

A Service of



Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Cheridito, Patrick; Ery, John; Wüthrich, Mario V.

Article

Assessing asset-liability risk with neural networks

Risks

Provided in Cooperation with: MDPI – Multidisciplinary Digital Publishing Institute, Basel

Suggested Citation: Cheridito, Patrick; Ery, John; Wüthrich, Mario V. (2020) : Assessing asset-liability risk with neural networks, Risks, ISSN 2227-9091, MDPI, Basel, Vol. 8, Iss. 1, pp. 1-17, https://doi.org/10.3390/risks8010016

This Version is available at: https://hdl.handle.net/10419/257971

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.



WWW.ECONSTOR.EU

https://creativecommons.org/licenses/by/4.0/

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.







Article Assessing Asset-Liability Risk with Neural Networks

Patrick Cheridito *, John Ery^D and Mario V. Wüthrich

RiskLab, ETH Zurich, 8092 Zurich, Switzerland; john.ery@math.ethz.ch (J.E.); mario.wuethrich@math.ethz.ch (M.V.W.)

* Correspondence: patrick.cheridito@math.ethz.ch

Received: 30 November 2019; Accepted: 4 February 2020; Published: 9 February 2020



Abstract: We introduce a neural network approach for assessing the risk of a portfolio of assets and liabilities over a given time period. This requires a conditional valuation of the portfolio given the state of the world at a later time, a problem that is particularly challenging if the portfolio contains structured products or complex insurance contracts which do not admit closed form valuation formulas. We illustrate the method on different examples from banking and insurance. We focus on value-at-risk and expected shortfall, but the approach also works for other risk measures.

Keywords: asset-liability risk; risk capital; solvency calculation; value-at-risk; expected shortfall; neural networks; importance sampling

1. Introduction

Different financial risk management problems require an assessment of the risk of a portfolio of assets and liabilities over a given time period. Banks, mutual funds and hedge funds usually calculate value-at-risk numbers for different risks over a day, a week or longer time periods. Under Solvency II, insurance companies have to calculate one-year value-at-risk, whereas the Swiss Solvency Test demands a computation of one-year expected shortfall. A determination of these risk figures requires a conditional valuation of the portfolio given the state of the world at a later time, called risk horizon. This is particularly challenging if the portfolio contains structured products or complicated insurance policies which do not admit closed form valuation formulas. In theory, the problem can be approached with nested simulation, which is a two-stage procedure. In the outer stage, different scenarios are generated to model how the world could evolve until the risk horizon, whereas in the inner stage, cash flows occurring after the risk horizon are simulated to estimate the value of the portfolio conditionally on each scenario; see, e.g., Lee (1998), Glynn and Lee (2003), Gordy and Juneja (2008), Broadie et al. (2011) or Bauer et al. (2012). While nested simulation can be shown to converge for increasing sample sizes, it is often too time-consuming to be useful in practical applications. A more pragmatic alternative, usually used for short risk horizons, is the delta-gamma method, which approximates the portfolio loss with a second order Taylor polynomial; see, e.g., Rouvinez (1997), Britten-Jones and Schaefer (1999) or Duffie and Pan (2001). If first and second derivatives of the portfolio with respect to the underlying risk factors are accessible, the method is computationally efficient, but its accuracy depends on how well the second order Taylor polynomial approximates the true portfolio loss. Similarly, the replicating portfolio approach approximates future cashflows with a portfolio of liquid instruments that can be priced efficiently; see e.g., Wüthrich (2016), Pelsser and Schweizer (2016), Natolski and Werner (2017) or Cambou and Filipović (2018). Building on work on American option valuation (see e.g., Carriere 1996; Longstaff and Schwartz 2001; Tsitsiklis and Van Roy 2001), Broadie et al. (2015) as well as Ha and Bauer (2019) have proposed to regress future cash flows on finitely many basis functions depending on state variables known at the risk horizon. This gives good results in a number of applications. However, typically, for it to work well, the basis functions have to be chosen well-adapted to the problem.

In this paper we use a neural network approach to approximate the value of the portfolio at the risk horizon. Since our goal is to estimate tail risk measures such as value-at-risk and expected shortfall, we employ importance sampling when training the networks and estimating the risk figures. In addition, we try different regularization techniques and test the adequacy of the neural network approximations using the defining property of the conditional expectation. Neural networks have also been used for the calculation of solvency capital requirements by Hejazi and Jackson (2017), Fiore et al. (2018) and Castellani et al. (2019). But they all exclusively focus on value-at-risk, do not make use of importance sampling and apply neural networks in a slightly different way. Hejazi and Jackson (2017) develop a neural network interpolation scheme within a nested simulation framework. Fiore et al. (2018) and Castellani et al. (2019) both use reduced-size nested Monte Carlo to generate training samples for the calibration of the neural network. Here, we directly regress future cash-flows on a neural network without producing training samples.

The remainder of the paper is organized as follows. In Section 2, we set up the underlying risk model, introduce the importance sampling distributions we use in the implementation of our method and recall how value-at-risk and expected shortfall can be estimated from a finite sample of simulated losses. Section 3 discusses the training, validation and testing of our network approximation of the conditional valuation functional. In Section 4 we illustrate our approach on three typical risk calculation problems: the calculation of risk capital for a single put option, a portfolio of different call and put options and a variable annuity contract with guaranteed minimum income benefit. Section 5 concludes the paper.

2. Asset-Liability Risk

We denote the current time by 0 and are interested in the value of a portfolio of assets and liabilities at a given risk horizon $\tau > 0$. Suppose all relevant events taking place until time τ are described by a *d*-dimensional random vector $X = (X_1, \ldots, X_d)$ defined on a measurable space (Ω, \mathcal{F}) that is equipped with two equivalent probability measures, \mathbb{P} and \mathbb{Q} . We think of \mathbb{P} as the real-world probability measure and use it for risk measurement. \mathbb{Q} is a risk-neutral probability measure specifying the time- τ value of the portfolio through

$$V = v(X) + \mathbb{E}^{\mathbb{Q}}\left[\sum_{i=1}^{I} \frac{N_{\tau}}{N_{t_i}} C_{t_i} \mid X\right],$$

where v is a measurable function from \mathbb{R}^d to \mathbb{R} describing the part of the portfolio whose value is directly given by the underlying risk factors X_1, \ldots, X_d ; C_{t_1}, \ldots, C_{t_I} are random cash flows occurring at times $\tau < t_1 < \cdots < t_I$; and $N_{\tau}, N_{t_1}, \ldots, N_{t_I}$ model the evolution of a numeraire process used for discounting.

Our goal is to compute $\rho(L)$ for a risk measure ρ and the time τ net liability L = -V, which can be written as

$$L = \mathbb{E}^{\mathbb{Q}}[Y \mid X]$$

for

$$Y = -v(X) - \sum_{i=1}^{I} \frac{N_{\tau}}{N_{t_i}} C_{t_i}.$$

Our approach works for a wide class of risk measures ρ . But for the sake of concreteness, we concentrate on value-at-risk and expected shortfall. We follow the convention of McNeil et al. (2015) and define value-at-risk at a level $\alpha \in (0, 1)$, such as 0.95, 0.99 or 0.995, as the left α -quantile

$$\operatorname{VaR}_{\alpha}(L) := \min \left\{ x \in \mathbb{R} : \mathbb{P}[L \le x] \ge \alpha \right\} = \sup \left\{ x \in \mathbb{R} : \mathbb{P}[L \ge x] > 1 - \alpha \right\}$$
(1)

$$\mathrm{ES}_{\alpha}(L) := \frac{1}{1-\alpha} \int_{\alpha}^{1} \mathrm{VaR}_{u}(L) du.$$
⁽²⁾

There exist different definitions of VaR and ES in the literature. But if the distribution function F_L of *L* is continuous and strictly increasing, they are all equivalent¹.

2.1. Conditional Expectations as Minimizing Functions

Let $\mathbb{P} \otimes \mathbb{Q}$ be the probability measure on \mathcal{F} given by

$$\mathbb{P}\otimes \mathbb{Q}[A] := \int_{\mathbb{R}^d} \mathbb{Q}[A \mid X = x] \, \pi(dx),$$

where π is the distribution of X under \mathbb{P} and $\mathbb{Q}[A \mid X = x]$ is a regular conditional version of \mathbb{Q} given X. We assume that Y belongs to $L^2(\Omega, \mathcal{F}, \mathbb{P} \otimes \mathbb{Q})$. Then L is of the form L = l(X) for a measurable function $l \colon \mathbb{R}^d \to \mathbb{R}$ minimizing the mean squared distance

$$\mathbb{E}^{\mathbb{P}\otimes\mathbb{Q}}\left[\left(l(X)-Y\right)^{2}\right];$$
(3)

see, e.g., Bru and Heinich (1985). Note that *l* minimizes (3) if and only if

$$l(x) = \underset{u \in \mathbb{R}}{\arg\min} \int_{\mathbb{R}} (u - y)^2 \mathbb{Q}[Y \in dy \mid X = x] \quad \text{for } \pi\text{-almost all } x \in \mathbb{R}^d$$

where $\mathbb{Q}[Y \in dy \mid X = x]$ is a regular conditional \mathbb{Q} -distribution of Y given X. This shows that l is unique up to π -almost sure equality and can alternatively be characterized by

$$l(x) = \underset{u \in \mathbb{R}}{\operatorname{arg\,min}} \int_{\mathbb{R}} (u - y)^2 \mathbb{Q}[Y \in dy \mid X = x] \quad \text{for } \nu\text{-almost all } x \in \mathbb{R}^d$$

for any probability measure ν on \mathbb{R}^d that is equivalent to π . In particular, if \mathbb{P}^{ν} is the probability measure on $\sigma(X)$ under which X has distribution ν , and Y is in $L^2(\Omega, \mathcal{F}, \mathbb{P}^{\nu} \otimes \mathbb{Q})$, then l can be determined by minimizing

$$\mathbb{E}^{\mathbb{P}^{\nu} \otimes \mathbb{Q}} \left[\left(l(X) - Y \right)^2 \right] \tag{4}$$

instead of (3). Since we are going to approximate the expectation (4) with averages of Monte Carlo samples, this will give us some flexibility in simulating (X, Y).

2.2. Monte Carlo Estimation of Value-at-Risk and Expected Shortfall

Let X^1, \ldots, X^n be independent \mathbb{P} -simulations of *X*. By $X^{(1)}, \ldots, X^{(n)}$ we denote the same sample reordered so that

$$L^{(1)} = l(X^{(1)}) \ge \dots \ge L^{(n)} = l(X^{(n)}).$$

To obtain \mathbb{P} -simulation estimates of VaR_{α}(*L*) and ES_{α}(*L*), we apply the VaR and ES definitions (1)–(2) to the empirical measure $\frac{1}{n} \sum_{i=1}^{n} \delta_{I}(i)$. This yields

$$\widehat{\operatorname{VaR}}_{\alpha}(n) := L^{(j)} \quad \text{and} \quad \widehat{\operatorname{ES}}_{\alpha}(n) := \frac{1}{1-\alpha} \sum_{i=1}^{j-1} \frac{L^{(i)}}{n} + \left(1 - \frac{j-1}{(1-\alpha)n}\right) L^{(j)},\tag{5}$$

¹ More precisely, VaR_{α} is usually defined as an α - or $(1-\alpha)$ -quantile depending on whether it is applied to *L* or -L. So up to the sign convention, all VaR definitions coincide if F_L is strictly increasing. Similarly, different definitions of ES are equivalent if F_L is continuous.

where

$$j = \min \{i \in \{1, \dots, n\} : i/n > 1 - \alpha\};$$

see also McNeil et al. (2015). It is well known that if F_L is differentiable at $VaR_{\alpha}(L)$ with $F'_L(VaR_{\alpha}(L)) > 0$, then $\widehat{VaR}_{\alpha}(n)$ is an asymptotically normal consistent estimator of order $n^{-1/2}$; see, e.g., David and Nagaraja (2003). If in addition, *L* is square-integrable, the same is true for $\widehat{ES}_{\alpha}(n)$; see Zwingmann and Holzmann (2016).

2.3. Importance Sampling

Assume now that *X* is of the form X = h(Z) for a transformation $h \colon \mathbb{R}^k \to \mathbb{R}^d$ and a *k*-dimensional random vector *Z* with density $f \colon \mathbb{R}^k \to \mathbb{R}_+$. To give more weight to important outcomes, we introduce a second density *g* on \mathbb{R}^k satisfying g(z) > 0 for all $z \in \mathbb{R}^k$, where f(z) > 0. Let Z^g be a *k*-dimensional random vector with density *g* and $Z^{g,1}, \ldots, Z^{g,n}$ independent simulations of Z^g . We reorder the simulations such that

$$L^{g,(1)} = l \circ h(Z^{g,(1)}) \ge \dots \ge L^{g,(n)} = l \circ h(Z^{g,(n)})$$

and consider the random weights

$$w^{(i)} = \frac{f(Z^{g,(i)})}{ng(Z^{g,(i)})}.$$

Since $f(Z^g)/g(Z^g)$ is integrable, one obtains from the law of large numbers that for every threshold $x \in \mathbb{R}$,

$$\sum_{i=1}^{n} \mathbb{1}_{\{l \circ h(Z^{g,(i)}) \ge x\}} w^{(i)}$$
(6)

is an unbiased consistent estimator of the exceedance probability

$$\mathbb{P}[L \ge x] = \mathbb{P}[l \circ h(Z) \ge x] = \mathbb{E}\left[\mathbb{1}_{\{l \circ h(Z^g) \ge x\}} \frac{f(Z^g)}{g(Z^g)}\right]$$

If in addition, *g* can be chosen so that $f(Z^g)/g(Z^g)$ is square-integrable, it follows from the central limit theorem that (6) is asymptotically normal with a standard deviation of order $n^{-1/2}$. In any case, $\sum_{i=1}^{n} w^{(i)}$ converges to 1 and the random measure $\sum_{i=1}^{n} w^{(i)} \delta_{loh(Z^{g,(i)})}$ approximates the distribution of *L*. Accordingly, we adapt the \mathbb{P} -simulation estimators (5) by replacing i/n with $\sum_{m=1}^{i} w^{(m)}$. This yields the *g*-simulation estimates

$$\widehat{\text{VaR}}^{g}_{\alpha}(n) := L^{g,(j)} \quad \text{and} \quad \widehat{\text{ES}}^{g}_{\alpha}(n) := \frac{1}{1-\alpha} \sum_{i=1}^{j-1} w^{(i)} L^{g,(i)} + \left(1 - \frac{1}{1-\alpha} \sum_{i=1}^{j-1} w^{(i)}\right) L^{g,(j)} L^{g,(j)}$$

where

$$j = \min\left\{i \in \{1, \dots, n\} : \sum_{m=1}^{i} w^{(m)} > 1 - \alpha\right\}.$$

If the α -quantile l_{α} of $L = l \circ h(Z)$ were known, the exceedance probability $\mathbb{P}[L \ge l_{\alpha}]$ could be estimated by means of (6) with $x = l_{\alpha}$. To make the procedure efficient, one would try to find a density g on \mathbb{R}^k from which it is easy to sample and such that the variance

$$\operatorname{Var}\left(1_{\{l\circ h(Z^g)\geq l_{\alpha}\}}\frac{f(Z^g)}{g(Z^g)}\right) = \mathbb{E}\left[1_{\{l\circ h(Z^g)\geq l_{\alpha}\}}\frac{f(Z^g)^2}{g(Z^g)^2}\right] - \mathbb{E}\left[1_{\{l\circ h(Z^g)\geq l_{\alpha}\}}\frac{f(Z^g)}{g(Z^g)}\right]^2,$$

becomes as small as possible. Since

$$\mathbb{E}\left[\mathbf{1}_{\{l \circ h(Z^{g}) \ge l_{\alpha}\}} \frac{f(Z^{g})}{g(Z^{g})}\right] = \mathbb{E}\left[\mathbf{1}_{\{l \circ h(Z) \ge l_{\alpha}\}}\right] = \mathbb{P}[L \ge l_{\alpha}]$$

does not depend on g, it can be seen that g is a good importance sampling (IS) density if

$$\mathbb{E}\left[\mathbf{1}_{\{l\circ h(Z^{g})\geq l_{\alpha}\}}\frac{f(Z^{g})^{2}}{g(Z^{g})^{2}}\right] = \mathbb{E}\left[\mathbf{1}_{\{l\circ h(Z)\geq l_{\alpha}\}}\frac{f(Z)}{g(Z)}\right]$$
(7)

is small. We use the same criterion as a basis to find a good IS density for estimating VaR_{α} and ES_{α} ; see, e.g., Glasserman (2003) for more background on importance sampling.

2.4. The Case of an Underlying Multivariate Normal Distribution

In many applications *X* can be modelled as a deterministic transformation of a random vector with a multivariate normal distribution. Examples are included in Section 4 below. In this case, it can be written as

$$X = u(AZ) \tag{8}$$

for a function $u: \mathbb{R}^p \to \mathbb{R}^d$, a $p \times k$ -matrix A and a k-dimensional random vector Z with a standard normal density f. To keep the problem tractable, we look for a suitable IS density g in the class of $N_k(m, I_k)$ -densities f_m with different mean vectors $m \in \mathbb{R}^k$ and covariance matrix equal to the k-dimensional identity matrix I_k . To determine a good choice of m, let $v \in \mathbb{R}^p$ be the vector with components

$$v_i = \begin{cases} 1 & \text{if } l \circ u \text{ is increasing in } y_i \\ -1 & \text{if } l \circ u \text{ is decreasing in } y_i \\ 0 & \text{else.} \end{cases}$$

If $A^T v = 0$, we choose m = 0. Otherwise, $v^T AZ / ||A^T v||_2$ is standard normal. Denote its α -quantile by z_{α} . Then Z falls into the region

$$D = \left\{ z \in \mathbb{R}^k : v^T A z / \|A^T v\|_2 \ge z_\alpha \right\}$$

with probability $1 - \alpha$, and $l \circ u(Az)$ tends to be large for $z \in D$. We choose *m* as the maximizer of f(z) subject to $z \in D$, which leads to

$$m = \frac{A^T v}{\|A^T v\|_2} z_\alpha.$$

It is easy to see that this yields

$$f(z) < f_m(z)$$
 for all $z \in D$.

So, if *D* is sufficiently similar to the region $\{z \in \mathbb{R}^k : l \circ u(Az) \ge l_{\alpha}\}$, it can be seen from (7) that this choice of IS density will yield a reduction in variance.

3. Neural Network Approximation

Usually, the distribution of the risk factor vector X is assumed to be known. For instance, in all our examples in Section 4, they are of the form (8). On the other hand, in many real-world applications, there is no closed form expression for the loss function $l: \mathbb{R}^d \to \mathbb{R}$ mapping X to L. Therefore, we approximate L = l(X) with $l^{\theta}(X)$ for a neural network $l^{\theta}: \mathbb{R}^d \to \mathbb{R}$; see e.g., Goodfellow et al. (2016). We concentrate on feedforward neural networks of the form

$$l^{\theta} = \psi \circ a_{J}^{\theta} \circ \varphi \circ a_{J}^{\theta} \circ \cdots \circ \varphi \circ a_{1}^{\theta},$$

where

- $q_0 = d$, $q_I = 1$, and q_1, \ldots, q_{I-1} are the numbers of nodes in the hidden layers $1, \ldots, J-1$;
- a_j^{θ} are affine functions of the form $a_j^{\theta}(x) = A_j x + b_j$ for matrices $A_j \in \mathbb{R}^{q_j \times q_{j-1}}$ and vectors $b_j \in \mathbb{R}^{q_j}$, for j = 1, ..., J;
- φ is a non-linear activation function used in the hidden layers and applied component-wise. In the examples in Section 4 we choose $\varphi = \tanh$;
- ψ is the final activation function. For a portfolio of assets and liabilities a natural choice is ψ = id. To keep the presentation simple, we will consider pure liability portfolios with loss L > 0 in all our examples below. Accordingly, we choose ψ = exp.

The parameter vector θ consists of the components of the matrices A_j and vectors b_j , j = 1, ..., J. So it lives in \mathbb{R}^q for $q = \sum_{j=1}^J q_j(q_{j-1} + 1)$. It remains to determine the architecture of the network (that is, J and $q_1, ..., q_{J-1}$) and to calibrate $\theta \in \mathbb{R}^q$. Then VaR and ES figures can be estimated as described in Section 2 by simulating $l^{\theta}(X)$.

3.1. Training and Validation

In a first step we take the network architecture $(J \text{ and } q_1, \ldots, q_{J-1})$ as given and try to find a minimizer of $\theta \mapsto \mathbb{E}[(l^{\theta}(X) - Y)^2]$, where the expectation is either with respect to $\mathbb{P} \otimes \mathbb{Q}$ or $\mathbb{P}^{\nu} \otimes \mathbb{Q}$ for an IS distribution ν on \mathbb{R}^d . To do that we simulate realizations (X^m, Y^m) , $m = 1, \ldots, M_1 + M_2$, of (X, Y) under the corresponding distribution. The first M_1 simulations are used for training and the other M_2 for validation. More precisely, we employ a stochastic gradient descent method to minimize the Monte Carlo approximation based on the training samples

$$\frac{1}{M_1} \sum_{m=1}^{M_1} \left(l^{\theta}(X^m) - Y^m \right)^2 \tag{9}$$

of $\mathbb{E}[(l^{\theta}(X) - Y)]$. At the same time we use the validation samples to check whether

$$\frac{1}{M_2} \sum_{m=M_1+1}^{M_1+M_2} \left(l^{\theta}(X^m) - Y^m \right)^2 \tag{10}$$

is decreasing as we are updating θ .

3.1.1. Regularization through Tree Structures

If the number q of parameters is large, one needs to be careful not to overfit the neural network. For instance, in the extreme case, the network could be so flexible that it can bring (9) down to zero even in cases where the true conditional expectation $l(X) = \mathbb{E}^{\mathbb{Q}}[Y \mid X]$ is not equal to Y. To prevent this, one can generate the training samples by first simulating N_1 realizations X^i of X and then for every X^i , drawing N_2 simulations $Y^{i,j}$ from the conditional distribution of Y given X^i . In the simple example of Section 4.1, we chose $N_2 = 1$. In Sections 4.2 and 4.3 we used $N_2 = 5$.

3.1.2. Stochastic Gradient Descent

In principle, one can use any numerical method to minimize (9). But stochastic gradient descent methods have proven to work well for neural networks. We refer to Ruder (2016) for an overview of

different (stochastic) gradient descent algorithms. Here, we randomly² split the M_1 training samples into *b* mini-batches of size *B*. Then we update θ based on the θ -gradients of

$$\frac{1}{B}\sum_{m=(i-1)B+1}^{iB} \left(l^{\theta}(X^m) - Y^m \right)^2, \quad i = 1, \dots, b.$$

We use batch normalization and Adam updating with the default values from TensorFlow.

After b gradient steps, all of the training data have been used once and the first epoch is complete. For further epochs, we reshuffle the training data, form new mini-batches and perform b more gradient steps. The procedure is repeated until the training error (9) stops to decrease or the validation error (10) starts to increase.

3.1.3. Initialization

We follow standard practice and initialize the parameters of the network randomly. The final operation of the network is

$$x \mapsto \psi(A_I x + b_I).$$

Since the network tries to approximate *Y*, and in all our examples below we use $\psi = \exp$, we initialize the last bias as

$$b_J^0 = \log\left(rac{1}{M_1}\sum_{m=1}^{M_1}Y^m
ight).$$

For the other parameters we use Xavier initialization; see Glorot and Bengio (2010).

3.2. Backtesting the Network Approximation

After having determined an approximate minimizer $\theta \in \mathbb{R}^q$ for a given network architecture, one can test the adequacy of the approximation $l^{\theta}(X)$ of the true conditional expectation $l(X) = \mathbb{E}^{\mathbb{Q}}[Y \mid X]$. The quality of the approximation depends on different aspects:

(i) *Generalization error*:

The true conditional expectation $\mathbb{E}^{\mathbb{Q}}[Y \mid X]$ is of the form l(X) for the unique³ measurable function $l: \mathbb{R}^d \to \mathbb{R}$ minimizing the mean squared distance $\mathbb{E}[(l(X) - Y)^2]$. To approximate l we choose a network architecture and try to find a $\theta \in \mathbb{R}^q$ that minimizes the empirical squared distance

$$\frac{1}{M_1} \sum_{m=1}^{M_1} \left(l^{\theta}(X^m) - Y^m \right)^2.$$
(11)

But if the samples (X^m, Y^m) , $m = 1, ..., M_1$, do not represent the distribution of (X, Y) well, (11) might not be a good approximation of the true expectation $\mathbb{E}\left[\left(l^{\theta}(X) - Y\right)^2\right]$.

(ii) Numerical minimization method:

The minimization of (11) usually is a complex problem, and one has to employ a numerical method to find an approximate solution θ . The quality of $l^{\theta}(X)$ will depend on the performance of the numerical method being used.

(iii) Network architecture:
 It is well known that feedforward neural networks with one hidden layer have the universal approximation property; see, e.g., Cybenko (1989); Hornik et al. (1989) or Leshno et al. (1993).

² If the training data is generated according to a tree structure as in Section 3.1.1, one can either group the simulations $(X^i, Y^{i,j}), i = 1, ..., N_1, j = 1, ..., N_2$, so that pairs with the same X^i -component stay together or not. In our implementations, both methods gave similar results.

³ More precisely, uniqueness holds if functions are identified that agree π -almost surely.

That is, they can approximate any continuous function uniformly on compacts to any degree of accuracy if the activation function is of a suitable form and the hidden layers contain sufficiently many nodes. As a consequence, $\mathbb{E}[(l^{\theta}(X) - l(X))^2]$ can be made arbitrarily small if the hidden layer is large enough and θ is chosen appropriately. However, we do not know in advance how many nodes we need. And moreover, feedforward neural networks with two or more hidden layers have shown to yield better results in different applications.

Since we simulate from an underlying model, we are able to choose the size M_1 of the training sample large and train extensively. In addition, for any given network architecture, we also evaluate the empirical squared distance (11) on the validation set (X^m, Y^m) , $m = M_1 + 1, ..., M_2$. So we suppose the generalization error is small and our numerical method finds a good approximate minimizer θ of (11). However, since we do not know whether a given network architecture is flexible enough to provide a good approximation to the true loss function l, we test for each trained network whether it satisfies the defining properties of a conditional expectation.

The loss function $l: \mathbb{R}^d \to \mathbb{R}$ is characterized by

$$\mathbb{E}[l(X)\xi(X)] = \mathbb{E}[Y\xi(X)]$$
(12)

for all measurable functions $\xi \colon \mathbb{R}^d \to \mathbb{R}$ such that $Y\xi(X)$ is integrable. Ideally, we would like l^{θ} to satisfy the same condition. But there will be an approximation error, and (12) cannot be checked for all measurable functions $\xi \colon \mathbb{R}^d \to \mathbb{R}$ satisfying the integrability condition. Therefore, we select finitely many measurable subsets $B_i \subseteq \mathbb{R}^d$, i = 1, ..., I. Then we generate M_3 more samples (X^m, Y^m) of (X, Y) and test whether the differences

(a)
$$\sum_{m=M_2+1}^{M_1+M_2+M_3} \left(l^{\theta}(X^m) - Y^m \right) / M_3$$

(b)
$$\sum_{m=M_2+1}^{M_1+\bar{M}_2+M_3} \left(l^{\theta}(X^m) - Y^m \right) l^{\theta}(X^m) / M_3$$

(c)
$$\sum_{m=M_2+1}^{M_1+M_2+M_3} (l^{\theta}(X^m) - Y^m) 1_{B_i}(X^m) / M_3$$

are sufficiently close to zero. If this is not the case, we change the network architecture and train again.

4. Examples

As examples, we study three different risk assessment problems from banking and insurance. For comparison we generated realizations (X^m, Y^m) of (X, Y) under $\mathbb{P} \otimes \mathbb{Q}$ as well as $\mathbb{P}^{\nu} \otimes \mathbb{Q}$, where ν is the IS distribution on \mathbb{R}^d obtained by changing the distribution of X. In all our examples, X has a transformed normal distribution as in Section 2.4. In each case we used 1.5 million simulations with mini-batches of size 10,000 for training, 500,000 simulations for validation and 500,000 for backtesting. After training and testing the network, we simulated another 500,000 realizations of X, once under \mathbb{P} and then under \mathbb{P}^{ν} , to estimate VaR_{99.5%}(L) and ES_{99%}(L).

We implemented the algorithms in Python. To train the networks we used the TensorFlow package.

4.1. Single Put Option

As a first example, we consider a liability consisting of a single put option with strike price K = 100and maturity T = 1/3 on an underlying asset starting from $s_0 = 100$ and evolving according to

$$dS_t = \mu S_t dt + \sigma S_t dW_t^{\mathbb{P}} = rS_t dt + \sigma S_t dW_t^{\mathbb{Q}}$$

for an interest rate r = 1%, a drift $\mu = 5\%$, a volatility $\sigma = 20\%$, a \mathbb{P} -Brownian motion $W^{\mathbb{P}}$ and a \mathbb{Q} -Brownian motion $W^{\mathbb{Q}}$. As risk horizon we choose $\tau = 1/52$. The time τ -value of this liability is

$$L = e^{-r(T-\tau)} \mathbb{E}^{\mathbb{Q}} \left[(K - S_T)^+ \mid S_\tau \right].$$

Using Itô's formula, one can write

$$S_{\tau} = s_0 \exp\left((\mu - \sigma^2/2)\tau + \sigma\sqrt{\tau}Z\right) \quad \text{and} \quad S_T = S_{\tau} \exp\left((r - \sigma^2/2)(T - \tau) + \sigma\sqrt{T - \tau}V\right),$$

where *Z* and *V* are two independent standard normal random variables under $\mathbb{P} \otimes \mathbb{Q}$. It is well-known that *L* is of the form $P(S_{\tau}, r, \sigma, T - \tau)$, where *P* is the Black–Scholes formula for a put option. This allows us to calculate reference values for VaR_{α}(*L*) and ES_{α}(*L*).

To train the neural network approximation of *L*, we simulated realizations of the pair (X, Y) for $X = S_{\tau}$ and $Y = e^{-r(T-\tau)}(K - S_T)^+$. For comparison, we first simulated according to $\mathbb{P} \otimes \mathbb{Q}$ and then according to $\mathbb{P}^{\nu} \otimes \mathbb{Q}$ for an IS distribution \mathbb{P}^{ν} . Clearly, *L* is decreasing in *Z*. Therefore, we chose \mathbb{P}^{ν} so as to make *Z* normally distributed with mean equal to the $1 - \alpha$ -quantile of a standard normal and variance 1. Since $x \mapsto P(x, r, \sigma, T - \tau)$ is a simple one-dimensional function, we selected a network with a single hidden layer containing 5 nodes. This is sufficient to approximate *P* and faster to train than a network with more hidden layers. We did not use the tree structure of Section 3.1.1 for training (that is, $N_2 = 1$) and trained the network over 40 epochs.

It can be seen in Figure 1 that the empirical squared distance on both, the training and validation data, decreases with a very similar decaying pattern. This provides a first validation of our approximation procedure.



Figure 1. Empirical squared distance without IS during training (left) and on the validation data (right).

Figure 2 shows the empirical evaluation of (a) and (b) of Section 3.2 on the test data after each training epoch.



Figure 2. Empirical evaluation of (a) (left) and (b) (right) of Section 3.2.

Similarly, Figure 3 illustrates the empirical evaluation of (c) of Section 3.2 on the test data after each training epoch for the sets

 $B_1 = \{x \in \mathbb{R} : x < s_{40\%}\}$ and $B_2 = \{x \in \mathbb{R} : x > s_{70\%}\}$,

where s_{β} denotes the β -quantile of $X = S_{\tau}$.



Figure 3. Empirical evaluation of (c) of Section 3.2 for the sets B_1 (left) and B_2 (right).

Training, validation and testing with IS worked similarly.

Once the network has been trained and tested, one can estimate VaR and ES numbers. Figure 4 shows our results for increasing sample sizes. The left panel shows our estimate of VaR_{99.5%}(*L*) without and with IS. Plugging the 0.5%-quantile of S_{τ} into the Black–Scholes formula gives a reference value of 8.3356. Our method yielded 8.3358 without and 8.3424 with importance sampling. The right panel shows our results for ES_{99%}(*L*). Transforming simulations of S_{τ} with the Black–Scholes formula and using the empirical ES estimate (5) resulted in a reference value of 8.509. Without importance sampling, the neural network learned a value of 8.456 versus 8.478 with importance sampling. It can be seen that in both cases, IS made the method more efficient. It has to be noted that for increasing sample sizes, the VaR and ES estimates converge to the corresponding risk figures in the neural network model, which are not exactly equal to their analogs in the correct model. But it can be seen that the blue lines are close to their final values after very few simulations.



Figure 4. Convergence of the empirical 99.5%-VaR (**left**) and 99%-ES (**right**) without IS (orange) and with IS (blue) compared to the reference values obtained from the Black–Scholes formula (green).

4.2. Portfolio of Call and Put Options

In our second example we introduce a larger set of risk factors. We consider a portfolio of 20 short call and put options on different underlying assets with initial prices $s_0^i > 0$ and dynamics

$$dS_t^i = \mu_i S_t^i dt + \sigma_i S_t^i dW_t^{\mathbb{P},i} = rS_t^i dt + \sigma_i S_t^i dW_t^{\mathbb{Q},i}$$

for \mathbb{P} -Brownian motions $W^{\mathbb{P},i}$ and \mathbb{Q} -Brownian motions $W^{\mathbb{Q},i}$, i = 1, ..., 20 such that $(W^{\mathbb{P},1}, ..., W^{\mathbb{P},20})$ is a multivariate Gaussian process under \mathbb{P} with an instantaneous correlation of 30% between different components and $(W^{\mathbb{Q},1}, ..., W^{\mathbb{Q},20})$ is a multivariate Gaussian process under \mathbb{Q} , also with instantaneous correlation of 30% between different components. We set $s_0^i = 100$ for all i = 1, ..., 20 and r = 1%. The drift and volatility parameters are assumed to be $\mu_i = \mu_{10+i} = (3 + i/2)\%$ and $\sigma_i = \sigma_{10+i} = (15 + i)\%$, i = 1, ..., 10. As in the first example, we choose a maturity of T = 1/3 and a risk horizon of

 $\tau = 1/52$. We assume all options have the same strike price K = 100. Then the time- τ value of the liability is

$$L = e^{-r(T-\tau)} \mathbb{E}\left[\sum_{i=1}^{10} (S_T^i - K)^+ + \sum_{i=11}^{20} (K - S_T^i)^+ \middle| X\right],$$

where *X* is the vector $(S_{\tau}^1, \ldots, S_{\tau}^{20})$.

In this example we trained a neural network with two hidden layers containing 15 nodes each. We first simulated according to $\mathbb{P} \otimes \mathbb{Q}$ and trained for 100 epochs. Figure 5 shows the decay of the empirical squared distance on the training and validation data set.



Figure 5. Empirical squared distance without IS during training (left) and on the validation data (right).

After training and testing the network under $\mathbb{P} \otimes \mathbb{Q}$, we did the same under $\mathbb{P}^{\nu} \otimes \mathbb{Q}$ for the IS distributions ν resulting from the procedure of Section 2.4 for $\alpha = 99.5\%$ (for VaR) and 99% (for ES). The two plots in Figure 6 show the empirical evaluations of (a) and (b) on the test data under the IS measure $\mathbb{P}^{\nu} \otimes \mathbb{Q}$ corresponding to $\alpha = 99\%$.



Figure 6. Empirical evaluation of (a) (**left**) and (b) (**right**) of Section 3.2 under $\mathbb{P}^{\nu} \otimes \mathbb{Q}$ corresponding to $\alpha = 99\%$.

As an additional test, we consider the two sets

$$B_1 = \{x \in \mathbb{R}^{20} : x_i > s_{20\%}^i \text{ and } x_{10+i} < s_{80\%}^{10+i} \text{ for } i = 1, 2, 3\}$$

and

$$B_2 = \{x \in \mathbb{R}^{20} : x_i < s_{80\%}^i \text{ and } x_{10+i} > s_{20\%}^{10+i} \text{ for } i = 1, 2, 3\},\$$

where s^i_{β} is the β -quantile of S^i_{τ} under $\mathbb{P}^{\nu} \otimes \mathbb{Q}$, and evaluate (c) on the test data generated under the IS distribution $\mathbb{P}^{\nu} \otimes \mathbb{Q}$ corresponding to $\alpha = 99\%$. The results are depicted in Figure 7.



Figure 7. Empirical evaluation of (c) of Section 3.2 for the sets B_1 (**left**) and B_2 (**right**) under $\mathbb{P}^{\nu} \otimes \mathbb{Q}$ corresponding to $\alpha = 99\%$.

After training and testing, we generated simulations of X to estimate VaR_{99,5%}(L) and ES_{99%}(L). The convergence for increasing sample sizes is shown in Figure 8. The reference values, 104.92 for 99.5%-VaR and 105.59 for 99%-ES, were obtained from the empirical estimates (5) by simulating S_{τ} and using the Black–Scholes formula for each of the 20 options. The neural network estimates of 99.5%-VaR were 104.56 without and 104.48 with IS. Those of 99%-ES were 105.03 without and 104.65 with IS. In both cases, IS made the procedure more efficient.



Figure 8. Convergence of the empirical 99.5%-VaR (**left**) and 99%-ES (**right**) without IS (orange) and with IS (blue) compared to the reference values obtained from the Black–Scholes formula (green).

4.3. Variable Annuity with GMIB

As a third example we study a variable annuity (VA) with guaranteed minimum income benefit (GMIB). We consider the contract analyzed by Ha and Bauer (2019) using polynomial regression. At time 0 the contract is sold to an x-year old policyholder. If she is still alive at maturity T, she can choose between the balance of an underlying investment account and a lifetime annuity. Therefore, in case of survival, the time-T value of the contract is

$$\max\left\{S_T, b a_{x+T}(T)\right\}$$
,

where S_T is the account value, *b* a guaranteed rate and $a_{x+T}(T)$ the time-*T* value of an annuity paying one unit of currency to a (x + T)-year old policyholder at times T + 1, T + 2, ... for as long as the person lives.

The contract is exposed to three types of risk: investment risk, interest rate risk and mortality risk. We suppose the log-account value $q_t = \log(S_t)$, the interest rate r_t and the mortality rate μ_{x+t} of our policyholder start from known constants q_0 , r_0 , μ_x and for $x + t \le 120$, have \mathbb{P} -dynamics

$$dq_t = \left(m - \frac{1}{2}\sigma_S^2\right)dt + \sigma_S dW_t^{\mathbb{P},S},$$
$$dr_t = \zeta(\gamma - r_t)dt + \sigma_r dW_t^{\mathbb{P},r},$$
$$d\mu_{x+t} = \kappa \mu_{x+t} dt + \sigma_\mu dW_t^{\mathbb{P},\mu},$$

for given parameters $m, \zeta, \gamma, \kappa, \sigma_S, \sigma_r, \sigma_\mu$ and \mathbb{P} -Brownian motions $W^{\mathbb{P},S}$, $W^{\mathbb{P},r}$ and $W^{\mathbb{P},\mu}$ forming a three-dimensional Gaussian process with instantaneous correlations ρ_{12} , ρ_{13} and ρ_{23} . We assume that our policyholder does not live longer than 120 years. Therefore, we set $\mu_{x+t} \equiv \infty$ for x + t > 120. The dynamics for $x + t \le 120$ under the risk-neutral probability \mathbb{Q} are assumed to be

$$dq_t = \left(r_t - \frac{1}{2}\sigma_S^2\right)dt + \sigma_S dW_t^{\mathbb{Q},S}$$
$$dr_t = \zeta(\bar{\gamma} - r_t)dt + \sigma_r dW_t^{\mathbb{Q},r},$$
$$d\mu_{x+t} = \kappa \mu_{x+t}dt + \sigma_\mu dW_t^{\mathbb{Q},\mu},$$

where $W^{\mathbb{Q},S}$, $W^{\mathbb{Q},r}$, $W^{\mathbb{Q},\mu}$ are \mathbb{Q} -Brownian motions constituting a three-dimensional Gaussian process with the same instantaneous correlations as the corresponding \mathbb{P} -Brownian motions. As Ha and Bauer (2019), we assume there is no risk premium for mortality and a constant risk premium λ for interest rate risk, such that $\bar{\gamma} = \gamma - \lambda \sigma_r / \zeta$. Provided that the policyholder is still alive at the risk horizon $\tau < T$, the value of the contract at that time is

$$L = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_{\tau}^{T} r_s + \mu_{x+s} ds} \max\left\{e^{q_T}, b \, a_{x+T}(T)\right\} \mid X\right],$$

where we denote $X = (q_{\tau}, r_{\tau}, \mu_{x+\tau})$. Discounting with $r_s + \mu_{x+s}$ takes into account that the policyholder might die between τ and T. On the other hand, a possible death time between 0 and τ is not considered. This results in a conservative estimate of the capital requirement for the issuer of the contract. Alternatively, one could model the loss as $I_A L$, where $A \subseteq \Omega$ is the event that the policyholder survives until time τ .

We follow Ha and Bauer (2019) and set x = 55, $\tau = 1$, T = 15, b = 10.792, $q_0 = 4.605$, m = 5%, $\sigma_S = 18\%$, $r_0 = 2.5\%$, $\zeta = 25\%$, $\gamma = 2\%$, $\sigma_r = 1\%$, $\lambda = 2\%$, $\mu_x = 1\%$, $\kappa = 7\%$, $\sigma_\mu = 0.12\%$, $\rho_{12} = -30\%$, $\rho_{13} = 6\%$, $\rho_{23} = -4\%$. Then

$$a_{x+T}(T) = \sum_{k=1}^{50} {}_k E_{x+T}(T),$$

where

$$_{k}E_{x+t}(t) = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_{t}^{t+k}r_{s}+\mu_{x+s}ds} \mid r_{t}, \mu_{x+t}\right]$$

is the time-*t* value of a pure endowment contract with maturity t + k. Since *r* and μ are affine, one has

$$_k E_{x+t}(t) = F(t,k,r_t,\mu_{x+t})$$

for the function $F(t, k, r_t, \mu_{x+t}) = A(t, t+k)e^{-B_r(t,t+k)r_t - B_\mu(t,t+k)\mu_{x+t}}$, with *A*, *B*_r, and *B*_µ as given in the Appendix of Ha and Bauer (2019). Moreover, one can write

$$L = {}_{T-\tau}E_{x+\tau}(\tau)\mathbb{E}^{\mathbb{Q}_E}[\max\left\{e^{q_T}, b\,a_{x+T}(T)\right\} \mid X]$$

for the probability measure \mathbb{Q}_E given by

$$\frac{d\mathbb{Q}_E}{d\mathbb{Q}} = \frac{\exp\left(-\int_0^T r_s + \mu_{x+s} ds\right)}{\mathbb{E}^{\mathbb{Q}} \exp\left(-\int_0^T r_s + \mu_{x+s} ds\right)}.$$

Under \mathbb{P} , $X = (q_{\tau}, r_{\tau}, \mu_{\tau})$ is a three-dimensional normal vector, and the conditional \mathbb{Q}_E -distribution of (q_T, r_T, μ_{x+T}) given X is normal too (the precise form of these distributions is given in the Appendix of Ha and Bauer (2019)). This makes it possible to efficiently simulate (X, Y) under $\mathbb{P} \otimes \mathbb{Q}_E$, where

$$Y = F(\tau, T - \tau, r_{\tau}, \mu_{x+\tau}) \max \left\{ e^{q_T}, b \sum_{k=1}^{50} F(T, k, r_T, \mu_{x+T}) \right\}.$$

To approximate L we chose a network with two hidden layers containing 4 nodes each and trained it for 40 epochs. Figure 9 shows the empirical squared distance without IS on the training and validation data set.



Figure 9. Empirical squared distance without IS during training (left) and on the validation data (right).

The panels in Figure 10 illustrate the empirical evaluations of the test criteria (a) and (b) from Section 3.2.



Figure 10. Empirical evaluation of (a) (**left**) and (b) (**right**) of Section 3.2 under $\mathbb{P} \otimes \mathbb{Q}_E$.

To test (c) from Section 3.2 we considered the sets

 $B_1 = \left\{ x \in \mathbb{R}^3 : x_1 > q_{70\%} \text{ and } x_2 < r_{30\%} \right\} \text{ and } B_2 = \left\{ x \in \mathbb{R}^3 : x_1 < q_{30\%} \text{ and } x_2 > r_{70\%} \right\}$ where q_β and r_β denote the β -quantiles of q_τ and r_τ under $\mathbb{P} \otimes \mathbb{Q}_E$; see Figure 11.



Figure 11. Empirical evaluation of (c) of Section 3.2 for the sets B_1 (left) and B_2 (right) under $\mathbb{P} \otimes \mathbb{Q}_E$.

We also trained and tested under $\mathbb{P}^{\nu} \otimes \mathbb{Q}_E$ for an IS distribution ν on \mathbb{R}^d . Our results are reported in Figure 12. Our estimate of VaR_{99.5%}(*L*) was 138.64 without and 138.52 with IS. In comparison, the VaR estimate obtained by Ha and Bauer (2019) using 37 monomials and 40 million simulations is 139.74. Our estimates of ES_{99%}(*L*) came out as 141.12 without and 142.12 with IS. There exist no reference values for this case.



Figure 12. Convergence of the empirical 99.5%-VaR (**left**) and 99%-ES (**right**) without IS (orange) and with IS (blue) compared to the reference value from Ha and Bauer (2019) (green).

5. Conclusions

In this paper we have developed a deep learning method for assessing the risk of an asset-liability portfolio over a given time horizon. It first computes a neural network approximation of the portfolio value at the risk horizon. Then the approximation is used to estimate a risk measure, such as value-at-risk or expected shortfall, from Monte Carlo simulations. We have investigated how to choose the architecture of the network, how to learn the network parameters under a suitable importance sampling distribution and how to test the adequacy of the network approximation. We have illustrated the approach by computing value-at-risk and expected shortfall in three typical risk assessment problems from banking and insurance. In all cases the approach has worked efficiently and produced accurate results.

Author Contributions: P.C., J.E., and M.V.W. have contributed equally to this work. All authors have read and agreed to the published version of the manuscript.

Funding: As SCOR Fellow, John Ery thanks SCOR for financial support.

Conflicts of Interest: The authors declare no conflict of interest.

References

Bauer, Daniel, Andreas Reuss, and Daniela Singer. 2012. On the calculation of the solvency capital requirement based on nested simulations. *ASTIN Bulletin* 42: 453–99.

- Britten-Jones, Mark, and Stephen Schaefer. 1999. Non-linear Value-at-Risk. *Review of Finance, European Finance Association* 5: 155–80.
- Broadie, Mark, Yiping Du, and Ciamac Moallemi. 2011. Efficient risk estimation via nested sequential estimation. *Management Science* 57: 1171–94. [CrossRef]
- Broadie, Mark, Yiping Du, and Ciamac Moallemi. 2015. Risk estimation via regression. *Operations Research* 63: 1077–97. [CrossRef]
- Bru, Bernard, and Henri Heinich. 1985. Meilleures approximations et médianes conditionnelles. *Annales de l'I.H.P. Probabilités et Statistiques* 21: 197–224.
- Cambou, Mathieu, and Damir Filipović. 2018. Replicating portfolio approach to capital calculation. *Finance and Stochastics* 22: 181–203. [CrossRef]
- Carriere, Jacques F. 1996. Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* 19: 19–30. [CrossRef]
- Castellani, Gilberto, Ugo Fiore, Zelda Marino, Luca Passalacqua, Francesca Perla, Salvatore Scognamiglio, and Paolo Zanetti. 2019. *An Investigation of Machine Learning Approaches in the Solvency II Valuation Framework*. SSRN Preprint.
- Cybenko, G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems* 2: 303–14. [CrossRef]
- David, Herbert A., and Haikady N. Nagaraja. 2003. *Order Statistics*, 3rd ed. Wiley Series in Probability and Statistics. Hoboken: Wiley.
- Duffie, Darrell, and Jun Pan. 2001. Analytical Value-at-Risk with jumps and credit risk. *Finance and Stochastics* 5: 155–80. [CrossRef]
- Fiore, Ugo, Zelda Marino, Luca Passalacqua, Francesca Perla, Salvatore Scognamiglio, and Paolo Zanetti. 2018. Tuning a deep learning network for Solvency II: Preliminary results. *Mathematical and Statistical Methods for Actuarial Sciences and Finance: MAF 2018.* Berlin: Springer.
- Glasserman, Paul. 2003. *Monte Carlo Methods in Financial Engineering*. Stochastic Modelling and Applied Probability. Berlin: Springer.
- Glorot, Xavier, and Yoshua Bengio 2010. Understanding the difficulty of training deep feedforward neural networks. *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* 9: 249–56.
- Glynn, Peter, and Shing-Hoi. Lee. 2003. Computing the distribution function of a conditional expectation via Monte Carlo: Discrete conditioning spaces. *ACM Transactions on Modeling and Computer Simulation* 13: 238–58.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. Deep Learning. Cambridge: MIT Press.
- Gordy, Michael B., and Sandeep Juneja. 2008. Nested simulation in portfolio risk measurement. *Management Science* 56: 1833–48. [CrossRef]
- Ha, Hongjun, and Daniel Bauer. 2019. A least-squares Monte Carlo approach to the estimation of enterprise risk. Working Paper. Forthcoming.
- Hejazi, Seyed Amir, and Kenneth R. Jackson. 2017. Efficient valuation of SCR via a neural network approach. *Journal of Computational and Applied Mathematics* 313: 427–39. [CrossRef]
- Hornik, Kurt, Maxwell Stinchcombe, and Halbert White. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks* 2: 359–66. [CrossRef]
- Lee, Shing-Hoi. 1998. Monte Carlo Computation of Conditional Expectation Quantiles. Ph.D. thesis, Stanford University, Stanford, CA, USA.
- Leshno, Moshe, Vladimir Lin, Allan Pinkus, and Shimon Schocken. 1993. Multilayer feedforward networks with a non-polynomial activation function can approximate any function. *Neural Networks* 6: 861–67. [CrossRef]
- Longstaff, Francis, and Eduardo Schwartz. 2001. Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies* 14: 113–47. [CrossRef]
- McNeil, Alexander J., Rüdiger Frey, and Paul Embrechts. 2015. *Quantitative Risk Management: Concepts, Techniques and Tools*, rev. ed. Princeton: Princeton University Press.
- Natolski, Jan, and Ralf Werner. 2017. Mathematical analysis of replication by cashflow matching. *Risks* 5: 13. [CrossRef]
- Pelsser, Antoon, and Janina Schweizer. 2016. The difference between LSMC and replicating portfolio in insurance liability modeling. *European Actuarial Journal* 6: 441–94. [CrossRef] [PubMed]
- Rouvinez, Christophe. 1997. Going Greek with VaR. Risk Magazine 10: 57-65.
- Ruder, Sebastian. 2016. An overview of gradient descent optimization algorithms. arXiv arXiv:1609.04747.

Tsitsiklis, John, and Benjamin Van Roy. 2001. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* 12: 694–703. [CrossRef] [PubMed]

Wüthrich, Mario V. 2016. *Market-Consistent Actuarial Valuation*. EAA Series. Berlin: Springer. Zwingmann, Tobias, and Hajo Holzmann. 2016. Asymptotics for the expected shortfall. *arXiv* arXiv:1611.07222.



 \odot 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0/).