UCLouvain

FACULTE DES SCIENCES

Faculté des sciences

Techniques de couverture profonde pour les contrats en unités de compte.

Auteur·es : Stéphanelle Odetta Takougue Ngamgo

Promoteur·rices : Karim Barigou

Lecteur·rices : Donation Hainaut

Année académique 2024-2025

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
FACULTÉ DES SCIENCES
ÉCOLE DE STATISTIQUE, BIOSTATISTIQUE ET SCIENCES ACTUARIELLES

Deep Hedging Techniques for equity-linked life insurance contracts

Stephanelle Odetta Takougue Ngamgo
Supervisor : Karim Barigou
Reader : Donatien Hainaut

Master thesis
Academic Year 2024-2025

Abstract

This thesis explores the application of deep learning techniques to the hedging of equity-linked life insurance contracts with long-term guarantees, focusing in particular on Guaranteed Minimum Death Benefits (GMDB). Such contracts expose insurers to multiple sources of risk, including equity market volatility, interest rate fluctuations, and mortality uncertainty.

We present a unified simulation framework that combines Black–Scholes dynamics for equity prices, Hull–White dynamics for stochastic interest rates, and deterministic (Gompertz–Makeham) and stochastic Ornstein–Uhlenbeck mortality models. Within this setting, we compare traditional Local Risk Minimization (LRM) and a Long Short-Term Memory (LSTM) network, both implemented under the deep hedging methodology.

Our numerical experiments show that deep hedging, the LSTM architecture, consistently reduces hedging errors and tail risks (measured by Value-at-Risk and Conditional Value-at-Risk) relative to the classical benchmark. These findings underscore the ability of neural networks to capture path-dependency and complex interactions in incomplete insurance markets.

It offers actuaries and regulators some important pointers as to how deep learning is likely to be a key part of the future when it comes to insurance risk management and prudential supervision.

Acknowledgment

First and foremost, I would like to thank my supervisor, Professor Karim Barigou, for his availability, and very helpful, clever advice during this thesis. His guidance and kindness have allowed me to conduct this research under the best conditions.

I would also like to thank all the faculty members of the Institute of Statistics, Biostatistics, and Actuarial Sciences at UCLouvain for the quality of the training and their contribution to the research field.

I am also deeply grateful to my close friends for their constant presence, their encouragement, and their ability to bring balance and joy into my life.

Lastly, there are no words that will express my gratitude to my family for their unwavering love, patience, and trust in me during the time of this study. Special thanks to my big brothers, Hermann Verde and Judicaël Briand, for their academic guidance. Without them, none of this would have been possible.

And the best for last: thank you, God.

Contents

Introduction

Life insurance products have become increasingly innovative in recent years, particularly with the development of equity-linked contracts, such as Unit-Linked Life Insurance policies. These products allow policyholders to invest in financial markets while benefiting from embedded guarantees that protect them against downside risk. Typical guarantees include the Guaranteed Minimum Death Benefit (GMDB), which pays a minimum amount upon death, and the Guaranteed Minimum Accumulation Benefit (GMAB), which guarantees a minimum return at maturity. From the insurer's perspective, these features transform the contract into a complex contingent claim whose value and risk depend on financial markets, interest rates, and mortality developments.

Managing these risks is a central concern in actuarial practice. In classical frameworks, financial risks are typically modeled using the Black-Scholes model, where stock prices are assumed to follow a geometric Brownian motion. Mortality risk is often assumed to be deterministic and governed by standard laws such as Gompertz-Makeham. These assumptions provide analytical tractability, allowing the use of closed-form formulas or PDE methods for pricing and hedging. However, real markets are incomplete. Although mortality rates may exhibit dynamics and random shocks (e.g., pandemics, medical gain), interest rates have stochastic behavior. Imperfections within such settings bring about the inability to replicate perfectly, so standard delta hedging also becomes suboptimal.

To address these challenges, several extensions have been proposed. On the actuarial side, stochastic mortality models such as the Cox-Ingersoll-Ross (CIR) process or Ornstein–Uhlenbeck processes introduce randomness in the force of mortality to account for longevity risk and demographic uncertainty (Luciano Luca and Vigna (2012) [7]). On the financial side, stochastic interest rate models such as Hull-White provide more realistic dynamics for long-dated insurance liabilities. These richer frameworks allow for more accurate valuation but require more advanced hedging techniques.

In incomplete markets, hedging is often approached through the theory of risk-minimizing strategies. Local Risk Minimization (LRM), introduced by Föllmer and Schweizer (1991) [8], focuses on minimizing the variance of the hedging error over time. Although LRM is not riskless, it suffices as a systematic and theoretically valid approach that can be formulated via minimal hypotheses. It is based on model specification, although still effective in the presence of non-linear payoffs or intricate relationships among risks.

In parallel, machine learning has opened new avenues in financial and actuarial modeling. Neural networks, particularly deep learning architectures, have shown remarkable capabilities in function approximation, time series modeling, and dynamic decision-making. Recent contributions have therefore turned to machine learning approaches. Buehler et al. [4] introduced the *deep hedging* framework, which leverages neural networks to approximate optimal dynamic strategies under general market conditions. Following this line of research, Barigou et al. [2] applied reinforcement learning techniques to life insurance guarantees, demonstrating that neural networks can effectively handle the joint presence of financial and biometric risks, and outperform classical strategies in incomplete markets.

Several studies have extended the deep hedging framework to insurance-related products. For instance, Carbonneau (2021) [1] applies deep reinforcement learning to the hedging of variable annuities with embedded guarantees, such as Guaranteed Minimum Maturity Benefits (GMMB) with ratchet features, which can be represented as lookback put options. His numerical experiments show that deep hedging strategies, especially when trained with non-quadratic loss functions such as downside-risk or CVaR penalties, significantly outperform traditional quadratic-risk benchmarks, leading to both reduced tail-risk exposures and improved expected hedging outcomes.

This thesis contributes to this line of research by comparing classical and deep learning-based hedging strategies in the context of GMDB guarantees. We simulate a realistic environment combining equity risk (modeled by Black-Scholes dynamics), stochastic interest rates (via the Hull-White model), and both deterministic and stochastic mortality models (using the Gompertz law and an Ornstein–Uhlenbeck process). Three types of hedging strategies are implemented: local risk minimization, as a benchmark, and deep hedging using a Long Short-Term Memory (LSTM) network. The models are evaluated in terms of their ability to reduce hedging error, Value-at-Risk (VaR), and Conditional Value-at-Risk (CVaR).

The central research question proposed in the thesis is: *How to hedge the Guaranteed Minimum Death Benefits (GMDB) of insurers in an incomplete market with stochastic interest rates and mortality rates?*. Addressing such a question demands more than traditional methods for e.g., LRM and should be tackled in a data-driven approaches.

The originality of this work lies in the development of a unified simulation framework that integrates equity, interest-rate, and mortality risks, and in the application of Long Short-Term Memory (LSTM) networks to learn dynamic hedging strategies. We show that these models not only reduce mean squared hedging errors, but also significantly improve tail-risk measures such as Conditional Value-at-Risk (CVaR), which are directly relevant to solvency regulation.

The remainder of this thesis is organized as follows. Chapter 1 develops the mathematical foundations, introducing the discrete-time market setting, risk measures, and the universal approximation capabilities of neural networks. Chapter 2 presents the Black–Scholes model as the theoretical benchmark for pricing and hedging, and discusses its limitations when applied to real markets. Chapter 3 extends the analysis to equity-linked life insurance contracts, formulating the hedging problem in incomplete markets and presenting the deep hedging framework, together with its numerical implementation and performance analysis. Finally, the conclusion summarizes the main findings and suggests avenues for future research.

Chapter 1

Discrete Market Models and Neural Hedging Foundations

This chapter provides the mathematical background to assess hedging strategies in financial markets, preparing for their implementation with neural networks. We begin by modeling a discrete-time market under realistic constraints, incorporating transaction costs and time-varying information sets. Unlike classical models that presume market completeness and the presence of an equivalent martingale measure (EMM), we deliberately leave these assumptions open.

We cast the problem as an optimal hedging problem in the agent-asset framework, in which the agent aims to hedge against a liability modeled by a terminal random variable Z with a self-financing trading strategy in d assets. This environment includes realistic factors, such as trading constraints and market imperfections. We define a terminal P&L including gains, losses, and execution costs. This general setup enables us to construct risk-based optimization problems, which we can numerically solve later on with the help of neural networks. This chapter thus lays out the key mathematical tools and modeling assumptions necessary for the deep hedging framework developed throughout the thesis.

1.1 Discrete time market

We concentrate on a discrete-time financial market, in the sense that prices and transaction values are observed, adapted, or simulated at specific points in time. This is a reasonable assumption given that they fluctuate constantly.

Let us consider a finite time horizon T and trading dates:

$$0 = t_0 < t_1 < \cdots < t_n = T.$$

We work on a finite probability space $(\Omega, \mathscr{F}, \mathbb{P})$, where $\Omega = \{\omega_1, \ldots, \omega_m\}$, and $\mathbb{P}$ is the physical probability measure. This space is equipped with a filtration $\mathscr{F} := \{\mathscr{F}_{t_k}\}_{k=0}^{n}$, which represents the flow of market information available to investors over time.
We introduce a $\mathbb{R}^d$-valued stochastic process $S = (S_k)_{k=0}^{n}$, representing the prices of d hedging instruments. At each time t_k, the available market information is encoded in a variable $I_k \in \mathbb{R}^r$, which may include asset prices, liquidity signals, regulatory limits, or other relevant features. The sequence $I = (I_k)_{k=0}^{n}$ generates the filtration $\mathscr{F}$, thereby guiding trading decisions at each time point.

Importantly, we do not assume the existence of an equivalent martingale measure (EMM), nor the continuous availability of all hedging instruments. Recall that a martingale is a process whose conditional expectation equals its current value, making it a "fair game." An EMM is typically used in risk-neutral pricing, but deep hedging instead minimizes risk under the real-world measure.

The liability to hedge is modeled as an $\mathscr{F}_T$-measurable random variable Z. The agent selects a self-financing trading strategy $\delta = (\delta_k)_{k=0}^{n-1}$, where each $\delta_k = (\delta_k^1, \ldots, \delta_k^d) \in \mathbb{R}^d$ represents the asset holdings at time t_k. For convenience, we define $\delta_{-1} = \delta_n := 0$, implying full liquidation at maturity.

Trading strategies are constrained: $\delta_k \in \mathscr{H}_k$, where $\mathscr{H}_k$ is the image of a measurable map $H_k : \mathbb{R}^{d(k+1)} \to \mathbb{R}^d$, continuous and $\mathscr{F}_k$-adapted, with $H_k(0) = 0$. These constraints reflect market frictions for instance, an option available only after three months would enforce $\delta_k = 0$ for earlier times.

If we first ignore transaction costs, the portfolio value at time T is:

$$PL_T = -Z + p_0 + (\delta \cdot S)_T, \quad \text{where } (\delta \cdot S)_T := \sum_{k=0}^{n-1} \delta_k \cdot (S_{k+1} - S_k).$$

To account for execution costs, we assume that trading a position $n \in \mathbb{R}^d$ at time t_k incurs a cost $c_k(n)$. The total cost over the strategy δ is then:

$$C_T(\delta) := \sum_{k=0}^{n} c_k(\delta_k - \delta_{k-1}).$$

The resulting P&L becomes:

$$PL_T(Z, p_0, \delta) := -Z + p_0 + (\delta \cdot S)_T - C_T(\delta).$$

The cost functions $c_k : \mathbb{R}^d \to \mathbb{R}_+$ are assumed to be non-negative, upper semi-continuous, and normalized so that $c_k(0) = 0$. These assumptions ensure the costs behave predictably and avoid undesirable jumps.

Finally, we assume the strategy is self-financing: trading adjustments are made solely from proceeds or losses of the assets held, without injecting additional capital. This is particularly relevant in options and derivatives hedging. While this problem can be interpreted in the language of reinforcement learning with states, actions, and rewards we adopt the formalism of mathematical finance.

To assess and compare trading strategies, we will next describe a framework built based on convex risk measures. These measures quantify risk-adjusted performance and and will be used as objective functions in our optimization.

1.2 Convex Risk Measure

In incomplete markets with frictions, it is necessary to define an optimality criterion. This allows us to determine an acceptable price for a contingent claim. This price represents the minimum cash injection necessary to add to a position to implement the optimal hedge, ensuring that the overall position meets the acceptability standards based on the agent's preferences. Our objective is to find the minimum capital required to make a position acceptable according to the risk measure. This acceptability is defined using convex risk measures. The main properties of such a functional are the following (Buelher et al (2019)) [4] :

Definition 1.1. *Assume that $X, X_1, X_2 \in \mathscr{X}$ represent asset positions (i.e., $-X$ is a liability). We call $\rho : \mathscr{X} \to \mathbb{R}$ a convex risk measure if it is:*

1. *Monotone decreasing: if $X_1 \geq X_2$ then $\rho(X_1) \leq \rho(X_2)$.*
   A more favorable position requires less cash injection.

2. *Convex: $\rho(\alpha X_1 + (1-\alpha)X_2) \leq \alpha\rho(X_1) + (1-\alpha)\rho(X_2)$ for $\alpha \in [0,1]$.*
   Diversification works.

3. *Cash-Invariant: $\rho(X+c) = \rho(X) - c$ for $c \in \mathbb{R}$.*
   Adding cash to a position reduces the need for more by as much. In particular, this means that $\rho(X + \rho(X)) = 0$, i.e., $\rho(X)$ is the least amount c that needs to be added to the position X in order to make it acceptable in the sense that $\rho(X+c) \leq 0$.

We call ρ normalized if $\rho(0) = 0$.
For $X \in \mathscr{X}$ consider the optimization problem

$$\pi(X) := \inf_{\delta \in \mathscr{H}} \rho(X + (\delta \cdot S)_T - C_T(\delta)). \tag{1.1}$$

with $\mathscr{H}$ the set of restricted trading strategies.

The monotonicity property means that a position X_1 which is uniformly more favorable than X_2 requires less additional capital. Convexity formalizes the benefit of diversification. Cash-invariance reflects that cash is risk-free and offsets risk capital one-to-one.

The quantity $\pi(X)$ is the minimum initial capital that must be injected into position X when using an optimal strategy $\delta^* \in \mathscr{H}$, so that the overall position becomes acceptable according to ρ.

Proposition 1.2. *π is monotone decreasing and cash-invariant. If, in addition, $C_T(\cdot)$ and $\mathscr{H}$ are convex, then the functional π is a convex risk measure.*

You can find the detailed demonstration in the section 3 of *H. a. G. Buehler Lukas and Teichmann* [4] .

1.2.1 Capital Requirement and Optimal Hedging

We define an optimal hedging strategy $\delta \in \mathscr{H}$ and interpret $\rho(-Z)$ as the minimal capital that must be added to the risky position $-Z$ to render it acceptable under the convex risk measure ρ. Consequently, $\pi(-Z)$ represents the minimal amount of capital required by the agent to make the final position acceptable when hedging optimally, while $\pi(0)$ corresponds to the capital needed in the absence of any exposure. Recall that a position is considered acceptable if it achieves a suitable trade-off between return and risk, in accordance with operational and regulatory constraints.

The indifference price $p(Z)$ is the amount of cash (or capital) an agent must charge to be indifferent between taking a risky position Z or not taking it. The formula is given by : $\pi(-Z + p_0) = \pi(0)$, with p_0 defined as the solution. By cash invariance, this is equivalent to taking $p_0 := p(Z)$, where

$$p(Z) := \pi(-Z) - \pi(0). \tag{1.2}$$

Without trading restrictions and transaction costs, this price coincides with the price of a replicating portfolio (if it exists). A replicating portfolio is a portfolio that reproduces the cash flow (or financial results) of another asset or position. In this context, it means that the replicating portfolio generates exactly the same financial results as the risky position Z. If such a portfolio exists, it allows perfect hedging of the risk associated with the position.

We can approximate optimal hedging strategies in a setting where the price p_0 is given exogenously: Let a loss function $L : \mathbb{R} \to [0, \infty)$. Suppose $p_0 > 0$ is given, for example being the result of trading derivatives in the market at competitive prices, without taking into account risk-management. The agent aims to minimize her loss at maturity, i.e. she defines an optimal hedging strategy as a minimizer to

$$\inf_{\delta \in \mathscr{H}} \mathbb{E}\left[L(-Z + p_0 + (\delta \cdot S)_T - C_T(\delta))\right].$$

The choice of the loss function is based on the objectives and the risk aversion of the hedger. As an example we have the mean-square error $L(x) = x^2$ or else the semi-mean square error $L(x) = x^2 \mathbf{1}_{\{x>0\}}$ which only penalize hedging losses.

1.2.2 Optimized Certainty Equivalents

Let $U(x) = -e^{-\lambda x}$ denote the exponential utility function. The certainty equivalent (or indifference price) $q(Z)$ solves:

$$q(Z) := \sup_{\delta \in \mathscr{H}} \mathbb{E}\left[U\left(q(Z) - Z + (\delta \cdot S)_T - C_T(\delta)\right)\right] = \sup_{\delta \in \mathscr{H}} \mathbb{E}\left[U\left((\delta \cdot S)_T - C_T(\delta)\right)\right].$$

Lemma 1.3. *Let ρ be the entropic risk measure:*

$$\rho(X) = \frac{1}{\lambda} \log \mathbb{E}[\exp(-\lambda X)].$$

Then, the indifference price $p(Z)$ defined by (1.2) *coincides with $q(Z)$.*

1.2.3 Risk Measures via Loss Functions

Definition 1.4. *Let $\ell : \mathbb{R} \to \mathbb{R}$ be a continuous, convex, and non-decreasing function (loss function). The associated risk measure is:*

$$\rho(X) := \inf_{w \in \mathbb{R}} \left\{w + \mathbb{E}[\ell(-X - w)]\right\}, \quad X \in \mathscr{X}.$$

Such measures are called *optimized certainty equivalents*, and can accommodate a wide range of utility or loss preferences. For examples, see Section 3.3 in [4].

Having introduced the concept of convex risk measures as a framework for quantifying and minimizing financial risk, we now explore how neural networks can be used to approximate the corresponding optimal hedging strategies. The following section presents the foundations of neural networks and their relevance in the context of finance and insurance.

1.3 Neural Networks in finance and insurance

In recent years, the insurance and banking industries have increasingly relied on data to support risk management, pricing, and customer profiling. Questions such as determining a client's profile or estimating the probability of accident involvement have become central to insurers. On the other hand, banks assess credit risk. What is the probability that a client fails to repay his debt? These challenges have accelerated the integration of data science and artificial intelligence into financial decision-making processes.

With the fast evolution of artificial intelligence, it has become a must in risk management. As we outlined in the previous section and will further develop in the next chapter, standard hedging techniques often struggle in incomplete markets, particularly when faced with high-dimensional, path-dependent, or non-linear dynamics. Deep hedging methods, based on neural networks, provide a flexible alternative capable of learning optimal strategies from data.

Deep hedging uses neural networks to estimate future prices of financial assets under different market states and manages portfolios based on trading strategies that minimize a specified loss. The cost of hedging can then be decreased by insurers, increasing the degree of protection from unfavorable market scenarios.

Biological Inspiration and Neural Network Architecture In biology, neural cells (neurons) work in parallel and reorganize during the learning phase. Neurons receive input signals throughout their dendritic trees (*Deep Learning for insurance and finance* Hainaut (2024) [3]). After receiving it, it "activates" and transmits a signal through its axon. This is the logic that inspired the development of artificial neural networks. A neural network is a set of interconnected neurons with infinitely many possible layouts. Every neural is usually represented by a mathematical function that takes several entries, combines them with weights, applies a function of activation, and produces an output. Deep neural networks have many hidden layers that help us learn from hierarchical representations of complex data. Globally, an artificial neural network can be viewed as a composition of both linear and nonlinear functions, potentially repeated, which maps an input x to an output y.

Artificial neural networks are structured as layers of interconnected neurons, each performing a basic transformation. Deep-learning architectures, in general, have many layers of hidden units that facilitate learning progressively more abstract representations of data. At their core, these networks can be expressed as a composition of linear and nonlinear functions, which map inputs x to outputs y, with parameters trained to minimize a predefined cost function over a dataset.

1.3.1 Feed-forward neural networks

There are different types of neural network attached to specific tasks. A Feed-forward neural network (FFNN) is a parametrized composite function that maps input to output vectors through the composition of a sequence of hidden layers (*A. Carbonneau* [1]). Each hidden layer performs an affine transformation followed by the application of a non-linear activation function. Consider each node as its linear regression model consisting of input, weighting, bias, and output data. Weights help to identify the importance of a variable.

A FFNN is characterized by a diffusion of unidirectional flux. Moreover perceptrons, also known as feed-forward networks, consist of neurons arranged in interconnected layers and are primarily used for classification or regression. Perceptrons, first introduced by Rosenblatt (1958), use supervised learning, requiring prior identification of relevant variables and expected outputs. For example, predicting car accident frequency with a perceptron involves segmenting explanatory variables beforehand (Denuit, Hainaut & Truffin (2019))[10]. Now we will introduce the mathematical definition [1].

Definition 1.5. *A FFNN $F_\theta : \mathbb{R}^{d_0} \to \mathbb{R}^{\tilde{d}}$ with L hidden layers has the following representation:*

$$F_\theta(X) := o \circ h_L \circ \cdots \circ h_1,$$
$$h_l(X) := \phi(W_l X + b_l), \quad l = 1, \ldots, L,$$

where $W_l \in \mathbb{R}^{d_l \times d_{l-1}}$ and $b_l \in \mathbb{R}^{d_l \times 1}$ are respectively known as the weight matrix and bias vector of the l-th hidden layer h_l. The function g is non-linear and is applied to each scalar given as input and $o : \mathbb{R}^{d_L} \to \mathbb{R}^{\tilde{d}}$ is the output function which applies an affine transformation to the output of the last hidden layer h_L and possibly also a nonlinear transformation with the same range as F_θ.

Furthermore, the trainable parameters θ is the set of all weight matrices and bias vectors that are learned (i.e., fitted in statistical terms) by minimizing a specified cost function.

The function ϕ is a univariate and continuous activation function. The most known are functions are the sigmoïd and the tangent hyperbolic respectively given by $\phi(x) = \frac{1}{1+\exp(-x)}$ and $\phi(x) = \frac{\exp(x)-\exp(-x)}{\exp(x)+\exp(-x)}$. The activation functions are essential for approximating functions.

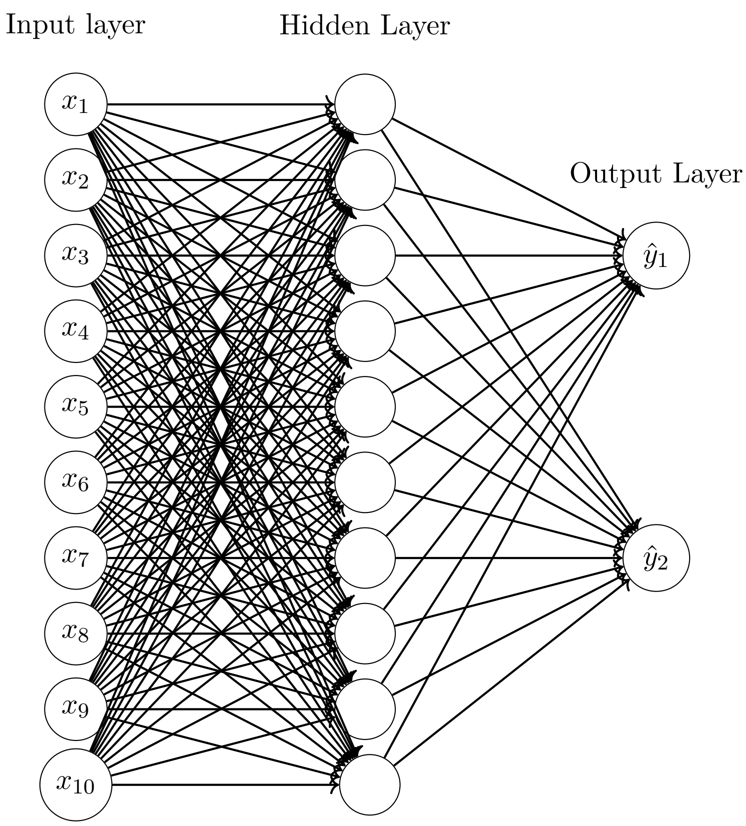

Figure 1.1: a neural network with one hidden layers, architecture 10-10-2

Figure 1.1 represent a neural network with one hidden layer. The input nodes represent the input variables and the output nodes represent the responses.

The universal approximation properties of deep feedforward networks are presented in Section 1.3.3. They are very effective at approximating hedging strategies. They quantify the minimum connectivity required within the network to approximate all functions of a given class with a specified level of precision. This establishes a universal relationship between network connectivity and the complexity of the approximate function class. Let introduce a specific type of recurrent neural network.

1.3.2 Long Short-Term Memory (LSTM)

Recurrent Neural Networks (RNNs) are a type of neural network that are especially good at modeling sequential data. They have recurrent connectives through which the network can store and carry information from a previous step to another. These architectures have been successful in areas such as Natural Language Generation, Video Processing, and other sequence based problems. However, traditional RNNs have difficulty modeling long-range dependencies because of two issues: the vanishing gradient problem, where gradients decay exponentially and prevent learning arising from long sequences, and the exploding gradient problem, where gradients become unbounded, causing unstable training.

In order to mitigate these problems, more advanced RNN structures, such as Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU), have been proposed. In this thesis, we will concentrate on LSTM networks: indeed they work better at learning long-term dependencies and, as we will demonstrate with the help of *A. Carbonneau* [1], to be particularly successful in the hedging of path-dependent financial derivatives. Below, we give a precise definition according to the formulation in *Effective statistical learning methods for actuaries III* [10].

Definition 1.6. *Let $p \in \mathbb{N}$ denote the number of input lags and q the number of hidden neurons. An LSTM model of order (p, q) takes the form:*

$$x_{t+1} = f_{\Omega}(x_t, y_{t-1}, c_t) + w_t,$$

where w_t is a zero-mean stochastic process, and f_{Ω} encodes the memory cell dynamics and output gates based on past observations.

LSTM networks are structured to capture long-range dependencies by containing a memory cell that only partially forgets and inputs information via gated pathways. These gates are modulated by sigmoid activation function. This disposition is especially useful in financial domain where hedging strategies need to dynamically adjust to the time-dependent behavior of market factors. LSTM trained with the past can learn to optimally respond to path-dependent exposures, leading to better facilitation of financial risk management and a more robust portfolio.

1.3.3 Neural networks as universal approximation

A foundational result in the theory of neural networks is the universal approximation theorem. The next theorem (Hornik (19910 [11]) ensures that, under mild assumptions, neural networks can approximate any continuous function on a compact domain to arbitrary precision. We denote by $\mathscr{N}^{\sigma}_{\infty,d_0,d_1}$ the set of neural networks mapping from $\mathbb{R}^{d_0} \to \mathbb{R}^{d_1}$ with activation function σ.

Theorem 1.7. *Suppose σ is bounded and non-constant. The following statements hold:*

- *For any finite measure μ on $(\mathbb{R}^{d_0}, \mathscr{B}(\mathbb{R}^{d_0}))$ and $1 \leq p < \infty$, the set $\mathscr{N}^{\sigma}_{\infty,d_0,d_1}$ is dense in $L^p(\mathbb{R}^{d_0}, \mu)$.*
- *If in addition $\sigma \in C(\mathbb{R})$, then $\mathscr{N}^{\sigma}_{\infty,d_0,d_1}$ is dense in $C(\mathbb{R}^{d_0})$ for the topology of uniform convergence on compact sets.*

This property justifies the use of neural networks to approximate optimal hedging strategies, which are typically complex and non-linear. Moreover, in financial settings, relevant state variables often lie on lower-dimensional manifolds, which enhances the effectiveness of learning algorithms. It is sufficient to approximate the hedging strategies on this low dimensional subset.

Since each component of the $\mathbb{R}^{d_1}$-valued neural network is $\mathbb{R}$-valued, it is straightforward to generalize this property to $\mathscr{N}_{\infty,d_0,d_1}$ with $d_1 > 1$. We define the following notation $\mathscr{N}\mathscr{N}_{\infty,d_0,d_1} := \mathscr{N}\mathscr{N}^{\sigma}_{\infty,d_0,d_1}$ with σ a fixed activation function. Let

$$\{\mathscr{N}\mathscr{N}_{M,d_0,d_1}\}_{M\in\mathbb{N}}$$

be a sequence of subsets of $\mathscr{N}_{\infty,d_0,d_1}$ with these properties:

- $\mathscr{N}\mathscr{N}_{M,d_0,d_1} \subset \mathscr{N}\mathscr{N}_{M+1,d_0,d_1}$ for all $M \in \mathbb{N}$
- $\bigcup_{M\in\mathbb{N}} \mathscr{N}\mathscr{N}_{M,d_0,d_1} = \mathscr{N}\mathscr{N}_{\infty,d_0,d_1}$
- for any $M \in \mathbb{N}$, one has $\mathscr{N}\mathscr{N}_{M,d_0,d_1} = \{F^{\theta} : \theta \in \Theta_{M,d_0,d_1}\}$ with $\Theta_{M,d_0,d_1} \subset \mathbb{R}^q$ for some $q \in \mathbb{N}$ (depending on M).

Motivated by the universal approximation results stated above, we now consider neural network hedging strategies. We therefore assume the activation function is bounded and non-constant, in accordance with the conditions of Theorem 1.7.

In what follows, we exploit this expressive power to train neural networks that approximate hedging strategies under realistic market conditions. Having established the theoretical tools, we now turn to the modeling of financial markets and classical hedging methods in the next chapter.

Chapter 2

Black-Scholes Model

Let us consider a probability space $(\Sigma, \mathscr{F}, \mathbb{P})$ equipped with a filtration $F = (\mathscr{F}_t)_{\{0 \leq t \leq T\}}$, where $T < \infty$. In this chapter, we introduce the Black–Scholes model as the theoretical benchmark for hedging in financial markets. We assume a complete, arbitrage-free market with constant risk-free rate r, where all investors share the same information and trading opportunities. These assumptions imply the existence of a unique risk-neutral probability measure $\mathbb{Q}$, equivalent to the real-world measure $\mathbb{P}$. The objectives of this chapter are twofold: (i) to present Black–Scholes as the foundation of classical hedging theory, and (ii) to highlight its practical limitations, which motivate the use of deep hedging in subsequent chapters.

2.1 Hedging in Black-Scholes

The Black-Scholes model was developed by integrating the delta-hedging concept. Its major contribution lies in demonstrating that, under idealized assumptions such as continuous rebalancing and absence of market frictions, the risk associated with issuing a derivative can be entirely eliminated by dynamically replicating its payoff. Let us give more details in the next section.

2.1.1 Theoretical Model

In the early 1970s, Fischer Black, Myron Scholes, and Robert Merton achieved a major breakthrough in the pricing of European stock options (J. C a. B. Hull Sankarshan(2016) [9]). They developed the Black–Scholes model, which became the theoretical foundation for the pricing and hedging of derivatives. Black-Scholes is based on stochastic processes to model changes in the underlying asset's price. More precisely, it assumes that the price of the asset follows a geometric Brownian movement, described by a stochastic differential equation (SDE).

Let S_t denote the price of a risky asset at time $t \geq 0$. Under the risk-neutral measure $\mathbb{Q}$, the price process follows the stochastic differential equation:

$$\frac{dS_t}{S_t} = r\,dt + \sigma\,dW_t^{\mathbb{Q}},$$

where r is the constant risk-free rate, $\sigma > 0$ the constant volatility, and $W_t^{\mathbb{Q}}$ a standard Brownian motion under $\mathbb{Q}$.

This equation describes how the asset's price evolves with its expected return and volatility, with a random term to capture the uncertainty of price fluctuations.

Now, suppose that f is the price of a call option or other derivative contingent on S. The variable f must be some function of S and t that must satisfy Black-Scholes Partial Differential Equation :

$$\frac{\partial f}{\partial t} + rS\frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} = rf.$$

The PDE is derived using no-arbitrage methodologies and dynamic hedging principles. It describes how the fair price of a derivative evolves in a frictionless, complete market.

The closed-form expressions for European call and put options are as follows. Define:

$$d_1 = \frac{\ln\left(\frac{S_t}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}, \quad d_2 = d_1 - \sigma\sqrt{T-t},$$

and let $\Phi(\cdot)$ denote the standard normal cumulative distribution function. Then the prices of European call and put options are:

$$C = S_t\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2),$$

$$P = Ke^{-r(T-t)}\Phi(-d_2) - S_t\Phi(-d_1),$$

where K is the strike price and S_T the stock price at maturity T.

The Black-Scholes model is revolutionary, but it suffers from a serious deficiency when applied in the "real" financial markets :

- **Constant volatility assumption:** The model assumes a constant volatility, which is not true in practice. Volatility can be time-varying and depends on both the strike and maturity. A popular class of stochastic volatility models such as the Heston model was proposed to address this limitation.

- **No price jumps:** Black-Scholes does not include the nature of sudden, unpredictable events in the market that cause prices to abruptly jump (Jump Risk). Changes in asset prices resulting from economic announcements, financial crashes, or bankruptcies can happen out of the blue. Such effects call for models like Merton's jump-diffusion model to better capture them.

- **Rational market assumption:** The model is based on the theory of rational markets, largely ignoring behavioral biases and irrational investors, which can lead to price discrepancies. The assumption of fully rational investors can be questioned, as market behavior is often influenced by psychological biases.

- **Continuous hedging impracticality:** Continuous-time hedging is inapplicable in real markets due to transaction and limited liquidity factors. In practice, hedging must be done discretely and at a cost.

Recent models use deep learning to dynamically adapt to market conditions, without making overly restrictive assumptions about volatility or yield distribution. Before introducing them we first present the traditional method of Delta Hedging.

2.1.2 Delta Hedging

Delta hedging relies on the Black–Scholes model and assumes continuous rebalancing to neutralize risk. It measures the sensitivity of an option's price to changes in the underlying asset price via the *Delta* (Δ), defined as:

$$\Delta = \frac{\partial C}{\partial S},$$

where C is the option price and S the underlying asset price. For example, a delta of 0.7 for a call option means that a small change in the stock price will change the option price by approximately 70% of that movement [9].

In our numerical experiments, the delta of a European call option is computed using the explicit Black–Scholes formula:

$$\Delta_{\text{call}}(t, S_t) = \Phi(d_1), \quad d_1 = \frac{\ln(S_t/K) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \tag{2.1}$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function.

For a European put option, the delta is given by a similar explicit formula:

$$\Delta_{\text{put}}(t, S_t) = 1 - \Phi(d_1).$$

At each rebalancing date t_k, $\Delta_{\text{call}}(t_k, S_{t_k})$ is updated using this formula and the portfolio is adjusted accordingly. In a frictionless market, this perfectly replicates the option payoff. In reality, continuous hedging is not possible due to transaction costs, discrete trading, and market frictions. Delta also assumes small price changes, ignoring jump risk. These limitations motivate alternative approaches, such as *deep hedging*, which can adapt to realistic market conditions.

2.2 Delta Hedging vs Deep Hedging

We compare the theoretical Black–Scholes delta hedging strategy to a neural network–based approach in a discretized setting over 30 trading days ($T = 30/365$, $dt = 1/365$) for a European call option with $S_0 = K = 1$, $r = 0$, and $\sigma = 0.2$.

The network $f_\theta : \mathbb{R}^4 \to \mathbb{R}$ learns an additive correction to the Black–Scholes delta:

$$\hat{\Delta}_{t_k} = \Delta^{BS}_{t_k} + \tanh\Big(f_\theta(x_{t_k})\Big),$$

with input vector:

$$x_{t_k} = \Big(\log(S_{t_k}/K),\ T - t_k,\ \sigma,\ \hat{\Delta}_{t_{k-1}}\Big),$$

where $\hat{\Delta}_{t_{k-1}}$ is the previous hedge ratio and $\Delta^{BS}_{t_k}$ the Black–Scholes delta is defined in (2.1). The parameters θ are trained to minimize the expected value of a chosen risk measure applied to the P&L distribution:

$$\theta^* = \arg\min_\theta\ \mathbb{E}\Big[\ell\Big(\text{P\&L}_\theta\Big)\Big],$$

where ℓ can be:

- *Entropic loss*: $\ell_{\text{ent}}(p) = -\mathbb{E}[-e^{-p}]$,

- *Conditional Value-at-Risk (CVaR)* at level α:

$$\ell_{\mathrm{CVaR}_\alpha}(p) = -\frac{1}{\alpha}\,\mathbb{E}\Big[p \cdot \mathbf{1}_{\{p<q_\alpha\}}\Big]\,.$$

The architecture is a fully-connected feed-forward network with 5 layers: $N_0 = 4$ (input), four hidden layers with 32 neurons each (ReLU activations), and an output layer ($N_5 = 1$) with a tanh activation to bound hedge adjustments. This balance ensures expressive power while limiting overfitting.

The resulting profit and loss P&L is the difference between the final value of the hedging portfolio and the option payoff, including trading gains and transaction costs. Performance is evaluated in two settings: **without transaction costs** and **with transaction costs**.

2.2.1 P&L without Transaction Costs

When no transaction costs are incurred, the portfolio is rebalanced at each trading date to match the theoretical Black-Scholes delta without any penalty. The final portfolio value at maturity T is given by:

$$V_T = \varDelta_T S_T + B_T,$$

where:

- $\varDelta_T$ is the final position in the underlying asset,
- B_T is the cash account at time T,
- S_T is the underlying price at maturity.

The resulting profit and loss (P&L) from hedging is then:

$$\mathrm{P\&L}_T = V_T - \mathrm{Payoff}$$

where the payoff typically corresponds to $(S_T - K)_+$ for a European call option with strike K.

In an idealized market perfectly matching the Black-Scholes assumptions, the hedging error would theoretically converge to zero as trading becomes continuous.

2.2.2 P&L with Transaction Costs

In real markets, each rebalancing operation incurs transaction costs, typically proportional to the traded notional. Let $c > 0$ denote the proportional transaction cost rate (e.g., $c = 0.001$ corresponds to a 0.1% fee per transaction).

The cost of rebalancing at trading date t_k is modeled by:

$$\mathrm{Cost}_k = cS_{t_k}|\varDelta_k - \varDelta_{k-1}|,$$

where:

- $\varDelta_k$ is the new position in the underlying asset after rebalancing at time t_k,
- $\varDelta_{k-1}$ is the previous position at time t_{k-1},

- S_{t_k} is the underlying asset price at time t_k,
- $|\cdot|$ denotes the absolute value (as both buying and selling incur costs).

The total transaction costs incurred over the option's lifetime are given by:

$$\sum_{k=0}^{N-1} \mathrm{Cost}_k = c \sum_{k=0}^{N-1} S_{t_k} |\Delta_k - \Delta_{k-1}|.$$

Consequently, the final profit and loss (P&L) accounting for transaction costs is expressed as:

$$\mathrm{P\&L}_T = V_T - \mathrm{Payoff} - c \sum_{k=0}^{N-1} S_{t_k} |\Delta_k - \Delta_{k-1}|,$$

where the last term represents the cumulative impact of transaction fees resulting from re-balancing the portfolio.

This difference between the frictionless and costly world is crucial for a correct interpretation of the numerical experiments on delta-hedging vs. deep hedging performance, which we present in the following section.

2.3 Numerical Analysis

Before training our model, we compare its profit and loss (P&L) distribution to the one of the Black-Scholes strategy. Figure 2.1 shows the histograms obtained from 50,000 simulated trajectories. This distribution (Figure 2.1) is not optimal since the feed-forward network is not yet trained.

The classical Black-Scholes strategy (in blue) shows a distribution centered around a stable loss of about -0.025, which is expected if transaction costs are not taken into account. In contrast, the network strategy (in red) shows greater variability and lower efficiency at this stage because the model has not yet learned how to dynamically adjust the coverage ratio. The mean is slightly closer to zero, indicating both improvement potential and instability typical of an untrained model. This step highlights the opportunities for progress that can be achieved through the deep learning approach compared to the traditional hedging model.

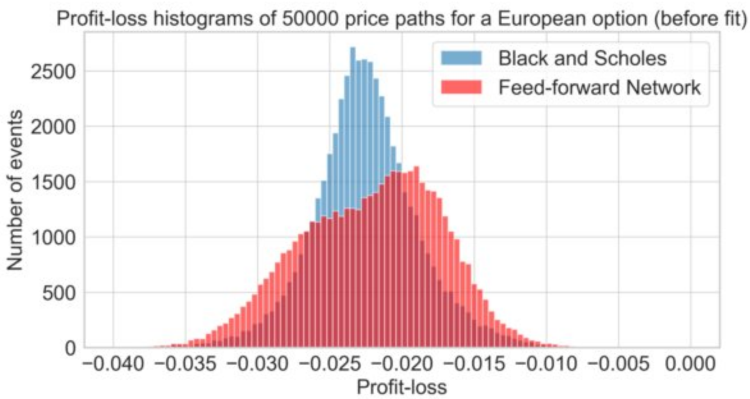

Figure 2.1: Black-Scholes Vs Untrained Feed forward network

This preliminary result highlights the potential improvement that can be achieved after training, which we now address using two distinct loss functions: entropic loss and Conditional Value at Risk (CVaR).

2.3.1 Entropy Loss

Entropic loss is based on the exponential utility function, which more heavily penalizes significant negative outcomes, reflecting risk-averse preferences. It is defined as:

$$L_{\text{ent}}(\text{P\&L}_\theta) = -\mathbb{E}[u(\text{P\&L}_\theta)], \quad \text{where} \quad u(x) = -\exp(-x).$$

Here, $\text{P\&L}_\theta$ is computed from the discrete-time hedging strategy over a grid $0 = t_0 < \cdots < t_N = T$:

$$\text{P\&L}_\theta = \sum_{k=0}^{N-1} \hat{\Delta}_{t_k} \left(S_{t_{k+1}} - S_{t_k}\right) - c \sum_{k=0}^{N-1} \left|\hat{\Delta}_{t_k} - \hat{\Delta}_{t_{k-1}}\right| S_{t_k} - \text{Payoff}(S_T),$$

where c is the proportional transaction cost rate and $\hat{\Delta}_{t_k}$ is the hedge ratio given by:

$$\hat{\Delta}_{t_k} = \Delta_{t_k}^{BS} + \tanh\Big(f_\theta(x_{t_k})\Big).$$

The network output $\tanh(f_\theta(x_{t_k}))$ provides an additive adjustment to the Black–Scholes delta $\Delta_{t_k}^{BS}$, directly influencing each term in the P&L calculus. Entropy loss training looks θ that maximize the expected exponential utility of the P&L, thereby improving the strategy's robustness to adverse scenarios.

To assess the effectiveness of our feed-forward neural network model, we compare its performance to that of the classical Black-Scholes strategy under two scenarios: without transaction costs (Figure 2.2) and with transaction costs set at 0.001 (Figure 2.3). The performance is evaluated using the P&L distribution over 50,000 simulated price trajectories.

We chose to evaluate the statistical value of the two models in Table 2.1 for the following values of α: $\alpha \in \{0.5, 0.75, 0.95, 0.99\}$.

Table 2.1: Comparison between the theoretical Black-Scholes and the Feed-forward model with the entropy loss function (without transaction cost)

Statistics	**Black-Scholes**	**Feed-Forward**
Mean	-0.02253	-0.02211
Standard Deviation	0.00366	0.00431
VaR(50%)	-0.02262	-0.02205
VaR(75%)	-0.02466	-0.02468
VaR(95%)	-0.02838	-0.02859
VaR(99%)	-0.03189	-0.03179
VaR(99.5%)	-0.03327	-0.03322

Without transaction costs, the Black-Scholes model remains a benchmark for hedging performance. As shown in Figure 2.2, both methods yield narrow and centered P&L

distributions, confirming the theoretical expectation that, under ideal conditions, Black-Scholes provides an almost perfect replication strategy. The feed-forward network shows a slightly improved average P&L (from -0.02253 to -0.02211), indicating a modest reduction in systematic losses, although this comes with a slightly higher standard deviation. Overall, in frictionless markets, neural networks do not provide significant gains over the analytical model.

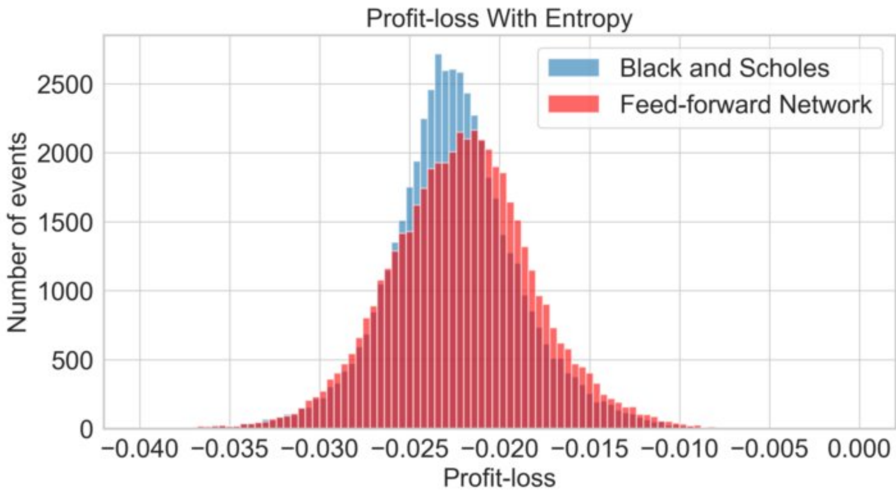

Figure 2.2: Without transaction cost.

When a transaction cost of 0.001 is introduced (Figure 2.3), the contrast becomes more relevant. The Black-Scholes strategy, which does not account for market frictions, suffers from a pronounced systematic loss, as reflected by the leftward shift in its P&L distribution. By contrast, the neural network model adapts its hedging decisions to the presence of transaction costs, dynamically reducing rebalancing frequency and avoiding unnecessary adjustments. Overall, these experiments illustrate the ability of deep hedging methods to outperform classical delta hedging in incomplete markets, particularly when frictions are present.

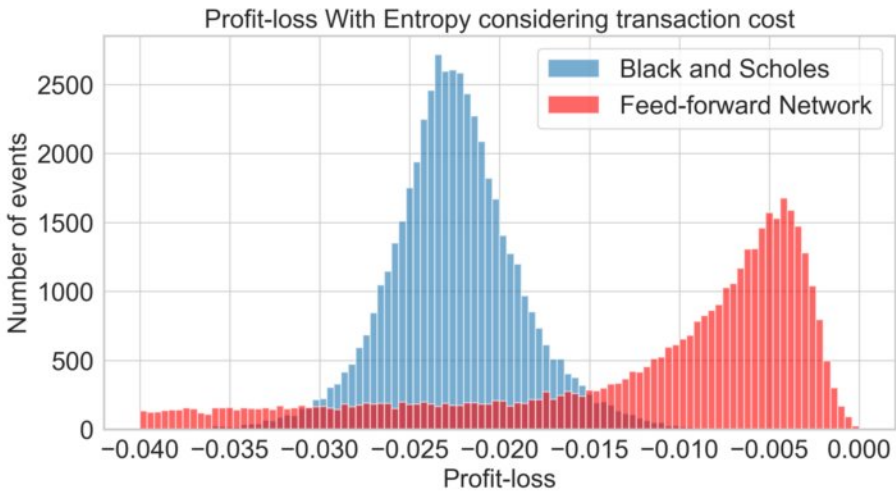

Figure 2.3: Transaction cost of 0.001.

To conclude,

- In a frictionless market, the Black-Scholes model is still efficient and near optimal, and the neural model only gives minor improvements.
- In the presence of market frictions, the feed-forward network significantly dominates the Black-Scholes strategy by adjusting the hedge ratios to reduce the cumulative costs.

Having evaluated the performance of the feed-forward network under entropic loss, we now turn our attention to another loss function that explicitly targets tail risk: the Expected Shortfall (ES) also known as the Conditional Value at Risk (CVaR). Unlike entropy, which aims to reduce the overall dispersion of P&L, the Expected Shortfall explicitly focuses on the most unfavorable losses, that is, the left tail of the distribution. This makes it a risk management choice criterion, particularly suitable in contexts where tolerance for extreme losses is limited, such as for regulated financial institutions.

2.3.2 Conditional Value at Risk

The Conditional Value-at-Risk (CVaR) measures the average loss in the worst α fraction of outcomes, thereby capturing both the frequency and severity of extreme events. It is often preferred to the Value-at-Risk (VaR), which only sets a loss threshold without accounting for tail magnitude.

For a random variable X and confidence level $\alpha \in (0, 1)$, the CVaR is:

$$\mathrm{CVaR}_\alpha(X)) = \mathbb{E}[X \mid X \leq q_\alpha(X)],$$

where $q_\alpha(X)$ is the α-quantile of X.

In our framework, X corresponds to $\mathrm{P\&L}_\theta$, the profit-and-loss of the hedging strategy with parameters θ. The optimization problem becomes:

$$\theta^* = \arg\min_\theta \; -\mathrm{CVaR}_\alpha(\mathrm{P\&L}_\theta) = \arg\min_\theta \; -\frac{1}{\alpha}\,\mathbb{E}\Big[\mathrm{P\&L}_\theta \; \mathbf{1}_{\{\mathrm{P\&L}_\theta < q_\alpha(\mathrm{P\&L}_\theta)\}}\Big] \,.$$

The negative sign indicates that maximizing CVaR i.e., improving the worst α fraction of P&L outcomes is equivalent to minimizing this loss function during training.

Table 2.2: Comparison between the theoretical Black-Scholes and the Feed-forward model with CVaR_{50} loss function (without transaction cost)

Statistics	**Black-Scholes**	**Feed-Forward**
Mean	-0.02253	-0.02210
Standard Deviation	0.00366	0.00370
VaR(50%)	-0.02262	-0.02207
VaR(75%)	-0.02466	-0.02422
VaR(95%)	-0.02838	-0.02817
VaR(99%)	-0.03189	-0.03195
VaR(99.5%)	-0.03327	-0.03350

The performance of these strategies is compared in Table 2.2. The Feed-Forward network achieves a slightly higher average P&L (-0.02210 vs. -0.02253) with similar variability (standard deviation 0.00370 vs. 0.00366). At moderate confidence levels (50% and 75%), it shows an advantage in Value-at-Risk (VaR), while at extreme levels (99% and 99.5%) the difference with Black–Scholes remains negligible.

Figure 2.4 confirms these results: for low α, the two distributions are close; as α increases, the neural network's distribution exhibits slight skewness, reflecting its adaptive rebalancing.

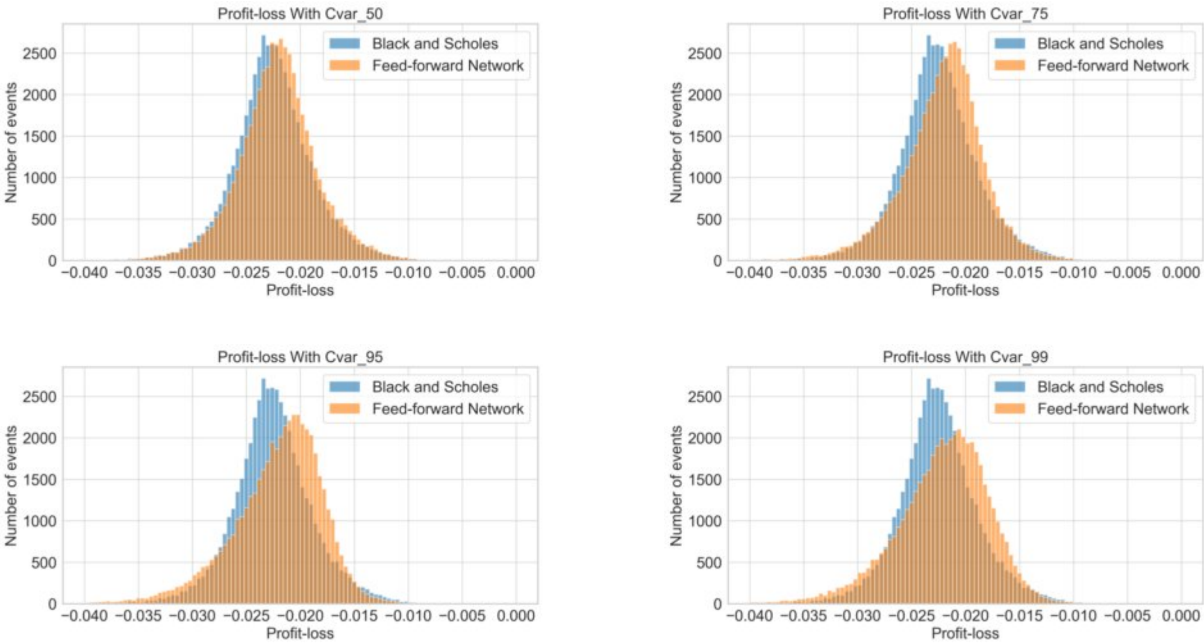

Figure 2.4: Without transaction cost.

Under transaction costs of 0.001 (Figure 2.5), the differences are amplified. While the Black–Scholes strategy remains rigid and suffers systematic losses from frequent rebalancing, the neural network reduces exposure to adverse tail events, especially at CVaR_{95} and CVaR_{99}, concentrating P&L around small losses or modest gains. This tail truncation results from fewer portfolio adjustments, lowering transaction costs while staying responsive to market volatility.

These results show that deep hedging's flexibility allows it to align with specific risk management objectives. In particular, optimizing directly on CVaR enables explicit tail-risk control, an appealing feature for institutions facing strict solvency and capital constraints.

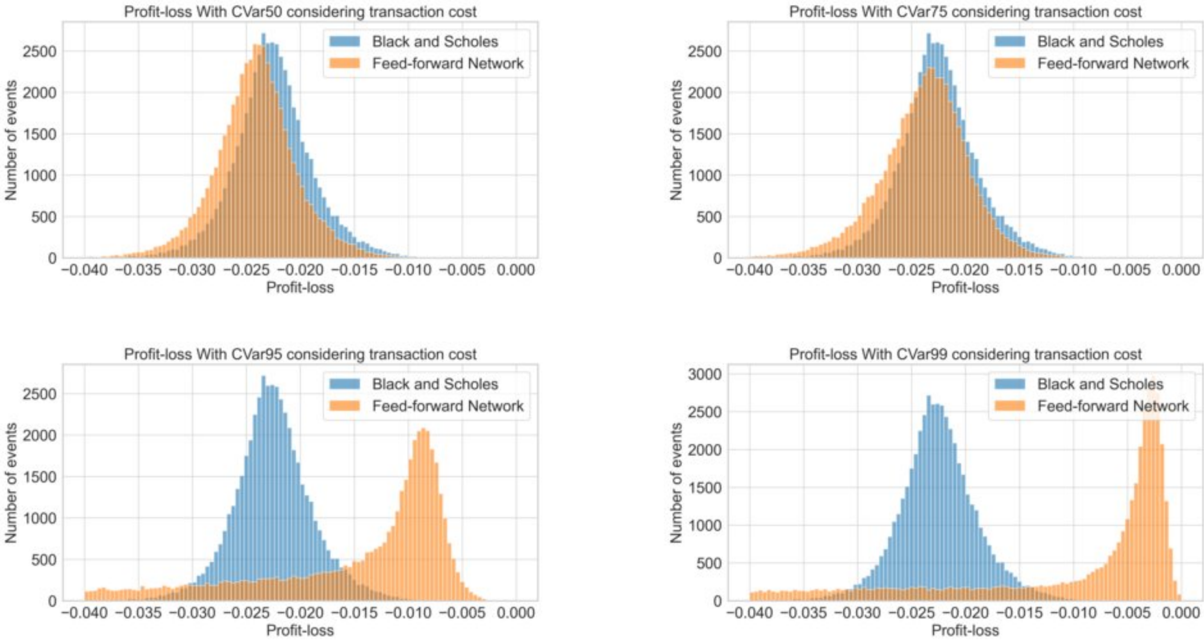

Figure 2.5: Transaction cost of 0.001.

We demonstrate that a neural network model can reproduce the behavior of the Black-Scholes framework in frictionless markets, but outperforms it when transaction costs are introduced. The rigidity of delta-hedging, which requires frequent rebalancing, leads to large cumulative costs and systematic losses. By contrast, the neural network dynamically adjusts its hedging frequency and magnitude, reducing unnecessary trades and thereby improving both mean performance and tail-risk measures.
This highlights a practical advantage of deep hedging: while it cannot outperform classical theory in idealized settings, it becomes superior in realistic environments where frictions, costs, and constraints matter precisely the context faced by insurers.

A key insight from this analysis is the influence of the chosen loss function on the resulting hedging strategy:

- **Entropy** encourages stable average performance with narrow profit and loss distributions,
- **CVaR** focuses on minimizing extreme losses to enhance robustness in adverse scenarios.

In teaching itself to adapt its hedging strategy dynamically, the neural network balances between this two-schools dilemma, demonstrating once more how artificial intelligence might be the next level of risk management for today.

Theoretically, delta-hedging strategies within the BS framework are ideal in idealized circumstances; however, the use of real insurance products, especially equity-linked life insurance contracts, usually generates complications. Such products generate both financial (equity and interest rate risk) and insurance (mortality and longevity) risks, which are stochastic and typically dependent. In addition, these policies commonly include path-dependent guarantees, such as minimum death or maturity benefits.

The markets in which these products are hedged are also incomplete, marked by transaction costs, discrete trading opportunities, stochastic volatility, and potential price jumps. In such environments, classical delta-hedging becomes suboptimal and can leave insurers exposed to significant residual risk.

These considerations call for more advanced hedging techniques capable of:

- capturing multiple sources of risk,
- adapting dynamically to changing market conditions,
- and optimizing risk-adjusted performance under convex risk measures rather than relying purely on replication.

In the next chapter, we shift our focus to equity-linked life insurance contracts and explore how deep learning methods can enhance their pricing and hedging strategies in practice.

Chapter 3

Equity-linked life insurance contracts

Today, with the increasing complexity of life insurance products (especially equity-linked contracts with embedded guarantees), valuation and risk management have become a major issue. These contracts expose insurers to a wide range of risks, including market fluctuations, interest rate volatility, and biometric uncertainties such as mortality and longevity.

In this chapter, we study sophisticated hedging strategies for these contracts, with a particular emphasis on Guaranteed Minimum Death Benefits (GMDB). We propose a unified simulation framework that integrates stochastic mortality and interest rate dynamics, and compare two approaches: the classical *local risk minimization* (LRM) technique, and the state-of-the-art *deep hedging* approach implemented via Long Short-Term Memory (LSTM) neural networks. LSTM is selected in view of its capability to learn the temporal dependencies involved in dynamic hedging for long-term, path-dependent guarantees. We are interested in whether these models can lower hedging error, especially during adverse market moves.

3.1 General Context

Variable Annuities (VAs), popularly known as unit-linked life insurance contracts, are unique products that mix an investment in the capital market with insurance protection. These products give consumers the chance to benefit from financial markets while providing some degree of protection against unintended events for policyholders. These guarantees, referred to collectively as GMxB (*Guaranteed Minimum x Benefit*), provide downside protection where x indicates the type of insured event: death, survival, income, or withdrawals. The four most widely studied GMxB guarantees are the following (D. Hainaut (2024)) [5]):

GMDB – Guaranteed Minimum Death Benefit

The GMDB ensures that, in the event of death before maturity T, the beneficiary receives at least a predefined guaranteed amount K, even if the account value S_t has declined. The benefit payable at time of death $t \in [0, T]$ is:

$$\text{GMDB}(t) = \max(S_t, K).$$

The expected present value of the GMDB liability can be expressed as:

$$\text{EPV}_{\text{GMDB}} = \sum_{k=1}^{N} {}_{k-1|1}q_x \cdot \mathbb{E}\left[e^{-\int_0^{k/12} r_t\,dt} \cdot \max(K - S_{k/12}, 0)\right],$$

where ${}_{k-1|1}q_x$ is the probability that an individual aged x dies in the month $(k-1, k]$, r_t is the instantaneous risk-free rate, and $N = 12T$ is the total number of months until maturity.

GMAB – Guaranteed Minimum Accumulation Benefit

The GMAB guarantees a minimum benefit at the end of the contract term, provided if the policyholder survives. At maturity T, the insured receives:

$$\text{GMAB}(T) = \max(S_T, K).$$

From the insurer's perspective, this guarantee embeds a short European put option. It is particularly relevant for policyholders focused on capital preservation. Its expected present value is given by:

$$\text{EPV}_{\text{GMAB}} = {}_Tp_x \cdot \mathbb{E}\left[e^{-\int_0^T r_t dt} \cdot \max(K - S_T, 0)\right],$$

where ${}_Tp_x$ is the survival probability to time T.

GMIB – Guaranteed Minimum Income Benefit

The GMIB provides a guaranteed annuitization option at time T, allowing the policyholder to convert the greater of the account value S_T or a guaranteed capital K into an income stream. Assuming the annuity is a temporary life annuity-due over n years, the guaranteed income is:

$$c = \frac{\max(S_T, K)}{\ddot{a}_{x+T}^{(n)}},$$

where $\ddot{a}_{x+T}^{(n)}$ is the actuarial present value of a unit annuity payable for n years to an individual aged $x + T$. The expected present value of the guarantee becomes:

$$\text{EPV}_{\text{GMIB}} = {}_Tp_x \cdot \mathbb{E}\left[e^{-\int_0^T r_t dt} \cdot \max\left(\frac{K}{\ddot{a}_{x+T}^{(n)}} - S_T, 0\right)\right].$$

GMWB – Guaranteed Minimum Withdrawal Benefit

The GMWB allows the policyholder to make regular withdrawals, up to a guaranteed minimum amount, regardless of the fund's performance. If $c_{\min}$ is the annual guaranteed withdrawal and $\gamma_k F_k$ is the contractual withdrawal at time k, then:

$$c_k = \max(\gamma_k F_k, c_{\min}),$$

where F_k is the account value at time k. The present value of total GMWB payouts is:

$$\text{EPV}_{\text{GMWB}} = \sum_{k=1}^{T} {}_kp_x \cdot \mathbb{E}\left[e^{-\int_0^k r_t dt} \cdot \max(\gamma_k F_k - c_{\min}, 0)\right],$$

subject to the constraint that total withdrawals do not exceed the initial guarantee base.

These are guarantees that are valuable to policyholders who want protection against market declines or longevity risk. But for insurers, their existence complicates pricing and risk management. The payoff structure is akin to that of exotic options (e.g., lookback options, puts, swaptions), preventing realistic market-based closed-form pricing.

Moreover, financial markets are incomplete due to the presence of non-tradable risks such as stochastic mortality, stochastic interest rates, and transaction costs. In such contexts, traditional hedging methodologies relying on sensitivities (like delta, rho) cannot be optimal. To overcome these limitations, alternative approaches have been developed:

- **Local risk minimization:** aims to reduce the conditional variance of the hedging error at each rebalancing step;
- **Global risk minimization:** seeks to minimize the expected (semi-)variance of the terminal hedging error using quadratic or non-quadratic loss functions;
- **Deep hedging with LSTM networks:** employs recurrent neural architectures designed to capture temporal dependencies in sequential data, enabling the learning of near-optimal dynamic strategies in high-dimensional stochastic environments [6].

This chapter explores the pricing and hedging of GMxB products in a setting of incomplete markets, which are described by a complete modeling of stochastic market, interest rate, as well as mortality, and which also explicitly accounts for these interactions. We compare monolithic strategies (e.g., using local risk minimization) with sophisticated deep learning strategies of LSTM type based cast in relation with the reduction of hedging error within realistic market constraints.

3.2 Model Specifications

In this section, we introduce the key stochastic models governing the three main risk drivers market, mortality, and interest rates underlying the valuation and hedging of GMxB guarantees. These components are jointly simulated to reflect the economic and biometric uncertainty that insurers face, particularly in long-term and capital-intensive contracts.

3.2.1 Stock Price Dynamics – Black-Scholes Model

The price process of the underlying risky asset S_t is assumed to follow a geometric Brownian motion under the risk-neutral measure $\mathbb{Q}$, in line with the Black-Scholes framework:

$$dS_t = S_t \left(r_t \, dt + \sigma_s \, dW_t^S \right), \tag{3.1}$$

where:

- S_t is the asset price at time t,
- r_t is the instantaneous short rate,
- $\sigma_s > 0$ is the constant volatility of the asset,

- W_t^S is a standard Brownian motion under $\mathbb{Q}$.

This model assumes frictionless markets, continuous trading, and lognormally distributed asset prices.

3.2.2 Mortality Dynamics – Gompertz Makeham

The selection of an appropriate mortality law for the purpose of modeling mortality for long-term financial products, such as equity-linked life insurance contracts is an important issue. The Gompertz-Makeham law models the age-specific mortality rate as the sum of:

- A constant age-independent risk (**Makeham term**): external causes of death.
- An exponentially increasing age-dependent risk (**Gompertz term**): biological aging.

The instantaneous mortality rate (force of mortality) at age x is given by:

$$\mu(x) = \lambda + Ce^{Dx}$$

where:

- $\lambda \geq 0$: constant (Makeham) component,
- $C > 0$: initial Gompertz coefficient,
- $D > 0$: exponential rate of aging.

The corresponding survival function is:

$$S(x) = \exp\left(-\lambda x - \frac{C}{D}\left(e^{Dx} - 1\right)\right)$$

The Gompertz-Makeham model is particularly well-suited for long-term insurance modeling due to its analytical tractability and biologically interpretable parameters. Its parameters can be calibrated from national mortality tables or fitted directly on cohort-specific data. However, this model tends to be less accurate at extreme ages and relies on assumptions such as population homogeneity and time-invariant parameters.

3.2.3 Interest Rate Dynamics – Hull-White Model

To capture the dynamics of interest rates in a realistic and analytically tractable way, we use the extended Vasicek model, known as the Hull–White model. The short rate r_t evolves according to the following stochastic differential equation:

$$dr_t = a(b - r_t)dt + \eta\, dW_t, \tag{3.2}$$

where:

- $a > 0$ is the mean-reversion speed,
- b is the long-term mean level,
- η is the volatility of the short rate,

- W_t is a Brownian motion under the risk-neutral measure $\mathbb{Q}$.

This model is capable of reproducing a wide range of term structures and allows for negative interest rates due to the Gaussian nature of r_t.

An important feature of the Hull-White model is that it can be exactly fitted to the initial zero-coupon yield curve. It takes log-normally distributed rates, and this gives the flexibility for interest rates to be negative, which corresponds to market phenomenology from the last years [6].

These three processes are co-simulated to produce plausible economic scenarios. They provide a basis for assessing the fair value of guarantees and testing the efficacy of hedging strategies in various market environments.

3.3 Risk Minimization in Incomplete Markets

Two main paradigms guide the construction of hedging strategies: local and global risk minimization. Local hedging, as implemented in the Föllmer– Schweizer framework [8], minimizes the conditional variance of the hedging error at each time step, focusing on short-term error control. In contrast, global hedging aims to minimize a risk criterion (e.g., variance, CVaR) evaluated on the entire trajectory until maturity. Deep hedging belongs to this second category, enabling the use of non-quadratic loss functions and richer information sets for long-term optimization.

3.3.1 Local Risk Minimization (LRM)

In incomplete markets, perfect replication of contingent claims is generally impossible. A well-known alternative is the *local risk minimization* (LRM) approach introduced by Föllmer and Schweizer (1991). Unlike *global* approaches, which minimize a risk measure applied to the *final* hedging error, LRM proceeds backwards in time and minimizes the conditional variance of the *incremental* hedging error at each step.

Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a probability space with a discrete-time filtration $\mathbb{F} = (\mathscr{F}_t)_{t=0}^T$. Let $S = (S_t)_{t=0}^T$ denote the price process of the risky asset and $H \in L^2(\Omega, \mathscr{F}_T, \mathbb{P})$ a square-integrable contingent claim. A predictable strategy is given by $\phi = (\phi_t)_{t=0}^{T-1}$, with portfolio value:

$$V_t := V_0 + \sum_{s=0}^{t-1} \phi_s (S_{s+1} - S_s).$$

In LRM, the optimization is *local* in time: at each date t, given the future portfolio V_{t+1}, we choose ϕ_t to minimize the conditional mean-squared *incremental cost*:

$$\min_{\phi_t} \mathbb{E}\left[(C_t(\phi_t))^2 \,\middle|\, \mathscr{F}_t \right],$$

where the cost process $C_t(\phi_t)$ is defined as:

$$C_t(\phi_t) = V_{t+1} - \phi_t S_{t+1} - B_{t+1},$$

and B_t denotes the risk-free asset position. This local minimization is mathematically equivalent to the Galtchouk–Kunita–Watanabe decomposition:

$$H = H_0 + \sum_{t=0}^{T-1} \phi_t (S_{t+1} - S_t) + L_T,$$

where L is a martingale orthogonal to S under $\mathbb{P}$. The optimal ϕ_t minimizes the local conditional variance of the hedging error increment and satisfies the *mean-self-financing* property:

$$\mathbb{E}[V_{t+1} - V_t \mid \mathscr{F}_t] = \phi_t \, \mathbb{E}[S_{t+1} - S_t \mid \mathscr{F}_t].$$

This backward, step-by-step optimization contrasts with global approaches, where ϕ is chosen to minimize a single loss function applied to $H - V_T$. In both cases, however, the portfolio dynamics V_t are defined in the same way; it is only the optimization criterion that differs.

While LRM provides a theoretically grounded and computationally efficient benchmark, it is inherently tied to a quadratic, stepwise criterion. For long-term guarantees or path-dependent products involving multiple sources of risk, global risk minimization can offer greater flexibility by directly targeting a terminal risk measure such as variance or CVaR.

3.3.2 Global Risk Minimization

Because certain sources of risk cannot be hedged, contingent claims are generally not perfectly replicable. The aim then becomes to minimize the residual risk of the hedging strategy at maturity rather than to eliminate it entirely.

Global risk minimization focuses on the terminal hedging error over the whole time horizon, rather than on step-by-step adjustments as in local risk minimization. The goal is to find a strategy ϕ that minimizes a chosen risk measure applied to the terminal portfolio value V_T^ϕ, which is computed using the same portfolio dynamics as in LRM.

Two main classes of criteria are commonly used:

(a) Global Quadratic Hedging. The most classical approach is to minimize the expected squared terminal hedging error:

$$\min_{\phi} \; \mathbb{E}\left[\left(H - V_T^\phi\right)^2\right],$$

where:

- H is the liability to be hedged (e.g., the payoff of an option or an insurance guarantee),
- V_T^ϕ is the terminal value of the portfolio obtained by following strategy ϕ.

This quadratic criterion gives the same weight to losses and gains in excess of the target value.

A modern implementation of this idea is provided by the *Quadratic Deep Hedging* framework developed by Bühler et al. (2019) [4]. In this approach, the strategy ϕ is parameterized by a neural network with parameters θ, and the objective becomes:

$$\mathscr{L}(\theta) = \mathbb{E}\left[\left(V_T^\theta - H\right)^2\right].$$

(b) Global Non-Quadratic Hedging. In many practical contexts — particularly for insurers — losses are more critical than gains. It then becomes relevant to use asymmetric or convex loss functions, such as:

- the semi-variance,
- the Conditional Value-at-Risk (CVaR), which measures the average loss beyond a given quantile,
- convex risk measures such as the entropic risk.

These alternative criteria allow the hedging strategy to reflect a risk-averse or regulatory perspective. For instance, CVaR is widely used in prudential frameworks such as Solvency II to explicitly capture tail risk.

Traditional approaches such as LRM rely on a *local* quadratic criterion, whereas *global* methods directly minimize the overall terminal risk according to the chosen measure, which may be quadratic or not. Deep learning–based frameworks fall naturally into this global setting, as they can optimize any differentiable risk measure while accommodating realistic market features.

3.4 Pricing Hedging Instruments

To hedge the financial and biometric risks resulting from equity-linked life insurance products, insurers may rely on structured derivatives such as mortality-linked securities and interest-rate derivatives. These instruments allow non-tradable risks arising from stochastic mortality and stochastic interest rates to be partially replicated. In our simulation framework (Section 3.7), we consider three main hedging instruments: the stock itself, zero-coupon bonds, and survival bonds.

3.4.1 Pricing Mortality Derivatives: Mortality Bonds

Mortality-linked securities are designed to transfer mortality or longevity risk from insurers to capital markets. The two main instruments in this category are mortality swaps and mortality bonds. Here, we focus on the latter.

A **mortality bond** is a financial instrument issued via a special purpose vehicle (SPV)and sold to institutional investors as a means of achieving diversification because certain mortality risks are uncorrelated with traditional financial risk. Investors collect income on a schedule throughout the life of the bond, but some or all of their principal may be lost if a mortality index goes above a certain level. This arrangement offers insurers a convenient tool for transferring risk, and investors access to a non-financial risk and perhaps a bit more yield.

In our framework, we adopt a simplified version of a mortality bond with maturity T, notional value N, and a trigger mechanism based on the mortality intensity at time T. Let $\mu(x+T)$ denote the mortality intensity at age $x+T$, and μ_{trigger} the threshold. The payoff is modeled as:

$$\text{Payoff}_{\text{mort bond}} = \mathbf{1}_{\{\mu(x+T)>\mu_{\text{trigger}}\}} \cdot N \cdot e^{-rT},$$

where r is the constant risk-free rate. This binary design captures the essence of catastrophe bonds; in practice, more realistic payoffs are often proportional to a mortality index [7].

Following Luciano et al. (2012) [7], mortality bonds can be priced under the risk-neutral measure $\mathbb{Q}$. Let I_T be a mortality index at time T, and K the trigger threshold. The time-t price is:

$$\Pi(t,T) = \mathbb{E}^{\mathbb{Q}}\left[e^{-r(T-t)} \cdot \mathbf{1}_{\{I_T<K\}} \mid \mathscr{F}_t\right].$$

If mortality follows a deterministic law such as Gompertz–Makeham, closed-form expressions are available. Under stochastic mortality models (e.g., Ornstein–Uhlenbeck dynamics), I_T is random and valuation requires Monte Carlo simulations. This distinction allows us to assess how the mortality modeling choice impacts valuation and hedging performance.

3.4.2 Pricing Interest Rate Derivatives: Bonds

Life insurance liabilities are highly sensitive to interest rate fluctuations due to their long-term duration. To mitigate this exposure, insurers use interest rate derivatives, primarily zero-coupon bonds and survival bonds.

Zero-Coupon Bonds

A zero-coupon bond (ZCB) pays one unit of currency at maturity T, with no interim coupons. Its price under $\mathbb{Q}$ is:

$$P(t,T) = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T r_s\,ds} \mid \mathscr{F}_t\right],$$

where r_t is the short rate.

Under the Hull–White model (Section 3.2), the ZCB price has a closed form:

$$P(t,T) = A(t,T) \cdot \exp\Big(-B(t,T)\,r_t\Big),$$

with

$$\begin{aligned} B(t,T) &= \frac{1-e^{-a(T-t)}}{a}, \\ A(t,T) &= \exp\left((B(t,T)-(T-t)) \cdot \frac{a^2 b - 0.5\eta^2}{a^2} - \frac{\eta^2}{4a^3}\left(1-e^{-a(T-t)}\right)^2\right), \end{aligned}$$

where a is the mean-reversion speed, b the long-term mean, and η the short-rate volatility.

Survival Zero-Coupon Bonds

To hedge both mortality and interest-rate risks simultaneously, we introduce *survival zero-coupon bonds* (SZCBs). These contracts pay one unit at maturity T provided the policyholder is still alive. Their price under the risk-neutral measure is:

$$\Pi(t,T) = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T r_s\,ds} \cdot \mathbf{1}_{\{\tau>T\}} \mid \mathscr{F}_t\right],$$

where r_t is the short rate and τ denotes the time of death.

Assuming independence between financial and biometric risks, this decomposes as:

$$\Pi(t,T) = P(t,T) \cdot S(t,T),$$

with $P(t,T)$ the risk-free zero-coupon bond price and $S(t,T)$ the conditional survival probability from t to T.

This product is particularly relevant for GMDB hedging: its payoff structure naturally combines financial discounting and survival probabilities, two key risk drivers of guaranteed benefits.

In practice, the closed-form Hull–White formula provides analytical benchmarks for $P(t,T)$, while Monte Carlo simulations of interest-rate and mortality dynamics generate realistic SZCB price distributions. This dual representation makes SZCBs a natural instrument for assessing the incremental value of biometric-linked assets in incomplete markets.

3.5 Simulation Framework

We now describe the simulation setup used to implement and evaluate hedging strategies for unit-linked contracts with GMDB guarantees. Two main approaches are compared:

- the benchmark *local risk minimization* (LRM) strategy, obtained via sequential regressions (Föllmer–Schweizer);
- *deep hedging* strategies, based on recurrent neural networks (LSTM), trained end-to-end on simulated trajectories.

All simulations are performed in discrete time with monthly rebalancing over a 20-year horizon, producing rebalancing steps $N = 240$. A total of $n = 10{,}000$ Monte Carlo paths are simulated for both training and testing.

3.5.1 Mortality Modeling and Applications

We first consider a Gompertz–Makeham law, calibrated to Belgian data between ages 65 and 90. The empirical mortality intensity is:

$$\mu(x) \approx -\ln(1 - q_x),$$

with q_x the observed death probability at age x.

Parameters are estimated by minimizing the squared distance between observed and theoretical intensities:

$$\min_{\lambda,c,d} \sum_{x=x_{\min}}^{x_{\max}} \left(\mu_{\text{obs}}(x) - \lambda - ce^{dx}\right)^2,$$

yielding (for ages 65–90):

$$\lambda = 0.002279, \quad c = 0.000001, \quad d = 0.138340.$$

Although deterministic, this model permits the construction of survival trajectories and survival zero-coupon bonds. Monthly paths are simulated by drawing deaths at each step from the conditional probability

$$q_t = 1 - \exp(-\mu(x+t)\Delta t)\,, \quad \Delta t = \tfrac{1}{12}.$$

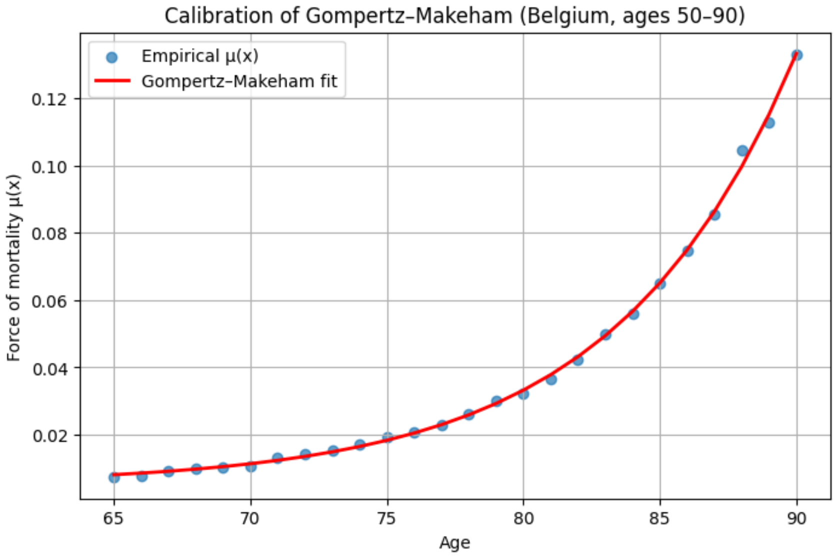

Figure 3.1: Calibration of the Gompertz–Makeham model to Belgian mortality data (ages 65–90).

Mortality bond pricing under parameter uncertainty. To illustrate mortality-linked assets, we also consider a stylized zero-coupon mortality bond paying $N = 100$ at $T = 20$ if the mortality intensity at age $x+T$ falls below a trigger $\mu_{\text{trigger}} \in \{0.1, 0.2, 0.3\}$. Uncertainty is modeled by perturbing the slope parameter:

$$d_{\text{sim}} \sim \mathcal{N}(d, \sigma_d^2), \qquad \sigma_d = 0.005.$$

The terminal intensity is then

$$\mu(x+T) = \lambda + ce^{d_{\text{sim}}(x+T)}.$$

In our calibration, trigger breaches are negligible, so the mortality bond price coincides with the discounted risk-free payoff. This illustrates the limitation of purely deterministic laws for generating stochastic payouts, motivating stochastic extensions. We will explore an extension in Section 3.8

3.5.2 Interest-Rate Modeling and Applications

Interest-rate dynamics follow the one-factor Hull–White model:

$$dr_t = a(b - r_t)\,dt + \sigma_r\,dW_t, \quad a = 10\%,\ \sigma_r = 1\%.$$

Exact discretization is used for monthly simulation. Zero-coupon bonds are priced analytically via:

$$P(t,T) = A(t,T)\exp(-B(t,T)r_t),$$

with A and B in closed form.

In our framework, Hull–White bonds play a dual role:

- benchmark risk-free instruments,
- hedging tools against financial risk in GMDB liabilities.

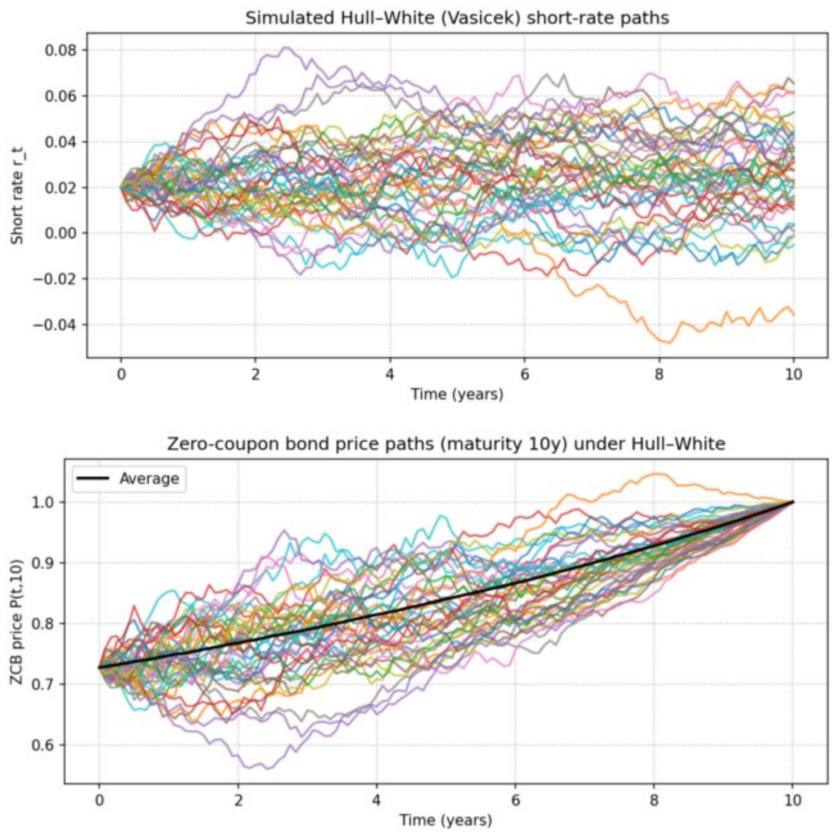

Figure 3.2: Simulated Hull–White short-rate paths (top) and corresponding zero-coupon bond prices (bottom).

3.5.3 Payoff Structure

We consider a pure GMDB contract with payoff at maturity:

$$H_T = \mathbf{1}_{\{\tau \leq T\}} \cdot \max(K - S_T, 0),$$

with $K = 100$. This structure mirrors a European put, activated only if death occurs before maturity. Deferring payments to T ensures consistency in valuation and hedging.

3.5.4 Evaluation Metrics

The hedging portfolio value V_T generates a terminal profit-and-loss:

$$\text{P\&L} = V_T - H_T.$$

Performance is assessed via:

- Root Mean Squared Error (RMSE):

$$\text{RMSE} = \sqrt{\mathbb{E}[(V_T - H_T)^2]},$$

  also used as the LSTM training loss;

- Value-at-Risk ($\text{VaR}_{95}\%$):

$$\text{VaR}_{95\%} = \inf\{x \mid \mathbb{P}(\text{P\&L} \leq x) \geq 0.05\};$$

- Conditional Value-at-Risk ($\text{CVaR}_{95}\%$):

$$\text{CVaR}_{95\%} = \mathbb{E}[\text{P\&L} \mid \text{P\&L} \leq \text{VaR}_{95\%}].$$

The RMSE reflects average hedge accuracy, while the VaR and CVaR capture tail risks, aligning with solvency-driven risk management.

3.6 Model Architectures

We compare two families of hedging strategies:

Local Risk Minimization (LRM). At each step, hedge ratios are computed by regressing discounted future liabilities onto asset returns (Föllmer–Schweizer decomposition). This method is linear, interpretable, and computationally efficient, but lacks memory and nonlinear interactions.

Recurrent Deep Hedging (LSTM). The LSTM network takes in sequences of market features (equity, interest rate, mortality survival probabilities) and outputs hedge ratios. Its gating mechanism with hidden and cell states allows path-dependency and long-term memory, which can be adapted for long-term insurance contracts. The training aims at minimizing the MSE of P&L and includes regularizations (L2 and turnover penalties) to maintain realistic strategies.

To conclude:

- **Dynamic LRM:** Linear, interpretable, and computationally efficient; adapts at each step but without temporal memory or nonlinear interactions.
- **Dynamic LSTM:** Nonlinear, recurrent; incorporates temporal memory to capture path-dependent effects in long-term contracts.

3.7 Results and Interpretation

This section presents the quantitative outcomes of the hedging strategies under deterministic mortality, focusing on the Gompertz–Makeham specification. Both the traditional Local Risk Minimization (LRM) and the deep hedging approach with LSTM networks are compared in terms of average hedging error and tail risk measures, complemented by graphical analyses of the profit-and-loss distributions.

Table 3.1: Performance of Local Risk Minimization (LRM) under Gompertz–Makeham mortality

Strategy	RMSE (€)	VaR 95% (€)	CVaR 95% (€)
Stock Only	25.83	-54.44	-64.37
Stock + Interest-Rate	25.59	-54.98	-64.87
Stock + Interest-Rate + Mortality	24.92	-55.07	-65.32

Table 3.2: Performance of Deep Hedging – LSTM under Gompertz–Makeham mortality

Strategy	RMSE (€)	VaR 95% (€)	CVaR 95% (€)
Stock Only	20.45	-34.64	-39.36
Stock + Interest-Rate	19.94	-30.63	-37.48
Stock + Interest-Rate + Mortality	19.88	-30.71	-37.71

The results in Tables 3.1 and 3.2, combined with the graphical evidence in Figures 3.3–3.7, reveal consistent patterns in the hedging efficiency under Gompertz–Makeham mortality.

First, the LSTM-based deep hedging model achieves a markedly lower RMSE across all strategies (values around €20) compared to the LRM benchmark (values around €25). This reduction in RMSE indicates that the neural network learns to control average hedging error more effectively, even when only equity is available.

Second, the risk of extreme losses ($\mathrm{VaR}_{95\%}$ and $\mathrm{CVaR}_{95\%}$) is substantially reduced under LSTM. For example, the worst-case tail loss ($\mathrm{CVaR}_{95\%}$) improves from roughly €-65 under LRM to about €-38–€-39 under LSTM, representing a tail risk reduction of nearly 40%. The barplots (Figures 3.4 and 3.5) confirm this improvement across strategies.

Third, the marginal role of additional hedging instruments is evident. For LRM, adding interest-rate or mortality-linked products only slightly reduces RMSE and tail losses, suggesting limited diversification benefits in this setting. For LSTM, the gains from extending the hedging set are also modest, but the grouped histograms (Figure 3.6) highlight that the neural network systematically produces narrower P&L distributions and lighter left tails compared to LRM, irrespective of the instrument set.

Finally, the distributional shape observed in the individual histograms (Figure 3.7) confirms that LSTM delivers more symmetric and concentrated P&L profiles, whereas LRM distributions remain more dispersed with heavier negative skewness. This contrast illustrates the ability of deep hedging to better adapt to non-linear risk factors embedded in the mortality-linked payoff.

From a practical point of view, the reduction of tail risk measures such as VaR and CVaR implies a lowering of the Solvency Capital Requirement (SCR) under Solvency II, which improves solvency ratios and frees up capital. Deep hedging also adapts to transaction costs by reducing unnecessary rebalancing, which limits systematic losses and improves profitability. By explicitly targeting downside risk, neural networks provide insurers stronger protection against extreme scenarios and greater financial resilience.

Overall, under Gompertz–Makeham mortality, deep hedging via LSTM consistently outperforms LRM by reducing both average hedging errors and exposure to extreme losses, while the addition of fixed-income and mortality-linked hedges provides only incremental improvements.

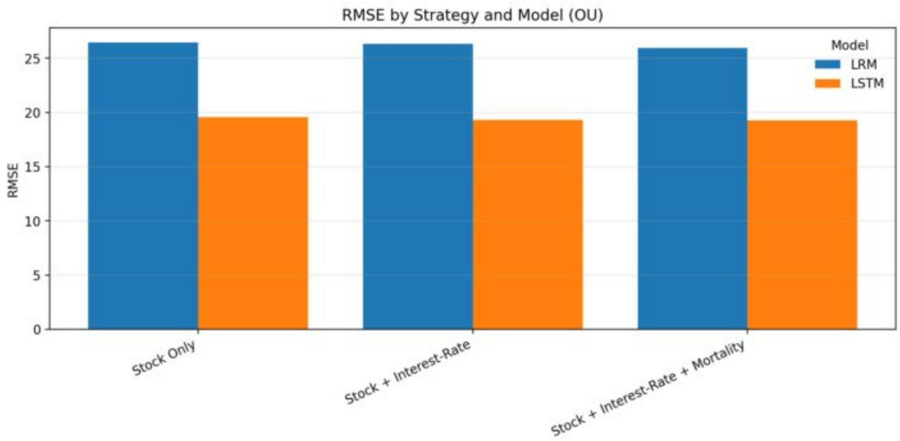

Figure 3.3: Gompertz–Makeham — RMSE par stratégie et modèle.

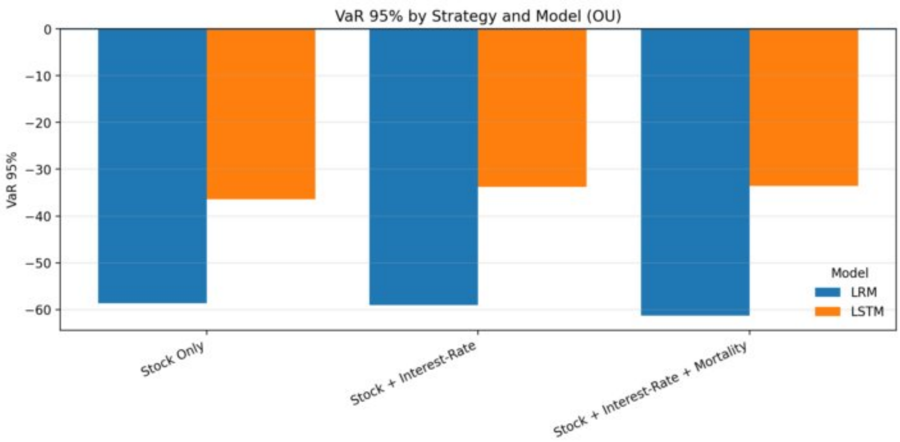

Figure 3.4: Gompertz–Makeham — $\mathrm{VaR}_{95\%}$ par stratégie et modèle.

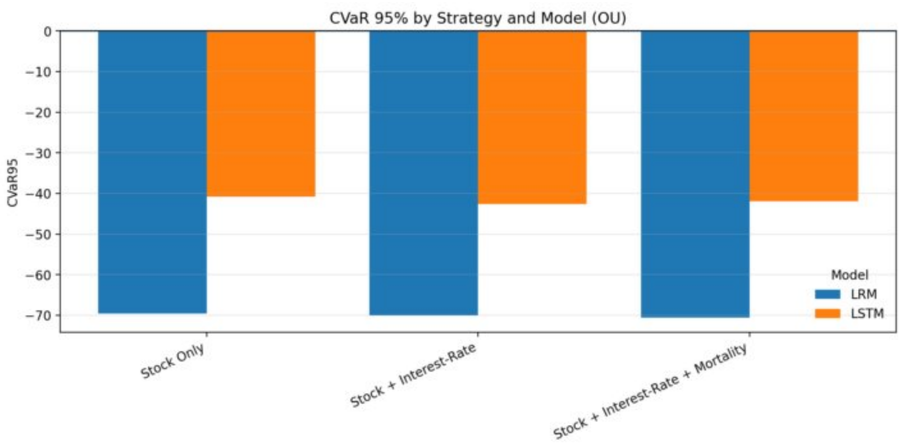

Figure 3.5: Gompertz–Makeham — $\mathrm{CVaR}_{95\%}$ par stratégie et modèle.

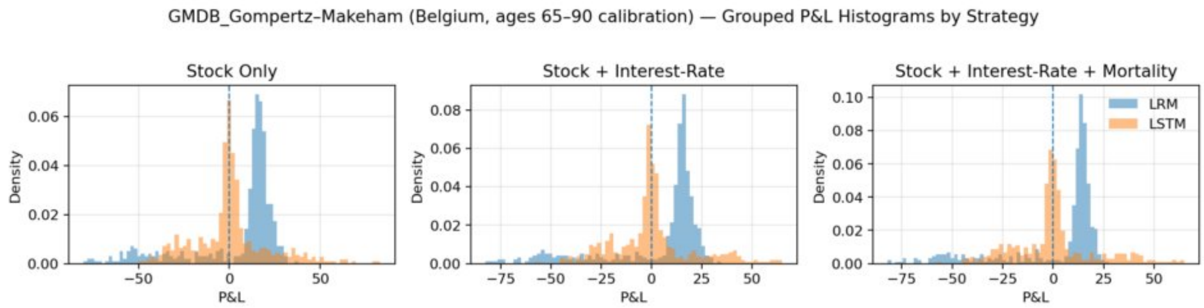

Figure 3.6: Gompertz–Makeham (Belgium, ages 65–90 calibration) — Histogrammes groupés des P&L par stratégie.

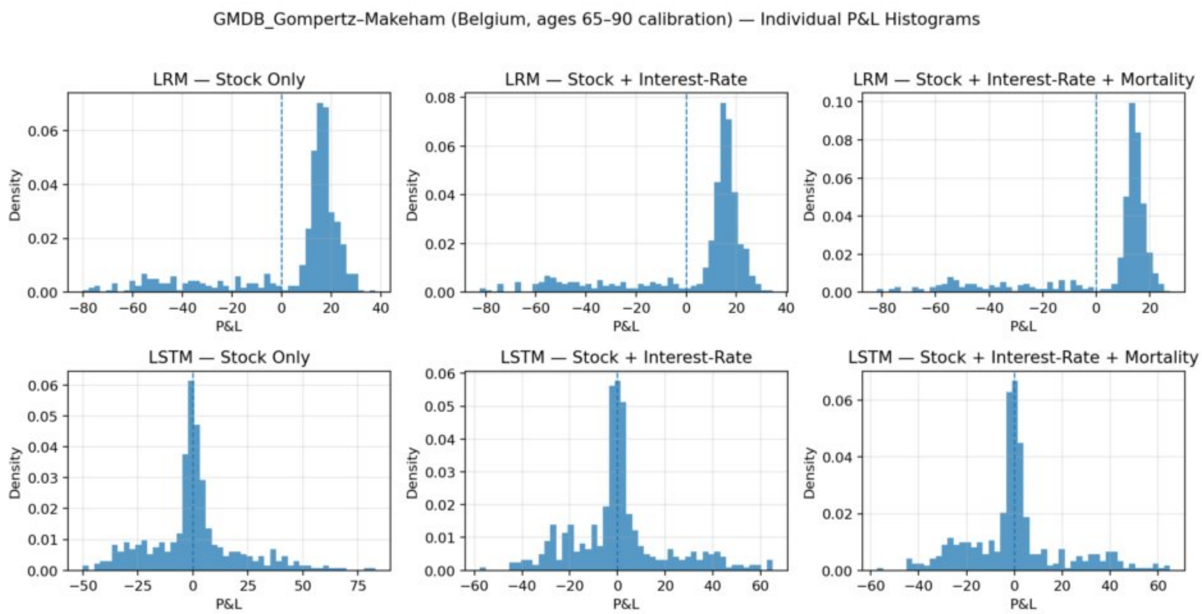

Figure 3.7: Gompertz–Makeham (Belgium, ages 65–90 calibration) — Histogrammes individuels des P&L (par modèle et stratégie).

These results have clear practical actuarial implications. First, the consistently decreasing RMSE obtained by the LSTM-based strategy shows that deep hedging significantly reduces average hedging errors in comparison to Local Risk Minimization. Notably, the enhancements in tail-risk measures (VaR, CVaR) indicate that insurers who pursue those strategies would experience fewer extreme losses. This leads to a lower Solvency Capital Requirements under Solvency II, with capital being driven largely by tail events, and more stable liability valuation under IFRS 17. In other words, deep hedging not only minimizes the portfolio risk, but it also enhances regulatory compliance and financial soundness in adverse cases.

3.8 Extension to Stochastic Mortality Modeling

In the previous sections, mortality dynamics were assumed to follow a deterministic Gompertz–Makeham law, i.e., survival probabilities are fully determined once the age of the policyholder is known. While this assumption ensures analytical tractability and facilitates calibration, it fails to capture key features of real-world longevity risk such as stochastic shocks, long-term mortality improvements, or adverse epidemics.

To overcome this limitation, we introduce stochastic mortality dynamics through an Ornstein–Uhlenbeck (OU) process for the mortality intensity μ_t, following the specification of [7]:

$$d\mu_t = \kappa(\theta - \mu_t)\, dt + \sigma\, dW_t, \tag{3.3}$$

where:

- $\kappa > 0$: speed of mean reversion,
- $\theta > 0$: long-term mortality level,
- $\sigma > 0$: volatility of mortality shocks,
- W_t: standard Brownian motion under the risk-neutral measure.

The OU specification generates Gaussian fluctuations around the mean level θ. Although non-negativity of μ_t is not strictly guaranteed, in practical calibrations the probability of negative intensities is negligible, and the model is widely adopted for its analytical tractability. In particular, closed-form expressions for conditional survival probabilities can be derived, making the OU framework attractive for pricing and hedging mortality-linked claims.

3.8.1 Integration into the Deep Hedging Framework

In our implementation, we rely on the calibration of [7], based on Italian males aged 65 in 2010, which yields:

$$\kappa = 10.94\%, \quad \sigma = 0.07\%, \quad \mu_0 = \theta = 0.885\%.$$

The model generates monthly mortality intensities through exact simulation of the OU process. At each rebalancing date, the survival status is updated by drawing from the conditional death probability:

$$q_t = 1 - e^{-\mu_t \Delta t}, \qquad \Delta t = \tfrac{1}{12}.$$

The GMDB payoff is then evaluated by combining the insured's survival status with the equity-linked guarantee, ensuring consistency with the contract specification.

This stochastic framework allows the hedging engine to incorporate both systematic longevity trends and random mortality shocks. In particular, it captures the dispersion of survival paths and provides a richer testing ground for mortality-linked hedging instruments such as survival zero-coupon bonds.

Remark on Comparability. The Gompertz–Makeham law used in earlier sections is calibrated to Belgian population mortality between ages 65 and 90, representing a deterministic baseline. By contrast, the OU model is implemented here with parameters from Luciano et al. (2012), reflecting a specific Italian cohort (males born in 1945, aged 65 in 2010). As a result, the absolute levels of mortality intensities and survival probabilities are not directly comparable across the two models. Our objective, however, is not to contrast Belgian versus Italian mortality, but rather to assess the methodological impact of adopting a deterministic versus stochastic mortality framework on pricing and hedging performance.

3.8.2 Hedging Performance under Mortality Uncertainty

Table 3.3: Performance of Local Risk Minimization (LRM) under OU mortality dynamics

Strategy	**RMSE (€)**	**VaR 95% (€)**	**CVaR 95% (€)**
Stock Only	13.41	-27.04	-52.88
Stock + Interest-Rate	13.37	-27.15	-52.97
Stock + Interest-Rate + Mortality	13.26	-28.15	-53.45

Table 3.4: Performance of Deep Hedging – LSTM under OU mortality dynamics

Strategy	**RMSE (€)**	**VaR 95% (€)**	**CVaR 95% (€)**
Stock Only	12.93	-24.73	-48.98
Stock + Interest-Rate	12.85	-25.97	-48.22
Stock + Interest-Rate + Mortality	12.88	-25.68	-48.44

Tables 3.3 and 3.4, supported by Figures 3.8–3.12, provide insights into hedging efficiency when mortality follows an Ornstein–Uhlenbeck (OU) process. Several observations emerge.

First, both models achieve lower RMSE levels compared to the Gompertz–Makeham setting, reflecting reduced variability in hedging errors when mortality shocks are mean-reverting. The LSTM approach again outperforms LRM, with RMSE values around €12.9 versus €13.3–13.4, albeit with a narrower performance gap than in the deterministic mortality case.

Second, the analysis of extreme losses confirms the advantage of LSTM. The $\mathrm{CVaR}_{95\%}$ improves from about €-53 under LRM to approximately €-48–€-49 under LSTM, representing a tail-risk reduction of roughly 10%. This relative improvement is smaller than in the Gompertz–Makeham framework, suggesting that the presence of stochastic mean-reversion in mortality dynamics naturally dampens tail risk, leaving less scope for neural networks to provide large incremental gains.

Third, the contribution of additional hedging instruments remains limited. In both LRM and LSTM, adding interest-rate and mortality-linked products yields only marginal improvements in RMSE and tail metrics. The grouped histograms (Figure 3.11) confirm that the dominant driver of hedging performance is the model type (LSTM vs LRM), rather than the composition of the hedging set.

Finally, the individual histograms (Figure 3.12) reveal that LSTM distributions remain more concentrated and exhibit lighter left tails compared to LRM. However, the contrast is less pronounced than under Gompertz–Makeham mortality, consistent with the fact that stochastic mortality with mean reversion already mitigates extreme deviations in payoff distributions.

In summary, under OU mortality, deep hedging with LSTM continues to outperform the LRM by diminishing both average errors and left- and right-tail risk; however, the relative gain is less drastic than under deterministic mortality, which emphasizes the stabilizing influence of the mean-reverting behavior in longevity dynamics.

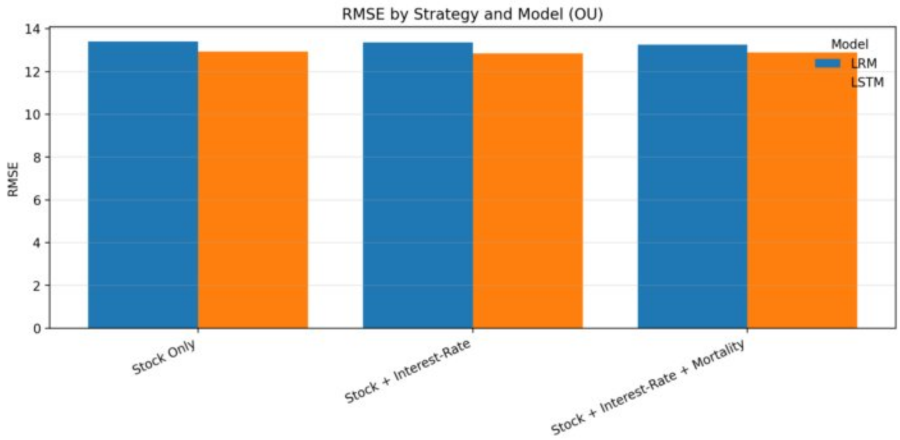

Figure 3.8: OU mortality — RMSE par stratégie et modèle.

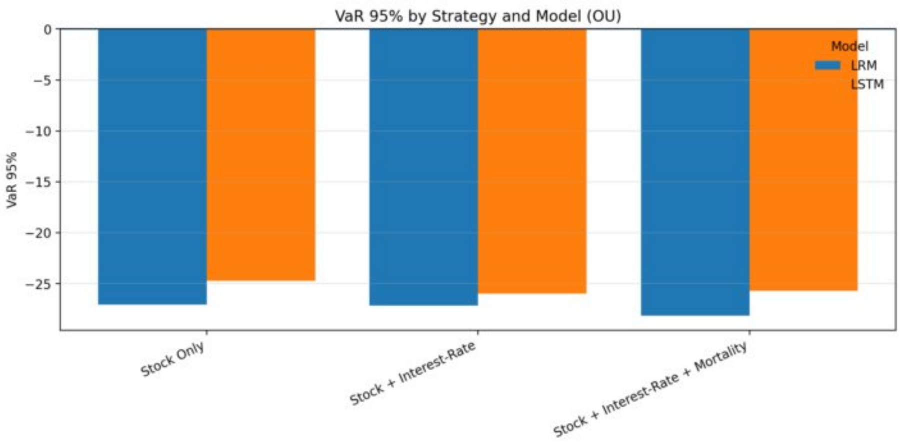

Figure 3.9: OU mortality — $\mathrm{VaR}_{95\%}$ par stratégie et modèle.

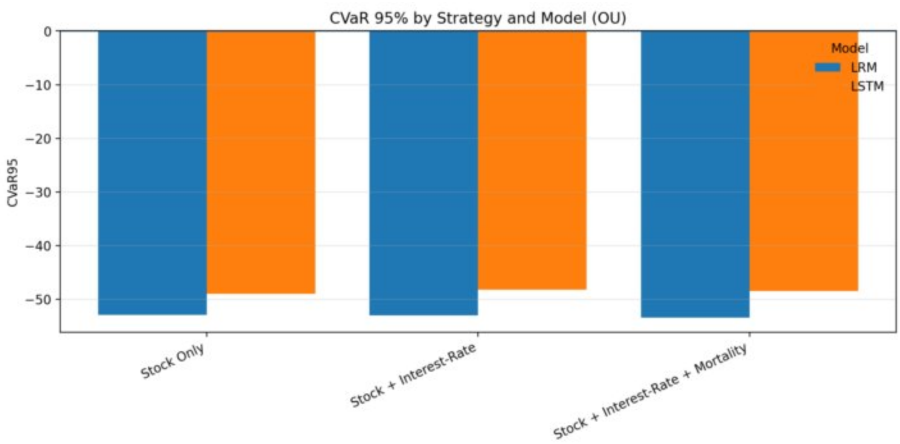

Figure 3.10: OU mortality — $\mathrm{CVaR}_{95\%}$ par stratégie et modèle.

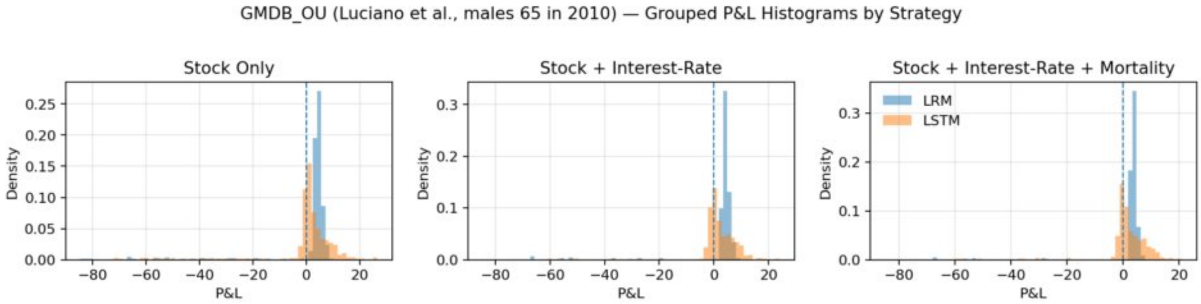

Figure 3.11: OU mortality (Luciano et al., males 65 in 2010) — Histogrammes groupés des P&L par stratégie.

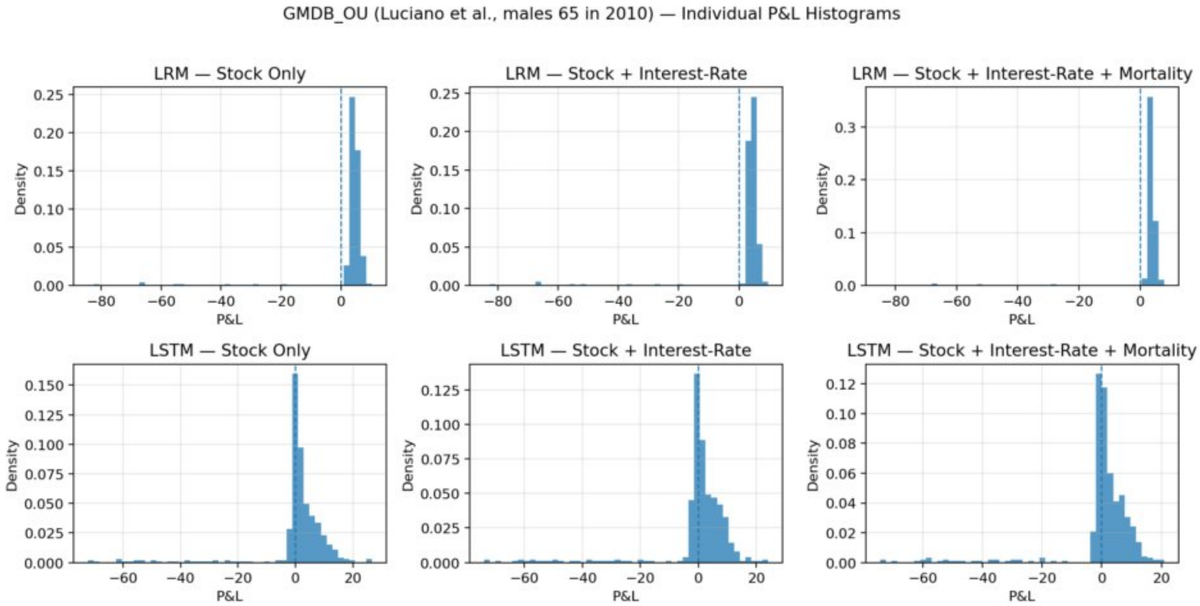

Figure 3.12: OU mortality (Luciano et al., males 65 in 2010) — Histogrammes individuels des P&L (par modèle et stratégie).

The inclusion of stochastic mortality dynamics in the form of an Ornstein–Uhlenbeck process extends the framework in order to incorporate demographic shocks, such as pandemics, unforeseen medical progress, or social changes. While the Gompertz–Makeham model provides analytical tractability, it cannot capture sudden or structural deviations in mortality patterns. Our results show that the LSTM strategy adapts more effectively to this additional uncertainty than traditional LRM, confirming the flexibility of neural networks in environments with both financial and biometric risks. This illustrates the potential of deep hedging to provide robust protection for insurers against longevity and mortality shocks that cannot be hedged using standard deterministic assumptions.

Limitations. Despite the promising results, several limitations must be acknowledged:

- **Computational cost:** The deep hedging framework relies on intensive Monte Carlo simulations and neural network training, which may be computationally expensive and sensitive to hyperparameter choices.

- **Interpretability:** While local risk minimization (LRM) yields explicit hedge ratios with a clear financial meaning, LSTM-based strategies are harder to interpret, which could hinder their adoption in regulated environments.
- **Market incompleteness:** The limited gains from adding mortality and interest-linked instruments illustrate the structural incompleteness of insurance markets, where some risks remain unhedgeable irrespective of the chosen approach.
- **Modeling simplifications:** Our analysis is based on simplified dynamics (e.g., one-factor Hull–White, Gompertz–Makeham or OU mortality). Extending the framework to richer stochastic models could affect both the performance and robustness of the hedging strategies.

This chapter highlighted the complexity of pricing and hedging equity-linked life insurance contracts with embedded guarantees such as GMDB. By jointly modeling market, mortality, and interest rate dynamics, we constructed a simulation framework able to capture the main drivers of long-term insurance risk.

We compared two competing methods: classical local risk minimization (LRM) and deep hedging with LSTM neural networks. Our findings indicate that although the LRM is a transparent and computationally efficient benchmark, it has difficulty in hedging errors when there are path-dependent and nonlinear risks. In comparison, the average hedging error (RMSE) and tail risk measures (VaR and CVaR) are lowered in a uniform manner by LSTM networks under all scenarios, even when only a sparse number of hedging instruments are on offer. Nevertheless, the marginal value of including assets with death or interest exposure is small in both frameworks, due to the incompleteness of insurance markets.

Overall, these findings illustrate the potential of deep learning to enhance hedging efficiency for unit-linked contracts, but also underline the challenges of model interpretability and computational cost. In practice, the choice between classical and neural hedging techniques will depend on the insurer's trade-off between transparency, efficiency, and operational feasibility. This sets the stage for further research into hybrid strategies, robustness analysis, and regulatory integration of advanced machine learning tools in actuarial risk management.

Conclusion

The primary objective of this thesis was to investigate and compare simple strategies for hedging Guaranteed Minimum Death Benefits (GMDB) embedded in equity-linked life insurance contracts with long-term guarantees. These products expose insurers simultaneously to equity risk, interest rate risk, and mortality risk. Managing these risks jointly is particularly challenging, since insurance markets are structurally incomplete and do not allow for perfect replication.

Specifically, we constructed a simulation-based framework integrating three primary forms of risk: Black-Scholes dynamics for stock prices, the Hull-White model for stochastic short rates (and their impacts on interest rates), and life mortality risk (which could be characterized either deterministically via the Gompertz law or stochastically using an Ornstein–Uhlenbeck process).

Within this setting, we studied two hedging methods: local risk minimization (LRM) as a classical benchmark, and a neural network with a Long Short-Term Memory (LSTM) network. Although the numbers speak for themselves, deep hedging techniques can retrieve good-quality hedging strategies from empirical data. It is worth noting that the LSTM hedging strategy significantly outperformed LRM methods in terms of such risk measures as RMSE, Value-at-Risk (VaR), and Conditional Value-at-Risk (CVaR) to reduce sporadically the forecasting error. This indicates that, because LSTMs are adept at modeling time series and learning when certain events happen over time, they are well-placed to manage long-dated insurance liabilities with their path-dependent risks. This benefit of using deep learning methods was even more apparent when considered under stochastic mortality, since classical models cannot account for the danger of unexpected longevity shocks.

Besides the actual results, this thesis also demonstrates the importance of the fusion of classical actuarial models with new machine learning methods. The hedging strategy was extended to the stochastic interest rate and mortality derivatives (namely zero-coupon bonds and mortality-linked securities) so that the different risk factors were more properly hedged in a realistic manner to cover the multiple drivers of risk. The flexibility of deep learning models was also particularly helpful in handling these multi-factor dynamics.

Nevertheless, this study is also limited. It depends entirely on simulated data and market assumptions that are far from fact. Real-world data, however, can be both scarce and noisy or incomplete, especially for mortality-linked financial instruments. Deep learning models are also super flexible but can be hard to interpret and overfit if not well regularized. It can be computationally expensive to train neural networks, particularly when using high-dimensional or path-dependent instruments.

Despite these limitations, the work presented here offers several lessons. First, it suggests that deep hedging strategies can outperform classical ones in markets where risks are not hedgeable by simple means. Second, it shows that the introduction of additional hedging instruments such as mortality bonds and interest-rate derivatives can significantly reduce residual hedging error.

This work can be extended in several directions. First, incorporating richer financial dynamics, such as multi-factor interest rate models or stochastic volatility with jumps, would provide a more realistic description of market risk. Second, integrating dependency structures between risk factors for instance, equity and mortality risks through copulas or latent factor models could better capture the joint impact of economic shocks such as financial crises or pandemics. Third, extending the framework to other guarantees, such as Guaranteed Minimum Income Benefits (GMIB) or Guaranteed Minimum Withdrawal Benefits (GMWB), would enhance its applicability to products with richer path-dependencies and optionality.

Further research should also address practical implementation issues. Accounting explicitly for transaction costs, liquidity constraints, and regulatory capital requirements would improve the model's realism and relevance for insurers. Applying the framework to real-world data, such as bond index fund returns and actual mortality experience, would provide an empirical assessment of its effectiveness. Finally, standard valuation frameworks under Solvency II or IFRS 17 could be supplemented with deep hedging techniques, offering insurers powerful new tools to improve both capital efficiency and regulatory compliance. Another promising avenue is the exploration of explainability methods for neural networks, which could help bridge the gap between performance and interpretability, thereby facilitating the adoption of deep hedging strategies in regulated environments.

Finally, this thesis truly reveals that actuarial modeling and deep learning are not a zero-sum game; rather, their conjunction provides more effective tools in hedging intricate insurance risks. Much remains to be done, but the best is likely yet to come with a hopeful future of finely balancing between empirical and theoretical lines in the new generation of actuarial software.

Bibliography

[1] A. Carbonneau, *Deep hedging of long-term financial derivatives*, Insurance: Mathematics and Economics **99** (2021), 327–340.

[2] K. a. D. Barigou Łukasz, *Pricing equity-linked life insurance contracts with multiple risk factors by neural networks*, Journal of Computational and Applied Mathematics **404** (2022), 113922.

[3] D. Hainaut, *LDATS2310 : Deep learning for insurance and finance*, Université Catholique de Louvain, Louvain-La-Neuve, 2023.

[4] H. a. G. Buehler Lukas and Teichmann, *Deep hedging*, Quantitative Finance **19** (2019), no. 8, 1271–1291.

[5] D. Hainaut, *LACTU 2170 : Stochastic finance*, Université Catholique de Louvain, Louvain-La-Neuve, 2024.

[6] D. a. G. Doyle Chris, *Using neural networks to price and hedge variable annuity guarantees*, Risks **7** (2018), no. 1, 1.

[7] E. a. R. Luciano Luca and Vigna, *Delta–gamma hedging of mortality and interest rate risk*, Insurance: Mathematics and Economics **50** (2012), no. 3, 402–412.

[8] H. a. S. Föllmer Martin, *Hedging of contingent claims under incomplete information*, Applied Stochastic Analysis (1991), 389–414.

[9] J. C a. B. Hull Sankarshan, *Options, futures, and other derivatives*, Pearson Education India, 2016.

[10] M. a. H. Denuit Donatien and Trufin, *Effective statistical learning methods for actuaries III*, Springer, 2019.

[11] K. Hornik, *Approximation capabilities of multilayer feedforward networks*, Neural networks **4** (1991), no. 2, 251–257.

[12] *Tables de mortalité et espérance de vie*, 2025.

Appendix

Chapter 3

Table 3.5: Belgian Mortality rates from age 49-71 (2024)

Exact age (x)	**Mortality rate (T_x)**
49	0.001554
50	0.002006
51	0.001814
52	0.002207
53	0.002414
54	0.002216
55	0.003074
56	0.003076
57	0.003544
58	0.003108
59	0.004003
60	0.004552
61	0.004669
62	0.005391
63	0.005825
64	0.006385
65	0.007381
66	0.007890
67	0.009165
68	0.009838
69	0.010440
70	0.010833
71	0.013105

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
Faculté des sciences
Place des sciences, 2 bte L6.06.01, 1348 Louvain-la-Neuve, Belgique | **www.uclouvain.be/sc**