Pricing equity-linked life insurance contracts with multiple risk factors by neural networks

Karim Barigou, Lukasz Delong

Pricing equity-linked life insurance contracts with multiple risk factors by neural networks

Karim Barigou*[1] and Łukasz Delong†[2]

[1]ISFA, Univ Lyon 1, UCBL, LSAF EA2429, F-69007, Lyon, France
[2]SGH Warsaw School of Economics, Collegium of Economic Analysis, Institute of Econometrics, Niepodległości 162, Warsaw 02-554, Poland

Version: November 1, 2021

Abstract

This paper considers the pricing of equity-linked life insurance contracts with death and survival benefits in a general model with multiple stochastic risk factors: interest rate, equity, volatility, unsystematic and systematic mortality. We price the equity-linked contracts by assuming that the insurer hedges the risks to reduce the local variance of the net asset value process and requires a compensation for the non-hedgeable part of the liability in the form of an instantaneous standard deviation risk margin. The price can then be expressed as the solution of a system of non-linear partial differential equations. We reformulate the problem as a backward stochastic differential equation with jumps and solve it numerically by the use of efficient neural networks. Sensitivity analysis is performed with respect to initial parameters and an analysis of the accuracy of the approximation of the true price with our neural networks is provided.

Keywords: Equity-linked contracts; Neural networks; BSDEs with jumps; Stochastic mortality; Heston stochastic volatility; Hull-White stochastic interest rates.

*Karim.Barigou@univ-lyon1.fr (Corresponding author)
†lukasz.delong@sgh.waw.pl

1 Introduction

An equity-linked life insurance contract is a contract between a policyholder and an insurance company which presents insurance and investment features (Palmer (2006), Bauer et al. (2008)). Typically, the policyholder pays either a single premium or a stream of periodic premiums during an accumulation phase. In return, the insurer guarantees a stream of periodic payments starting either immediately or at a future date. For investors, equity-linked life insurance contracts allow to participate in the equity market and offer higher returns than fixed annuity contracts. Moreover, they offer protection against downside shocks in the financial market by the presence of diverse guarantees. In the context of increasing life expectancies and the population facing the need for sustainable income, the demand for guaranteed income after retirement offered by equity-linked contracts is important (Haefeli (2013)).

Pricing of equity-linked life insurance contracts is a classical theme in the actuarial literature. Typically, this involves the choice of a financial and actuarial model to capture the financial and insurance risks. Based on no-arbitrage arguments and using the standard assumption that the financial and insurance risks are independent, these two components are merged together by a product of risk-neutral expectations (see Fung et al. (2014), Da Fonseca and Ziveyi (2015), Ignatieva et al. (2016) among others). However, this approach does not address the fundamental issue of incompleteness arising from the insurance and the financial market. Since the risks cannot be perfectly hedged, the valuation is non-linear and we should proceed in two-steps by disentangling the hedgeable and non-hedgeable part of the liability (Pelsser and Stadje (2014), Dhaene et al. (2017) and Barigou et al. (2019)).

We follow the pricing approach from Delong et al. (2019*b*) where the authors derived a system of partial differential equations (PDEs) for the fair valuation of a stream of payments contingent on financial and insurance risk factors. In Delong et al. (2019*b*), three risk factors were considered and numerical methods for solving the PDEs were not discussed. In practice, we would like to apply a valuation operator which takes into account multiple correlated risk factors. In this paper we deal with six stochastic financial and insurance risk factors: interest rate (two factors), equity, volatility, unsystematic and systematic mortality. The interest rate is modelled with two-factor Hull-White model, the equity is modelled with Heston model where the volatility is modelled with Cox-Ingersoll-Ross model, the lifetimes of policyholders are modelled, conditionally on the force of mortality, with independent exponential distributions and the force of mortality is modelled with Feller process. We assume that equity-linked contracts include two types of minimum guarantees: death benefits in case of the deaths of the policyholder and survival benefits in case the policyholder is still alive at the contract termination. In such a framework, we can derive a system of PDEs for the fair valuation of a portfolio of equity-linked contracts. Solving such a system of

PDEs seems to be challenging due to a large number of state variables and a large number of PDEs which reflects the number of policies in the insurance portfolio. It is known that finite difference methods for PDEs in high dimensions are not feasible due to the curse of dimensionality. In our case, we face an additional difficulty if we valuate a large portfolio of equity-linked contracts, since we have to solve recursively a large number of PDEs, where the solution of one equation (for the number of policyholders equal to n) is used as the input to the next equation (for the number of policyholders equal to $n+1$).

In recent years, some papers have considered the pricing of equity-linked contracts by neural networks. Doyle and Groendyke (2019) explored the use of neural networks to price and hedge equity-linked contracts. They showed that neural networks offer an important computational gain compared to crude Monte-Carle methods, however they did not consider stochastic mortality neither stochastic volatility in their framework. In Hejazi and Jackson (2016), the authors proposed a neural network approach to price equity-linked contracts for large insurance portfolios. Their framework prices some representative contracts via Monte-Carlo and uses a spatial interpolation method to price the whole portfolio.

Another stream of literature investigated the use of neural networks to accelerate PDE solvers. Han et al. (2018), Weinan et al. (2017), Beck et al. (2019) take advantage of reinforcement learning to speed up solving high-dimensional partial differential equations by using a so-called *Deep BSDE* method. Moreover, Chan-Wai-Nam et al. (2019) proposed an improved version of the Deep BSDE algorithm using different network architectures and parameterizations.

We see three contributions of this paper. First, we consider the pricing of equity-linked contracts in a general incomplete financial-actuarial market with stochastic interest rate, equity, volatility and mortality. In the spirit of Delong et al. (2019*b*), we derive a system of multi-dimensional non-linear PDEs and the optimal hedging strategy for our fair valuation problem – this is our analytical contribution. Second, we present a numerical contribution: using the representation of the PDEs as a BSDE with jumps, we show how to solve numerically the pricing problem by neural networks. In particular, we show how the jump component in the BSDE, which reflects the number of PDEs in the system of equations and the number of policyholders in the portfolio, can be efficiently handled. Finally, we discuss the sensitivity of the price to model parameters and study the accuracy of our approach to standard pricing methods. We emphasize that the resolution of the system of non-linear PDEs by neural networks presented in this paper is general and is not limited to life insurance.

We would also like to point out that BSDEs with jumps, which we apply in this paper to represent the solution to our system of non-linear PDEs, are usually used to represent solutions to non-linear partial integro-differential equations with Lévy jump measures. Re-

cently, Castro (2021) considered a neural network to solve a PIDE in a similar fashion to ours, in the sense that an additional network is trained for the jump component. However, our system of non-linear PDEs and our jump components are different from the single PIDE in Castro (2021).

The paper is structured as follows. In Section 2, we present in detail the financial and insurance model. Section 3 derives the pricing PDEs, discusses their representation as a BSDE with jumps and the numerical implementation by neural networks. In Section 4, we present numerical results and Section 5 provides concluding remarks.

2 Financial and Insurance Model

Throughout the paper, we work on a probability space $(\Omega, \mathbb{F}, \mathbb{P})$ equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{0\leq t\leq T}$ and a finite time horizon $T < \infty$. We consider a general financial market consisting of a bank account, a bond and an equity. We use a two-factor interest rate model to describe the short-rate process and Heston model to describe the dynamics of the equity price and its volatility. The insurance risk is modeled by a jump process with stochastic force of mortality.

2.1 The financial market

We assume that the dynamics of the risk-free rate under an equivalent martingale measure $\mathbb{Q}$ is given by the G2++ two-additive-factor gaussian model:

$$\begin{aligned} r(t) &= \psi(t) + x(t) + y(t), \quad r(0) = r_0, \\ dx(t) &= -ax(t)dt + \sigma_x dW_1^{\mathbb{Q}}(t), \quad x(0) = 0, \\ dy(t) &= -by(t)dt + \sigma_y dW_2^{\mathbb{Q}}(t), \quad y(0) = 0, \end{aligned}$$

where ψ is a deterministic function such that the initial term structure predicted by the model is fitted to the market term structure. The G2++ model is equivalent to two-factor Hull-White model under appropriate re-parametrization (see e.g. Section 4.2.5 in Brigo and Mercurio (2007)). For hedging in incomplete markets, models need to be defined under the real-world measure $\mathbb{P}$. We follow the approach from Ait-Sahalia and Kimmel (2010) and introduce the dynamics

$$\begin{aligned} r(t) &= \psi(t) + x(t) + y(t), \quad r(0) = r_0, \\ dx(t) &= \big(\delta_x\sigma_x - ax(t)\big)dt + \sigma_x dW_1(t), \quad x(0) = 0, \\ dy(t) &= \big(\delta_y\sigma_y - by(t)\big)dt + \sigma_y dW_2(t), \quad y(0) = 0, \end{aligned}$$

where we use constant risk premiums δ_x, δ_y to switch from the martingale measure to the real-world measure.

The dynamics of the risk-free bank account is given by

$$dR(t) = R(t)r(t, x(t), y(t))dt, \quad R(0) = 1,$$

where we emphasize that the risk-free rate r depends on $(t, x(t), y(t))$. In our G2++ model, the dynamics of the bond price $P(t) = \mathbb{E}^{\mathbb{Q}}[e^{-\int_t^{T^*} r(s)ds}|\mathcal{F}_t]$, which matures at time $T^* \geq T$, is given by

$$dP(t) = P(t)\Big(\big(r(t, x(t), y(t)) + \zeta(t)\big)dt + A(t)dW_1(t) + B(t)dW_2(t)\Big),$$

for some given deterministic functions A, B (see e.g. Section 4.2.3 in Brigo and Mercurio (2007)):

$$A(t) = -\sigma_x \frac{1 - e^{-a(T^*-t)}}{a}, \quad B(t) = -\sigma_y \frac{1 - e^{-b(T^*-t)}}{b},$$

and the risk premium ζ:

$$\zeta(t) = A(t)\delta_x + B(t)\delta_y.$$

The risk premium ζ should be positive, hence the risk premiums δ_x, δ_y should be negative (as always, we assume that σ_x, σ_y are positive).

We assume that the equity follows the Heston stochastic volatility model (Heston (1993)):

$$\begin{aligned} dS(t) &= S(t)\Big(\big(r(t, x(t), y(t)) + \gamma\sqrt{v(t)}\big)dt + \sqrt{v(t)}dW_3(t)\Big), \quad S(0) = 1, \\ dv(t) &= k(\eta - v(t))dt + \sigma_v\sqrt{v(t)}dW_4(t), \quad v(0) = v_0, \end{aligned}$$

where we also assume a constant risk premium γ for the equity. We assume that $v_0 > 0$, $k > 0$ and $2k\eta \geq \sigma_v^2$ so that v is strictly positive.

The policyholder pays a single premium for an equity-linked life insurance policy. The initial premium is invested in the policyholder's account value and the policyholder's funds are invested in the bond and the equity. The policyholder's account dynamics is then given by

$$\begin{aligned} dF(t) &= uF(t)\Big(\big(r(t, x(t), y(t)) + \zeta(t)\big)dt + A(t)dW_1(t) + B(t)dW_2(t)\Big) \\ &\quad + (1-u)F(t)\Big(\big(r(t, x(t), y(t)) + \gamma\sqrt{v(t)}\big)dt + \sqrt{v(t)}dW_3(t)\Big) - cF(t)dt, \\ F(0) &= F_0, \end{aligned}$$

where F_0 is the initial premium paid by the policyholder, u is a fixed percentage of the account invested in the bond, and c is a constant fee deducted by the insurer to cover the

financial guarantees and insurance protection embedded in the contract. The portfolio of equity-linked contracts, which we introduce in the next section, is homogeneous in the sense that all policyholders invest the same single premium, follow the same investment strategy and have the same guarantees and insurance protection, i.e. F_0, u and c are equal and fixed for all policyholders. We point out that the value of the fee c results from pricing the contract. In this paper we assume that the fee c has already been set by a pricing actuary and we calculate the price of the contracts given c . However, we can also find c such that the price of the portfolio of equity-linked contracts is zero and we can call such a fee - a fair fee.

Finally, all Brownian motions involved in the interest rate and the equity models are correlated with $dW_i(t)dW_j(t) = \rho_{ij}dt$ for $i, j = 1, \ldots, 4$.

2.2 The insurance portfolio

The portfolio consists of n policies. All policyholders have the same age and are entitled to two types of benefits: a death benefit D paid at the moment that the insured dies (provided he dies in $[0, T]$) and a survival benefit S paid at terminal time T if the insured survives that time. The benefits are contingent on the values of the policyholder's account F, hence we will use the notations $D(t, F(t))$ and $S(F(T))$ for death and survival benefits, respectively.

We assume that the lifetimes of the policyholders $(\tau_k)_{k=1,\ldots,n}$ at policy inception are, conditional on the mortality intensity, independent and identically distributed with mortality intensity λ, i.e. we assume that

$$\mathbb{P}\big(\tau_k > t|(\lambda(s), 0 \leq s \leq t)\big) = e^{-\int_0^t \lambda(s)ds}.$$

The spot mortality intensity at calendar time t of a head aged x at time 0 is denoted by $\lambda(t)$. Following Luciano et al. (2008) and Luciano et al. (2012), we assume that $\lambda(t)$ follows the Feller process without mean reversion:

$$d\lambda(t) = q\lambda(t)dt + \sigma_\lambda\sqrt{\lambda(t)}dW_5(t), \quad \lambda(0) = \lambda_0, \tag{2.1}$$

where $\lambda_0 > 0$ and $q > 0$. The mortality intensity can attain zero value but Luciano and Vigna (2008) showed that the probability of such an event is negligible for calibrated mortality data. The choice of this model is motivated by its parsimony-few parameters to be calibrated and its ability to fit cohort life tables due to the lack of mean reversion. Moreover, under the dynamics (2.1), the survival probability can be obtained in the closed-form:

$$\mathbb{P}\big(\tau_k > t\big) = \mathbb{E}\left[e^{-\int_0^t \lambda(s)ds}\right] = e^{\beta(t)\lambda(0)}, \tag{2.2}$$

with $\beta(t) = \frac{1-e^{bt}}{b_1+b_2e^{bt}}, b = -\sqrt{q^2+2\sigma_\lambda^2}, b_1 = \frac{b+q}{2}, b_2 = \frac{b-q}{2}$.

The lives of the policyholders are dependent since they are all affected by the path of $\lambda(t)$ (systematic mortality risk). However, conditionally to the sample path followed by $\lambda(t)$, the lives are independent (unsystematic mortality risk). Similar as in credit risk, we can compute the probability that two policyholders survive (see e.g. Section 9.5 in McNeil et al. (2005)). Using the conditional independence, we obtain:

$$\begin{aligned}\mathbb{P}\left(\tau_i > T, \tau_j > T\right) &= \mathbb{E}\left[\mathbb{P}\big(\tau_i > T, \tau_j > T | (\lambda(s), 0 \le s \le T)\big)\right] \\ &= \mathbb{E}\left[\mathbb{P}\big(\tau_i > T | (\lambda(s), 0 \le s \le T)\big) \times \mathbb{P}\big(\tau_j > T | (\lambda(s), 0 \le s \le T)\big)\right] \\ &= \mathbb{E}\left[e^{-2\int_0^T \lambda(s)ds}\right] = e^{\tilde{\beta}(T)\lambda(0)}, \quad i \neq j,\end{aligned}$$

where $\tilde{\beta}(T) = \frac{2\left(1-e^{\tilde{b}T}\right)}{\tilde{b}_1+\tilde{b}_2 e^{\tilde{b}T}}, \tilde{b} = -\sqrt{q^2+4\sigma_\lambda^2}, \tilde{b}_1 = \frac{\tilde{b}+q}{2}, \tilde{b}_2 = \frac{\tilde{b}-q}{2}$, which is an increasing function of σ_λ. Therefore, the higher the volatility of the mortality intensity, the higher the systematic risk that both policyholders survive.

We introduce the processes N and J which count the number of deaths and in-force policies in the insurance portfolio, respectively:

$$\begin{aligned}N(t) &= \sum_{k=1}^{n} \mathbf{1}\{\tau_k \le t\}, \quad N(0) = 0, \\ J(t) &= n - N(t), \quad J(0) = n.\end{aligned}$$

In addition, the compensated counting process:

$$\tilde{N}(t) = N(t) - \int_0^t (n - N(s-))\lambda(s)ds, \tag{2.3}$$

is a martingale, and it will be used in further derivations. Finally, in the remainder of the paper, we assume that the processes which drive the insurance risk (N, λ) are independent of the financial market.

The financial and insurance models considered in this paper are very popular in actuarial applications. Let us remark that different models can also be used and the methods presented below still hold.

3 Pricing equity-linked life insurance

Even though the insurer can invest in a bank account, a bond or an equity, the insurance liabilities are not perfectly hedgeable due to the presence of the stochastic volatility risk and the independent actuarial risks. Therefore, there is no unique pricing technique for insurance contracts and one needs to disentangle the hedgeable and non-hedgeable part of the liability. For instance, Bayraktar et al. (2009) used the instantaneous Sharpe ratio and Liang and Lu

(2017) considered the principle of equivalent utility. In this paper, we follow the approach of Delong et al. (2019*b*) and use the fair valuation operator with the instantaneous standard deviation risk margin. Our method for pricing can be described in the following steps:

1. We set up a hedging portfolio composed of a bond, an equity and a bank account. To find the optimal investments, we minimize the local variance (quadratic variation) of the net asset value process, which is defined as the excess value of the hedging portfolio over the fair price of the liability. In case of a complete market, the minimized local variance is zero. However, the incompleteness of the market leads to a residual risk, as measured by the remaining local variance.

2. The insurer asks a compensation for this residual non-hedgeable risk and we assume that it is priced via an instanteneous standard deviation risk margin.

3. The price can then be expressed as the solution to a system of non-linear PDEs. The system of PDEs describes the prices given the number of in-force policies and the number of equations is equal to the number of policies in the portfolio,

4. In order to solve these PDEs, we reformulate the PDEs as a forward-backward stochastic differential equation with jumps and approximate the price using deep neural networks.

3.1 Fair pricing with the instantaneous standard deviation risk margin

We assume that the insurer can invest in the bank account, the bond and the equity in order to hedge the financial guarantees and the benefits embedded in the contracts. Let θ_1 denote the amount of money invested in the bond, and θ_2 denote the amount of money invested in the equity. The remaining amount is invested in the bank account. We assume that θ_1 and θ_2 are Markov control strategies and are functions of $(t, x(t), y(t), F(t), v(t), \lambda(t), J(t))$. Let $V^\theta := (V^\theta(t), 0 \leq t \leq T)$ denote the self-financing hedging portfolio under a strategy $\theta = (\theta_1, \theta_2)$. The dynamics of the insurer's hedging portfolio is given by

$$\begin{aligned} dV^\theta(t) &= \theta_1(t)\Big(\big(r(t,x(t),y(t)) + \zeta(t)\big)dt + A(t)dW_1(t) + B(t)dW_2(t)\Big) \\ &\quad +\theta_2(t)\Big(\big(r(t,x(t),y(t)) + \gamma\sqrt{v(t)}\big)dt + \sqrt{v(t)}dW_3(t)\Big) \\ &\quad +(V^\theta(t) - \theta_1(t) - \theta_2(t))r(t,x(t),y(t))dt + J(t-)cF(t)dt - D(t,F(t))dN(t), \\ V^\theta(0) &= nF(0). \end{aligned}$$

Moreover, let us remark that the survival benefits $J(T)S(F(T))$ are paid at the terminal time T and therefore, do not affect the insurer's hedging portfolio.

In order to simplify the derivations, we introduce the 5-dimensional Itô process $Z(t) = (x(t), y(t), F(t), v(t), \lambda(t))$ which models the five risk factors:

$$dZ(t) = \mu(t, Z(t))dt + \sigma(t, Z(t))dW(t), \tag{3.1}$$

with

$$\mu(t, Z(t)) = \begin{pmatrix} \delta_x \sigma_x - a x(t) \\ \delta_y \sigma_y - b y(t) \\ \Big(r(t, x(t), y(t)) - c + u\zeta(t) + (1-u)\gamma\sqrt{v(t)} \Big) F(t) \\ k(\eta - v(t)) \\ q\lambda(t) \end{pmatrix}, \tag{3.2}$$

and

$$\sigma(t, Z(t)) = \begin{pmatrix} \sigma_x & 0 & \dots & \dots & 0 \\ 0 & \sigma_y & 0 & \dots & 0 \\ uF(t)A(t) & uF(t)B(t) & (1-u)F(t)\sqrt{v(t)} & 0 & 0 \\ 0 & \dots & 0 & \sigma_v\sqrt{v(t)} & 0 \\ 0 & \dots & \dots & 0 & \sigma_\lambda\sqrt{\lambda(t)} \end{pmatrix}. \tag{3.3}$$

In the expression (3.1), $W(t)$ is a 5-dimensional correlated Brownian motion with correlation matrix $\mathcal{Q}$, i.e. $dW_i(t)dW_j(t) = \rho_{ij}dt$ where ρ_{ij} is the (i,j)-component of $\mathcal{Q}$ (for $i, j = 1, \dots, 5$).

Let $\varphi := (\varphi(t), 0 \le t \le T)$ denote the continuous-time fair valuation operator representing the fair price of the liabilities. We assume that φ is a function of $(t, x(t), y(t), F(t), v(t), \lambda(t), J(t))$ and is denoted by $\varphi^{J(t)}(t, Z(t))$. We introduce the process $X(t) = \varphi^{J(t)}(t, Z(t)) - V^\theta(t)$ representing the difference between the price of the liabilities and the value of the insurer's hedging portfolio. Consequently, the process $NAV(s) = -X(s)$ determines the excess of the assets over the technical provision, and is called the net asset value. Let $\Theta(s, Z(s)) = (\theta_1(s)A(s), \theta_1(s)B(s), \theta_2(s)\sqrt{v(s)}, 0, 0)$ denote a 5-dimensional vector related to the hedging strategy. By application of the multidimensional Itô's lemma, we find the dynamics:

$$\begin{aligned}
dX(s) = &\Big\{ X(s) r(s, x(s), y(s)) + \varphi_t^{J(s-)}(s, Z(s)) + \nabla\varphi^{J(s-)}(s, Z(s)) \cdot \mu(s, Z(s)) - cJ(s-)F(s) \\
&- \varphi^{J(s-)}(s, Z(s)) r(s, x(s), y(s)) + \Big(\varphi^{J(s-)-1}(s, Z(s)) + D(s, F(s)) - \varphi^{J(s-)}(s, Z(s)) \Big) J(s-)\lambda(s) \\
&- \theta_1(s)\zeta(s) - \theta_2(s)\gamma\sqrt{v(s)} \Big\} ds \\
&+ \frac{1}{2} \operatorname{Tr}\Big(\sigma(s, Z(s)) \mathcal{Q} \sigma^\intercal(s, Z(s)) \operatorname{Hess}_z \varphi^{J(s-)}(s, Z(s)) \Big) ds \\
&+ \Big(\varphi^{J(s-)-1}(s, Z(s)) + D(s, F(s)) - \varphi^{J(s-)}(s, Z(s)) \Big) d\tilde{N}(s) \\
&+ \Big(\nabla\varphi^{J(s-)}(s, Z(s)) \sigma(s, Z(s)) - \Theta(s, Z(s)) \Big) dW(s).
\end{aligned}$$

Here, $\nabla\varphi$ and $\mathrm{Hess}_z\,\varphi$ denote the gradient and the Hessian of the function φ with respect to Z, Tr denotes the trace of a matrix.

We are interested in finding the valuation operator φ which satisfies the property:

$$\lim_{h\to 0}\left\{\frac{\mathbb{E}\Big[X(t+h)e^{-\int_t^{t+h}r(s)ds}-X(t)\Big|Z(t)=z,J(t)=k\Big]}{h}\right.$$
$$\left.+\frac{RM\Big[X(t+h)e^{-\int_t^{t+h}r(s)ds}-X(t)\Big|Z(t)=z,J(t)=k\Big]}{h}\right\}=0. \tag{3.4}$$

As the one-period risk margin, we consider the standard deviation:

$$\begin{aligned}&RM\Big[X(t+h)e^{-\int_t^{t+h}r(s)ds}-X(t)\Big|Z(t)=z,J(t)=k\Big]\\ &=\ \frac{1}{2}\alpha\sqrt{Var\Big[X(t+h)e^{-\int_t^{t+h}r(s)ds}-X(t)\Big|Z(t)=z,J(t)=k\Big]}\sqrt{h}.\end{aligned}$$

The parameter α can be interpreted as the insurer's risk aversion coefficient towards the non-hedgeable risks.

We can choose the hedging strategy θ so that the insurer does not charge a risk margin for the hedgeable financial risk. The optimal investments in the bond and the equity are then determined by minimizing the instantaneous (local) variance of the net asset value process. Therefore, we need to compute the quadratic variation $[X]_s$ and find the bivariate process $(\theta_1(s),\theta_2(s))$ which minimizes this quantity. The quadratic variation is given by

$$[X]_T=\int_0^T\underbrace{\Big(\nabla\varphi(s,Z(s))\sigma(s,Z(s))-\Theta(s,Z(s))\Big)\mathcal{Q}\Big(\nabla\varphi(s,Z(s))\sigma(s,Z(s))-\Theta(s,Z(s))\Big)^{\intercal}}_{g(s,\theta_1(s),\theta_2(s))}ds$$
$$+\Big(\varphi^{J(s-)-1}(s,Z(s))+D(s,F(s))-\varphi^{J(s-)}(s,Z(s))\Big)dN(s). \tag{3.5}$$

The optimal hedging strategy is then obtained by solving the first-order conditions:

$$\begin{aligned}\frac{\partial g(s,\theta_1,\theta_2)}{\partial\theta_1}&=-2\sum_{j=1}^{4}\Big(\rho_{1j}A(s)+\rho_{2j}B(s)\Big)\Big(\nabla\varphi(s,Z(s))\sigma(s,Z(s))-\Theta(s,Z(s))\Big)_j\\ &=0,\\ \frac{\partial g(s,\theta_1,\theta_2)}{\partial\theta_2}&=-2\sum_{j=1}^{4}\rho_{3j}\sqrt{v(s)}\Big(\nabla\varphi(s,Z(s))\sigma(s,Z(s))-\Theta(s,Z(s))\Big)_j\\ &=0.\end{aligned}$$

We remark that the index $j=5$ does not appear in the derivatives since the mortality intensity is not correlated with the financial risk factors. The first-order conditions yield a

system of two equations with two unknowns for which the solutions are given by

$$\begin{aligned}
\theta_1^*(t,Z(t)) &= u\varphi_f(t,Z(t))F(t) \\
&\quad + \frac{\varphi_x(t,Z(t))\sigma_x\rho^x(t) + \varphi_y(t,Z(t))\sigma_y\rho^y(t) + \varphi_v(t,Z(t))\sigma_v\sqrt{v(t)}\rho^v(t)}{A(t)\rho^x(t) + B(t)\rho^y(t)}, \\
\theta_2^*(t,Z(t)) &= (1-u)\varphi_f(t,Z(t))F(t) \\
&\quad + \frac{\varphi_x(t,Z(t))\sigma_x}{\sqrt{v(t)}}\left(\rho_{31} - (\rho_{31}A(t) + \rho_{32}B(t))\,\frac{\rho^x(t)}{A(t)\rho^x(t) + B(t)\rho^y(t)}\right) \\
&\quad + \frac{\varphi_y(t,Z(t))\sigma_y}{\sqrt{v(t)}}\left(\rho_{32} - (\rho_{31}A(t) + \rho_{32}B(t))\,\frac{\rho^y(t)}{A(t)\rho^x(t) + B(t)\rho^y(t)}\right) \\
&\quad + \varphi_v(t,Z(t))\sigma_v\left(\rho_{34} - (\rho_{31}A(t) + \rho_{32}B(t))\,\frac{\rho^v(t)}{A(t)\rho^x(t) + B(t)\rho^y(t)}\right),
\end{aligned}$$

with some deterministic functions, which take into account the correlation between the different risk factors:

$$\begin{aligned}
\rho^x(t) &= A(t) + \rho_{12}B(t) - \rho_{31}\rho_{13}A(t) - \rho_{31}\rho_{23}B(t), \\
\rho^y(t) &= B(t) + \rho_{12}A(t) - \rho_{32}\rho_{13}A(t) - \rho_{32}\rho_{23}B(t), \\
\rho^v(t) &= \rho_{14}A(t) + \rho_{24}B(t) - \rho_{34}\rho_{13}A(t) - \rho_{34}\rho_{23}B(t).
\end{aligned}$$

As it is usually the case, it turns out that the optimal hedging strategy appears as a delta-hedging strategy involving the sensitivities of the liabilities to the underlying financial risk factors and adjusted for the dependence structure.

Plugging the optimal parameters into the quadratic variation (3.5) leads to

$$\begin{aligned}
[X]_T^* &= \int_0^T \Big(\nabla\varphi(s,Z(s))\sigma(s,Z(s)) - \Theta^*(s,Z(s))\Big)\mathcal{Q}\Big(\nabla\varphi(s,Z(s))\sigma(s,Z(s)) - \Theta^*(s,Z(s))\Big)^\intercal ds \\
&+ \Big(\varphi^{J(s-)-1}(s,Z(s)) + D(s,F(s)) - \varphi^{J(s-)}(s,Z(s))\Big)dN(s),
\end{aligned}$$

which represents the remaining quadratic variation. We remark that the vector $\Big(\nabla\varphi\sigma - \Theta^*\Big)$ can be written as $\Big(\nabla\varphi\sigma^*\Big)$ for an approriate matrix σ^*.[1] Similarly, the risk-adjusted drift μ^* can be obtained from

$$\nabla\varphi^{J(s-)}(s,Z(s))\cdot\mu^*(s,Z(s)) = \nabla\varphi^{J(s-)}(s,Z(s))\cdot\mu(s,Z(s)) - \theta_1^*(s)\zeta(s,x(s),y(s)) - \theta_2^*(s)\gamma\sqrt{v(s)}.$$

[1] Given the optimal parameters, it is easy to decompose Θ^* as $\Theta^* = \nabla\varphi\tilde{\sigma}$. Then the appropriate matrix is $\sigma^* = \sigma - \tilde{\sigma}$ since $\nabla\varphi(\sigma - \tilde{\sigma}) = \nabla\varphi\sigma - \Theta^*$.

If we consider the gradient component related to f, we find that

$$\begin{aligned}\varphi_f^{J(s-)}(s,Z(s))\mu^{*,f}(s,Z(s)) &= \varphi_f^{J(s-)}(s,Z(s))\left[\left(r(s,x(s),y(s))-c+u\zeta(s)+(1-u)\gamma\sqrt{v(s)}\right)F(s)\right]\\ &\quad-\varphi_f^{J(s-)}(s,Z(s))uF(s)\zeta(s)\\ &\quad-\varphi_f^{J(s-)}(s,Z(s))(1-u)F(s)\gamma\sqrt{v(s)}\\ &= \varphi_f^{J(s-)}(s,Z(s))\left(r(s,x(s),y(s))-c\right).\end{aligned}$$

Hence, the risk-adjusted drift of the fund $F(t)$ is the risk-free rate $r(t,x(t),y(t))$ minus the fee rate c, and the equity risk-premium is cancelled.

Under the optimal hedging strategy, we find that

$$\begin{aligned}&\lim_{h\to 0}\frac{\mathbb{E}\big[X(t+h)e^{-r(t)h}-X(t)|Z(t)=z,J(t)=k\big]}{h}\\ =\;& \varphi_t^k(t,z)+\nabla\varphi^k(t,z)\cdot\mu^*(t,z)+\frac{1}{2}\operatorname{Tr}\left(\sigma(t,z)\mathcal{Q}\sigma^\intercal(t,z)\operatorname{Hess}_z\varphi^k(t,z)\right)\\ &-ckf+\Big(\varphi^{k-1}(t,z)+D(t,f)-\varphi^k(t,z)\Big)k\lambda-\varphi^k(t,z)r(t,x,y),\end{aligned}\tag{3.6}$$

and

$$\begin{aligned}&\lim_{h\to 0}\frac{\operatorname{Var}\big[X(t+h)e^{-r(t)h}-X(t)|Z(t)=z,J(t)=k\big]}{h}\\ &=\Big(\nabla\varphi^k(t,z)\sigma^*(t,z)\Big)\mathcal{Q}\Big(\nabla\varphi^k(t,z)\sigma^*(t,z)\Big)^\intercal+\Big(\varphi^{k-1}(t,z)+D(t,f)-\varphi^k(t,z)\Big)^2k\lambda.\end{aligned}\tag{3.7}$$

The last quantity represents the remaining quadratic variation and quantifies the non-hedgeable risks of the net asset value process X. We can conclude that the continuous-time fair valuation operator φ solves the following system of non-linear PDEs:

$$\begin{cases}\varphi_t^k(t,z)+\nabla\varphi^k(t,z)\cdot\mu^*(t,z)+\frac{1}{2}\operatorname{Tr}\left(\sigma(t,z)\mathcal{Q}\sigma(t,z)^\intercal\operatorname{Hess}_z\varphi^k(t,z)\right)\\ -ckf+\Big(\varphi^{k-1}(t,z)+D(t,f)-\varphi^k(t,z)\Big)k\lambda(t)-\varphi^k(t,z)r(t,x,y)\\ +\alpha\sqrt{\Big(\nabla\varphi^k(t,z)\sigma^*(t,z)\Big)\mathcal{Q}\Big(\nabla\varphi^k(t,z)\sigma^*(t,z)\Big)^\intercal+\Big(\varphi^{k-1}(t,z)+D(t,f)-\varphi^k(t,z)\Big)^2k\lambda}=0,\\ \varphi^k(T,z)=kS(f),\end{cases}\tag{3.8}$$

for $k\in\{0,\dots,n\}$. For each k, we have to deal with a 5-dimensional PDE. The number of PDEs in the system (3.8) is equal to $n+1$.

3.2 The PDE, FBSDE and FSDE representations

It is well known that there is a close relation between partial differential equations (PDEs) and backward stochastic differential equations (BSDEs). The solutions $(\varphi^k(t,z))_{k=0,\dots,n}$ to

the PDEs (3.8) satisfy the following forward-backward stochastic differential equation (FB-SDE) with jumps (see e.g. Barles et al. (1997) or Chapter 4 in Delong (2013)):

$$\begin{aligned}
&\varphi^{J(t)}(t,Z(t)) = J(T)S(F(T)) \\
&+\int_t^T \Upsilon^{J(s-)}\Big(s,Z(s),\varphi^{J(s-)}(s,Z(s)),\varphi^{J(s-)-1}(s,Z(s))\Big)ds + \int_t^T D(s,F(s))J(s-)\lambda(s)ds \\
&\qquad - \int_t^T \Big(cJ(s-)F(s) + \varphi^{J(s-)}(s,Z(s))r(s,x(s),y(s))\Big)ds \\
&\qquad\quad - \int_t^T \nabla\varphi^{J(s-)}(s,Z(s))\sigma(s,Z(s))dW(s) \\
&\qquad\qquad - \int_t^T \Big(\varphi^{J(s-)-1}(s,Z(s)) - \varphi^{J(s-)}(s,Z(s))\Big)d\tilde{N}(s), \quad (3.9)
\end{aligned}$$

$$\begin{aligned}
Z(t) &= Z(0) + \int_0^t \mu^*(s,Z(s))ds + \int_0^t \sigma(s,Z(s))dW(s), \\
dJ(t) &= -dN(t), \quad (3.10)
\end{aligned}$$

where

$$\begin{aligned}
&\Upsilon^{J(s-)}\Big(s,Z(s),\varphi^{J(s-)}(s,Z(s)),\varphi^{J(s-)-1}(s,Z(s))\Big) \\
&\quad = \alpha\Big\{\Big(\nabla\varphi^{J(s-)}(s,Z(s))\sigma^*(s,Z(s))\Big)\mathcal{Q}\Big(\nabla\varphi^{J(s-)}(s,Z(s))\sigma^*(s,Z(s))\Big)^{\intercal} \\
&\qquad + \Big(\varphi^{J(s-)-1}(s,Z(s)) + D(s,F(s)) - \varphi^{J(s-)}(s,Z(s))\Big)^2 J(s-)\lambda(s)\Big\}^{1/2}. \quad (3.11)
\end{aligned}$$

The equation (3.9) is a one-dimensional BSDE with jumps where the generator and the terminal condition depend on a process satisfying a forward stochastic differential equation. In our case, the FSDE is a 6-dimensional equation given by (3.10), where five components are diffusions and one component is driven by a step process. The BSDE (3.9) can also be written as a FSDE:

$$\begin{aligned}
&\varphi^{J(t)}(t,Z(t)) = \varphi^n(0,Z(0)) \\
&-\int_0^t \Upsilon^{J(s-)}\Big(s,Z(s),\varphi^{J(s-)}(s,Z(s)),\varphi^{J(s-)-1}(s,Z(s))\Big)ds - \int_0^t D(s,F(s))J(s-)\lambda(s)ds \\
&\qquad + \int_0^t \Big(cJ(s-)F(s) + \varphi^{J(s-)}(s,Z(s))r(s,x(s),y(s))\Big)ds \\
&\qquad\quad + \int_0^t \nabla\varphi^{J(s-)}(s,Z(s))\sigma(s,Z(s))dW(s) \\
&\qquad\qquad + \int_0^t \Big(\varphi^{J(s-)-1}(s,Z(s)) - \varphi^{J(s-)}(s,Z(s))\Big)d\tilde{N}(s). \quad (3.12)
\end{aligned}$$

Starting from an initial price of the liabilities: $\varphi^n(0, Z(0))$ (i.e. for n policyholders at time $t = 0$ with initial values of the risk factors: $z = Z(0)$), the FSDE (3.12) describes the price dynamics of the insurance liabilities through time.

The function Υ in (3.12) represents the price for the non-hedgeable risks: the first term in (3.11) is the residual risk from the diffusion component (interest rate, volatility, equity and systematic mortality risks) and the second term in (3.11) is the residual risk from the jump component (unsystematic mortality risk). We can notice that the risk margin for the non-hedgeable financial risks takes into account the correlation matrix between these risks. The second term in (3.12) represents the price for the expected death benefits. The first two terms in (3.12) decrease the price as they are released over time. The next two terms increase the price. The third term stands for the continuous fees paid by the policyholders to cover the guarantees and benefits (since the fees to hedge the future liabilities are collected by the insurer, the price of the future liabilities must increase over time). The fourth term describes the effect of interest accrual and the time value of money on the price dynamics. The fifth term shows the impact of the diffusion component. Finally, the last term depicts the consequence of the death of a policyholder on the price (there is one policyholder less $J(s-) \mapsto J(s) = J(s-) - 1$).

In order to derive a numerical algorithm to determine the price at time 0, i.e $\varphi^n(0, Z(0))$, we treat $\varphi^n(0, Z(0))$ as parameters for the moment and view the FSDE (3.12) as a way of computing the values of $\varphi^{J(t)}(t, Z(t))$ knowing the initial price $\varphi^n(0, Z(0))$ and the dynamics of φ. We consider a temporal discretization for the risk factors and the FSDE (3.12). For a set of time steps $0 = t_0 < t_1 < \ldots < t_N = T$, we consider the simple Euler scheme for $n = 1, \ldots, N-1$:

$$Z(t_{n+1}) - Z(t_n) \approx \mu^*(t_n, Z(t_n))\Delta t_n + \sigma(t_n, Z(t_n))\Delta W(t_n),$$

for the diffusion process (3.1) and

$$\tilde{N}(t_{n+1}) - \tilde{N}(t_n) = N(t_{n+1}) - N(t_n) - J(t_n)\lambda(t_n)\Delta t_n,$$

for the compensated jump process (2.3). For the FSDE, the Euler scheme has the following

form:

$$
\begin{aligned}
&\varphi^{J(t_{n+1})}(t_{n+1}, Z(t_{n+1})) = \varphi^{J(t_n)}(t_n, Z(t_n)) \\
&- \alpha\Big\{\Big(\nabla\varphi^{J(t_n)}(t_n, Z(t_n))\sigma^*(t_n, Z(t_n))\Big)\mathcal{Q}\Big(\nabla\varphi^{J(t_n)}(t_n, Z(t_n))\sigma^*(t_n, Z(t_n))\Big)^\intercal \\
&+ \Big(\varphi^{J(t_n)-1}(t_n, Z(t_n)) + D(t_n, F(t_n)) - \varphi^{J(t_n)}(t_n, Z(t_n))\Big)^2 J(t_n)\lambda(t_n)\Big\}^{1/2}\Delta t_n \\
&- D(t_n, F(t_n))J(t_n)\lambda(t_n)\Delta t_n \\
&+ \Big(cJ(t_n)F(t_n) + \varphi^{J(t_n)}(t_n, Z(t_n))r(t_n, x(t_n), y(t_n))\Big)\Delta t_n \\
&+ \nabla\varphi^{J(t_n)}(t_n, Z(t_n))\sigma(t_n, Z(t_n))\Delta W(t_n) \\
&+ \Big(\varphi^{J(t_n)-1}(t_n, Z(t_n)) - \varphi^{J(t_n)}(t_n, Z(t_n))\Big)\Delta\tilde{N}(t_n)
\end{aligned}
\tag{3.13}
$$

where

$$\Delta t_n = t_{n+1} - t_n, \quad \Delta W(t_n) = W(t_{n+1}) - W(t_n), \quad \Delta\tilde{N}(t_n) = \tilde{N}(t_{n+1}) - \tilde{N}(t_n).$$

To move from time t_n to time t_{n+1}, there are two unknown functions: the gradient of the price, $\nabla\varphi^{J(t_n)}(t_n, Z(t_n))$, and the price with one policyholder less, $\varphi^{J(t_n)-1}(t_n, Z(t_n))$. Following the work of Weinan et al. (2017) and Chan-Wai-Nam et al. (2019) (we remark that in both cases, the authors only consider BSDEs without jumps), we approximate both functions by the use of neural networks.

3.3 Feed-forward neural networks as function approximators

Hereafter, we briefly introduce feed-forward neural networks (short *neural networks*). In regression theory, networks can be seen as a broad generalization of generalized linear models (GLMs) where the linear predictor is replaced by a non-linear one.

Let us assume that the input is in dimension d_0 (the state variable x) and the output is in dimension d_1 (the number of value functions to estimate). The network is characterized by a number of layers $L + 1 \in \mathbb{N}\backslash\{1, 2\}$ with $m_\ell, \ell = 0, \ldots, L$, the number of neurons (units or nodes) on each layer: the first layer is the input layer with $m_0 = d_0$, the last layer is the output layer with $m_L = d_1$, and the $L - 1$ layers between are called hidden layers, where we choose for simplicity the same dimension $m_\ell = m, \ell = 1, \ldots, L - 1$. A network is a function from $\mathbb{R}^{d_0}$ to $\mathbb{R}^{d_1}$ defined as the composition

$$x \in \mathbb{R}^{d_0} \mapsto \mathcal{N}(x) = A_L \circ \varrho \circ A_{L-1} \circ \ldots \circ \varrho \circ A_1(x) \in \mathbb{R}^{d_1}$$

Here, $A_\ell, \ell = 1, \ldots, L$, are affine transformations represented by

$$A_\ell(x) = \mathcal{W}_\ell x + \beta_\ell,$$

for a matrix of weights $\mathcal{W}_\ell$ and a bias term β_ℓ. The non-linear function $\varrho : \mathbb{R} \to \mathbb{R}$ is called the activation function and is applied component-wise on the outputs of A_ℓ, i.e., $\varrho(x_1, \ldots, x_m) = (\varrho(x_1), \ldots, \varrho(x_m))$. Standard examples of activation functions are the sigmoid, the ReLU, the elu and tanh. For a general introduction on neural networks, we refer to Goodfellow et al. (2016).

The universal approximation theorem of Hornik et al. (1989) states that networks can approximate any continuous function on a compact support arbitrarily well if we allow for arbitrarily many neurons $q_1 \in \mathbb{N}$ in the hidden layer.

In the Euler scheme (3.13) of the FBSDE, we approximate the gradient and the price with one less policyholder by two respective neural networks. Moreover, in practice, prices need to be computed several times for different input parameters. Therefore, as suggested in Weinan et al. (2017), we determine the price at time 0 in a region of input parameters instead of a single space-point by including a third neural network. Overall, we have the three following networks as depicted in Figure 1:

N1. The price at time 0, $\varphi^{J(0)}(0, Z(0))$, is approximated by a neural network $\mathcal{N}_1 : (z, k) \in \mathbb{R}^6 \mapsto \varphi^k(0, z) \in \mathbb{R}$ with parameters ϕ:

$$\varphi^{J(0)}(0, Z(0)) = \mathcal{N}_1^\phi(Z(0), J(0)).$$

N2. The gradient, $\nabla\varphi^{J(t_n)}(t_n, Z(t_n))$, is approximated by a second neural network $\mathcal{N}_2 : (t, z, k) \in \mathbb{R}^7 \mapsto \nabla\varphi^k(t, z) \in \mathbb{R}^5$ with parameters χ:

$$\nabla\varphi^{J(t_n)}(t_n, Z(t_n)) = \mathcal{N}_2^\chi(t_n, Z(t_n), J(t_n)).$$

N3. The price with one less policyholder, $\varphi^{J(t_n)-1}(t_n, Z(t_n))$, is approximated by a third neural network $\mathcal{N}_3 : (t, z, k-1) \in \mathbb{R}^7 \mapsto \varphi^{k-1}(t, z) \in \mathbb{R}$ with parameters ψ:

$$\varphi^{J(t_n)-1}(t_n, Z(t_n)) = \mathcal{N}_3^\psi(t_n, Z(t_n), J(t_n) - 1).$$

By the use of the Euler scheme (3.13) and the neural networks, we can obtain a Monte-Carlo approximation $\widehat{\varphi}$ of $\varphi^{J(T)}(T, Z(T))$ which depends on the parameters of the neural networks (ϕ, χ, ψ). These parameters are then estimated in order to minimize the quadratic loss function:

$$l(\phi, \chi, \psi) = \mathbb{E}\left[\left|J(t_N)S(F(t_N)) - \widehat{\varphi}^{J(t_N)}(t_N, Z(t_N))\right|^2\right]. \tag{3.14}$$

The parameters can now be optimized by stochastic gradient descent-type (SGD) algorithm as it is usually done to train neural networks. In this paper, we use the Adaptive Moment Estimation (Adam) optimizer proposed by Kingma and Ba (2014) and the algorithm is implemented with the Keras R package (Chollet et al. (2017)), which is a user-friendly API to TensorFlow.

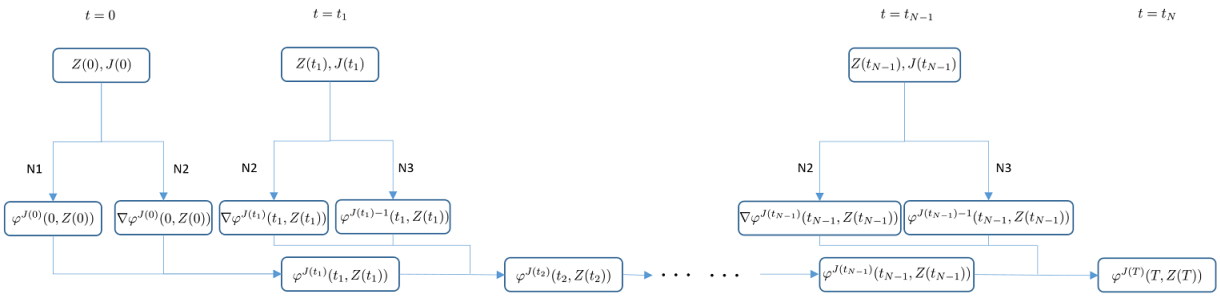

Figure 1: Illustration of the network architecture for solving the system of non-linear PDEs with N time steps. The network N1 stands for the price at time 0, the network N2 approximates the gradient and the network N3 is related to the jump component and approximates the price with one policyholder less.

Remark 3.1.

- One can notice either from the PDEs (3.8) or the FSDE (3.12) that we face a recursive problem: in order to determine the price of the insurance liabilities for n initial policyholders, we need to know the price for $k \in \{0, \dots, n-1\}$. One possibility of implementation would have been to solve n recursive optimizations with one neural network for the gradient in each optimization. Instead, we consider only one optimization with two neural networks (one for the gradient, one for the jump component) in order to reduce the computational cost of the problem.

- Contrary to Weinan et al. (2017), we consider one single network for all time steps instead of one neural network per time step in order to reduce the numbers of parameters to be estimated. Chan-Wai-Nam et al. (2019) showed that this neural network architecture significantly improves the precision of the results and the stability of the algorithms.

4 Numerical results

This section presents numerical results of our neural networks to price equity-linked contracts with death and survival benefits. In Section 4.1, we briefly discuss our input parameters and the hyperparameters of the neural networks. In Section 4.2, we solve the system of PDEs (3.8) with our neural networks and analyse the impact of various model parameters on the price of equity-linked contracts. In Section 4.3, we assess the accuracy of our approach for some specific contracts. Let us point out that we only investigate the initial price of the contracts, i.e. at time $t = 0$.

4.1 Specification of the insurance-financial model and neural networks

For the financial and insurance model, we consider a realistic set of model parameters which are consistent with previous studies:

- For the two-factor interest rate model, we consider

$$a = 0.2770,\ \sigma_x = 0.0118,\ b = 0.0551, \sigma_y = 0.0136,$$

  which are based on Russo and Torri (2019). Moreover, we chose: $\delta_x = \delta_y = -0.1, \psi(t) = 0.02$.

- For the Heston stochastic volatility model, the parameters are

$$\kappa = 0.0231,\ \eta = 0.9052, \sigma_v = 0.1434, \gamma = 0.0113,$$

  which are consistent with Li et al. (2016).

- For the force of mortality dynamics, we choose the parameters from Luciano and Vigna (2008) which correspond to UK individuals aged 65 at time 0:

$$q = 0.11,\ \sigma_\lambda = 0.007.$$

- The correlation matrix for the interest rate factors, the equity, the volatility and the mortality intensity is given by

$$\mathcal{Q} = \begin{pmatrix} 1 & -0.4 & 0.35 & 0 & 0 \\ -0.4 & 1 & 0.08 & 0 & 0 \\ 0.35 & 0.08 & 1 & -0.3 & 0 \\ 0 & 0 & -0.3 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \tag{4.1}$$

  which is consistent with Grzelak et al. (2011).

Additional input parameters are presented in Table 1. We point out that the pricing neural network should not only deliver one price but a surface of prices which is a function of the interest rate factors, the premium, the volatility, the number of policyholders and the force of mortality at time $t = 0$. Hence, we are interested in determining the price denoted by $\varphi^{J(0)}(0, Z(0))$. Here, $Z(0) = (x(0), y(0), F(0), v(0), \lambda(0))$ stands for the set of initial values for the interest rate factors, the premium, the volatility and the force of mortality respectively, and $J(0)$ stands for the number of policyholders at the inception. In our numerical example

Parameter	Value	Parameter	Value
c	0.01	$x(0)$	0
$D^\star$	1.02	$y(0)$	0
$S^\star$	1.02	$F(0)$	1
T	1	$v(0)$	0.1
Δt	0.01	$\lambda(0)$	0.015
α	0.1	$J(0)$	100

Table 1: The base values of the initial parameters.

we investigate ranges of input parameters by changing the values for $v(0), F(0), \lambda(0), J(0)$ specified in Table 1 by $+/-25\%$. For the interest rates, we consider the ranges $x(0) \in [0, 0.01], y(0) \in [0, 0.01]$.

For the neural network architecture, different hyperparameters need to be chosen. Each of the neural networks **N1**, **N2**, **N3** defined in Section 3.3 consists of two hidden layers of 20 neurons. For the activation function, we considered the ELU (exponential linear unit) which tends to accelerate the learning process (Clevert et al. (2015)). Moreover, neural networks tend to have convergence issues when their inputs are not scaled and centered, and perform best when the inputs are in the range $\sim [-1, 1]$. Therefore, we considered a min-max normalization of the inputs before passing them to the neural networks. Hence, we applied the function

$$g : \mathbb{R} \mapsto [-1, 1] : g(x) = 2 \times \left(\frac{x - \min(x, 0)}{\max(x, 0) - \min(x, 0)} \right) - 1.$$

Finally, we chose a batch size and a number of epochs equal to 200.[2] We considered 10000 sample paths for each risk factor and the default learning rate of the Adam optimizer. We tried different hyperparameters and we did not observed significant improvements in training of the neural networks.

4.2 Pricing in the general case and sensitivity analysis

Each equity-linked life insurance contract provides insurance death and survival benefits linked to the policyholder's account value. These financial guarantees ensure that the benefits are not lower than the minimal guarantee levels stipulated at the inception of the contract. We set the benefits:

$$D(t, F(t)) = (D^* - F(t))_+, \quad S(F(T)) = (S^* - F(T))_+.$$

[2]The computation time for the training of the price surface takes about 10 min on a computer with an Intel Core i7-4970MQ processor running at 2.90 GHz with 16 GB of RAM.

The contract payoffs if the policyholder's account value goes below D^* at the time of death or below S^* in case of survival at the end of the contract. Hence, we are interested in the fair valuation of the financial guarantees embedded in the insurance contracts.

4.2.1 Neural network's convergence diagnostics

First, we consider the convergence of our algorithm for one price with the initial values $Z(0) = (0, 0, 1, 0.1, 0.015)$ and $J(0) = 100$ policyholders from Table 1. Figure 2 represents the convergence of the price and the mean square error against the number of epochs used for training the neural networks. Let us recall that our objective in the reinforcement learning is to minimize the mean square error (3.14), hence we present the mean square error loss function to validate the convergence. We observe that, with the initialization of our algorithm by default, the price reaches a plateau after around 100 epochs. After 100 epochs, the MSE continues to slightly decreases but without significant impact on the price. We arrive at the price of the portfolio of the equity-linked contracts equal to $\varphi^{J(0)}(0, Z(0)) = 6.069$.

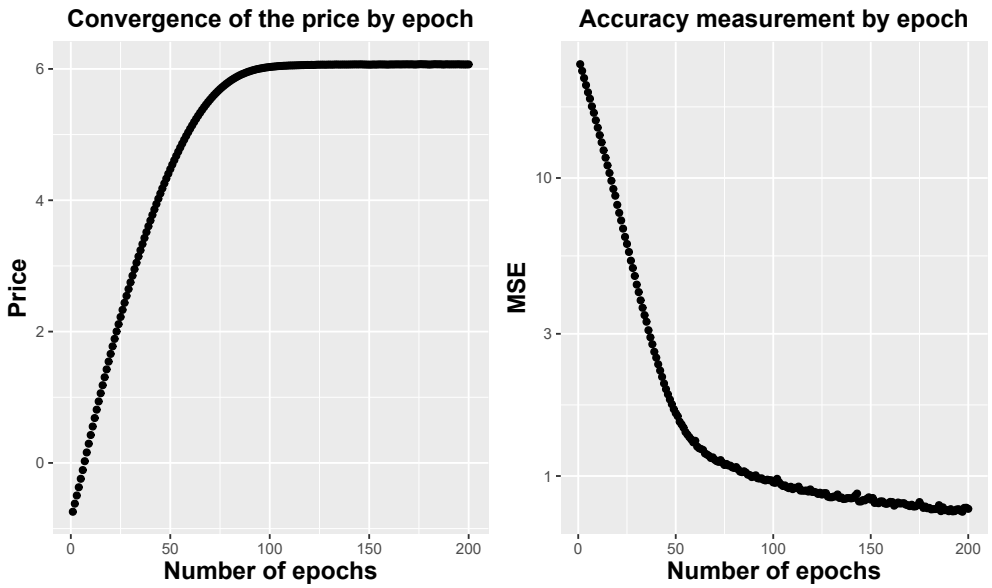

Figure 2: Diagnostics of the neural network training. On the left: Convergence of the price at time 0 per epoch. On the right: Mean Square Error per epoch.

Now, we study the convergence of the algorithm for a surface of prices for different sets of initial parameters $Z(0)$ and $J(0)$. The sets were constructed by taking the values given in Table 1 and changing $v(0), F(0), \lambda(0), J(0)$ by $+/- 25\%$ and considering $x(0) \in [0, 0.01], y(0) \in [0, 0.01]$. Figure 3 shows the price convergence for 5 random choices of initial sets of parameters. We observe that the machine learns the price essentially in the first 100 epochs and then stabilizes. We remark that the price of the portfolio of the equity-linked

contracts can possibly be negative as the contracts contain a fee rate. If the price is negative, it means that the fees deducted from the policyholder's account are too high in comparison with the guarantees embedded in the contract and the insurer expects to earn a profit.

From Figures 2-3 and similar analysis we performed, we can conclude that our algorithm is stable and the convergence is achieved.

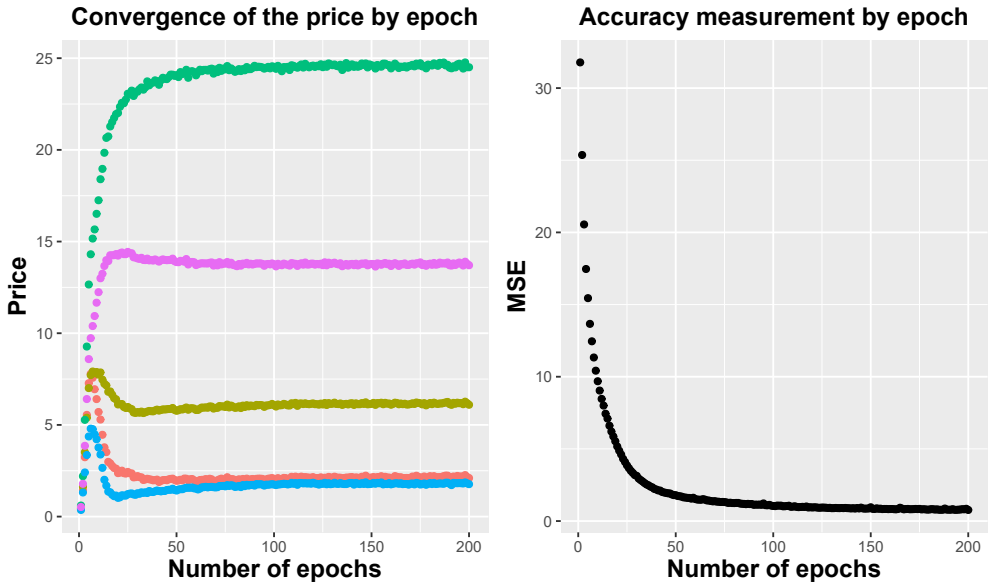

Figure 3: On the left: Convergence of the price at time 0 during the neural network training for five randomly chosen initial parameters. On the right: Mean Square Error per epoch during the neural network training of the surface of prices.

4.2.2 Price sensitivity analysis

Hereafter, we illustrate the impact of various model parameters on the price of the portfolio of the equity-linked contracts. One of the purpose of this exercise is to validate that our PDEs provide a reasonable pricing operator.

Figure 4 represents the effect of the initial interest rate factors $x(0)$ and $y(0)$ on the price. As expected, the price decreases as the interest rate factors increase due to the time value of money. In particular, if both factors increase from 0 to 0.01, the price decreases by 16%.

Figure 5 shows the effect of varying the initial premium $F(0)$ and the number of policyholders $J(0)$ on the price. Obviously, the price of the portfolio of the equity-linked contracts increases with the number of policyholders since more policyholders mean more benefits to pay. Moreover, we note that the price of the portfolio of the equity-linked contracts decreases with the initial premium. Indeed, the increase of the initial fund value decreases the probability that the fund will go below the death guarantee level D^* and the survival guarantee

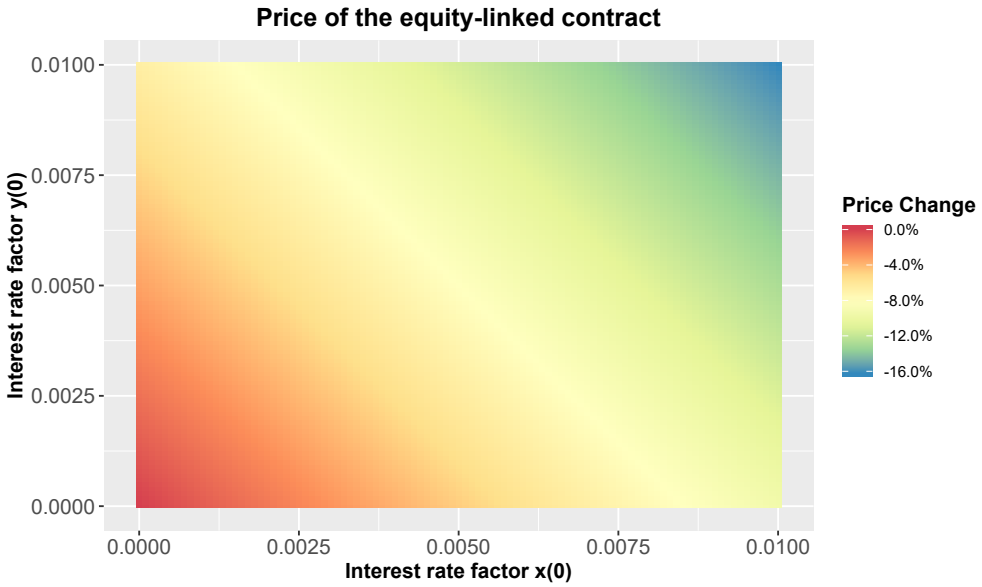

Figure 4: The price of the portfolio of equity-linked contracts for a range of the initial interest rate factors $x(0) \in [0; 0.01]$ and $y(0) \in [0; 0.01]$.

level S^*, and therefore decreases the price of the contract. We can observe that the price can double if the initial premium decreases by 10% as depicted in Figure 5.

Figure 6 depicts the impact of varying the initial volatility parameter $v(0)$ and the initial force of mortality $\lambda(0)$ on the price of the portfolio of the equity-linked contracts. As expected, we note that the price is an increasing function of the initial volatility. A highly volatile market environment implies that the stock and the fund values are less predictable, which will result in a higher risk charge required by the insurer in order to cover the guarantees. When the initial force of mortality increases, the value of the death benefit increases while the value of the survival benefit decreases as policyholders are more likely to die. On Figure 6, the price is a fairly constant function of the initial force of mortality, showing that the value reduction of the survival benefit compensates the value increase of the death benefit, at least for $D^* = S^* = 1.02$ as in our numerical example.

We also investigate the impact of the correlation between the fund and the interest rate factor x as well as the correlation between the fund and the volatility on the price of the liabilities. The results are presented in Table 2. As expected, the lower the correlation coefficient, the lower the price for the portfolio of the equity-linked contracts. This property reflects diversification effect between the financial risks and is consistent with the findings reported in Gudkov et al. (2019) and Dai et al. (2015).

Finally, we test the effect of the risk aversion coefficient α. We expect that the more risk averse the insurer is, the larger the risk margin he will keep to cover the non-hedgeable risks

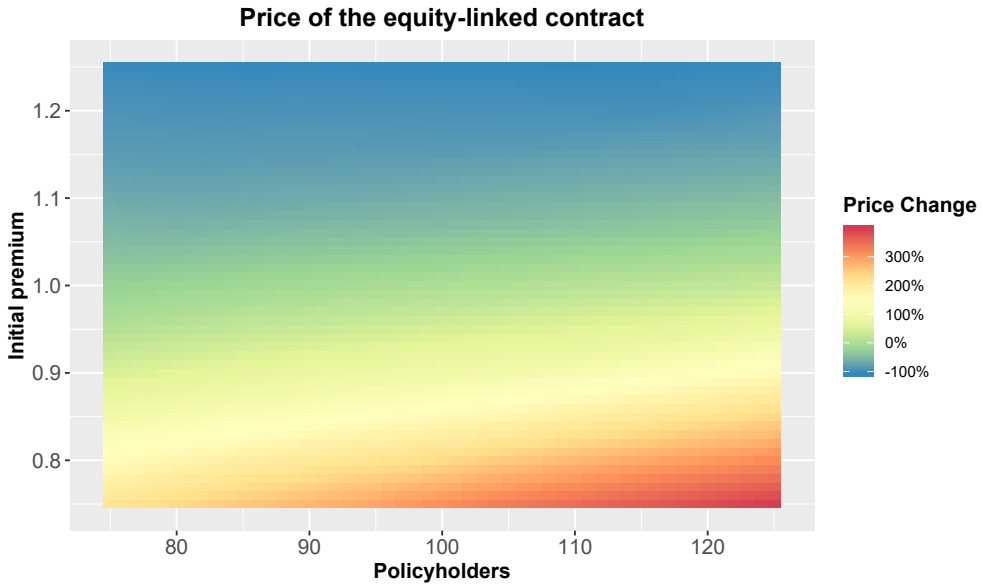

Figure 5: The price of the portfolio of equity-linked contracts for a range of the initial premiums $F(0) \in [0.75; 1.25]$ and the numbers of policyholders $J(0) \in [75; 125]$.

Change of the correlation	Change of the price	
	ρ_{xf}	ρ_{fv}
$+0.1$	$+0.106\%$	$+0.131\%$
-0.1	-0.248%	-0.447%

Table 2: The percentage change of the price of the portfolio of equity-linked contracts for a correlation change by ± 0.1.

and the larger the price of the portfolio of the equity-linked contracts. This is confirmed in Figure 7. We note that the relation between the price and the risk aversion α is not linear. The case $\alpha = 0$ corresponds to the best estimate of the liability, which we discuss in the next subsection, when the insurer does not set any risk margin for the non-hedgeable risks.

4.3 Numerical comparisons

Hereafter, we assess the accuracy of our neural network approach by comparing it with Monte-Carlo pricing for two particular cases.

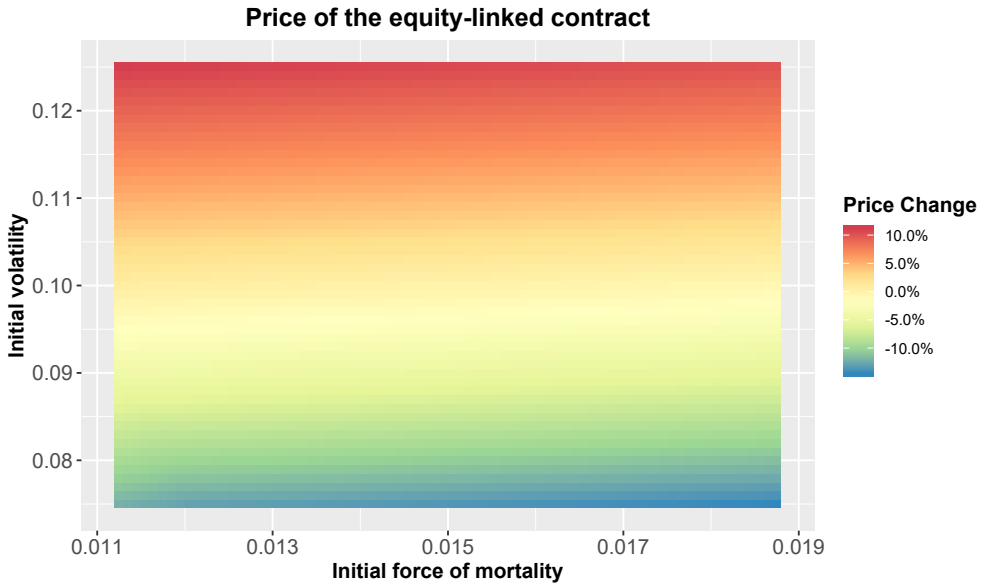

Figure 6: The price of the portfolio of equity-linked contracts for a range of the initial volatilities $v(0) \in [0.075; 0.125]$ and the forces of mortality $\lambda(0) \in [0.01125; 0.01875]$.

4.3.1 Best Estimate Liability

In insurance, the Best Estimate Liability (BEL) plays a central role in the valuation of insurance liabilities. It is defined as the expected present value of the future cashflows discounted by the risk-free rate. For our portfolio of the equity-linked life insurance contracts, the BEL is given by

$$\begin{aligned} BEL \;=\; & \mathbb{E}\left[e^{-\int_t^T r(\tau)d\tau} J(T) S(F(T)) + \int_t^T e^{-\int_t^s r(\tau)d\tau} D(s, F(s)) J(s) \lambda(s) ds \right. \\ & \left. - \int_t^T e^{-\int_t^s r(\tau)d\tau} c J(s) F(s) \Big| J(t), x(t), y(t), F(t), v(t), \lambda(t) \right] \end{aligned} \tag{4.2}$$

where the first term is the expected survival benefits, the second term - the expected death benefits and the third term - the expected fees. By the Feynman-Kac formula, the BEL is also solution to the following system of PDEs:

$$\begin{cases} \varphi_t^k(t,z) + \nabla \varphi^k(t,z) \cdot \mu^*(t,z) + \dfrac{1}{2} \operatorname{Tr}\left(\sigma(t,z) \mathcal{Q} \sigma(t,z)^\intercal \operatorname{Hess}_z \varphi^k(t,z) \right) \\ \quad - ckf + \Big(\varphi^{k-1}(t,z) + D(t,f) - \varphi^k(t,z) \Big) k\lambda(t) - \varphi^k(t,z) r(t,x,y) = 0, \\ \varphi^k(T,z) = kS(f), \end{cases} \tag{4.3}$$

for $k \in \{0, \ldots, n\}$.

The system (4.3) corresponds to our pricing equation (3.8) without the risk margin (with $\alpha = 0$). To assess the accuracy of our neural network, we compare the price obtained by an

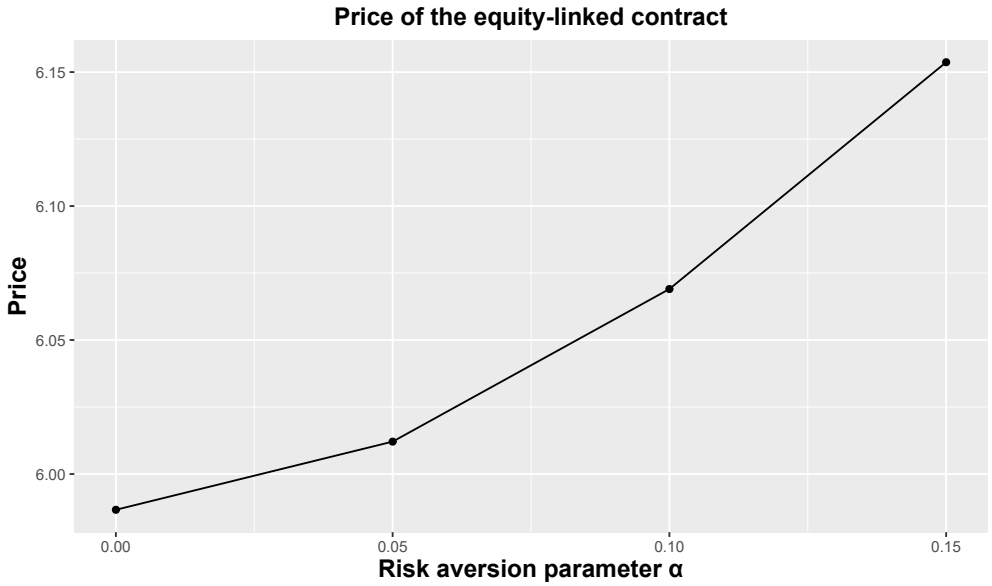

Figure 7: The price of the portfolio of equity-linked contracts in terms of the risk aversion parameter α.

approximation of the BEL (4.2) by 200.000 Monte-Carlo simulations and the price obtained by our neural network using the same parameters as in Table 1. Figure 8 shows that there is a good convergence of the neural network to the Monte-Carlo price. After 100 epochs, we obtain a relative error of c.a. 0.1%.

We could include the value of the BEL as an initial point from which we start to calibrate our neural networks for the price of the portfolio of unit-linked contracts since we expect that the true price with a risk margin for $\alpha > 0$ should be close to the price with $\alpha = 0$. Our tests showed that such an approach does not speed up the calibration of our neural networks mainly due to the fact that we still have to initiate the neural network for the gradient from a random starting point.

4.3.2 Pricing a portfolio of GMMB contracts in the Black-Scholes framework

In the following example, we consider a portfolio of equity-linked contracts in which only a survival benefit is guaranteed. Such contracts are often called Guaranteed Minimum Maturity Benefits (GMMB). We consider a simplified financial-actuarial market, namely the Black-Scholes financial market model and a jump process $J(t)$ with constant mortality λ to model the deaths of policyholders. Under these assumptions, the pricing PDEs (3.8) reduce

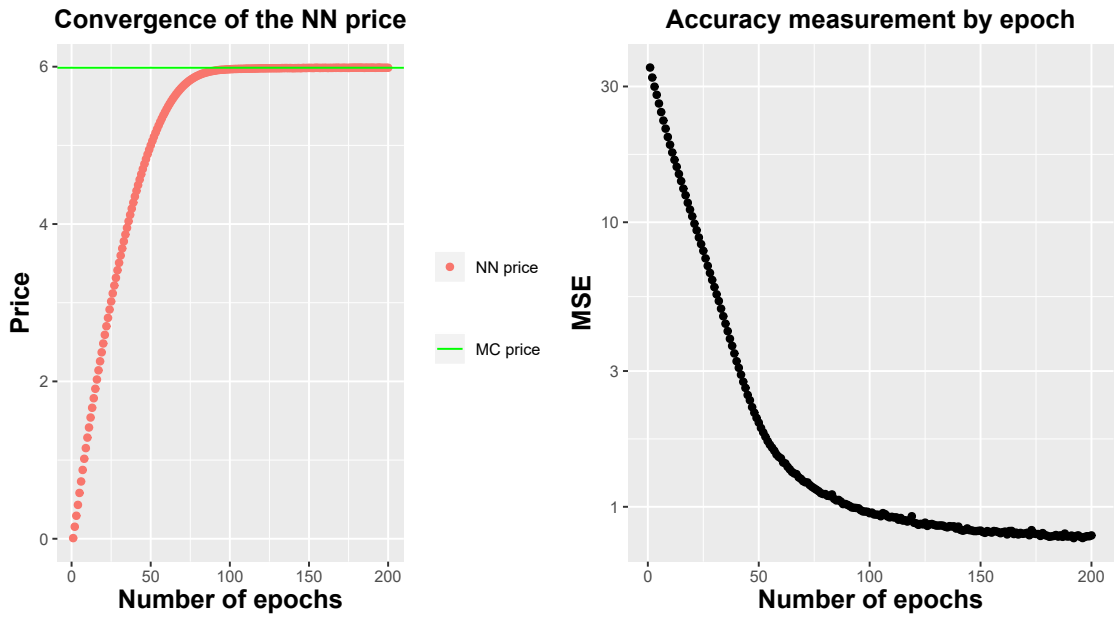

Figure 8: Neural Network price convergence to the Monte-Carlo price for the BEL and the Mean Square Error per epoch

to

$$\begin{cases} \varphi_t^k(t,f) + \varphi_f^k(t,f) f r + \frac{1}{2}\varphi_{ff}^k(t,f) f^2 \sigma_f^2 - \varphi^k(t,f) r \\ + \left(\varphi^{k-1}(t,f) - \varphi^k(t,f)\right) k\lambda \left(1 - \frac{\alpha}{\sqrt{k\lambda}}\right) = 0, \\ \varphi^k(T,f) = kS(f), \end{cases} \tag{4.4}$$

for $k \in \{0, \ldots, n\}$.

We remark that this PDE was considered in (Delong et al. (2019a), Example 4.5). The first part of the PDE (4.4) exactly corresponds to the standard Black-Scholes PDE but the second part adds a risk margin for the uncertainty about the number of survivals which cannot be completely hedged. Applying the Feynman-Kac formula, the solution of the PDEs (4.4) is given by

$$\varphi^n(0, F(0)) = \mathbb{E}^{\tilde{\mathbb{Q}}}\left[e^{-rT} J(T) S(F(T))\right], \tag{4.5}$$

where under $\tilde{\mathbb{Q}}$, the jump process has now a stochastic mortality intensity given by

$$\tilde{\lambda}(t) = \lambda\left(1 - \frac{\alpha}{\sqrt{J(t)\lambda}}\right).$$

We choose the parameters: $r = 0.02, F(0) = 1, \sigma_f = 0.1, S^* = 1.02, T = 1, n = 100, \alpha = 0.1$. Under these parameters, $\tilde{\lambda}(t)$ is strictly positive which allows us to use a step process to simulate the number of deaths under the measure $\tilde{\mathbb{Q}}$.

Figure 9 compares the Monte-Carlo price obtained by approximating (4.5) with 200.000 simulations and the price obtained by neural networks by solving the PDEs (4.4). We clearly observe a fast convergence after only 25 epochs and a relative error of c.a. 0.2% after 100 epochs.

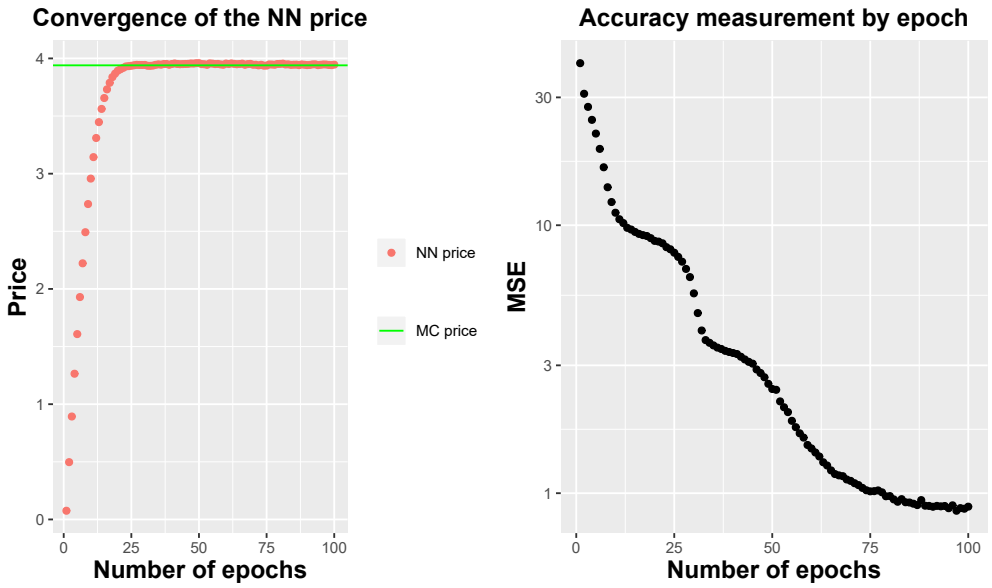

Figure 9: Neural Network price convergence to the Monte-Carlo price for the GMMB and Mean Square Error per epoch.

Summing up, both examples show that our algorithm is accurate. The remaining error is due to discretization steps applied in reinforcement learning, when we look for the optimal strategy in a discrete time, and Monte-Carlo simulations of discretized stochastic processes. Moreover, we note that our remaining MSE is of the order $[10^{-2}; 10^{-1}]$ which is consistent with orders obtained in Chan-Wai-Nam et al. (2019). The second example in this section also shows that our approach to handling the jump component performs very well.

5 Conclusion

In this paper, we have derived a partial differential equation for fair pricing of equity-linked life insurance contracts in a general financial-actuarial market with stochastic interest rate, equity, volatility and mortality. We assumed that the insurer can hedge its liabilities with a bank account, a bond and an equity and requires a compensation for non-hedgeable risks in the form of an instantaneous standard deviation risk margin. We derived a system of PDEs which described the fair valuation operator. We used the connection with BSDEs with jumps and proposed an efficient neural network architecture to solve the PDEs. The

network consists essentially of two subnetworks: one to approximate the gradient component and one to approximate the jump component. We also included a third subnetwork for the price at time 0 which allows us to derive the price for a range of initial parameters rather than just one price.

Numerical results provided convergence diagnostics, price sensitivities to various model parameters and some comparisons with Monte-Carlo methods. In particular, we observed a fast convergence and stabilization of the price after a few dozens of epochs and a good approximation compared to prices obtained by Monte-Carlo simulations. In terms of sensitivity, we proved that the calibrated price agreed with intuition.

Our paper provides a stable computational algorithm for fair pricing of insurance liabilities in incomplete markets which can be used for a wide variety of financial and insurance products. Moreover, our neural network offers potential other applications in fields where a system of non-linear PDEs naturally arises, such as quantum mechanics, game theory, dynamic programming and many more.

6 Acknowledgements

The authors would like to thank the Editor and the anonymous referee who provided useful and detailed comments that substantially improved the current manuscript. Karim Barigou acknowledges the financial support of the Joint Research Initiative on "Mortality Modeling and Surveillance" funded by AXA Research Fund.

References

Ait-Sahalia, Y. and Kimmel, R. L. (2010), 'Estimating affine multifactor term structure models using closed-form likelihood expansions', *Journal of Financial Economics* **98**(1), 113–144.

Barigou, K., Chen, Z. and Dhaene, J. (2019), 'Fair dynamic valuation of insurance liabilities: Merging actuarial judgement with market-and time-consistency', *Insurance: Mathematics and Economics* **88**, 19–29.

Barles, G., Buckdahn, R. and Pardoux, E. (1997), 'Backward stochastic differential equations and integral-partial differential equations', *Stochastics: An International Journal of Probability and Stochastic Processes* **60**(1-2), 57–83.

Bauer, D., Kling, A. and Russ, J. (2008), 'A universal pricing framework for guaranteed minimum benefits in variable annuities', *Astin Bulletin* **38**(02), 621–651.

Bayraktar, E., Milevsky, M. A., Promislow, S. D. and Young, V. R. (2009), 'Valuation of mortality risk via the instantaneous sharpe ratio: applications to life annuities', *Journal of Economic Dynamics and Control* **33**(3), 676–691.

Beck, C., Weinan, E. and Jentzen, A. (2019), 'Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations', *Journal of Nonlinear Science* **29**(4), 1563–1619.

Brigo, D. and Mercurio, F. (2007), *Interest rate models-theory and practice: with smile, inflation and credit*, Springer Science & Business Media.

Castro, J. (2021), 'Deep learning schemes for parabolic nonlocal integro-differential equations', *arXiv preprint arXiv:2103.15008* .

Chan-Wai-Nam, Q., Mikael, J. and Warin, X. (2019), 'Machine learning for semi linear pdes', *Journal of Scientific Computing* **79**(3), 1667–1712.

Chollet, F., Allaire, J. et al. (2017), 'R interface to keras', `https://github.com/rstudio/keras`.

Clevert, D.-A., Unterthiner, T. and Hochreiter, S. (2015), 'Fast and accurate deep network learning by exponential linear units (elus)', *arXiv preprint arXiv:1511.07289* .

Da Fonseca, J. and Ziveyi, J. (2015), 'Valuing variable annuity guarantees on multiple assets', *Scandinavian Actuarial Journal* pp. 1–22.

Dai, T.-S., Yang, S. S. and Liu, L.-C. (2015), 'Pricing guaranteed minimum/lifetime withdrawal benefits with various provisions under investment, interest rate and mortality risks', *Insurance: Mathematics and Economics* **64**, 364–379.

Delong, Ł. (2013), *Backward stochastic differential equations with jumps and their actuarial and financial applications*, Springer.

Delong, Ł., Dhaene, J. and Barigou, K. (2019*a*), 'Fair valuation of insurance liability cashflow streams in continuous time: Applications', *ASTIN Bulletin: The Journal of the IAA* **49**(2), 299–333.

Delong, Ł., Dhaene, J. and Barigou, K. (2019*b*), 'Fair valuation of insurance liability cashflow streams in continuous time: Theory', *Insurance: Mathematics and Economics* .

Dhaene, J., Stassen, B., Barigou, K., Linders, D. and Chen, Z. (2017), 'Fair valuation of insurance liabilities: Merging actuarial judgement and market-consistency', *Insurance: Mathematics and Economics* **76**, 14–27.

Doyle, D. and Groendyke, C. (2019), 'Using neural networks to price and hedge variable annuity guarantees', *Risks* **7**(1), 1.

Fung, M. C., Ignatieva, K. and Sherris, M. (2014), 'Systematic mortality risk: An analysis of guaranteed lifetime withdrawal benefits in variable annuities', *Insurance: Mathematics and Economics* **58**, 103–115.

Goodfellow, I., Bengio, Y. and Courville, A. (2016), *Deep learning*, MIT press.

Grzelak, L. A., Oosterlee, C. W. and Van Weeren, S. (2011), 'The affine heston model with correlated gaussian interest rates for pricing hybrid derivatives', *Quantitative Finance* **11**(11), 1647–1663.

Gudkov, N., Ignatieva, K. and Ziveyi, J. (2019), 'Pricing of guaranteed minimum withdrawal benefits in variable annuities under stochastic volatility, stochastic interest rates and stochastic mortality via the componentwise splitting method', *Quantitative Finance* **19**(3), 501–518.

Haefeli, D. (2013), 'Variable annuities-an analysis of financial stability', *The Geneva Association* .

Han, J., Jentzen, A. and Weinan, E. (2018), 'Solving high-dimensional partial differential equations using deep learning', *Proceedings of the National Academy of Sciences* **115**(34), 8505–8510.

Hejazi, S. A. and Jackson, K. R. (2016), 'A neural network approach to efficient valuation of large portfolios of variable annuities', *Insurance: Mathematics and Economics* **70**, 169–181.

Heston, S. L. (1993), 'A closed-form solution for options with stochastic volatility with applications to bond and currency options', *The review of financial studies* **6**(2), 327–343.

Hornik, K., Stinchcombe, M., White, H. et al. (1989), 'Multilayer feedforward networks are universal approximators.', *Neural networks* **2**(5), 359–366.

Ignatieva, K., Song, A. and Ziveyi, J. (2016), 'Pricing and hedging of guaranteed minimum benefits under regime-switching and stochastic mortality', *Insurance: Mathematics and Economics* **70**, 286–300.

Kingma, D. P. and Ba, J. (2014), 'Adam: A method for stochastic optimization', *arXiv preprint arXiv:1412.6980* .

Li, D., Rong, X. and Zhao, H. (2016), 'Optimal reinsurance and investment problem for an insurer and a reinsurer with jump-diffusion risk process under the heston model', *Computational and Applied Mathematics* **35**(2), 533–557.

Liang, X. and Lu, Y. (2017), 'Indifference pricing of a life insurance portfolio with risky asset driven by a shot-noise process', *Insurance: Mathematics and Economics* **77**, 119–132.

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

Luciano, E., Spreeuw, J. and Vigna, E. (2008), 'Modelling stochastic mortality for dependent lives', *Insurance: Mathematics and Economics* **43**(2), 234–244.

Luciano, E. and Vigna, E. (2008), 'Mortality risk via affine stochastic intensities: calibration and empirical relevance', *Belgian Actuarial Bulletin* **8**(1), 5–16.

McNeil, A. J., Frey, R. and Embrechts, P. (2005), *Quantitative Risk Management: Concepts, Techniques and Tools-revised edition*, Princeton university press.

Palmer, B. (2006), 'Equity-indexed annuities: Fundamental concepts and issues', *Work* .

Pelsser, A. and Stadje, M. (2014), 'Time-consistent and market-consistent evaluations', *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics* **24**(1), 25–65.

Russo, V. and Torri, G. (2019), 'Calibration of one-factor and two-factor hull–white models using swaptions', *Computational Management Science* **16**(1-2), 275–295.

Weinan, E., Han, J. and Jentzen, A. (2017), 'Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations', *Communications in Mathematics and Statistics* **5**(4), 349–380.