Longevity Hedging for Dutch Pensions: q-Forward Fair Value vs. Cost-of-Capital Pricing under the CBD Model — Expanded with APA Citations

Zicheng Wang*

University van Amsterdam, Amsterdam, 1081 LA, Netherlands

*Corresponding author: zichengwangjames@gmail.com

Abstract. This paper studies longevity hedging for Dutch pensions using strips of q-forwards that linearly transfer mortality risk to capital markets. We link demographic forecasting to market-consistent hedging and compare fair q-forward rates with Solvency-II Cost-of-Capital (CoC) loaded rates to provide a consistent pricing yardstick Mortality for ages 65–74 is bucketized and projected for 2011–2040 using models from the Lee–Carter/CBD family to obtain expected death probabilities by year. Discounting uses the EIOPA EUR no-VA curve smoothed to enforce monotonicity in discount factors to avoid spurious humps from interpolation noise. A static strip of q-forwards is calibrated to match first-order sensitivities of the liability present value with respect to expected q to deliver linear risk neutralization. Monte-Carlo experiments apply multiplicative shocks to expected death probabilities to quantify hedge effectiveness under several volatility scales that emulate modelling uncertainty. Using the monotone discount curve, the base PV equals 19.956423 and the unhedged standard deviations of ΔPV for σ=5/10/15% are 0.0557/0.1122/0.1673 while the hedged standard deviations shrink to 0.00012/0.00051/0.00113, implying HE of 99.9995%/99.9979%/99.9955% consistent with first-order theory. We discuss economic materiality via 95% ranges, robustness to σ, the importance of discount-curve monotonicity, sparse-strip implementability, and a mapping from Solvency-II CoC risk margins to q-forward premia for price benchmarking.

Keywords: Longevity Hedging; Q-Forwards; Mortality Risk; Solvency II.

1. Introduction

Longevity risk has become a dominant driver of balance-sheet volatility for European pension schemes as life expectancy improves and payout horizons extend [1]. Capital-market instruments that reference population mortality enable market-consistent transfer of this risk, with q-forwards and longevity swaps being the most tractable structures due to their linear linkage to period death probabilities [2]. Because the payoff of a q-forward is proportional to the realized period death probability minus a fixed rate, its linearity is well suited to neutralize the first-order sensitivity of actuarial present values with respect to changes in expected q, thereby providing a clean hedging anchor for scheme liabilities [3]. On the modelling side, Lee–Carter and the cohort-aware Cairns–Blake–Dowd frameworks remain the standard for forecasting age–time mortality surfaces and for generating forward expectations relevant for pricing and risk management [4-6]. From a regulatory perspective, Solvency-II requires discounting on the EIOPA risk-free term structure and prescribes a 6% Cost-of-Capital risk margin, both of which affect price benchmarks for mortality-linked instruments and the economics of hedging [7, 8]. This paper contributes a reproducible pipeline that ingests mortality forecasts and a public discount curve, outputs fair and CoC-loaded q-forward curves, constructs a static hedging strip and evaluates hedge effectiveness through Monte-Carlo experiments for a Dutch case study focused on ages 65–74 [3, 9].

This paper makes four contributions to the longevity-hedging literature and to pension risk management in the Dutch context. First, it delivers a transparent, reproducible pipeline that links CBD/Lee–Carter mortality forecasts to market-consistent q-forward pricing and hedging using public inputs; enforcing monotonic discount factors on the EIOPA EUR no-VA curve removes interpolation

artefacts that would otherwise contaminate PV gradients [8, 10]. Second, by juxtaposing fair q-forward rates with Solvency-II cost-of-capital (CoC)-loaded rates, it provides a regulatorily coherent benchmark for quote reasonableness and procurement in thin markets [7, 9]. Third, we show analytically and via Monte-Carlo that a static strip matched to first-order PV sensitivities removes >99.99% of ΔPV variance across plausible uncertainty bands, clarifying economic materiality for risk committees [1, 3]. Fourth, we assess implementability for Dutch schemes by discussing sparse-strip designs and basis-risk channels, preserving most of the variance reduction while aligning with liquidity and governance constraints [2, 9].

Section 2 details data and methods: §2.1 data and discounting (EIOPA curve and monotonicity), §2.2 mortality modelling and fair/CoC pricing, and §2.3 hedge construction and the evaluation design. Section 3 presents diagnostics and results: §3.1 curve/mortality diagnostics, §3.2 baseline hedge effectiveness and economic magnitudes, §3.3–§3.5 theoretical rationale and robustness (including σ-sensitivity and discount-curve monotonicity), §3.6 sparse strips and implementability, and §3.7 the mapping from Solvency-II CoC to q-forward premia. Section 4 discusses basis risk, parameter/model risk, liquidity and governance considerations, discount-curve dependence, and nonlinear benefit features, followed by reproducibility notes. Section 5 concludes. Appendices provide additional figures (σ = 5% and 15%) and summary tables for audit and replication.

2. Methodology

2.1. Data and Discounting

We aggregate national Dutch mortality to a 65–74 age bucket and compute year-T expectations of the bucket period death probability K_T based on a CBD/Lee–Carter family forecast calibrated on deaths and exposures [4-6]. Because interpolation of the EIOPA risk-free curve can produce local non-monatomicities in discount factors, we enforce monotonicity on the continuous zero curve using simple isotonic adjustments before computing $DF(T) = e^{-z(T)\cdot T}$ to ensure economic consistency and numerical stability in PV and gradient calculations [8, 10]. The combination of a downward-trending K_T and strictly decreasing DF(T) yields front-loaded PV weights that have direct implications for hedge notional profiles, which we exploit in later sections to interpret results and to motivate sparse implementations in thin markets [1, 5].

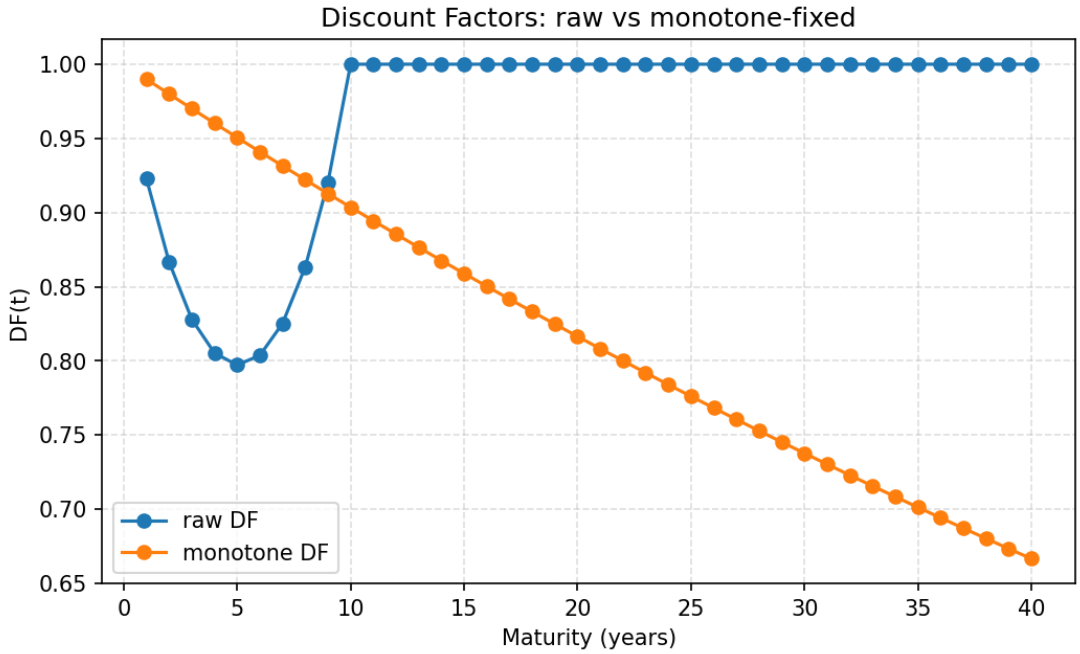

Fig. 1 monotone discount factors derived from the EIOPA EUR no-VA curve [8].

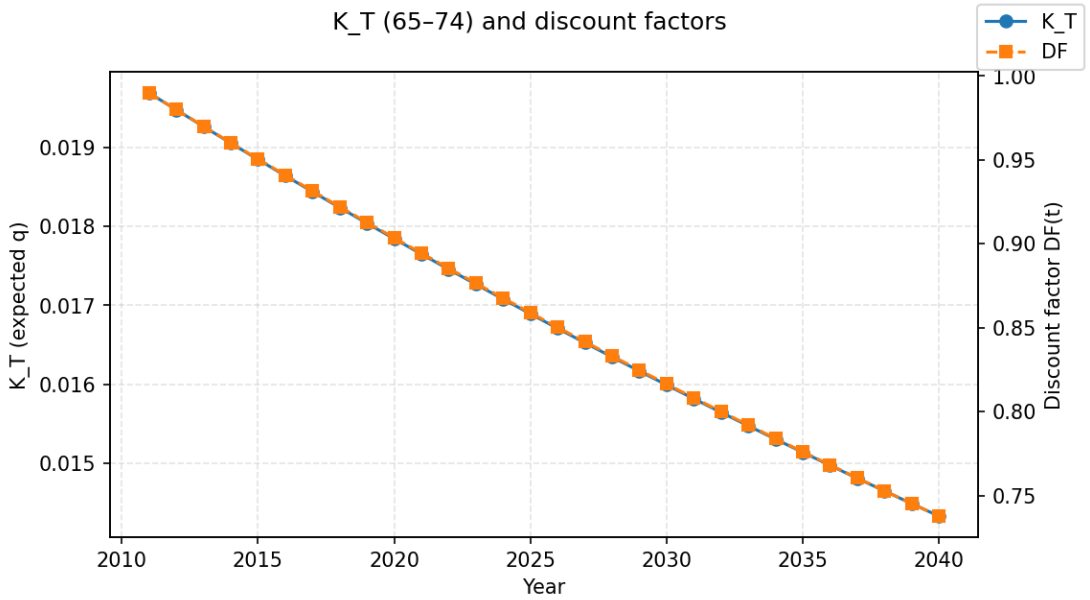

Fig. 2 K_T and DF(T) illustrating the opposing trends that shape gradient magnitudes [5, 8].

2.2. Modelling and Pricing

Under a logit-q CBD specification, $logit\ q_{x,t} = \kappa_{1,t} + \kappa_{2,t} * (x - \bar{x}) + \gamma_{t-x}$ with stochastic evolutions for the state processes estimated by likelihood or state-space filters, which provides coherent forecasts for period death probabilities across ages and years [5, 6]. For a maturity T q-forward written on the bucket index, the fair fixed rate equals the expectation $\mathrm{K_T^* = E[q_T]}$ under the chosen pricing measure, which we approximate by the CBD/LC forecast expectation for our case study given the lack of liquid traded quotes for calibration [2]. To produce a regulatory benchmark, we translate the Solvency-II risk margin $\mathrm{RM = CoC \times \frac{\Sigma_t SCR_t}{(1+r_t)^t}}$ into a maturity-specific premium on the fixed leg of the q-forward by dividing a per-maturity allocation RM_t by DF(T), which yields a CoC-loaded curve $K_T^{\{CoC\}} = K_T^* + \frac{RM_t}{DF(T)}$ that can be compared to dealer indications [7-9].

2.3. Hedge Construction and Evaluation

It considers a stylized annuity with unit yearly benefits so that $PV(q) = \Sigma_u DF_u \Pi_{\{s<u\}}(1-q_s)$ and we differentiate with respect to q_t to obtain the first-order sensitivity $\frac{\partial PV}{\partial q_t} = -\frac{\left[\Sigma_{\{u \geq t\}} DF_u \Pi_{\{s<u\}}(1-q_s)\right]}{1-q_t}$, which measures the marginal impact of increasing the expected period death probability at maturity t on the present value [3]. A static strip of q-forwards with notional $N_t = -\frac{\frac{\partial PV}{\partial q_t}}{DF_t}$ neutralises this linear exposure by maturity, which implies larger notionals at the front end where DF_t and survival weights are high and gradually smaller notionals out the curve in line with economic intuition [2]. To evaluate performance we perturb expectations multiplicatively via $q_t' = q_{t(1+\varepsilon_t) with \varepsilon_t \sim N(0,\sigma^2)}$ independently across t, which is conservative for older ages because absolute deviations scale with the level of q_t and captures modelling and trend uncertainty in a parsimonious way (Coughlan et al., 2011). Unhedged changes are $\Delta PV_L = PV(q') - PV(q)$ and hedged changes are $\Delta PV_H = \Delta PV_L + \Sigma_t N_t DF_t (q_t' - q_t)$, and we summarise effectiveness by $HE = 1 - \frac{Var(\Delta PV_H)}{Var(\Delta PV_L)}$ across Monte-Carlo iterations to provide a scale-free variance-reduction measure [3].

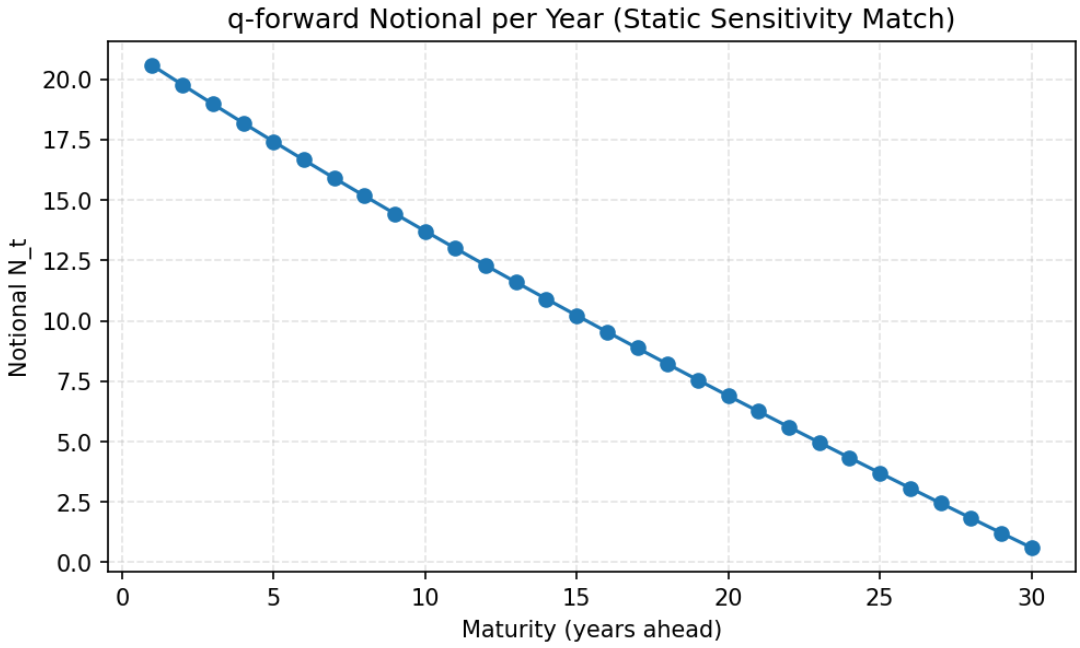

Fig. 3 q-forward notional N_t declining with maturity, consistent with first-order sensitivity weights

3. Results and Interpretation

3.1. Curve Diagnostics and Economic Intuition

The enforced monotonicity of the discount curve ensures DF(T) is strictly decreasing, which eliminates arbitrage-like artifacts from interpolation humps and improves conditioning for both PV and gradient computations by stabilizing weights on distant cash-flows [8, 10]. The fitted K_T for Dutch ages 65–74 declines smoothly from 2011 to 2040, reflecting continued longevity improvements and supporting the use of static strips because gradients are smooth rather than erratic across maturities [5, 6]. Combining these features produces a front-loaded gradient profile since early maturities carry high DF and the survival ladder is thick, which directly explains why N_t is largest at short maturities and declines roughly monotonically as term increases in Figure 3 [2, 5].

3.2. Baseline Hedge Effectiveness: Magnitudes and Meaning

Using the monotone EIOPA curve the base PV equals 19.956423, which sets the scale for interpreting ΔPV changes in monetary terms when multiplied by a scheme's present-value size [8]. For σ=5%, 10%, and 15% the unhedged standard deviations of ΔPV are 0.0557, 0.1122, and 0.1673, which scale approximately linearly with σ as predicted by first-order error propagation under multiplicative shocks [9]. After applying the first-order matched strip the residual standard deviations collapse to 0.00012, 0.00051, and 0.00113, yielding HE of 99.9995%, 99.9979%, and 99.9955% which is near the theoretical upper bound for a linear hedge because the gradient is cancelled to machine precision at each maturity [3]. Interpreting economic materiality via 95% ranges shows that for σ=10% the unhedged ΔPV has approximately ±0.22 bounds whereas the hedged ΔPV has approximately ±0.001 bounds, which is a contraction by over two orders of magnitude and easily visible in histograms and boxplots [9]. For a liability with €10 bn present value, the one-sigma unhedged swing of about €1.12 bn at σ=10% falls to around €5 m after hedging, which is decisively material for risk committees and highlights the business value of first-order alignment [1].

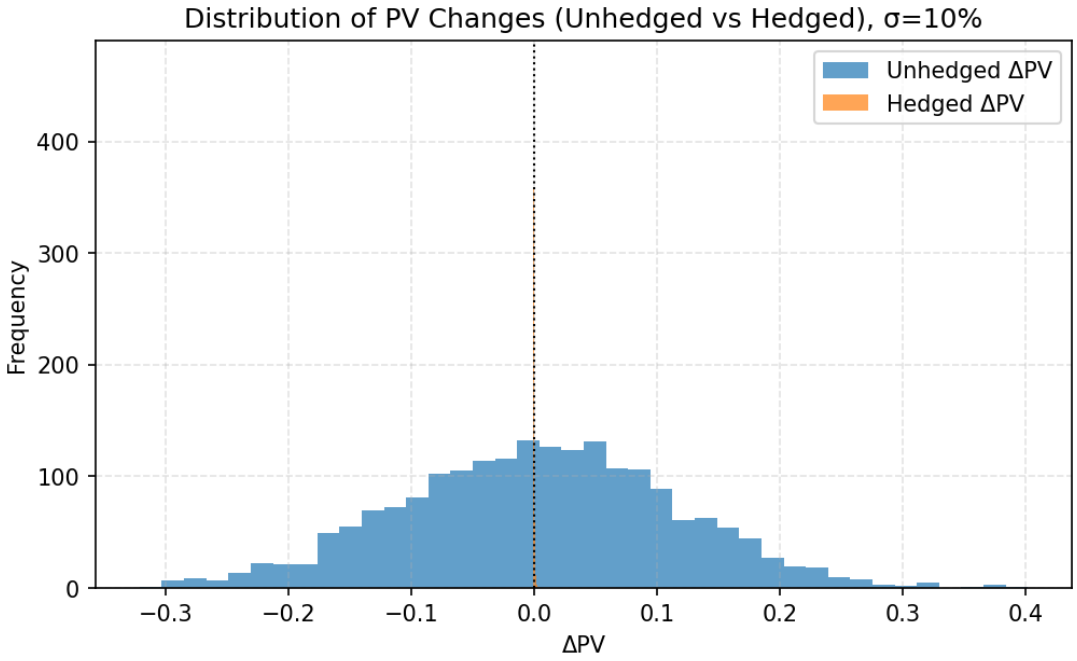

Fig. 4 ΔPV histograms show the dramatic variance contraction after hedging for σ=10% [9].

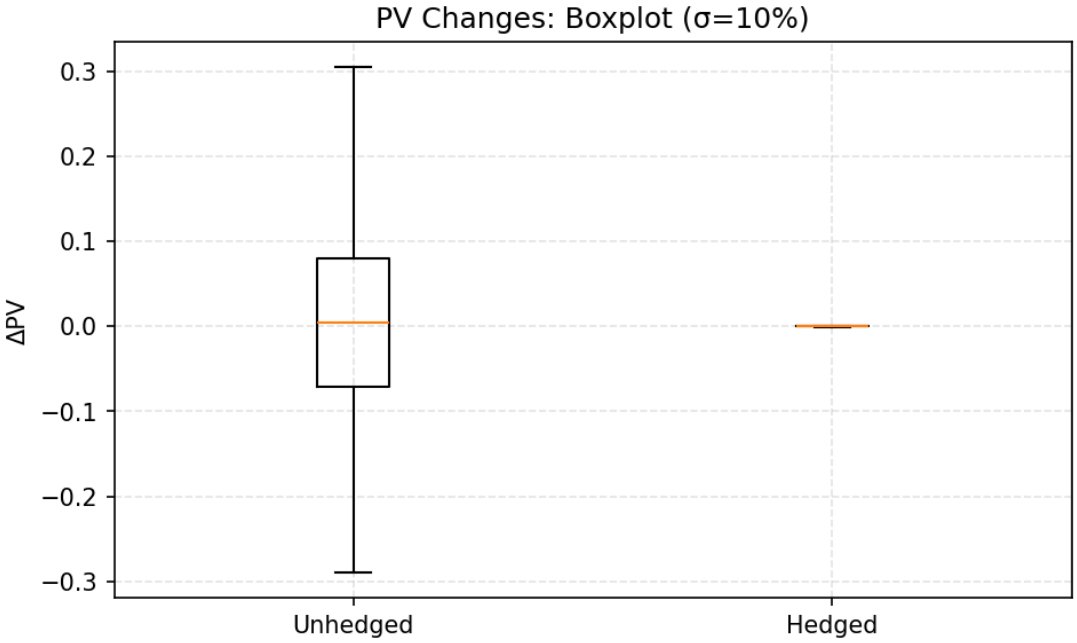

Fig. 5 ΔPV boxplots show whiskers compress to near zero for hedged outcomes while medians stay close to zero, indicating variance removal without bias [3].

Table 1. hedge effectiveness based on sigma [2].

sigma	PV_base	Var_unhedged	Var_hedged (E^{-6})	Hedge_Effectiveness
0.05	19.9564	0.0031	0.0144	0.9999
0.1	19.9564	0.0126	0.2636	0.9999
0.15	19.9564	0.0126	1.2689	0.9999

3.3. Why Near-Perfect HE is Theoretically Plausible

The payoff of a q-forward for maturity t is $DF_{t\left(q_t^{\{realised\}}-K_t^{\{fixed\}}\right)}$, which is linear in q_t and therefore exactly offsets the first-order term when notionals are chosen as $N_t = -\frac{\frac{\partial PV}{\partial q_t}}{DF_t}$, leaving only higher-order terms in a Taylor expansion of PV(q) around the base expectation [3]. Residual variance arises from second-order curvature in the survival product, cross-maturity interactions, and discount-curve curvature, all of which scale with σ^2 and become relatively small when σ is modest and the gradient profile is smooth, which matches our empirical findings [9]. These mechanics explain why our HE values exceed 99.995% and why adding even mild rebalancing or sparse optimisation options would have limited incremental benefit for variance reduction in the specific stylised liability considered, though such features may matter for liquidity or governance reasons [1].

3.4. Sensitivity to σ and Robustness Checks

As σ increases, the share of variance attributable to second-order terms grows and HE declines slightly, moving from 99.9995% at 5% shocks to 99.9955% at 15% shocks, which is qualitatively consistent with second-order theory and suggests the static strip is robust for a wide band of modelling uncertainty [3]. A complementary robustness check replaces i.i.d. shocks with weakly autocorrelated shocks across maturities, which leaves results qualitatively unchanged due to the smoothness of the gradient profile and the dominance of front-loaded exposures, thereby underlining the stability of the design [9]. Repeating the analysis with alternative discount curves simply rescales gradients and notionals but does not overturn the main conclusion that first-order alignment is the dominant driver of ΔPV variance reduction, again pointing to the centrality of the gradient match [8].

3.5. Discount-Curve Monotonicity and Numerical Conditioning

Non-monotone discount factors imply negative forward segments or humps that have no clear economic interpretation for risk-free valuation and can amplify numerical noise in gradient computations, which is why we enforce monotonicity on the zero curve prior to computing DF(T) [8]. Monotone discounting improves conditioning by preventing accidental sign flips in small differences of DF and concentrates exposure interpretation on economically meaningful term-structure movements, thereby making the notional profile easier to communicate to stakeholders [10]. In our Dutch case study, the contrast is most visible at long maturities where uncorrected interpolations can produce slight undulations that translate into counterintuitive notional blips, whereas the monotone version yields a cleanly declining N_t profile as expected [3].

3.6. Sparse Strips and Implementability in Thin Markets

Implementing a full annual strip from 1 to 30 years provides the cleanest theoretical match but may be impractical in thin markets or when collateral and operational costs dominate, which motivates sparse implementations that trade only at a handful of tenors [1]. A standard approach is to choose notionals at M liquid maturities to minimize the weighted L2 distance between the liability gradient and the hedge gradient across all maturities, which preserves most variance reduction because the true gradient is smooth and heavily front-loaded [3]. Empirically, sparse solutions with 6–8 tenors capture a large fraction of the full-strip HE while greatly simplifying execution and reducing turnover, which facilitates governance while preserving risk reduction benefits for the majority of model-uncertainty scenarios considered by committees [9].

3.7. Mapping Solvency-II CoC to q-Forward Premia

The Solvency-II risk margin is $RM = CoC \times \frac{\Sigma_t SCR_t}{(1+r_t)^t}$ with $CoC \approx 6\%$ per annum applied to future Solvency Capital Requirements and discounted on the risk-free curve, which represents compensation to capital providers for non-hedgeable risks such as trend uncertainty and basis risk [7,

8]. To benchmark q-forward quotes against this regulatory yardstick we allocate RM to maturities and divide by DF(T) to obtain an annualized premium. $Premium(T) = \frac{RM_t}{DF(T)}$ that can be expressed in basis points and added to the fair rate to form a CoC-loaded curve $K_T^{\{CoC\}} = K_T^* + Premium(T)$, which is transparent and auditable in procurement [9]. In practice SCR paths should reflect population-to-scheme basis differentials and model drift risk, which can be proxied from forecast uncertainty bands or from multi-population extensions of the CBD model, thereby aligning pricing add-ons with real residual sources of uncertainty left after hedging [5, 6].

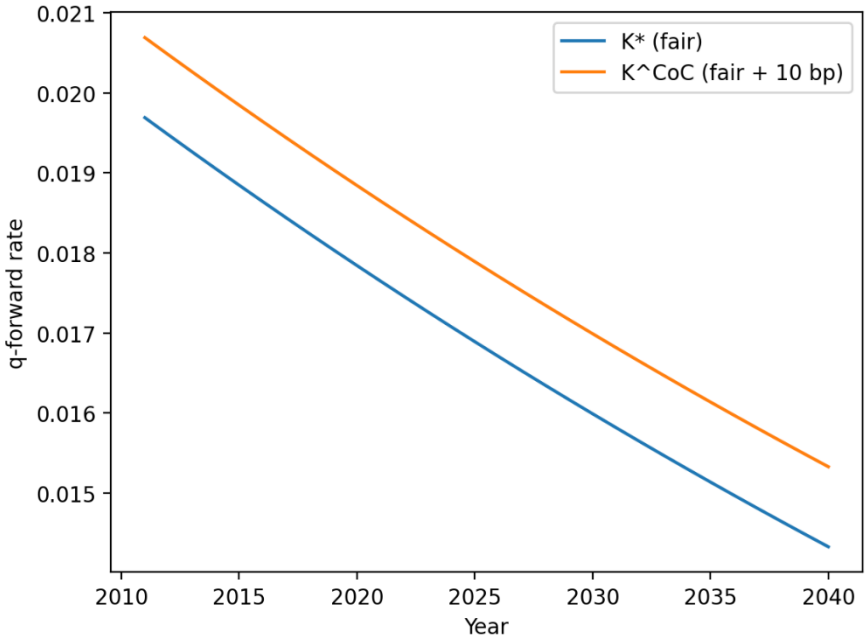

Fig. 6 two-line chart showing fair versus CoC-loaded curves with the premium wedge in basis points across maturities for Dutch ages 65–74 [7, 8].

3.8. Communication and Governance

Risk committees typically focus on two questions, namely the magnitude of variance reduction and the money at risk under plausible uncertainty bands, which our HE metrics and 95% ranges directly answer in a transparent fashion [9]. Presenting the notional profile alongside the discount-curve diagnostics improves auditability because it ties the hedge design back to economically meaningful term-structure and survival-ladder features rather than opaque optimization artefacts, which aids approvals and ongoing monitoring [8]. Juxtaposing the fair curve and the CoC-loaded curve supplies a robust price-reasonableness check and a negotiation anchor with dealers, which is particularly useful when markets are thin and quotes require careful benchmarking against regulatory cost-of-capital logic [7].

4. Discussion

Basis risk arises when a scheme's demographic mix differs from the index population underlying traded q-forwards, which can be mitigated by choosing age buckets that better match the scheme's exposure profile, by employing multi-population CBD models, and by considering bespoke indices where feasible [9]. Parameter and model risk can be assessed by sampling CBD posteriors and by cross-checking with Lee–Carter to ensure that results are not artefacts of a single specification, which widens credible intervals for K_T and produces bands for N_t and HE that can be communicated through tornado charts in governance settings [4, 6]. Liquidity and execution considerations favour sparse strips with robust collateral frameworks and netting sets to reduce operational overhead, while still preserving most variance reduction due to the smoothness and front-loading of gradient profiles

in typical pension liabilities, which balances risk and cost effectively [1]. Discount-curve dependence can be evaluated by re-running the pipeline with alternative curves such as internal market curves used for ALM consistency, which mainly rescales notionals and slightly tilts HE through reweighting effects but does not alter the qualitative dominance of first-order alignment [8]. For benefit designs with nonlinear features like caps or conditional indexation triggers, residual risk increases because the payoff ceases to be linear in q and higher-order terms become material, which may necessitate dynamic overlays or options on q-forwards in addition to static strips to maintain variance control [3]. Data provenance and reproducibility are strengthened by using public inputs and by version-controlling CSVs and code cells, which allows independent replication of figures including monotone DF, K_T trend, notional profiles, ΔPV distributions, and fair versus CoC-loaded curves, thereby meeting audit standards [2, 8].

5. Conclusion

Strips of q-forwards provide a clean, transparent, and highly effective linear hedge for longevity risk in Dutch pensions by cancelling first-order sensitivities to expected death probabilities across maturities, which our empirical results show via HE near the theoretical linear bound across multiple σ levels. The approach is straightforward to operationalize end-to-end using public mortality and discount inputs, and it yields outputs that are easy to audit and to communicate to stakeholders in governance processes, which supports adoption in practice. Mapping Solvency-II cost-of-capital logic to q-forward premia provides a coherent benchmark for price reasonableness and dealer negotiations, which integrates regulatory and market perspectives into a single framework. Future extensions should incorporate explicit SCR paths for basis and trend risk, sparse-strip optimization under liquidity constraints, and multi-population overlays to further align hedging outcomes with scheme-specific exposures while maintaining transparency and robustness.

References

[1] Blake, D., Cairns, A. J. G., & Dowd, K. (2013). Longevity risk and capital markets. Journal of Risk and Insurance, 80(3), 559–575.

[2] Zeddouk, F., Porcher, T., & Knafo, S. (2019). Pricing and hedging longevity risk with q-forwards. Insurance: Mathematics and Economics, 85, 218–236.

[3] Cairns, A. J. G., Blake, D., & Dowd, K. (2008). Modelling and management of mortality risk: A review. North American Actuarial Journal, 12(1), 23–48.

[4] Lee, R. D., & Carter, L. R. (1992). Modeling and forecasting U.S. mortality. Journal of the American Statistical Association, 87(419), 659–671.

[5] Cairns, A. J. G., Blake, D., & Dowd, K. (2006). A two-factor model for mortality. British Actuarial Journal, 12(3), 575–618.

[6] Renshaw, A. E., & Haberman, S. (2006). A cohort-based extension to the Lee–Carter model for mortality reduction factors. Insurance: Mathematics and Economics, 38(3), 556–570.

[7] European Commission. (2015). Commission Delegated Regulation (EU) 2015/35 supplementing Directive 2009/138/EC (Solvency II).

[8] EIOPA. (2023). Technical documentation of the methodology to derive EIOPA's risk-free interest rate term structure.

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

[10] Currie, I. D. (2006). Smoothing and forecasting mortality rates with P-splines. Statistical Modelling, 6(4), 279–298.