Annals of Actuarial Science (2026), 1–21
doi:10.1017/S174849952610044X

ORIGINAL RESEARCH PAPER

A natural hedging framework for longevity risk with graphical risk assessment

Lydia J. Gabric[1] and Kenneth Q. Zhou[2]

[1]School of Mathematical and Statistical Sciences, Arizona State University, USA; and [2]Department of Statistics and Actuarial Science, University of Waterloo, Canada
Corresponding author: Kenneth Q. Zhou; Email: kenneth.zhou@uwaterloo.ca

(Received 21 January 2026 UTC; revised 18 August 2026 UTC; accepted 26 August 2026 UTC)

Abstract

Natural hedging allows life insurers to manage longevity risk internally by offsetting the opposite exposures of life insurance and annuity liabilities. Although many studies have proposed natural hedging strategies under different settings, calibration methods, and mortality models, a unified framework for constructing and evaluating such hedges remains undeveloped. While graphical risk assessment has been explored for index-based longevity hedges, no comparable metric exists for natural hedging. This paper proposes a structured natural hedging framework paired with a graphical risk metric for hedge evaluation. The framework integrates valuation, calibration, and evaluation, while the graphical metric provides intuitive insights into residual dependencies and hedge performance. Applied to multiple hedging scenarios, the proposed methods demonstrate flexibility, interpretability, and practical value for longevity risk management.

Keywords: graphical risk metric; hedge effectiveness; longevity risk; natural hedging; stochastic mortality models

1. Introduction

The uncertainty surrounding future mortality trends, known as longevity risk, has been extensively studied in the context of life insurance valuation and risk management. In general, there are two categories of longevity risk management solutions: internal and external. External solutions transfer longevity risk from life insurers to capital market investors through index-based hedging instruments and standardized contracts. Internal solutions, in contrast, utilize the insurer's own portfolio structure to mitigate longevity and mortality risks. Among these, natural hedging has received considerable attention for its intuitive mechanism of balancing opposite exposures within life insurance and annuity portfolios.

The seminal work of Cox and Lin (2007) introduced the concept of natural hedging by leveraging the inverse relationship between life insurance and annuity products to diversify mortality and longevity risks. When mortality changes, the values of life insurance and annuity liabilities typically move in opposite directions, allowing the insurer to offset these changes and reduce overall risk exposure. Since Cox and Lin (2007), numerous studies have explored how this relationship can be more effectively utilized and implemented in different settings to manage mortality and longevity risks.

For instance, Lin and Tsai (2014) proposed an immunization-based approach to construct natural hedges using sensitivity measures. Zhu and Bauer (2014) explored a non-parametric mortality model and found that natural hedging performance may be less effective compared with parametric models. More recently, Yang *et al.* (2019) combined internal natural hedges with external

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Institute and Faculty of Actuaries. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

index-based hedges under a unified framework. Chen *et al.* (2024) analyzed product demand and supply considerations in natural hedging, while Cupido *et al.* (2024) studied the impact of spatial dependence on hedge performance. These studies collectively demonstrate the growing importance of natural hedging in longevity risk management.

Despite the aforementioned advancements, several challenges remain in constructing natural hedging effectively. First, identifying appropriate life insurance products to offset the mortality risk of annuity products can be complex. Second, determining the optimal quantity of life insurance products to effectively hedge requires detailed mathematical calibration. Third, evaluating hedge effectiveness is complicated due to residual risks such as calibration risk, structural basis risk, and model risk. Although prior research has attempted to address aspects of these challenges, there is still no unified framework that allows the systematic construction, calibration, and comparison of natural hedges across different models and methods.

In addition, existing studies primarily rely on numerical risk measures to evaluate hedging performance, which may not fully capture the joint behavior of liabilities and hedging instruments. This shortfall could understate the strength of the inverse relationship or the offsetting effects in natural hedging. To address these gaps, this paper develops a unified natural hedging framework that integrates both numerical and graphical evaluation methods that can serve as a practical "toolkit" for constructing and analyzing naturally hedged portfolios. The framework provides a consistent process for valuation, calibration, and evaluation, while the accompanying graphical risk metric offers visual insights into the resulting hedge performance. The contributions of this paper are threefold.

The first contribution of this paper is the development of a risk management framework specifically designed for natural hedging. While frameworks for constructing and evaluating external longevity hedges have been proposed in prior studies, such as the index-based hedging procedure of Coughlan *et al.* (2011) and the decomposition of hedge effectiveness in Cairns *et al.* (2014), a comparable framework dedicated to natural hedging has not been established. To fill this gap, we propose a three-step hedging framework that extends the structure of Cairns *et al.* (2014) and adapts it for natural hedging.

The proposed framework consists of three stages: valuation, calibration, and evaluation. In the first stage, we derive formulas for valuing life insurance and annuity products under a stochastic mortality setting. In the second stage, we consider multiple foundational hedge calibration techniques categorized by their objectives, including variance minimization, duration matching, and delta hedging. In the third stage, hedge evaluation combines risk measures with graphical tools to provide a comprehensive and interpretable assessment of hedging effectiveness. The integration of graphical evaluation within the framework leads to our second contribution.

The second contribution of this paper is the introduction of a graphical risk metric specifically tailored to natural hedging. Graphical tools have been explored in the context of standardized longevity hedging, such as the graphical basis risk metrics proposed by Chan *et al.* (2016) and extended by Sherris *et al.* (2020), as well as by Blake *et al.* (2008), Dowd *et al.* (2010), and Li and Liu (2021). Because no comparable metric has been developed for evaluating natural hedges, we propose a graphical risk metric that visualizes the relationship between life annuity and insurance portfolios, allowing for a direct and intuitive evaluation of hedge performance.

The proposed graphical risk metric is constructed using a series of joint prediction regions that represent potential hedge outcomes across multiple confidence levels. To facilitate practical application, we further develop an interpretation procedure that enables users to visually assess hedge effectiveness and compare alternative hedging strategies. Visual elements such as shaded regions and reference lines highlight deviations from expected outcomes and the degree of diversification effects. This graphical approach provides a transparent and intuitive means to diagnose whether hedging shortfalls arise from under-hedging, over-hedging, or a lack of hedging effect.

The third contribution of this paper is the numerical implementation of the proposed natural hedging framework and graphical risk metric across three practical applications using simulated

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

mortality data. First, we demonstrate how the framework can identify trade-offs when selecting an insurance portfolio as the hedging instrument. Second, we compare alternative calibration techniques to determine the optimal allocation of the hedging portfolio. Third, we evaluate model risk by comparing two stochastic mortality models and a non-parametric approach for a given portfolio configuration. The results from these illustrations show that the proposed framework and graphical risk metric together serve as a flexible and interpretable tool for constructing and evaluating natural hedges.

The remainder of the paper is organized as follows. Section 2 outlines the components of the natural hedging framework and demonstrates its use through a toy example. Section 3 describes the construction and interpretation of the graphical risk metric. Section 4 applies the framework and graphical metric to three practical illustrations. Finally, Section 5 concludes the paper and discusses possible extensions and limitations.

2. The natural hedging framework

Consider a life insurer issuing both life annuity and life insurance products. Let $\mathcal{A}$ and $\mathcal{I}$ denote the present value random variables of the cash flows generated by the annuity and insurance portfolios, respectively. To hedge the longevity risk of $\mathcal{A}$, the insurer employs a natural hedging strategy resulting in a hedged position

$$\mathcal{P}(h) = \mathcal{A} + \mathcal{L}(h),$$

where $\mathcal{L}(h) = h \cdot \mathcal{I}$ represents the calibrated life insurance portfolio, and h is the hedge ratio indicating the amount of $\mathcal{I}$ needed for the hedged position.

To support a systematic implementation of such natural hedging, we develop a three-step framework in this section. We briefly summarize the three steps as follows:

1. *Valuation of $\mathcal{A}$ and $\mathcal{I}$*: Derive the present value of the annuity and insurance liabilities and obtain their expected values under stochastic mortality assumptions.
2. *Calibration of h*: Determine the hedge ratio that achieves a specified objective, such as variance minimization, duration matching, or delta-neutral.
3. *Evaluation of $\mathcal{P}(h)$*: Assess the performance of the hedged portfolio through both numerical and graphical metrics, such as variance, value-at-risk (VaR), or histograms.

The rest of this section examines each step of the proposed framework in detail and concludes with a simple illustrative example.

2.1 Valuation

To construct a natural hedge, the insurer must first quantify the present value of the annuity and insurance liabilities. Under stochastic mortality modeling, these liabilities are random variables as their values depend on uncertain future survival outcomes. Thus, the first step of our natural hedging framework is to define and value these present value random variables as the foundation for establishing the hedge.

Let $\mathcal{S}_x(T)$ be the probability that an individual aged x at time 0 will survive another T years. This probability can be expressed as:

$$\mathcal{S}_x(T) = \prod_{s=1}^{T} (1 - q_{x+s-1,s}), \tag{1}$$

where $q_{x,s}$ is the probability that an individual aged x at the beginning of year s dies during year s. In a stochastic setting, $\mathcal{S}_x(T)$ is a random variable because $q_{x,s}$, for $s = 1, \ldots, T$, is uncertain due

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

to random mortality fluctuations. We denote its expectation by $S_x(T) := \mathbb{E}[\mathcal{S}_x(T)]$. The procedure for generating realizations of $\mathcal{S}_x(T)$ is provided in Appendix B of the Supplementary Material.

2.1.1 Annuity liabilities

Recall that we denote $\mathcal{A}$ as the present value random variable of the cash flows generated by the annuity portfolio under stochastic mortality. Assume that the insurer's annuity portfolio consists of n_A distinct product types (e.g., annuities issued at different ages). If a constant force of interest δ is used, then $\mathcal{A}$ is given by

$$\mathcal{A} = \sum_{j=1}^{n_A} \omega_j^A \sum_{k=\tau_j}^{\tau_j+t_j-1} c_j \cdot e^{-\delta k} \cdot \mathcal{S}_{x_j}(k),$$

where ω_j^A is the weight of the j-th annuity product, τ_j is its deferral period, t_j is the number of payments, c_j is the payment size, and x_j is the issuing age. The weights satisfy $\sum_{j=1}^{n_A} \omega_j^A = 1$ and $0 < \omega_j^A \leq 1$ for all j. We remark that ω_j^A and c_j serve distinct roles and allow us to distinguish between portfolio allocation decisions (captured by ω_j^A) and contractual product specifications (captured by c_j). A similar argument applies to insurance liabilities in the next subsection.

It follows that the expectation of $\mathcal{A}$ represents the expected present value of the entire annuity portfolio liabilities, which can be written as:

$$A := \mathbb{E}[\mathcal{A}] = \sum_{j=1}^{n_A} \omega_j^A \sum_{k=\tau_j}^{\tau_j+t_j-1} c_j \cdot e^{-\delta k} \cdot S_{x_j}(k). \tag{2}$$

The distribution of $\mathcal{A}$ and the value of A can be obtained using Monte Carlo simulations under stochastic mortality modeling.

2.1.2 Insurance liabilities

Recall again that we denote $\mathcal{I}$ as the present value random variable of the cash flows generated by the insurance portfolio under stochastic mortality. If the insurer has n_I distinct insurance product types in the portfolio and applies a constant force of interest δ, then $\mathcal{I}$ is given by

$$\mathcal{I} = \sum_{j=1}^{n_I} \omega_j^I \sum_{k=0}^{t_j-1} b_j \cdot e^{-\delta(k+1)} \cdot (\mathcal{S}_{x_j}(k) - \mathcal{S}_{x_j}(k+1)),$$

where ω_j^I is the weight of the j-th insurance product, t_j is its term length, b_j is the benefit amount, and x_j is the issuing age. The weights again satisfy $\sum_{j=1}^{n_I} \omega_j^I = 1$ and $0 < \omega_j^I \leq 1$ for all j.

Similar to the annuity liabilities, the expectation of $\mathcal{I}$ is the expected present value of the entire insurance portfolio liabilities, which can be written as:

$$I := \mathbb{E}[\mathcal{I}] = \sum_{j=1}^{n_I} \omega_j^I \sum_{k=0}^{t_j-1} b_j \cdot e^{-\delta(k+1)} \cdot (S_{x_j}(k) - S_{x_j}(k+1)). \tag{3}$$

The distribution of $\mathcal{I}$ and the value of I can also be obtained using Monte Carlo simulations under stochastic mortality modeling.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

Table 1. Summary of the notations used in the natural hedging framework

Description	Present value (PV)	Expected PV	Mean-adjusted PV
Annuity portfolio	$\mathcal{A}$	A	$\tilde{\mathcal{A}} = \mathcal{A} - A$
Insurance portfolio	$\mathcal{I}$	I	$\tilde{\mathcal{I}} = \mathcal{I} - I$
Calibrated portfolio	$\mathcal{L}(h)$	$L(h)$	$\tilde{\mathcal{L}}(h) = \mathcal{L}(h) - L(h)$
Combined portfolio	$\mathcal{P}(h)$	$P(h)$	$\tilde{\mathcal{P}}(h) = \mathcal{P}(h) - P(h)$

2.1.3 Summary of notations

Having derived the annuity and insurance portfolios, we now summarize the notations before presenting the next two steps of the framework. The hedge ratio is h, the calibrated insurance portfolio is $\mathcal{L}(h) = h \cdot \mathcal{I}$ with its expected present value being $L(h) = h \cdot I$, and the combined portfolio (i.e., the hedged position) is $\mathcal{P}(h) = \mathcal{A} + \mathcal{L}(h)$ with its expected present value being $P(h) = A + L(h)$.

For ease of reference, Table 1 summarizes the key notations used throughout the paper and also introduces the mean-adjusted present values that will play a central role in the rest of the paper. Note that, in Table 1, the calibrated portfolio refers to the insurance portfolio after applying the hedge ratio h, the combined portfolio refers to the annuity and calibrated insurance portfolios taken together, and the mean-adjusted present values in the last column represent deviations from the corresponding expected present values.

2.2 Calibration

We now turn to the calibration step of the framework. The hedge ratio h determines how the insurance liabilities offset annuity liabilities, and calibrating h is therefore a critical step that directly affects the effectiveness of the natural hedge. In this section, we categorize calibration methods into three groups: optimization, immunization, and Greek neutral. Before outlining a specific technique from each category to be implemented later in this paper, we briefly review the main approaches and related literature.

The first category is optimization, which calibrates h by minimizing a chosen risk measure. Cupido *et al.* (2024) implemented a risk-minimization approach to construct a naturally hedged portfolio that accounts for spatial dependencies between different populations. Yang *et al.* (2019) proposed a unified hedging framework that minimizes changes in the insurer's profit function for a portfolio combining internal natural hedges and external index-based hedges. Zhu and Bauer (2014) applied a financial hedging approach to minimize economic capital, as originally introduced by Zhu and Bauer (2011).

The next category is immunization, which seeks to minimize portfolio liability sensitivity to changes in mortality rates. Li and Luo (2012) proposed a key q-Duration framework that matches sensitivities in mortality rates to hedge longevity risk with q-forward contracts. Tsai *et al.* (2010) introduced a conditional VaR minimization approach that optimizes the insurer's product mix and compared it with the modified duration matching method of Wang *et al.* (2010). Chen *et al.* (2024) generalized the approach of Tsai *et al.* (2010) to incorporate both insurance supply and demand, while Gatzert and Wesker (2014) extended the method of Wang *et al.* (2010) to the entire insurance perspective, building on the framework of Gatzert and Wesker (2012).

Lastly, Greek neutral techniques aim to reduce portfolio liabilities with respect to changes in model-specific mortality quantities. Cairns (2013) introduced delta-nuga hedging as a robust approach to address recalibration risk in portfolios with index-based hedging instruments. Liu and Li (2017) expanded the delta-nuga method of Cairns (2013) through a generalized state-space hedging framework. To hedge changes in the reserves of a naturally hedged portfolio, Luciano *et al.* (2017) developed delta-gamma hedging based on first- and second-order approximations.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

Jevtić and Regis (2015) applied this approach to assess the solvency of a naturally hedged portfolio from an asset-liability perspective.

We remark that the above review is not intended to be exhaustive but rather to provide a context for the calibration step of the natural hedging framework. A simple benchmark is the uncalibrated hedge, where the hedge ratio is set to $h = 1$, and the insurer makes no adjustment to the insurance liabilities. Additionally, setting $h = 1$ may be appropriate when an insurer has already constructed a naturally hedged portfolio with pre-determined annuity and insurance liabilities. Lastly, under a discrete-time stochastic mortality model, the hedge ratio will generally need to be calculated using Monte Carlo simulations. In the remainder of this section, we focus on one representative technique from each category, which will later be used in our numerical implementation.

2.2.1 Optimization

The aim of optimization is to calibrate the hedge ratio h by minimizing a chosen risk measure of the hedged position. Formally, this optimization problem can be written as:

$$\min_{h\in\mathbb{R}} \rho(\mathcal{P}(h)),$$

where ρ denotes a selected risk measure, and $\mathcal{P}(h) = \mathcal{A} + h \cdot \mathcal{I}$ is the hedged portfolio.

When ρ is taken to be variance, the problem reduces to variance minimization, a well-established approach in longevity risk management. In this case, Cairns *et al.* (2014) showed that the optimal hedge ratio $h^{(VM)}$ has a closed-form solution given by

$$h^{(VM)} = -\frac{\mathrm{Cov}(\mathcal{A}, \mathcal{I})}{\mathrm{Var}(\mathcal{I})}. \tag{4}$$

The corresponding hedged portfolio is denoted by $\mathcal{P}(h^{(VM)}) = \mathcal{A} + h^{(VM)}\mathcal{I}$. This case will be used in the numerical illustrations later in the paper.

2.2.2 Immunization

Immunization was originally introduced by Redington (1952) for hedging interest rate risk. Wang *et al.* (2010) adapted it for natural hedging, where the goal is to neutralize the sensitivity of annuity and insurance liabilities to changes in mortality. The mortality duration of the combined portfolio $P(h)$ with respect to the force of mortality μ is defined as:

$$D^P_\mu = \frac{dP(h)}{d\mu} = \frac{dA}{d\mu} + \frac{dL(h)}{d\mu} = \frac{dA}{d\mu} + h\frac{dI}{d\mu}.$$

The optimal hedge ratio under this approach is obtained by setting $D^P_\mu = 0$.

To approximate the mortality durations in D^P_μ, we define the modified durations of the annuity and insurance portfolios as $D^A_\mu = \frac{A^+ - A^-}{2\epsilon}$ and $D^I_\mu = \frac{I^+ - I^-}{2\epsilon}$, where ϵ is a small constant, $A^\pm$ and $I^\pm$ are the expected present values of the annuity and insurance liabilities, respectively, evaluated under adjusted mortality $\mu \pm \epsilon$. These quantities capture how much the expected present values change in response to a small shift in mortality. The resulting hedge ratio is

$$h^{(DM)} = -\frac{D^A_\mu}{D^I_\mu}. \tag{5}$$

The corresponding hedged portfolio is $\mathcal{P}(h^{(DM)}) = \mathcal{A} + h^{(DM)}\mathcal{I}$. This portfolio matches the modified durations of the annuity and insurance portfolios, thereby aligning their sensitivities to mortality fluctuations.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

2.2.3 Greek neutral

Greek neutral methods extend the notion of Greeks to mortality and longevity hedging. Luciano *et al.* (2012) first introduced this idea, and Zhou and Li (2021) further developed it for index-based longevity hedging, where the longevity delta was defined as a sensitivity measure with respect to the period effect of a stochastic mortality model. We now adapt this approach to natural hedging.

Consider the Lee–Carter (LC) model (Lee and Carter, 1992):

$$\ln(m_{x,t}) = \alpha_x + \beta_x \kappa_t,$$

where $m_{x,t}$ is the central death rate for age x in year t, and α_x, β_x, and κ_t are the LC parameters, with κ_t following a random walk with drift. The longevity delta measures the first-order sensitivity of the survival probability $S_x(T)$ to changes in the current period effect κ_0 and is defined as $\Delta_x(T) := \frac{\partial S_x(T)}{\partial \kappa_0}$. Assuming a constant force of mortality between integer ages, Zhou and Li (2021) showed that $\Delta_x(T)$ can be written as:

$$\Delta_x(T) = -\sum_{s=1}^{T} \beta_{x+s-1}\, \mathbb{E}\big[\exp(Y_x(s) - Z_x(T))\big],$$

where $Y_x(s) = \alpha_{x+s-1} + \beta_{x+s-1}\kappa_s$ and $Z_x(T) = \sum_{s=1}^{T} \exp(Y_x(s))$, with $\Delta_x(0) = 0$ because $S_x(0) = 1$.

For the natural hedging framework, the longevity deltas of the annuity and insurance portfolios from Section 2.1 can then be derived as $\Delta^A = \sum_{j=1}^{n_A} \omega_j^A \sum_{k=\tau_j}^{\tau_j+t_j-1} c_j e^{-\delta k} \Delta_{x_j}(k)$ and $\Delta^I = \sum_{j=1}^{n_I} \omega_j^I \sum_{k=0}^{t_j-1} b_j e^{-\delta(k+1)} (\Delta_{x_j}(k) - \Delta_{x_j}(k+1))$. A delta-neutral natural hedge would require $\Delta^A + h\Delta^I = 0$, and thus the hedge ratio is

$$h^{(DN)} = -\frac{\Delta^A}{\Delta^I}. \tag{6}$$

The corresponding delta-neutral hedged portfolio is $\mathcal{P}(h^{(DN)}) = \mathcal{A} + h^{(DN)}\mathcal{I}$. Lastly, we note that longevity delta is model-dependent. While we have illustrated it under the LC model, the delta-neutral approach can be applied to other stochastic mortality models by deriving the corresponding longevity deltas.

2.3 Evaluation

The last step of the proposed natural hedging framework is evaluation. In this step, the insurer evaluates the effectiveness of a hedged position $\mathcal{P}(h)$ using both numerical and graphical assessment tools. We discuss both types of assessments in this subsection and consider them in the toy example shown in the next subsection.

A numerical risk measure for natural hedging can be defined in general as a mapping,

$$\rho : \mathcal{P}(h) \mapsto \mathbb{R},$$

which assigns a real-valued assessment to the hedged position $\mathcal{P}(h)$. A fundamental example is the portfolio variance, given by $\mathrm{Var}(\mathcal{P}(h)) = \mathrm{Var}(\mathcal{A}) + h^2\mathrm{Var}(\mathcal{I}) + 2h\,\mathrm{Cov}(\mathcal{A}, \mathcal{I})$. If the insurer focuses on downside risk, the VaR can be considered, defined as $\mathrm{VaR}_\alpha(\mathcal{P}(h)) = \inf\{x \in \mathbb{R} : F_{\mathcal{P}(h)}(x) \geq \alpha\}$ at confidence level $\alpha \in (0,1)$, where $F_{\mathcal{P}(h)}$ is the distribution function of $\mathcal{P}(h)$. A related measure is expected shortfall (ES), given by $\mathrm{ES}_\alpha(\mathcal{P}(h)) = \frac{1}{1-\alpha}\int_\alpha^1 \mathrm{VaR}_u(\mathcal{P}(h))du$, which captures the average loss in the tail beyond VaR. For a natural hedge, these numerical risk measures can provide complementary perspectives on the risk profile of the hedged position.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

Table 2. The proposed natural hedging framework

Step	Details
Step 1: Valuation	Derive and calculate the present values of annuity and insurance liabilities to be included in the hedge
Step 2: Calibration	Determine the hedge ratio, either set to 1 for an uncalibrated hedge or calculated under a chosen calibration method
Step 3: Evaluation	Assess the resulting hedged positions using numerical risk measures and graphical risk assessment tools
Projected mortality	Mortality scenarios are required in each of the valuation, calibration, and evaluation steps

We make two remarks on numerical risk measures. First, beyond variance, VaR, and ES, other measures may also be considered, and they can be applied to transformations of the hedged position. For example, one may evaluate hedge effectiveness using $\mathrm{VaR}_\alpha(\tilde{\mathcal{P}}(h))$, where $\tilde{\mathcal{P}}(h)$ is the mean-adjusted hedged portfolio given in Table 1. Second, we do not aim to identify a single numerical risk measure that is universally applicable to hedged portfolios. In optimization, the chosen objective is a natural candidate, but in immunization or Greek neutral approaches no obvious measure exists. In such cases, it is important to consider multiple numerical risk measures and complementary graphical assessments.

In conjunction with numerical measures, graphical risk metrics can further provide visual insights into the distribution of the hedged position. A common approach is to plot histograms of $\mathcal{P}(h)$ or $\tilde{\mathcal{P}}(h)$, and overlay numerical risk measures such as VaR and ES. Other useful visualizations include empirical cumulative distribution functions, Q-Q plots, or kernel density estimates. While these graphical tools are widely used, they remain limited in the depth of analysis they provide. In Section 3, we extend beyond such standard methods and propose a new graphical risk metric tailored for natural hedging.

In summary, the last step of our framework incorporates both numerical and graphical risk assessments to evaluate hedged positions, aiming to provide a comprehensive view of their risk profile. However, relying solely on numerical measures can overlook important aspects of hedge performance, and simple graphical diagnostics may also be insufficient. To illustrate these issues, the next subsection presents a toy example that demonstrates the implementation of the proposed natural hedging framework and motivates the new graphical risk metric introduced in Section 3.

2.4 Framework summary and toy example

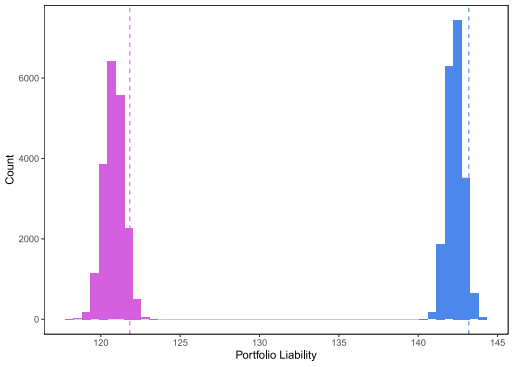
We now summarize the proposed natural hedging framework and illustrate its implementation through a toy example. Table 2 provides a compact reference for the three steps of the framework. Under stochastic mortality modeling, each step requires future mortality projections. In practice, Monte Carlo simulations are used to compute present values in the valuation step, hedge ratios in the calibration step, and risk measures in the evaluation step. The mortality simulation methods need not be identical across steps, which in turn allows model uncertainty to be present in the natural hedge. The mortality simulation methods considered in this paper are presented in Appendix B of the Supplementary Material.

To demonstrate the implementation of the proposed framework, we now present a toy example. Consider a life insurer with an annuity portfolio consisting of a single 20-year deferred 20-term life annuity issued to a 45-year-old individual, paying \$20 annually at the beginning of each year after the deferral period. To hedge this liability, the insurer issues a 30-year term life insurance policy with a death benefit of \$250, payable at the end of the year of death, to an individual aged either 40 or 50. For both issuing ages, the hedge ratio is set to $h = 1$, yielding two uncalibrated hedges. The interest rate is assumed to be $i = 4\%$ (or $\delta = \ln(1.04)$). We note that all

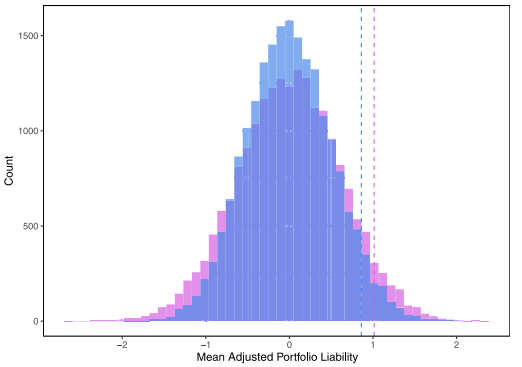
Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

Table 3. Numerical risk measures for the annuity portfolio $\mathcal{A}$, the two insurance portfolios $\mathcal{L}_1$ and $\mathcal{L}_2$, and the hedged portfolios $\mathcal{P}_1$ and $\mathcal{P}_2$

Portfolio	Mean	Variance	$\text{VaR}_{0.95}$	$\text{VaR}_{0.95} - \text{Mean}$
$\mathcal{A}$	98.81	1.61	100.84	2.03
$\mathcal{L}_1$	22.00	1.01	23.69	1.68
$\mathcal{L}_2$	43.52	2.67	46.27	2.75
$\mathcal{P}_1$	120.81	0.39	121.83	1.01
$\mathcal{P}_2$	142.33	0.27	143.19	0.86



(a) Liabilities for each hedged portfolio with pink and blue dashed lines denoting the $\text{VaR}_{0.95}$ for $\mathcal{P}_1$ (pink) and $\mathcal{P}_2$ (blue), respectively.



(b) Mean-adjusted liabilities for each hedged portfolio with pink and blue dashed lines denoting the $\text{VaR}_{0.95}$ for $\tilde{\mathcal{P}}_1$ and $\tilde{\mathcal{P}}_2$, respectively.

Figure 1. Graphical risk assessment for the hedged portfolios $\mathcal{P}_1$ (pink) and $\mathcal{P}_2$ (blue).

portfolios assumed are composed of a large number of identical contracts so that individual mortality randomness is fully diversified, and only systematic mortality risk arising from uncertain future mortality trends is being hedged.

Following the framework, Step 1 derives the present values of the annuity $\mathcal{A}$ and the two insurance portfolios $\mathcal{I}_1$ and $\mathcal{I}_2$ for ages 40 and 50, respectively. Step 2 constructs the calibrated insurance portfolios $\mathcal{L}_1(1) = \mathcal{I}_1$ and $\mathcal{L}_2(1) = \mathcal{I}_2$ by setting $h_1 = h_2 = 1$. Step 3 evaluates the hedged portfolios $\mathcal{P}_1 = \mathcal{A} + \mathcal{I}_1$ and $\mathcal{P}_2 = \mathcal{A} + \mathcal{I}_2$ using both numerical and graphical risk assessments. Mortality scenarios are generated using the bootstrapping method described in Appendix B.2 of the Supplementary Material, based on US male mortality data (ages 40–99, years 1970–2018) obtained from the Human Mortality Database.[1]

Table 3 reports the mean, variance, $\text{VaR}_{0.95}$, and the excess of $\text{VaR}_{0.95}$ over the mean for the annuity portfolio $\mathcal{A}$, the two insurance portfolios $\mathcal{L}_1$ and $\mathcal{L}_2$, and the combined hedged portfolios $\mathcal{P}_1$ and $\mathcal{P}_2$. The variance is greatly reduced when either $\mathcal{L}_1$ or $\mathcal{L}_2$ is added to $\mathcal{A}$, reflecting the hedging effect of combining portfolios. Between the two hedged portfolios, $\mathcal{P}_2$ has the smaller variance and excess VaR, while $\mathcal{P}_1$ achieves the lower mean and lower VaR. The mixed results underscore the complexity of assessing hedge effectiveness in natural hedging and motivate the need for further graphical risk assessments.

Figure 1 shows a graphical assessment of the hedged portfolios using histograms. Panel (a) shows the empirical distributions of $\mathcal{P}_1$ and $\mathcal{P}_2$, with dashed lines indicating their $\text{VaR}_{0.95}$.

[1]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 (data downloaded on 2 January 2024).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

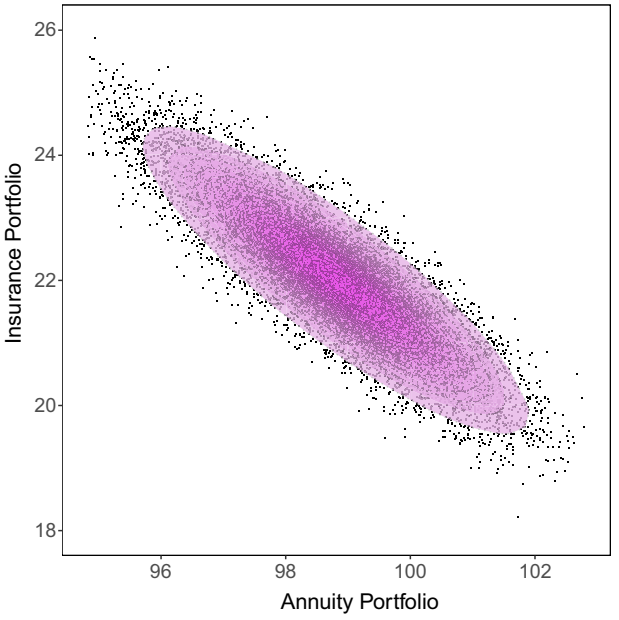


Figure 2. Graphical risk metric illustrating the joint distribution of the annuity portfolio $\mathcal{A}$ and the insurance portfolio $\mathcal{L}_1$ for the hedged position $\mathcal{P}_1$ from the toy example.

Consistent with the numerical results, $\mathcal{P}_1$ has a smaller $\text{VaR}_{0.95}$, reflecting lower liabilities compared with $\mathcal{P}_2$. Panel (b) displays the empirical distributions of the mean-adjusted portfolios $\tilde{\mathcal{P}}_1$ and $\tilde{\mathcal{P}}_2$ (defined in Table 1), again with dashed lines marking their $\text{VaR}_{0.95}$. In this panel, $\mathcal{P}_2$ shows smaller deviations from its mean than $\mathcal{P}_1$, confirming the lower variance observed in the numerical measures.

This toy example illustrates how different numerical and graphical risk assessments can lead to inconclusive results about hedge effectiveness. More importantly, such assessments cannot capture the underlying relationship between the annuity portfolio $\mathcal{A}$ and the insurance portfolio $\mathcal{L}$ within the hedged position $\mathcal{P}$. For example, when the realized liability of $\mathcal{P}$ exceeds its $\text{VaR}_{0.95}$, it is unclear whether this outcome reflects weak offsetting effect between $\mathcal{A}$ and $\mathcal{L}$, or whether both portfolios simultaneously contribute to large losses. Answering such questions can help insurers to make more informed decisions in portfolio selection and in the design of natural hedges, which we address in Section 3.

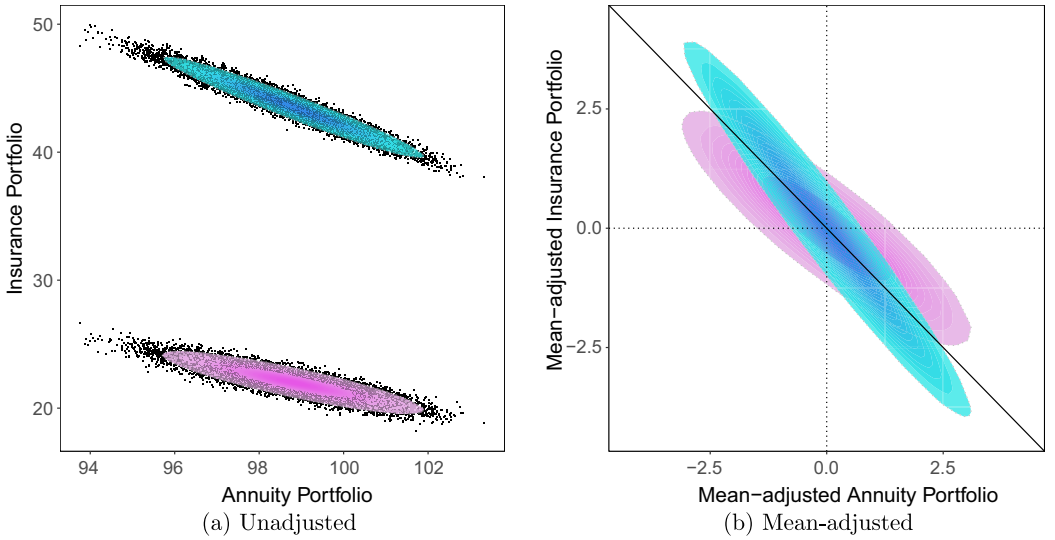
3. The graphical risk metric

This section introduces a new graphical risk metric for natural hedging. The metric provides a deeper analysis of the hedged position by examining the joint distribution of the underlying annuity and insurance portfolios on a two-dimensional scale. In the following, we first describe how this graphical risk metric is constructed and then develop an interpretation procedure that enables an enhanced risk assessment of the hedged position. To illustrate the construction and interpretation, we continue to use the toy example outlined in Section 2.4.

3.1 Constructing the metric

The construction of our graphical risk metric is inspired by the work of Chan *et al.* (2016), who introduced a visual tool for evaluating population basis risk in index-based longevity hedges. We adapt this approach to the context of natural hedging and construct a two-dimensional assessment

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X



(a) Unadjusted (b) Mean-adjusted

Figure 3. Graphical risk metric illustrating the effectiveness of the hedged positions $\mathcal{P}_1$ (pink) and $\mathcal{P}_2$ (blue) from the toy example.

of the joint distribution of the underlying annuity and insurance portfolios, rather than focusing solely on the hedged position. This representation enables direct visualization of the dependence structure between the annuity and insurance liabilities.

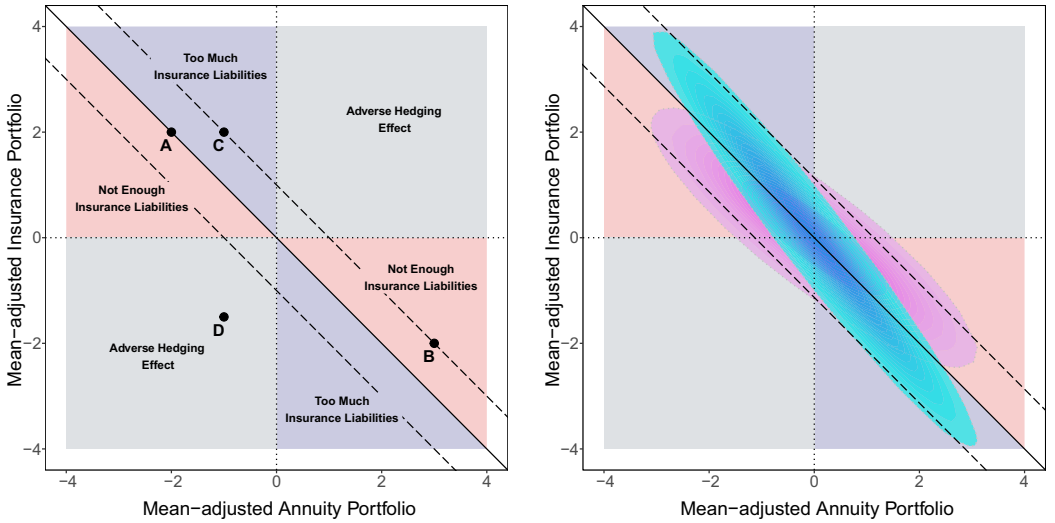
Figure 2 illustrates the result of applying the metric to the hedged position $\mathcal{P}_1$ from the toy example. Details on how this figure is constructed are provided in Appendix A of the Supplementary Material. The x- and y-axes represent the present values of the annuity portfolio $\mathcal{A}$ and the insurance portfolio $\mathcal{L}_1$, respectively. Each black dot corresponds to a realization of $(\mathcal{A}, \mathcal{L}_1)$, while the shaded pink regions denote joint prediction regions at different confidence levels. The figure reveals a clear inverse relationship between $\mathcal{A}$ and $\mathcal{L}_1$, forming a downward-sloping pattern in both the point cloud and the shaded regions.

The graphical risk metric illustrated in Figure 2 enables a visual evaluation of the hedge effectiveness of a single portfolio. The same construction can be extended to compare multiple hedged portfolios, which is particularly useful when several hedging opportunities are available. Figure 3 illustrates such an extension by jointly displaying the annuity portfolio $\mathcal{A}$ against two insurance portfolios, $\mathcal{L}_1$ and $\mathcal{L}_2$, corresponding to the hedged positions $\mathcal{P}_1$ and $\mathcal{P}_2$ from the toy example. This extension allows for a direct comparison of alternative natural hedges under a single visual platform.

Panel (a) of Figure 3 plots the annuity portfolio $\mathcal{A}$ against the corresponding insurance portfolios $\mathcal{L}_1$ and $\mathcal{L}_2$. As the annuity is identical across both portfolios, the larger insurance liabilities of $\mathcal{L}_2$ (relative to $\mathcal{L}_1$) result in higher total liabilities, consistent with the numerical measures reported in Table 3. This outcome is expected, as the life insurance policy from $\mathcal{P}_2$ is issued to a 50-year-old individual, whose higher mortality risk will lead to greater liabilities than the 40-year-old policy from $\mathcal{P}_1$.

To enable a fairer comparison, Panel (b) of Figure 3 presents the mean-adjusted annuity and insurance liabilities, represented by $(\tilde{\mathcal{A}}, \tilde{\mathcal{L}}_1)$ and $(\tilde{\mathcal{A}}, \tilde{\mathcal{L}}_2)$, whose joint prediction regions are centered at the origin. The solid black line with slope -1 represents the *benchmark* line, corresponding to scenarios in which deviations in the annuity and insurance liabilities perfectly offset (i.e., $\tilde{\mathcal{A}} + \tilde{\mathcal{L}} = 0$), resulting in a "perfect" hedge. Building on this benchmark, we next introduce an interpretative procedure to analyze potential hedging outcomes based on the mean-adjusted portfolios.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S1748499526100044X



(a) Interpretation (b) Application

Figure 4. Interpretation procedure for the proposed graphical risk metric and its application to the hedged positions $\mathcal{P}_1$ (pink) and $\mathcal{P}_2$ (blue) from the toy example.

3.2 Interpreting the metric

To gain a comprehensive visual understanding of the risk profile underlying a hedged position, we now develop an interpretation aid for the proposed graphical risk metric. This aid functions as a diagnostic tool for analyzing the joint behavior of the annuity and insurance portfolios. By representing all present value realizations in the $(\tilde{\mathcal{A}}, \tilde{\mathcal{L}})$ plane, it enables visual assessment of how effectively deviations in the insurance liability offset those in the annuity liability.

Figure 4 illustrates our interpretation procedure. Panel (a) defines the shaded regions and benchmark lines used to classify hedging outcomes, where the shading corresponds to different types of hedging outcome and the dashed lines indicate the magnitude of deviation from a perfect hedge. To clarify the interpretation of these shaded regions, we provide brief explanations of the scenarios marked by Points A, B, C, and D in Figure 4(a):

- **Point A**: The annuity liability is \$2 below its expected value, while the insurance liability is \$2 above. The surplus in the insurance liability fully offsets the annuity shortfall, representing a perfectly hedged scenario. All points along the benchmark line correspond to this ideal outcome.
- **Point B**: The annuity liability is \$3 above its expectation, but the insurance liability is only \$2 below. The shortfall in the insurance liability leaves a net liability surplus of \$1 for the hedged position, corresponding to the case of *Not Enough Insurance Liabilities* (red region).
- **Point C**: The annuity liability is \$1 below its expected value, while the insurance liability is \$2 above. The excess in the insurance liability leads to a net liability surplus of \$1 for the hedged position, representing the case of *Too Much Insurance Liabilities* (blue region).
- **Point D**: Both the annuity and insurance liabilities fall below their expectations, by \$1 and \$1.5 respectively. This scenario creates a combined liability deficit of \$2.5, indicating the case of *Adverse Hedging Effect* (gray region), where losses occur simultaneously in both portfolios.

The two dashed lines in Figure 4(a) represent equal magnitudes of deviation from the expected total liability, corresponding to either a surplus or a deficit of \$1. Because Point D lies below the lower dashed line, it immediately means that it has a larger magnitude of deviation than Points B and C. Although Points B and C both have a surplus of \$1, they arise from different causes; Point

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

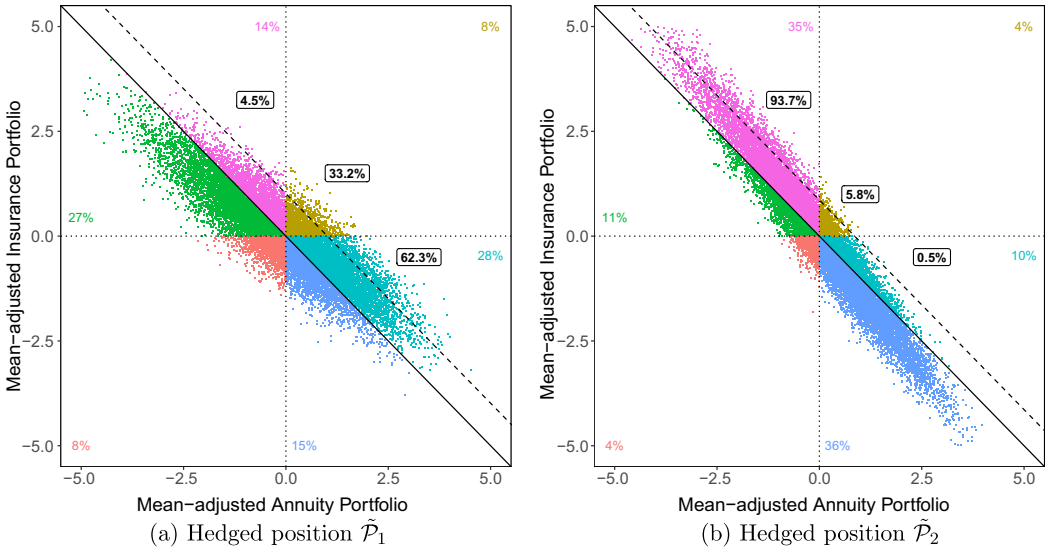


Figure 5. Simulated outcomes of $\tilde{\mathcal{P}}_1$ and $\tilde{\mathcal{P}}_2$, colored by hedging outcome with the corresponding proportion shown in each region; the dashed line denotes the mean-adjusted $\text{VaR}_{0.95}$, and the boxed black percentages indicate the share of each outcome exceeding the mean-adjusted $\text{VaR}_{0.95}$.

B has insufficient insurance liabilities, whereas Point C has excessive insurance liabilities. This distinction would not be apparent in a one-dimensional graphical assessment, such as a histogram. Figure 5 extends this idea by showing simulated realizations colored by hedging outcome types.

Figure 5 presents 20,000 simulated realizations of $(\tilde{\mathcal{A}}, \tilde{\mathcal{L}}_1)$ and $(\tilde{\mathcal{A}}, \tilde{\mathcal{L}}_2)$ in Panels (a) and (b), respectively. Points are colored by the type of hedging outcome. Pink and blue indicate *Too Much Insurance Liabilities*, leading to a surplus and deficit in total liability, respectively. Green and teal denote *Not Enough Insurance Liabilities*, leading to a deficit and surplus in total liability, respectively. Red and gold represent *Adverse Hedging Effect*, where both the insurance and annuity portfolios deviate in the same direction. The dashed lines mark the mean-adjusted $\text{VaR}_{0.95}$ of the corresponding hedged positions. Any realizations above these lines exceed the respective $\text{VaR}_{0.95}$ values reported in Table 3. Lastly, we caution that the proportion of realizations in each colored region reflects the probability of the corresponding hedging outcome, but does not directly indicate the magnitude of the associated financial gain or loss.

Comparing the simulated hedging outcomes in Figure 5, $\tilde{\mathcal{P}}_1$ shows more realizations associated with *Not Enough Insurance Liabilities* (green and teal) and *Adverse Hedging Effect* (red and gold) than $\tilde{\mathcal{P}}_2$. Conversely, $\tilde{\mathcal{P}}_2$ exhibits a higher frequency of *Too Much Insurance Liabilities* (pink and blue), indicating that its insurance exposure tends to overcompensate for annuity deviations. Among the outcomes that exceed the mean-adjusted $\text{VaR}_{0.95}$ as marked by the dashed lines, those from $\tilde{\mathcal{P}}_1$ primarily arise from *Not Enough Insurance Liabilities* (teal) and *Adverse Hedging Effect* (gold), while those from $\tilde{\mathcal{P}}_2$ are mostly driven by *Too Much Insurance Liabilities* (pink).

We now return to Figure 4(b), which applies the interpretation aid to compare $\tilde{\mathcal{P}}_1$ and $\tilde{\mathcal{P}}_2$ by overlaying their joint prediction regions. The dashed lines above and below the solid benchmark line represent equal magnitudes of surplus and deficit in the hedged position. It is clear that the joint prediction region of $\tilde{\mathcal{P}}_1$ extends further into the area associated with *Not Enough Insurance Liabilities*, while $\tilde{\mathcal{P}}_2$ covers more of the *Too Much Insurance Liabilities* area. More importantly, the joint prediction region of $\tilde{\mathcal{P}}_1$ extends above the upper dashed line, indicating more severe (liability surplus) scenarios than that of $\tilde{\mathcal{P}}_2$.

In conclusion, the proposed graphical risk metric complements numerical evaluation by exposing the underlying hedging behavior. Specifically, $\mathcal{P}_2$ achieves stronger offsets between annuity

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

Table 4. Natural hedging framework applied to Illustration 1

Step	Details
Step 1: Valuation	Derive the present values of the annuity portfolio (issuing ages 40–60, deferred to age 65, 35 annual payments of \$10,000) and three candidate insurance portfolios with a \$750,000 death benefit: $\mathcal{I}_1$ (issuing ages 40–60, whole life), $\mathcal{I}_2$ (issuing ages 40–49, whole life), and $\mathcal{I}_3$ (issuing ages 40–49, 20-year term)
Step 2: Calibration	Calibrate hedge ratios h_1, h_2, and h_3 using the variance-minimization method
Step 3: Evaluation	Construct the hedged portfolios $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_3$, and evaluate them comparatively using both numerical risk measures and the proposed graphical risk metric
Projected mortality	Mortality scenarios are generated using the bootstrapping method described in Appendix B.2 of the Supplementary Material, applied consistently across the valuation, calibration, and evaluation steps

Table 5. Calibrated hedge ratios and numerical risk measures for Illustration 1

Portfolio	Hedge ratio	Variance	$\text{VaR}_{0.95}$	Mean-adjusted $\text{VaR}_{0.95}$
$\mathcal{P}_1 = \mathcal{A} + h_1 \cdot \mathcal{I}_1$	0.267	34,090	133,295	290
$\mathcal{P}_2 = \mathcal{A} + h_2 \cdot \mathcal{I}_2$	0.275	112,252	123,496	563
$\mathcal{P}_3 = \mathcal{A} + h_3 \cdot \mathcal{I}_3$	0.311	875,566	90,413	1,571

and insurance liabilities but is more prone to excess insurance exposure, whereas $\mathcal{P}_1$ faces greater downside risk from insufficient offsetting. These findings highlight how the graphical risk metric reveals important behavioral differences between hedged positions. In the next section, we apply this graphical framework to a broader set of numerical illustrations to further explore its practical use and interpretive value.

4. Numerical illustrations

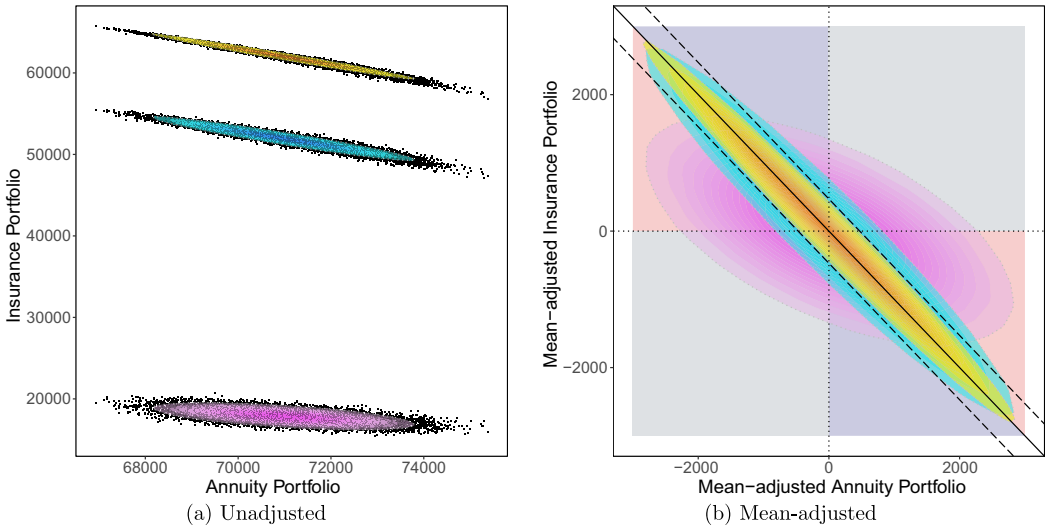
We now apply the natural hedging framework and the graphical risk metric to address stylized questions in natural hedging using public mortality data. Three illustrations are considered: (1) identifying the most effective insurance portfolio, (2) selecting the optimal hedge calibration technique, and (3) evaluating model risk arising from alternative mortality projections. Similar to the toy example, all portfolios in these illustrations are also assumed to be sufficiently large so that only systematic mortality risk remains. Each illustration follows the same framework as in Section 2 and employs the graphical risk metric developed in Section 3. Appendix C of the Supplementary Material provides the details of the annuity and insurance portfolios considered in this section.

4.1 Illustration 1: insurance portfolio selection

Our first illustration investigates how to identify an effective insurance portfolio for natural hedging. Consider an insurer managing an annuity portfolio issued to individuals aged 40–60, with \$10,000 annual payments starting at age 65. The insurer considers three candidate insurance portfolios, each providing a death benefit of \$750,000 payable at the end of the year of death: (1) whole life policies issued to individuals aged 40–60 ($\mathcal{I}_1$), (2) whole life policies issued to individuals aged 40–49 ($\mathcal{I}_2$), and (3) 20-year term policies issued to individuals aged 40–49 ($\mathcal{I}_3$). The annual interest rate is assumed to be 4%, and the limiting age is set at 100. Table 4 summarizes how the natural hedging framework is applied to this illustration. The three steps are applied consistently, with the only difference arising from the features of the insurance portfolios.

Table 5 reports the calibrated hedge ratios and numerical risk measures for the three hedged portfolios. Based on the variance and the mean-adjusted $\text{VaR}_{0.95}$, $\mathcal{P}_1$ outperforms $\mathcal{P}_2$ and $\mathcal{P}_3$, whereas the $\text{VaR}_{0.95}$ indicates that $\mathcal{P}_3$ performs best. To complement these numerical results and reveal the underlying hedging behavior, Figure 6 presents the graphical risk metric for the three

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X



(a) Unadjusted (b) Mean-adjusted

Figure 6. Graphical risk metric applied to the hedged positions $\mathcal{P}_1$ (yellow), $\mathcal{P}_2$ (blue), and $\mathcal{P}_3$ (pink) from Illustration 1.

hedged portfolios. In Panel (a), the unadjusted version shows that the three portfolios differ in the magnitude of their calibrated insurance liabilities, with $\mathcal{P}_1$ having the largest and $\mathcal{P}_3$ the smallest. The mean-adjusted version in Panel (b) shows that $\tilde{\mathcal{P}}_1$ aligns most closely with the benchmark line, indicating stronger offsetting between annuity and insurance liabilities, $\tilde{\mathcal{P}}_2$ exhibits greater variability with more outcomes falling into the *Not Enough Insurance Liabilities* and *Too Much Insurance Liabilities* regions, and $\tilde{\mathcal{P}}_3$ displays the widest dispersion and more frequent *Adverse Hedging Effect*.

In summary, this illustration demonstrates how the proposed framework and graphical risk metric together facilitate the selection of an appropriate insurance portfolio for natural hedging. Consistent with expectations, $\mathcal{P}_1$ with whole life policies issued to the same age range as the annuity portfolio provides the most effective hedge. Restricting the insurance portfolio to younger issuing ages, as in $\mathcal{P}_2$, increases variability and weakens the offset between annuity and insurance liabilities. Further shortening the policy term to 20 years, as in $\mathcal{P}_3$, amplifies this effect and leads to more outcomes with *Adverse Hedging Effect*. These results highlight that when issuing ages or policy terms deviate from those of the underlying annuity portfolio, the effectiveness of natural hedging reduces and the residual risk becomes more pronounced.

4.2 Illustration 2: hedge calibration

The second illustration examines how different hedge calibration methods affect the effectiveness of a natural hedge. Consider an insurer managing an annuity portfolio issued to individuals aged 40–60, with \$10,000 annual payments starting at age 65 for 20 years. The insurer constructs an insurance portfolio of whole life policies issued to the same age range and applies three different calibration methods for determining the hedge ratio: variance-minimization (VM), duration-matching (DM), and delta-neutral (DN). Table 6 summarizes how the natural hedging framework is applied to this illustration. The valuation and evaluation steps remain identical across the three methods, with the only difference lying in the calibration step.

Table 7 reports the calibrated hedge ratios and numerical risk measures for the three portfolios. The hedge ratios range between 0.189 and 0.211, with the smallest from duration matching and the largest from delta neutral. As expected, $\mathcal{P}^{(VM)}$ achieves the lowest variance while $\mathcal{P}^{(DN)}$ exhibits the highest risk across all measures. Figure 7 compares the three hedged portfolios using

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S1748499526100044X

Table 6. Natural hedging framework applied to Illustration 2

Step	Details
Step 1: Valuation	Derive the present values of the annuity portfolio (issuing ages 40–60, deferred to 65, 20 annual payments of $10,000) and the insurance portfolio of whole life policies with a $750,000 death benefit issued to individuals aged 40–60
Step 2: Calibration	Calibrate the hedge ratio using three different methods: $h^{(VM)}$ from the variance-minimization method, $h^{(DM)}$ from the duration-matching method, and $h^{(DN)}$ from the delta-neutral method
Step 3: Evaluation	Construct the hedged portfolios $\mathcal{P}^{(VM)}$, $\mathcal{P}^{(DM)}$, and $\mathcal{P}^{(DN)}$ and evaluate them comparatively using both numerical risk measures and the proposed graphical risk metric
Projected mortality	Mortality scenarios are generated using the Lee-Carter model described in Appendix B.1 of the Supplementary Material, applied consistently across the valuation, calibration, and evaluation steps

Table 7. Calibrated hedge ratios and numerical risk measures for Illustration 2

Portfolio	Hedge ratio	Variance	$\text{VaR}_{0.95}$	Mean-adjusted $\text{VaR}_{0.95}$
$\mathcal{P}^{(VM)} = \mathcal{A} + h^{(VM)} \cdot \mathcal{I}$	0.198	4,816	108,670	113
$\mathcal{P}^{(DM)} = \mathcal{A} + h^{(DM)} \cdot \mathcal{I}$	0.189	5,744	106,614	123
$\mathcal{P}^{(DN)} = \mathcal{A} + h^{(DN)} \cdot \mathcal{I}$	0.211	6,642	111,690	132

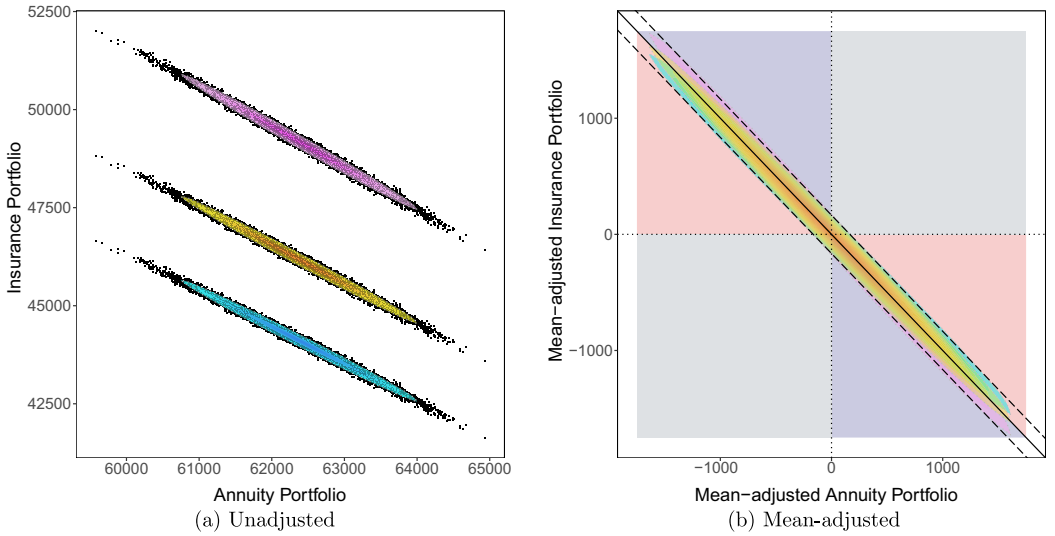


Figure 7. Graphical risk metric applied to the hedged positions $\mathcal{P}^{(VM)}$ (yellow), $\mathcal{P}^{(DM)}$ (blue), and $\mathcal{P}^{(DN)}$ (pink) from Illustration 2.

the graphical risk metric. In Panel (a), the unadjusted portfolios show that $\mathcal{P}^{(DM)}$ requires the least amount of insurance liabilities, while $\mathcal{P}^{(DN)}$ requires the most, consistent with their respective hedge ratios. Panel (b) displays the mean-adjusted portfolios, where $\tilde{\mathcal{P}}^{(VM)}$ lies closest to the benchmark line, indicating balanced offsetting between annuity and insurance liabilities. In contrast, $\tilde{\mathcal{P}}^{(DM)}$ tilts toward the *Not Enough Insurance Liabilities* region, and $\tilde{\mathcal{P}}^{(DN)}$ toward the *Too Much Insurance Liabilities* region, revealing systematic under- and over-hedging tendencies.

In summary, this illustration shows how the choice of hedge ratio calibration method influences both numerical performance and hedging behavior. The variance-minimization method

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

Table 8. Natural hedging framework applied to Illustration 3

Step	Details
Step 1: Valuation	Derive the present values of the annuity portfolio (issuing ages 40–60, deferred to age 65, 35 annual payments of \$10,000) and the whole-life insurance portfolio with a \$750,000 death benefit issued to the same ages
Step 2: Calibration	Calibrate the hedge ratio h using the variance-minimization method with mortality scenarios generated from the LC model
Step 3: Evaluation	Construct and compare the hedged portfolios $\mathcal{P}^{(LC)}$, $\mathcal{P}^{(CBD)}$, and $\mathcal{P}^{(BS)}$, which are evaluated using mortality scenarios generated by the LC, CBD, and BS models, respectively
Projected mortality	Mortality scenarios are generated from the LC model in the calibration step, and from the LC, CBD, and BS models in the evaluation step. The mortality generating processes are provided in Appendix B of the Supplementary Material

achieves the most balanced hedge, with the lowest variance and a symmetric distribution of hedging outcomes. The duration-matching approach tends to under-hedge, whereas the delta-neutral approach tends to over-hedge, leaving asymmetric residual risk exposure in the hedged position. The numerical and graphical results jointly provide a coherent view of the trade-offs underlying different calibration methods.

4.3 Illustration 3: model risk

Our final illustration examines the impact of model risk on natural hedging effectiveness. Suppose an insurer uses the LC model to calibrate a natural hedge but is concerned that actual mortality may not be accurately described by it. To investigate the risk that the mortality model assumed for calibration differs from the one generating actual future mortality, the insurer considers two additional models for evaluation: the Cairns–Blake–Dowd (CBD) model and a non-parametric bootstrapping (BS) model. The goal is to assess whether the hedged position constructed by the LC model is robust to alternative mortality generating methods.

Table 8 summarizes how the natural hedging framework is applied to this setting. The annual interest rate is again assumed to be 4%, and the limiting age is set at 100. Before presenting the numerical results, we briefly remark on model risk in general. A mortality model is suitable for a given application if it can adequately capture the aspects of mortality dynamics that are relevant to the question at hand. In particular, a model used for long-term stochastic projections may require different features from one used for short-term projections. This distinction is especially relevant to natural hedging, where the annuity and insurance liabilities involve different projection horizons. For illustration purposes, we only consider two well-known parametric models and one non-parametric model, which might lead to a misrepresentation of the true level of model risk. We therefore caution the reader that model risk can be substantial but can also be mitigated through careful model selection.

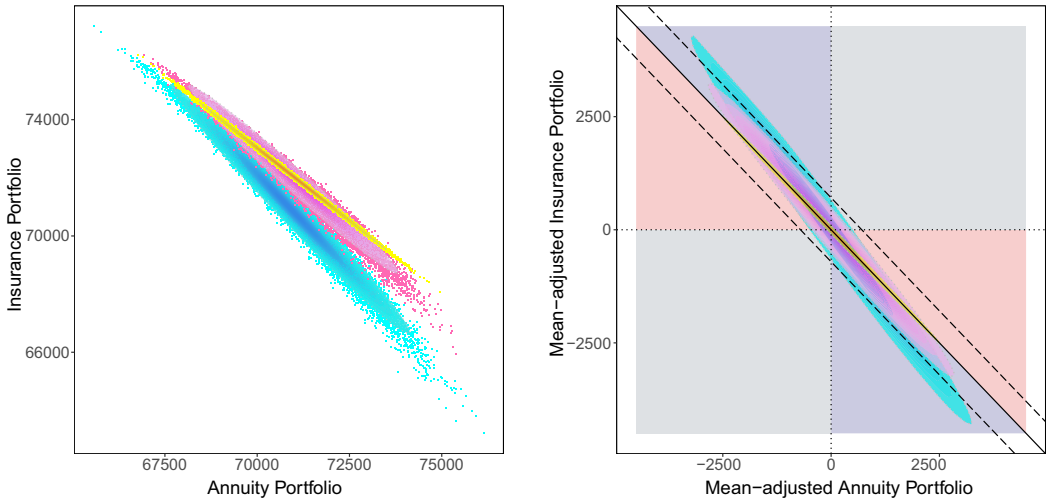
Table 9 reports the calibrated hedge ratio and numerical risk measures for the three portfolios. The hedge ratio is identical across all portfolios ($h = 0.310$). When evaluated by the same model used in the calibration step, $\mathcal{P}^{(LC)}$ exhibits the lowest variance and mean-adjusted $\text{VaR}_{0.95}$. However, when mortality is generated from an alternative model, $\mathcal{P}^{(CBD)}$ shows substantially higher variability, while $\mathcal{P}^{(BS)}$ lies between $\mathcal{P}^{(LC)}$ and $\mathcal{P}^{(CBD)}$ in both variance and mean-adjusted $\text{VaR}_{0.95}$. These results suggest that when a natural hedge is both calibrated and evaluated under the LC model, its hedge effectiveness may be significantly overestimated.

Figure 8 compares the three portfolios using the graphical risk metric. In Panel (a), the unadjusted version shows that $\mathcal{P}^{(LC)}$ exhibits the smallest joint prediction region, while $\mathcal{P}^{(CBD)}$ and $\mathcal{P}^{(BS)}$ display much wider dispersion in both the annuity and insurance portfolios. The

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

Table 9. Calibrated hedge ratio and numerical risk measures for Illustration 3

Portfolio	Hedge ratio	Variance	$\text{VaR}_{0.95}$	Mean-adjusted $\text{VaR}_{0.95}$
$\mathcal{P}^{(LC)}$	0.310	1,628	143,093	67
$\mathcal{P}^{(CBD)}$	0.310	245,611	142,538	697
$\mathcal{P}^{(BS)}$	0.310	66,711	143,290	350



(a) Unadjusted (b) Mean-adjusted

Figure 8. Graphical risk metric applied to the hedged positions $\mathcal{P}^{(LC)}$ (yellow), $\mathcal{P}^{(CBD)}$ (blue), and $\mathcal{P}^{(BS)}$ (pink) in Illustration 3.

mean-adjusted version in Panel (b) reveals that $\tilde{\mathcal{P}}^{(LC)}$ lies closest to the benchmark line, whereas both $\tilde{\mathcal{P}}^{(CBD)}$ and $\tilde{\mathcal{P}}^{(BS)}$ tilt toward the *Too Much Insurance Liabilities* region. This suggests that the LC-calibrated natural hedge tends to require an excessive amount of insurance liabilities, resulting in suboptimal hedging outcomes when actual mortality deviates from the LC model. This is an insight that cannot be easily inferred from the numerical results alone.

From Figure 8, we see that $\tilde{\mathcal{P}}^{(CBD)}$ extends beyond the dashed lines, while $\tilde{\mathcal{P}}^{(BS)}$ remains within them, indicating that the CBD model implies more volatile and extreme hedging outcomes than the non-parametric BS approach. To further illustrate their differences, Figure 9 plots simulated realizations of $\tilde{\mathcal{P}}^{(CBD)}$ and $\tilde{\mathcal{P}}^{(BS)}$, colored by hedging outcome. Both portfolios show a concentration of realizations in the *Too Much Insurance Liabilities* region. However, when compared against their respective mean-adjusted $\text{VaR}_{0.95}$ thresholds (marked by the dashed lines), the exceedances for $\tilde{\mathcal{P}}^{(CBD)}$ occur mainly due to *Too Much Insurance Liabilities*, whereas for $\tilde{\mathcal{P}}^{(BS)}$, the exceedances arise across all three hedging outcome types.

In summary, this illustration demonstrates that model risk can significantly distort the perceived effectiveness of a natural hedge. When the mortality model used for evaluation differs from the one used in calibration, the hedged position exhibits suboptimal performance and excess volatility under our natural hedging framework. The proposed graphical risk metric complements traditional numerical measures by revealing how these distortions arise from differing mortality dynamics and by exposing variations in hedging outcome types that numerical summaries alone would overlook.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S1748499526100044X

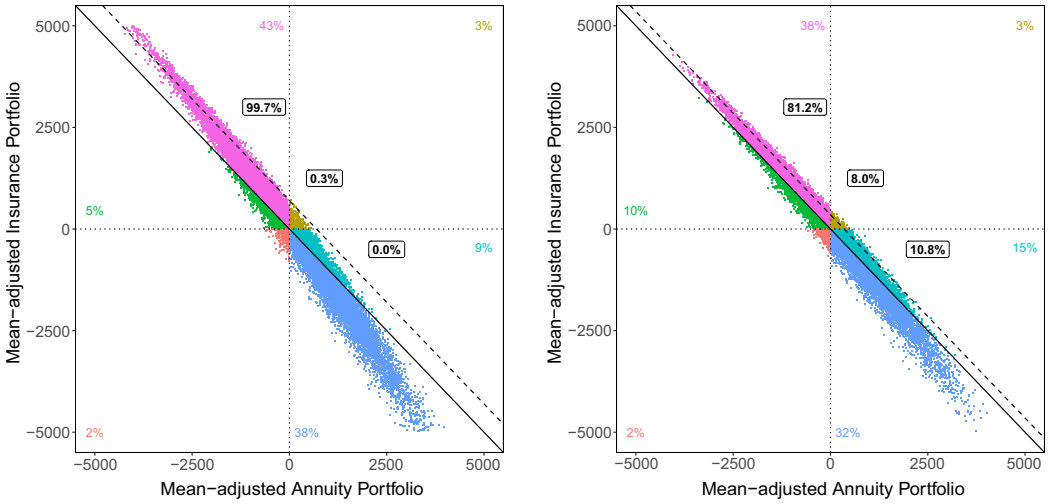


Figure 9. Simulated outcomes of $\tilde{\mathcal{P}}^{(CBD)}$ and $\tilde{\mathcal{P}}^{(BS)}$, colored by hedging outcome with the corresponding proportion shown in each region; the dashed line denotes the mean-adjusted $\text{VaR}_{0.95}$, and the boxed black percentages indicate the share of each outcome exceeding the mean-adjusted $\text{VaR}_{0.95}$.

5. Conclusion

This paper developed a comprehensive framework for constructing and evaluating natural hedging strategies, complemented by a graphical risk metric specifically designed for visual assessment of natural hedges. The proposed framework unifies the valuation, calibration, and evaluation steps of natural hedging to provide a structured process for comparing different portfolio settings, calibration techniques, and mortality scenarios. The graphical risk metric provides a new evaluation dimension by visualizing the joint distribution of annuity and insurance liabilities and by distinguishing different hedging outcome types. Together, these risk management tools form an integrated approach for analyzing longevity and mortality risks underlying a life insurer's balance sheet.

The proposed natural hedging framework and graphical risk metric offer several practical benefits. The framework enables consistent implementation of natural hedging across multiple policy types, calibration techniques, and model assumptions. It is also highly flexible and can accommodate various existing or newly developed natural hedging strategies beyond those demonstrated in this paper. The graphical risk metric complements the framework as an innovative evaluation tool for assessing and comparing multiple natural hedges when making hedging decisions. It further enhances interpretability by revealing whether poor performance arises from insufficient or excessive diversification, or from a lack of offsetting effects between liabilities.

Using three numerical illustrations, we demonstrated the capability and diagnostic value of the proposed methods. The results show that the framework can identify trade-offs among different insurance portfolios, assess calibration techniques across variance-minimization, duration-matching, and delta-neutral approaches, and evaluate hedge robustness under alternative mortality scenario generators. In all cases, the graphical risk metric reveals dependencies and asymmetries in hedge performance that are otherwise hidden within numerical risk measures. Although the underlying risks differ across illustrations, the proposed methods effectively address each scenario, demonstrating their flexibility and applicability under simulated future mortality.

This study has several limitations that call for further development. First, future research could extend the framework to incorporate stochastic interest rates, dynamic hedging strategies, and alternative mortality models. Second, the graphical risk metric can be enhanced by asymmetric or

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

non-convex joint prediction regions to capture skewed risk profiles and sudden mortality shocks. Finally, the natural hedging framework could be expanded to integrate solvency capital requirements and to support intuitive visualization for regulatory solvency assessments. Together, these potential directions would further advance the integration of natural hedging within longevity risk management research.

Finally, we note some practical considerations that may affect the implementation of natural hedging in practice. First, insurers may not always be able to adjust their insurance portfolio composition freely. For instance, selling additional life insurance policies in a specific age group may be constrained by market demand, underwriting capacity, or regulatory standards. Second, in certain jurisdictions, a significant portion of mortality or longevity gains may need to be distributed to policyholders due to surplus participation requirements, which limit the insurer's ability to compensate its associated annuity losses. In conclusion, the hedge ratios obtained from our framework should be interpreted as theoretical benchmarks rather than practical implementation guidelines.

Acknowledgments. The authors declare that generative artificial intelligence (AI) tools were not used in conceptualizing, designing, or conducting this research, including data collection, methodology development, analysis, or interpretation of results. Generative AI tools were employed exclusively for language enhancement and readability improvement. The authors assume full responsibility for the intellectual content and accuracy of the article.

Supplementary material. The supplementary material for this article can be found at: https://doi.org/10.1017/S174849952610044X.

Data availability statement. All numerical illustrations are based on publicly available mortality data from the Human Mortality Database (www.mortality.org). The code and data used in this study are available from the corresponding author upon reasonable request.

Funding statement. This research was supported by the *Natural Sciences and Engineering Research Council of Canada* under Grant No. RGPIN-2025-04157 and DGECR-2025-00488, and the *Society of Actuaries* under the Hickman Scholars program.

Competing interests. The authors declare that they have no competing interests.

References

Blake, **D.**, **Dowd**, **K.**, & **Cairns**, **A. J.** (2008). Longevity risk and the Grim Reaper's toxic tail: The survivor fan charts. *Insurance: Mathematics and Economics*, **42**(3), 1062–1066.

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

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

Chan, **W.-S.**, **Li**, **J. S.-H.**, **Zhou**, **K. Q.**, & **Zhou**, **R.** (2016). Towards a large and liquid longevity market: A graphical population basis risk metric. *The Geneva Papers on Risk and Insurance-Issues and Practice*, **41**(1), 118–127.

Chen, **H.**, **Gao**, **J.**, & **Zhu**, **W.** (2024). A unified framework for insurance demand and mortality immunization. *North American Actuarial Journal*, **28**(2), 469–490.

Coughlan, **G. D.**, **Khalaf-Allah**, **M.**, **Ye**, **Y.**, **Kumar**, **S.**, **Cairns**, **A. J.**, **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.

Cox, **S. H.**, & **Lin**, **Y.** (2007). Natural hedging of life and annuity mortality risks. *North American Actuarial Journal*, **11**(3), 1–15.

Cupido, **K.**, **Jevtić**, **P.**, **Regis**, **L.**, & **Zhou**, **K. Q.** (2024). Spatial natural hedging: A general framework with application to the mortality of U.S. states. *Scandinavian Actuarial Journal*, **2024**(10), 1036–1064.

Dowd, **K.**, **Blake**, **D.**, & **Cairns**, **A. J.** (2010). Facing up to uncertain life expectancy: The longevity fan charts. *Demography*, **47**, 67–78.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X

Gatzert, **N.**, & **Wesker**, **H.** (2012). The impact of natural hedging on a life insurer's risk situation. *The Journal of Risk Finance*, **13**(5), 396–423.

Gatzert, **N.**, & **Wesker**, **H.** (2014). Mortality risk and its effect on shortfall and risk management in life insurance. *Journal of Risk and Insurance*, **81**(1), 57–90.

Jevtić, **P.**, & **Regis**, **L.** (2015). Assessing the solvency of insurance portfolios via a continuous-time cohort model. *Insurance: Mathematics and Economics*, **61**, 36–47.

Lee, **R. D.**, & **Carter**, **L. R.** (1992). Modeling and forecasting US mortality. *Journal of the American Statistical Association*, **87**(419), 659–671.

Li, **J. S.-H.**, & **Liu**, **Y.** (2021). Recent declines in life expectancy: Implication on longevity risk hedging. *Insurance: Mathematics and Economics*, **99**, 376–394.

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

Lin, **T.**, & **Tsai**, **C. C.-L.** (2014). Applications of mortality durations and convexities in natural hedges. *North American Actuarial Journal*, **18**(3), 417–442.

Liu, **Y.**, & **Li**, **J. S.-H.** (2017). The locally linear Cairns–Blake–Dowd model: A note on delta–nuga hedging of longevity risk. *ASTIN Bulletin: The Journal of the IAA*, **47**(1), 79–151.

Luciano, **E.**, **Regis**, **L.**, & **Vigna**, **E.** (2012). Delta–Gamma hedging of mortality and interest rate risk. *Insurance: Mathematics and Economics*, **50**(3), 402–412.

Luciano, **E.**, **Regis**, **L.**, & **Vigna**, **E.** (2017). Single- and cross-generation natural hedging of longevity and financial risk. *Journal of Risk and Insurance*, **84**(3), 961–986.

Redington, **F. M.** (1952). Review of the principles of life-office valuations. *Journal of the Institute of Actuaries*, **78**(3), 286–340.

Sherris, **M.**, **Xu**, **Y.**, & **Ziveyi**, **J.** (2020). Cohort and value-based multi-country longevity risk management. *Scandinavian Actuarial Journal*, **2020**(7), 650–676.

Tsai, **J. T.**, **Wang**, **J. L.**, & **Tzeng**, **L. Y.** (2010). On the optimal product mix in life insurance companies using conditional value at risk. *Insurance: Mathematics and Economics*, **46**(1), 235–241.

Wang, **J. L.**, **Huang**, **H.**, **Yang**, **S. S.**, & **Tsai**, **J. T.** (2010). An optimal product mix for hedging longevity risk in life insurance companies: The immunization theory approach. *Journal of Risk and Insurance*, **77**(2), 473–497.

Yang, **S. S.**, **Huang**, **H.-C.**, & **Yeh**, **Y.-Y.** (2019). Optimal longevity hedging framework for insurance companies considering basis and mispricing risks. *Journal of Risk and Insurance*, **86**(3), 783–805.

Zhou, **K. Q.**, & **Li**, **J. S.-H.** (2021). Longevity Greeks: What do insurers and capital market investors need to know? *North American Actuarial Journal*, **25**(supp. 1), S66–S96.

Zhu, **N.**, & **Bauer**, **D.** (2011). Applications of forward mortality factor models in life insurance practice. *The Geneva Papers on Risk and Insurance-Issues and Practice*, **36**(4), 567–594.

Zhu, **N.**, & **Bauer**, **D.** (2014). A cautionary note on natural hedging of longevity risk. *North American Actuarial Journal*, **18**(1), 104–115.

Cite this article: Gabric LJ and Zhou KQ (2026). A natural hedging framework for longevity risk with graphical risk assessment, *Annals of Actuarial Science*, 1–21. https://doi.org/10.1017/S174849952610044X

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 13 Nov 2026 at 03:03:22, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S174849952610044X