



HAL
open science

Backward Hedging for American Options with Transaction Costs

Ludovic Goudenège, Andrea Molent, Antonino Zanette

► **To cite this version:**

Ludovic Goudenège, Andrea Molent, Antonino Zanette. Backward Hedging for American Options with Transaction Costs. 2023. hal-04308637

HAL Id: hal-04308637

<https://hal.science/hal-04308637v1>

Preprint submitted on 27 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Backward Hedging for American Options with Transaction Costs

LUDOVIC GOUDENÈGE*

ANDREA MOLENT †

ANTONINO ZANETTE‡

Abstract

In this article, we introduce an algorithm called Backward Hedging, designed for hedging European and American options while considering transaction costs. The optimal strategy is determined by minimizing an appropriate loss function, which is based on either a risk measure or the mean squared error of the hedging strategy at maturity. The proposed algorithm moves backward in time, determining, for each time-step and different market states, the optimal hedging strategy that minimizes the loss function at the time the option is exercised, by assuming that the strategy used in the future for hedging the liability is the one determined at the previous steps of the algorithm. The approach avoids machine learning and instead relies on classic optimization techniques, Monte Carlo simulations, and interpolations on a grid. Comparisons with the Deep Hedging algorithm in various numerical experiments showcase the efficiency and accuracy of the proposed method.

Keywords: Hedging, Greeks, Transaction costs, American options, Heston model

*Fédération de Mathématiques de CentraleSupélec - CNRS FR3487, France - ludovic.goudenege@math.cnrs.fr

†Dipartimento di Scienze Economiche e Statistiche, Università degli Studi di Udine, Italy - andrea.molent@uniud.it

‡Dipartimento di Scienze Economiche e Statistiche, Università degli Studi di Udine, Italy - antonino.zanette@uniud.it

1 Introduction

The groundbreaking work of Black and Scholes [5], followed by Merton [25], enabled the objective determination of the price of a derivative product as the monetary amount required by an agent for the exact replication of its payoff. This definition assumes that exact replication is achievable, aiming to eliminate any potential risk from holding a short position on the derivative. To achieve this, the authors make several assumptions, the most apparent being the possibility of continuously rebalancing the hedging portfolio. This assumption, however, is not feasible in practice due to physical constraints, making it necessary to rebalance the portfolio at a set of discrete times. These considerations also apply to advanced models, which incorporate additional risk factors, such as Heston's stochastic volatility model [18]. The infeasibility of continuous-time hedging causes these models to transform into approximate versions of themselves, which are associated with incomplete markets: exact replication of derivative products is generally not possible.

In real-world applications, practitioners typically use continuous-time models to calculate the hedging strategy, despite the impossibility of continuous-time hedging. This is done to exploit the computational power of such models and, specifically, to derive hedging strategies. Furthermore, they include a suitable margin in the price of the derivative to account for the replication error due to discrete-time hedging based on a continuous-time model.

Another limitation to the use of standard theoretical models in a practical context is the presence of frictions, primarily the transaction costs that arise during the construction and rebalancing of the hedging portfolio. Naturally, these costs, while making the replication procedure more expensive on one hand, also discourage the rebalancing prescribed by the theoretical model on the other. Consequently, it is essential to strike a balance in order to minimize the overall cost of the strategy without significantly compromising its effectiveness.

The first studies related to a pricing model in a imperfect market were developed by Duffie and Richardson [14]. They introduce the mean-variance hedging problem in continuous time, when the hedger can only trade futures contracts on an asset which is positively correlated with the underlying and so they consider the problem of determining the optimal hedge. Schweizer [30, 31] extends the results of Duffie and Richardson and obtains the first ground-breaking result by solving the hedging problem for a general contingent claim when a mean-squared error loss function is considered. Subsequently, Föllmer [15, 16] studies theoretical properties of hedging strategies when the loss to minimize is given by the VaR or the CVaR at maturity, computed from the perspective of the hedger. Other authors focus on the development of numerical algorithms for computing optimal hedging strategies. For example, Potter et al. [27] develop an algorithm based on Monte Carlo simulations to determine a hedging strategy that minimizes the quadratic error of hedging between two rebalancing steps. Pochart and Bouchaud [26] continue this analysis, by considering minimization of a local risk function such as Var and CVaR, also including transaction costs through a first order approximation. The use of convex risk measures for option pricing has been studied also by other authors. For example, Xu [33] investigates from a theoretical perspective the evaluation of financial option by risk measure pricing rules. Also Klöppel and Schweizer [22] develop a theoretical investigation and find interesting results on time-consistency about families of conditional convex risk measures. To the same scope, some authors study asymptotic behaviour of prices and strategies. With this regard, Whalley and Wilmott [32] develop a model for pricing European options in the presence of rehedging transaction costs by means of asymptotic analysis. Kallsen and Muhle-Karbe [20] obtain asymptotic formulas for utility indifference prices and hedging strategies in the presence of small transaction costs. Similarly, Bouchard et al. [7] obtain explicit formulae for option hedging in a market with proportional transaction costs.

Recently, some authors have proposed to use modern reinforcement machine learning methods to compute

an optimal hedging strategy through neural networks. One of the main results in this sense is due to Buehler et al. [10]. They address the problem of hedging under conditions of incomplete markets with transaction costs by proposing an algorithm called Deep Hedging: the optimal strategy is determined by training a complex neural network so that it approximates an optimal hedging strategy in terms of an appropriate risk measure. Their approach is greek-free in that the hedging strategy is not determined from the greeks, but rather as the minimizer of a loss function, that is as a solution to a variational problem. The resulting algorithm for finding the optimal hedging strategy has many merits: it is model free, it can manage transaction costs, it can also include, in the feature set, additional information such as trading signals, news analytics or past hedging decisions. This new algorithm has enjoyed considerable success in the financial landscape and beyond. For example, Horvat et al [19] investigate the hedging of short positions on options under rough volatility models again by using the Deep Hedging algorithm. Furthermore, the Deep Hedging algorithm has been recently employed in the insurance field by Carbonneau [11], with the scope of hedging long term insurance products such as variable annuities.

Despite its merits, there are also some elements that limit the use of the Deep Hedging algorithm: in order to obtain reasonable results, the calculation time required is generally high as the proposed neural network optimises the whole hedging strategy over the entire historical series passed as input. Training such a complex neural network requires considerable numerical effort, and sometimes many iterations are required to converge to the optimal strategy. Moreover, neural networks behave as a black box, which may meet with resistance from regulators (see e.g. Lin [24]). The Deep Hedging algorithm, as it stands, is specifically designed for European options and it is not suitable for handling American options.

In the context of American options, we would like to mention the work by Becker et al. [4], which employs neural networks for pricing and hedging these options.

In this paper, motivated by the algorithm introduced by Buehler et al. [10], we propose an alternative approach that avoids the use of machine learning and relies on traditional optimization and interpolation techniques. We construct an appropriate discretization of the possible market states by using a grid of values, and our algorithm moves backward in time to compute the optimal strategy from each individual state. Consequently, we term this the *Backward Hedging* (BH) algorithm. The optimal strategy is determined by calculating the optimal quantities of hedging instruments to include in the portfolio, minimizing the desired loss function when applied to forward Monte Carlo simulations generated from the state under consideration. Our proposed method proves to be efficient and reliable, with easily interpretable results. We stress out that, because of lack of time-consistency for most of risk measures (see e.g. Boda and Filar [6], Cheridito and Stadje [12], Cui et al. [13]), the problem solved by the Backward Hedging algorithm is not equivalent to the global one, which is solved by the Deep Hedging algorithm, but it can be considered as an alternative problem that best reflects the behavior followed in practice by hedgers, as well as an approximation of the global problem. Furthermore, the use of a backward approach also allows for the valuation of American options. With respect to American options, it is important to note that the approach proposed by Becker et al. [4] differs from ours, as they rely on the exercise strategy derived from the continuous-time model, rather than adopting the worst-case strategy for the hedger, as we suggest in our proposal.

Our investigation includes numerical tests in the one-dimensional Heston model, primarily considered in [10]. The results obtained are consistent with those deduced from the Deep Hedging algorithm and are competitive in terms of quality, confirming the validity of our proposed procedure. Specifically, although the solution of the backward problem is usually not optimal for the global problem, numerical results show that the hedging strategy returned by our algorithm is effective and sometimes even a better minimizer for the global problem than the one obtained with the Deep Hedging algorithm. It is worth noting that Buehler et al. [10] demonstrate the applicability of the Deep Hedging algorithm in the context of a multidimensional

Heston model. Our proposed algorithm, however, cannot efficiently address this case due to the curse of dimensionality, which can only be solved by machine learning-based algorithms. Nevertheless, we believe that our approach may be useful for practitioners when considering low-dimensional problems.

The paper is organized as follows. Section 2 introduces the pricing framework of the Deep Hedging algorithm. Section 3 discuss the proposed new algorithm. Section 4 is devoted to the experimental results. Finally, the conclusions follow in Section 5.

2 Deep Hedging framework

The hedging model presented in this manuscript is developed within the framework discussed in the Buehler et al. paper [10]. The Deep Hedging algorithm is also used as a point of comparison for our algorithm; therefore, we begin our discussion by briefly presenting it. Although the algorithm is highly adaptable to different stochastic models, we present it with reference to Heston’s (one-dimensional) stochastic model, which is also studied in the original manuscript. It is worth noting that in the original version of the algorithm, the authors assume both the drift rate and the interest rate to be zero. Here, we present an updated version of the Deep Hedging algorithm, which works for non-zero drift and interest rate.

The Deep Hedging algorithm allows one to approximate optimal hedging strategies for managing a short position on a European option. The search for an optimal strategy is carried out by exploiting neural networks and it is based on the minimization of a certain loss function. The optimal strategy ultimately determines the option price as the amount of money required to make such a strategy acceptable according to a risk profile. This model makes it possible to determine the value of an option even when perfect hedging is not possible, also when trading costs are considered.

2.1 Settings

We consider a financial market with finite time horizon T and trading dates $0 = t_0 < t_1 < \dots < t_N = T$, with $t_n = n \cdot \Delta t$ for $n = 0, \dots, N$. Let S^1 be the stochastic process representing the price of a stock and V the associated variance process. These two stochastic processes are supposed to evolve, under a real-world probability \mathbb{P} , according to the Heston model, that is the following stochastic differential equations hold:

$$\begin{cases} S_t^1 = \mu S_t^1 dt + \sqrt{V_t} S_t^1 dB_t, & S_0^1 = s_0, \\ V_t = \alpha(b - V_t) dt + \sigma \sqrt{V_t} dW_t, & V_0 = v_0, \end{cases}$$

where B and W are one-dimensional Brownian motions with their correlation ρ_{BW} a real constant, just as $\mu, \alpha, b, \sigma, s_0$ and v_0 . In the original version of the manuscript [10], the authors consider a market with a zero drift and interest rate, that is $\mu = r = 0$, however, here, following Rémillard and Rubenthaler [29], we allow these values to be non-null. Unlike process S^1 , process V is not directly tradable, nonetheless we assume that an idealized variance swap is available in the market, and we denote its price by S^2 . In particular, the maturity of the variance swap is the same as the considered derivative, i.e. T , and if we define the filtration $\mathcal{F}_t := \sigma((S_s^1, V_s) : s \in [0, t])$, then

$$S_t^2 := e^{-r(T-t)} \cdot \mathbb{E}_{\mathbb{Q}} \left[\int_0^T V_s ds \mid \mathcal{F}_t \right], \quad t \in [0, T].$$

with \mathbb{Q} a risk-neutral probability. Moreover, if we assume the market price of risk for the variance process to be zero, one can prove that the idealized price of the variance swap writes down

$$S_t^2 = e^{-r(T-t)} \cdot \left[\int_0^t V_s \, ds + L(t, V_t) \right]$$

where the function L is defined by

$$L(t, v) = \frac{v - b}{\alpha} \left(1 - e^{-\alpha(T-t)} \right) + b(T - t).$$

Then, S^1 and S^2 are the prices of tradeable assets that can be used to set up a portfolio to hedge a short position on a derivative.

Now, let us consider a European option, whose payoff is a random variable \mathcal{F}_T measurable, denoted by $\varphi(S_T^1)$. We assume the position of an economic agent who sold the option, so $\varphi(S_T^1)$ is a liability, and we have to hedge such a short position.

2.2 Trading strategies

In order to hedge a liability $\varphi(S_T^1)$ at T , one trades S^1 and S^2 by using an \mathbb{R}^2 -valued \mathcal{F} -adapted stochastic process $\delta = (\delta_n)_{n=0, \dots, N-1}$ with $\delta_n = (\delta_n^1, \delta_n^2)$. Please observe that, δ_n^i denotes the agent's holdings of the i -th asset, after rebalancing the portfolio at time t_n . For simplicity of notation, we may also define $\delta_{-1}^i = \delta_N^i := 0$ for $i = 1, 2$, which implies an empty portfolio at beginning and full liquidation at maturity.

Let p_0 denote the cost for building the portfolio at t_0 , so p_0 is the amount of money required in payment by the buyer of the option. In a market without transaction costs the agent's global profit and loss PL_0^T at time T is thus given by

$$PL_0^T(\varphi(S_T^1), p_0, \delta) = -\varphi(S_T^1) + p_0 e^{rT} + (\delta \cdot S)_0^N,$$

where

$$(\delta \cdot S)_{n_1}^{n_2} := \sum_{k=n_1}^{n_2-1} \left[\delta_k^1 \left(S_{t_{k+1}}^1 - S_{t_k}^1 e^{r\Delta t} \right) + \delta_k^2 \left(S_{t_{k+1}}^2 - S_{t_k}^2 e^{r\Delta t} \right) \right] e^{(T-t_{k+1})r}$$

is the profit due to the hedging strategy, with respect to the time intervals $[t_k, t_{k+1}]$ for $k = n_1, \dots, n_2 - 1$, capitalised to maturity.

Trading costs can also be included in the model. We assume that any trading activity causes proportional costs as follows: if the hedger decides to buy or sell $\alpha \in \mathbb{R}$ units of S^i at time t_k , then he incurs in a transaction cost which amounts to $\varepsilon \cdot S_k^i \cdot |\alpha|$, with ε the proportional cost for trading 1 currency unit of asset. The capitalised to maturity cost of the trading strategy δ , with respect to trading times from t_{n_1} to t_{n_2} , is therefore given by

$$C_{n_1}^{n_2}(\delta) := \sum_{k=n_1}^{n_2} \varepsilon \left[S_{t_k}^1 |\delta_k^1 - \delta_{k-1}^1| + S_{t_k}^2 |\delta_k^2 - \delta_{k-1}^2| \right] e^{(T-t_k)r}.$$

We stress out that, according to the original paper, $C_0^N(\delta)$ includes the cost of setting up the portfolio at time 0 and the costs of divestment at time T . As a result, when transaction costs are applied, the agent's terminal portfolio value at T reads out

$$PL_0^T(\varphi(S_T^1), p_0, \delta) := -\varphi(S_T^1) + p_0 e^{rT} + (\delta \cdot S)_0^N - C_0^N(\delta).$$

Such a random variable can be considered as the error of the trading strategy, and one seeks to minimize it.

2.3 Hedging and pricing

The basic situation to be considered is given by a complete market with continuous-time trading, no transaction costs, and unconstrained hedging. In this particular case, for any liability $\varphi(S_T^1)$, there exists a unique replication strategy δ and a unique fair price $p_0 \in \mathbb{R}$ such that

$$-\varphi(S_T^1) + p_0 e^{rT} + (\delta \cdot S)_0^N = 0$$

holds \mathbb{P} -a.s.. More generally, the optimal hedging strategy δ and the option price p_0 can be defined as the solution of the following problem: computing $p_0 \in \mathbb{R}$ and $\delta \in \mathcal{H}$ such that

$$\mathbb{E} \left[\left(-\varphi(S_T^1) + p_0 e^{rT} + (\delta \cdot S)_0^N \right)^2 \right] = 0, \quad (2.1)$$

where \mathcal{H} stands for the set of all possible trading strategies.

In an incomplete market with frictions, perfect hedging is (usually) no more possible, so, it is no longer possible to obtain a strategy that generates an almost certainly zero error. Consequently, the new objective is to minimize the error of the hedging strategy, so, one considers the problem of minimizing the mean squared error of the strategy:

$$\inf_{p_0 \in \mathbb{R}} \inf_{\delta \in \mathcal{H}} \mathbb{E} \left[\left(-\varphi(S_T^1) + p_0 e^{rT} + (\delta \cdot S)_0^N - C_0^N(\delta) \right)^2 \right]. \quad (2.2)$$

The solution p_0 of (2.2) can be regarded as the fair hedging price, but it does not take into account the non-eliminable risk faced by the hedger. In this case, one has to specify a certain criterion that defines an acceptable “minimal price” for any liability, a price that includes compensation for the risk assumed by the hedger. Such a price, can be defined as the minimum amount of money to add to the hedging portfolio at the launch of the contract, so that, when the optimal strategy is employed, the overall position becomes acceptable at maturity. The criterion for identifying what is an acceptable position is defined by means of a convex risk measure ρ , such that $\rho(0) = 0$. For example, hereafter, we consider ρ as the α -level CVaR, that is $\rho(X) = \text{CVaR}_\alpha(X)$. The problem to solve is then computing

$$\inf_{\delta \in \mathcal{H}} \rho \left(-\varphi(S_T^1) + (\delta \cdot S)_0^N - C_0^N(\delta) \right), \quad (2.3)$$

so that the hedger’s price writes down

$$\hat{p}_0 = e^{-rT} \inf_{\delta \in \mathcal{H}} \rho \left(-\varphi(S_T^1) + (\delta \cdot S)_0^N - C_0^N(\delta) \right). \quad (2.4)$$

Let $\hat{\delta} \in \mathcal{H}$ be the optimal strategy, solution of problem (2.2) or (2.3), according to the evaluation approach under investigation. Please, observe that the optimal hedging strategy $\hat{\delta}_n^i$ at time t_n can be seen as the outputs of a function that depends on filtration at time t_n : $\hat{\delta}_n^i = \hat{\delta}_n^i(S_0^1, S_0^2, \dots, S_n^1, S_n^2)$. Accordingly, problems (2.2) and (2.3) can be interpreted as variational problems, which can be tackled by exploiting a reinforcement learning algorithm based on neural networks.

2.4 The Deep Hedging algorithm

Neural networks can be used to solve variational problems such as (2.2) and (2.3) by choosing a suitable loss function.

The first step consists in simulating the paths of the underlyings (S^1, S^2) at discrete times t_n . These simulations are used to train a semi-recurrent neural network: at each time-step t_n , for $n = 0, \dots, N - 1$

it takes as inputs the 4 values of $S_{t_n}^1$, $S_{t_n}^2$, δ_{n-1}^1 and δ_{n-1}^2 and it outputs δ_n^1 and δ_n^2 . The network is trained to minimize the loss function, which value is determined based on the simulations of the paths. At the end of this process, the network returns a hedging strategy, denoted by $\hat{\delta}$, that approximates the solution to the problems (2.2) or (2.3). Specifically, the weight matrices and biases of the neurons are the parameters of the neural network but also the initial values δ_0^1 and δ_0^2 are added among the parameters to be optimized. Now, let θ be the vector all the parameters of the neural network. So, θ is computed as a minimum point of a suitable loss function $\ell(\theta)$. While problem (2.3) is considered, the loss function is $\ell(\theta) = \rho(-\varphi(S_T^1) + (\delta \cdot S)_0^N - C_0^N(\delta))$. As far as problem (2.2) is considered, also p_0 is included among the parameters to be optimized and the loss function is set as $\ell(\theta) = \mathbb{E} \left[(-\varphi(S_T^1) + p_0 e^{rT} + (\delta \cdot S)_0^N - C_0^N)^2 \right]$. The neural network is then trained by using the ADAM algorithm (Kingma and Ba [21]). Finally, when considering problem (2.3), the trained neural network is tested on out-of-sample dataset to determine the price p_0 , so that the measurement of network performance is not altered by overfitting.

Once the neural network has been trained, it returns the following outputs: the option price p_0 and the initial values $\hat{\delta}_0^1$ and $\hat{\delta}_0^2$ as well as it can be employed to calculate delta values at each possible trading time-step. This means that the neural network can be used for hedging purposes: given a path $(S_{t_n}^1, S_{t_n}^2)_{n=0, \dots, N}$ or even a partial path $(S_{t_k}^1, S_{t_k}^2)_{k=0, \dots, n}$ (with $n \leq N$), the neural network returns its optimal values for $(\hat{\delta}_k^1, \hat{\delta}_k^2)_{k=0, \dots, n}$, so one can effectively exploit the Deep Hedging algorithm for hedging a short position.

3 Backward Hedging

The Deep Hedging algorithm relies on neural networks to achieve global optimisation of the hedging strategy. This means that the algorithm simultaneously optimises all the parameters θ that define the entire hedging strategy. This approach results in a single high-dimensional optimisation problem.

The algorithm we propose, on the other hand, solves a large number of optimisation problems in only two variables, that is (δ_n^1, δ_n^2) . Starting from a specific configuration of the market state, it determines the optimal delta values for that specific configuration. Thanks to linear interpolation, the hedging strategy is extended to the entire space of interest. We emphasise that, unlike other works already present in the literature (see e.g. Potter et al. [27]) that consider the loss between two trading times, the loss function to be optimised is evaluated at maturity, i.e. it considers the entire hedging strategy, from the considered time-step until maturity.

Finally, we would like to point out that the pseudocodes of the main versions of the Backward Hedging algorithm, discussed below, are included in the concluding Appendix A.

3.1 The backward approach

We begin our discussion by presenting the proposed hedging model for the special case $\varepsilon = 0$, that is no transaction costs and loss function based on the risk measure ρ , as in problem (2.3).

The hedging strategy we propose is determined by proceeding backward in time, starting from time t_{N-1} , to the initial instant $t=0$. At the generic hedging instant t_n , the amount of hedging instruments to be held in the portfolio δ_n^1 and δ_n^2 are determined by minimizing an appropriate measure of risk relative to the future hedging error, that is, between t_n and T . Specifically, the Backward Hedging strategy $\bar{\delta}$ at time t_n is defined as the solution of the following problem:

$$(\bar{\delta}_n^1, \bar{\delta}_n^2) = \arg \min_{(x,y) \in \mathbb{R}^2} \rho \left(-\varphi(S_T^1) + (\bar{\delta} \cdot S)_{n+1}^N + \left[x \left(S_{t_{n+1}}^1 - S_{t_n}^1 e^{r\Delta t} \right) + y \left(S_{t_{n+1}}^2 - S_{t_n}^2 e^{r\Delta t} \right) \right] e^{(T-t_{n+1})r} \right). \quad (3.1)$$

Let us remark that in equation (3.1), at time t_n , $S_{t_n}^1$ and $S_{t_n}^2$ are measurable, while $S_{t_k}^1$ and $S_{t_k}^2$ for $k > n$ are random variables. We observe that problem (3.1) does not involve the previous hedging strategy, but only the future hedging strategy, i.e. $(\bar{\delta}_k^1, \bar{\delta}_k^2)$ for $k = n+1, \dots, N$. We stress out that excluding past history seems reasonable from a practical perspective since an hedger decides his strategy looking only at future results. In fact, the gain and the losses already realised through the hedging strategy are measurable values and, due to the cash invariance property of risk measures, their inclusion/exclusion does not change the solution to problem (3.1). Moreover, problem (3.1) is suitable to be tackled through a backward procedure: if one solves this kind of problems backward in time, future hedging strategy should be available as referred to time steps already treated. Furthermore, the applicability of a backward approach is also ensured by the fact that Heston's model is Markovian, so the future evolution of processes, conditioned by the state of the present market, is independent of the past. We can therefore observe that values $(\bar{\delta}_k^1, \bar{\delta}_k^2)$ only depend on the current state of the market. Once one has solved problem (3.1) for $t_n = t_N, \dots, 0$, the Backward Hedging option price can be computed as

$$\bar{p}_0 = e^{-rT} \rho \left(-\varphi(S_T^1) + (\bar{\delta} \cdot S)_0^N - C_0^N(\bar{\delta}) \right). \quad (3.2)$$

Before presenting the resolution of the problem (3.1) in Heston's model, we discuss the relation between the strategies $\hat{\delta}$ and $\bar{\delta}$, that is the solution of the global hedging problem (2.3) and of the backward problem (3.1), respectively. We now show that the two strategies, in general, are not equivalent.

3.2 A key example

We consider, as an example, a discrete time model, with only 3 trading times, namely $t_0 = 0, t_1 = 1$ and $t_2 = 2$. For this example only, the market includes a single underlying with initial value $S_0 = 1$. The underlying can assume 5 equiprobable values at time t_1 , which are $S_1^i = S_0 \cdot u^{2i-6}$, for $i = 1, \dots, 5$ and $u = \exp(0.2\sqrt{1/4})$. Moreover, for each value of S_1^i , there are 5 equiprobable continuation values, namely $S_2^{i,j} = S_1^i \cdot u^{2j-6}$, for $j = 1, \dots, 5$. So there are 5 possible paths from t_0 to t_1 and 25 possible paths from t_0 to t_2 . We stress out that these values are the same as those generated at time-steps 0, 4 and 8 by the CRR model with $T = 2$ and $N = 8$ discretization steps, with the only difference being that, to simplify the calculations, we assume transaction probabilities between two consecutive instants all equal to 1/5. Also, for simplicity, we assume the interest rate r to be zero.

In this model, an hedging strategy consists of 6 values: δ_0 , i.e. the amount of stocks to include in the hedging portfolio at time $t = 0$, and $\delta_{1,i} = \delta_1(S_1^i), i = 1, \dots, 5$, i.e. the amount of stocks to include in the hedging portfolio at time t_1 according to the 5 possible values of the underlying at that time.

Now, we consider a short position on a strangle, i.e. a short position on a call and on a put, with strikes $K_C = 0.8$ and $K_P = 1.3$ respectively. In our example, we want to hedge the strangle so as to minimize the 40%-CVaR. For the global-optimal strategy $\hat{\delta}$, this consists of choosing the 6 hedging parameters so as to minimize the CVaR relative to profit-and-loss between t_0 and t_2 , which we denote by CVaR_0^2 . For the backward strategy $\bar{\delta}$, on the other hand, this consists of determining at time t_1 , for each of the possible values of the underlying S_1^i , the hedging parameter $\bar{\delta}_1(S_1^i)$ so as to minimize the CVaR relative to the profit-and-loss between t_1 and t_2 , which we denote by CVaR_1^2 , and then the hedging parameter $\bar{\delta}_0$ at time t_0 so as to minimize the CVaR_0^2 between t_0 and t_2 , being subject to the constraint of using $\bar{\delta}_1(S_1^i)$ as the hedging

Index	$\hat{\delta}$	$\bar{\delta}$
0	0.1468	0.1468
1, 1	0.9217	0.9217
1, 2	0.5203	0.6351
1, 3	0.0650	0.2850
1, 4	-0.3733	-0.2353
1, 5	-0.8671	-0.5567

Table 1: The hedging strategies $\hat{\delta}$ and $\bar{\delta}$.

strategy at time t_1 . Let us point that, since we are assuming equal transaction probabilities and 40% as CVaR level, then CVaR_0^2 is equal to the opposite of the average of the 10 lower profit-and-loss outcomes between t_0 and t_2 , while CVaR_1^2 is equal to the opposite of the average of the 2 lower profit-and-loss outcomes between t_1 and t_2 .

The two optimal strategies are shown in Table 1, while Table 2 reports the computations for those strategies. Specifically, the first three columns show the 5 possible values of S_1 , the 25 final values of S_2 and their payoffs, respectively. The next columns report the profit and loss values for each of the 25 possible trajectories, calculated by applying the 2 strategies under consideration. These profit-and-loss are calculated between t_0 and t_2 , denoted by PL_0^2 , or between t_1 and t_2 , denoted by PL_1^2 . Values in bold are the worst values, and they contribute to the calculation of CVaR values. For both the two strategies, CVaR_1^2 is reported for each of the 5 possible values of S_1 , while the two values for CVaR_0^2 are reported at the bottom of the Table.

We can observe that the two strategies are successful in their respective problems: $\hat{\delta}$ generates a smaller CVaR_0^2 than $\bar{\delta}$, while $\bar{\delta}$ generates smaller (or equal) CVaR_1^2 for all possible values of S_1 . We can therefore conclude that the two strategies $\hat{\delta}$ and $\bar{\delta}$, and therefore the problems (2.3) and (3.1), are not equivalent.

3.3 Comparison between strategies

At this point, it is worth asking which of the two strategies is the most correct to consider for practical purposes. The global optimal strategy, pursued by the Deep-Hedging algorithm, hides a pitfall.

The strategy $\hat{\delta}$ is optimal at time zero, that is between $t_0 = 0$ and T , but only at time zero and has no optimality property at times $t_n > 0$, between $t_n > 0$ and T : in fact, as shown by the previous example, the strategy $\hat{\delta}$ is not optimal at time t_1 , as it does not always minimize CVaR_1^2 , in as exceeded by $\bar{\delta}$. If we assume that, at any time $t_n > 0$, the hedger wishes to minimize the CVaR between t_n and T , then he will be led to abandon strategy $\hat{\delta}$ that he had decided at time 0, in favor of a more performing one. The initial valuation is therefore not consistent with what is expected of the hedger: the option is priced based on a strategy which will not be followed afterwards. This annoying phenomenon is related to the absence of time-consistency for CVaR, and more generally for the most common risk measures.

The $\bar{\delta}$ strategy, on the other hand, at time t_n is optimal given the strategy that will be followed subsequently, therefore it is not convenient to modify it, unlike $\hat{\delta}$, and it is less penalizing in future times. Once established, it should be maintained (unless it is replaced with another $\hat{\delta}$ strategy recomputed to be optimal between t_n and T , which however will be disregarded in subsequent times $t_k > t_n$). This aspect could be appreciated by the supervisory bodies which must validate the hedging procedure.

Finally, we point out that Basak and Chabakauri [3] and subsequently Cui et al. [13], have come to similar

			Strategy $\hat{\delta}$		Strategy $\bar{\delta}$	
S_1	S_2	Payoff	PL_0^2	PL_1^2	PL_0^2	PL_1^2
1.4918	2.2255	1.4255	-0.6771	-0.7493	-0.6771	-0.7493
	1.8221	1.0221	-0.6455	-0.7177	-0.6455	-0.7177
	1.4918	0.6918	-0.6196	-0.6918	-0.6196	-0.6918
	1.2214	0.5000	-0.6771	-0.7493	-0.6771	-0.7493
	1.0000	0.5000	-0.8811	-0.9533	-0.8811	-0.9533
CVaR $_1^2$				0.8513		0.8513
1.2214	1.8221	1.0221	-0.6770	-0.7095	-0.6081	-0.6406
	1.4918	0.6918	-0.5186	-0.5511	-0.4876	-0.5201
	1.2214	0.5000	-0.4675	-0.5000	-0.4675	-0.5000
	1.0000	0.5000	-0.5827	-0.6152	-0.6081	-0.6406
	0.8187	0.5000	-0.6770	-0.7095	-0.7232	-0.7557
CVaR $_1^2$				0.7095		0.6982
1.0000	1.4918	0.6918	-0.6599	-0.6599	-0.5517	-0.5517
	1.2214	0.5000	-0.4856	-0.4856	-0.4369	-0.4369
	1.0000	0.5000	-0.5000	-0.5000	-0.5000	-0.5000
	0.8187	0.5000	-0.5118	-0.5118	-0.5517	-0.5517
	0.6703	0.6297	-0.6511	-0.6511	-0.7236	-0.7236
CVaR $_1^2$				0.6555		0.6376
0.8187	1.2214	0.5000	-0.6769	-0.6503	-0.6214	-0.5948
	1.0000	0.5000	-0.5943	-0.5677	-0.5693	-0.5427
	0.8187	0.5000	-0.5266	-0.5000	-0.5266	-0.5000
	0.6703	0.6297	-0.6009	-0.5743	-0.6214	-0.5948
	0.5488	0.7512	-0.6770	-0.6504	-0.7143	-0.6877
CVaR $_1^2$				0.6504		0.6412
0.6703	1.0000	0.5000	-0.8343	-0.7859	-0.7319	-0.6835
	0.8187	0.5000	-0.6771	-0.6287	-0.6310	-0.5826
	0.6703	0.6297	-0.6781	-0.6297	-0.6781	-0.6297
	0.5488	0.7512	-0.6942	-0.6458	-0.7319	-0.6835
	0.4493	0.8507	-0.7074	-0.6590	-0.7760	-0.7276
CVaR $_1^2$				0.7225		0.7056
CVaR $_0^2$			0.7180		0.7314	

Table 2: Comparison between hedging strategies $\hat{\delta}$ and $\bar{\delta}$.

considerations in the context of portfolio allocation: although the global-optimal strategy first appears to have strong global investment performance, investors would deviate from such a strategy throughout the investment process.

For the reasons previously mentioned, we propose to use $\bar{\delta}$ in place of $\hat{\delta}$.

3.4 European option with no transaction costs, based on risk measure

We discuss the numerical procedure we propose to solve problem (3.1). The issues that one has to face are controlling the hedging strategy for future time-steps, computing the objective function, that is the risk measure, and optimizing the objective function.

First of all, we propose to estimate the risk measure ρ through (forward) M Monte Carlo simulations, from t_n to T , of the underlying processes S^1 and S^2 conditioned by \mathcal{F}_{t_n} . In order to simulate future values of S^1, S^2 , one only needs to know $S_{t_n}^1, V_{t_n}$ and $I_{t_n} = \int_0^{t_n} V_s ds$ in place of all filtration \mathcal{F}_{t_n} , so the market state can be described by this triple. Please, note that knowing $(S_{t_n}^1, S_{t_n}^2)$ is not sufficient to simulate future values $\{(S_{t_k}^1, S_{t_k}^2) | \mathcal{F}_{t_n}, k = n+1, \dots, N\}$, as, to simulate the underlying $S_{t_k}^1$, one must first simulate the volatility V_{t_k} , and the initial value of the volatility V_{t_n} cannot be deduced from $(S_{t_n}^1, S_{t_n}^2)$. Therefore we also generate the values of I and V , and denote the simulated values with

$$\mathcal{S}_n = \{(S_{k,m}^1, S_{k,m}^2, I_{k,m}, V_{k,m}), k = n+1, \dots, N, m = 1, \dots, M\}.$$

Now, for each time-step t_n , we consider a grid of values $\mathcal{G}_n = \mathcal{G}_n^S \times \mathcal{G}_n^I \times \mathcal{G}_n^V$ which is defined as follows: \mathcal{G}_n^S is a uniform mesh of N_S points for $S_{t_n}^1$, \mathcal{G}_n^I is a uniform mesh of N_I points for I_{t_n} and \mathcal{G}_n^V is a mesh of N_V points for V_{t_n} . The maximum and minimum values of \mathcal{G}_n^S and \mathcal{G}_n^I are set by computing appropriate quantiles (in our tests we consider 0.1% and 99.9% quantiles) for $S_{t_n}^1$ and I_{t_n} . Those quantiles are estimated from M_0 Monte Carlo simulations of the Heston model, starting from the initial values $S_0^1 = s_0$ and $V_0 = v_0$, which we denote with

$$\hat{\mathcal{S}}_n = \left\{ (\hat{S}_{k,m}^1, \hat{I}_{k,m}), k = 1, \dots, N, m = 1, \dots, M_0 \right\}.$$

Simulations of the stochastic processes are obtained through the hybrid Monte Carlo-Tree method introduced by Briani et al. [8, 9], together with the trinomial tree introduced by Goudenège et al. [17], by using N_T simulation time-steps (N_T must be a multiple of N , so that the simulated values at the trading epochs can be obtained by subsampling). This simulation method is particularly interesting in that the possible states for the random variable V_{t_n} are discretized through the use of a trinomial tree \mathcal{T} that exactly reproduces the first and second moments of the transition distribution of the process V , thus leading to accurate results with only a few time-discretization steps. This trinomial tree is also used to define the grid \mathcal{G}_n^V as the values of V associated to the tree nodes at time t_n . Moreover, it is worth observing that the simulated values of V always belong to the grid. Finally, we also set $\mathcal{G}_0 = \{s_0\} \times \{0\} \times \{v_0\}$ for the market state at inception.

Problem (3.1) is solved for each point of \mathcal{G} backward in time so, at time t_n , the hedging strategy is known at future time-steps $\{t_k\}_{k=n+1, \dots, N-1}$, but only for those values of $(S_{t_k}^1, I_{t_k}, V_{t_k})$ that are points of \mathcal{G}_k . However, the simulated values of $S_{t_k}^1$ and I_{t_k} for the estimation of the risk measure may not belong to the grid \mathcal{G}_k and therefore the hedging strategy would not be known for those special values. This problem is overcome by using linear interpolation, which allows one to extend the hedging strategy from \mathcal{G}_k to all \mathbb{R}^3 . In this regard, we point out that it is necessary to resort to an interpolation in the S and I coordinates only, as the simulated values of V belong to the grid thanks to the use of the trinomial tree \mathcal{T} for simulation. We also point out that although other forms of interpolation are possible (see, e.g., Akima [1, 2] for Modified Akima interpolation or Press et al. [28] for Bicubic Interpolation), linear interpolation is chosen for its low computational cost. This element is decisive as the proposed algorithm requires the calculation of a large

number of interpolations. As a result, we can represent the hedging strategy as a collection of function $\{(\bar{\delta}_n^1, \bar{\delta}_n^2) : \mathbb{R}^3 \rightarrow \mathbb{R}^2\}$ such that $(\bar{\delta}_n^1, \bar{\delta}_n^2)(s, i, v)$ are the optimal hedging values when $S_{t_n}^1 = s, I_{t_n} = i$ and $V_{t_n} = v$.

Now that we have clarified how to calculate the objective function, we must solve the optimization problem (3.1), thus determining the specific values $(\bar{\delta}_n^1, \bar{\delta}_n^2)$ for the grid point under evaluation. This is done by exploiting a standard numerical procedure. Specifically, we used the Simplex Search method of Lagarias et al. [23], which is an efficient derivative-free algorithm and it is particularly suitable for problems of this kind.

Summing up, the Backward Hedging algorithm proceeds backward in time by solving problems (3.1) for all nodes in the grid, starting from time t_{N-1} down to time t_0 . Once the hedging strategy has been determined for all trading times, we simulate a new sample of M_1 Monte Carlo simulations

$$\bar{S}_0 = \{(\bar{S}_{k,m}^1, \bar{S}_{k,m}^2, \bar{I}_{k,m}, \bar{V}_{k,m}), k = 1, \dots, N, m = 1, \dots, M_1\}$$

and we apply the algorithm on this new out-of-sample data. The risk measure determined on this new sample defines the initial option price \bar{p}_0 with respect to the Backward Hedging algorithm.

3.5 Loss as mean squared error

The Backward Hedging algorithm can also be employed to tackle problem (2.2), by a backward approach. Again, let us again assume that no transaction costs are applied. It is easy to prove that the optimal value of p_0 according to (2.2) is given by

$$p_0 = -e^{-rT} \mathbb{E} [-\varphi(S_T^1) + (\delta \cdot S)_0^N], \quad (3.3)$$

and for this choice of p_0 , we obtain

$$\mathbb{E} [(-\varphi(S_T^1) + p_0 e^{rT} + (\delta \cdot S)_0^N)^2] = \text{Var} [-\varphi(S_T^1) + (\delta \cdot S)_0^N],$$

thus (2.2) reduces to solving

$$\inf_{\delta \in \mathcal{H}} \text{Var} [-\varphi(S_T^1) + (\delta \cdot S)_0^N]. \quad (3.4)$$

Now, let us consider the backward version of problem (3.4), that is at time t_n one has to solve the following problem

$$(\bar{\delta}_n^1, \bar{\delta}_n^2) = \arg \min_{(x,y) \in \mathbb{R}^2} \text{Var} [-\varphi(S_T^1) + (\bar{\delta} \cdot S)_{n+1}^N + \left[x \left(S_{t_{n+1}}^1 - S_{t_n}^1 e^{r\Delta t} \right) + y \left(S_{t_{n+1}}^2 - S_{t_n}^2 e^{r\Delta t} \right) \right] e^{(T-t_{n+1})r}]. \quad (3.5)$$

Problem (3.5) is of the same type of (3.1), with the only difference being that ρ is replaced with Var , so one can tackle it again with the Backward Hedging algorithm, as described in the previous Subsection.

3.6 American options

American options can be exercised at any moment before maturity. With the scope of pricing, we approximate such a feature with a weaker one: we suppose that the option can be exercised at any trading time. If the option is exercised, the hedger pays the payoff to the buyer, otherwise he rebalances the hedging portfolio. Pricing requires to model the exercising strategy, which depends on the buyer. Here, we assume the so called *worst-hedging-case*, which means that the option is exercised in the worst case for the hedger, that is as soon

as the payoff is higher than the hedging cost, which corresponds to the continuation cost. Specifically, the option price \bar{p}_n at time t_n is define as a function of $S_{t_n}^1, I_{t_n}, V_{t_n}$, again as the cost for hedging the position, from t_n up the the exercise time τ_n that is

$$\bar{p}_n(S_{t_n}^1, I_{t_n}, V_{t_n}) = \inf_{\delta \in \mathcal{H}} \rho \left(\left(-\varphi(S_{\tau_n}^1) + (\bar{\delta} \cdot S)_0^{\tau_n/\Delta t} \right) e^{-r(\tau_n - t_n)} \right). \quad (3.6)$$

Specifically, τ_n is the stopping time defined as

$$\tau_n = \min \{ \tau \in \{t_n, \dots, T\} \text{ s.t. } \varphi(S_\tau^1) \geq \bar{p}_n(S_\tau^1, I_\tau, V_\tau) \}. \quad (3.7)$$

and so we can tackle such a problem again by backward induction. In particular, the hedging strategy is computed by solving the following problem

$$\begin{aligned} (\bar{\delta}_n^1, \bar{\delta}_n^2) = \arg \min_{(x,y) \in \mathbb{R}^2} \rho \left(-\varphi(S_{\tau_n}^1) + (\bar{\delta} \cdot S)_{n+1}^{\tau_n/\Delta t} + \right. \\ \left. \left[x \left(S_{t_{n+1}}^1 - S_{t_n}^1 e^{r\Delta t} \right) + y \left(S_{t_{n+1}}^2 - S_{t_n}^2 e^{r\Delta t} \right) \right] e^{r(\tau_n - t_{n+1})} \right). \end{aligned} \quad (3.8)$$

Again, such a problem is solved backward in time, by exploiting Monte Carlo simulations and linear interpolation. In particular, the calculation of the exercise strategy requires the comparison between the payoff and option price, so, in this case, linear interpolation is used for both interpolating the hedging strategy and the option price at each trading time.

3.7 Transaction costs

Let us consider proportional transaction on traded assets as in Buehler et al. [10]. The inclusion of transaction costs complicates the evaluation of the hedging strategy because the optimal strategy at one instant is closely related to the strategy employed at the earlier time-step.

We start by considering the problem of hedging a European derivative, based on the risk measure ρ , when proportional transaction costs are due. The following equation defines the Backward Hedging strategy at time t_n :

$$\begin{aligned} (\bar{\delta}_n^1, \bar{\delta}_n^2) = \arg \min_{(x,y) \in \mathbb{R}^2} \rho \left(-\varphi(S_T^1) + (\bar{\delta} \cdot S)_{n+1}^N + C_{n+1}^N(\bar{\delta}) \right. \\ \left. + \left[x \left(S_{t_{n+1}}^1 - S_{t_n}^1 e^{r\Delta t} \right) + y \left(S_{t_{n+1}}^2 - S_{t_n}^2 e^{r\Delta t} \right) \right] e^{(T-t_{n+1})r} \right. \\ \left. + \varepsilon \left[S_{t_n}^1 |x - \delta_{n-1}^1| + S_{t_n}^2 |y - \delta_{n-1}^2| \right] e^{(T-t_n)r} \right). \end{aligned} \quad (3.9)$$

Similarly to what we did for the no-transaction-costs case, we replace conditioning with respect to \mathcal{F}_{t_n} with a discrete set of data, i.e. $S_{t_n}^1 = s, I_{t_n} = i, V_{t_n} = v, \delta_{n-1}^1 = d^1, \delta_{n-1}^2 = d^2$.

We stress out that, in this case, we cannot eliminate the formula's dependence on $(\delta_{n-1}^1, \delta_{n-1}^2)$, which, proceeding backward, will be computed at the next time-step as $(\bar{\delta}_{n-1}^1, \bar{\delta}_{n-1}^2)$, with the only exception for $\delta_0^1 = \delta_0^2 = 0$. So, we cannot apply the backward approach directly.

To overcome this obstacle, we consider a 5 dimensional grid, which discretizes the conditioning event. Specifically, for each time-step t_n , we consider a grid of values $\mathcal{G}_n = \mathcal{G}_n^S \times \mathcal{G}_n^I \times \mathcal{G}_n^V \times \mathcal{G}_n^{\delta^1} \times \mathcal{G}_n^{\delta^2}$ which is defined as follows: $\mathcal{G}_n^S, \mathcal{G}_n^I$ and \mathcal{G}_n^V are defined as in Section 3.4, while $\mathcal{G}_n^{\delta^1}$ and $\mathcal{G}_n^{\delta^2}$ are two uniform mesh of N_{δ^1} and N_{δ^2} points that represent the values for δ_{n-1}^1 and δ_{n-1}^2 respectively. Specifically, the maximum and minimum values of $\mathcal{G}_n^{\delta^1}$ and $\mathcal{G}_n^{\delta^2}$ are determined as specific quantiles (in our tests we consider 0.1% and 99.9% quantiles) of δ_{n-1}^1 and δ_{n-1}^2 when $\varepsilon = 0$. Basically, those quantiles are computed by exploiting Monte

Carlo simulations of the Backward Hedging algorithm for the case without transaction costs, starting from time $t = 0$. The interval for the grids $\mathcal{G}_n^{\delta^1}$ and $\mathcal{G}_n^{\delta^2}$ are appropriately extended to include the zero value, which is the delta value at the beginning and at the end of the hedging strategy.

The Backward Hedging algorithm for this case is developed following the same valuation principle adopted for the case without transaction costs: it proceeds backward in time and, at each time-step t_n , for each point of the 5-dimensional, it determines the optimal values of δ_n^1 and δ_n^2 through a series of Monte Carlo simulations up to maturity (or to the exercise time with respect to American options), considering the values associated with the starting grid point as the market state parameters. Again, linear interpolation is used to extend the hedging strategy from the grid nodes to the entire \mathbb{R}^5 space. Furthermore, in order to favour the stability of the algorithm in the case of extreme trajectories (that is simulated values are not inside in the grid), the values predicted by the interpolation for δ_k^1 and δ_k^2 for $k = n + 1, \dots, N$, are restricted to the discretization interval of $\mathcal{G}_k^{\delta^1}$ and $\mathcal{G}_k^{\delta^2}$ respectively.

Remark 1. In all cases considered in the preceding sections, for each time-step, the calculation of the hedging strategies for each grid point are independent of each other, so this step can be parallelised, thus reducing the total calculation time of the Backward Hedging algorithm.

4 Numerical results

In this Section we discuss some numerical results related to the proposed algorithm. Specifically, we compute the optimal hedging strategy for a European or American put option, considering either the minimization of the mean square error of hedging or the CVaR at maturity as the objective, including (i.e. $\varepsilon = 1\%$) or not transaction costs (i.e. $\varepsilon = 0$). Moreover, we consider different settings for N , the number of trading dates, and M , the number of Monte Carlo simulations used for solving the optimization problems that define the hedging strategy. In particular, in order to obtain reasonable calculation times, we have considered values $M = 10^3, 10^4, 10^5$. We first consider the model parameters employed for a test by Buehler et al. [10], and secondly the same parameters but with positive interest rate and drift. The employed values are reported in Table 3 together with the put option parameters, while Table 4 displays the parameters for the Backward Hedging algorithm.

The following tables show the hedging prices at time t_0 in the various cases considered, obtained by means of the Backward Hedging algorithm (BH), which has been implemented in the MATLAB environment. For the sake of comparison, these prices were also calculated using our implementation of the Deep Hedging algorithm (DH), which has been implemented in Python by us. The parameters for the Deep Hedging algorithm are listed in Table 5. We point out that, for the Deep Hedging algorithm, there are two parameters that govern the number of Monte Carlo simulations: M the number of simulations for training the neural network and M_1 the number of out-of-sample simulations performed at the end to generate the option price. For the Backward Hedging algorithm, to these two parameters, we add M_0 , the number of simulations used to determine the grid that discretizes the market states. For the sake of comparison, we also report the computational times, although these times should be taken with some detachment as they refer to codes developed in different environments. Furthermore, the implementation of the Deep Hedging algorithm was carried out by us and there is probably room for improvement.

The timed comparisons were carried out on a personal computer equipped with an Intel i5-1035G1 processor and 8 GB of RAM, and only one core was used in order to stabilize the calculation times as much as possible.

We emphasise that the prices shown in the following tables are determined by considering the profit and

loss of the hedging strategy from time zero to maturity, thus representing an opportunity to compare the two hedging strategies on the same objective functions, $\rho(PL_0^T)$ and $\text{Var}(PL_0^T)$. In this regard, DH's strategy should be the optimum point in the set of hedging strategies for the global problem, while BH's strategy can be considered as a solution to the same problem, which includes the constraint of given by the recursive scheme that defines the backward strategy. Then, through these numerical tests, we can compare the two strategies and see how optimal they are for the global problem.

Symbol	Meaning	Value
S_0	Initial spot value	100
r	Risk free i.r.	0, 0.1
μ	Drift	0, 0.1
v_0	Vol at beginning	0.04
α	Mean reversion speed	1
b	Long run vol	0.04
σ	Vol of Vol	2
ρ_{BW}	Correlation	-0.7
T	Maturity	1.0
K	Strike	100

Table 3: Parameters of the Heston model and of the put option.

Symbol	Meaning	Value
M_0	Number of MC paths for computing quantiles	10^6
M	Number of MC paths for backward step	$10^3, 10^4, 10^5$
M_1	Number of MC paths for out-of-sample step	10^6
N	Number of trading dates	4, 8, 16
N_T	Number of time-steps for simulation	16
N_S	Number of points to discretize S	15
N_I	Number of points to discretize the integral of V	2
N_{δ^1}	Number of points to discretize δ_{n-1}^1	3
N_{δ^2}	Number of points to discretize δ_{n-1}^2	3

Table 4: Backward Hedging parameters.

Symbol	Meaning	Value
M_0	Number of MC paths for computing quantiles	10^6
M	Number of MC paths for backward step	$10^3, 10^4, 10^5$
N	Number of trading dates	4, 8, 16
N_T	Time-steps for simulation	64
N_N	Number of neurons per layer	17
N_L	Number of hidden layers	2
$\sigma(x)$	Activation function	$\max(x, 0)$

Table 5: Deep Hedging parameters.

4.1 Mean squared error

In this Section we consider the minimization of the mean squared hedging error as target. Tables 6 and 7 show the results for the European and the American put option for $r = \mu = 0$, while Tables 8 and 9 report the corresponding results for the $r = \mu = 0.1$ case. We also show the price in Heston’s continuous model, obtained through the Hybrid Tree-PDE algorithm by Briani et al. [8, 9]. This price serves as a benchmark for the case where transaction costs are zero. No benchmark is available for positive transaction costs.

In all the considered cases, the prices returned by the BH algorithm are very stable with respect to M and the results are consistent with both the benchmarks (when available) and the results returned by the DH algorithm (for the European option). When $r = \mu = 0$, the American price is very similar to the European price, but when $r = \mu = 0.10$, the difference is clear. In the latter case, we observe that the price increases as the number of instants of hedging increases because the number of possibilities for exercising the option increases.

To conclude, the proposed algorithm proves to be very efficient and stable.

Benchmark	$\varepsilon = 0$			$\varepsilon = 1\%$			
	M	N	3.7165	4	8	16	
BH	10^3	3.7361 (7)	3.7518 (10)	3.8075 (19)	5.2907 (97)	5.6050 (377)	7.0452 (1501)
	10^4	3.7024 (8)	3.7173 (15)	3.6962 (32)	5.1429 (359)	5.3127 (2059)	5.4350 (8841)
	10^5	3.7081 (16)	3.7031 (44)	3.7061 (130)	5.0779 (2548)	5.1680 (16919)	5.2412 (75182)
DH	10^3	3.3484 (200)	2.6859 (622)	2.6354 (721)	5.3223 (172)	6.3477 (624)	8.0656 (1188)
	10^4	3.6130 (991)	3.5731 (2153)	3.4796 (4152)	5.4080 (2752)	6.9781 (3007)	10.5884 (4918)
	10^5	3.7019 (8741)	3.7241 (19587)	3.7426 (51194)	5.3400 (18890)	6.6993 (50125)	9.3661 (92425)

Table 6: Prices for a European put option with $r = \mu = 0$, while minimizing loss based on mean squared error.

Benchmark		$\varepsilon = 0$			$\varepsilon = 1\%$		
		M	N				
Benchmark		3.7165			-		
		4	8	16	4	8	16
	10^3	3.7360 (8)	3.7513 (16)	3.8066 (49)	5.2907 (86)	5.6603 (382)	6.9647 (1618)
BH	10^4	3.7022 (12)	3.7172 (35)	3.6958 (111)	5.0861 (407)	5.3101 (2377)	5.5282 (10195)
	10^5	3.7081 (42)	3.7038 (197)	3.7057 (872)	5.1669 (3336)	5.1404 (20906)	5.3706 (88598)

Table 7: Prices for an American put option with $r = \mu = 0$, while minimizing loss based on mean squared error.

Benchmark		$\varepsilon = 0$			$\varepsilon = 1\%$		
		M	N				
Benchmark		2.1334			-		
		4	8	16	4	8	16
	10^3	2.1537 (6)	2.1470 (10)	2.2142 (18)	2.9850 (96)	3.1667 (391)	3.9847 (1499)
BH	10^4	2.1392 (8)	2.1547 (14)	2.1412 (31)	3.0313 (364)	3.0702 (2072)	3.0806 (8895)
	10^5	2.1462 (16)	2.1401 (43)	2.1410 (147)	2.9695 (2674)	2.8845 (17137)	2.9150 (74737)
	10^3	1.3124 (208)	1.3044 (396)	1.1886 (677)	2.3709 (440)	3.9069 (506)	5.2301 (800)
DH	10^4	1.9725 (944)	1.9670 (1981)	1.7516 (4399)	3.4638 (900)	4.7569 (2427)	4.6196 (4678)
	10^5	2.1270 (11319)	2.0983 (20367)	2.1339 (41625)	3.2011 (27554)	3.9313 (64813)	4.1068 (85591)

Table 8: Prices for a European put option with $r = \mu = 0.1$, while minimizing loss based on mean squared error.

Benchmark		$\varepsilon = 0$			$\varepsilon = 1\%$		
		M	N				
Benchmark		2.8355			-		
		4	8	16	4	8	16
	10^3	2.6580 (6)	2.7427 (10)	2.8339 (20)	3.5394 (87)	3.8889 (387)	4.7330 (1423)
BH	10^4	2.6238 (8)	2.7503 (16)	2.7745 (39)	3.4659 (444)	3.8864 (2945)	3.7912 (10535)
	10^5	2.6264 (20)	2.7376 (64)	2.7730 (203)	3.5296 (4140)	3.5533 (27944)	3.6406 (93940)

Table 9: Prices for an American put option with $r = \mu = 0.1$, while minimizing loss based on mean squared error.

4.2 CVaR error

In this Section we consider the minimization of the 95%-CVaR of the hedging strategy at maturity as target. Before discussing the results, we point out that, in order to facilitate comparison between the proposed algorithms, in this case we have also included, for the DH algorithm only, the case $M = 10^6$, which is necessary to obtain competitive results with BH.

Tables 10 and 11 show the results for the European and the American put option for $r = \mu = 0$, while Tables 12 and 13 report the corresponding results for the $r = \mu = 0.1$ case. No benchmark is available in this case. In all cases considered, the prices returned by the BH algorithm are very stable with respect to M . As far

as European options are considered, the price decreases as N , the number of possible portfolio rebalances, increases, since the higher this number, the smaller the error committed at the end of the hedging period and the lower the relative CVaR. In fact, when one aims to minimise a risk measure of replication error, unlike mean squares error minimisation, the lower the price obtained, the better the strategy.

As observed in the previous numerical experiments, discussed in Subsection 4.1, in the presence of transaction costs, the option price increases, as expected. As noted in the previous section, when $r = \mu = 0$, the American price is very similar to the European price, but when $r = \mu = 0.10$, the difference is clear. In the latter case, we observe that the price decreases: while the increase in trading dates favours the optimal exercise of the option that raises the price, it also favours the development of an effective hedging strategy that lowers the price. Between these two contrasting effects, the second seems to prevail. As shown in Tables 10 and 12, the comparison with DH shows that BH's proposed strategy is also very efficient in dealing with the global problem: for all the parameter configurations considered, the price of the BH strategy is very similar to the price of DH's strategy, as their difference is small, and in some cases even lower. Moreover, the computational time tends to reward BH. We also observe that the run times of BH are more regular than those of DH, as the stop criterion employed for training the neural network seems to be more sensitive to specific input data. In particular, DH times seem to increase up to $M = 10^5$, whereas for $M = 10^6$ they are of the same order of magnitude as for $M = 10^5$. In conclusion, BH again proves to be very efficient and stable.

To complete our analysis, we present Figure 4.1 which shows a plot of the optimal values of $\bar{\delta}_3^1$ and $\bar{\delta}_3^2$, the hedging strategy at time $t_3 = 0.75$, computed through the Backward Hedging algorithm, as a function of the hedging strategy at time $t_2 = 0.5$, (δ_2^1, δ_2^2) . In particular, we consider an American put option with $N = 4$ hedging steps, $S_{0.75} = 100$, $I_{0.75} = 0.04 \cdot 0.75$ and $V_{0.75} = 0.04$, with $M = 10^4$ Monte Carlo simulations. Hedging is based on the CVaR minimization and it employs $\varepsilon = 1\%$ proportional transaction cost. To improve the graphic detail, the values N_{δ_1} and N_{δ_2} have been set equal to 11.

The graphs suggest that both $\bar{\delta}_3^1$ and $\bar{\delta}_3^2$ increase as δ_2^1 and δ_2^2 increase, with a greater dependence on δ_2^1 . This fact probably results from the low absolute cost in trading S^2 due to the small price of S^2 with respect of S^1 . To confirm this, in Figure 4.2 we can observe a graph similar to that of Figure 4.1, in which the relative cost of trading S^2 has been raised to 10%, while that for S^1 has remained 1%. In this case, two increasing graphs are still observed, but the dependence on δ_2^1 is more relevant, in particular for $\bar{\delta}_3^2$.

		$\varepsilon = 0$			$\varepsilon = 1\%$		
M	N	4	8	16	4	8	16
BH	10^3	10.4559 (10)	8.7008 (14)	7.6950 (24)	11.5441 (133)	10.3812 (502)	9.3263 (1978)
	10^4	10.1111 (35)	8.6460 (74)	7.4112 (143)	11.5375 (467)	10.3645 (2510)	9.3284 (10388)
	10^5	10.0945 (351)	8.6725 (762)	7.4855 (1566)	11.5099 (3440)	10.3386 (19424)	9.3564 (80715)
DH	10^3	20.7569 (704)	22.8564 (794)	45.1245 (1908)	25.8685 (748)	25.4389 (305)	24.9047 (500)
	10^4	18.4303 (2001)	20.9419 (1539)	24.5623 (5676)	20.3021 (1748)	27.6064 (784)	25.3484 (1455)
	10^5	12.1416 (32783)	10.7884 (52974)	12.6447 (45811)	11.7335 (4667)	12.0040 (9600)	13.7020 (11239)
	10^6	9.9972 (21818)	9.2974 (42418)	8.3096 (84488)	11.5243 (28314)	9.9993 (45034)	9.0129 (100586)

Table 10: Hedging cost for a European put option with $r = \mu = 0$, while minimizing loss based on CVaR.

		$\varepsilon = 0$			$\varepsilon = 1\%$		
M	N	4	8	16	4	8	16
	10^3	10.4560 (10)	8.7006 (14)	7.6962 (27)	11.5435 (119)	10.3831 (499)	9.3273 (2087)
BH	10^4	10.1111 (35)	8.6462 (73)	7.4082 (148)	11.5382 (476)	10.3588 (2733)	9.3308 (11618)
	10^5	10.0944 (355)	8.6724 (759)	7.4855 (1625)	11.5117 (3721)	10.3395 (21980)	9.3630 (91381)

Table 11: Hedging cost for an American put option with $r = \mu = 0$, while minimizing loss based on CVaR.

		$\varepsilon = 0$			$\varepsilon = 1\%$		
M	N	4	8	16	4	8	16
	10^3	6.5798 (13)	5.3969 (13)	4.4896 (23)	7.4402 (139)	6.1939 (497)	5.3032 (1957)
BH	10^4	6.5619 (35)	5.2388 (70)	4.3412 (140)	7.3739 (448)	6.0801 (2443)	5.3044 (10104)
	10^5	6.5139 (359)	5.2283 (803)	4.3347 (1545)	7.3460 (3354)	6.0491 (20161)	5.3076 (79368)
	10^3	18.1152 (388)	23.8998 (522)	24.1038 (1530)	24.5469 (424)	27.4145 (802)	29.8575 (1689)
DH	10^4	17.5067 (1620)	19.2940 (1818)	18.3218 (3710)	18.2993 (2031)	18.8594 (1892)	21.2792 (5720)
	10^5	10.2705 (26602)	12.2695 (11261)	8.1678 (67796)	10.3475 (42343)	8.1222 (53501)	7.9165 (105889)
	10^6	6.6167 (22974)	5.8382 (31135)	5.3762 (42442)	6.9460 (19964)	5.9089 (62982)	5.2575 (110851)

Table 12: Hedging cost for a European put option with $r = \mu = 0.1$, while minimizing loss based on CVaR.

		$\varepsilon = 0$			$\varepsilon = 1\%$		
M	N	4	8	16	4	8	16
	10^3	8.2635 (9)	7.0993 (14)	6.0820 (30)	8.7475 (117)	7.7933 (438)	7.3570 (1729)
BH	10^4	8.3280 (34)	6.9922 (71)	6.0175 (143)	8.6478 (454)	7.9860 (2332)	7.4585 (9742)
	10^5	8.2982 (341)	6.9864 (731)	5.9918 (1530)	8.8270 (3539)	8.0841 (19797)	7.5112 (77681)

Table 13: Hedging cost for an American put option with $r = \mu = 0.1$, while minimizing loss based on CVaR.




Figure 4.1: Optimal values of $\bar{\delta}_3^1$ (right) and $\bar{\delta}_3^2$ (left) as a function of (δ_2^1, δ_2^2) .




Figure 4.2: Optimal values of $\bar{\delta}_3^1$ (right) and $\bar{\delta}_3^2$ (left) as a function of (δ_2^1, δ_2^2) . Proportional cost for trading S^1 is 1%, while the proportional cost for trading S^2 has been raised to 10%.

5 Conclusions

In this paper, we have discussed a method for identifying the most efficient hedging strategies for derivative products, including both European and American options, while taking into account transaction costs. The hedging strategy, defined as the quantities of assets held in the hedging portfolio, is determined to minimize a risk measure or the mean squared error of the hedging error at the time the option is exercised. We calculate the optimal strategy for each point on a grid by using Monte Carlo simulations and then we extend the strategy to other points via interpolation. Although the proposed strategy is not optimal at time zero for the global optimization problem, it offers, at each time-step, a strategy that is optimal given the strategy that will be followed later, so it is easier to follow in practice, pleasing both practitioners and supervisors. The proposed method proves to be effective and reliable, offering clear interpretation of the results. The algorithm has the potential to be valuable for practitioners managing option positions on a single underlying asset.

References

- [1] H. Akima. A method of bivariate interpolation and smooth surface fitting based on local procedures. *Communications of the ACM*, 17(1):18–20, 1974.
- [2] H. Akima. A method of bivariate interpolation and smooth surface fitting for irregularly distributed data points. *ACM Transactions on Mathematical Software (TOMS)*, 4(2):148–159, 1978.
- [3] S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *The Review of Financial Studies*, 23(8):2970–3016, 2010.
- [4] S. Becker, P. Cheridito, and A. Jentzen. Pricing and hedging American-style options with deep learning. *Journal of Risk and Financial Management*, 13(7):158, 2020.
- [5] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.
- [6] K. Boda, J. A. Filar, et al. Time consistent dynamic risk measures. *Mathematical Methods of Operations Research*, 63(1):169–186, 2006.
- [7] B. Bouchard, L. Moreau, and H. M. Soner. Hedging under an expected loss constraint with small transaction costs. *SIAM Journal on Financial Mathematics*, 7(1):508–551, 2016.
- [8] M. Briani, L. Caramellino, and A. Zanette. A hybrid approach for the implementation of the Heston model. *IMA Journal of Management Mathematics*, 28(4):467–500, 2017.
- [9] M. Briani, L. Caramellino, and A. Zanette. A hybrid tree/finite-difference approach for Heston-Hull-White-type models. *Journal of Computational Finance*, 21(3):1–45, 2017.
- [10] H. Buehler, L. Gonon, J. Teichmann, and B. Wood. Deep Hedging. *Quantitative Finance*, 19(8):1271–1291, 2019.
- [11] A. Carbonneau. Deep hedging of long-term financial derivatives. *Insurance: Mathematics and Economics*, 99:327–340, 2021.

- [12] P. Cheridito and M. Stajde. Time-inconsistency of var and time-consistent alternatives. *Finance Research Letters*, 6(1):40–46, 2009.
- [13] X. Cui, J. Gao, Y. Shi, and S. Zhu. Time-consistent and self-coordination strategies for multi-period mean-conditional value-at-risk portfolio selection. *European Journal of Operational Research*, 276(2):781–789, 2019.
- [14] D. Duffie and H. R. Richardson. Mean-variance hedging in continuous time. *The Annals of Applied Probability*, pages 1–15, 1991.
- [15] H. Föllmer and P. Leukert. Quantile hedging. *Finance and Stochastics*, 3:251–273, 1999.
- [16] H. Föllmer and P. Leukert. Efficient hedging: cost versus shortfall risk. *Finance and Stochastics*, 4:117–146, 2000.
- [17] L. Goudenège, A. Molent, and A. Zanette. Gaussian process regression for pricing variable annuities with stochastic volatility and interest rate. *Decisions in Economics and Finance*, 44:57–72, 2021.
- [18] S. L. Heston. 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, 1993.
- [19] B. Horvath, J. Teichmann, and Ž. Žurič. Deep Hedging under rough volatility. *Risks*, 9(7):138, 2021.
- [20] J. Kallsen and J. Muhle-Karbe. Option pricing and hedging with small transaction costs. *Mathematical Finance*, 25(4):702–723, 2015.
- [21] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [22] S. Klöppel and M. Schweizer. Dynamic indifference valuation via convex risk measures. *Mathematical Finance*, 17(4):599–627, 2007.
- [23] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright. Convergence properties of the Nelder–Mead simplex method in low dimensions. *SIAM Journal on optimization*, 9(1):112–147, 1998.
- [24] T. C. Lin. Artificial intelligence, finance, and the law. *Fordham L. Rev.*, 88:531, 2019.
- [25] R. C. Merton. Theory of rational option pricing. *The Bell Journal of economics and management science*, pages 141–183, 1973.
- [26] B. Pochart and J.-P. Bouchaud. Option pricing and hedging with minimum local expected shortfall. *Quantitative Finance*, 4(5):607–618, 2004.
- [27] M. Potters, J.-P. Bouchaud, and D. Sestovic. Hedged Monte-Carlo: low variance derivative pricing with objective probabilities. *Physica A: Statistical Mechanics and its Applications*, 289(3-4):517–525, 2001.
- [28] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007.
- [29] B. Rémillard and S. Rubenthaler. Optimal hedging in discrete and continuous time. *Available at SSRN 1522090*, 2009.

- [30] M. Schweizer. Mean-variance hedging for general claims. *The annals of applied probability*, pages 171–179, 1992.
- [31] M. Schweizer. Variance-optimal hedging in discrete time. *Mathematics of Operations Research*, 20(1):1–32, 1995.
- [32] A. E. Whalley and P. Wilmott. An asymptotic analysis of an optimal hedging model for option pricing with transaction costs. *Mathematical Finance*, 7(3):307–324, 1997.
- [33] M. Xu. Risk measure pricing and hedging in incomplete markets. *Annals of Finance*, 2(1):51–71, 2006.

A Pseudocode

This Appendix contains pseudocodes of the Backward Hedging algorithm for the valuation of European and American options.

Algorithm 1 Backward Hedging for European options

Set up volatility tree \mathcal{T} by following [17]

if $\varepsilon > 0$ **then**

run this algorithm with $\varepsilon = 0$ to obtain the parameters for the 5D grids \mathcal{G}_n for $n = 1, \dots, N - 1$

else

simulate M_0 random paths $(S_{n,m}^1, I_{n,m})$ for $k = 1, \dots, N$

for $n = N - 1, \dots, 0$ **do**

compute quantiles from $(\tilde{S}_{n,m}^1, \tilde{I}_{n,m})$ and set the 3D grid \mathcal{G}_n from quantiles and \mathcal{T}

end for

end if

for $n = N - 1, \dots, 0$ **do**

for all points G of \mathcal{G}_n (if $\varepsilon = 0$ then $G = (s, i, v)$ else $G = (s, i, v, d^1, d^2)$) **do**

simulate M random paths $(S_{k,m}^1, S_{k,m}^2, I_{k,m}, V_{k,m})$ for $k = n + 1, \dots, N$, according to $S_{t_n}^1 = s, I_{t_n} = i, V_{t_n} = v$

while the Simplex Search method has not converged **do**

change $(\delta_{n,m}^1, \delta_{n,m}^2)$ according to the Simplex Search method

for $m = 1, \dots, M$ (for each simulated path) **do**

for $k = n + 1, \dots, N$ **do**

compute the hedging strategy $(\bar{\delta}_{k,m}^1, \bar{\delta}_{k,m}^2)$ by linear interpolation over \mathcal{G}_k (given $(\delta_{k-1,m}^1, \delta_{k-1,m}^2)$ if $\varepsilon > 0$)

end for

compute the hedging error $-\varphi(S_{T,m}^1) + (\bar{\delta} \cdot S_m)_n^N + C_n^N(\bar{\delta})$

end for

estimate $\rho(-\varphi(S_T^1) + (\bar{\delta} \cdot S)_n^N + C_n^N(\bar{\delta}))$ from simulations

end while

end for

if $n > 0$ **then**

compute and store coefficients for linear interpolation of $\bar{\delta}_n^1$ and $\bar{\delta}_n^2$ over \mathcal{G}_n

end if

end for

Simulate M_1 random paths $(S_{n,m}^1, S_{n,m}^2, I_{n,m}, V_{n,m})$ for $n = 1, \dots, N$

for $m = 1, \dots, M$ (for each simulated path) **do**

compute the hedging strategy $(\bar{\delta}_{n,m}^1, \bar{\delta}_{n,m}^2)$ for $n = 0, \dots, N$, the payoff $\varphi(S_T^1)_m$, and the hedging error $-\varphi(S_T^1)_m + (\bar{\delta} \cdot S_m)_n^N + C_n^N(\bar{\delta})$

end for

if $\varepsilon > 0$ **then**

return estimate $\bar{p}_0 = \rho(-\varphi(S_T^1) + (\bar{\delta} \cdot S)_0^N + C_0^N(\bar{\delta})) e^{-rT}$

else

return estimate $\bar{p}_0 = \rho(-\varphi(S_T^1) + (\bar{\delta} \cdot S)_0^N) e^{-rT}$ and, for $n = 1, \dots, N - 1$, the quantiles of $S_{t_n}^1, I_{t_n}, \bar{\delta}_n^1$ and $\bar{\delta}_n^2$ based on the out of sample simulations

end if

Algorithm 2 Backward Hedging for American options

Set up volatility tree \mathcal{T} by following [17]
if $\varepsilon > 0$ **then**
 run this algorithm with $\varepsilon = 0$ to obtain the parameters for the 5D grids \mathcal{G}_n for $n = 1, \dots, N - 1$
else
 simulate M_0 random paths $(S_{n,m}^1, I_{n,m})$ for $k = 1, \dots, N$
 for $n = N - 1, \dots, 0$ **do**
 compute quantiles from $(S_{n,m}^1, I_{n,m})$ and set the 3D grid \mathcal{G}_n from quantiles and \mathcal{T}
 end for
end if
for $n = N - 1, \dots, 0$ **do**
 for all points G of \mathcal{G}_n (if $\varepsilon = 0$ then $G = (s, i, v)$ else $G = (s, i, v, d^1, d^2)$) **do**
 simulate M random paths $(S_{k,m}^1, S_{k,m}^2, I_{k,m}, V_{k,m})$ for $k = n + 1, \dots, N$, according to $S_{t_n}^1 = s, I_{t_n} = i, V_{t_n} = v$
 while the Simplex Search method has not converged **do**
 change $(\delta_{n,m}^1, \delta_{n,m}^2)$ according to the Simplex Search method
 for $m = 1, \dots, M$ (for each simulated path) **do**
 compute worst hedging stopping time $\tau_{n,m}$ by comparing future payoff $\varphi(S_{k,m}^1)$ and hedging cost \bar{p}_k (estimated by interpolation)
 for $k = n + 1, \dots, \tau_{n,m}/\Delta t$ **do**
 compute the hedging strategy $(\bar{\delta}_{k,m}^1, \bar{\delta}_{k,m}^2)$ by linear interpolation over \mathcal{G}_k (given $(\delta_{k-1,m}^1, \delta_{k-1,m}^2)$ if $\varepsilon > 0$)
 end for
 compute the hedging error $-\varphi(S_{\tau_{n,m},m}^1) + (\bar{\delta} \cdot S_m)_n^{\tau_{n,m}/\Delta t} + C_n^{\tau_{n,m}/\Delta t}(\bar{\delta})$
 end for
 estimate $\bar{p}_n = \rho\left(-\varphi(S_T^1) + (\bar{\delta} \cdot S)_n^{\tau_{n,m}/\Delta t} + C_n^{\tau_{n,m}/\Delta t}(\bar{\delta})\right) e^{-r(\tau_{n,m}-t_n)}$ from simulations
 end while
 end for
 if $n > 0$ **then**
 compute and store coefficients for linear interpolation of $\bar{\delta}_n^1, \bar{\delta}_n^2$ and \bar{p}_n over \mathcal{G}_n
 end if
end for
Simulate M_1 random paths $(S_{n,m}^1, S_{n,m}^2, I_{n,m}, V_{n,m})$ for $n = 1, \dots, N$
for $m = 1, \dots, M$ (for each simulated path) **do**
 compute worst hedging stopping time $\tau_{n,m}$ by comparing payoff and hedging cost \bar{p}_n (estimated by interpolation)

 for $k = n + 1, \dots, \tau_{n,m}/\Delta t$ **do**
 compute the hedging strategy $(\bar{\delta}_{k,m}^1, \bar{\delta}_{k,m}^2)$ by linear interpolation over \mathcal{G}_k (given $(\bar{\delta}_{k-1,m}^1, \bar{\delta}_{k-1,m}^2)$ if $\varepsilon > 0$)
 end for
 compute the hedging error $-\varphi(S_{\tau_{n,m},m}^1) + (\bar{\delta} \cdot S_m)_n^{\tau_{n,m}/\Delta t} + C_n^{\tau_{n,m}/\Delta t}(\bar{\delta})$
end for
if $\varepsilon > 0$ **then**
 return estimate $\bar{p}_0 = \rho\left(-\varphi(S_T^1) + (\bar{\delta} \cdot S)_0^{\tau_0/\Delta t} + C_0^{\tau_0/\Delta t}(\bar{\delta})\right) e^{-r\tau_0}$
else
 return estimate $\bar{p}_0 = \rho\left(-\varphi(S_0^1) + (\bar{\delta} \cdot S)_0^{\tau_0/\Delta t}\right) e^{-r\tau_0}$ and, for $n = 1, \dots, N - 1$, the quantiles of $S_{t_n}^1, I_{t_n}, \bar{\delta}_n^1$ and $\bar{\delta}_n^2$ based on the out of sample simulations
end if
