DISCUSSION PAPER PI-2603

Protection Growth, Balance-Sheet Capacity and the Economics of Risk Transfer.

Valeria D'Amato

September 2026

ISSN 1367-580X

The Pensions Institute
Bayes Business School
City, University of London
106 Bunhill Row
London EC1Y 8TZ
UNITED KINGDOM
http://www.pensions-institute.org/

Developing the Mortality Risk Market

Protection Growth, Balance-Sheet Capacity and the Economics of Risk Transfer

A Cedant-Side Framework for European Life Insurers

Valeria D'Amato

Sapienza University of Rome

Professor of Financial and Actuarial Mathematics

Discussion paper - September 2026

The quantitative calibration in this paper is illustrative and uses no proprietary insurer data. The framework is designed for subsequent portfolio-specific calibration using internal exposure, reinsurance and capital information.

Abstract

Mortality risk transfer is often approached from the perspective of available instruments, reinsurance structures or capital-market capacity. This paper adopts a cedant-side perspective and identifies the conditions under which growth in the primary protection market generates a mortality exposure large and systematic enough to create an economically meaningful demand for additional risk-bearing capacity. The framework separates diversifiable mortality variation from common mortality shocks, introduces Systematic Mortality Intensity (SMI), and measures the insurer's Net Demographic Exposure (NDE) after allowing for the interaction between mortality-sensitive and longevity-sensitive liabilities. These quantities feed into a Mortality Risk Capacity Frontier (MRCF), which characterises the combinations of protection scale, systematic dependence and balance-sheet capacity that can be supported before the preferred allocation of risk changes. The associated thresholds P^RI, P^K and P^dagger identify, respectively, the activation of conventional risk sharing, the capital-only scale limit and the point at which broader capacity becomes economically relevant. An illustrative calibration shows that reinsurance can more than double the sustainable protection scale relative to full retention, while a broader capacity set can increase it further. The analysis is extended to regime-dependent mortality dynamics, under which shifts in systematic mortality intensity move the capacity frontier even if nominal protection volume is unchanged. The paper also distinguishes technical transferability from economic investability: broader institutional participation requires sufficiently pure exposure, reliable data, transparent modelling, manageable basis risk and maturity alignment. The resulting development chain runs from protection demand and product design to mortality exposure, risk measurement, capacity allocation and, where economically justified, additional sources of risk-bearing capacity.

Keywords: mortality risk; life insurance; protection market; systematic mortality; natural hedging; reinsurance; risk-bearing capacity; insurance-linked risk; capital allocation.

Acknowledgement

The paper benefited from the discussion generated by the mortality-risk panel at the Twenty-First International Longevity Risk and Capital Markets Solutions Conference (L21), hosted by Sapienza University of Rome in September 2026. The arguments and results presented here are the author's own and do not represent the views of any insurer, reinsurer, investor or supervisory authority.

1. Introduction

The economic relevance of mortality risk for a life insurer is inseparable from the market in which that risk is originated. Mortality exposure is generated through the sale of protection products and is shaped by customer demand, distribution capacity, underwriting, product design and the demographic composition of insured lives. The allocation of mortality risk across retention, natural diversification, reinsurance and other sources of risk-bearing capacity therefore becomes economically material only as the primary protection market generates an exposure of sufficient scale and systematic content.

This starting point changes the natural ordering of mortality-risk analysis. The literature has established sophisticated models for stochastic mortality and has shown that demographic risks can in principle be hedged, reinsured or transferred through capital-market structures (Lee and Carter, 1992; Blake and Burrows, 2001; Cowley and Cummins, 2005; Lin and Cox, 2005; Cairns et al., 2006b). Technical feasibility, however, does not by itself establish a cedant-side business case. A transfer mechanism creates economic value only when the underlying exposure is material, the balance-sheet cost of retention is sufficiently high, and the all-in cost of alternative capacity is competitive.

The distinction is especially important for mortality risk because portfolio growth has two opposing effects. The idiosyncratic component of individual mortality is progressively diversified as the number of insured lives increases. Common mortality shocks, by contrast, survive pooling and can create correlated claims across large portfolios. As protection business grows, the proportion of total variance attributable to the common component therefore rises, creating a diversification floor that cannot be removed simply by adding policyholders.

A second balance-sheet mechanism arises from the coexistence of mortality-sensitive and longevity-sensitive liabilities. Higher mortality increases death-benefit claims but can reduce the value of annuity-type liabilities. This creates natural demographic diversification (Cox and Lin, 2007). The offset is rarely exact, because protection and annuity books may refer to populations with different age, socioeconomic, geographic and underwriting characteristics (Li and Haberman, 2015). The economically relevant object is consequently the net demographic exposure of the insurer rather than gross population mortality alone.

The same demographic shock can therefore generate very different economic losses across insurers. Aggregate excess mortality does not reveal which cohorts are affected, the sums at risk attached to them, or the degree of concentration in a given protection portfolio. Cohort mismatch also matters for hedging: a population index may provide a transparent reference while only imperfectly matching the cedant's actual liability profile. Data granularity and basis risk are thus part of the economics of mortality-risk transfer, not secondary implementation details.

The paper develops a cedant-side framework around three connected objects. Systematic Mortality Intensity (SMI) measures the proportion of aggregate mortality variance that is systematic. Net Demographic Exposure (NDE) measures the direction and scale of the insurer's first-order demographic exposure after allowing for mortality-longevity offsets. The Mortality Risk Capacity Frontier (MRCF) maps the combinations of exposure and balance-sheet capacity at which the preferred risk allocation changes.

The model produces three economically interpretable thresholds. P^K is the maximum protection scale supportable under a capital-only configuration. P^RI marks the activation of conventional risk sharing when its marginal cost becomes preferable to additional internal retention. P^dagger identifies the point at which a broader source of capacity enters the optimal allocation. These thresholds are endogenous to protection growth, systematic mortality intensity, natural hedging and transfer costs; they are not imposed by the existence of a particular instrument.

The framework is deliberately neutral about the final transfer technology. Reinsurance is the natural first external layer because it combines risk-bearing capacity with underwriting expertise, portfolio information and customisation. Broader institutional capacity may become relevant as exposures become larger, more measurable and more transparent. The resulting market-development sequence is therefore protection demand, origination, exposure measurement, capacity allocation, standardisation and transparency, and economic investability.

The paper also allows mortality dynamics to vary across latent regimes. A shift from a normal to a high-systematic-mortality regime raises the common component of risk, moves the capacity frontier inward and can change the preferred allocation even if the nominal volume of protection business is unchanged. This extension links stochastic mortality modelling directly to a balance-sheet decision rather than treating regime changes as a forecasting exercise alone.

The contribution is both theoretical and operational. The model provides a compact language for discussing the business case for mortality-risk transfer from the perspective of a protection writer. The numerical illustration shows how the capacity frontier and protection-scale multipliers can be computed. The empirical implementation protocol is designed for public mortality data and can subsequently be calibrated to proprietary insurer data, including product-level exposures, reinsurance treaties, capital costs and candidate index structures.

2. Related literature and contribution

The paper lies at the intersection of stochastic mortality modelling, natural demographic hedging, insurance risk allocation and biometric-risk transfer. These strands are well developed individually, but they are less often connected through the endogenous growth of the primary protection portfolio.

Stochastic mortality modelling provides the statistical foundation. Lee and Carter (1992) introduced a parsimonious common-factor representation of mortality evolution. The literature subsequently developed Poisson, multi-factor, cohort, multi-population and alternative stochastic specifications (Brouhns et al., 2002; Cairns et al., 2006a; Renshaw and Haberman, 2006; Cairns et al., 2009; Plat, 2009; D'Amato et al., 2011; Hunt and Blake, 2015; Villegas et al., 2018). Booth and Tickle (2008) provide a broader review of mortality modelling and forecasting, while Basellini et al. (2023) document the continuing centrality of Lee-Carter-type structures and the importance of model choice, age patterns and forecast uncertainty.

For the present problem, the key statistical distinction is not between one mortality forecasting model and another, but between diversifiable and systematic components of aggregate claims. A common mortality factor creates covariance across insured lives. As portfolio size grows, the idiosyncratic term scales linearly with exposure while the systematic term scales quadratically. This simple asymptotic property is the foundation of SMI and the diversification floor developed below.

Natural hedging forms the second strand. Cox and Lin (2007) show how life insurance and annuity exposures can offset because mortality changes affect their liabilities in opposite directions. Li and Haberman (2015) demonstrate that the degree of risk reduction depends materially on the populations involved. This paper incorporates the same economic intuition through NDE and uses it as an input into capacity allocation rather than as a stand-alone product-mix problem.

A third literature studies mortality and longevity securitisation. Survivor bonds, mortality-linked securities, mortality-market pricing frameworks and related instruments demonstrate that demographic risks are technically transferable (Blake and Burrows, 2001; Cowley and Cummins, 2005; Lin and Cox, 2005; Cairns et al., 2006b; Blake et al., 2008; Dawson et al., 2010). More recent work extends this logic to pandemic-related mortality bonds (Chen et al., 2025). Index-based hedges also make basis risk central to contract design (Cairns, 2013). The present paper is complementary but intentionally upstream: it establishes the conditions under which a cedant has an economic reason to seek additional capacity before specifying the instrument used to provide it.

Regime-dependent mortality provides the fourth strand. Mortality dynamics may change across latent states, particularly when shocks are temporary, persistent or structurally different from normal improvement patterns. Milidonis et al. (2011) show the relevance of mortality regimes for pricing. The extension in Section 8 uses the same idea to show how regime shifts alter SMI, capital requirements and the MRCF.

The paper's contribution is therefore to link protection-market scale, systematic mortality dependence, natural hedging and capacity allocation in one cedant-side framework. This yields an economic distinction between technical transferability and investability, identifies endogenous activation thresholds, and provides a transparent route from public-data analysis to insurer-specific calibration.

3. Protection scale and systematic mortality risk

3.1 Aggregate protection claims

Consider a one-period protection portfolio with N insured lives. Let A_i denote the amount at risk on life i and D_i the death indicator over the risk horizon. Aggregate claims are

$$C_N = \sum_{i=1}^{N} A_i D_i \tag{1}$$

Mortality dependence is represented through a common factor Z. Conditional on Z, individual deaths are assumed independent, with conditional mortality probability

$$q_i(Z) = Pr(D_i = 1 \mid Z) \tag{2}$$

The law of total variance separates aggregate claims into an idiosyncratic component and a systematic component:

$$Var(C_N) = V_N^I + V_N^S \tag{3}$$

$$V_N^I = \sum_i A_i^2 E\left[q_i(Z)(1 - q_i(Z))\right], V_N^S = Var\left[\sum_i A_i q_i(Z)\right] \tag{4}$$

The first term reflects conditional individual variation and is diversifiable. The second is generated by the common mortality factor and persists at portfolio level. This decomposition is deliberately model-agnostic: Z may be a Lee-Carter factor, a common shock factor, a regime indicator, or a vector of demographic drivers.

3.2 Homogeneous benchmark and diversification floor

To obtain transparent comparative statics, let total protection scale be P = NA and adopt the reduced-form variance representation

$$V^I(P) = a\,P, V^S(P) = b\,P^2, a > 0, b \geq 0 \tag{5}$$

The total claim variance and relative volatility are therefore

$$Var\left[C(P)\right] = a\,P + b\,P^2 \tag{6}$$

$$RV(P) = \frac{\sqrt{Var\left[C(P)\right]}}{P} = \sqrt{\frac{a}{P} + b} \tag{7}$$

As protection scale grows, idiosyncratic volatility per unit of exposure disappears but systematic volatility remains. The diversification floor is

$$DF = \lim_{P \to \infty} RV(P) = \sqrt{b} \tag{8}$$

The parameter b is therefore economically important even when small. It governs the irreducible volatility of a large protection book and is the channel through which common mortality shocks translate portfolio growth into capacity pressure.

3.3 Systematic Mortality Intensity

Systematic Mortality Intensity is defined as the share of aggregate mortality variance attributable to the common component:

$$SMI(P) = \frac{V^S(P)}{V^I(P) + V^S(P)} = \frac{bP}{a + bP} \tag{9}$$

SMI is bounded between zero and one and provides a scale-dependent measure of how much of aggregate protection risk survives pooling.

Proposition 1. Scale effect on systematic mortality.

For a>0 and b>0, SMI(P) is strictly increasing in P and converges to one as P tends to infinity. Hence portfolio growth reduces relative idiosyncratic volatility while increasing the fraction of aggregate variance that is systematic.

Proof. Differentiation gives dSMI/dP = ab/(a+bP)² > 0, and the limiting result follows directly from (9).

3.4 Capital requirement under full retention

For a transparent benchmark, unexpected mortality capital is approximated by a quantile multiple of aggregate standard deviation:

$$K_\alpha(P) = z_\alpha \sqrt{aP + bP^2} \tag{10}$$

where z_α is the chosen confidence multiplier. If the internal capacity allocated to the mortality book is $\bar{K}$, the maximum protection scale under full retention is the positive solution of K_α(P)=$\bar{K}$:

$$P^K = \frac{-a + \sqrt{a^2 + 4b\left(\acute{K}/z_\alpha\right)^2}}{2b}, b>0 \tag{11}$$

When b tends to zero, (11) converges to the familiar linear-diversification limit P^K=($\bar{K}$/z_α)²/a. For b>0, systematic dependence reduces the protection scale that can be supported by a fixed capital allocation.

4. Natural hedging and Net Demographic Exposure

4.1 Mortality and longevity sensitivities

Let V_t^M denote the value of mortality-sensitive liabilities and V_t^L the value of longevity-sensitive liabilities. Let Z_t be a common mortality state variable normalised so that an increase corresponds to adverse mortality. The first-order sensitivities satisfy

$$\frac{\partial V_t^M}{\partial Z_t} > 0, \frac{\partial V_t^L}{\partial Z_t} < 0 \tag{12}$$

The net demographic sensitivity is

$$\Delta_t^D = \frac{\partial V_t^M}{\partial Z_t} + \frac{\partial V_t^L}{\partial Z_t} \tag{13}$$

Normalising by own funds OF_t gives the Net Demographic Exposure

$$NDE_t = \frac{\Delta_t^D}{OF_t} \tag{14}$$

NDE_t>0 denotes a balance sheet dominated, at first order, by adverse-mortality exposure. NDE_t<0 denotes a longevity-dominant position. NDE_t near zero indicates a first-order offset, not the absence of demographic risk. Nonlinearities, cohort mismatch, basis risk and multiple mortality factors can leave material residual exposure even when the scalar first-order measure is close to zero.

4.2 Protection growth and the direction of NDE

Let B(P) denote the mortality sensitivity generated by the protection portfolio, with B'(P)>0 and B''(P)≤0. Holding the longevity-sensitive book and own funds fixed over the local comparison,

$$NDE(P) = \frac{B(P) + \Delta^L}{OF}, \Delta^L < 0 \tag{15}$$

Proposition 2. Protection growth shifts the demographic balance.

Under B'(P)>0, fixed Δ^L and OF>0, NDE(P) is increasing in protection scale. Natural hedging can delay the point at which mortality risk becomes capacity-relevant, but it cannot be assumed to eliminate the systematic component generated by a growing protection book.

Proof. Differentiation gives dNDE/dP=B'(P)/OF>0.

4.3 Cohort mismatch and multi-factor residual risk

A scalar NDE is most informative when a common mortality factor captures the main balance-sheet movement. In a multi-factor setting, let s_M and s_L denote vectors of liability sensitivities to demographic factors and define s_net=s_M+s_L. With factor covariance matrix Σ_Z, residual demographic risk is summarised by

$$R_{dem} = \sqrt{s_{net}' \Sigma_Z s_{net}} \tag{16}$$

This formulation makes cohort mismatch explicit. Protection and annuity books can have offsetting exposure to a broad mortality trend but materially different sensitivity to age-specific, socioeconomic, geographic or cause-specific shocks. Aggregate population mortality is therefore not a sufficient statistic for insurer loss.

5. The Mortality Risk Capacity Frontier

5.1 Capacity allocation problem

Let E denote the systematic mortality exposure that remains after allowing for diversification and natural demographic offsets. The cedant allocates E across internal retention r, reinsurance q and a broader external capacity channel x:

$$r + q + x = E, r, q, x \geq 0, q \leq \acute{Q}, x \leq \acute{X} \tag{17}$$

The economic allocation minimises the all-in cost of bearing the exposure:

$$V(E, C) = min_{r,q,x} \left\{ C_I(r; C) + C_R(q) + C_X(x) \right\} \tag{18}$$

$$\textit{subject to (17).} \tag{19}$$

C is the internal balance-sheet capacity available to the mortality portfolio. C_I includes the marginal economic cost of internal capital and concentration. C_R includes reinsurance premium, counterparty and transaction costs. C_X is the all-in cost of broader capacity, including structuring, basis-risk and liquidity or maturity costs where relevant. The formulation does not assume that x is an insurance-linked security; it represents any additional balance sheet outside the conventional cedant-reinsurance allocation.

5.2 Marginal cost conditions

For an interior allocation, the Karush-Kuhn-Tucker conditions imply equality of active marginal costs:

$$MC_I(r; C) = MC_R(q) = MC_X(x) = \lambda \tag{20}$$

where λ=∂V/∂E is the shadow value of one additional unit of mortality-risk capacity. If a capacity constraint binds, the associated multiplier measures the economic value of relaxing that constraint.

5.3 Definition of the MRCF

The Mortality Risk Capacity Frontier is the exposure level at which internal retention ceases to be the sole optimal allocation:

$$E^{\star}(C)=inf\left\{E:q^{\star}(E,C)+x^{\star}(E,C)>0\right\} \tag{21}$$

If the first available transfer layer has entry marginal cost

$$MC_T=min\left\{C_R^{'}(0),C_X^{'}(0)\right\} \tag{22}$$

the frontier condition is

$$MC_I\left(E^{\star}(C);C\right)=MC_T \tag{23}$$

Proposition 3. Capacity shifts the frontier outward.

Suppose MC_I is increasing in retained exposure and decreasing in internal capacity C. Then dE*/dC>0. Greater internal capacity allows a larger systematic mortality exposure to be retained before transfer becomes economically efficient.

Proof. Apply the implicit-function theorem to MC_I(E*,C)−MC_T=0. Since ∂MC_I/∂E>0 and ∂MC_I/∂C<0, dE*/dC=−(∂MC_I/∂C)/(∂MC_I/∂E)>0.

5.4 Reinsurance-first allocation and broader capacity

Proposition 4. Reinsurance-first condition.

If C_R′(0)<C_X′(0), reinsurance is the first external layer to enter the optimal allocation. Broader capacity enters only after reinsurance marginal cost rises to the relevant competing margin or reinsurance capacity becomes binding.

This ordering captures an important institutional feature of mortality risk. Reinsurance combines capital with underwriting knowledge and portfolio-specific structuring, while an external investor generally requires a more standardised and transparent representation of the risk. The model therefore permits broader capacity to enter endogenously without assuming that it replaces reinsurance.

5.5 Protection-scale thresholds

Mapping systematic exposure E back to protection scale P yields the three thresholds used in the numerical analysis:

- P^RI: the smallest protection scale for which conventional external risk sharing is active in the optimum.
- P^K: the capital-only scale limit defined by (11).
- P^dagger: the smallest protection scale for which the broader capacity channel x is active in the optimum.

An auxiliary threshold P^H denotes the point at which the reinsurance layer reaches its calibrated capacity $\bar{H}$. The threshold ordering is economically informative: it indicates whether a growing protection book first encounters internal-capital scarcity, reinsurance scarcity, or a relative-price condition that makes broader capacity attractive.

5.6 Endogenous protection growth and the joint cedant decision

The fixed-exposure allocation in Sections 5.1-5.5 determines where a given systematic exposure should sit. To make the scale of the protection business itself endogenous, let G(P) denote the economic contribution generated by protection growth before mortality-risk-bearing costs. For the benchmark, a parsimonious concave specification is

$$G(P)=\beta P-\frac{\eta}{2}P^2,\beta>0,\eta>0 \tag{24}$$

Let H and X denote, respectively, the protection-scale-equivalent amounts transferred to conventional reinsurance and to a broader source of capacity. The benchmark retained unexpected-capital proxy is

$$EC(P,H,X)=z_{\alpha}\left(1-\frac{H+X}{P}\right)\sqrt{aP+bP^2} \tag{25}$$

Equation (25) is deliberately stylised: it assumes that transferred capacity reduces the benchmark mortality exposure proportionally. A portfolio-specific implementation would replace this mapping with the actual capital response to treaty terms, attachment points, portfolio composition and basis risk.

The all-in transfer costs are represented by convex functions

$$C_{RI}(H)=c_1 H+\frac{c_2}{2}H^2, C_X(X)=d_1 X+\frac{d_2}{2}X^2 \quad (26)$$

The joint cedant decision is then to choose protection scale and capacity allocation so as to maximise the economic value of the protection strategy:

$$J(P,H,X)=G(P)-C_{RI}(H)-C_X(X)-\gamma\, EC(P,H,X) \quad (27)$$

$$EC(P,H,X)\le \acute{K}, 0\le H\le \acute{H}, X\ge 0, H+X\le P \quad (28)$$

The coefficient γ converts retained unexpected capital into an economic cost. The capital constraint and the convex transfer costs make the demand for external capacity endogenous: transfer is not valuable because an instrument exists, but because the marginal value of protection growth eventually exceeds what can be absorbed efficiently through internal retention alone. The generic allocation problem of Section 5 remains the conceptual MRCF; equations (24)-(28) provide the concrete benchmark used in Section 7.

6. From transferability to economic investability

The capacity model establishes when the cedant values additional risk-bearing capacity. It does not imply that such capacity will be supplied. Investor-side feasibility requires a second condition: the mortality exposure must be represented in a form that can be identified, modelled and priced outside the originating insurance balance sheet.

Risk purity is central to this transition. An investor seeking demographic diversification needs a clear definition of the reference population, event, trigger, loss function and risks retained by the cedant. Additional exposures embedded unintentionally in the structure reduce the ability to identify the intended mortality risk and can weaken its portfolio-diversification value.

Basis risk creates a structural trade-off. A broad population index improves standardisation and comparability but can diverge from the actual insured portfolio. A customised indemnity structure can track the cedant's loss more closely while sacrificing standardisation, potentially reducing comparability and market liquidity. This is analogous to the index-basis-risk problem in longevity hedging (Cairns, 2013), but extreme mortality events can magnify cohort and age-profile mismatch.

Data quality therefore becomes part of risk quality. Reliable mortality histories, exposure information, cohort definitions and modelling assumptions reduce both parameter uncertainty and ambiguity about the asset being transferred. Institutional infrastructure, contractual clarity and model governance are consequently economic inputs to investability rather than merely operational requirements.

Maturity alignment is equally important. The period for which a cedant needs protection must be compatible with the duration for which a capital provider is willing to commit funds. A structure with attractive expected economics but an unsuitable investment horizon may fail to mobilise capacity. The investability condition is therefore joint: risk purity, transparency, basis risk, expected compensation and duration must be mutually consistent.

The distinction between technical transferability and economic investability is important for market development. Mortality risk can be technically transferable before either side has an economic reason to transact. A sustainable market emerges only when cedant-side capacity value and investor-side investability overlap.

7. Numerical illustration and capacity multipliers

7.1 Benchmark calibration

The numerical benchmark is designed to illustrate the mechanics of the framework rather than reproduce a specific insurer's portfolio. Protection scale is normalised, internal capacity is set to one unit, and the variance coefficients generate a modest idiosyncratic term together with a non-zero systematic floor. The reduced-form growth and transfer-cost coefficients implement the joint decision problem in equations (24)-(28).

Parameter	Symbol	Benchmark
Idiosyncratic variance coefficient	a	0.020

Systematic variance coefficient	b	0.0015
Capital confidence multiplier	z_α	3.000
Internal mortality capacity	$\bar{K}$	1.000
Protection value coefficient	β	0.180
Protection curvature coefficient	η	0.012
Reinsurance linear cost coefficient	c_1	0.015
Reinsurance convex cost coefficient	c_2	0.010
Reinsurance capacity	$\bar{H}$	4.000
Broader-capacity linear cost coefficient	d_1	0.045
Broader-capacity convex cost coefficient	d_2	0.005
Capital-cost coefficient	γ	0.020

Table 1. Illustrative benchmark parameters.

7.2 Capital-only benchmark

With a=0.020, b=0.0015, z_α=3 and $\bar{K}$=1, equation (11) gives

$$P^K = 4.220 \tag{29}$$

This value is the benchmark scale sustainable under full retention. At P=P^K, the systematic share of variance is already economically meaningful even though b is small. The result illustrates why the absolute size of b cannot be judged independently of portfolio scale.

Protection scale P	SMI(P)	Relative volatility
4.220	0.240	0.0790
9.562	0.418	0.0599
11.208	0.457	0.0573

Table 2. Scale effect on systematic mortality intensity in the benchmark.

7.3 Reinsurance and broader capacity

The calibrated optimisation is evaluated under three nested capacity sets: full retention; retention plus reinsurance; and retention plus reinsurance plus broader capacity. The results are

Capacity configuration	Optimal P*	Reinsurance H*	Broader capacity X*	Objective J*
Full retention	4.220	0.000	0.000	0.633
Reinsurance available	9.562	4.000	0.000	1.013
Broader capacity available	11.208	3.797	1.594	1.037

Table 3. Illustrative optimisation results. All quantities are normalised benchmark units.

For protection scales below P^K, the benchmark transfer prices make full retention preferable and the internal capital constraint is slack. At P^K=4.220 the constraint first binds; hence, in this calibration, the conventional risk-sharing activation threshold coincides with the capital-only threshold: P^RI=P^K=4.220. This equality is calibration-specific rather than structural.

The broader-capacity activation point can be derived from marginal transfer costs. With X=0, broader capacity first competes with reinsurance when C_RI'(H)=C_X'(0), giving

$$H^{\dagger} = \frac{d_1 - c_1}{c_2} = 3.000 \tag{30}$$

Combining this marginal-cost condition with the binding capital constraint gives

$$EC\left(P^{\dagger},3,0\right)=\acute{K}\Rightarrow P^{\dagger}=8.339 \tag{31}$$

If only reinsurance is available, its calibrated limit H̄=4 is reached when

$$EC\left(P^{H},\acute{H},0\right)=\acute{K}\Rightarrow P^{H}=9.562 \tag{32}$$

The ordering P^RI=P^K < P^dagger < P^H is economically informative. Broader capacity becomes price-relevant before conventional reinsurance capacity is exhausted. At the joint optimum P*=11.208, the allocation H*=3.797 and X*=1.594 equalises the active marginal transfer costs while the capital constraint remains binding. The broader channel therefore substitutes partially for reinsurance rather than simply adding capacity on top of a fixed reinsurance quantity.

7.4 Protection Capacity Multipliers

To express the economic effect in a scale-free form, define the Protection Capacity Multiplier associated with reinsurance as

$$PCM_{RI}=\frac{P^{\star}_{RI}}{P^{\star}_{retention}} \tag{33}$$

and the Extended Protection Capacity Multiplier as

$$EPCM=\frac{P^{\star}_{RI+X}}{P^{\star}_{retention}} \tag{34}$$

The benchmark values are

$$PCM_{RI}=\frac{9.562}{4.220}\approx 2.27\,,EPCM=\frac{11.208}{4.220}\approx 2.66 \tag{35}$$

The interpretation is direct. Relative to the capital-only benchmark, reinsurance more than doubles the economically supportable protection scale; access to a broader capacity set raises the scale further. These multipliers are not universal estimates. They are comparative-statical outputs of the benchmark and will change with the mortality dependence structure, reinsurance terms, capital costs, basis risk and the pricing of broader capacity.

7.5 Sensitivity logic

The model produces four robust directional results. A higher systematic coefficient b raises SMI and moves the capital frontier inward. Greater internal capacity K̄ shifts the frontier outward. Lower reinsurance marginal cost or higher reinsurance capacity increases PCM_RI. Lower all-in cost of broader capacity, including basis-risk and structuring costs, reduces P^dagger and increases EPCM. These comparative statics provide the bridge from a stylised calibration to an insurer-specific business case.

7.6 Public-data implementation protocol

A public-data implementation can be built using the Human Mortality Database (HMD) for a set of mature European insurance markets. The pre-specified design uses Italy, France, Spain, the Netherlands, Austria, Belgium, and England & Wales; ages 40-94; a 1990-2019 estimation window; and 2020-2022 as an explicit held-out stress period. The purpose is not to proxy a proprietary insured portfolio perfectly, but to estimate the common and idiosyncratic components of mortality variation and to stress the capacity model under empirically grounded mortality dynamics.

The implementation proceeds in four stages:

1. Estimate age-specific mortality dynamics and extract one or more common mortality factors for each population, using a transparent benchmark model and robustness checks across alternative stochastic mortality specifications.
2. Map the estimated factor structure into portfolio-weighted death-claim exposure for synthetic protection portfolios with alternative age and sum-at-risk profiles.
3. Estimate the implied systematic component and compute SMI, the diversification floor and regime-specific capital requirements across protection scales.
4. Feed those quantities into the MRCF and compare the resulting activation thresholds across markets, portfolio compositions and mortality regimes.

The public-data exercise is intentionally separated from firm-specific calibration. Insured populations differ from national populations through underwriting, socioeconomic selection, product mix and geography. The HMD analysis therefore supports model validation and comparative stress testing; it does not substitute for internal claims and exposure data.

8. Regime-dependent mortality and capacity

8.1 Latent mortality states

Let S_t take values in a finite set of mortality regimes {1,...,m} and evolve as a Markov chain with transition matrix Π=(p_ij), where p_ij=Pr(S_{t+1}=j|S_t=i). The reduced-form variance coefficients become state dependent:

$$Var\left[C(P) \mid S_t = s\right] = a_s P + b_s P^2 \tag{36}$$

The state-dependent systematic intensity and capital requirement are

$$SMI(P \mid s) = \frac{b_s P}{a_s + b_s P} \tag{37}$$

$$K_\alpha(P \mid s) = z_\alpha \sqrt{a_s P + b_s P^2} \tag{38}$$

A high-systematic-mortality state is represented by a larger b_s and, where appropriate, a different idiosyncratic coefficient a_s. This structure can represent periods in which deaths become more correlated across lives, even when the shock is temporary.

8.2 State-dependent frontier

The capacity problem becomes state dependent:

$$V(E, C, s) = min\left[C_I(r; C, s) + C_R(q, s) + C_X(x, s)\right] \tag{39}$$

subject to the same exposure-allocation constraint. The MRCF is now E*(C,s). A regime change can therefore alter the preferred allocation without any change in nominal protection volume.

Proposition 5. Systematic regimes move the capital threshold inward.

For fixed a_s, z_α and $\bar{K}$, an increase in b_s lowers P^K_s. Hence a regime characterised by stronger common mortality dependence reduces the scale supportable by a fixed internal capital allocation.

> The result follows directly from K_α(P|s)=$\bar{K}$: for each positive P, K_α is increasing in b_s, so the inverse scale threshold is decreasing in b_s.

8.3 Regime uncertainty and forward-looking capacity

At time t, the relevant capacity decision may depend on the transition probabilities rather than the current state alone. A simple forward-looking measure is the conditional expected capital requirement

$$E_t\left[K_\alpha(P, S_{t+1})\right] = \sum_s p_{S_t, s} K_\alpha(P \mid s) \tag{40}$$

Stress testing can place additional weight on adverse regimes or evaluate the frontier under each state separately. The main economic implication is unchanged: systematic mortality is not only a level parameter. It can be a state variable, and a protection strategy that appears capacity-efficient in a normal regime may become transfer-intensive after a structural change in mortality dependence.

9. Discussion and market implications

The framework produces a scale-dependent interpretation of mortality-risk transfer. At low protection volumes, idiosyncratic variation is largely poolable and the systematic component can remain small relative to existing balance-sheet capacity. In this region, sophisticated transfer structures can be technically feasible without creating material cedant-side value. As protection portfolios expand, systematic exposure becomes progressively more important and the marginal economics of capacity change.

This mechanism places the primary protection market at the beginning of the development chain. Protection scale is not an arbitrary balance-sheet parameter: it is the outcome of customer demand, distribution capability, underwriting and product design. Protection demand can also be state dependent. Systemic health events can increase the salience of mortality and health protection, so the same broad risk environment that raises claims in the short term can influence the subsequent scale of protection business. The current model treats P as a scale state variable rather than modelling household demand explicitly, but the economic interpretation is endogenous to the primary market.

Distribution matters because protection products are more difficult to communicate than simple investment propositions. The customer must understand the protection need, the insured event, the economic value of coverage and contractual limitations. A growing protection market therefore requires product expertise, appropriate training and sustained support for distribution networks. Protection can also be embedded within savings or investment-oriented propositions, so mortality exposure can expand through product architecture as well as through stand-alone pure-risk contracts.

The second implication concerns natural hedging. A diversified life group may hold mortality-sensitive and longevity-sensitive liabilities simultaneously, but gross offsets should not be confused with effective hedges. Cohort composition determines the transmission from demographic shock to economic loss. This is why NDE is best interpreted together with multi-factor residual exposure and portfolio-specific stress testing.

The third implication concerns the ordering of capacity. Reinsurance remains the natural first external layer because it combines balance-sheet capacity with underwriting knowledge and customisation. Broader sources of institutional capacity become potentially valuable when protection exposure is sufficiently material, systematic and transparent. The numerical benchmark illustrates this complementarity: the joint optimum uses less reinsurance than the reinsurance-only capacity limit while supporting a larger protection scale.

The fourth implication is the distinction between transferability and investability. Capital availability alone does not generate investor demand. Mortality-linked exposure must be sufficiently pure, transparent and standardisable to be evaluated within an institutional portfolio. Reference population, trigger, basis risk, modelling assumptions and investment horizon are therefore joint determinants of market depth.

The fifth implication concerns data. Data quality is part of risk quality because it determines both how accurately the cedant can estimate the exposure and how clearly an external risk bearer can understand what is being assumed. Developed markets with reliable mortality statistics, strong legal infrastructure and transparent modelling standards have an advantage in transforming insurance exposure into an investable risk representation.

The framework does not imply that mortality risk must evolve into a capital-market asset class. It establishes a more limited and testable proposition: protection growth can endogenously generate a demand for additional risk-bearing capacity, and broader capacity becomes economically relevant only where the cedant-side value of transfer overlaps with an investable representation of the risk.

9.1 Implementation pathway for a large diversified life insurer

A portfolio-specific application requires internal data at a level that allows the public mortality factor to be mapped into actual economic exposure. The minimum calibration set includes:

- exposure counts and sums at risk by age, sex, geography, product family and underwriting class;
- death claims and expected mortality bases, including experience studies and catastrophe or pandemic adjustments;
- mortality-sensitive and longevity-sensitive liability sensitivities required to estimate NDE and multi-factor residual demographic exposure;
- current reinsurance structures, attachment points, limits, pricing, counterparty constraints and retrocession assumptions where relevant;
- allocated own funds, internal capital charges and the marginal economic cost of increasing retained mortality exposure;
- candidate reference indices or trigger definitions and an explicit estimate of population and portfolio basis risk;
- product-growth scenarios covering stand-alone protection, credit or mortgage protection, embedded death benefits and relevant riders;
- distribution and underwriting scenarios that change the age and socioeconomic composition of new business; and
- liquidity, duration, legal and transaction-cost assumptions for any broader capacity structure.

These inputs permit the abstract thresholds to be translated into management quantities: the protection scale consistent with a given capital budget, the marginal value of additional reinsurance, the shadow value of reinsurance capacity, the conditions under which broader capacity adds value, and the sensitivity of each result to adverse mortality regimes.

9.2 Insurer-specific pilot calibration

A portfolio-specific implementation should begin with a deliberately limited pilot. The purpose is not to reproduce the entire mortality-risk architecture of a diversified life insurer, nor to design a transfer instrument, but to determine whether the mortality-capacity mechanism identified in this paper is economically material in an actual protection portfolio.

The initial perimeter should comprise one protection portfolio, or a small set of sufficiently homogeneous portfolios, for which exposures, sums at risk, claims experience and reinsurance arrangements can be reconstructed consistently. Interpretability is more important than breadth at this stage. A narrowly defined portfolio provides a more informative test of systematic mortality exposure than a broad aggregation in which demographic composition, underwriting and product effects cannot be separated.

The first step is to map insured exposure by age, sex, geography, underwriting class, product family and sums at risk. This is required because population mortality and insured loss are not equivalent. The same population mortality shock can generate materially different economic losses depending on the demographic and financial composition of the insured book.

The second step is to link internal claims experience to external mortality factors. Public mortality data can be used to identify common demographic movements, while internal exposure and claims data determine how strongly the insurer's portfolio loads on those movements. The objective is not to assume that insured mortality replicates national mortality, but to estimate the transmission from common mortality shocks to portfolio losses.

This provides the empirical basis for estimating the idiosyncratic and systematic components of claim variance and the associated Systematic Mortality Intensity. The purpose of this first calibration is diagnostic: it establishes whether the systematic component remains economically small at current scale or becomes increasingly relevant under plausible protection-growth scenarios.

The analysis should then incorporate the insurer's demographic balance sheet. Mortality-sensitive protection liabilities and longevity-sensitive liabilities should be evaluated jointly in order to estimate Net Demographic Exposure and the residual demographic position after natural hedging. Gross mortality and longevity exposures should not be assumed to offset automatically, because the effectiveness of the hedge depends on the correspondence between the underlying populations.

The existing reinsurance programme and internal capital position complete the traditional capacity assessment. Reinsurance should ultimately be represented through actual treaty limits, attachment points, pricing and counterparty conditions, while internal capital should reflect the insurer's own marginal economic cost of retained mortality exposure. For an initial pilot, however, a reduced representation of principal retentions, limits and marginal pricing can provide a sufficiently informative first test.

The output of the exercise should be a set of management quantities rather than an instrument valuation. These include portfolio-specific SMI, the diversification floor, NDE, the capital-only threshold P^K, the point at which additional reinsurance becomes economically valuable, the marginal value of further reinsurance capacity and the position of projected protection volumes relative to the Mortality Risk Capacity Frontier.

The calibration should be repeated under alternative protection-growth scenarios and adverse mortality regimes. This allows management to determine whether the portfolio remains comfortably within existing capacity, approaches the frontier as protection expands, or becomes materially more capacity intensive when the common component of mortality increases.

A first-stage exercise therefore requires only a restricted dataset: aggregated exposures and sums at risk, historical and expected death claims, principal reinsurance terms, an internal mortality-capital measure, a first approximation of mortality- and longevity-sensitive liabilities where feasible, and commercially relevant protection-growth scenarios. This reduced dataset cannot support the design or pricing of a transaction. It is sufficient, however, to answer the more fundamental question of whether additional mortality risk-bearing capacity is likely to become economically relevant at the scale to which the insurer intends to grow.

9.3 Decision rules and evidence of a business case

The existence of systematic mortality risk does not by itself establish a business case for additional risk transfer. Nor does the technical possibility of constructing a mortality-linked exposure. A cedant-side business case exists only when changing the allocation of mortality risk creates measurable economic value for the insurer.

At low protection volumes, idiosyncratic mortality is strongly diversified and the common component may remain small relative to available capital. If existing reinsurance is also abundant and competitively priced, the efficient response may simply be to retain the exposure within the existing insurance and reinsurance structure. In this region there is mortality risk, but there is no material business case for expanding the capacity set.

An emerging business case arises when protection growth changes the marginal economics of mortality exposure. SMI may rise, residual NDE may remain material after natural hedging, capital utilisation may move closer to the insurer's internal limit and the marginal value of reinsurance may increase. At this stage, the relevant issue is capacity optimisation rather than market creation. The insurer should first test whether the constraint can be addressed more efficiently through retention, conventional reinsurance, treaty redesign, underwriting, portfolio composition or product architecture.

A stronger business case arises only when three conditions coincide. Residual systematic mortality exposure must be material after diversification and natural hedging. The existing combination of internal capital and conventional reinsurance must become economically constraining, whether through marginal cost, concentration, capacity limits, counterparty constraints or other frictions. Finally, an alternative source of capacity must be able to absorb the residual exposure at a competitive all-in cost.

The relevant comparison is not simply the quoted price of external capacity. The all-in economic cost must include basis risk, transaction costs, legal and operational costs and any structural inefficiencies created by an imperfect transfer. Even a competitive price is not sufficient if the residual exposure cannot be represented in a sufficiently identifiable, measurable and transparent form. Reference population, trigger definition, basis risk, modelling assumptions and maturity therefore remain part of the business-case assessment rather than secondary implementation details.

The empirical test is consequently symmetric. The framework must be capable of rejecting the case for broader capacity. If residual mortality exposure is immaterial, internal capital remains abundant, reinsurance remains efficient or basis risk makes external transfer uneconomic, the appropriate conclusion is that no broader-capacity business case currently exists. Conversely, further market development becomes economically justified only when material cedant demand for additional capacity coincides with an exposure that can be represented and priced by another risk-bearing balance sheet.

This distinction is particularly important because a strong business case for additional reinsurance does not necessarily imply a business case for broader institutional capital. The framework is therefore intended to identify where the insurer sits along the capacity-development path, rather than to presuppose its destination.

9.4 Strategic implications for protection growth

The economic role of mortality-risk capacity extends beyond the management of an existing book. If protection demand is expected to grow, capacity can become an input into the insurer's commercial strategy.

In a static interpretation, reinsurance is primarily a mechanism for reducing the risk of an existing portfolio. In a dynamic protection strategy, additional capacity may also create value by allowing the insurer to originate profitable new protection business without increasing retained mortality exposure, concentration or capital utilisation beyond economically acceptable levels. The relevant comparison is therefore between the protection scale that can be supported under alternative capacity configurations.

This is the economic interpretation of the Protection Capacity Multiplier introduced in Section 7. The multiplier measures the extent to which a change in capacity supports additional protection business rather than merely transferring part of the existing exposure. The same reasoning underlies the Extended Protection Capacity Multiplier when a broader capacity source is economically available.

This interpretation distinguishes between defensive and growth-enabling risk transfer. Defensive transfer reduces volatility or capital intensity on an existing book. Growth-enabling transfer creates value because it allows economically attractive protection business to expand beyond the scale that could otherwise be supported efficiently.

The decision is nevertheless broader than a choice between retention and transfer. Management can also modify product mix, underwriting, geographic composition, distribution strategy and reinsurance architecture. The economically preferred strategy is the combination that maximises the value of protection growth after accounting for the mortality exposure created by that growth and the cost of absorbing it.

This matters because identical premium growth can generate very different mortality-capacity profiles. Stand-alone term-life protection, credit and mortgage protection, employee benefits, riders and death benefits embedded in savings products can differ substantially in age composition, sums at risk, underwriting intensity and geographic concentration. Protection strategy and capacity strategy should therefore not be determined independently.

The framework developed in this paper implies that a broader mortality-risk market should not be considered separately from the underlying protection market. Its economic foundation lies in the demand for protection and in the exposures generated by satisfying that demand. As protection business expands, insurers originate larger mortality exposures. Pooling progressively reduces policy-level variation, while common mortality shocks remain. Reinsurance can absorb part of the resulting net exposure. If protection growth then makes concentration, accumulation or capital costs sufficiently material, the economic value of additional capacity may increase.

Broader risk sharing is therefore not an independent objective. It is a possible consequence of the economics of protection growth. The strongest evidence for the development of a broader mortality-risk market would not be the existence of a new financial instrument, but evidence that protection growth generates a persistent and economically material demand for additional risk-bearing capacity that cannot be supplied as efficiently through existing retention and reinsurance arrangements.

This preserves the cedant-side logic of the paper: the market follows the exposure; the exposure does not follow the instrument.

10. Conclusions

This paper develops a cedant-side framework for analysing the relationship between protection-market growth, systematic mortality exposure and the allocation of risk-bearing capacity. Its central contribution is not the design of a mortality-linked instrument, but a method for determining whether the growth of protection business creates an endogenous demand for additional capacity and for identifying the conditions under which the efficient allocation of mortality risk should change.

The starting point is the distinction between diversifiable and systematic mortality risk. As the number of insured lives increases, idiosyncratic mortality variation is progressively pooled. Common mortality shocks, however, remain and can generate correlated claims across large portfolios. The Systematic Mortality Intensity introduced in this paper provides a scale-dependent measure of the proportion of mortality variance that survives pooling and therefore establishes a direct link between protection growth and residual mortality exposure.

Population mortality alone is not sufficient to determine the economic impact on an insurer. Age composition, sums at risk, geography, underwriting selection, product mix and concentration determine how a demographic shock is transmitted to insured losses. For this reason, the framework complements SMI with Net Demographic Exposure, which captures the first-order interaction between mortality-sensitive and longevity-sensitive liabilities. Natural diversification can be economically valuable, but its effectiveness depends on the degree of correspondence between the underlying insured populations and should not be inferred from gross balance-sheet positions alone.

These exposures feed into the Mortality Risk Capacity Frontier. The MRCF identifies the combinations of protection scale, systematic mortality exposure and balance-sheet capacity under which the preferred allocation among internal retention, conventional reinsurance and broader sources of capacity changes. The associated thresholds provide an economically interpretable way to identify when internal capacity becomes constraining, when reinsurance begins to create marginal value and when a broader capacity set may become competitive.

The numerical calibration illustrates the mechanism rather than providing a market estimate. Under the benchmark assumptions, conventional reinsurance increases the optimal protection scale materially relative to full retention, while the availability of a broader source of capacity allows a further increase. The relevant result is not the numerical value of the multipliers themselves, but the fact that the value of risk transfer is endogenous to the scale and composition of the protection portfolio. Additional capacity can therefore have a growth-enabling role as well as a defensive risk-management role.

This distinction has important strategic implications. Reinsurance should not be evaluated only by the amount of existing risk that it removes from the balance sheet. It can also be assessed by the additional protection business that it allows the insurer to originate while remaining within an economically acceptable capital and concentration envelope. The same logic can be applied to broader capacity, but only after all-in costs, basis risk, structural complexity and investor-side feasibility have been incorporated.

The paper consequently separates three questions that are often conflated. The first is whether mortality exposure can be measured with sufficient reliability. The second is whether the cedant has an economically material need for additional risk-bearing capacity. The third is whether such capacity can be supplied at a competitive all-in cost by a balance sheet outside the existing retention and reinsurance structure. A sustainable business case requires these conditions to align.

This provides an important discipline for market development. The existence of systematic mortality risk does not imply that a broader mortality-risk market is necessary, and the technical ability to construct a transferable exposure does not establish that such a structure creates value for the cedant. The framework can therefore produce a negative result: if residual mortality exposure remains modest, existing capital is sufficient and reinsurance remains abundant and competitively priced, there may be no current economic case for expanding the risk-bearing set.

Conversely, a stronger business case emerges when protection growth produces a material residual exposure, when internal or conventional reinsurance capacity becomes increasingly costly or constrained, and when the residual risk can be represented in a sufficiently transparent, measurable and robust form. Under these conditions, broader institutional capacity may complement rather than replace conventional reinsurance.

The regime-dependent extension reinforces this conclusion. The capacity frontier is not determined solely by the nominal size of the protection book. Changes in the common component of mortality can move the frontier even when business volume remains unchanged. A portfolio that appears comfortably absorbable in a normal mortality regime may therefore become substantially more capacity intensive following a structural change in mortality dependence.

The empirical implementation proposed in this paper is deliberately sequential. A limited first-stage calibration can establish whether the mortality-capacity mechanism is material using a restricted set of exposure, claims, capital and reinsurance data. A more detailed second stage can then incorporate portfolio granularity, actual treaty structures, internal economic-capital costs and, only where justified, candidate forms of broader capacity.

This approach makes the framework directly testable using insurer-specific data. It allows management to determine whether the binding constraint on protection growth lies in customer demand, product and distribution design, internal capital, reinsurance capacity, the representation of residual mortality exposure or the availability of external risk-bearing capacity.

The central economic question is therefore not whether mortality risk can be transferred. It is whether protection growth creates an exposure sufficiently large, systematic and economically material that transferring part of that risk becomes more valuable than retaining it within the existing insurance and reinsurance system.

If that condition is met, the development of additional mortality-risk capacity should be understood not as an objective in itself, but as a consequence of the underlying economics of the protection market.

Appendix A. Additional analytical results

A.1 Diversification floor from a common-factor representation

A complementary representation starts from per-policy loss Y_i=μ+β_Z S+η_i, where S is a common mortality shock with variance σ_S² and η_i are independent idiosyncratic terms with variance σ_η². For the average loss Ȳ_N=N^{-1} Σ_iY_i,

$$Var\left(\acute{Y}_N\right)=\beta_Z^2\sigma_S^2+\frac{\sigma_\eta^2}{N} \tag{A1}$$

Hence lim_{N→∞}Var(Ȳ_N)=β_Z²σ_S². Equation (A1) is the micro-foundation of the reduced-form diversification floor in (8).

The corresponding finite-N systematic share is

$$SMI_N=\frac{\beta_Z^2\sigma_S^2}{\beta_Z^2\sigma_S^2+\sigma_\eta^2/N} \tag{A2}$$

A.2 Shadow price of reinsurance capacity

Let μ_R≥0 be the KKT multiplier on q≤Q̄. When the reinsurance capacity constraint binds, μ_R measures the reduction in minimum total risk-bearing cost produced by a marginal increase in Q̄:

$$\mu_R=-\frac{\partial V}{\partial \acute{Q}}>0 \tag{A3}$$

This shadow price provides a direct economic diagnostic. A high μ_R indicates that the cedant is not primarily constrained by the technical transferability of mortality risk but by the amount or marginal price of conventional capacity available at the relevant exposure level.

A.3 Transaction-level interpretation

The frontier formulation can also be expressed as a transaction-level transfer decision. Let L^N be the net demographic loss before external transfer and L^res the residual loss after reinsurance and any broader hedge. An all-in transfer cost C(h) includes premium, structuring, counterparty, basis-risk and transaction costs. The economic value of transfer can be written as

$$B(h)=\eta_K\left[K\left(L^N\right)-K\left(L^{res}\right)\right)+\lambda_\rho\left[\rho\left(L^N\right)-\rho\left(L^{res}\right)\right)-C(h) \tag{A4}$$

The optimal hedge h* maximises B(h). This representation is equivalent in economic content to the marginal-cost allocation of Section 5: transfer is valuable when the reduction in capital and retained risk exceeds the all-in cost of obtaining the capacity.

Appendix B. Public-data estimation design

The HMD-based design is intended as a reproducible proof of concept. For each population, annual deaths and exposures are converted into age-specific central death rates for ages 40-94. The benchmark common-factor model is estimated on 1990-2019, while 2020-2022 is retained as a held-out stress period rather than absorbed into the fitting window. A benchmark common-factor mortality model is estimated on log mortality rates, with alternative Poisson and multi-factor specifications used as robustness checks. The common component is then propagated through synthetic protection portfolios using age-specific sums at risk.

The core outputs are not mortality forecasts for their own sake. They are inputs into capacity analysis: estimated a and b coefficients; SMI(P); the diversification floor; portfolio-specific claim variance; regime-conditioned capital requirements; and the resulting MRCF thresholds. Robustness analysis should vary the fitting window, age range, portfolio weights and treatment of exceptional mortality years, while preserving a genuine out-of-sample stress check wherever possible.

Because national-population and insured-population mortality differ, the public-data stage is best interpreted as a market-level calibration of systematic dependence. Firm-specific implementation then replaces synthetic weights with actual insured exposures and introduces underwriting selection, policy terms, reinsurance and capital costs.

Appendix C. Insurer-specific calibration roadmap

Stage	Internal / external inputs	Primary model output
1. Exposure map	Policies, sums at risk, age/sex/geography, underwriting class, product family	Portfolio-weighted mortality exposure
2. Demographic balance	Protection and annuity liability sensitivities; own funds	NDE and residual multi-factor exposure
3. Systematic risk	Historical claims plus external mortality factors	SMI, diversification floor, stress regimes
4. Reinsurance	Treaties, limits, attachment, pricing, counterparty terms	C_R(q), $\bar{Q}$, reinsurance shadow price
5. Internal capital	Allocated capacity and marginal cost of capital	C_I(r;C), P^K, MRCF
6. Broader capacity	Candidate triggers, index data, structure, duration, investor pricing	C_X(x), basis risk, P^dagger
7. Growth scenarios	New-business plans and distribution/product mix	PCM, EPCM under strategic scenarios

Table C1. Practical calibration sequence.

Appendix D. Main notation

Symbol	Meaning
P	Protection portfolio scale
C_N	Aggregate protection claims
Z	Common mortality factor
a	Idiosyncratic variance coefficient
b	Systematic variance coefficient
SMI	Systematic Mortality Intensity
DF	Diversification floor
K_α(P)	Mortality capital requirement at confidence level α
$\bar{K}$	Internal mortality-risk capacity
P^K	Capital-only protection-scale threshold
V^M,V^L	Mortality- and longevity-sensitive liability values
NDE	Net Demographic Exposure
OF	Own funds
E	Systematic exposure to be allocated
r	Internal retention
q / H	Reinsurance allocation
x / X	Broader external-capacity allocation
$\bar{Q}$ / $\bar{H}$	Reinsurance capacity limit
$\bar{X}$	Broader-capacity limit
MRCF	Mortality Risk Capacity Frontier
P^RI	Reinsurance / conventional risk-sharing activation threshold
P^dagger	Broader-capacity activation threshold
PCM_RI	Protection Capacity Multiplier with reinsurance
EPCM	Extended Protection Capacity Multiplier
S_t	Latent mortality regime
Π	Regime transition matrix

References

Basellini, U., Camarda, C.G. and Booth, H. (2023). Thirty years on: A review of the Lee-Carter method for forecasting mortality. International Journal of Forecasting, 39(3), 1033-1049. https://doi.org/10.1016/j.ijforecast.2022.11.002

Blake, D. and Burrows, W. (2001). Survivor bonds: Helping to hedge mortality risk. Journal of Risk and Insurance, 68(2), 339-348. https://doi.org/10.2307/2678106

Blake, D., Cairns, A.J.G. and Dowd, K. (2008). The birth of the life market. Asia-Pacific Journal of Risk and Insurance, 3(1), 1-32. https://doi.org/10.2202/2153-3792.1027

Booth, H. and Tickle, L. (2008). Mortality modelling and forecasting: A review of methods. Annals of Actuarial Science, 3(1-2), 3-43. https://doi.org/10.1017/S1748499500000440

Brouhns, N., Denuit, M. and Vermunt, J.K. (2002). A Poisson log-bilinear regression approach to the construction of projected lifetables. Insurance: Mathematics and Economics, 31(3), 373-393. https://doi.org/10.1016/S0167-6687(02)00185-3

Cairns, A.J.G. (2013). Robust hedging of longevity risk. Journal of Risk and Insurance, 80(3), 621-648. https://doi.org/10.1111/j.1539-6975.2013.01525.x

Cairns, A.J.G., Blake, D. and Dowd, K. (2006a). A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. Journal of Risk and Insurance, 73(4), 687-718. https://doi.org/10.1111/j.1539-6975.2006.00195.x

Cairns, A.J.G., Blake, D. and Dowd, K. (2006b). Pricing death: Frameworks for the valuation and securitization of mortality risk. ASTIN Bulletin, 36(1), 79-120. https://doi.org/10.2143/AST.36.1.2014145

Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., Epstein, D., Ong, A. and Balevich, I. (2009). A quantitative comparison of stochastic mortality models using data from England and Wales and the United States. North American Actuarial Journal, 13(1), 1-35. https://doi.org/10.1080/10920277.2009.10597538

Chen, Z., Li, H., Mao, Y. and Zhou, K.Q. (2025). Learning from COVID-19: A catastrophe mortality bond solution in the post-pandemic era. Insurance: Mathematics and Economics, 123, 103113. https://doi.org/10.1016/j.insmatheco.2025.103113

Cowley, A. and Cummins, J.D. (2005). Securitization of life insurance assets and liabilities. Journal of Risk and Insurance, 72(2), 193-226. https://doi.org/10.1111/j.1539-6975.2005.00121.x

Cox, S.H. and Lin, Y. (2007). Natural hedging of life and annuity mortality risks. North American Actuarial Journal, 11(3), 1-15. https://doi.org/10.1080/10920277.2007.10597464

D'Amato, V., di Lorenzo, E., Haberman, S., Russolillo, M. and Sibillo, M. (2011). The Poisson log-bilinear Lee-Carter model: Applications of efficient bootstrap methods to annuity analyses. North American Actuarial Journal, 15(2), 315-333. https://doi.org/10.1080/10920277.2011.10597623

Dawson, P., Dowd, K., Cairns, A.J.G. and Blake, D. (2010). Survivor derivatives: A consistent pricing framework. Journal of Risk and Insurance, 77(3), 579-596. https://doi.org/10.1111/j.1539-6975.2010.01356.x

HMD. Human Mortality Database. Max Planck Institute for Demographic Research (Germany), University of California, Berkeley (USA), and French Institute for Demographic Studies (France). Available at www.mortality.org.

Hunt, A. and Blake, D. (2015). On the structure and classification of mortality models. Pensions Institute Discussion Paper PI-1506, Cass Business School, City University London.

Lee, R.D. and Carter, L.R. (1992). Modeling and forecasting U.S. mortality. Journal of the American Statistical Association, 87(419), 659-671. https://doi.org/10.1080/01621459.1992.10475265

Li, J.S.-H. and Haberman, S. (2015). On the effectiveness of natural hedging for insurance companies and pension plans. Insurance: Mathematics and Economics, 61, 286-297. https://doi.org/10.1016/j.insmatheco.2015.01.009

Lin, Y. and Cox, S.H. (2005). Securitization of mortality risks in life annuities. Journal of Risk and Insurance, 72(2), 227-252. https://doi.org/10.1111/j.1539-6975.2005.00122.x

Milidonis, A., Lin, Y. and Cox, S.H. (2011). Mortality regimes and pricing. North American Actuarial Journal, 15(2), 266-289. https://doi.org/10.1080/10920277.2011.10597621

Plat, R. (2009). On stochastic mortality modeling. Insurance: Mathematics and Economics, 45(3), 393-404. https://doi.org/10.1016/j.insmatheco.2009.08.006

Renshaw, A.E. and Haberman, S. (2006). A cohort-based extension to the Lee-Carter model for mortality reduction factors. Insurance: Mathematics and Economics, 38(3), 556-570. https://doi.org/10.1016/j.insmatheco.2005.12.001

Villegas, A.M., Kaishev, V.K. and Millossovich, P. (2018). StMoMo: An R package for stochastic mortality modeling. Journal of Statistical Software, 84(3), 1-38. https://doi.org/10.18637/jss.v084.i03