as it serves as one of the main building blocks in the construction of the benefit at risk. VaR, put simply, is a probability-based measure of loss potential. It quantifies the maximum loss that an investor will not exceed with a certain probability and for a given horizon.

Let $X(\tau)$ be a continuous loss random variable (where losses are represented by positive numbers) over a given horizon that we denote by τ . The likelihood of observing various potential losses is represented by the density of $X(\tau)$. The cumulative distribution function or cdf $F_{X(\tau)}(x)$, on the other hand, denotes the probability that the loss does not exceed x :

$$F_{X(\tau)}(x) = \mathbb{P}(X(\tau) \leq x).$$

Mathematically, the quantile function (also known as the percentile function or the percent-point function) is the inverse of the cumulative distribution function. Given a probability level p , it identifies the maximum loss level x such that

$$\mathbb{P}(X(\tau) \leq x) = p.$$

The value at risk at probability level $p \in (0,1)$ is the p -quantile of $X(\tau)$, or

$$F_{X(\tau)}(\text{VaR}_p[X(\tau)]) = p \quad \Leftrightarrow \quad \text{VaR}_p[X(\tau)] = F_{X(\tau)}^{-1}(p).$$

For instance, a one-week 95% VaR of \$100 means that there is a 95% probability that the investor's losses will not exceed \$100 over the next week; this is equivalent to a 5% chance that the losses will actually surpass \$100.

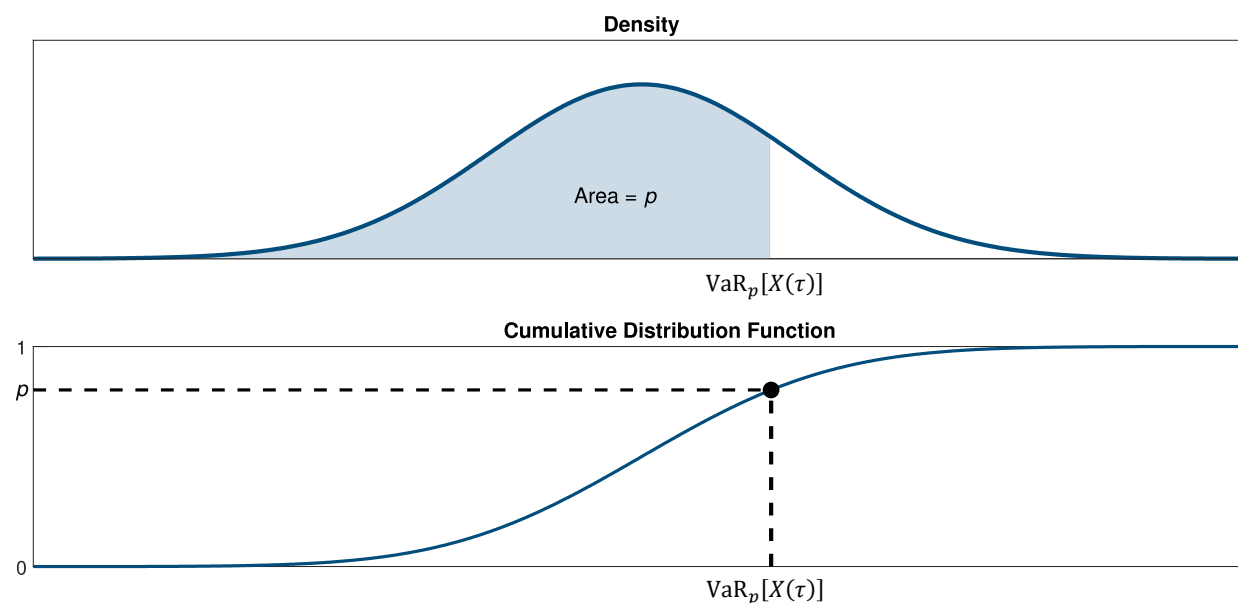
We illustrate the relationship between the density, the cumulative distribution function, and the value at risk in Figure 1. The top panel shows the likelihood of observing various potential losses. The area under this curve (in blue) represents the probability of observing a loss that is below a given quantity; that is, $\mathbb{P}(X(\tau) \leq x)$, which is also the value of the cdf of $X(\tau)$ at x , denoted by $F_{X(\tau)}(x)$ above and illustrated in the bottom panel of Figure 1. If we want to know the loss level x associated with an area p under the density curve, we can simply read off the x -coordinate of the point on the cumulative density function that has height p . This quantity is $\text{VaR}_p[X(\tau)]$ because

$$F_{X(\tau)}(\text{VaR}_p[X(\tau)]) = F_{X(\tau)}(F_{X(\tau)}^{-1}(p)) = p.$$

To summarize, the loss level associated with an area p under the density curve and to the left is $\text{VaR}_p[X(\tau)]$, and the cdf of $X(\tau)$ evaluated at $\text{VaR}_p[X(\tau)]$ is p . By inverting the cdf, one can then get the quantile function; that is, a function that takes a probability level as an input (i.e., p) and gives the loss level associated with this probability (i.e., $\text{VaR}_p[X(\tau)]$).

Figure 1.

RELATIONSHIP BETWEEN DENSITY, CUMULATIVE DISTRIBUTION FUNCTION, AND VALUE AT RISK.



This figure shows the relationship between the density (top panel), the cumulative distribution function $F_{X(\tau)}$ (bottom panel), and the value at risk. The density represents the likelihood of observing various losses; the area under this curve and to the left of a given point gives the value of the cdf at that point. The value at risk is obtained by evaluating the inverse of the cumulative distribution function—the quantile function—for a given probability level (i.e., p in this case).

During the 1990s, VaR became one of the most widely adopted ways to measure risk.⁶ It is still extensively used today in risk management, even though the measure has been controversial, especially because of misinterpretation and misuse in the past.⁷

2.2 BENEFIT AT RISK

Like VaR, BaR has a probability level and a horizon, denoted by p and τ , respectively. However, instead of investigating profits and losses, the focus is on benefits. The notion of benefits is very broad; one could be interested in benefits received at a specific point in time, whereas others could focus on the average benefit or the minimum benefit over a given period. The BaR end user makes this choice. We call this input the *benefit statistic*, and we denote the corresponding random variable by $\beta(\tau)$.

⁶ See Holton (2002) for the history of VaR, tracing back its origins to as early as 1922.

⁷ The value at risk has been under fire since it moved from trading desks into the public eye in 1994. For instance, VaR was the subject of a very public debate between Philippe Jorion and Nassim Taleb in 1997. See Jorion (1997) for more details.

Another key input to the BaR is the *comparator*, denoted here by C , which acts as a benchmark for the benefit level $\beta(\tau)$. The comparator could be the current level of benefits at the time the BaR is being calculated, the expected future level of benefits, or the total lifetime benefits, for instance.

In all generality, the BaR at probability level p for benefit statistic $\beta(\tau)$ and comparator C is defined as

$$\text{BaR}_p[C - \beta(\tau)] = F_{C-\beta(\tau)}^{-1}(p). \quad (1)$$

The expression $C - \beta(\tau)$ represents the amount by which the benefit $\beta(\tau)$ falls short of the comparator C . The BaR is the p -quantile of this shortfall; it measures the member's risk of benefit loss relative to the comparator.

For instance, if the benefit statistic is chosen to be the benefit received one year hence, and the comparator—say, the current benefit—is \$1,000, then saying that *the 95% BaR is \$100* means that the member will receive a benefit higher than \$900 next year with a 95% probability.

The following two subsections are dedicated to specific choices of benefit statistics and comparators that are relevant to members.

2.3 BAR AS A BUDGETING TOOL

Arrangements without a hard guarantee lead to the potential for benefit level changes. As members budget for their expenditures, they need to understand the extent to which benefits might be reduced.

We conceptualize budgeting in retirement as the process of planning discretionary expenditures given projected retirement income and fixed expenses. This exercise has a short-term focus due to the uncertainty in both the projected income and the expenses themselves. Specifically, as people age, their needs and desires can change significantly, and so do their spending patterns. Blanchett (2014) observes a “retirement spending smile” in aggregate data in the US, with spending in real terms being higher near retirement, then decreasing, and finally increasing again towards the end of life. In financial planning circles, this pattern is interpreted in the context of three distinct phases of retirement:⁸

- The “go-go” years when retirees have the freedom and ability to pursue many activities—this time is characterized by additional spending on travel, hobbies, and leisure.
- The “slow-go” years, characterized by a slowdown in activities and travel, translating into more modest discretionary spending.
- The “no-go” years when health-related expenses tend to rise.

Although these phases are quite pronounced in the aggregate data on spending, there are a wide range of experiences at the individual level, with the transition from one phase to the next generally being dictated by factors outside of retirees' control rather than being planned ahead of time. As a result, individual budgeting is best constrained to the short term (maximum five years), with necessary adjustments made as circumstances change.⁹

The BaR can be a useful metric for retirees engaged in such a short-term budgeting exercise. Specifically, as they plan for the next five years, retirees may wish to line up their expenditures with projected income in an adverse (rather

⁸ These phases are aligned with how the US healthcare system is operated. Canadian retirees also go through similar phases, with health-related expenses potentially rising at advanced ages due to long-term care needs that are not covered under Canada's universal publicly funded healthcare system.

⁹ In practice, some lifetime pension pools and risk-shared pension plans involve smoothing and administrative lags that defer recognition of experience gains and losses. A 5-year horizon is long enough to reflect the consequences of updating the benefit less often than every year or introducing an administrative lag between the end of the year and the benefit change.

than neutral or best-estimate) scenario. This rationale allows them to pay for all their expenses, even when benefits dip lower than expected. To capture this, the benefit statistic in the BaR metric is set to the minimum benefit observed over a given horizon of $\tau = 5$ years in this subsection.¹⁰

The next step is choosing a comparator. We assume that members would compare potential benefits to their current benefit level $B(0)$.¹¹ This choice is justified from a habit formation perspective: in the short term and for budgeting purposes, members care most about shortfalls relative to the current level (see Pollak, 1970; MacDonald et al., 2013).

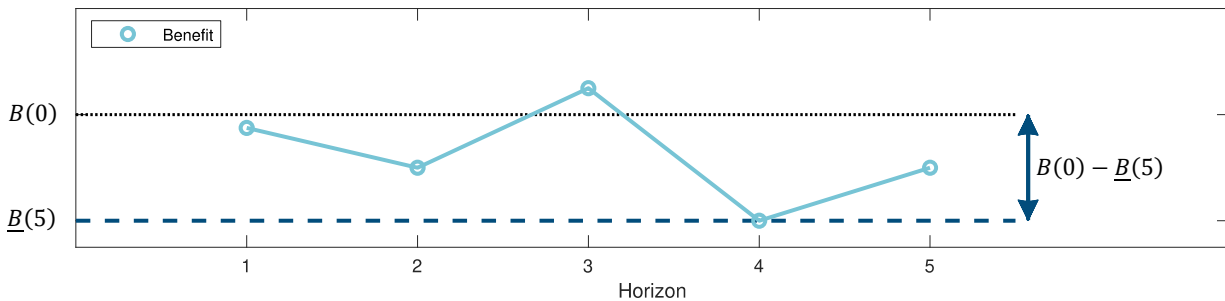
Let $B(t)$ represent the time- t benefit, where time 0 corresponds to the time the BaR is being computed. For illustration, we assume benefits are revised once a year, with no lag between the end of the year and payment of the adjusted benefit. The minimum benefit over a horizon of τ years is then given by

$$\underline{B}(\tau) = \min_{t \in \{1, \dots, \tau\}} B(t),$$

and the resulting worst shortfall over this horizon under a single scenario is $B(0) - \underline{B}(\tau)$.¹²

Figure 2 shows a hypothetical benefit path over a five-year period (light blue circles). The benefit departs from its current value, $B(0)$, and changes over time. For this specific path, the minimum (or worst) benefit in the next five years is observed in year 4 (dashed blue line). The difference between the current benefit (dotted black line) and this minimum represents the member's greatest benefit shortfall over the next five years relative to the comparator under this particular scenario.

Figure 2.
HYPOTHETICAL BENEFIT PATH AND ITS MINIMUM $\underline{B}(5)$.



This figure shows a hypothetical benefit path over a five-year period (light blue circles) along with the current benefit $B(0)$ (dotted black line), and the minimum benefit over the period $\underline{B}(5)$ (dashed blue line). The difference between the current benefit and the minimum benefit over a five-year horizon is the amount by which the minimum benefit $\underline{B}(5)$ falls short of the comparator $B(0)$.

Note that, while the value of the comparator is known *ex ante*, $\underline{B}(5)$ is a random quantity whose value becomes known only *ex post* (i.e., at the end of the horizon). Different hypothetical benefit paths give rise to different shortfalls. When budgeting, it is helpful to have a sense of how severe these shortfalls can get in extreme scenarios. This is what the quantity mBaR(5) attempts to capture: it is the biggest shortfall observed over the next five years relative to the current benefit in an adverse scenario that represents a 1-in-40 event (2.5% probability). Looking at it another way,

¹⁰ Five years is obviously a subjective choice; shorter or longer horizons could have been selected based on the pool, its members' risk aversion, and their investment horizon. We let the end users determine the horizon that fits their needs and context.

¹¹ In our setting, the current benefit corresponds to the starting or initial benefit. In all generality, however, the BaR measures can be calculated at any time after the inception of the pool, and the current benefit is therefore interpreted as the most recent benefit that was paid to members.

¹² Note that we do not include in our calculation the current benefit (i.e., that at time 0) because it is known and does not involve any uncertainty.

the biggest benefit shortfall observed over the next five years is projected to be smaller than $mBaR(5)$ 39 times out of 40. The probability level 97.5% is chosen to be conservative but not overly cautious.^{13,14}

Mathematically, we replace the benefit statistic $\beta(\tau)$ with the minimum $\underline{B}(\tau)$ in Equation (1) to obtain

$$mBaR(5) \equiv BaR_{97.5\%}[B(0) - \underline{B}(5)] = F_{B(0) - \underline{B}(5)}^{-1}(0.975), \quad (2)$$

where $mBaR(5)$ stands for the five-year minimum benefit at risk, while assuming a probability level of 97.5%, and the current benefit as the comparator.

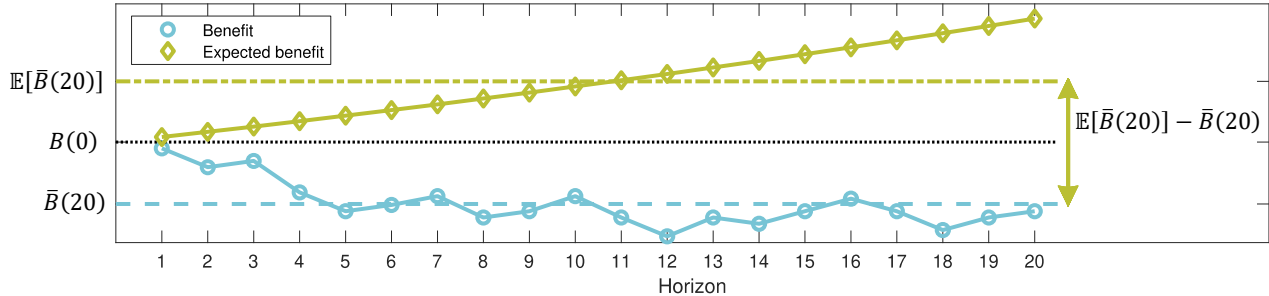
Section 4 applies this BaR measure to lifetime pension pools and provides robustness tests to assess some of the subjective choices made above.

2.4 BAR AS A DECISION-MAKING TOOL

Another use of BaR, from the members' perspective, is as a decision-making tool. Members could be interested in comparing different arrangements before entering any of them, and this comparison could include an assessment of the potential for benefit shortfalls under each arrangement.

The relevant horizon for decision-making purposes is longer than the budgeting horizon considered in Section 2.3. We employ a horizon of 20 years, which is close to an average life expectancy at retirement. One could select a slightly longer horizon if needed, especially in light of the longevity improvements observed over the last decades; we expect that changing the horizon from 20 years to 25 or 30 years will not impact the measure in a material way.

Figure 3.
HYPOTHETICAL BENEFIT PATH AND ITS AVERAGE $\bar{B}(20)$.



This figure shows a hypothetical benefit path over a 20-year period (light blue circles) along with the current benefit $B(0)$ (dotted black line), the expected benefit path (green diamonds), the expected average benefit over the next 20 years $\mathbb{E}[\bar{B}(20)]$ (dot-dashed green line), and the average benefit over the period $\bar{B}(20)$ (dashed light blue line). The difference between the expected average benefit and the average benefit over a 20-year horizon is the amount by which the average benefit $\bar{B}(20)$ falls short of the comparator $\mathbb{E}[\bar{B}(20)]$.

¹³ There is obviously a relationship between the probability level and the horizon considered. Indeed, a longer horizon should be coupled with a lower level and a shorter horizon with a high level. As explained in Dhaene et al. (2008), a good approximation rule for the risk measure probability level suitable for a horizon τ is $p = p_{\text{annual}}^{\tau}$, where p_{annual} is the annual probability level. Note that assuming that $p_{\text{annual}} = 99.5\%$ (i.e., the level used for solvency capital requirements under Solvency II) yields an approximate five-year level of about 97.5%—the level selected for the BaR investigated in this subsection.

¹⁴ Even though this level fits the Solvency II framework, we recognize that end users might want to apply different probability levels that fit their purpose and needs. We encourage practitioners to find probability levels that are consistent with their communication and disclosure goals.

Over a longer horizon, there are two main reasons the benefits may change. On the one hand, the arrangement may target a non-level (i.e., either increasing or decreasing) expected benefit pattern.¹⁵ On the other hand, there is statistical uncertainty around this target: actual outcomes will differ from expected, regardless of whether the expected benefits have an increasing, level, or decreasing pattern. To construct a shortfall measure that considers both of these aspects, we focus on the average benefit level over the entire horizon and compare the “actual average” benefit against the “expected average” benefit. In other words, the average observed benefit over our horizon becomes the benefit statistic, and the expected average benefit becomes the comparator.

Finally, for the probability level, we use 90%, which represents a one-in-ten event.¹⁶

Mathematically, the average benefit over a horizon of τ years is given by the expression

$$\bar{B}(\tau) = \frac{1}{\tau} \sum_{t=1}^{\tau} B(t),$$

and the average expected benefit is $\mathbb{E}[\bar{B}(\tau)]$.

Replacing the benefit statistic $\beta(\tau)$ by the average $\bar{B}(\tau)$ and the comparator C with $\mathbb{E}[\bar{B}(\tau)]$ in Equation (1) yields

$$\text{aBaR}(20) \equiv \text{BaR}_{90\%}[\mathbb{E}[\bar{B}(20)] - \bar{B}(20)] = F_{\mathbb{E}[\bar{B}(20)] - \bar{B}(20)}^{-1}(0.9), \quad (3)$$

where $\text{aBaR}(20)$ stands for the 20-year average benefit at risk, while assuming a probability level of 90%.

Figure 3 mirrors Figure 2 but for the average instead of the minimum; it shows a hypothetical benefit path over a 20-year horizon (light blue circles). From this realization, we can obtain the *average* observed benefit over these 20 years, denoted by $\bar{B}(20)$ (dashed light blue line). The (theoretical) expected benefit is reported with green diamonds in Figure 3, and its *average* level over the next 20 years is indicated by the dashed green line. This latter quantity is then compared to the average observed benefit, with the difference between the two representing the amount lost by the member when compared to what he was supposed to obtain (on an expected value basis) over the long term.

Once again, the shortfall $\mathbb{E}[\bar{B}(20)] - \bar{B}(20)$ is a random quantity that varies from scenario to scenario. The metric $\text{aBaR}(20)$ represents the shortfall in an adverse 20-year scenario that is likely to materialize one time out of 10.¹⁷ In other words, in 10% of the possible future scenarios, we expect a shortfall of $\text{aBaR}(20)$ or more.

2.5 LIMITATIONS OF BaR

The benefit at risk suffers from limitations similar to those of the value at risk:¹⁸

- It provides a false sense of security and has a narrow focus. The actual benefit loss can be higher than BaR. Unfortunately, many end users could think of BaR as *the most a member could lose*. In reality, this number can be far from the worst benefit loss a member can come across.

¹⁵ For example, in the context of lifetime pension pools, an increasing expected benefit pattern could arise if the hurdle rate is selected to be significantly lower than the expected portfolio returns. A decreasing expected benefit pattern could arise if the hurdle rate is higher than the expected portfolio returns.

¹⁶ This is consistent with an annual probability level of 99.5%, as prescribed by Solvency II, since $0.995^{20} \approx 0.9$.

¹⁷ Assuming a member was supposed to receive \$10,000 on average every year and their $\text{aBaR}(20)$ is \$1,000, then they should expect to receive a benefit of less than \$9,000 on average every year for the next 20 years in the worst one out of ten future scenarios (such as the one reported in Figure 3).

¹⁸ Another commonly known problem for the VaR is that it is not subadditive (see, e.g., Artzner et al., 1999, for more details). In the context of BaR, however, this problem is less of an issue.

- It is only as good as the assumptions used to compute the measure. The BaR measure depends on some subjective inputs—the probability level, the horizon, the benefit statistic, and the comparator—as well as the assumptions used in the calculation of the benefits—especially the variability of investment returns. A change in these assumptions and inputs will change the end result.
- It should not be used as an objective in an optimization exercise as it can produce sub-optimal decisions. It is well-known that using the VaR—and by extension, the BaR—in a risk minimization problem yields counterintuitive results: by not bringing in the magnitude of the losses once you exceed the cutoff probability, you expose members to very, very low benefits.¹⁹ We therefore have some reservations about the use of BaR by actuaries in the context of design; it would need to be used with other measures to capture all the facets of benefit risk.

Further to the last bullet, one would need to be extremely careful when using BaR in the context of design, because it does not tell the whole story.

- Some designs might produce a lower (and therefore more desirable) BaR but do so at the expense of some other feature of the benefit distribution. For instance, a lifetime pension pool with a low hurdle rate creates a pattern of increasing benefits, which lowers BaR; however, such an arrangement would also produce a considerably lower current benefit. In these cases, considering only the BaR without also taking into account the expected level of benefits might be misleading.
- The benefit at risk does not comment on the variability of the benefits from one year to the next. Indeed, BaR considers only the distribution of benefits across all possible scenarios at specific times and not how benefits might evolve as a function of time along a single scenario.

Despite the limitations and criticisms mentioned above, we believe that BaR is a great device when used appropriately—as a budgeting and decision-making tool for members, for instance.

¹⁹ In a different context, Basak and Shapiro (2001) note that managers using VaR for portfolio allocation often invest in more risky portfolios than managers who do not use VaR as a risk assessment tool.

Section 3: Lifetime Pension Pools

As mentioned in the introduction, some lifetime pension pool designs already exist. This report's goal is not to assess and compare the different designs but rather to show how to compute BaR in the context of a given hypothetical pool design. We select a very simple structure for which benefits are adjusted annually based on asset returns and the realized mortality experience. The operation of the pool is similar to those explained in Piggott et al. (2005) in the context of group self-annuitization plans. It is also reminiscent of the benefit update rule used by the College Retirement Equities Funds (CREF, 2022a,b).

This plan is far from being general and does not encompass all possible lifetime pension pools. Rather, we wish to focus on a simple version of the pool—closed membership group and no systematic mortality improvements—to gain better insights into the behaviour of BaR in the context of these arrangements.²⁰

In addition to providing technical details on the stylized lifetime pension pools investigated in this report, this section also gives an example of how benefits are determined. The last part of this section discusses mortality pooling—one of the main features of these lifetime pension pools.

3.1 SIMPLE FRAMEWORK FOR VARYING ANNUITY PAYOUTS

We consider a simple, stylized lifetime pension pool in the spirit of Piggott et al. (2005), Valdez et al. (2006), Qiao and Sherris (2013), and Hanewald et al. (2013).^{21,22} We follow the main steps used to update benefits in these plans.

We use \mathcal{L}_t to represent the set of survivors at time t ; that is, $k \in \mathcal{L}_t$ if and only if the k^{th} life is alive at time t . Assume that at time 0, a pool of annuitants decides on the amount they wish to invest in a (closed) pool; $A_k(0)$ denotes the initial investment made by member k , so that the total asset value is given by

$$A(0) = \sum_{k \in \mathcal{L}_0} A_k(0).$$

For each member, the current benefit of $B_k(0)$ (or initial benefit in this case) is obtained from

$$B_k(0) = \frac{A_k(0)}{\ddot{a}_{x_k}}, \quad (4)$$

where \ddot{a}_{x_k} is the price of an annuity-due for the k^{th} life, who is assumed to be aged x_k at inception; that is,

$$\ddot{a}_{x_k} = \sum_{s=0}^{\infty} {}_s p_{x_k} \exp(-sh),$$

using standard actuarial notation. Note that, for convenience's sake, we assume that the stylized lifetime pension pool pays benefits at the beginning of the year.

²⁰ Simple modifications of the basic plan can allow for an open membership group and systematic mortality improvements. We comment on these two modifications in the present report without implementing these alterations.

²¹ This stylized lifetime pension pool is also called nominal-payout method in the literature (see, e.g., Sabin and Foreman, 2016). The rationale of this method is to update benefits by distributing the decedents' accounts to survivors in proportion to each surviving member's nominal payout.

²² Sabin (2010) and Donnelly (2015) showed that this scheme introduces a bias that favours some members over others; that is, the expected value of the payout is better for some members than that implied by actuarial neutrality, and worse for others. It has been argued in these studies that this bias is very small if the number of members is large.

In this stylized lifetime pension pool, the annuity price is obtained based on two main assumptions:

1. A fixed, constant life table that is representative of systematic mortality in the pool.²³
2. A fixed interest rate—the so-called (continuously compounded) hurdle rate—denoted by h .

In this report, we use continuously compounded rates for mathematical convenience. We acknowledge that some practitioners might be more familiar with effective annual rates. In this case, one can simply replace the discounting factors using the following relationship:

$$\exp(-sh) = (1 + \tilde{h})^{-s},$$

where $\tilde{h} = \exp(h) - 1$. These substitutions should not change the results presented in this report, assuming that continuously compounded rates are appropriately converted into effective annual rates.

Note that, unlike in Section 2, we add subscripts to the benefits to make it clear that the benefits can be different from one member to another in this section.

We now present a general rule for the development of future benefit payments in the case where the actuarial survival pattern is different than expected, and the realized rate of return on the asset portfolio is different from the hurdle rate.

At time 0, the payment to each survivor is given by Equation (4). At time 1, however, the total asset value becomes

$$A(1) = \left(A(0) - \sum_{k \in \mathcal{L}_0} B_k(0) \right) \exp(r_1^{\text{PF}}) = \sum_{k \in \mathcal{L}_0} B_k(0) (\ddot{a}_{x_k} - 1) \exp(r_1^{\text{PF}}),$$

where r_1^{PF} is the time-1 continuously compounded rate of return realized on the asset portfolio.²⁴ The equation above is obtained retrospectively: simply put, it is a roll-forward of assets using actual benefits and actual investment returns. One can also obtain a value for $A(1)$ in a prospective fashion:

$$A(1) = \sum_{k \in \mathcal{L}_1} B_k(1) \ddot{a}_{x_{k+1}} = \sum_{k \in \mathcal{L}_1} (\alpha_1 B_k(0)) \ddot{a}_{x_{k+1}},$$

where α_1 is the time-1 adjustment applied to the time-0 benefit; that is, $B_k(1) = \alpha_1 B_k(0)$. Simply put, the value of the adjusted future benefits should be equal to the accumulated value of the assets. Equating the retrospective and prospective values for $A(1)$ gives

$$\alpha_1 = \left(\frac{\sum_{k \in \mathcal{L}_0} B_k(0) (\ddot{a}_{x_k} - 1)}{\sum_{k \in \mathcal{L}_1} B_k(0) \ddot{a}_{x_{k+1}}} \right) \exp(r_1^{\text{PF}}),$$

which can be further simplified by using the recursive relationship between \ddot{a}_{x_k} and $\ddot{a}_{x_{k+1}}$:

²³ We assume that this is known and fixed. In practical application, systematic mortality might change, and the lifetime pension pool operator should adjust the life table from time to time. In this case, the benefits will need to be adjusted. Piggott et al. (2005) make this adjustment by cohort, whereas Qiao and Sherris (2013) propose a common adjustment for all members.

²⁴ To use effective annual rates instead, $\exp(r_t^{\text{PF}})$ should be replaced with $(1 + \tilde{r}_t^{\text{PF}})$, where $\tilde{r}_t^{\text{PF}} = \exp(r_t^{\text{PF}}) - 1$ is the effective annual rate equivalent to the continuously compounded rate r_t^{PF} .

$$\ddot{a}_{x_k} - 1 = p_{x_k} \ddot{a}_{x_{k+1}} \exp(-h).$$

Indeed, this simplification yields the following expression:

$$\begin{aligned} \alpha_1 &= \left(\frac{\sum_{k \in \mathcal{L}_0} B_k(0) p_{x_k} \ddot{a}_{x_{k+1}}}{\sum_{k \in \mathcal{L}_1} B_k(0) \ddot{a}_{x_{k+1}}} \right) \times \exp(r_t^{\text{PF}} - h) \\ &= \text{MEA}_1 \times \text{IEA}_1. \end{aligned}$$

This formula says that the adjustment applied at time-1 is the product of the time-1 mortality experience adjustment factor, MEA_1 , and the time-1 investment experience adjustment, IEA_1 . This rationale holds at any time, so the time- t adjustment is given by

$$\alpha_t = \left(\frac{\sum_{k \in \mathcal{L}_{t-1}} B_k(t-1) p_{x_{k+t-1}} \ddot{a}_{x_{k+t}}}{\sum_{k \in \mathcal{L}_t} B_k(t-1) \ddot{a}_{x_{k+t}}} \right) \times \exp(r_t^{\text{PF}} - h),$$

leading to the following benefit update rule:

$$\begin{aligned} B_k(t) &= B_k(t-1) \times \left(\frac{\sum_{j \in \mathcal{L}_{t-1}} B_j(t-1) p_{x_{j+t-1}} \ddot{a}_{x_{j+t}}}{\sum_{j \in \mathcal{L}_t} B_j(t-1) \ddot{a}_{x_{j+t}}} \right) \times \exp(r_t^{\text{PF}} - h) \\ &= B_k(t-1) \times \text{MEA}_t \times \text{IEA}_t, \end{aligned} \quad (5)$$

for member k . Notice that if every member retires at the same age x , the mortality experience adjustment simplifies to

$$\text{MEA}_t = \frac{p_{x+t-1}}{p_{x+t-1}^*}, \quad \text{where } p_{x+t-1}^* = \frac{\sum_{j \in \mathcal{L}_t} B_j(t-1)}{\sum_{j \in \mathcal{L}_{t-1}} B_j(t-1)}, \quad (6)$$

which is the actual survival rate weighted by the benefit payment.

This conceptual framework could be extended in multiple dimensions (e.g., open membership group, systematic mortality improvements). In the present report, we focus on closed pools, but nothing prevents one from including new entrants in the set of survivors. To do so, the asset value would need to be adjusted by adding the contributions from new entrants to the pool and adjusting the set of survivors accordingly.

3.2 EXAMPLE

To better understand the operation of these stylized lifetime pension pools, we now consider a simple numerical example. We consider 100 new members at time 0, each aged 65.

Let us assume that the systematic mortality is modelled by the female CPM 2014 table without generational adjustments and that the (continuously compounded) hurdle rate is set to 4.5%.²⁵ The actuarial present value of the annuity-due associated with these assumptions is $\ddot{a}_{65} = 14.3410$.

If each member deposits \$143,410 in the pool, bringing the total assets at inception to \$14,341,000, this initial asset value allows every member to receive a benefit of

²⁵ We do not include generational adjustments for convenience; the results are not affected by this simplification.

$$B_k(0) = \frac{143,410}{14,3410} = 10,000, \quad \forall k \in \mathcal{L}_0,$$

at inception. After paying these benefits, the total asset value drops to

$$14,341,000 - \sum_{k \in \mathcal{L}_0} B_k(0) = 14,341,000 - (100)(10,000) = 13,341,000.$$

This amount is then invested for a year in a portfolio earning an uncertain rate of return. Suppose that the continuously compounded rate of return realized in the first year is $r_1^{\text{PF}} = 3\%$; this leads to a time-1 asset value of

$$A(1) = 13,341,000 \times \exp(0.03) = 13,747,294.$$

Next, assume that five members die during the first year; the time-1 post-redistribution benefit for each surviving member is therefore given by

$$B_k(1) = B_k(0) \frac{A(1)}{\sum_{j \in \mathcal{L}_1} B_j(0) \ddot{a}_{x_{j+1}}} = \frac{13,747,294}{(100 - 5)(14.0339)} = 10,311, \quad \forall k \in \mathcal{L}_1,$$

where $\ddot{a}_{66} = 14.0339$ as per the CPM 2014 life table. In other words, the time-1 benefit is the time-1 asset value divided across the survivors (while accounting for their new annuity factor and given that all members are now aged 66).

We obtain identical adjustments when using the formula developed in Section 3.1:

$$\begin{aligned} \text{MEA}_1 &= \left(\frac{\sum_{k \in \mathcal{L}_0} B_k(0) p_{65} \ddot{a}_{66}}{\sum_{k \in \mathcal{L}_1} B_k(0) \ddot{a}_{66}} \right) = \left(\frac{0.9944}{\left(\frac{100 - 5}{100} \right)} \right) = 1.0467, \\ \text{IEA}_1 &= \exp(r_1^{\text{PF}} - h) = \exp(0.03 - 0.045) = 0.9851, \end{aligned}$$

leading to an adjustment factor of $\alpha_1 = \text{MEA}_1 \times \text{IEA}_1 = 1.0311$, and a benefit of

$$B_k(1) = B_k(0) \times \alpha_1 = 10,000 \times 1.0311 = 10,311.$$

Even though investment returns were below the hurdle rate—which should have caused a decrease in the time-1 benefits—more members than expected died during the first year. This ultimately leads to an adjustment factor that is higher than 1, meaning that the time-1 benefits are larger than those paid at time 0.

By applying the same logic recursively, one can obtain the benefit for each year and member once the realized mortality and rate of return are observed.

3.3 ADJUSTMENT FACTOR FOR SYSTEMATIC MORTALITY IMPROVEMENTS

As mentioned above, the equations of Section 3.1 could be altered to allow for systematic mortality improvements. To show the impact of this modification on our basic setup, let us first introduce some notation. First, the time- t actuarial present value of the annuity-due is given by $\ddot{a}_{x_k}^{[t]}$ for a member aged x_k . Similarly, the one-year survival probability measured at time t for a member aged x_k is given by $p_{x_k}^{[t]}$.

In this case, the mortality adjustment factor becomes

$$\text{MEA}_t = \left(\frac{\sum_{k \in \mathcal{L}_{t-1}} B_k(t-1) p_{x_{k+t-1}}^{[t-1]} \ddot{a}_{x_{k+t}}^{[t-1]}}{\sum_{k \in \mathcal{L}_t} B_k(t-1) \ddot{a}_{x_{k+t}}^{[t]}} \right).$$

The modified version of the mortality experience adjustment contains sums of annuity factors under the old and the new mortality assumptions, weighted by the number of retirees at each age and the amount of their benefits, thus accounting for any expected changes in systematic mortality.

3.4 DISCUSSION ON MORTALITY POOLING

Interestingly, most of the literature focuses on the effect of (idiosyncratic) mortality on the benefit stream (see, e.g., Piggott et al., 2005, Qiao and Sherris, 2013, Olivieri and Pittacco, 2020), giving less regard to the impact of the realized rate of return on the asset portfolio. We take a different route: we assume at first that the pool size is sufficiently large and that all starting benefits are the same so that the impact of potentially imperfect mortality pooling is negligible.^{26,27} This is equivalent to assuming that the idiosyncratic mortality risk is fully diversified. In practical terms, this also means that $p_{x+t-1}^* = p_{x+t-1}$, and that $\text{MEA}_t = 1$ for all t in a closed membership pool for which all entrants have the same age x at inception.

The largest possible reduction in benefits due to mortality experience arises in a year when none of the members die. Since expected mortality rates are quite low early in retirement—especially in comparison to the potential reduction in benefits due to a stock market crash—setting $\text{MEA}_t = 1$ cannot materially understate mortality risk for a younger group of members. This assumption is somewhat unrealistic in closed pools with older members, and we recommend using it carefully when assessing the actual risk in these pools. To introduce the benefit at risk in the context of lifetime pension pools, however, the additional complication of having mortality risk is unnecessary as it will only have a second-order impact on our end results.²⁸

Based on this simplification, the only variable impacting the benefit level in large pools is the portfolio return. Using a simple recursion, one can write the time- t benefit level from Equation (5) as

$$B_k(t) = B_k(0) \prod_{s=1}^t \exp(r_s^{\text{PF}} - h) = B_k(0) \exp\left(\sum_{s=1}^t r_s^{\text{PF}} - th\right). \quad (7)$$

Equation (7) will be used to understand the evolution of benefits in Sections 4.1 to 4.4 of this report. In Section 4.5, we will consider a simple idiosyncratic mortality assumption to assess its impact on large and small pools.

²⁶ According to Sabin and Forman (2016), who investigate similar pools, the payout volatility depends almost exclusively on the investment strategy as soon as the membership is large enough (i.e., more than 1,000 members).

²⁷ This approach was also used in Bernhardt and Donnelly (2019) in their article on modern tontines.

²⁸ The results of Section 4.5 show that idiosyncratic mortality, when adequately diversified and when using the right mortality table, has a minimal impact on the benefit adjustments. Asset returns are therefore the main determinant of the benefit adjustments for large pools.

Section 4: Benefit at Risk in Lifetime Pension Pools

In this section, we apply the concept of BaR to the stylized lifetime pension pool introduced in Section 3. We begin by providing an analytical exploration of the two specific benefit at risk measures—mBaR(5) and aBaR(20)—described above. Then, in Section 4.2, we provide some sensitivity testing and robustness analyses to understand if the two proposed measures are robust to changes in some of the inputs (i.e., the hurdle rate and the asset allocation). Section 4.3 assesses the impact of inflation on the benefits and redefines BaR in real terms. Then, in the fourth part of this section, we discuss another robustness test in which we consider a slightly more realistic risky asset return distribution obtained by bootstrapping. This new assumption allows us to comment on the robustness of the analytical expressions obtained in Section 4.1. Finally, in the last subsection, we discuss the impact of including idiosyncratic mortality risk in the BaR calculation.

Again, we stress that the two specific BaR measures introduced in this report are based on subjective choices relating to budgeting and decision-making concerns. These choices are far from being unique; indeed, there is nothing stopping end users from creating their own BaR metrics that fit their purpose. Note that the BaR measure selected in the end might also depend on the pool design.

4.1 ANALYTICAL EXPLORATION

For illustrative purposes, we begin our analysis with the stylized plan described in Section 3.1, assuming a closed membership group. Members join the pool at 65 years old. They all put the same initial asset $A_k(0)$ in the fund, for $k \in \mathcal{L}_0$. These two assumptions make the benefits equal for each member; that is,

$$B(t) \equiv B_k(t), \quad \forall k \in \mathcal{L}_t.$$

As the first step in our analysis, we assume that the continuously compounded asset portfolio returns r_t^{PF} are normally distributed with a mean of $\nu \equiv \omega\mu + (1 - \omega)r$ and a standard deviation of $\zeta = \omega\sigma$, where the ω could be loosely interpreted as the proportion invested in the risky asset, μ the risky asset average return, σ its standard deviation, and r the risk-free rate.²⁹

Because the sum of independent normal random variables is also normal, we rewrite the time- t benefit from Equation (7) as

$$B(t) = B(0) \exp(R_t),$$

where R_t , the cumulative return on the asset portfolio in excess of the hurdle rate, is normally distributed with mean $t(\nu - h) = t(\omega\mu + (1 - \omega)r - h)$ and standard deviation $\sqrt{t}\zeta$.

In this section, we use the following parameters: an average risky asset return of 7%, a risky asset standard deviation of 15%, and a risk-free rate of 2%.³⁰ The proportion invested in the risky asset is set to 50%. The hurdle rate is set to the expected portfolio return at first; that is, $h = \nu$. We assume that each lifetime pension pool member deposits

²⁹ This approximate distribution is obtained by combining the gross return on the risky asset (i.e., a stock index in this case) and the gross return on the risk-free asset such that

$$\begin{aligned} r_t^{\text{PF}} &= \log(\omega e^{s_t} + (1 - \omega)e^r) \\ &\approx \omega s_t + (1 - \omega)r, \end{aligned}$$

using a first-order Taylor expansion, where s_t is the risky asset continuously compounded return. If the time- t risky asset return is normally distributed with a mean of μ and a standard deviation of σ , then the portfolio return is approximately normal with a mean of $\omega\mu + (1 - \omega)r$ and a standard deviation of $\omega\sigma$.

³⁰ These numbers are consistent with the Canadian economy. Between 1990 and 2021, the average excess return (including dividends) on the S&P/TSX Composite index was about 5% and its volatility was 15%. By adding the risk-free rate of 2% to the average excess return, we obtain an average return of 7% on the risky asset.

\$143,410 in the fund. The price of the initial annuity is $\ddot{a}_{65} = 14.3410$ using the female CPM 2014 table without generational adjustments, resulting in an initial benefit of \$10,000.

4.1.1 MINIMUM BENEFIT AT RISK


The minimum BaR—used for budgeting purposes—is based on the minimum benefit over the next τ years. To compute its value, we therefore need to find the cdf of the current benefit, $B(0)$, minus the minimum, $\underline{B}(\tau)$, which is the comparator minus the benefit statistic for the minimum BaR.³¹ Using the normality assumption stated above, we can find the cdf of this random variable in semi-closed form, up to an integral. The five-year minimum BaR in this context is thus given by

$$\text{mBaR}(5) = F_{B(0) - \underline{B}(5)}^{-1}(0.975),$$

where the cdf of $B(0) - \underline{B}(5)$, denoted by $F_{B(0) - \underline{B}(5)}$, is given in Appendix A.1.³²

Figure 4 reports the minimum BaR measure for different horizons. For a horizon of five years—the one used to define $\text{mBaR}(5)$ in Section 2—the measure is equal to \$2,895. This number means that, in one out of 40 future scenarios, a member should expect that their worst annual benefit shortfall experienced over the next five years (relative to the \$10,000 current benefit) will exceed \$2,895. In most scenarios (39 times out of 40), the worst annual benefit shortfall will be less than \$2,895, meaning that the annual benefit the member receives should be more than \$7,105 in each of the next five years in these scenarios.

Figure 4.
MINIMUM BENEFIT AT RISK AS A FUNCTION OF THE HORIZON.



This figure reports the minimum benefit at risk as a function of the horizon. The BaR level is set to 97.5%, and the hurdle rate is set to the expected portfolio return; that is, 4.5%. Half of the portfolio is invested in the risky asset, and the other half is invested in the risk-free asset. We assume that the retiree joins the pool at 65 with \$143,410, so the current benefit is set to \$10,000.

The horizon seems to matter here. For short horizons, minimum BaRs vary, with a value of \$2,292 for a horizon of three years and \$3,355 for a horizon of seven years. While we chose a five-year horizon, other short horizons might also be adequate in this context.³³

³¹ Appendix B.1 considers a different version of the minimum BaR for which the comparator is set to an exogenous quantity determined by the members. For instance, this amount could be a target level of benefits or an external benchmark.

³² In this study, we rely on quadrature methods to obtain this cdf. The cdf is inverted numerically using the Newton-Raphson method to obtain the mBaR.

³³ Some insights from a recent Society of Actuaries report support a horizon of five years in the context of budgeting, although the actual concern of the survey was more about lump sum losses and less about drops in annual income. See Greenwald Research (2022) for more details.

4.1.2 AVERAGE BENEFIT AT RISK

The distribution of the average benefit is more complicated as it involves a sum of dependent random variables. When the continuously compounded returns are assumed to be normally distributed, each year's investment adjustment IEA_t follows a lognormal distribution. Consequently, the cumulative adjustments to each year's benefit in Equation (7) also follow a lognormal distribution (although the cumulative adjustments are not independent). The sum or average of all the adjustments does not exactly follow a lognormal distribution. Nonetheless, it has been argued in the literature that one can approximate such distributions by moment-matched lognormal distributions, henceforth called the Fenton-Wilkinson approximation (see Fenton, 1960, for a description of the method and Dufresne, 2008, for a review of approximation methods for sums of lognormal random variables).


Appendix A.2 shows closed-form expressions for the first two moments of $\bar{B}(\tau)$. Note that these are obtained recursively. Assuming that the average benefit $\bar{B}(\tau)$ is given by a lognormal random variable with parameters $\mu_{\bar{B}(\tau)}$ and $\sigma_{\bar{B}(\tau)}$, we have that

$$\begin{aligned} \text{aBaR}(20) &= F_{\mathbb{E}[\bar{B}(20)] - \bar{B}(20)}^{-1}(0.90) = \mathbb{E}[\bar{B}(20)] - F_{\bar{B}(20)}^{-1}(0.10) \\ &= \exp(\mu_{\bar{B}(20)} + \sigma_{\bar{B}(20)}^2/2) - \exp(\mu_{\bar{B}(20)} + \sigma_{\bar{B}(20)} \Phi^{-1}(0.10)). \end{aligned}$$

This formula for $\text{aBaR}(20)$ is simply the difference between the mean and the 10th percentile of a lognormal distribution with the same mean and standard deviation as the average of 20-year average benefits.³⁴

Figure 5 shows average BaRs for different horizons, assuming a probability level of 90%. The parameters used to generate this figure are the same as those used in the previous figure. As expected, average BaRs are lower than those for the minimum benefits in Figure 4. For a 20-year horizon, the 90% average BaR is \$2,180, meaning that in one out of ten future scenarios, we expect members to lose more than \$2,180 annually when compared to what they were supposed to obtain, on average. In other words, there is a 10% chance that members lose more than \$2,180 per year in benefits by investing in this pool.³⁵

Figure 5.
AVERAGE BENEFIT AT RISK AS A FUNCTION OF THE HORIZON.



This figure reports the average benefit at risk as a function of the horizon. The BaR level is set to 90%, and the hurdle rate is set to the expected portfolio return; that is, 4.5%. Half of the portfolio is invested in the risky asset, and the other half is invested in the risk-free asset. We assume that the retiree joins the pool at 65 with \$143,410, so the current benefit is set to \$10,000.

³⁴ The lognormal approximation is very precise in our context and yields results in a split second, which is a lot faster than brute force Monte Carlo methods that could be used to compute the benefit at risk.

³⁵ Similar to Appendix B.1, Appendix B.2 introduces a version of aBaR that uses an exogenous comparator.

4.2 SENSITIVITY TESTING AND ROBUSTNESS ANALYSES

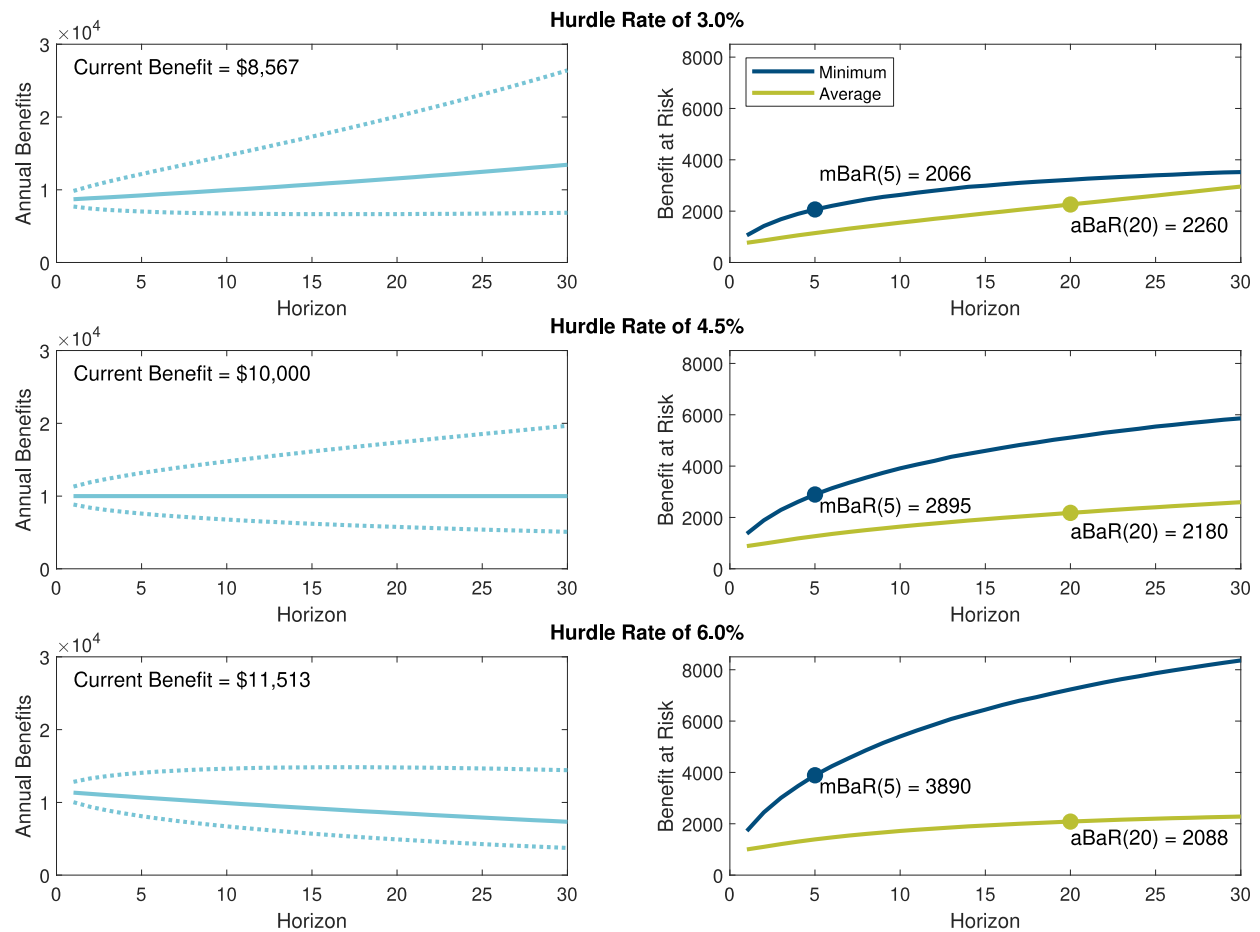
Section 4.1 showed minimum and average BaRs for different horizons, with a focus on the five-year horizon for mBaR and the 20-year horizon for aBaR. This section assesses the impact of other assumptions selected in the calculation above; it gives the reader an overview of how BaR behaves with respect to its main inputs (namely, the hurdle rate and the asset allocation strategy).

At this stage, we stress that although reducing BaR can be seen as a good thing, generally speaking, this measure only captures one dimension (i.e., shortfall risk). One should therefore be extremely careful when drawing general conclusions from these sensitivity tests—especially from a design perspective.

4.2.1 HURDLE RATE

The hurdle rate significantly impacts the annual benefits paid to retirees—a lower hurdle rate reduces the current benefit but increases the likelihood of future benefit increases, and vice versa.

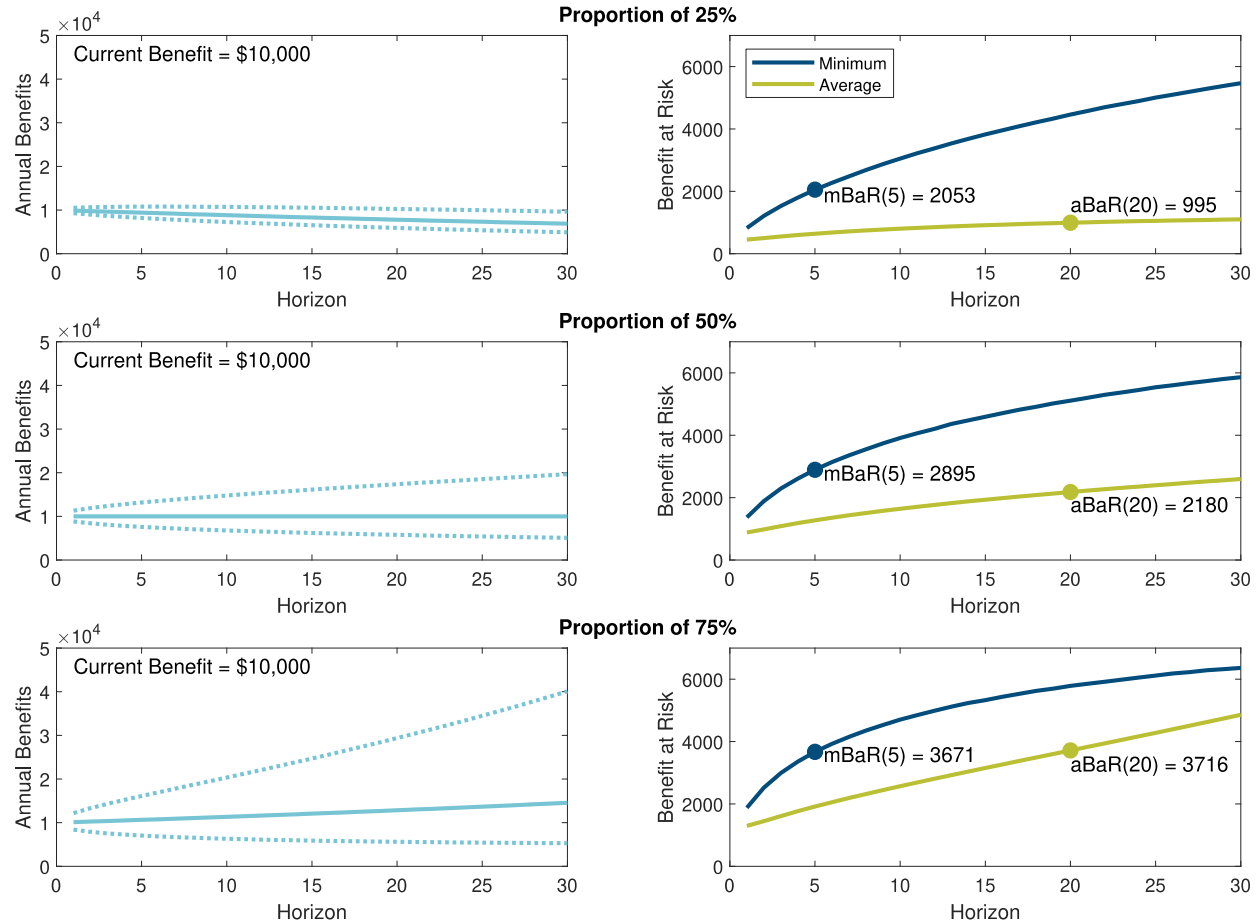
Figure 6.
ANNUAL BENEFIT FUNNEL OF DOUBT AND BENEFIT AT RISK FOR DIFFERENT HURDLE RATES.



The left panels of this figure report funnels of doubt for the annual benefits. We show the median (solid line) as well as 5th and 95th quantiles (dotted lines). We do so for three different hurdle rates: 3.0% (top panel), 4.5% (middle panel), and 6.0% (bottom panel). The right panels of this figure show the minimum BaR at the 97.5% level (blue lines) and the average benefit BaR at the 90% level (green lines). We assume a constant asset allocation strategy (i.e., $\omega = 0.5$).

The left panels of Figure 6 report funnels of doubt (median as well as 5th and 95th quantiles) for the annual benefits. When the hurdle rate is low (i.e., 3.0%), the current benefit is also low (about \$1,400 lower when compared to the base case of Section 4.1). To compensate for this lower current benefit, the benefit tends to increase over time; the median annual benefit increases to more than \$13,000 after 30 years. A hurdle rate of 4.5% corresponds to our base case investigated in Section 4.1; overall, the median benefits are constant at \$10,000, which is the current benefit. For a hurdle rate of 6.0%, we witness a decreasing pattern of annual benefits, on average. Note that this decreasing pattern is compensated by a higher current benefit, which is \$1,500 above the base case level.

Figure 7.
ANNUAL BENEFIT FUNNEL OF DOUBT AND BENEFIT AT RISK FOR DIFFERENT ASSET ALLOCATION STRATEGIES.



The left panels of this figure report funnels of doubt for the annual benefits. We show the median (solid line) as well as 5th and 95th quantiles (dotted lines). We do so for three different asset allocation strategies (i.e., the proportion of assets invested in risky assets): 25% (top panel), 50% (middle panel), and 75% (bottom panel). The right panels of this figure show the minimum BaR at the 97.5% level (blue lines) and the average benefit BaR at the 90% level (green lines). We assume a constant hurdle rate of 4.5% throughout this figure.

The right panels of Figure 6 show minimum BaRs (blue lines) and average BaRs (green lines) for different maturities. Note that, in Figure 6, the comparators for mBaR and aBaR are adjusted in each case to reflect the benefits; in other words, $B(0)$ and $\mathbb{E}[\bar{B}(20)]$ are updated to be consistent with the selected hurdle rate. These two measures are given for the three hurdle rates studied in this section (i.e., 3.0%, 4.5%, and 6.0%). The mBaR measures are relatively low for the low hurdle rate: the five-year mBaR is \$2,066, leading to a minimum benefit higher than \$6,500 in about 97.5% of the future scenarios over the next five years. When the hurdle rate is set at 4.5%—our base case—we obtain a minimum BaR that is higher and increases more steeply as a function of the horizon than that obtained with a hurdle rate of 3.0%. Specifically, for a horizon of five years, we have an mBaR of about \$2,900. This result is to be expected:

with a hurdle rate of 3.0%, we anticipate that the lower current amount is compensated by future higher benefits, whereas a hurdle rate of 4.5% increases the current benefit received by the member but decreases the likelihood of future benefit increases. As the mBaR comparator is the current benefit, having a lower hurdle rate forces $B(0)$ to be lower, reducing the measure.

The final case (bottom-right panel of Figure 6) considers a hurdle rate of 6.0%. In this case, the current benefit starts high but is expected to decrease on average. Indeed, the five-year minimum BaR is \$3,890. Relative to the high current benefit of \$11,513, a member can stand to lose significantly more.

The aBaR values are less sensitive to the choice of hurdle rates than the mBaR values; interestingly, they hover between \$2,088 and \$2,260. This result is a by-product of the 20-year horizon selected in Section 2. Having a longer horizon would increase the impact of the difference in hurdle rates, as shown in Figure 6: low hurdle rates lead to large average BaRs when the horizon is long (i.e., 30 years). The lower hurdle rate does a good job of preserving assets for late in life but sacrifices near-term benefits to do so.

4.2.2 ASSET ALLOCATION

We now turn to asset allocation and its impact on the benefit at risk. Specifically, we change the proportion ω invested in the risky asset (which is assumed to be a stock index in this case), while keeping all other inputs (including the hurdle rate) constant.

The left panels of Figure 7 report funnels of doubt for the annual benefits for three different proportions: 25% (top panel), 50% (middle panel), and 75% (bottom panel). As expected, having riskier assets leads to more uncertainty in the benefit level: the width of the funnel of doubt is small if only 25% of the assets are invested in the risky asset; it becomes much wider when 75% of the portfolio is invested in the risky asset.

The right panels of Figure 7 show minimum and average BaRs over different horizons and for the three asset allocation strategies mentioned above. As expected, a smaller investment in risky assets over short horizons leads to smaller BaR values. For mBaR(5), decreasing the allocation from 50% to 25% reduces the benefit at risk by about \$800; an increase in the allocation from 50% to 75% has the opposite effect on the risk measure—it increases mBaR(5) by \$800.

The 20-year average BaRs are highly impacted by the investment strategy—more so than the impact of the hurdle rate on aBaR(20). The measure of interest decreases by \$1,200 when the allocation to risky assets is reduced from 50% to 25% and increases by about \$1,500 when it is raised to 75%.

Investing a smaller proportion of the members' assets in risky assets is clearly beneficial as it reduces the BaR level in most cases. There is a trade-off, however: it also negatively impacts the potential upside: for instance, the 95th quantile of the 30-year benefit distribution is about \$10,000 if $\omega = 25\%$ but becomes \$40,000 if $\omega = 75\%$.

4.2.3 COMBINED EFFECTS OF THE HURDLE RATE AND THE ASSET ALLOCATION

Sections 4.2.1 and 4.2.2 reported on the impact of the hurdle rate and the asset allocation on the benefit at risk measures, respectively, and concluded that both the minimum and average BaRs are affected by changes in these assumptions. Yet, when varying the asset allocation, it is hard to grasp if the BaRs are changing as a result of the increase of the benefit risk profile or because of the difference between the expected return on the asset portfolio and the hurdle rate that impacts the general trend in future benefits. This section addresses this issue by investigating BaR as a function of the asset allocation ω and the so-called margin, defined as the difference between the expected return on the asset portfolio ν and the hurdle rate h .

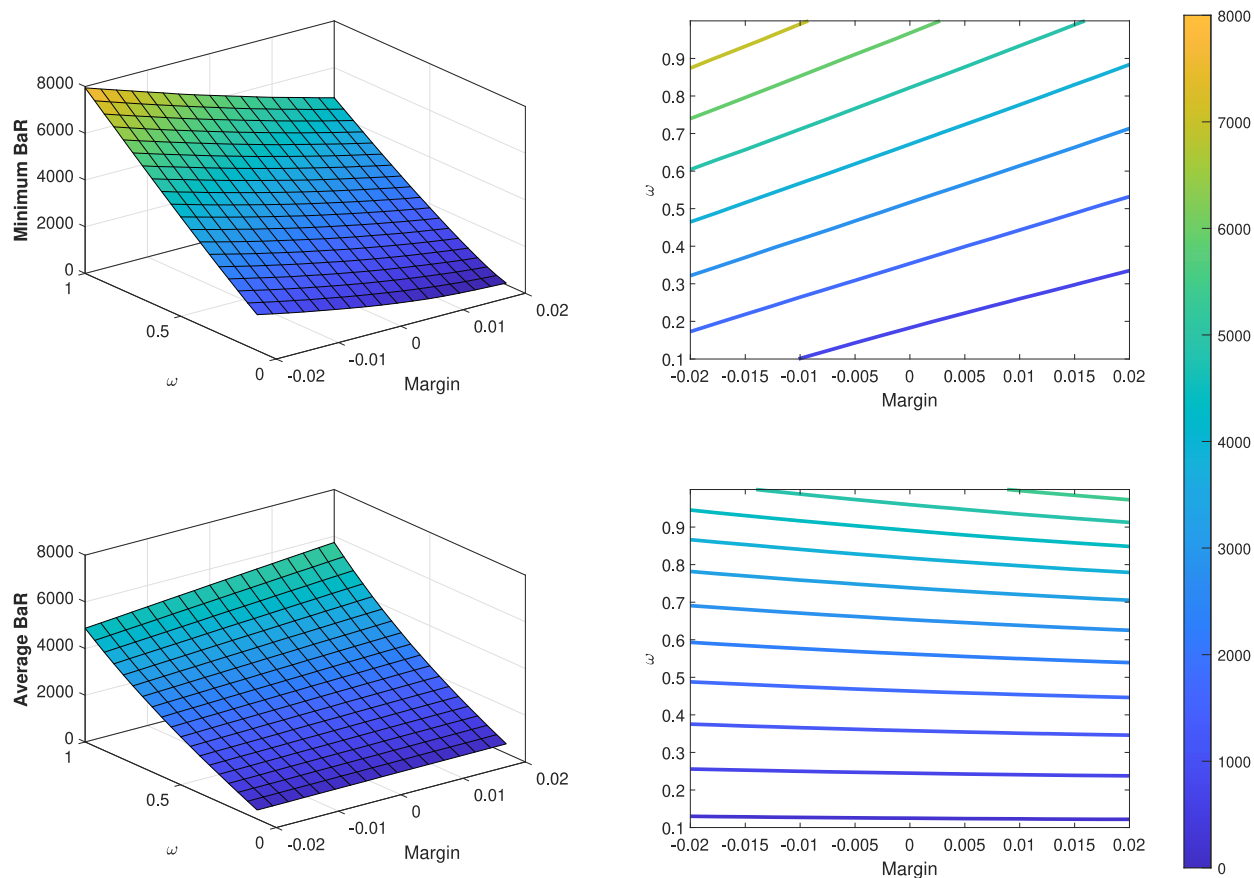
A margin of zero means that the expected return on the portfolio and the hurdle rate are equal, leading to level median benefits (i.e., like those presented in the middle panels of Figures 6 and 7). A positive margin implies that the

assets earn more on average than the rate used to set the current benefit and that future benefits will tend to increase, generally speaking. In contrast, a negative margin suggests the opposite (i.e., larger current benefit and decreasing future benefits).

The change in the margin and in asset allocation strategy affects not only the streams of benefits but also the comparators used to compute mBaR and aBaR. The current benefit changes as a result of a change in the margin, which in turn impacts mBaR. aBaR's comparator is also impacted as the average expected benefit is a function of both the margin and the asset allocation.

The left panels of Figure 8 report minimum and average BaR surfaces for different asset allocation strategies and margins. The minimum BaR (top-left panel) uses a horizon of five years, whereas the average BaR (bottom-left panel) employs a horizon of 20 years, both consistent with the definitions of Section 2. The right panels of Figure 8 disclose information similar to that presented in the left panels but in a different way: the two panels show contour plots—constant slices of the surface rendered on a two-dimensional plane—associated with the surfaces.

Figure 8.
MINIMUM AND AVERAGE BENEFIT AT RISK SURFACES AND CONTOUR PLOTS FOR DIFFERENT ASSET ALLOCATION STRATEGIES AND MARGINS BETWEEN THE EXPECTED RETURN ON THE ASSET PORTFOLIO AND THE HURDLE RATE.



The top-left (bottom-left) panel of this figure reports the five-year minimum BaR (20-year average BaR) surface as a function of the asset allocation parameter ω and the margin defined as the difference between the expected return on the asset portfolio and the hurdle rate. The top-right (bottom-right) panel shows the contour plot associated with the top-left (bottom-left) panel. Contours are defined as constant slices of the surface rendered on a two-dimensional plane.

The mBaR(5) increases as a function of the weight invested in the risky asset ω and decreases as a function of the margin. This behaviour is to be expected:

1. If the portfolio includes a higher proportion of risky assets, the future benefits are less certain, leading to higher values of minimum BaRs.
2. A negative margin implies a hurdle rate that is larger than the expected return on the asset portfolio, leading to a decreasing benefit pattern, all other things being equal. As the mBaR comparator is the current level of benefit, starting with a high benefit value and expecting a decrease seems worse—especially in the budgeting context—than starting with a lesser benefit and expecting an increase. This leads to higher mBaR values and captures the members’ budgeting risk when the design embeds likely benefit reductions.

Interestingly, the 20-year average BaR—used as a decision-making tool—behaves very differently. For portfolios invested less heavily in the risky asset (i.e., below 50% invested in the risky asset), the aBaR(20) is virtually immune to changes in the margin. Over the 20-year horizon, having a high current benefit followed by smaller future benefits, on average, is somewhat equivalent to a low current benefit followed by larger future benefits. This conclusion is a by-product of the comparator used in calculating the aBaR; that is, the average benefit over the next 20 years.

When more than half of the fund is invested in the risky asset, the average BaR tends to be higher for large, positive margins. This result is mainly driven by an increase in benefit risk via a large proportion of the fund being invested in risky assets and the fact that benefits tend to be larger when the margin is large (i.e., when the hurdle rate is low).

4.3 NOMINAL VERSUS REAL TERMS

Some members might be interested in understanding their effective purchasing power and how inflation will impact their benefits. This is of paramount importance as level benefits, such as those obtained with an annuity with fixed payments, reduce a member’s purchasing power over the years.

Fortunately, the BaR definitions introduced in Section 2 can be easily adjusted to account for inflation. Specifically, the time- t benefit in time-0 dollars (i.e., real terms) is given by:

$$B^*(t) = B(t) \exp\left(-\sum_{s=1}^t q_s\right),$$

where q_t is the time- t inflation rate. To account for inflation, the five-year minimum BaR of Equation (2) can be adapted in the following way:

$$\text{mBaR}^*(5) = F_{\underline{B}^*(5)}^{-1}(0.975),$$

where $\underline{B}^*(\tau) = \min_{t \in \{1, \dots, \tau\}} B^*(t)$.

Similarly, the 20-year average BaR of Equation (3), is given by

$$\text{aBaR}^*(20) = F_{\mathbb{E}[\bar{B}^*(20)]}^{-1}(0.9),$$

where $\bar{B}^*(\tau) = \frac{1}{\tau} \sum_{t=1}^{\tau} B^*(t)$ and the average expected benefit in real terms is $\mathbb{E}[\bar{B}^*(\tau)]$.

To illustrate the impact of considering benefits in real terms in a concise manner, we assume a constant inflation rate of 2%, which is in line with the Bank of Canada target and the average inflation between 1990 and 2021 in Canada. In

fact, the Bank of Canada aims to keep inflation at the 2% midpoint of an inflation target range of 1 to 3%.³⁶ All the other assumptions are the same as in Section 4.1.

We consider seven different hurdle rates, from 3.0% to 6.0%, with increments of 0.5%. For each hurdle rate, we compute the current benefit (still using \$143,410 as the current assets owned by each member), the average lifetime benefits (i.e., over the next 50 years), their standard deviation, and the two BaR measures introduced in Section 2 in both nominal and real terms.

These results are presented in Table 1. First, there is a positive relationship between the current benefit and the hurdle rate—as expected. The average lifetime benefits have a negative relationship with the hurdle rate: low hurdle rates build in increases in the benefits, on average, and high hurdle rates exhibit negative trends in the benefits. The standard deviation follows the same trend: lower hurdle rates lead to higher lifetime benefits and higher standard deviations in the benefits. These conclusions are valid both in nominal and real terms, although real average lifetime benefit averages and standard deviations are lower than those obtained with nominal benefits.

Table 1.

BENEFIT AT RISK AS A FUNCTION OF THE HURDLE RATE IN NOMINAL AND REAL TERMS.

Hurdle Rate	Current Benefit	Nominal Terms				Real Terms			
		Mean	Std. Dev.	mBaR(5)	aBaR(20)	Mean	Std. Dev.	mBaR*(5)	aBaR*(20)
3.0%	8,567	13,942	3,937	2,066	2,260	8,106	2,068	2,619	1,755
3.5%	9,035	12,741	3,513	2,327	2,235	7,562	1,877	2,907	1,741
4.0%	9,513	11,684	3,142	2,603	2,208	7,080	1,709	3,211	1,725
4.5%	10,000	10,752	2,818	2,895	2,180	6,652	1,560	3,534	1,708
5.0%	10,496	9,932	2,534	3,209	2,150	6,271	1,430	3,870	1,689
5.5%	11,001	9,208	2,286	3,540	2,119	5,933	1,314	4,216	1,670
6.0%	11,513	8,569	2,068	3,890	2,088	5,631	1,212	4,584	1,650

This table reports the current benefit, statistics on the average lifetime benefit (i.e., over the next 50 years), and benefit at risk measures in nominal and real terms. The mBaR(5) and mBaR*(5) use a level of 97.5%, and the aBaR(20) and aBaR*(20) use a level of 90%. We assume a constant asset allocation strategy (i.e., $\omega = 0.5$). The real dollar values are obtained in terms of time-0 dollars.

The mBaR measure in real terms also increases as a function of the hurdle rate. Low hurdle rates lead to increasing benefit patterns and a lesser likelihood of low benefits when compared to the current benefit (i.e., the comparator in the context of mBaR). For high hurdle rates, it is the opposite. The mBaR(5) in real terms is higher than that in nominal terms as inflation will reduce the members' purchasing power over time, making their benefits lower in real terms.

The results of Table 1 might have some budgeting implications for members: low hurdle rates lead to lesser benefit risk (especially in the left tail), which translates into small BaR values in real terms. In fact, the increasing trend in the benefits cancels some of the loss in purchasing power due to the positive inflation.

Regarding the aBaR(20), lower hurdle rates lead to slightly higher risk measures, both in nominal and real terms. The difference among all the average benefit at risk measures seems rather small as a function of the hurdle rate, however. Indeed, as mentioned above, having a high current benefit followed by lower future benefits should be somewhat similar, on average, to having a low current benefit followed by higher future benefits. This conclusion does not change for the average BaR in real terms.


³⁶ To understand the risk associated with inflation rate changes, one should include an inflation model in their framework. Inflation is indeed an integral part of most economic scenario generators proposed since the 1980s (see, e.g., Wilkie, 1986, 1995; Ahlgrim et al., 2005; Bégin, 2021).

4.4 BOOTSTRAPPING BACKTESTING

A key assumption in the BaR calculation is the distribution used for the risky asset annual returns. In Section 4 so far, we have relied on a normal distribution with a mean return of 7% and a standard deviation of 15%, consistent with the Canadian economy between 1990 and 2021. Normality is a common assumption for continuously compounded annual returns in the pension context, but this is clearly not the only one used in the literature.³⁷

To assess the impact of this assumption on our results, we select another, non-parametric assumption for comparison. Specifically, we rely on block bootstrap to generate risky asset returns. The rationale behind this methodology is to rely on past data and to resample—bootstrap—the past realizations to create new series of returns (see Appendix C for more details). We use 100,000 replications in our calculation.³⁸

Figure 9.
ANNUAL BENEFIT FUNNEL OF DOUBT FOR BLOCK BOOTSTRAPPED AND NORMAL RETURNS.



This figure reports funnels of doubt for the annual benefits. We show the median (solid line) as well as 5th and 95th quantiles (dotted lines). We do so for two different types of risky asset returns. First, we assume returns coming from the block bootstrap method as explained in Appendix C based on S&P/TSX Composite returns. Second, we assume a normal distribution with mean of 7% and a standard deviation of 15%, as done in Sections 4.1, 4.2, and 4.3. The hurdle rate is set to 4.5%. Half of the portfolio is invested in the risky asset, and the other half is invested in the risk-free asset. We assume that the retiree joins the pool at 65 with \$143,410, so the current benefit is set to \$10,000.


Figure 9 mimics the left panels of Figures 6 and 7: it gives the annual benefit funnel of doubt obtained with the block bootstrap method and that obtained with independent and identically distributed (iid) normal random variables (i.e., the method explained in Section 4.1). Interestingly, there is a slight difference between the two random benefit streams: the normal assumption leads to a wider funnel of doubt than that obtained by bootstrapping. This result is somewhat counterintuitive as one would expect realized returns to have fatter tails than those implied by normality (see, e.g., Figure 15 in Appendix C).

This intuition would be accurate if the returns were truly iid; however, annual S&P/TSX Composite returns exhibit negative autocorrelation (about -20%), meaning that positive annual returns tend to be followed by negative annual returns, and vice versa. This negative autocorrelation, when aggregated, signifies less risk in the benefit stream and thus narrower funnels of doubt. Note, however, that the difference between the 5th quantiles reported in Figure 9—block bootstrap- and normal distribution-based—is tiny (especially when compared to the difference between the 95th quantiles), meaning that this change in distribution should not have a material impact on our BaR measures.

³⁷ If continuously compounded returns are normally distributed, then the (gross) effective annual rates are lognormally distributed.

³⁸ Minimum and average BaRs are obtained by evaluating empirically Equations (2) and (3) based on the 100,000 benefit paths generated by Monte Carlo simulation. The benefit paths are obtained by applying Equation (7) to each simulated risky asset path.

Figure 10.
BENEFIT AT RISK FOR BLOCK BOOTSTRAPPED RETURNS.



This figure reports the minimum benefit at risk (blue line) and average benefit at risk (green line) as a function of the horizon. The mBaR level is set to 97.5%, the aBaR level to 90%, and the hurdle rate to 4.5%. Half of the portfolio is invested in the risky asset, and the other half is invested in the risk-free asset. We assume that the retiree joins the pool at 65 with \$143,410, so the current benefit is set to \$10,000. The risky asset returns are generated using the block bootstrap method explained in Appendix C.

Figure 10 reports the minimum and average BaRs obtained from bootstrapped returns; these values are virtually identical to those presented in Figures 4 and 5. This implies that the results obtained under normality in Sections 4.1 to 4.3 are robust to the risky asset return distribution assumption and that the analytical expressions derived in Appendix B and Section 4.1 could still act as legitimate and convenient guidelines in realistic contexts.

4.5 IMPACT OF IDIOSYNCRATIC MORTALITY

In Section 3.4, we assumed no mortality risk and derived benefit update equations based on this assumption. We now wish to assess the impact of this assumption by relaxing it. To do so and for simplicity's sake, we consider only idiosyncratic mortality risk in this study.³⁹ If we assume that all members have the same age and put in the same initial asset amount at first—the assumptions used so far—we have the following benefit update rule from Equation (6):

$$B(t) = B(t-1) \frac{p_{x+t-1}}{p_{x+t-1}^*} \exp(r_t^{\text{PF}} - h), \quad (8)$$

where $p_{x+t-1}^* = \frac{\sum_{k \in \mathcal{L}_t} B(t-1)}{\sum_{k \in \mathcal{L}_{t-1}} B(t-1)} = \frac{n(\mathcal{L}_t)}{n(\mathcal{L}_{t-1})}$ and $n(A)$ is the number of items in set A (i.e., cardinality). To account for idiosyncratic mortality, we therefore need to keep track of the members in \mathcal{L}_t for $t \in \{0, 1, \dots\}$. Assume a pool of N members so that $\mathcal{L}_0 = \{1, 2, \dots, N\}$. Using recursion, we can then define \mathcal{L}_t from \mathcal{L}_{t-1} by keeping track of the decedents and the survivors. For instance, if the k^{th} member is alive at time $t-1$ (i.e., $k \in \mathcal{L}_{t-1}$), then there are two possibilities for their status at time t :

1. They will survive until at least time t with probability p_{x+t-1} ; in this case, $k \in \mathcal{L}_t$.
2. They will die between $t-1$ and t with probability $q_{x+t-1} = 1 - p_{x+t-1}$; in this case, $k \notin \mathcal{L}_t$.

This behaviour is reminiscent of a Bernoulli random variable. Indeed, if $k \in \mathcal{L}_{t-1}$, then

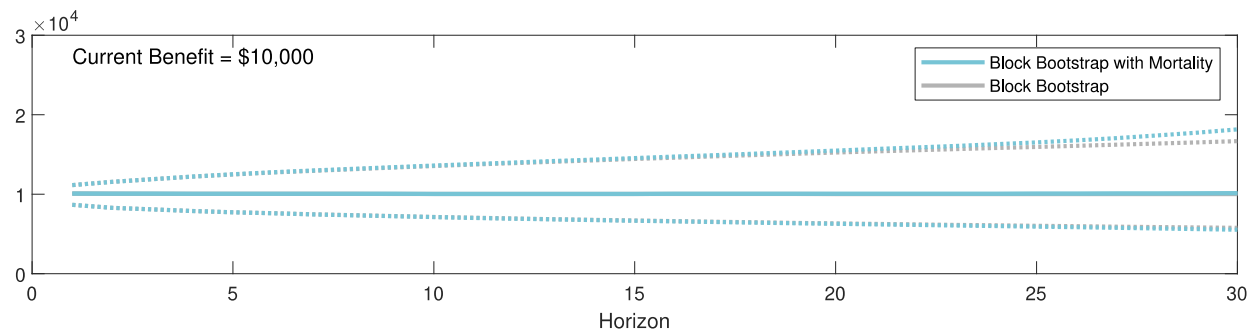
³⁹ Systematic mortality risk is also relevant but out of the scope of this report. To account for this risk, users must rely on deterministic generational life tables or stochastic mortality models allowing for longevity improvements (see Lee and Carter, 1992; Cairns et al., 2006, for examples of such mortality models).

$$\begin{cases} k \in \mathcal{L}_t & \text{with probability } p_{x+t-1} \\ k \notin \mathcal{L}_t & \text{with probability } 1 - p_{x+t-1} \end{cases}$$

where success (first outcome) is survival and failure (second outcome) is death. Using this rationale, we generate 100,000 mortality scenarios to capture the mortality uncertainty and use them in concert with Equation (8) to obtain benefit stream scenarios, similar to the process explained in Footnote 38.⁴⁰

Figure 11 reports funnels of doubt for the annual benefits (i.e., median as well as 5th and 95th quantiles of the benefit distribution) for two different mortality assumptions and for bootstrapped risky asset returns. First, we assume idiosyncratic mortality risk as explained above and obtain benefits by using Equation (8). Second, we assume no mortality risk, as done in Sections 4.1, 4.2, 4.3, and 4.4 and presented in Equation (7). We assume a very small group of 100 members with initial age of 65 in this illustration. The medians for both cases are virtually identical and stay at \$10,000 in both cases (i.e., the current benefit). Yet the benefit stream allows for more variability when idiosyncratic mortality is present, especially at older ages (i.e., after about 20 years after inception); this is expected as older members have higher death probabilities, which brings more variability in the benefits. Not surprisingly, the 5th quantiles are somewhat similar, meaning that low benefits are not drastically different when idiosyncratic mortality is accounted for—and this holds even at older ages.

Figure 11.
ANNUAL BENEFIT FUNNEL OF DOUBT FOR BLOCK BOOTSTRAPPED RETURNS WITH IDIOSYNCRATIC MORTALITY.



This figure reports funnels of doubt for the annual benefits. We show the median (solid line) as well as 5th and 95th quantiles (dotted lines). We do so for two different types of mortality assumption. First, we assume idiosyncratic mortality risk as explained in Section 4.5. Second, we assume no mortality risk. The hurdle rate is set to 4.5%. Half of the portfolio is invested in the risky asset and the other half in the risk-free asset. We assume that the retiree joins the pool at 65 with \$143,410, so that the current benefit is set to \$10,000. The risky asset returns are generated using the block bootstrap method. The current pool size is set to 100 members.

Figure 12 shows the minimum and average benefit at risk for different horizons, again using 100 members at inception and a current benefit of \$10,000. Interestingly, this plot is almost identical to Figure 10; the mBaR(5) estimate is \$2,943 when accounting for mortality, which is only \$15 higher than the mBaR(5) presented in Figure 10 (i.e., without mortality). The aBaR(20) shows a similar behaviour: it is \$2,028 when mortality is present in the model and \$2,002 when it is not. These observations imply that idiosyncratic mortality risk is a second-order consideration in lifetime pension pools for ages below 95 years old, consistent with the conclusions of Sabin and Forman (2016) and the conjecture of Section 3.4.

⁴⁰ Generating Bernoulli random variables with probability p is very easy. In fact, one only needs a (continuous) uniform random number over $[0,1]$, let us say U ; then, success is obtained if $U \leq p$, and failure if $U > p$.


Table 2.
BENEFIT AT RISK AS A FUNCTION OF THE INITIAL POPULATION SIZE AND THE CURRENT AGE OF THE MEMBERS.

Initial Population Size	mBaR(5)				aBaR(20)
	At 65	At 75	At 85	At 95	At 65
10	2,992	3,101	3,572	5,099	2,258
50	2,949	2,972	3,101	3,734	2,051
100	2,943	2,956	3,030	3,402	2,028
250	2,933	2,936	2,971	3,132	2,011
500	2,932	2,937	2,955	3,042	2,009
1,000	2,931	2,935	2,945	2,986	2,005
∞ (No Mortality)	2,928	2,928	2,928	2,928	2,002

This table reports the minimum and average benefit at risk as a function of the initial population size. The mBaR level is set to 97.5%, the aBaR level to 90%, and the hurdle rate to 4.5%. Half of the portfolio is invested in the risky asset, and the other half is invested in the risk-free asset. We assume that the current benefit is set to \$10,000. This table assumes different starting ages for the members when computing mBaR(5): 65, 75, 85, and 95 years old.

Based on the same assumptions above, Table 2 supplements the information in Figure 12. Specifically, it presents mBaR(5) at 65, 75, 85, and 95 years old, and aBaR(20) at 65 years old for different initial population sizes and a current benefit of \$10,000. For members aged 65 and 75 years old, the population size does not seem to impact the mBaR(5) estimates: for large cohorts, the minimum BaRs are the same as that computed without mortality risk, and for small cohorts, the increase in the mBaR(5) measure is tiny. As we consider older members, the size of the pool matters more: for 95-year-old members, the BaR almost doubles for pools of size 10. As expected, mBaR(5) becomes smaller when the size increases. We conjecture that this gets worse as members age past 95 years—something that is not outside the realm of possibility given current mortality improvement trends. Indeed, maintaining large open pools is paramount to diversifying idiosyncratic mortality risk.⁴¹

Figure 12.
BENEFIT AT RISK FOR BLOCK BOOTSTRAPPED RETURNS WITH IDIOSYNCRATIC MORTALITY.



This figure reports the minimum (blue line) and average benefit at risk (green line) as a function of the horizon. The mBaR level is set to 97.5%, the aBaR level to 90%, and the hurdle rate to the expected portfolio return; that is, 4.5%. Half of the portfolio is invested in the risky asset, and the other half is invested in the risk-free asset. We assume that the retiree joins the pool at 65 with \$143,410, so the current benefit is set to \$10,000. This figure assumes idiosyncratic mortality risk.

The impact of including mortality for older ages and small (closed) pools is palpable. Note that, in closed membership groups, having a small group is a matter of time: at some point, the group will become smaller as members die. Perhaps a better way to handle small groups is by creating open membership groups—allowing members of different

⁴¹ Bernhardt and Donnelly (2021) investigate the question of membership size and income stability in a recent paper. They also provide guidance in terms of pool size.

ages to join the plan every year. This comes with additional challenges (e.g., fairness and equity among generations, more complicated benefit update rules). Yet, this would solve the issues related to idiosyncratic mortality.

Another important aspect related to mortality that one should not lose track of is the impact of systematic mortality risk. Even though this risk has not been investigated formally in the report, it could significantly impact members' benefits and BaR measures. As mentioned in Footnote 39, a proper investigation of systematic mortality risk requires a model that accounts for generational mortality.⁴²

⁴² Potential adverse selection issues related to anti-selection could also generate additional systematic mortality risk. For more details about adverse selection in lifetime pension pools, see Valdez et al. (2006).

Section 5: Concluding Remarks and Further Developments


Daykin noted in 2004 that “the future probably lies in the development of different forms of risk-sharing between pensioners and annuity providers” (Daykin, 2004), meaning that future annuity products and arrangements will likely transfer more flexibility and risk to members than has traditionally been the case. Over the last two decades, we witnessed numerous pushes in this direction with the introduction of new legislation permitting lifetime pension pools, their launch in the marketplace, and the publication of several research papers on the topic.

As these new pools become a reality, actuaries and plan sponsors need to be able to understand the risks within these pools. They also need to communicate risk to members in a meaningful way, so that members can make appropriate decisions about retirement income. This is very challenging in practice as (1) members tend to be financially illiterate (see, e.g., Lusardi and Mitchell, 2011), and (2) actuaries struggle with communication when dealing with non-expert audiences (see, e.g., SOA, 2002).

This report addressed part of this issue by proposing a new collection of measures for member communication and disclosure. Specifically, we suggested two meaningful applications of the benefit at risk concept: the minimum BaR for budgeting purposes and the average BaR for decision-making purposes. We also assessed their behaviour when inputs used in its calculation are changed (e.g., horizon, hurdle rate, asset allocation strategy, risky asset return distribution, and idiosyncratic mortality). The measures were tested in the context of lifetime pension pools; however, these BaR measures could also be used in other contexts where benefits are unknown and fluctuating. This includes drawdown strategies for individual accounts and non-indexed or escalating annuities that do not track inflation. We leave these applications for future research.

The present report is a first attempt at communicating risk-related information to members; there is still a lot of work to be done to make these pools accessible to the general public and ensure their success.

We noted in the report some limitations of the proposed measures. Notably, using BaR in other contexts than communication and disclosure might be misguided as it does not tell the whole story. It does not mean, however, that it cannot be used in concert with other metrics—capturing other risk dimensions relevant to members—to identify optimal pool design features (e.g., the hurdle rate policy and the smoothing provisions).^{43,44} We leave this interesting question for future research.



⁴³ In the report, we considered a constant hurdle rate, yet it would be possible to consider a variable hurdle rate in practice, which could be linked to either nominal or real return expectations. On the one hand, adjustments to the hurdle rate may introduce additional year-to-year volatility in benefits. On the other hand, not adjusting the hurdle rate when return expectations change may diminish fairness between cohorts.

⁴⁴ Smoothing refers to the deferred recognition of gains and losses in the benefit payouts. Delayed recognition could give participants a longer time to adjust their consumption to changes in pension income. Nonetheless, it can also introduce value transfers among members (see, e.g., Cui et al., 2011; Kocken, 2012). Smoothing may also affect the ability of the pool to attract new participants.

Acknowledgments

The researchers' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the Project Oversight Group for their diligent work overseeing, reviewing, and editing this report for accuracy and relevance.

Project Oversight Group members:

Gavin S. Benjamin, FSA, FCIA
David R. Cantor, ASA, EA, CFA, FRM
Jean-Pierre Canuel, FSA, FCIA
Douglas P. Chandler, FSA, FCIA
Stephen A. Eadie, FSA, FCIA
Richard K. Fullmer, CFA, MSc
Hrvoje Lakota, FSA, FCIA, CFA
Dany Pineault, FSA, FCIA

At the Society of Actuaries Research Institute:

Steven C. Siegel, ASA, MAAA, SOA Research Actuary
Barbara Scott, Senior Research Administrator

At the Canadian Institute of Actuaries:

Esteban Rivera, Project Manager, Research and Professional Practice

The authors acknowledge the research support of Nikhil Kapoor.

Appendix A: Proofs

A.1 COMPARATOR MINUS MINIMUM BENEFIT DISTRIBUTION

Let us assume the comparator in the context of the minimum BaR is the current benefit $B(0)$, so that the random variable of interest is $B(0) - \underline{B}(\tau)$.

Starting from

$$\underline{B}(\tau) = \min_{t \in \{1, \dots, \tau\}} B(t),$$

we have that

$$\begin{aligned} F_{B(0) - \underline{B}(\tau)}(b) &= \mathbb{P}(B(0) - \underline{B}(\tau) \leq b) = \mathbb{P}(\underline{B}(\tau) \geq B(0) - b) \\ &= \mathbb{P}((B(1) \geq B(0) - b) \wedge \dots \wedge (B(\tau) \geq B(0) - b)), \end{aligned}$$

because if the minimum is larger than $B(0) - b$, each benefit needs to be larger than $B(0) - b$. By replacing the benefits by their definition, this could be rewritten as

$$\begin{aligned} &\mathbb{P}((B(1) \geq B(0) - b) \wedge \dots \wedge (B(\tau) > B(0) - b)) \\ &= \mathbb{P}\left(\left(r_1^{\text{PF}} \geq \log(B(0) - b) - \log(B(0)) + h\right) \wedge \left(r_2^{\text{PF}} \geq \log(B(0) - b) - \log(B(0)) + 2h - r_1^{\text{PF}}\right)\right. \\ &\quad \left. \wedge \dots \wedge \left(r_\tau^{\text{PF}} \geq \log(B(0) - b) - \log(B(0)) + \tau h - r_1^{\text{PF}} - \dots - r_{\tau-1}^{\text{PF}}\right)\right). \end{aligned}$$

This probability calculation could be expressed as a τ -dimensional integral:

$$\int_{K_1}^{\infty} \int_{K_2 - r_1^{\text{PF}}}^{\infty} \dots \int_{K_\tau - r_1^{\text{PF}} - \dots - r_{\tau-1}^{\text{PF}}}^{\infty} \left(\prod_{t=1}^{\tau} \phi(r_t^{\text{PF}}; \nu, \zeta) \right) dr_\tau^{\text{PF}} \dots dr_2^{\text{PF}} dr_1^{\text{PF}},$$

where $\phi(x; m, s)$ is the probability density function of a normal random variable with mean m and standard deviation s evaluated at x . Moreover, $K_t = \log(B(0) - b) - \log(B(0)) + th$.

We can numerically evaluate this expression via quadrature or Monte Carlo simulation.

A.2 COMPARATOR MINUS AVERAGE BENEFIT DISTRIBUTION

Let us rewrite the average benefit as

$$\begin{aligned} \bar{B}(\tau) &= \frac{B(0)}{\tau} \left(\sum_{s=1}^{\tau} \exp\left(\sum_{u=1}^s r_u^{\text{PF}} - sh\right) \right) = \frac{B(0)}{\tau} \sum_{s=1}^{\tau} \exp\left(\sum_{u=1}^s \tilde{r}_u^{\text{PF}}\right) \\ &= \frac{B(0)}{\tau} \left(e^{\tilde{r}_1^{\text{PF}}} \left(1 + e^{\tilde{r}_2^{\text{PF}}} \left(1 + e^{\tilde{r}_3^{\text{PF}}} \left(1 + e^{\tilde{r}_4^{\text{PF}}} (1 + \dots) \right) \right) \right) \right), \end{aligned}$$

where, in this case, each return in excess of the hurdle rate, $\tilde{r}_u^{\text{PF}} = r_u^{\text{PF}} - h$, is normally distributed with a mean of $\nu - h$ and standard deviation ζ . They are also independent from one another.

The expected value of $\bar{B}(\tau)$ can be found recursively because of independence. Starting from $Y_{\tau+1} = 1$, we can define $Y_s = 1 + e^{\tilde{r}_s^{\text{PF}}} Y_{s+1}$ for $s = 2, \dots, \tau - 1$. Then, the expected value of Y_s is

$$\mathbb{E}[Y_s] = \mathbb{E}\left[1 + e^{\tau_s^{\text{PF}}} Y_{s+1}\right] = 1 + e^{(v-h) + \frac{\zeta^2}{2}} \mathbb{E}[Y_{s+1}].$$

The expected value of $\bar{B}(\tau)$ is thus given by

$$\mathbb{E}[\bar{B}(\tau)] = \frac{B(0)}{\tau} \mathbb{E}[Y_1], \quad \text{where } \mathbb{E}[Y_1] = e^{(v-h) + \frac{\zeta^2}{2}} \mathbb{E}[Y_2].$$

The second uncentred moment is obtained in a similar way, via recursions:

$$E[Y_s^2] = \mathbb{E}\left[\left(1 + e^{\tau_s^{\text{PF}}} Y_{s+1}\right)^2\right] = 1 + 2e^{(v-h) + \frac{\zeta^2}{2}} \mathbb{E}[Y_{s+1}] + e^{2(v-h) + 2\zeta^2} \mathbb{E}[Y_{s+1}^2],$$

with

$$\mathbb{E}[\bar{B}(\tau)^2] = \left(\frac{B(0)}{\tau}\right)^2 \mathbb{E}[Y_1^2], \quad \text{where } \mathbb{E}[Y_1^2] = e^{2(v-h) + 2\zeta^2} \mathbb{E}[Y_2^2].$$

The variance is thus obtained as $\text{Var}[\bar{B}(\tau)] = \mathbb{E}[\bar{B}(\tau)^2] - \mathbb{E}[\bar{B}(\tau)]^2$, and the parameters $\mu_{\bar{B}(\tau)}$ and $\sigma_{\bar{B}(\tau)}$ are obtained via moment-matching, by solving the following equations:

$$\begin{aligned} \mathbb{E}[\bar{B}(\tau)] &= \exp\left(\mu_{\bar{B}(\tau)} + \frac{\sigma_{\bar{B}(\tau)}^2}{2}\right), \\ \text{Var}[\bar{B}(\tau)] &= (\exp(\sigma_{\bar{B}(\tau)}^2) - 1) \exp(2\mu_{\bar{B}(\tau)} + 2\sigma_{\bar{B}(\tau)}^2). \end{aligned}$$

Recall that for the average BaR, the comparator is the average benefit $\mathbb{E}[\bar{B}(\tau)]$. So, the cdf of interest is the following:

$$F_{\mathbb{E}[\bar{B}(\tau)] - \bar{B}(\tau)}(b) = \mathbb{P}(\mathbb{E}[\bar{B}(\tau)] - \bar{B}(\tau) \leq b) = \mathbb{P}(\bar{B}(\tau) \geq \mathbb{E}[\bar{B}(\tau)] - b) = 1 - F_{\bar{B}(\tau)}(\mathbb{E}[\bar{B}(\tau)] - b),$$

which can be rewritten as a function of the cdf of $\bar{B}(\tau)$. Note that the latter cdf is easy to compute as $\bar{B}(\tau)$ is assumed to be (approximately) lognormal.

Appendix B: Exogenous Comparators

B.1 REDEFINITION OF MINIMUM AND AVERAGE BENEFIT AT RISK WITH EXOGENOUS COMPARATORS

In this appendix, we consider a slightly different definition of the minimum and average BaRs that consider exogenous comparators set by the members. This alternative definition of the benefit at risk could be helpful if members have a target in mind or an external benchmark.

B.1.1 MINIMUM BENEFIT AT RISK WITH EXOGENOUS COMPARATOR

Instead of using the current benefit level $B(0)$ as the comparator, we employ an exogenous amount of K in this appendix. Assuming that the minimum benefit over a horizon of τ years is still given by $\underline{B}(\tau)$, the resulting worst shortfall over this horizon under a single scenario is $K - \underline{B}(\tau)$. For a horizon of 5 years—same as proposed in Section 2—we obtain the following measure:

$$\text{mBaR}_e(5) \equiv \text{BaR}_{97.5\%}[K - \underline{B}(5)] = F_{K - \underline{B}(5)}^{-1}(0.975),$$

if the probability level is kept at 97.5%.

To compute its value, we now need to find the cdf of the exogenous comparator, K , minus the minimum, $\underline{B}(5)$. Using the normality assumption of Section 4.1, we can find the cdf of this random variable in semi-closed form, up to an integral. This derivation is indeed very similar to that presented in Appendix A.1 and omitted in the interest of space.

B.1.2 AVERAGE BENEFIT AT RISK WITH EXOGENOUS COMPARATOR

Similar to the case above, we replace the average expected benefit, $\mathbb{E}[\bar{B}(\tau)]$, by an exogenous amount of K , yielding a new measure that is given as follows:

$$\text{aBaR}_e(20) \equiv \text{BaR}_{90\%}[K - \bar{B}(20)] = F_{K - \bar{B}(20)}^{-1}(0.9),$$

while again assuming a probability level of 90% and a horizon of 20 years. Fortunately, the equations developed in Section 4.1.2 can be adapted to account for this new comparator:

$$\text{aBaR}_e(20) = K - F_{\bar{B}(\tau)}^{-1}(0.10) = K - \exp\left(\mu_{\bar{B}(20)} + \sigma_{\bar{B}(20)} \Phi^{-1}(0.10)\right).$$

B.2 EXAMPLES WITH EXOGENOUS COMPARATORS

To illustrate the different BaR definitions of Appendix B.1, we rerun the tests of Sections 4.2.1 and 4.2.2 on the robustness of the hurdle rate and asset allocation, respectively. This should be enough for the reader to get a sense of the differences between the measures of Appendix B.1 and those of Section 2.

B.2.1 HURDLE RATE

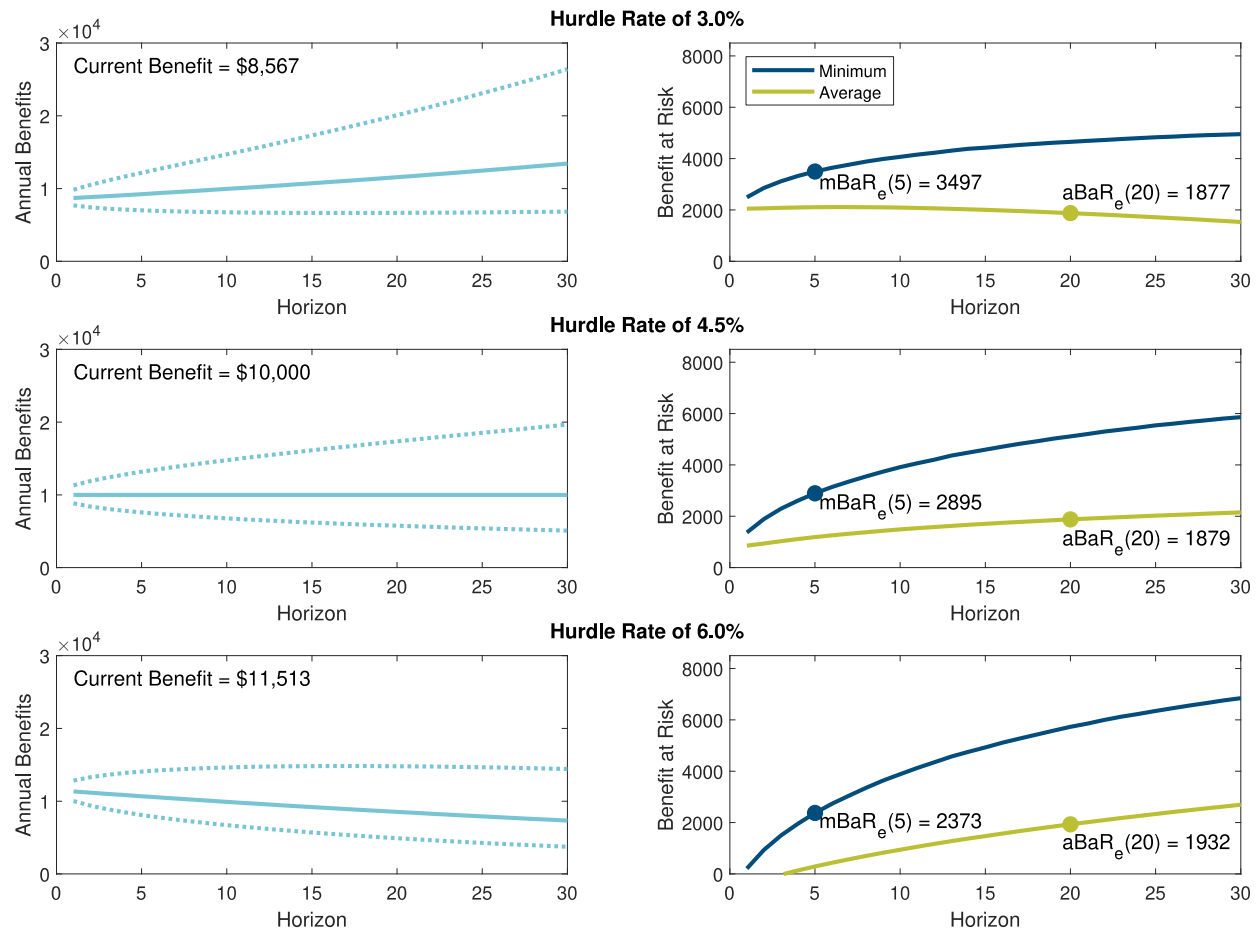
The left panels of Figure 13 report funnels of doubt (median as well as 5th and 95th quantiles) for the annual benefits; these panels are identical to the left panes of Figure 6.

The right panels of Figure 13 show minimum BaRs (blue lines) and average BaRs (green lines) with an exogenous comparator of \$10,000 and for different maturities. These two measures are again given for the three hurdle rates (i.e., 3.0%, 4.5%, and 6.0%). Interestingly, the relationship between the minimum BaR with an exogenous comparator of \$10,000 and the hurdle rate is the opposite of that reported in Figure 6: the measure reduces when the hurdle rate increases, and vice versa. This is a by-product of having smaller future benefits $B(t)$ when the hurdle rate is low,

leading to a lower minimum benefit over the next five years and a greater shortfall relative to the fixed comparator, K .

Figure 13.

ANNUAL BENEFIT FUNNEL OF DOUBT AND BENEFIT AT RISK WITH AN EXOGENOUS COMPARATOR OF \$10,000 AND FOR DIFFERENT HURDLE RATES.



The left panels of this figure report funnels of doubt for the annual benefits. We show the median (solid line) as well as 5th and 95th quantiles (dotted lines). We do so for three different hurdle rates: 3.0% (top panel), 4.5% (middle panel), and 6.0% (bottom panel). The right panels of this figure show the minimum BaR at the 97.5% level (blue lines) and the average benefit BaR at the 90% level (green lines), both using \$10,000 as an exogenous comparator. We assume a constant asset allocation strategy (i.e., $\omega = 0.5$).

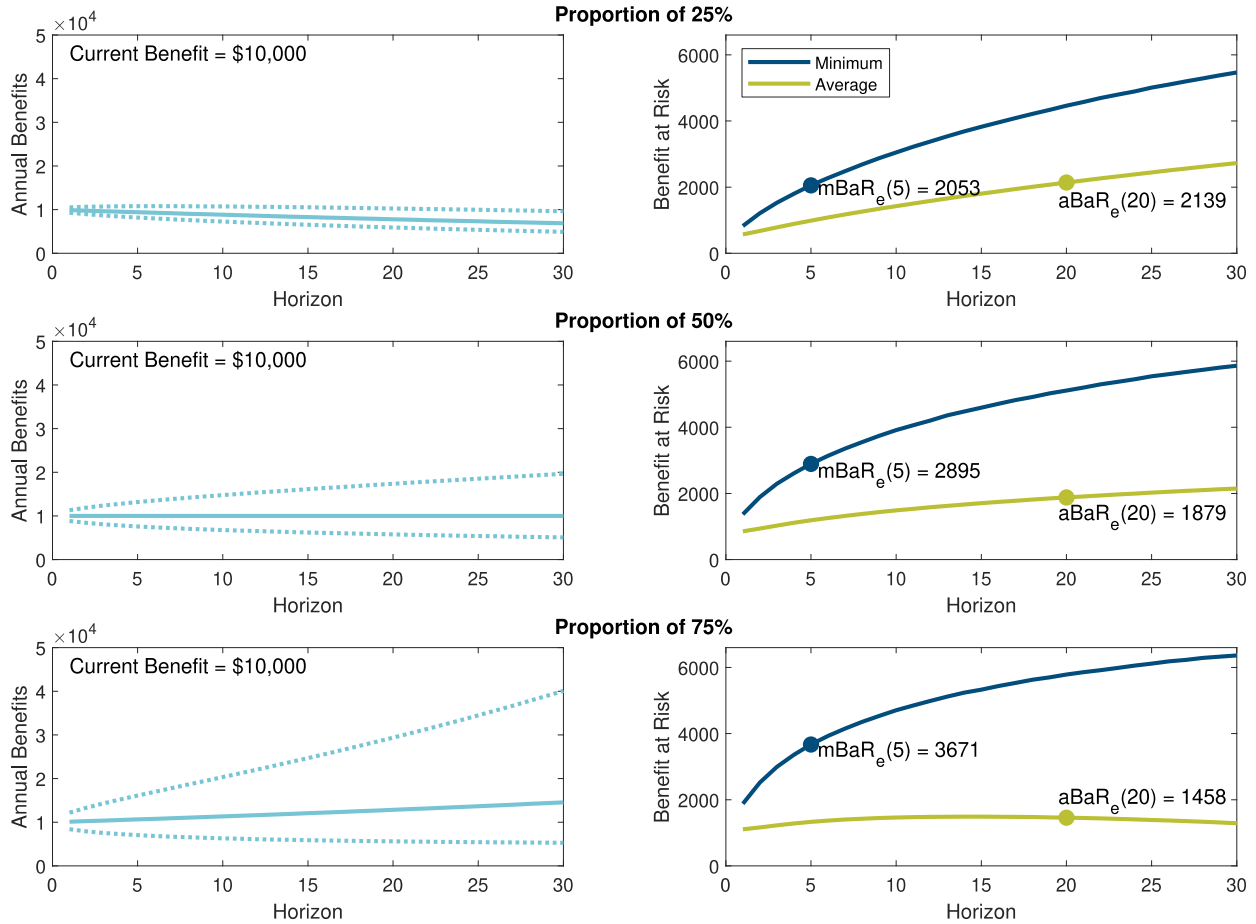
The average BaR with an exogenous comparator of \$10,000 also behaves differently than that presented in Figure 6; instead of decreasing as a function of the hurdle rate, the new BaR measure introduced in Appendix B.1.2 increases for larger hurdle rates. This does not come as a surprise: a higher hurdle rate implies a decreasing pattern of benefits, which in turn decreases the average benefit, making the difference between the exogenous comparator K and the benefit statistic—the average benefit in this case—larger. For a lower hurdle rate, it is the opposite: an increasing benefit pattern yields more benefits in the future, a larger average, and a smaller benefit at risk.

B.2.2 ASSET ALLOCATION

We now turn to asset allocation and its impact on the benefit at risk in the case where the comparator is set to an exogenous amount of K . Again, as in Section 4.2.2, we change the proportion ω invested in the risky asset (which is assumed to be a stock index in this case), while keeping all other inputs (including the hurdle rate) constant.

Figure 14.

ANNUAL BENEFIT FUNNEL OF DOUBT AND BENEFIT AT RISK WITH AN EXOGENOUS COMPARATOR OF \$10,000 AND FOR DIFFERENT ASSET ALLOCATION STRATEGIES.



The left panels of this figure report funnels of doubt for the annual benefits. We show the median (solid line) as well as 5th and 95th quantiles (dotted lines). We do so for three different asset allocation strategies (i.e., the proportion of assets invested in risky assets): 25% (top panel), 50% (middle panel), and 75% (bottom panel). The right panels of this figure show the minimum BaR at the 97.5% level (blue lines) and the average benefit BaR at the 90% level (green lines), both using \$10,000 as an exogenous comparator. We assume a constant hurdle rate of 4.5% throughout this figure.

The left panels of Figure 14 report funnels of doubt for the annual benefits for three different proportions: 25% (top panel), 50% (middle panel), and 75% (bottom panel). These, again, are the same funnels of doubt already reported in Figure 7.

The right panels of Figure 14 show minimum and average BaRs with an exogenous comparator of \$10,000 over different horizons and for the three asset allocation strategies mentioned above. As expected, the minimum BaRs of Figure 7 are identical to those provided in Figure 14 (see blue lines). Indeed, the current benefit was \$10,000 for the three proportions in Figure 7, which is the same as the selected K in this example.

The average BaRs with an exogenous comparator, however, are different from the aBaR measures presented in Figure 7; their behaviour is reversed. For the aBaR defined in Section 2.4, the shortfall risk increases when the proportion invested in the risky asset increases. For the case of average BaRs with an exogenous comparator of \$10,000, the shortfall risk decreases as the proportion invested in the risky asset is increased. Indeed, as the average benefit tends

to increase when the proportion invested in the risky asset is higher, the difference between the comparator of \$10,000 and the average benefit is expected to be reduced.

Appendix C: Block Bootstrapping

This appendix explains the block bootstrap methodology used in this report. Specifically, we follow the approach put forward in Portfolio Solution Group (2021). We use realized excess returns (computed as the monthly total return on the asset minus the monthly short rate, both taken at the same time), to which we add the (monthly) risk-free rate assumed in the modelling for consistency's sake. The latter risk-free rate is assumed to be 2% on an annual basis throughout this report (see Section 4).

Unlike common practice that uses individual monthly returns, we rely on multi-month blocks. The blocks' lengths are chosen randomly from an exponential distribution with a mean of 24 months in this application. For each path, we:


1. Select randomly a starting month between January 1990 and July 2021 using a uniform distribution.
2. Generate a block length using an exponential distribution.
3. Copy the appropriate block of excess returns.
4. Repeat until the resulting path is of the desired length.

This model-free, non-parametric approach preserves some of the real-world behaviour, such as fat tails, skewness, autocorrelation of returns and squared returns, and time-varying correlations.

Figure 15 reports an example of the block bootstrap distribution obtained by using S&P/TSX Composite excess returns (in blue). For comparison, the observed distribution obtained from S&P/TSX Composite excess returns (in green) and a moment-matched normal density (in light blue) are also shown in Figure 15. The observed distribution is in line with that obtained with the block bootstrap approach. The moment-matched normal is far from that obtained with block bootstrap, especially in the left tails.

Figure 15.

ANNUALIZED EXCESS RISKY ASSET RETURNS BASED ON THE BLOCK BOOTSTRAP METHOD.



This figure presents the observed distribution of the S&P/TSX Composite (continuously compounded) excess returns (green histogram) as well as those obtained using the block bootstrap methodology (blue line) and a moment-matched normal distribution with a mean of 5% and a standard deviation of 15% as in Section 4.1 (light blue line).

References

- Ahlgrim, K. C., D'Arcy, S. P., and Gorvett, R. W. (2005). Modeling financial scenarios: A framework for the actuarial profession. *Proceedings of the Casualty Actuarial Society*, 92(177): 177–238.
- ACPM. (2017). *Decumulation, the Next Critical Frontier: Improvements for Defined Contribution and Capital Accumulation Plans*. Association of Canadian Pension Management.
- Artzner, P., Delbaen, F., Eber, J. M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228.
- Australian Government Actuary. (2014). *Towards More Efficient Retirement Income Products*. Australian Government Actuary.
- Basak, S., and Shapiro, A. (2001). Value-at-risk-based risk management: Optimal policies and asset prices. *Review of Financial Studies*, 14(2), 371–405.
- Bégin, J. F. (2021). On complex economic scenario generators: Is less more? *ASTIN Bulletin*, 51(3), 779–812.
- Bernhardt, T., and Donnelly, C. (2019). Modern tontine with bequest: Innovation in pooled annuity products. *Insurance: Mathematics and Economics*, 86, 168–188.
- Bernhardt, T., and Donnelly, C. (2021). Quantifying the trade-off between income stability and the number of members in a pooled annuity fund. *ASTIN Bulletin*, 51(1), 101–130.
- Blanchett, D. (2014). Exploring the retirement consumption puzzle. *Journal of Financial Planning*, 27(5), 34–42.
- Boyle, P., Hardy, M., Mackay, A., and Saunders, D. (2015). *Variable Payout Annuities*. Society of Actuaries.
- Cairns, A. J., Blake, D., and Dowd, K. (2006). A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *Journal of Risk and Insurance*, 73(4), 687–718.
- Chen, A., and Rach, M. (2019). Options on tontines: An innovative way of combining tontines and annuities. *Insurance: Mathematics and Economics*, 89, 182–192.
- Chen, A., Hieber, P., and Klein, J. K. (2019). Tonuity: A novel individual-oriented retirement plan. *ASTIN Bulletin*, 49(1), 5–30.
- Chen, A., Rach, M., and Sehner, T. (2020). On the optimal combination of annuities and tontines. *ASTIN Bulletin*, 50(1), 95–129.
- Chen, A., Qian, L., and Yang, Z. (2021). Tontines with mixed cohorts. *Scandinavian Actuarial Journal*, 2021(5), 437–455.
- Chen, A., Nguyen, T., and Sehner, T. (2022). Unit-linked tontine: Utility-based design, pricing and performance. *Risks*, 10(4), 78.
- CREF. (2022a). *Prospectus*. College Retirement Equities Fund. Available online at <http://connect.rightprospectus.com/TIAA/TADF/194408803/S?site=VA&fundclass=RPVA>. Accessed on June 27, 2022.
- CREF. (2022b). *Statement of additional information*. College Retirement Equities Fund. Available online at <http://connect.rightprospectus.com/TIAA/TADF/194408803/S?site=VA&fundclass=RPVA>. Accessed on June 27, 2022.
- Cui, J., De Jong, F., and Ponds, E. (2011). Intergenerational risk sharing within funded pension schemes. *Journal of Pension Economics & Finance*, 10(1), 1–29.
- Daykin, C. (2004). *Annuities and alternative ways of providing retirement income*. Presented to the Pensions, Benefits and Social Security Seminar Part of the IACA, PBSS and IAAust Colloquium.
- Donnelly, C., Guillén, M., and Nielsen, J. P. (2013). Exchanging uncertain mortality for a cost. *Insurance: Mathematics and Economics*, 52(1), 65–76.
- Donnelly, C., Guillén, M., and Nielsen, J. P. (2014). Bringing cost transparency to the life annuity market. *Insurance: Mathematics and Economics*, 56, 14–27.
- Donnelly, C. (2015). Actuarial fairness and solidarity in pooled annuity funds. *ASTIN Bulletin*, 45(1), 49–74.
- Dufresne, D (2008). Sums of lognormals. *Proceedings of the 43rd Actuarial Research Conference at the University of Regina*.
- Dhaene, J., Goovaerts, M., and Van Weert, K. (2008). Some comments on QIS3. *Zavarovalniski Horizonti*, 3, 73–87.

- Fenton, L. (1960). The sum of lognormal probability distributions in scatter transmission systems. *IRE Transactions on Communications Systems*, 8(1), 57–67.
- Forman, J. B., and Sabin, M. J. (2015). Tontine pensions. *University of Pennsylvania Law Review*, 163(3), 755.
- Fullmer, R. K. (2019). *Tontines: A Practitioner's Guide to Mortality-Pooled Investments*. CFA Institute Research Foundation.
- Fullmer, R. K., and Sabin, M. J. (2019). Individual tontine accounts. *Journal of Accounting and Finance*, 19(8), 2158–3625.
- Fullmer, R. K., and Forman, J. B. (2022). State-sponsored pensions for private-sector workers. In Mitchell, O. (Ed.), *New Models for Managing Longevity Risk: Public-Private Partnerships* (pp. 171–206). Oxford University Press: Oxford, UK.
- Greenwald Research. (2022). *2021 Retirement Risk Survey, Report of Findings*. Society of Actuaries.
- Hanewald, K., Piggott, J., and Sherris, M. (2013). Individual post-retirement longevity risk management under systematic mortality risk. *Insurance: Mathematics and Economics*, 52(1), 87–97.
- Horneff, W. J., Maurer, R. H., Mitchell, O. S., and Stamos, M. Z. (2010). Variable payout annuities and dynamic portfolio choice in retirement. *Journal of Pension Economics & Finance*, 9(2), 163–183.
- Holton, G.A. (2002). History of value-at-risk. *Working Paper*.
- Iwry, J. M., Haldeman, C., Gale, W., and John, D. (2020). *Retirement Tontines: Using a Classical Finance Mechanism as an Alternative Source of Retirement Income*. The Brookings Institution.
- Jorion, P. (1997). In defense of VaR. *Derivatives Strategy*, 2(4), 20–23.
- Kocken, T. P. (2012). Pension liability measurement and intergenerational fairness: Two case studies. *Rotman International Journal of Pension Management*, 5(1):16–24.
- Lee, R. D., and Carter, L. R. (1992). Modeling and forecasting US mortality. *Journal of the American Statistical Association*, 87(419), 659–671.
- Lusardi, A., and Mitchell, O. S. (2011). Financial literacy around the world: An overview. *Journal of Pension Economics & Finance*, 10(4), 497–508.
- MacDonald, B. J., Jones, B., Morrison, R. J., Brown, R. L., and Hardy, M. (2013). Research and reality: A literature review on drawing down retirement financial savings. *North American Actuarial Journal*, 17(3), 181–215.
- Milevsky, M. A. (2015). *King William's Tontine: Why the Retirement Annuity of the Future Should Resemble its Past*. Cambridge University Press.
- Milevsky, M. A., and Salisbury, T. S. (2015). Optimal retirement income tontines. *Insurance: Mathematics and Economics*, 64, 91–105.
- Milevsky, M. A., and Salisbury, T. S. (2016). Equitable retirement income tontines: Mixing cohorts without discriminating. *ASTIN Bulletin*, 46(3), 571–604.
- OECD. (2022). *Recommendation of the council for the good design of defined contribution pension plans*. Organization for Economic Cooperation and Development. Available online at <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0467>. Accessed on June 30, 2022.
- Olivieri, A., and Pitacco, E. (2020). Linking annuity benefits to the longevity experience: Alternative solutions. *Annals of Actuarial Science*, 14(2), 316–337.
- Piggott, J., Valdez, E. A., and Detzel, B. (2005). The simple analytics of a pooled annuity fund. *Journal of Risk and Insurance*, 72(3), 497–520.
- Pollak, R. A. (1970). Habit formation and dynamic demand functions. *Journal of Political Economy*, 78(4), 745–763.
- Portfolio Solution Group (2021). *What's the Worst That Should Happen?* Applied Quantitative Research.
- Qiao, C., and Sherris, M. (2013). Managing systematic mortality risk with group self-pooling and annuitization schemes. *Journal of Risk and Insurance*, 80(4), 949–974.
- Sabin, M. J. (2010). Fair tontine annuity. *Working Paper*.
- Sabin, M. J., and Forman, J. B. (2016). The analytics of a single-period tontine. *Working Paper*.
- SOA. (2002). *Market Opportunities Research*. Society of Actuaries.
- Shemtob, M. (2021). CRITs: An alternative retirement income approach. *Retirement Section News of the Society of Actuaries*, September 2021.

- Stemtob, M. (2022). Time to take the plunge. *Contingencies*. September/October 2022, 32–36.
- Stamos, M. Z. (2008). Optimal consumption and portfolio choice for pooled annuity funds. *Insurance: Mathematics and Economics*, 43(1), 56–68.
- Valdez, E. A., Piggott, J., and Wang, L. (2006). Demand and adverse selection in a pooled annuity fund. *Insurance: Mathematics and Economics*, 39(2), 251–266.
- Wilkie, A. D. (1986). A stochastic investment model for actuarial use. *Transactions of the Faculty of Actuaries*, 39, 341–403.
- Wilkie, A. D. (1995). More on a stochastic asset model for actuarial use. *British Actuarial Journal*, 1(5), 777–964.

About the Society of Actuaries Research Institute

Serving as the research arm of the Society of Actuaries (SOA), the SOA Research Institute provides objective, data-driven research bringing together tried-and-true practices and future-focused approaches to address societal challenges and your business needs. The Institute provides trusted knowledge, extensive experience, and new technologies to help effectively identify, predict, and manage risks.

Representing the thousands of actuaries who help conduct critical research, the SOA Research Institute provides clarity and solutions on risks and societal challenges. The Institute connects actuaries, academics, employers, the insurance industry, regulators, research partners, foundations and research institutions, sponsors, and non-governmental organizations, building an effective network that provides support, knowledge, and expertise regarding the management of risk to benefit the industry and the public.

Managed by experienced actuaries and research experts from a broad range of industries, the SOA Research Institute creates, funds, develops, and distributes research to elevate actuaries as leaders in measuring and managing risk. These efforts include studies, essay collections, webcasts, research papers, survey reports, and original research on topics impacting society.

Harnessing its peer-reviewed research, leading-edge technologies, new data tools, and innovative practices, the Institute seeks to understand the underlying causes of risk and the possible outcomes. The Institute develops objective research spanning a variety of topics with its [strategic research programs](#): aging and retirement; actuarial innovation and technology; mortality and longevity; diversity, equity, and inclusion; health care cost trends; and catastrophe and climate risk. The Institute has a large volume of [topical research available](#), including an expanding collection of international and market-specific research, experience studies, models, and timely research.

Society of Actuaries Research Institute
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.soa.org

About the Canadian Institute of Actuaries

The Canadian Institute of Actuaries (CIA) is the qualifying and governing body of the actuarial profession in Canada. We develop and uphold rigorous standards, share our risk management expertise, and advance actuarial science to improve lives in Canada and around the world. Our more than 6,000 members apply their knowledge of math, statistics, data analytics, and business in providing services and advice of the highest quality to help Canadian people and organizations face the future with confidence.

VISION

Financial security for all Canadians.

MISSION

As the qualifying and governing body of the actuarial profession in Canada, we develop and uphold rigorous standards, share our risk management expertise, and advance actuarial science for the financial well-being of society.

VALUES

We care about principles. Professionalism defines us. We are trustworthy and accountable, and we uphold strict ethical principles.

We care about people. We foster a diverse and inclusive environment. We put the public interest ahead of our own.

We care about the future. Foresight is our strength. We anticipate risk to evolve with our changing environment.

Canadian Institute of Actuaries
360 Albert Street, Suite 1740
Ottawa, Ontario K1R 7X7
www.cia-ica.ca