A combined analysis of hedge effectiveness and capital efficiency in longevity hedging

Matthias Börger

Technische Hochschule Nürnberg Georg Simon Ohm, Nürnberg

Keßlerplatz 12, 90489 Nürnberg, Germany

E-mail: matthias.boerger@th-nuernberg.de; Phone: +49 (911) 5880-1755

Arne Freimann (corresponding author)

Institute for Financial and Actuarial Sciences (ifa), Ulm & Institute of Insurance, Ulm University

Lise-Meitner-Straße 14, 89081 Ulm, Germany

E-mail: a.freimann@ifa-ulm.de; Phone: +49 (731) 20 644-253

Jochen Ruß

Institute for Financial and Actuarial Sciences (ifa), Ulm & Institute of Insurance, Ulm University

Lise-Meitner-Straße 14, 89081 Ulm, Germany

E-mail: j.russ@ifa-ulm.de; Phone: +49 (731) 20 644-233

September 11, 2020

Abstract

By hedging longevity exposures, annuity providers can reduce both the uncertainty in future cash flows and capital charges in a cost efficient manner. We argue that a separate analysis of these two aspects cannot provide a full picture of the implications of longevity hedging, in particular when using index-based instruments.

Hence, we propose a stochastic modeling framework for a joint analysis of the risk-reducing effect and the economic impact of longevity hedges in terms of *hedge effectiveness* and *capital efficiency*, respectively. In an economic capital model under Solvency II, a wide selection of customized and index-based instruments is analyzed. We show that different hedging objectives require different instruments on different index populations and discuss the accompanying trade-off between hedge effectiveness and capital efficiency. While customized hedges naturally outperform their index-based counterparts in terms of hedge effectiveness, we show that cost efficient index-based designs may be more capital efficient.

JEL classification: G22; G23; G32

Keywords: Hedging longevity risk; Hedge effectiveness, Capital efficiency, Solvency regimes, Population basis risk

1 Introduction

The risk that policyholders live longer than anticipated is commonly referred to as 'longevity risk' and poses a major risk for annuity providers and pension funds. One possible approach for dealing with this risk is longevity hedging. At present, the longevity risk transfer market is dominated by insurance-based deals with reinsurers acting as main players. However, in light of the potential size of the global longevity market, the consensus among practitioners appears to be that the limited capacity of the reinsurance sector cannot permanently meet the rising demand for de-risking solutions (cf. Blake et al. (2019)).

While insurance-based deals are typically 'customized', i.e. tailored individually to the hedger's liability characteristics, they might not be appealing to other capital market investors as they require knowledge of the mortality characteristics of the specific underlying population. This can be avoided by linking the hedge payout to transparent mortality indices of a more general reference population. These index-based contracts can be appealing to a wider range of investors and the hedger may also benefit from greater liquidity and lower risk transfer premiums due to a more competitive market. However, from the hedger's perspective, index-based instruments give rise to population basis risk since the mortality experience of the portfolio population might deviate from that captured by the mortality indices.

The assessment of population basis risk calls for joint mortality modeling of the population to be hedged and the instrument's reference population. To this end, several authors have proposed suitable two-population mortality models and we refer to Haberman et al. (2014) or Villegas et al. (2017) for an overview. In the presence of population basis risk, several authors have addressed the calibration of standardized index-based hedging strategies by either relying on some sort of sensitivity matching strategy (e.g. Cairns (2013), Li and Hardy (2011), Li and Luo (2012), Liu and Li (2016), or Zhou and Li (2017)) or on a risk minimization strategy (see e.g. Cairns et al. (2014) or Coughlan et al. (2011)). In either case, the effectiveness of a hedging strategy is typically evaluated prospectively in a model-based simulation approach and quantified by means of risk measures as the achieved level of risk reduction. If the simulation model properly accounts for the imperfect correlation between the evolutions of the hedger's population and the hedge indices, index-based solutions perform less favorably compared to fully customized contracts. However, all previously mentioned studies have in common that they mainly focus on the reduction in variance and do not consider potential cost advantages of standardized contracts. Liu and Li (2019), Ngai and Sherris (2011) and Zhou and Li (2019) overcome this shortcoming by demonstrating that high hedging costs can make a hedge economically unappealing when relying on asymmetric risk measures such as expected shortfall or Value-at-Risk (VaR).

Under modern risk-based solvency regimes, such as Solvency II or the Swiss Solvency Test, the implications of hedging extend even further. Insurers are required to back their longevity business with adequate economic capital and the expected costs for providing future economic capital have to be reserved as a risk margin in addition to the best estimate liabilities. Longevity hedging reduces capital charges and hence creates value which gives a strong economic incentive for de-risking. Obviously, the economic attractiveness of a longevity transaction depends on two opposing effects: the generated cost of capital saving and the costs of hedging. Meyricke and Sherris (2014) address this trade-off by demonstrating that Solvency II capital requirements generate incentives to transfer short-term longevity risk while retaining longevity tail risk since the hedging costs tend to exceed the capital relief over longer time horizons. However, their analysis is limited to customized longevity swaps. The assessment of regulatory capital relief in the presence of population basis risk has recently been addressed by Cairns and El Boukfaoui (2019). While index-based hedges naturally provide lower (or equal) capital reliefs due to population basis risk, they may on the other hand offer cost advantages compared to customized deals. Therefore, both the achieved level of risk reduction and the generated cost of capital relief net of hedging costs should be considered for a well-informed hedging decision.

This paper complements the literature by contributing a stochastic modeling framework for a joint analysis of the risk-reducing effect and the economic impact of longevity hedging. More specifically, we use the following metrics to quantify the impact of hedging on the annuity provider's cash flow profile: *hedge effectiveness* measures the reduction in the uncertainty related to unexpected cash flows and *capital efficiency* captures the change in expected cash flows, which arises from the generated cost of capital relief in one direction and from the costs of hedging in the other direction. In a numerical case study, we consider an annuity provider who evaluates a wide selection of different customized and index-based hedging instruments. In particular, we consider instruments which are specifically designed for the purpose of reducing longevity capital charges. In a model-based simulation approach, we determine and simultaneously visualize the hedge effectiveness and capital efficiency for all considered hedging solutions. This can help hedgers to avoid suboptimal decisions and, more importantly, to identify suitable contract designs with regard to the strategic hedging objective.

Such a comprehensive analysis requires an extensive stochastic modeling framework which accounts for all relevant components of longevity risk, the costs of hedging, and the impact of hedging on capital charges within a realistic regulatory framework. For the individual subcomponents of our framework we build on several concepts and modeling techniques which have been developed by different authors. Since a detailed discussion would go beyond the scope of this paper, we limit ourselves to a concise summary and refer the interested reader to the cited works. Our modeling framework consists of the following subcomponents:

- *Economic capital model under Solvency II*: To compute longevity Solvency Capital Requirements (SCRs), the hedger either implements a stochastic (partial) internal model or alternatively applies the Solvency II standard formula for longevity risk. We determine the impact of longevity hedging on capital charges, in particular in the presence of population basis risk, under both SCR approaches and compare the results. As pointed out by Börger (2010), a proper implementation of a risk-based internal model requires the simulation of all components of longevity risk over a one-year horizon and subsequent valuation of liabilities based on potentially revised mortality assumptions.

- *Stochastic mortality model*: For the computation of capital charges, the mortality model needs to consist of two components: First, it requires a simulation model that jointly captures all relevant components of longevity risk, namely long-term mortality trend risk of the overall population, potentially differing mortality characteristics of the hedger's portfolio population, and idiosyncratic risk arising from a portfolio of limited size. This clear distinction allows us to evaluate different hedging instruments which give rise to varying levels of residual basis risk. Second, a valuation model is needed for pathwise derivations of best estimate mortality assumptions at future valuation dates.[1] We follow

[1]Also Cairns (2013), Cairns et al. (2014), Cairns and El Boukfaoui (2019), and Coughlan et al. (2011) clearly distinguish between a simulation model and a valuation model. Alternative mortality models for the computation of longevity capital requirements have been proposed by Börger et al. (2014), Plat (2011), Richards et al. (2014, 2020), and Zeddouk and Devolder (2019).

Börger et al. (2019) and clearly distinguish between the unobservable actual mortality trend prevailing at a certain point in time and the mortality trend that an observer would estimate given the observed mortality evolution up to that point in time.

- *Pricing*: We assume that the counterparty demands a risk premium as compensation for taking longevity risk. We follow Freimann (2020) and apply a risk-adjusted pricing approach. For comparative studies of different pricing approaches for longevity-linked securities, we refer to Bauer et al. (2010) or Leung et al. (2018).

Our work contributes to the existing literature on longevity hedging in the following ways: First and foremost, we simultaneously quantify the risk-reducing effect and the economic impact of longevity hedges. To the best of our knowledge, all previous related studies are limited to one of these perspectives. We also contribute to the ongoing discussion on how to optimally engineer an index-based hedge given a specific strategic hedging objective. This is related to Liu and Li (2019) who propose a hedging strategy based on q-forwards for minimizing longevity VaR (but do not incorporate population basis risk) and Cairns and El Boukfaoui (2019) who address the optimal design of index-based contracts under economic capital perspectives in the presence of population basis risk. Extending their works, we additionally address the potential cost advantages of index-based designs and the accompanying conflict between mitigating (or even fully eliminating) population basis risk and higher hedging costs arising from customization. Second, we explicitly model the uncertainty in future capital charges by iteratively deriving longevity SCRs in line with Solvency II guidelines over subsequent one-year horizons. In contrast, previous works of Cairns and El Boukfaoui (2019), Liu and Li (2019), or Richards et al. (2020) focus on the SCR at the start of the simulation and on the calculation of capital relief at the time the hedge is put in place. We demonstrate that the uncertainty in future capital charges can be quite significant and can have a notable impact on longevity hedge effectiveness. Finally, this paper provides a profound comparison between capital reliefs derived under the Solvency II standard formula and under a risk-based (partial) internal model for longevity risk.

The remainder of the paper is organized as follows. In Section 2, we describe our model setup and provide a definition of hedge effectiveness and capital efficiency in the presence of uncertain future capital charges for longevity risk. In Section 3, we introduce the considered hedging instruments by specifying the underlying hedge payout structure and the population the instrument is linked to. Section 4 provides the numerical results. After a brief investigation of the hedger's initial unhedged situation, we first separately analyze the benefits of hedging, namely capital relief (as driver for capital efficiency) and risk reduction (as driver for hedge effectiveness). Then, we combine both aspects and address the trade-off between hedge effectiveness and capital efficiency which allows us to discuss the suitability of various hedge designs for different objectives. This is followed by a sensitivity analysis with respect to the modeling assumption for population basis risk and the structure of the underlying liabilities. Finally, Section 5 concludes.

2 Model setup

We now establish a setting for a combined analysis of hedge effectiveness and capital efficiency in longevity hedging. To focus on longevity risk, we assume that all market participants are fully hedged against interest rate and investment risk and only invest at the risk-free interest rate r, which is also used for discounting purposes. For the sake of simplicity, we assume that the hedger's book is closed to new business and we ignore any other source of uncertainty (such as operational risk or counterparty credit risk[2]). Hence, any variation in future cash flows is only due to changes in actual or expected mortality.

In Section 2.1, we start by describing our stochastic mortality modeling framework which accounts for future changes in actual mortality as well as for the accompanying changes in best estimate mortality assumptions. Afterwards, we introduce the liability to be hedged in Section 2.2 and describe the hedger's economic capital model in Section 2.3. Finally, in Section 2.4, we provide a formal definition of hedge effectiveness and capital efficiency.

2.1 Stochastic modeling framework for the components of longevity risk

In this section, we give an overview of our stochastic multi-population mortality model. Following Börger et al. (2019), we clearly distinguish between the unobservable actual mortality trend (AMT) prevailing at a certain point in time and the estimated mortality trend (EMT) that an observer would estimate given the observed mortality evolution up to that point in time. Hence, our model consists of two components: the so-called *AMT simulation model* for generating paths of future mortality and the so-called *EMT valuation model* for pathwise derivations of best estimate mortality assumptions at future valuation dates. We focus on the most important model features and refer to Appendix A.1, Börger et al. (2019), and Freimann (2020) for further technical details.

The hedger's book population is interpreted as a subset of a larger reference population R (typically the national population) with potentially differing mortality characteristics e.g. due to a specific socioeconomic structure or selection effects. Within the reference population, we consider N_{Sub} distinct and sufficiently large subpopulations of different socioeconomic status where all individuals within one subgroup are assumed to experience the same force of mortality. To model the individual mortality characteristics of the hedger's book population, we follow Haberman et al. (2014) and apply a so-called 'characterization approach' by assuming that the company has sufficient information to uniquely assign each policyholder to the appropriate subpopulation and that they can be identified with this subgroup throughout their lives.[3] This approach offers several advantages. Most importantly, it can be applied to portfolios of any size without mistaking unsystematic variations in the data arising from a small sample size for systematic mortality differentials. Also, the evolution of the socioeconomic book composition over time is adequately captured.

Within this setting, we model the random future evolution of the book population consistently to future mortality in the reference population and its socioeconomic subpopulations in a multi-population extension of the Cairns-Blake-Dowd (CBD) model, see Cairns et al. (2006). The AMT simulation model captures the following components of longevity risk:

- The long-term mortality trend risk of the overall population modeled via a stochastic process for future probabilities of death $q_{x,t}^{[R]}, t \geq 0$ for the reference population. To

[2]In reality, longevity hedges inevitably entail counterparty credit risk, which can be mitigated if not fully eliminated by collateralization. Biffis et al. (2016) find that the cost of posing collaterals for longevity swaps is comparable (and often even much smaller) than that observed in the interest rate swap market.

[3]Lu et al. (2014) find that migration between subpopulations in England does not significantly distort trends in socioeconomic mortality inequalities. For simplicity, we also assume that the hedger's book characteristics are completely captured by the considered subpopulations and we neglect the risk of misspecifying the appropriate subgroup for individual policyholders.

this end, we follow Börger and Schupp (2018) and model both CBD time processes via trend-stationary processes with piecewise linear trends and randomly changing slopes, the so-called AMTs, over time.

- The evolution of socioeconomic mortality differentials over time relative to the overall population mortality giving subpopulation-specific death probabilities $q_{x,t}^{[p]}, t \geq 0$ for each subpopulation $p \in \{1, \ldots, N_{Sub}\}$. In the base case, we follow Villegas and Haberman (2014) and rely on a multivariate random walk with drift (RWD) for the mortality differentials.

- Small sample risk due to a limited portfolio size by drawing realizations for survivors (conditional on realized mortality rates $q_{x,t}^{[p]}$) from a Binomial distribution.

A rigorous differentiation between the three components of longevity risk will be essential for a clear assessment of population basis risk in index-based hedges which are associated with the reference population or its subpopulations.

In our multi-population setup, the EMT valuation model[4] for pathwise derivations of best estimate mortality assumptions at future valuation dates needs to consist of two components:[5]

- A methodology for estimating the reference population's prevailing mortality level and trend for pathwise derivations of time-T best estimate mortality rates $\tilde{q}_{x,t}^{[R]}(T), t > T$.

- Inspired by Cairns and El Boukfaoui (2019), an additional approach for deriving prevailing subpopulation-specific experience ratios to account for differing mortality levels and trends relative to the reference population. For each subpopulation $p \in \{1, \ldots, N_{Sub}\}$, best estimate mortality rates $\tilde{q}_{x,t}^{[p]}(T), t > T$ are then derived by adjusting best estimate mortality rates of the reference population according to the prevailing experience ratios.

The AMT simulation model is used to generate sample paths of future mortality and to derive, among others, annuity payments and payouts of cash flow hedges, whereas the computation of best estimate liabilities and the derivation of hedge payoffs that build on future best estimate mortality assumptions need to be carried out under the EMT valuation model.[6]

2.2 Liability to be hedged

The liability to be hedged consists of a simplified portfolio of immediate or deferred life annuities, where all premiums have been paid upfront to the insurer. Starting at the retirement age x_R, these contracts pay one unit of currency at the beginning of each year until the beneficiary dies.[7] We do not consider fees or any further features (such as death benefits or a profit sharing mechanism).

We consider a single cohort of size N_{Book} with starting age x_0 at time $t = 0$. The number of initial policyholders from subpopulation $p \in \{1, \ldots, N_{Sub}\}$ is given by $B_{x_0,0}^{[p]} := \eta_p N_{Book} \in \mathbb{N}$

[4]Quantities that are derived under the EMT valuation model are denoted by $\tilde{}$ throughout the paper.

[5]Here, we implicitly assume that the whole data set is publicly available for all (sub-)populations and that it will be updated directly in each future year according to realized mortality.

[6]For a recent discussion on which of the two models is relevant for which kind of question, we refer to Börger et al. (2019).

[7]Assuming no concentration of risk by amounts, small sample risk is diversifiable in large portfolios.

with $\eta_1 + \cdots + \eta_{N_{Sub}} = 1$. The time-$t$ random present value of all future *unhedged liabilities* then reads as

$$L(t) := \sum_{p=1}^{N_{Sub}} L^{[p]}(t) := \sum_{p=1}^{N_{Sub}} \sum_{s>t} (1+r)^{-(s-t)} \, \mathbb{1}_{\{x_0+s \geq x_R\}} B^{[p]}_{x_0+s,s}, \qquad t \geq 0,$$

where $B^{[p]}_{x_0+s,s}$ denotes the number of survivors in the book population from subpopulation p aged $x_0 + s$ at time $s > 0$. The time-t *best estimate unhedged liabilities* are calculated using best estimate mortality derived under the EMT valuation model as

$$\widetilde{L}\,(t) := \sum_{p=1}^{N_{Sub}} B^{[p]}_{x_0+t,t} \sum_{s>t} (1+r)^{-(s-t)} \, \mathbb{1}_{\{x_0+s \geq x_R\}} \prod_{u=t}^{s-1} \left(1 - \tilde{q}^{[p]}_{x_0+u,u+1}(t)\right), \qquad t \geq 0.$$

With a hedge in place, the liability is netted with future hedging instrument cash flows giving the time-t random present value of the *hedged liabilities*

$$L_H(t) := L(t) - H(t), \qquad t \geq 0,$$

where $H(t)$ denotes the time-t random present value of all future cash flows from the hedge according to the underlying hedge structure, which will be introduced in Section 3. Analogously to the unhedged case, we calculate the time-t *best estimate hedged liabilities* as

$$\widetilde{L}_H\,(t) := \widetilde{L}\,(t) - \widetilde{H}\,(t)\,, \qquad t \geq 0,$$

where $\widetilde{H}(t)$ denotes the time-t best estimate of all future hedging instrument cash flows. We would like to stress that this quantity is derived on a best estimate basis and is not meant to represent the instrument's market value.

2.3 Economic capital model

Under modern risk-based solvency regimes, (re)insurance companies are required to provide SCRs for longevity risk. Under Solvency II, the SCR is defined as the 99.5% VaR of the basic own funds over a one-year horizon, where the basic own funds correspond to the difference between the market value of assets and the market value of liabilities. In principle, the SCR corresponds to the capital required to cover all losses which may occur over the following year at a confidence level of at least 99.5%. Following Börger (2010) and Börger et al. (2019), we assume that the evolution of assets (besides the longevity hedge instruments) is independent of realized mortality and thus does not contribute to the longevity SCR. Moreover, we assume that there is no loss-absorbing capacity of technical provisions.

For the determination of longevity SCRs, the company can choose between the following approaches:

- **Internal model**: The company might use a (partial) stochastic internal model to determine the longevity SCR according to the 99.5% VaR concept. As discussed by Börger (2010), longevity risk over a one-year horizon consists of two components: more annuitants than anticipated might survive the year or longevity assumptions might change over the year in an unfavorable direction. Typically, the latter is the more relevant factor. The SCR in year T, denoted as $SCR_L^{IM}(T)$, is defined as the 99.5th percentile of the change in best estimate liabilities from time T to $T+1$:

$$\frac{\widetilde{L}\,(T+1) + CF(T+1)}{1+r} - \widetilde{L}\,(T)\,,$$

where $CF(T+1)$ denotes the company's cash flows of the longevity prone business (benefits paid to the annuitants) between T and $T+1$. For the hedged position, $SCR_{L_H}^{IM}(T)$ is defined analogously by substituting the unhedged quantities with the hedged liabilities and hedged cash flows, respectively, where the latter additionally contain cash flows from the hedge contract over the year.

- **Standard formula**: Alternatively, companies are allowed to use a simplified standard formula as an approximation for the 99.5% VaR approach. Under this approach, the SCR in year T is determined as the change in best estimate liabilities due to a sudden and permanent longevity shock of 20% on best estimate probabilities of death for all ages, i.e.

$$SCR_L^{SF}(T) := \widetilde{L}\,(T\,|shock(20\%)\,) - \widetilde{L}\,(T)\,.$$

For the hedged position, $SCR_{L_H}^{SF}(T)$ is defined analogously based on the hedged liabilities. The standard formula has come under some criticism in the academic literature for its unrealistically simple structure, see e.g. Börger (2010). In particular, the structure of the uniform one-off shock irrespective of age and maturity does not appropriately reflect the true nature of longevity risk as a typically slowly accumulating demographic trend risk. In spite of these shortcomings, many companies still use the standard formula.

Since the SCR at some future point in time T obviously depends on the mortality evolution up to time T, we interpret it as a random variable. Given an outer simulation path containing realized mortality up to time T, the SCR (with or without a longevity hedge in place) can be computed conditional on this simulation path. The derivation of the 99.5th percentile in the internal model requires an additional inner Monte Carlo simulation, where mortality over a one-year horizon is simulated with the AMT simulation model and the best estimate (hedged) liabilities are reevaluated in the EMT valuation model. Within a two-level nested Monte Carlo simulation, entire distributions for the company's SCRs over time can be derived.

2.4 Hedging objectives

Our main objective is to investigate the impact of different hedging instruments on the annuity provider's future cash flow profile. In addition to annual benefit payments to surviving annuitants, the company has to compensate its shareholders for providing equity to cover its SCR. Due to the stochastic nature of future SCRs this additional cost of regulatory capital is also a random variable. This motivates the definition of the *adjusted unhedged liabilities* as

$$\Pi_L^M := L(0) + CoC_L^M, \qquad M \in \{IM, SF\},$$

where $CoC_L^M := \sum_{t\geq 0} \frac{r_{CoC} SCR_L^M(t)}{(1+r)^{t+1}}$ denotes the time zero random present value of all costs of capital for the unhedged liabilities which are either derived with the internal model ($M = IM$) or by applying the standard formula ($M = SF$). The cost of capital rate r_{CoC} reflects the return in excess of the risk-free rate which shareholders demand for providing equity. The best estimate adjusted unhedged liabilities are the initial portfolio reserve plus the expected cost of regulatory capital which has to be reserved in addition as a risk margin under Solvency II. For the hedged position, the *adjusted hedged liabilities* are defined analogously as

$$\Pi_{L_H}^M := L_H(0) + CoC_{L_H}^M, \qquad M \in \{IM, SF\}.$$

We quantify the effect of hedging on the adjusted liabilities by means of the following measures:

- **Capital efficiency**: The expected mitigating impact on the adjusted liabilities is denoted as *net cost of capital relief*

$$\begin{aligned} NReCoC^M(H) :=& \mathbb{E}\left(\Pi_L^M\right) - \mathbb{E}\left(\Pi_{L_H}^M\right) \\ =& \mathbb{E}\left(CoC_L^M\right) - \mathbb{E}\left(CoC_{L_H}^M\right) + \mathbb{E}(H(0)), \qquad M \in \{IM, SF\}, \end{aligned}$$

  where two opposing effects come into play:

  - On the one hand, hedging typically reduces the hedger's SCRs which in turn generates a positive expected *cost of capital relief* of

$$ReCoC^M(H) := \mathbb{E}\left(CoC_L^M\right) - \mathbb{E}\left(CoC_{L_H}^M\right) \geq 0, \qquad M \in \{IM, SF\}.$$

  - On the other hand, the expected present value of all hedging instrument cash flows $\mathbb{E}(H(0))$ is typically negative reflecting the absolute risk loading on top of the objective best estimate value charged by the counterparty for taking risk.

  In this setting, a company which is completely hedged against longevity risk would not have to provide any SCRs for longevity risk and the cost of capital would reduce to zero. If such a perfect hedge was offered on a best estimate basis (i.e. $\mathbb{E}(H(0)) = 0$), it would obviously provide the maximal net cost of capital relief of $\mathbb{E}\left(CoC_L^M\right)$. With regard to this benchmark, we define the *capital efficiency* of a hedge H as

$$CE^M(H) := \frac{NReCoC^M(H)}{\mathbb{E}\left(CoC_L^M\right)}, \qquad M \in \{IM, SF\}.$$

  We refer to a hedge as being capital efficient if $CE^M(H) > 0$, i.e. if the generated cost of capital saving exceeds the hedging costs. Hedge $H1$ is said to be more capital efficient than hedge $H2$ if $CE^M(H1) > CE^M(H2)$.

- **Hedge effectiveness**: We define the effectiveness of a hedge as the achieved relative risk reduction measured under a risk measure ρ in the centered adjusted liabilities:

$$HE_\rho^M(H) := 1 - \frac{\rho\left(\bar{\Pi}_{L_H}^M\right)}{\rho\left(\bar{\Pi}_L^M\right)} := 1 - \frac{\rho\left(\Pi_{L_H}^M - \mathbb{E}\left(\Pi_{L_H}^M\right)\right)}{\rho\left(\Pi_L^M - \mathbb{E}\left(\Pi_L^M\right)\right)}, \qquad M \in \{IM, SF\}.$$

  Obviously, a perfect hedge offers the maximal hedge effectiveness of one. We refer to a hedge $H1$ as being more effective than hedge $H2$ if $HE_\rho^M(H1) > HE_\rho^M(H2)$. We would like to stress that this definition of hedge effectiveness based on the adjusted liabilities explicitly considers the reduction in uncertainty regarding future costs of capital.

Even though both quantities allow for an intuitive interpretation, the most effective hedge does not necessarily provide the highest capital efficiency, especially when the hedge provider charges a non-zero risk premium. In this case, the company needs to find a reasonable trade-off between hedge effectiveness and capital efficiency depending on the strategic hedging objective.

3 Hedging instruments

In line with the current state of the longevity risk transfer market, we assume an incomplete and illiquid market where the hedger cannot resell or terminate a contract prior to maturity. We follow a widely used approach and define an equivalent risk-adjusted measure $\mathbb{Q}$ under which the price of any security is the expected value of its discounted payoff. This pricing measure assigns higher probability mass to scenarios which are unfavorable for a longevity hedge provider implying a longevity risk premium, the amount of which depends on the assumed market prices for the underlying longevity risk drivers. Overall, the underlying model structure is preserved under this change of measure. A comprehensive description of this pricing technique, including all technical details, can be found in Freimann (2020).

In this section, we describe the hedging instruments which consist of two components: an underlying index population (IP) (introduced in Section 3.1) and a basic hedge payout structure (defined in Section 3.2). By linking the same payout structure to different IPs, we construct different instruments which give rise to varying levels of population basis risk.

3.1 Index populations

Each hedging instrument is linked to one of the following IPs:

- $IP = \mathcal{B}$: The hedge is fully customized and directly linked to the survivors and the mortality experience in the hedger's book population.[8]
- $IP = \mathcal{S}$: The instrument is index-based and linked to the subpopulations. As opposed to the previous design, this gives rise to small sample risk due to a limited portfolio size.
- $IP = \mathcal{R}$: The hedge is index-based and linked to the reference population leaving the hedger with small sample risk as well as socioeconomic basis risk.

In the spirit of Cairns and El Boukfaoui (2019), we construct the index-based instruments ($IP = \mathcal{S}, \mathcal{R}$) by replacing customized quantities in the respective fully customized version by appropriate hedge indices. These index-based proxies are defined to match as closely as possible (in terms of magnitude and sensitivity to mortality) the fully customized quantities while only using the observable information from the underlying IP. While all information which is available at inception may be included in structuring the hedge contract, new mortality information which becomes available over the hedge horizon, but is not captured by the underlying IP cannot be included. In particular, hedges linked to the reference population build on the experience ratios at the incepction of the hedge contract in order to anticipate mortality differentials between the book and the reference population. The risk of changing experiences ratios over the course of the hedge cannot be hedged though. This construction of index-based hedges is described in detail in Appendix A.2.

3.2 Hedge payout structures

The time-t random present value of all future hedging instrument cash flows is given by

$$H^{[IP]}(t) := \sum_{s>t}^{\tau} (1+r)^{-(s-t)}\, h^{[IP]}(s) - (1+r)^{-(s-1-t)}\, p^{[IP]}(s-1), \qquad IP \in \{\mathcal{B}, \mathcal{S}, \mathcal{R}\},$$

[8] Even though we also denote the hedger's portfolio population as an *index* population for the sake of a consistent notation, the associated hedge instruments are meant to be indemnity-based.

where τ denotes the contract maturity, $h^{[IP]}(t), 0 < t \leq \tau$ represents the payoff to the hedger at time t according to the underlying hedge payout structure, and $p^{[IP]}(t), 0 \leq t < \tau$ denotes the path-dependent premium charged by the counterparty at time t. These will only become relevant for rolling hedge strategies. For all forward-type hedge structures, the risk premium is directly included in the hedge payout structure by means of risk-adjusted forward rates which are determined at time zero so that $\mathbb{E}^{\mathbb{Q}}\left(H^{[IP]}(0)\right) = 0$. For hedge payments that are based on current or future best estimate mortality assumptions, we assume that both parties agree on deriving the required values with the EMT valuation model.[9]

In what follows, we give a concise overview of the considered hedge payout structures. Further technical details are provided in Appendix A.2.

3.2.1 Longevity swaps

In a longevity swap, the hedger receives a sequence of cash flows corresponding to the underlying liability cash flows in exchange for a series of forward rates of the form

$$h^{[IP]}(t) := \mathbb{1}_{\{x_0+t \geq x_R\}} \left(S^{[IP]}_{x_0+t,t} - \mathbb{E}^{\mathbb{Q}}\left(S^{[IP]}_{x_0+t,t}\right)\right), \qquad 0 < t \leq \tau, \qquad IP \in \{\mathcal{B}, \mathcal{S}, \mathcal{R}\},$$

where $S^{[IP]}_{x_0+t,t}$ either represents the actual number of survivors aged $x_0 + t$ at time t (in case of $IP = \mathcal{B}$) or a survivor index as a proxy (in case of $IP = \mathcal{S}, \mathcal{R}$). Note that by construction an unlimited ($\tau = \infty$) fully customized ($IP = \mathcal{B}$) longevity swap provides a perfect hedge.

3.2.2 Annuity forwards

In an annuity forward, the hedger receives the best estimate present value (according to up-to-date mortality assumptions at maturity) of all future annuity payments in exchange for a fixed forward liability resulting in a single hedge payout of the form

$$h^{[IP]}(\tau) := \widetilde{L}^{[IP]}(\tau) - \mathbb{E}^{\mathbb{Q}}\left(\widetilde{L}^{[IP]}(\tau)\right), \qquad IP \in \{\mathcal{B}, \mathcal{S}, \mathcal{R}\},$$

where $\widetilde{L}^{[IP]}(\tau)$ either represents the actual time-τ best estimate liabilities (in case of $IP = \mathcal{B}$) or a liability index as a proxy (in case of $IP = \mathcal{S}, \mathcal{R}$). While this hedge (partly) transfers longevity risk up to time τ, the hedger is still exposed to the risk that longevity after time τ might not evolve as expected. As discussed by Börger et al. (2019), when assessing the effectiveness of this hedge in a stochastic simulation one needs to be aware that the hedge payout can only be based on the EMT which might deviate from the unobservable AMT.

3.2.3 Q-forwards

In a q-forward, the hedger exchanges a fixed forward rate against realized mortality rates resulting in a single hedge payout of the form

$$h^{[IP]}(\tau) := \mathbb{E}^{\mathbb{Q}}\left(Q^{[IP]}_{x_0+\tau,\tau}\right) - Q^{[IP]}_{x_0+\tau,\tau}, \qquad IP \in \{\mathcal{S}, \mathcal{R}\},$$

where $Q^{[IP]}_{x_0+\tau,\tau}$ represents the underlying mortality rates which are either linked to the subpopulations (in case of $IP = \mathcal{S}$) or to the reference population (in case of $IP = \mathcal{R}$). Since the

[9] In practice, the hedge instrument valuation model (which is specified in the hedge contract) may differ from the liability valuation model. For simplicity, we use the EMT valuation model for both purposes.

liability consists of a single cohort aged $x_0 + \tau$ at time τ, we focus on a single q-forward with same reference date and age. Since this payoff structure obviously differs from the liability, the hedger additionally faces so-called structural basis risk. The derivation of the optimal q-forward portfolio is described in Appendix A.2. For simplicity, we limit our analysis to a static setting where all hedge ratios are calibrated at time zero and remain fixed over the hedge horizon.

3.2.4 Rolling one-year call spread options

Even though all instruments introduced so far typically reduce SCRs through their risk-mitigating effect, their payoff structures are not primarily designed for this particular purpose. In line with the one-year view of Solvency II, we also consider a rolling hedge strategy based on one-year call spread options, which have been analyzed in the context of a regulatory capital model by Cairns and El Boukfaoui (2019).

Assume that at any point in time $0 \leq t < \tau$, the hedger enters into a one-year call spread option contract of the following form: If the underlying hedge index $\widetilde{X}^{[IP]}(t+1)$ at the end of the year exceeds a predefined attachment point $AP(t)$, the hedger receives a payment of

$$h^{[IP]}(t+1) := (EP(t) - AP(t)) \max\left\{0; \min\left\{\frac{\widetilde{X}^{[IP]}(t+1) - AP(t)}{EP(t) - AP(t)}; 1\right\}\right\}, \ IP \in \{\mathcal{B}, \mathcal{S}, \mathcal{R}\},$$

which increases linearly until the hedge index reaches the exhaustion point $EP(t)$. The hedge index is defined to match (in case of $IP = \mathcal{B}$) or to replicate as closely as possible (in case of $IP = \mathcal{S}, \mathcal{R}$) the random drivers in the company's SCR computation from time t to $t+1$. Hence, the hedge provides an offsetting payment if longevity over the year evolves in an unfavorable direction. For comparability, we use the same attachment and exhaustion points for all IPs:

- The attachment points are typically set above the current best estimate liabilities to ensure that the option is only triggered in scenarios of significant mortality improvements over the year. Exemplarily, we define the attachment points at time t via a uniform relative longevity stress of $ap \geq 0$ on current best estimate liabilities:

  $$AP(t) := (1+r)\widetilde{L}\left(t \,|shock(ap)\right).$$

- The exhaustion points at time t are defined as

  $$EP(t) := \begin{cases} \alpha^{ep}\text{-th percentile of } \widetilde{X}^{[\mathcal{B}]}(t+1), & \text{for an internal model design} \\ (1+r)\widetilde{L}\left(t \,|shock(ep)\right), & \text{for a standard formula design,} \end{cases}$$

  which are typically chosen below $\alpha^{ep} = 99.5\%$ or $ep = 20\%$, respectively, so that the instrument always pays out in full in the Solvency II stress scenarios.

In return, the hedge provider demands a premium of $p^{[IP]}(t) := (1+r)^{-1}\mathbb{E}^{\mathbb{Q}}\left(h^{[IP]}(t+1)|\mathcal{F}_t\right)$ at time t, where $\mathcal{F}_t$ contains the observable mortality information available to the hedge provider at time t.[10] Note that the hedger is subject to 'rolling risk' since the costs for the next hedge in the sequence might significantly rise in unfavorable scenarios.

[10]This observable information does not include the prevailing AMT to avoid the misestimation of rolling risk. Instead, the risk taker needs to build his pricing on the current EMT. Furthermore, he should take into account the uncertainty in this mortality trend estimate (see e.g. Börger and Schupp (2018)). However, to avoid over-complexity, we refrain from doing so and accept a potential underestimation of market prices at future valuation dates.

Description	**Parameter**	**Value**
Starting age	x_0	65
Retirement age	x_R	65
Initial book size	N_{Book}	10,000
Socioeconomic composition	η	$(0, 0, 30\%, 30\%, 40\%)$
Risk-free interest rate	r	2%
Cost of capital rate	r_{CoC}	6%
Market price of longevity risk	λ	30%
Risk measure	ρ	$\text{TVaR}_{90\%}$

Table 1: Model parameters in the base case.

4 Numerical results

For the following numerical application, we use the male population of England and Wales as the reference population and consider five subpopulations of different socioeconomic status (ordered from the most to the least deprived areas) based on the Index of Multiple Deprivation (IMD) for England.[11] Details on the underlying data sets, the applied calibration techniques, and the resulting model parameters can be found in Freimann (2020).

In the base case, we consider a portfolio of immediate life annuities consisting of 10,000 policyholders aged 65 at the beginning of the year 2017, which is set to $t = 0$. Regarding its socioeconomic structure, we assume a rather affluent book population consisting exclusively of the three most affluent IMD-subpopulations to provoke a considerable exposure to socioeconomic basis risk. Furthermore, we assume a risk-free interest rate of 2% and a cost of capital rate of 6%. Moreover, we calibrate the market price of longevity risk so that the costs for full securitization roughly match the expected cost of capital for keeping the risk under the internal model and obtain $\lambda = 30\%$ for all risk drivers, cf. Freimann (2020). Throughout this work, we use the 90% Tail-Value-at-Risk (TVaR) for the quantification of hedge effectiveness.[12] Table 1 summarizes the model parameters for the base case.

We perform a two-level nested Monte Carlo simulation with 10,000 outer sample paths. For each path and for every year, we rely on 10,000 inner one-year scenarios for the SCR computation in the internal model and for the associated derivation of exhaustion points for call spread contracts. For their pathwise pricing, we additionally use 1,000 risk-adjusted inner one-year scenarios. For consistency, we rely on the same scenarios for all hedging instruments.

4.1 Unhedged (adjusted) liabilities

We start with a brief investigation of the hedger's initial situation without hedging. Table 2 shows the mean and the 90% TVaR of the company's (centered) liabilities with and without an adjustment for future cost of regulatory capital. As anticipated, the adjustment for future capital charges increases the expected financial obligations. The expected cost of capital is more than twice as high under the standard formula than under the internal model since the standard formula produces rather high SCRs compared to the internal model, cf. Figure 1. This is in line with findings of Börger (2010) and can be traced back to structural shortcomings

[11] The data sets are kindly provided by the Human Mortality Database and the Office for National Statistics.

[12] We also performed our analyses under alternative risk measures such as variance, 95% TVaR, and 99.5% VaR and found no structurally different results.

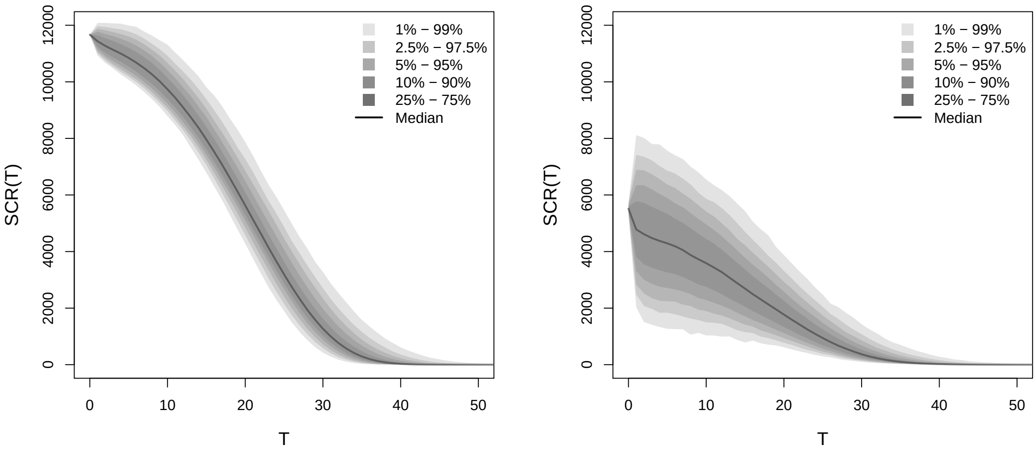

(a) Standard formula

(b) Internal model

Figure 1: Percentile plots of the company's SCRs over time.

	$L(0)$	Π_L^{SF}	Π_L^{IM}
$\mathbb{E}(\cdot)$	165,824	176,851	169,904
$\text{TVaR}_{90\%}(\cdot)$	10,074	11,305	11,477

Table 2: Mean and 90% TVaR of the hedger's unhedged (adjusted) liabilities.

of the standard formula's rather conservative longevity stress.[13]

The 90% TVaR indicates that the randomness in future SCRs also increases the company's overall longevity risk exposure. The uncertainty in future (cost of) regulatory capital is more pronounced under the internal model than under the standard formula since the standard formula's SCR at any point in time solely depends on the effect of a prescribed longevity shock on current best estimate liabilities. The SCR under the internal model on the other side is driven by the variability in best estimate mortality assumptions over a one-year horizon and thus mainly by the variability in one-year EMT changes, which can be quite significant. The reason for this is twofold. First, the variability in EMT changes from time t to $t+1$ depend on realized mortality up to time t, particularly after trend changes in previous years. Second, we account for parameter uncertainty in the trend change parameters at the start of each outer simulation path. Hence, the overall level of SCRs generally differs between the simulation paths depending on the drawn model parameters.[14]

[13]Of course, the SCRs derived under the internal model highly depend on the applied weighting in the EMT valuation model, which controls the adaption of best estimate mortality assumptions to observable changes in mortality patterns. To obtain objective SCRs, it is therefore important to determine the optimal weighting by means of a reasonable optimization criterion, see Börger et al. (2019).

[14]This is also the reason why the internal model's initial SCR lies noticeably above the median SCRs in the following years. The details for the modeling of parameter uncertainty can be found in Freimann (2020).

Instrument	Description	Parameter	Value
Longevity swaps	index population	IP	$\in \{\mathcal{B}, \mathcal{S}, \mathcal{R}\}$
	maturity	τ	$\in [25, \infty)$
Annuity forwards	index population	IP	$\in \{\mathcal{B}, \mathcal{S}, \mathcal{R}\}$
	maturity	τ	$\in [1, 25]$
Q-forwards	index population	IP	$\in \{\mathcal{S}, \mathcal{R}\}$
	maturity	τ	$\in [1, 25]$
Series of one-year call spread options (standard formula design)	index population	IP	$\in \{\mathcal{B}, \mathcal{S}, \mathcal{R}\}$
	maturity	τ	$\in [1, 25]$
	attachment point	ap	5% (shock)
	exhaustion point	ep	20% (shock)
Series of one-year call spread options (internal model design)	index population	IP	$\in \{\mathcal{B}, S, R\}$
	maturity	τ	$\in [1, 25]$
	attachment point	ap	1% (shock)
	exhaustion point	α^{ep}	99.5(th percentile)

Table 3: Overview of hedging instruments in the base case.

4.2 The drivers of hedge effectiveness and capital efficiency

Now, we introduce the hedges which are summarized in Table 3. For annuity forwards, q-forwards, and rolling call spread options, we consider contract maturities up to 25 years, i.e. until the underlying cohort reaches age 90, and for longevity swaps we focus on longer times to maturity beyond 25 years. Regarding the call spread option contracts, the hedger further needs to agree with the hedge provider on optimal attachment and exhaustion points. To simplify the discussion, we initially focus on one exemplary hedge structure for each SCR computation method: a design tailored to the internal model with an exhaustion point set at the relevant 99.5th percentile and an alternative parameter set of $(ap, ep) = (5\%, 20\%)$ for a standard formula design. We deliberately choose a higher attachment point of $ap = 5\%$ for the latter design (compared to $ap = 1\%$ for the former) to demonstrate the effect of higher attachment points. Nevertheless, we address the optimal choice of attachment and exhaustion points from a wider range of parameters in Section 4.3.

In this section, we start with a separate analysis of the two main effects of hedging: capital relief (as driver for capital efficiency) and risk reduction (as driver for hedge effectiveness). Subsequently, we analyze both aspects simultaneously in Section 4.3.

4.2.1 Capital relief: standard formula vs. internal model

As a first step, we analyze how selected hedges impact the hedger's longevity SCRs over time. Figure 2 shows percentile plots of the SCRs over time with selected fully customized ($IP = \mathcal{B}$) longevity hedges over $\tau = 25$ years in place under the internal model (upper row) and under the standard formula (lower row). Clearly, hedging only impacts the amount of regulatory capital over the presumed hedge horizon. Interestingly, the structures of the SCRs for the hedged positions clearly differ in terms of both level and variability not only among the instruments but also between the SCR computation methods.

First, we notice that the limited longevity swap provides a higher capital relief under the standard formula than under the internal model. This is due to the fact that the latter's SCRs are mainly driven by the risk that long-term mortality assumptions, which are not covered by

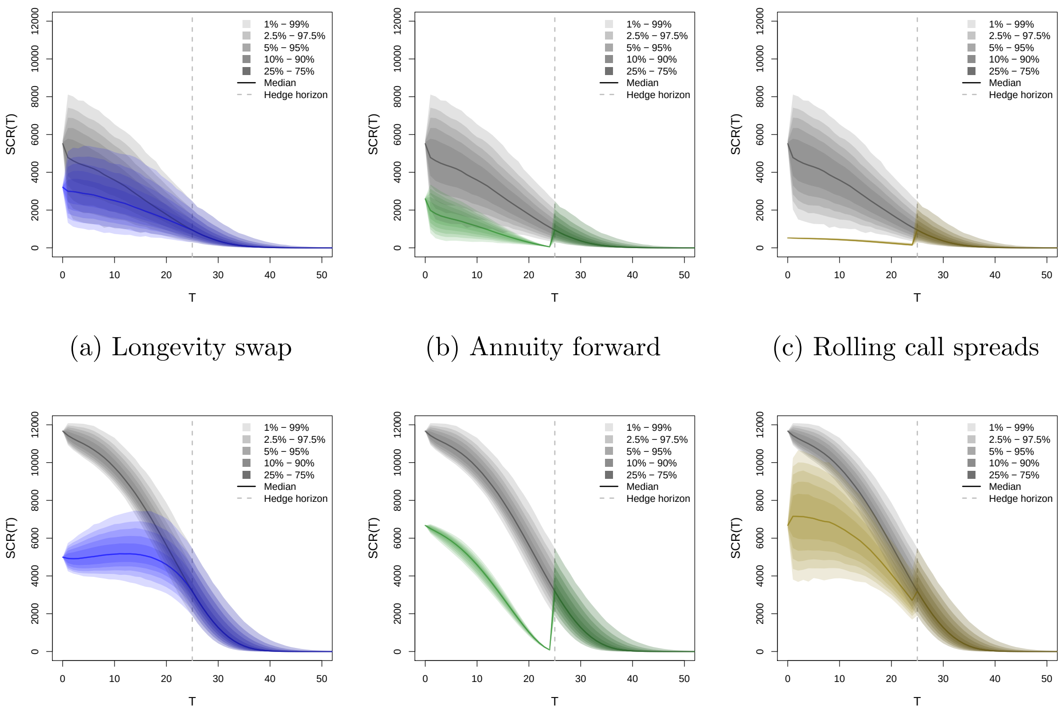

(a) Longevity swap (b) Annuity forward (c) Rolling call spreads

(d) Longevity swap (e) Annuity forward (f) Rolling call spreads

Figure 2: Percentile plots of SCRs over time derived under the internal model (upper row) and under the standard formula (lower row) without hedging (gray) and with different fully customized ($IP = \mathcal{B}$) longevity hedges over $\tau = 25$ years in place: longevity swap (blue, left panel), annuity forward (green, middle panel), and rolling call spread options designed for the internal model (gold, right panel).

a longevity swap over 25 years, need to be revised. The standard formula's uniform shock on the other side affects short and long-term mortality likewise resulting in a substantial capital relief for a limited longevity swap.

In contrast, the annuity forward covers all observable changes in mortality patterns over the hedge horizon. In particular, one year prior to maturity the instrument provides protection against all changes in best estimate mortality over the year and the hedger therefore only has to hold SCRs for the uncertain cash flows that are due at the end of the year. This explains why the annuity forward is quite effective in reducing SCRs towards the end of its term. Since similar observations can be made under both economic capital models, they seem to agree on an adequate relative capital relief for this instrument.

Finally, the rolling call spread portfolio which is tailored to the internal model shows the most striking behavior. Clearly, it serves its primary purpose of reducing longevity SCRs in the internal model. Since the attachment points are set moderately above the current best estimate liabilities, the hedger is still exposed to residual risk requiring the provision of SCRs. Nevertheless, these lie far below their unhedged counterparts and their variability is reduced substantially. However, the standard formula's SCRs for the hedged position are significantly larger and much more volatile. The reason for this phenomenon lies in the path-dependent exhaustion points which are calibrated with respect to the relevant 99.5th percentiles in the

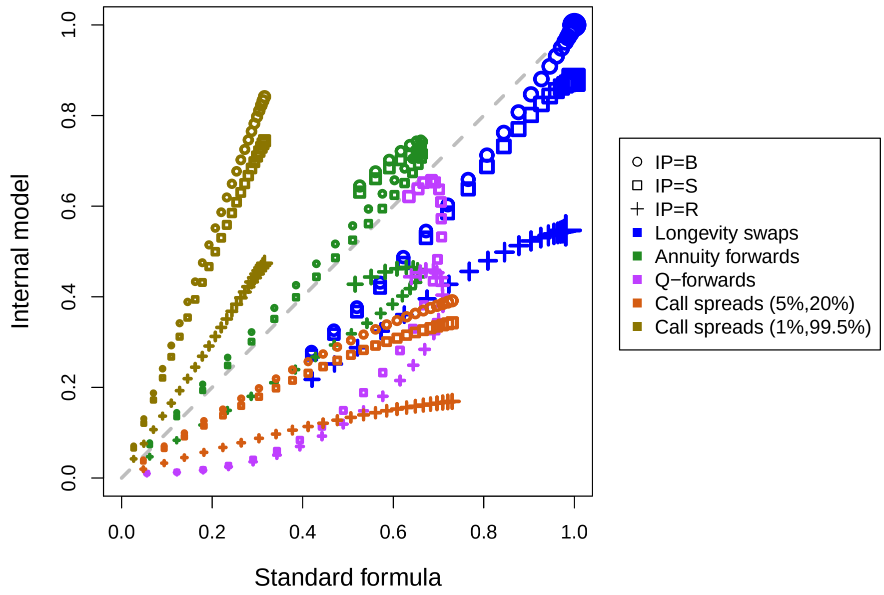

Figure 3: Overview of relative expected cost of capital relief under the internal model (y-axis) compared to the standard formula (x-axis). The dot size increases in the contract maturity.

internal model and therefore obviously not suited for the standard formula's shock approach.

Overall, the percentile plots illustrate that the standard formula and the internal model might produce discordant capital reliefs. Moreover, it can be seen that the considered hedges have structurally different impacts on future SCRs even though they are all directly linked to the hedger's book population. The matter becomes further complicated when we extend the analysis to index-based hedges. The hedger then further needs to accept an appropriate haircut for the reduction in hedge efficiency due to population basis risk.

To obtain a concise picture across all instruments, hedge horizons, and IPs under consideration, Figure 3 provides an overview of the proportionate expected cost of capital relief

$$\frac{ReCoC^{M}\left(H^{[IP]}\right)}{\mathbb{E}\left(CoC_{L}^{M}\right)}, \qquad M \in \{IM, SF\}, \qquad IP \in \{\mathcal{B}, \mathcal{S}, \mathcal{R}\}$$

for the hedges in Table 3 under the internal model on the y-axis compared to the standard formula on the x-axis. The dot size increases in the instrument's time to maturity and the underlying IP is visualized by the following symbols: circles for $IP = \mathcal{B}$, squares for $IP = \mathcal{S}$, and crosses for $IP = \mathcal{R}$ (this identification is used throughout the paper). If all points were exactly on the dashed gray diagonal, both SCR computation methods would yield concordant relative cost of capital reliefs. As we can clearly see in Figure 3, this is only the case for the unlimited fully customized longevity swap, which provides full cost of capital relief under both models, and approximately for a range of annuity forwards. For the other instruments, discordant capital reliefs between the standard formula and the internal model are the rule rather than the exception. More precisely, all longevity swaps and q-forwards over a limited hedge horizon provide a higher relative cost of capital relief under the standard formula. This effect is most pronounced for the rolling call spread portfolio which is tailored to the standard formula. Since the attachment points, which are specified via a 5% longevity shock, still lie far below the standard formula's 20%-stress, this hedging strategy provides an attractive cost of capital relief under the standard formula. However, the internal model recognizes that rather high attachment points leave a substantial part of longevity risk with the hedger and therefore allows a lower capital relief. We conclude that a hedging strategy which is constructed with regard to the standard formula does not necessarily perform equally well under a risk-based internal model. The rolling call spread portfolio tailored to the internal model marks the other extreme as discussed above.

Regarding the impact of the hedge horizon, we find that the resulting cost of capital relief increases in the time to maturity for longevity swaps and for rolling call spread portfolios. However, for annuity forwards and q-forwards there appears to be an optimal hedge horizon beyond which hedge efficiency is declining. As we will see in the next subsection, this is due to the fact that the degree of risk reduction increases in the time to maturity for longevity swaps and rolling call spread portfolios while it peaks somewhere around $\tau = 20$ for annuity forwards and q-forwards.

Finally, for any given hedge payout structure (i.e. dots of the same color) comparing the results among the IPs (i.e. among different symbols) yields valuable insights into the impact of socioeconomic basis risk and small sample risk. First, we note that the expected cost of capital relief under the standard formula does practically not differ among the IPs. Obviously, a prescribed longevity shock to all (sub-)populations regardless of the underlying risk drivers cannot detect population basis risk. The risk-based internal model however identifies that index-based hedges do not cover all components of longevity risk and the underlying IP consequently has a significant impact on capital relief: the fully customized design always provides the highest cost of capital relief, followed by the subpopulation-linked design, and the instruments linked to the reference population provide the lowest relief. Since the difference between the latter two is much more distinct than between the former two designs, socioeconomic basis risk seems to be of higher relevance than small sample risk for our model company. Indeed, the applied RWD for modeling socioeconomic mortality differentials in conjunction with a trend adjustment in the subpopulation-specific experience ratios produces 'haircuts' for population basis risk of up to 50% (relative to the fully customized counterpart) which might be regarded as rather drastic. We will continue this discussion in Subsection 4.4.1 by considering alternative modeling assumptions. Nevertheless, it should be kept in mind that we deliberately assumed a rather affluent socioeconomic structure to provoke a noticeable exposure to socioeconomic basis risk.

4.2.2 Risk reduction

We proceed by comparing the considered hedges in terms of the achieved level of risk reduction. The left panel of Figure 4 shows their hedge effectiveness depending on their time to maturity which is derived based on the adjusted liabilities (including future cost of capital) assuming use of the internal model.

We start by discussing the effectiveness of the fully customized designs. For longevity swaps, hedge effectiveness increases in the time to maturity since each additional year of coverage secures the annuity payments for another year. Since late payments at high ages generally pose the greatest risk, a long hedge horizon is required to reach a high hedge effectiveness. By construction, the unlimited fully customized longevity swap offers the maximal

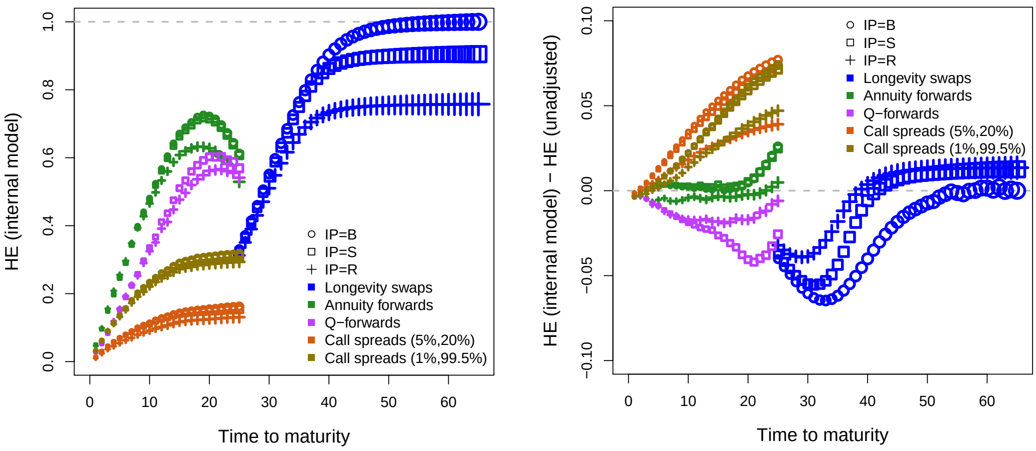

(a) $HE^{IM}_{TVaR_{90\%}}$ (b) $HE^{IM}_{TVaR_{90\%}} - HE^{-}_{TVaR_{90\%}}$

Figure 4: Hedge effectiveness for different times to maturity.

hedge effectiveness of one.

In contrast, mid-term contract maturities are sufficient to reach a substantial hedge effectiveness with q-forwards (of around 60%) and annuity forwards (of around 75%). In line with our findings in the previous subsection, we clearly see that the effectiveness of these instruments peaks somewhere around $\tau = 20$ which indicates the existence of an optimal hedge horizon. On the one side, short-term contracts leave a substantial part of longevity risk with the hedger since there is little time for longevity risk to accumulate. On the other side, linking the hedge payout to higher ages (and at the same time postponing the single hedge payout far into the future) obviously loses the connection to the hedger's liability structure.

Finally, we find that the rolling portfolios of one-year call spread options underperform in terms of hedge effectiveness. The reason lies in their pricing which is performed in each simulation path rather than at time zero. For the sake of illustration, consider a scenario in which longevity increases over time, for instance due to an unfavorable mortality trend change. As a consequence, longevity assumptions need to be revised and the current one-year hedge contract typically provides an offsetting payout. However, the premiums for the next contracts in the sequence typically rise as the hedge provider adapts to the updated mortality assumptions. These path-dependent hedge premiums reduce the effectiveness of this hedging strategy reflecting rolling risk. Moreover, it can be observed that higher attachment points lead to a reduction in hedge effectiveness since a greater portion of longevity risk is retained.

Regarding the impact of population basis risk, hedge effectiveness is declining for each component of longevity risk which is not covered. This effect is particularly pronounced for long-term longevity swaps: Starting from an effectiveness of 100% for an unlimited fully customized ($IP = \mathcal{B}$) contract, hedge effectiveness reduces to around 90% when small sample risk is retained (in case of $IP = \mathcal{S}$) and to 75% when also socioeconomic mortality differentials are no longer covered (in case of $IP = \mathcal{R}$). This illustrates the relevance of each component of longevity risk.

Now, assume for a moment that the hedger (inappropriately) measures hedge effectiveness based on the unadjusted liabilities regardless of any future capital charges. Since this commonly used approach does not account for the uncertainty in future cost of capital (and the reduction through the hedge overlay), the resulting hedge effectiveness may differ. The right panel of Figure 4 shows the absolute difference between the two quantities for the considered instruments, where a positive discrepancy means that hedge effectiveness is higher when derived based on the adjusted liabilities (including future cost of capital). It can clearly be observed that the allowance for stochastic cost of capital has a notable impact on longevity hedge effectiveness. If the uncertainty in future cost of capital is ignored, the effectiveness of mid-term longevity swaps with times to maturity between 25 and 40 years is overestimated by up to six percentage points. For longer hedge horizons, both approaches again coincide for the fully customized design and the effectiveness of long-term index-based longevity swaps is slightly higher when measured with respect to the adjusted liabilities. This is due to the fact that the uncertainty in future SCRs is mainly driven by mortality in the reference population, which is covered by index-based hedges. The q-forwards show a similar picture and the difference between the two approaches is moderate for annuity forwards. The largest discrepancy can be observed for the rolling call spread portfolios whose effectiveness is underestimated by up to seven percentage points when the uncertainty in future SCRs is ignored. As discussed in the previous subsection, these instruments are quite effective in reducing the magnitude and variability of future SCRs. Assessing their effectiveness without acknowledging the uncertainty in future cost of capital (and the reduction through hedging) clearly misses a key quality of these instruments and therefore underestimates their actual risk-reducing potential.

Overall, we conclude that hedge effectiveness might sometimes be underestimated, sometimes overestimated when the uncertainty in future cost of capital is ignored. The fact that misestimation can occur in both directions underlines the necessity to work with stochastic longevity SCRs for an objective analysis and comparison of different hedging instruments.

4.3 Combined analysis of hedge effectiveness and capital efficiency

So far, we have looked at the benefits of hedging, namely risk reduction and cost of capital relief, separately. In particular, we have seen that an unlimited fully customized longevity swap provides maximal hedge effectiveness and full cost of capital relief. If investors did not demand a risk premium, such a perfect hedge would obviously outperform any partial hedging solution in terms of both hedge effectiveness and capital efficiency. We now allow for a non-zero risk premium and incorporate the resulting costs for the hedger into our analysis.[15] As outlined in Section 3, we rely on a risk-adjusted pricing approach which compensates the hedge provider for the risks taken. Obviously, this approach does not account for all factors that might have an influence on market prices for longevity hedges. For instance, investors typically also consider administration costs, diversification benefits with other risks than longevity, or further strategic aspects. Nevertheless, it seems appropriate for our purposes to not explicitly deal with these issues.

The first subsection gives a brief overview of the costs of hedging. In the second subsection, we discuss the trade-off between hedge effectiveness and capital efficiency.

[15]Note that pricing already affected risk reduction and capital relief for rolling hedges in the previous section.

4.3.1 The costs of effective hedging

The left panel of Figure 5 shows the hedge effectiveness (y-axis) compared to the hedging costs (x-axis) for the considered instruments assuming a market price of longevity risk of $\lambda = 30\%$ for all systematic risk drivers. For each hedging instrument, the costs of hedging are derived as expected present value of all cash flows and typically have a negative sign reflecting the absolute risk loading on top of the objective best estimate value. Again, we derive hedge effectiveness based on the adjusted liabilities (including cost of capital) in the internal model.

When looking at the hedge payout structures individually, we observe that the applied risk-adjusted pricing approach produces reasonable prices in the sense that a higher hedge effectiveness generally comes with higher hedging costs. Regarding the considered IPs, index-based instruments linked to the reference population ($IP = \mathcal{R}$) are always offered at a lower price compared to their fully customized ($IP = \mathcal{B}$) and subpopulation-linked ($IP = \mathcal{S}$) counterparts. Intuitively, investors require higher compensation for instruments that also cover socioeconomic mortality differentials since they pose a systematic risk. Technically speaking, this is due to the risk adjustment for socioeconomic mortality differentials within the construction of the pricing measure. Furthermore, the costs for subpopulation-linked instruments coincide with those for their fully customized counterparts since we refrain from a further risk adjustment for diversifiable small sample risk. In reality, fully customized contracts might cause higher administration costs which the hedge provider might eventually pass on to the hedger.

Comparing all instruments, the costs for short-term annuity forwards and for series of rolling one-year call spreads appear disproportionately high given that they partly even exceed those for full securitization. Both phenomena can be traced back to the risk-adjustment for annual fluctuations around the prevailing AMT in the applied pricing approach. In our setup, an observer is not able to distinguish between a recent change in the long-term trend and a normal random fluctuation around it. For short-dated contracts which contain a feature that protects the hedger against variations in mortality assumptions, the counterparty is highly exposed to the risk that fluctuations around the AMT is mistaken for a longevity trend change and hence triggers a high hedge payout. Assigning higher probability mass to scenarios of unfavorable annual fluctuations therefore significantly raises the costs for short-dated annuity forwards and particularly for rolling hedge strategies that build on one-year contracts.

The right panel of Figure 5 shows the results assuming a reduced market price of risk for annual fluctuations around the AMT of $\lambda_\epsilon = 15\%$.[16] We find that the resulting hedging costs for annuity forwards and rolling call spread portfolios are now much more in line with the remaining instruments. We deduce that these contracts are only competitive if hedge providers reduce the risk adjustment for random fluctuations around the long-term mortality trend. Investors might be willing to do so if they are confident of writing this business over several consecutive years (for instance the hedger might commit to annual renewal of the contract). Over a multi-year horizon, the exposure to annual fluctuations diversifies. This is also the reason why the choice of the market price of risk for random fluctuations has a negligible impact on prices of long-dated contracts. Also note that the effectiveness of rolling hedge strategies slightly increases when reducing the hedging costs since rolling risk is mitigated. For these reasons, we use the reduced market price of risk for annual fluctuations

[16]The applied risk-adjusted pricing approach allows for assignment of individual market prices to the following risk drivers: unpredictable occurence, sign, and magnitude of future trend changes, annual fluctuations around the AMT, and subpopulation-specific mortality differentials, cf. Freimann (2020) for further details.

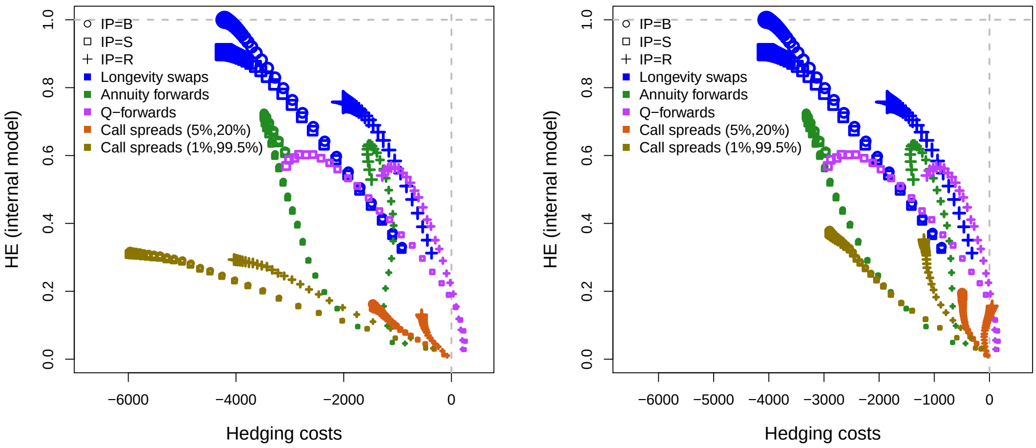

(a) $\lambda = 30\%$ for all risk drivers (b) Reduced $\lambda_\epsilon = 15\%$ for annual fluctuations

Figure 5: Hedging costs (x-axis) vs. hedge effectiveness (y-axis) in the base case.

throughout the remainder of this paper.

4.3.2 The trade-off between hedge effectiveness and capital efficiency

Figure 6 shows the hedge effectiveness of the considered instruments (y-axis) compared to their capital efficiency (x-axis) under the standard formula (left panel) and under the internal model (right panel). Recall from Section 2.4 that we quantify capital efficiency based on the expected cost of capital relief net of hedging costs, where the expected cost of capital for the unhedged liabilities serves as a benchmark. For the rolling call spread portfolios, we now focus on a fixed hedge horizon of 25 years and instead vary the attachment and exhaustion points. In more detail, the attachment points are increased from 1% to 10% (from red to orange) in steps of 1% for the standard formula design and we consider exhaustion points from the set $\{70\%, 75\%, 80\%, 85\%, 90\%, 95\%, 96\%, 97\%, 98\%, 99\%, 99.5\%\}$ (from yellow to gold) for the internal model design.

The key finding under the internal model is that there is no 'universally superior' hedging solution in the sense that no instrument simultaneously outperforms the others in terms of both hedge effectiveness and capital efficiency. Depending on the hedger's objective, different hedges might be suitable.

To begin with, consider a company which does not want to be exposed to longevity risk anymore. Aiming at maximal hedge effectiveness, it would certainly opt for the unlimited fully customized longevity swap. Since pricing is adapted to the internal model's risk assessment, the costs for this strategy practically offset the generated cost of capital saving resulting in a capital efficiency of approximately zero.

However, full risk reduction might not necessarily present the primary objective for all hedgers. Some might have an appetite to retain some longevity risk, especially if the costs for transferring certain parts of the risk exceed the expected cost of capital saving. For instance, if a fully customized longevity swap with a term of 35 years is considered, hedge effectiveness will naturally decline while capital efficiency will increase up to around 10%. These results are

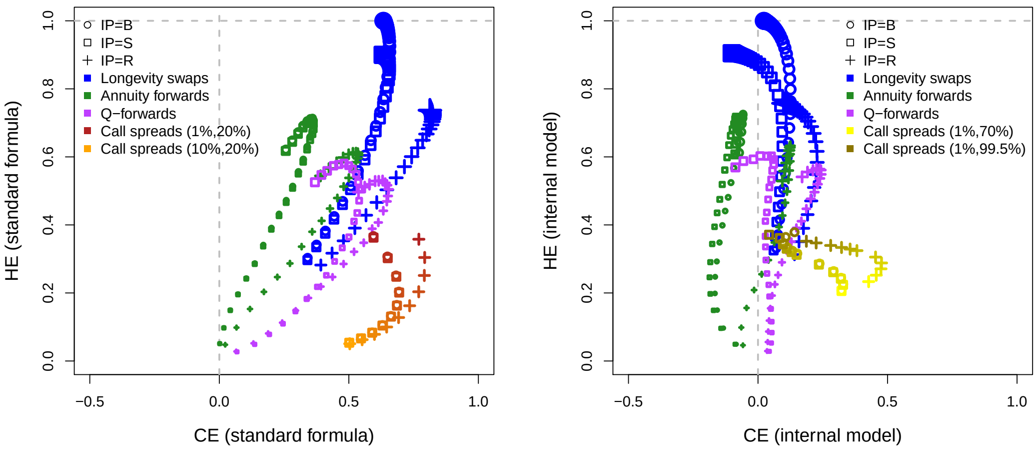

(a) Standard formula

(b) Internal model

Figure 6: Capital efficiency (x-axis) vs. hedge effectiveness (y-axis) in the base case.

in line with findings of Meyricke and Sherris (2014) who show that longevity hedging at high ages is generally not capital efficient under Solvency II. For terms below 35 years, however, capital efficiency is again declining since the cost of capital relief decreases more sharply than the hedging costs.

Comparing all instruments once again, we observe that an even higher capital efficiency can be reached by means of index-based instruments. Generally speaking, we identify a frontier of instruments which provide the highest capital efficiency for a given level of hedge effectiveness. This set includes long-dated fully customized longevity swaps (beyond 35 years), their index-based counterparts ($IP = \mathcal{R}$), and index-based q-forwards ($IP = \mathcal{R}$) over a hedge horizon of around 20 years. Eventually, the highest capital efficiency can be reached with a series of properly engineered one-year call spread options. For this purpose, an index-based strategy linked to the reference population ($IP = \mathcal{R}$) with exhaustion points set well below the relevant 99.5th percentiles appears to be most suitable. The reason for this is twofold. First, lowering the exhaustion points naturally reduces the hedging costs which in turn has a positive impact on capital efficiency. Second, we find that the negative impact of population basis risk on the cost of capital saving is negligible when the exhaustion points lie well below the relevant 99.5th percentiles (we will further discuss this in Subsection 4.4.1). This is in line with findings of Cairns and El Boukfaoui (2019), who demonstrate that index-based call spreads may provide the same capital relief as fully customized contracts when designed properly. Hence, these instruments allow the hedger to benefit from lower hedging costs for index-based deals without the adverse impact of material population basis risk on capital relief. Exhaustion points below the 85th percentile, however, do not provide sufficient capital relief anymore and are therefore less attractive. We conclude that call spread options are highly suited for reducing capital charges since they reduce longevity tail risk over a one-year horizon in a cost efficient manner.

In Figure 6, we also recognize that many hedges are dominated in the sense that alternative instruments can be found which offer a higher hedge effectiveness at the same capital efficiency or the other way around. First, we observe that all considered annuity forwards are outperformed by long-term longevity swaps. Certainly, one reason for this lies in the underlying portfolio of immediate annuities. Annuity forwards might be more suitable for a portfolio of deferred annuities which we analyze in Subsection 4.4.2. Second, we find that for most fully customized hedges an index-based alternative can be found which does not necessarily have the same payout structure but offers a higher capital efficiency at a comparable level of risk reduction. In fact, fully customized contracts are only indispensable when aiming at a high hedge effectiveness above 80%. A hedger who is willing to accept population basis risk might benefit from a higher capital efficiency since the additional costs for eliminating socioeconomic basis risk (and small sample risk) might exceed the resulting cost of capital relief. Of course, the outcome of this complex interplay highly depends on the market price of risk for socioeconomic mortality differentials. For lower values of this market price, the company would eventually decide in favor of the fully customized contract. A similar argument applies when considering the subpopulation-linked designs which are outperformed by their fully customized counterparts since they are offered at the same price but provide less protection and therefore also lower cost of capital relief. Again, beyond a certain threshold value for additional customization costs, index-based instruments associated with the subpopulations would become more capital efficient than their fully customized counterparts. As already stated, we refrain from explicitly dealing with potential further customization costs.

In the left panel of Figure 6, we obtain a completely different picture under the standard formula. First of all, it should be kept in mind that the assumed market price of longevity risk is still derived from the internal model's risk assessment, which generally produces lower capital charges compared to the standard formula as we have shown in Section 4.1. For this reason, the standard formula's conservative longevity stress gives a strong incentive for transferring longevity risk. However, the simplified standard formula clearly cannot provide a solid picture across all considered instruments. As already discussed in Subsection 4.2.1, it tends to overestimate the efficiency of longevity swaps, does not properly account for population basis risk, and allows for rather high capital reliefs for option-type contracts with high attachment points. In light of these shortcomings, the effects of hedging should also be analyzed under a risk-based internal model to obtain an adequate picture of the individual risk profile and to avoid suboptimal hedging decisions.

Overall, we conclude that hedgers might face a trade-off between hedge effectiveness and capital efficiency. In particular, differing hedging objectives typically require different types of instruments in terms of hedge horizon, payout structure, and underlying IP. While fully customized contracts naturally outperform their index-based counterparts in terms of hedge effectiveness, predominantly index-based designs offer the highest capital efficiency (at least for our model parametrization). The key finding that no instrument simultaneously outperforms the others in terms of both hedge effectiveness and capital efficiency underlines the importance of considering both aspects for a profound hedging decision.

4.4 Sensitivity analysis

We now analyze how sensitive our conclusions are with respect to different modeling assumptions.

4.4.1 Socioeconomic mortality differentials

The question of how to appropriately model population basis risk is a topic of ongoing discussions. When evaluating index-based hedges in a stochastic simulation, the results highly depend on the implied correlation between the mortality index and the mortality experience of the hedger's portfolio population. Often, mortality differentials between closely related populations are modeled by mean-reverting vector autoregressive (VAR) processes to enforce 'coherent', i.e non-diverging, mortality rates in the long run. However, the assumption of coherence is not always supported by data and might be criticized for being too strong, see for instance Hunt and Blake (2018) or Li et al. (2017).

In this subsection, we address these issues in the context of the risk-based internal model by alternatively assuming that socioeconomic mortality differentials are driven by a first-order VAR process.[17] As opposed to the previously applied difference-stationary RWD, temporary socioeconomic mortality differentials are now assumed to move around a long-term mean-reversion level.

We start by examining the adverse impact of population basis risk on capital relief in index-based hedging. To this end, we define the following relative 'haircut':

$$1 - \frac{ReCoC^{IM}\left(H^{[IP]}\right)}{ReCoC^{IM}\left(H^{[\mathcal{B}]}\right)}, \qquad IP \in \{\mathcal{S}, \mathcal{R}\},$$

which measures the proportionate loss in expected cost of capital relief due to population basis risk for an index-based hedge relative to its fully customized counterpart.

The left panel of Figure 7 shows the haircuts for the index-based instruments for which fully customized counterparts are available in the base case. For the sake of clarity, we only show the effect of lower exhaustion points for a 25-year rolling call spread portfolio linked to $IP = \mathcal{R}$ and refrain from showing the instruments tailored to the standard formula. As already indicated in Subsection 4.3.2, exhaustion points close to the relevant 99.5th percentile typically result in rather large haircuts whereas population basis risk vanishes as the exhaustion points are lowered (from gold to yellow). The figure also shows that the relative haircuts for socioeconomic basis risk are considerably larger than those for small sample risk (see also Subsection 4.2.1). In contrast, the right panel shows the haircuts under the assumption of level-stationary socioeconomic mortality differentials. Apparently, the haircuts for instruments linked to the reference population are significantly smaller compared to the base case. For instance, the haircuts for long-dated longevity swaps reduce from nearly 50% in the base case to just 20% in the right panel. Regarding the rolling call spread portfolios, population basis risk already becomes rather negligible for exhaustion points below the 95th percentile in the right panel, whereas choices below the 90th percentile are required to accomplish the same under the RWD. For the subpopulation-linked hedges, however, the haircuts remain roughly the same since these instruments do not give rise to socioeconomic basis risk. The haircuts among the IPs in the right panel show that socioeconomic basis risk is now less relevant than small sample risk. The observation that the order of relevance has changed

[17] To avoid overparametrization, we apply a restricted VAR(1) model that does not capture autoregressive effects betweeen the time process of distinct subpopulations. For the sake of consistency, we also adapt the counterpart in the EMT valuation model by restricting the derivation of subpopulation-specific experience ratios to the recalibration of mortality levels by means of a weighted average of the most recent data points. Again, we adapt the pricing to the internal model's risk assessment, which requires a market price of longevity risk of $\lambda = 40\%$ (and $\lambda_\epsilon = 20\%$, respectively).

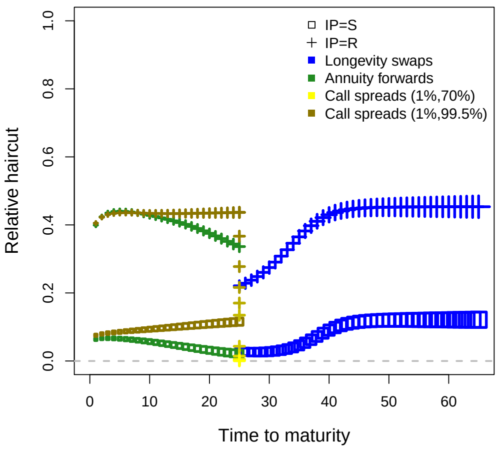

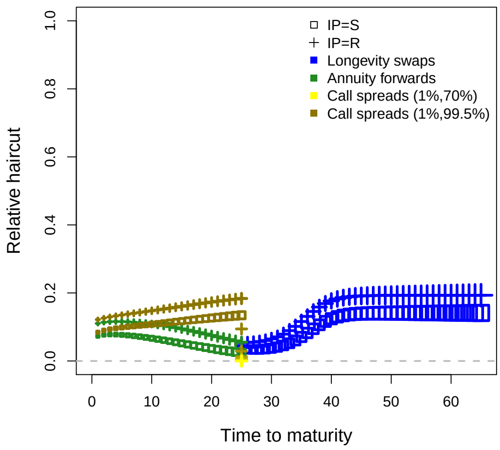

(a) RWD (trend and level adjustment)

(b) VAR(1) (level adjustment)

Figure 7: Relative haircuts for population basis risk under different modeling assumptions.

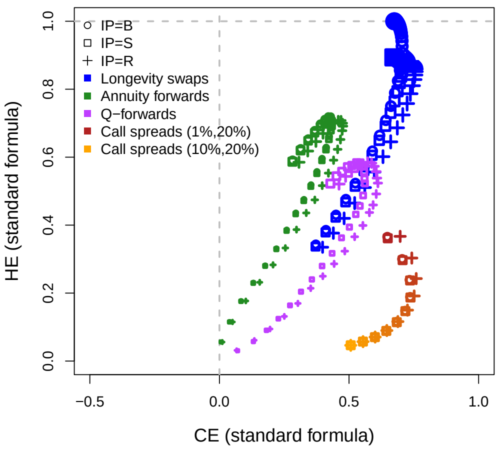

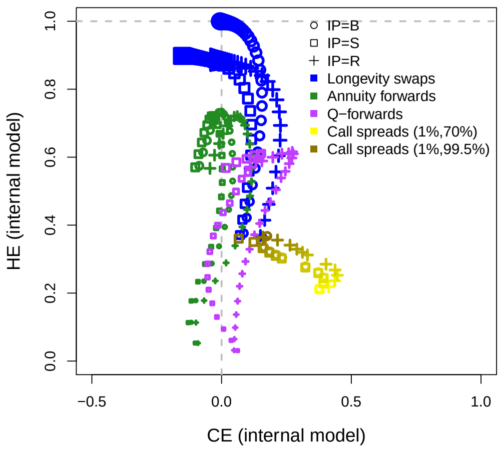

(a) Standard formula

(b) Internal model

Figure 8: Capital efficiency (x-axis) vs. hedge effectiveness (y-axis) assuming a VAR(1) process for modeling socioeconomic mortality differentials.

when imposing coherence on the model clearly indicates a significant model risk with regard to modeling socioeconomic mortality differentials.

This raises the question to what extent this issue impacts our previous results on hedge effectiveness and capital efficiency. For comparison, Figure 8 shows the corresponding charts obtained under the VAR(1) process. As expected, we observe that the distances between instruments linked to different IPs are smaller compared to the base case (in Figure 6) with respect to both dimensions. However, the modeling assumption for socioeconomic mortality differentials hardly affects the 'efficient frontier' in terms of shape and encased hedging instruments. In fact, for practically any strategic hedging objective, the hedger would arrive at the same optimal hedging decision as in the base case. While the hedger has to accept larger haircuts for socioeconomic basis risk under the RWD, the additional costs for customization are also higher since pricing is conducted via a risk-adjusted version of the underlying simulation model. Hence, the hedger can benefit from more attractive prices for index-based hedges which compensates for the larger haircuts. These findings suggest that, as long as the market can find a common agreement on how to model population basis risk, the question which modeling assumption will eventually prevail might after all be of secondary importance for the choice of optimal hedging strategies. Nevertheless, it remains highly relevant for pricing and for the computation of capital reliefs.

4.4.2 Starting age

Finally, we change the underlying liabilities by considering a portfolio of deferred annuities. Figure 9 shows the combined pictures of hedge effectiveness and capital efficiency for a lower starting age of $x_0 = 50$.[18] To account for the longer time horizon, we adjust the contract maturities by considering hedge horizons beyond 40 years for longevity swaps and up to 40 years for the remaining instruments. Since we do not obtain any new insights under the standard formula, we limit our discussion to the internal model.

Compared to the base case, we make three interesting observations: First, index-based longevity swaps do no longer offer a higher capital efficiency than fully customized contracts. As seen in the previous subsection, the hedger needs to accept rather large haircuts for population basis risk when considering long-term index-based longevity swaps. This effect is even more pronounced for longer-dated contracts, which are naturally required for a lower starting age. Second, only a substantially lower hedge effectiveness of around 40% can be reached with q-forwards compared to nearly 65% in the base case. Obviously, a hedge portfolio for long-term obligations requires more than a single q-forward to be effective. Finally, the hedge effectiveness of annuity forwards slightly increases since these value hedge agreements are structurally more suitable for deferred annuities.

So far, we have assumed that investors are willing to offer hedge contracts over any time to maturity. However, long-term investment horizons (for instance over more than half a century) might not be appealing to institutional investors. For simplicity, we also do not explicitly deal with interest rate risk and counterparty credit risk, which are typically more problematic for long-term contracts.

If we assumed that risk takers were only interested in investment horizons of at most 40 years, all longevity swaps (represented by blue dots) in Figure 9 would no longer be available. The only way to reach a hedge effectiveness of around 80% would be to enter into a customized annuity forward with a time to maturity of around 30 years, which generates an expected cost of capital relief that fully compensates for the hedging costs. In fact, the hedger may even benefit from a positive capital efficiency at the cost of some hedge effectiveness by considering an index-based deal. This type of instrument offers a high level of risk reduction compressed to a manageable contract duration in a capital efficient manner and might hence be suitable to reconcile hedger's and investor's interests.

[18] Note that we still assume that the annuity payout will begin at age 65.

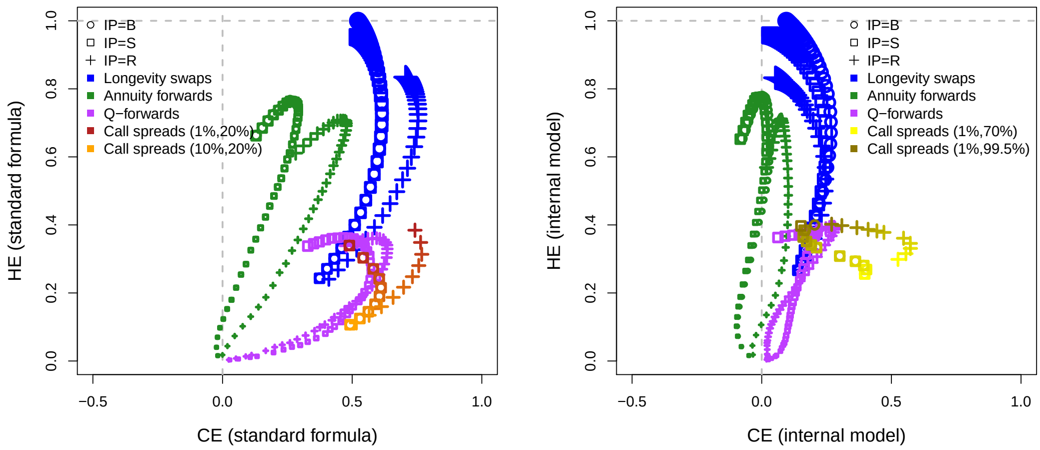

(a) Standard formula (b) Internal model

Figure 9: Capital efficiency (x-axis) vs. hedge effectiveness (y-axis) for a portfolio of deferred annuities with starting age $x_0 = 50$.

5 Conclusion

In this paper, we simultaneously and consistently analyze the two main benefits of transferring longevity risk under risk-based solvency regimes: a reduction in the uncertainty in future cash flows arising from uncertain future mortality (measured in terms of hedge effectiveness) and a cost efficient reduction in capital charges for longevity risk (measured in terms of capital efficiency). Unlike previous studies, we explicitly acknowledge the uncertainty regarding future cost of capital and incorporate its reduction resulting from hedging into the assessment of hedge effectiveness and capital efficiency. For our numerical analyses, we consider a wide selection of different hedging instruments in terms of hedge payout structure, time to maturity, and underlying population. In particular, a clear distinction between the different components of longevity risk within the underlying simulation model allows us to construct different hedge designs which give rise to varying levels of population basis risk.

We find that different hedging instruments have structurally different impacts on the hedger's economic capital in terms of both magnitude and variability. Moreover, we show that the Solvency II standard formula generates less consistent capital reliefs compared to those derived with a risk-based internal model. In particular, the standard formula's simplified one-off longevity stress does not properly account for population basis risk, overestimates the efficiency of limited longevity swaps, and assigns rather high capital reliefs to option-type contracts with high attachment points. Furthermore, we show that hedge effectiveness can change considerably if the uncertainty regarding future capital charges for longevity risk is properly accounted for. Analyses ignoring this source of uncertainty might systematically overestimate the effectiveness of limited longevity swaps, underestimate it for rolling call spread portfolios, and might consequently lead to suboptimal hedging decisions. The fact that misestimation can occur in both directions underlines the necessity to work with stochastic capital charges for a profound analysis and comparison of different hedging instruments.

Assuming the hedge counterparty demands a reasonable market price of longevity risk, we then conduct a combined study of hedge effectiveness and capital efficiency and find that generally no 'universally superior' hedging solution can be found. We visualize the implications of hedging and clearly identify the most efffective and the most capital efficient solutions as two extreme positions. Between them, a frontier of 'efficient' instruments can be identified which provide the highest capital efficiency for a given level of hedge effectiveness. From this set, the hedger may choose suitable contract designs that provide an optimal trade-off between hedge effectiveness and capital efficiency with regard to the strategic hedging objective. Furthermore, we address the benefits and costs of customization and the accompanying trade-off between mitigating (or even fully eliminating) population basis risk and lower costs of hedging. While customized hedges naturally outperform their index-based counterparts in terms of hedge effectiveness, in many cases cost efficient index-based designs provide the highest capital efficiency.

Finally, we demonstrate that altering the model with respect to socioeconomic mortality differentials can materially impact the haircut for population basis risk. However, the efficient frontier is mostly unaffected. It is slightly shifted, but includes virtually the same instruments for different levels of desired hedge effectiveness. These findings suggest that the model for population basis risk does only have a minor impact on finding optimal hedging instruments. To conclude, our findings should draw more attention to standardized instruments, which are expected to be more appealing to institutional investors.

References

Bauer, D., Börger, M., and Ruß, J. (2010). On the pricing of longevity-linked securities. *Insurance: Mathematics and Economics*, 46(1):139–149.

Biffis, E., Blake, D., Pitotti, L., and Sun, A. (2016). The cost of counterparty risk and collateralization in longevity swaps. *The Journal of Risk and Insurance*, 83(2):387–419.

Blake, D., Cairns, A. J. G., Dowd, K., and Kessler, A. R. (2019). Still living with mortality: The longevity risk transfer market after one decade. *British Actuarial Journal*, 24,e1:1–80.

Börger, M. (2010). Deterministic shock vs. stochastic value-at-risk — an analysis of the Solvency II standard model approach to longevity risk. *Blätter der DGVFM*, 31(2):225–259.

Börger, M., Fleischer, D., and Kuksin, N. (2014). Modeling the mortality trend under modern solvency regimes. *ASTIN Bulletin*, 44(1):1–38.

Börger, M. and Schupp, J. (2018). Modeling trend processes in parametric mortality models. *Insurance: Mathematics and Economics*, 78:369–380.

Börger, M., Schupp, J., and Ruß, J. (2019). It takes two: Why mortality trend modeling is more than modeling one mortality trend. Working Paper, Ulm University.

Cairns, A. J. G. (2013). Robust hedging of longevity risk. *The Journal of Risk and Insurance*, 80(3):621–648.

Cairns, A. J. G., Blake, D., and Dowd, K. (2006). A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *The Journal of Risk and Insurance*, 73(4):687–718.

Cairns, A. J. G., Dowd, K., Blake, D., and Coughlan, G. D. (2014). Longevity hedge effectiveness: A decomposition. *Quantitative Finance*, 14(2):217–235.

Cairns, A. J. G. and El Boukfaoui, G. (2019). Basis risk in index based longevity hedges: A guide for longevity hedgers. North American Actuarial Journal.

Coughlan, G. D., Khalaf-Allah, M., Ye, Y., Kumar, S., Cairns, A. J. G., Blake, D., and Dowd, K. (2011). Longevity hedging 101: A framework for longevity basis risk analysis and hedge effectiveness. *North American Actuarial Journal*, 15(2):150–176.

Freimann, A. (2020). Pricing longevity-linked securities in the presence of mortality trend changes. Working Paper, Ulm University.

Haberman, S., Kaishev, V., Millossovich, P., Villegas, A. M., Baxter, S., Gaches, A., Gunnlaugsson, S., and Sison, M. (2014). Longevity basis risk: A methodology for assessing basis risk. Institute and Faculty of Actuaries, Sessional Research Paper.

Hunt, A. and Blake, D. (2018). Identifiability, cointegration and the gravity model. *Insurance: Mathematics and Economics*, 78:360–368.

Leung, M., Fung, M. C., and O'Hare, C. (2018). A comparative study of pricing approaches for longevity instruments. *Insurance: Mathematics and Economics*, 82:95–116.

Li, J. S.-H., Chan, W.-S., and Zhou, R. (2017). Semicoherent multipopulation mortality modeling: The impact on longevity risk securitization. *The Journal of Risk and Insurance*, 84(3):1025–1065.

Li, J. S.-H. and Hardy, M. R. (2011). Measuring basis risk in longevity hedges. *North American Actuarial Journal*, 15(2):177–200.

Li, J. S.-H. and Luo, A. (2012). Key q-duration: A framework for hedging longevity risk. *ASTIN Bulletin*, 42(2):413–452.

Liu, Y. and Li, J. S.-H. (2016). The locally linear Cairns-Blake-Dowd model: A note on delta-nuga hedging of longevity risk. *ASTIN Bulletin*, 47(1):79–151.

Liu, Y. and Li, J. S.-H. (2019). An efficient method for mitigating longevity value-at-risk. *North American Actuarial Journal*.

Lu, J. L. C., Wong, W., and Bajekal, M. (2014). Mortality improvement by socio-economic circumstances in England (1982 to 2006). *British Actuarial Journal*, 19(1):1–35.

Meyricke, R. and Sherris, M. (2014). Longevity risk, cost of capital and hedging for life insurers under Solvency II. *Insurance: Mathematics and Economics*, 55:147–155.

Ngai, A. and Sherris, M. (2011). Longevity risk management for life and variable annuities: The effectiveness of static hedging using longevity bonds and derivatives. *Insurance: Mathematics and Economics*, 49(1):100–114.

Plat, R. (2011). One-year value-at-risk for longevity and mortality. *Insurance: Mathematics and Economics*, 49(3):462–470.

Richards, S. J., Currie, I. D., Kleinow, T., and Ritchie, G. P. (2020). Longevity trend risk over limited time horizons. *Annals of Actuarial Science*.

Richards, S. J., Currie, I. D., and Ritchie, G. P. (2014). A value-at-risk framework for longevity trend risk. *British Actuarial Journal*, 19(1):116–167.

Villegas, A. M. and Haberman, S. (2014). On the modeling and forecasting of socioeconomic mortality differentials: An application to deprivation and mortality in England. *North American Actuarial Journal*, 18(1):168–193.

Villegas, A. M., Haberman, S., Kaishev, V. K., and Millossovich, P. (2017). A comparative study of two-population models for the assessment of basis risk in longevity hedges. *ASTIN Bulletin*, 47(3):631–679.

Zeddouk, F. and Devolder, P. (2019). Pricing of longevity derivatives and cost of capital. *Risks*, 7(2):1–29.

Zhou, K. Q. and Li, J. S.-H. (2017). Dynamic longevity hedging in the presence of population basis risk: A feasibility analysis from technical and economic perspectives. *The Journal of Risk and Insurance*, 84(S1):417–437.

Zhou, K. Q. and Li, J. S.-H. (2019). Delta-hedging longevity risk under the M7–M5 model: The impact of cohort effect uncertainty and population basis risk. *Insurance: Mathematics and Economics*, 84:1–21.

Appendices

A Model specification

A.1 Specification of the multi-population AMT/EMT framework

Our relative multi-population extension of the CBD model assumes the following relation between the reference population and the subpopulations:

$$q_{x,t}^{[p]} = \phi\left(q_{x,t}^{[R]}, \xi_{x,t}^{[p]}\right) := \text{logit}^{-1}\left(\text{logit}\left(q_{x,t}^{[R]}\right) + \xi_{x,t}^{[p]}\right), \qquad p \in \{1, \ldots, N_{Sub}\},$$

where the function ϕ links the subpopulation-specific mortality rates to those of the reference population within the CBD model structure via

$$\xi_{x,t}^{[p]} := \kappa_t^{(1)[p]} + (x - \bar{x})\,\kappa_t^{(2)[p]}.$$

These so-called experience ratios capture differing mortality levels and trends in the subpopulations relative to the reference population. Throughout the paper, we extrapolate mortality rates up to a limiting age of 130 by following the inherent logit-extrapolation in the CBD model structure. To avoid an unrealistic reversal of the ranking among the subpopulations for highest ages, we apply the following adjustment:

$$q_{x,t}^{[p]} := \begin{cases} \max\left\{q_{x,t}^{[R]}; q_{x,t}^{[p]}\right\}, & \text{for } p = 1, 2 \\ \min\left\{q_{x,t}^{[R]}; q_{x,t}^{[p]}\right\}, & \text{for } p = 4, 5 \end{cases}$$

and leave the middle subpopulation unadjusted allowing it to take values above or below the reference population.

In the EMT valuation model, best estimate future mortality rates beyond time T are projected for the reference population by following a deterministic central path of mortality

$$\tilde{q}_{x,t}^{[R]}(T) := \text{logit}^{-1}\left(\tilde{\kappa}_T^{(1)[R]} + (t - T)\,\tilde{d}_T^{(1)[R]} + (x - \bar{x})\left(\tilde{\kappa}_T^{(2)[R]} + (t - T)\,\tilde{d}_T^{(2)[R]}\right)\right), \qquad t > T,$$

where $\tilde{\kappa}_T^{(i)[R]}, i = 1, 2$ denote the estimated mortality levels and $\tilde{d}_T^{(i)[R]}, i = 1, 2$ denote the EMT derived at time T. We follow Börger et al. (2019) and estimate these quantities for both time processes independently by means of a weighted linear regression on the most recent realizations of the CBD time processes.

Analogously, mortality rates for each subpopulation $p \in \{1, \ldots, N_{Sub}\}$ are projected as

$$\tilde{q}_{x,t}^{[p]}(T) := \phi\left(\tilde{q}_{x,t}^{[R]}(T), \tilde{\xi}_{x,t}^{[p]}(T)\right), \qquad t > T$$

based on the time-T best estimate experience ratios

$$\tilde{\xi}_{x,t}^{[p]}(T) := \tilde{\kappa}_T^{(1)[p]} + (t - T)\,\tilde{\nu}_T^{(1)[p]} + (x - \bar{x})\left(\tilde{\kappa}_T^{(2)[p]} + (t - T)\,\tilde{\nu}_T^{(2)[p]}\right), \qquad t > T,$$

which are also derived at time T from observed mortality in previous years and projected via a deterministic linear projection. In line with the underlying CBD model structure, these experience ratios consist of subpopulation-specific adjustments for different mortality levels $\tilde{\kappa}_T^{(i)[p]}, i = 1, 2$ and trends $\tilde{\nu}_T^{(i)[p]}, i = 1, 2$, respectively. For the sake of consistency, we apply the same weighted linear regression as for the reference population. For further technical details, we refer to Freimann (2020).

A.2 Construction of index-based hedging instruments

In this section, we provide further details on the hedging instruments which are introduced in Section 3.2. In particular, we construct index-based designs based on suitable hedge indices.

Mortality indices and q-forwards

For each subpopulation $p \in \{1, \dots, N_{Sub}\}$, we define the following mortality indices:

$$q_{x,t}^{[IP(p)]} := \begin{cases} q_{x,t}^{[p]}, & IP = \mathcal{S} \\ \phi\left(q_{x,t}^{[R]}, \tilde{\xi}_{x,t}^{[p]}(0)\right), & IP = \mathcal{R}, \end{cases}$$

which either represents the subpopulation-specific mortality rates (in case of $IP = \mathcal{S}$) or a proxy (in case of $IP = \mathcal{R}$) which is constructed by adjusting the mortality rates of the reference population by the initial experience ratios to account for anticipated mortality differences between the populations. The payout of the q-forward contract is then based on

$$Q_{x_0+\tau,\tau}^{[IP]} := \sum_{p=1}^{N_{Sub}} n_\tau^{[p]} q_{x_0+\tau,\tau}^{[IP(p)]}, \qquad IP \in \{\mathcal{S}, \mathcal{R}\},$$

where $n_\tau^{[p]} \geq 0$ represents the hedge ratio for subpopulation $p \in \{1, \dots, N_{Sub}\}$. For a given risk measure ρ and maturity τ, we determine the optimal hedge ratios via the objective function

$$\underset{n_\tau^{[p]}}{\operatorname{argmin}}\ \rho\left(L_H^{[p]}(0) - \mathbb{E}\left(L_H^{[p]}(0)\right)\right), \qquad p \in \{1, \dots, N_{Sub}\},$$

where $L_H^{[p]}(0) := L^{[p]}(0) - (1+r)^{-\tau}\, n_\tau^{[p]}\left(\mathbb{E}^{\mathbb{Q}}\left(q_{x_0+\tau,\tau}^{[IP(p)]}\right) - q_{x_0+\tau,\tau}^{[IP(p)]}\right)$. This objective function is solved numerically for each subpopulation $p \in \{1, \dots, N_{Sub}\}$ over all outer simulation paths.[19] To ensure comparability between the two designs, we determine the optimal hedge ratios for $IP = \mathcal{S}$ and choose the same hedge ratios for $IP = \mathcal{R}$.

Survivor indices and longevity swaps

The fully customized longevity swap is based on the actual number of living policyholders $S_{x_0+t,t}^{[\mathcal{B}]} := \sum_{p=1}^{N_{Sub}} B_{x_0+t,t}^{[p]}$ at time t. The index-based designs are then constructed by using the following survivor indices:

$$S_{x_0+t,t}^{[IP(p)]} := B_{x_0,0}^{[p]} \prod_{s=0}^{t-1}\left(1 - q_{x_0+s,s+1}^{[IP(p)]}\right), \qquad IP = \mathcal{S}, \mathcal{R},$$

which build on the previously introduced mortality indices.

[19]Note that the hedge ratios are defined based on the unadjusted liabilities instead of the adjusted liabilities, which constitutes a slight inconsistency but makes the derivation of the optimal hedge ratios practically feasible.

Liability indices and annuity forwards

The fully customized annuity forward is based on the actual time-τ best estimate liabilities

$$\widetilde{L}^{[\mathcal{B}]}(\tau) := \sum_{p=1}^{N_{Sub}} B^{[p]}_{x_0+\tau,\tau} \sum_{s>\tau} (1+r)^{-(s-\tau)} \mathbb{1}_{\{x_0+s\geq x_R\}} \prod_{u=\tau}^{s-1} \left(1 - \tilde{q}^{[p]}_{x_0+u,u+1}(\tau)\right).$$

For each subpopulation $p \in \{1, \ldots, N_{Sub}\}$, we define the following annuity indices:

$$\widetilde{a}^{[IP(p)]}_{x_0+\tau,\tau}(\tau) := \sum_{s>\tau} (1+r)^{-(s-\tau)} \mathbb{1}_{\{x_0+s\geq x_R\}} \prod_{u=\tau}^{s-1} \left(1 - \tilde{q}^{[IP(p)]}_{x_0+u,u+1}(\tau)\right), \qquad IP \in \{\mathcal{S}, \mathcal{R}\},$$

which denotes (in case of $IP = \mathcal{S}$) the time-τ best estimate value of a lifelong annuity of one paid annually in advance starting at the age of x_R for an individual from subpopulation $p \in \{1, \ldots, N_{Sub}\}$ aged $x_0 + \tau$ at time τ or a suitable proxy (in case of $IP = \mathcal{R}$). The index-based annuity forwards are then constructed by replacing the actual number of survivors by the survivor indices and the actual annuity value by the annuity indices:

$$\widetilde{L}^{[IP]}(\tau) := \sum_{p=1}^{N_{Sub}} S^{[IP(p)]}_{x_0+\tau,\tau}\, \widetilde{a}^{[IP(p)]}_{x_0+\tau,\tau}(\tau), \qquad IP \in \{\mathcal{S}, \mathcal{R}\}.$$

Call spread indices

Recall from Section 2.3 that the SCR at time t is basically determined by the randomness (or change due to a longevity shock) in

$$\begin{aligned}\widetilde{X}^{[\mathcal{B}]}(t+1) :=& CF(t+1) + \widetilde{L}(t+1) = \sum_{p=1}^{N_{Sub}} B^{[p]}_{x_0+t+1,t+1} \mathbb{1}_{\{x_0+t+1\geq x_R\}} \\ &+ \sum_{p=1}^{N_{Sub}} B^{[p]}_{x_0+t+1,t+1} \sum_{s>t+1} (1+r)^{-(s-(t+1))} \mathbb{1}_{\{x_0+s\geq x_R\}} \prod_{u=t+1}^{s-1} \left(1 - \tilde{q}^{[p]}_{x_0+u,u+1}(t+1)\right),\end{aligned}$$

on which the fully customized call spread contracts are based. The index-based designs are then constructed as follows:

- For the subpopulation-linked design, $\widetilde{X}^{[\mathcal{S}]}(t+1)$ is defined analogously by replacing the actual number of survivors in the book population $B^{[p]}_{x_0+t+1,t+1}$ by $B^{[p]}_{x_0+t,t}\left(1 - q^{[p]}_{x_0+t,t+1}\right)$ so that small sample risk over the one-year hedge horizon is no longer covered.
- Finally, $\widetilde{X}^{[\mathcal{R}]}(t+1)$ is defined structurally similarly to $\widetilde{X}^{[\mathcal{S}]}(t+1)$ by fixing the time-t experience ratios to also remove the uncertainty originating from the subpopulation over the year. More precisely, for each subpopulation $p \in \{1, \ldots, N_{Sub}\}$, realized mortality $q^{[p]}_{x_0+t,t+1}$ is replaced by $\phi\left(q^{[R]}_{x_0+t,t+1}, \tilde{\xi}^{[p]}_{x_0+t,t+1}(t)\right)$ and best estimate mortality $\tilde{q}^{[p]}_{x_0+u,u+1}(t+1)$ is replaced by $\phi\left(\tilde{q}^{[R]}_{x_0+u,u+1}(t+1), \tilde{\xi}^{[p]}_{x_0+u,u+1}(t)\right)$ for $u \geq t+1$.