

Deep portfolio optimization with stocks and options

Kristoffer Andersson and **Kees Oosterlee**¹

¹Utrecht University

TopQuants, Nov. 1st, 2023

Agenda

Stochastic control, optimal portfolio selection
Our trading strategy with options, stocks and bonds
Deep learning algorithm
Numerical examples

Stochastic control problem

Mean variance (MV) optimal asset optimization is well-known.

We focus on realistic asset dynamics, as well as objective functions, in line with rational preferences of an investor.

We explore how options can be used as complement to risky assets and bonds to improve the performance, for general objective functions.

Represent the strategy with a sequence of neural networks, with as the loss function an empirical objective function.

Optimization is performed only once for the entire problem. Time-consistent and time-inconsistent problems can be treated similarly.

Adding options and gain exibility

A trader is allowed to trade in a riskfree bond, N^{stocks} stocks, and N^{options} options.

$S = (S_t)_{t \in [0; T]}$ is an $\mathbb{R}^{N^{\text{stocks}}}$ valued time-continuous Markov process on a complete probability space $(\Omega; \mathcal{F}; \mathbb{P})$.

The bond is $B = (B_t)_{t \in [0; T]}$ and for $i \in \{1; 2; \dots; N^{\text{options}}\}$, $V^i(t; S_t; K)$ is an option with S as underlying (single stock or a basket of stocks), at time $t \in [0; T]$, terminating at T , with $K \in \mathbb{R}$ the strike price.

We set the initial values to unity at $t = 0$, i.e., for $j \in \{1; 2; \dots; N^{\text{stocks}}\}$ and $i \in \{1; 2; \dots; N^{\text{options}}\}$, we set $S_0^j = 1$, $V^i(0; S_0) = 1$ and $B_0 = 1$.

Stochastic control problem

$k = (k_t)_{t \in [0; T]}$ is the process describing the amount in stock and when $k = 0$, the amount in the bond.

Total wealth of the portfolio stemming from the stock and the bond holdings

$$x_t = {}^0_t B_t + \sum_{k=1}^{N^{\text{stocks}}} k_t S_t^k = A_t^0 + \sum_{k=1}^{N^{\text{stocks}}} A_t^k \quad (1)$$

Since the portfolio is **self-financing**,

$${}^0_t = \frac{1}{B(t)} x(t) \sum_{k=1}^{N^{\text{stocks}}} k_{(t)} S_{(t)}^k ; \quad (2)$$

where $(t) = \max_{s \leq t} \tau_j$, i.e., the most recent trading date.

The **return on investment** is then given by

$$R_{SB}(S; k) = \frac{x_T}{x_0} \quad (3)$$

Stochastic control problem

Denote the **amount of option** in the portfolio by x^i . So,

$$y_t = \sum_{i=1}^{N^{\text{options}}} x^i V^i(t; S_t; K^i);$$

Return on the investment from the static option position:

$$R_O(S; x) = \frac{y_T - y_0}{y_0} \quad (4)$$

Summing up (3) and (4), we obtain the **total return**

$$R(S; x; \pi) = R_{SB}(S; \pi) + R_O(S; x) \quad (5)$$

Market frictions

We add **transaction costs** as well as a **non-bankruptcy constraint** and for the trading strategies, given by and , we introduce **leverage constraints**.

In discrete time, the value of the stocks and bond can then be re-written a

$$x_{t_{n+1}} = x_{t_n} + \int_{t_n}^{t_{n+1}} (B_{t_{n+1}} - B_{t_n}) + \sum_{k=1}^{N^{\text{stocks}}} \gamma^k (S_{t_{n+1}}^k - S_{t_n}^k); \quad (6)$$

The sum of the **transaction costs** for stock k :

$$TC^k = \sum_{n=1}^N C e^{r(T - t_n)} (\gamma_{t_n}^k - \gamma_{t_{n-1}}^k) S_{t_n}^k; \quad (7)$$

where $100 - C \in \mathbb{R}_+$ is a percentage of the size of the transaction.

We do not pay transaction costs immediately, but instead at the end of the trading period, with appropriate interest rate.

Constraints:

No-bankruptcy: When the stocks plus bond value is non-positive, the portfolio is liquidated.

$$x_{t_{n+1}} = x_{t_n} + I_{f_{x>0g}}(x_{t_n}) \left(B_{t_{n+1}} - B_{t_n} \right) + \sum_{k=1}^{N^{\text{stocks}}} \phi_{t_n}^k (S_{t_{n+1}}^k - S_{t_n}^k) ; \quad (8)$$

where $I_{f_{x>0g}}(\cdot)$ is the indicator function.

No short-selling of stocks, for $t \in [0; T]$ and $1 \leq k \leq N^{\text{stocks}}$, $\phi_t^k \geq 0$.

No leverage: We cannot short sell the bond, for $t \in [0; T]$, $\phi_t^0 \geq 0$.

No bankruptcy: If $x_{t_n} \leq 0$, all positions are liquidated and for $t > t_n$, $x_t = x_{t_n}$.

Positivity of the bond and the stocks part of the portfolio $x_0 \geq 0$.

Objective function

A **good quality objective function** is able to represent the investor's preferences of how much risk would be acceptable for a certain level of potential profit.

To **penalize downside risk**, we maximize the average of the 10% worst outcomes; to **encourage upside potential**, we maximize the average of the 10% best outcomes. Expected shortfall can be defined as Value at Risk

$$ES_p^+(R) = E[R | R \geq VaR_p(R)]; \quad ES_p^-(R) = E[R | R \leq VaR_p(R)];$$

A typical objective function would then be

$$U = E[R] - \alpha_1 \text{Var}[R] + \alpha_2 ES_{p_1}^-(R) + \alpha_3 ES_{p_2}^+(R); \quad (9)$$

with $\alpha_1; \alpha_2; \alpha_3 \geq 0$ describing the risk preference and $p_1, p_2 \in (0; 1)$ controlling the sizes of the left and right tails.

Stochastic control problem

Figure: Example of probability density function for terminal wealth. Red, blue, green represent the lower expected shortfall, mean and higher expected shortfall. Gray area is the mean plus/minus the variance.

Full optimization problem

So far, θ and ϕ are the **trading strategies**.

We add the **set of strike prices** $K = (K^1; K^2; \dots; K^{\text{options}})$, as part of the trading strategy, $\theta = (\theta; \phi; K)$.

Figure: Returns against stock value at terminal time \bar{T} . Left: Return for investing in stock, buying one unit of option 1 and selling one unit of option 2. Right: Returns for three different combinations of the products; the red line is the classical bull-call spread.

Stochastic control problem

With objective function $U(\cdot) = u[L[R(S; \cdot)]]$, initial wealth $x_0^{IC} \in \mathbb{R}_+$ and the allowed trading strategies (taking all trading constraints into account).

$$\begin{aligned}
 & \text{maximize } U(\cdot); \quad \text{where} \\
 & R(S; \cdot) = R_{SB}(S; \cdot) + R_O(S; \cdot) \\
 & R_{SB}(S; \cdot) = x_T(S; \cdot) - x_0(\cdot) - \sum_{k=1}^N \text{TC}^k; \quad R_O(S; \cdot) = y_T(S; \cdot) - y_0(\cdot); \\
 & x_T(S; \cdot) = x_0 + \sum_{n=0}^N \int_{f(x) > 0} x_{t_n}(S; \cdot) \left(B_{t_{n+1}} - B_{t_n} \right) + \sum_{k=1}^N \int_{t_n}^{t_{n+1}} \sigma^k(S_{t_{n+1}} - S_{t_n}); \\
 & x_0 = x_0^{IC} - y_0(\cdot); \quad y_T(S; \cdot) = \sum_{i=1}^{N^{\text{options}}} V^i(T; S_T; K^i); \\
 & y_0(\cdot) = \sum_{i=1}^{N^{\text{options}}} V^i(0; S_0; K^i);
 \end{aligned}$$

(10)

Stochastic control problem

Given S_0 and assuming known drift and diffusion and jump coefficients, and J , we employ,

$$dS_t = \mu(t; S_t)dt + \sigma(t; S_t)dW_t + J(t; S_t)dX_t;$$

where X_t represents a jump process.

Let $t_N = T$ and for $0 \leq i \leq N-1$, $t_i < t_{i+1}$, generate $M \geq N_+$ samples of the N^{stocks} dimensional asset process S . Asset k , realization m , at time $t_n \in T_N$ is $S_{t_n}^k(m)$, etc.

We use **empirical distributions** for $\mathbb{P}[R(S; \cdot)]$ in a Monte Carlo fashion.

Discrete scheme is approximated by letting **deep neural networks** represent trading strategies and optimizing with a gradient-descent algorithm.

Neural network approximation

The trading strategy π is represented by a **sequence of neural networks**.

A neural network is a mapping, $(\cdot; \cdot): \mathbb{R}^{D^{in}} \rightarrow \mathbb{R}^{D^{out}}$, with θ containing all **trainable parameters** of the network.

Number of layers is $L \geq 2$; for layer l , the **number of nodes** is $N_l \geq 2$.

The **weight matrix**, between $l-1$ and l , is $w_l \in \mathbb{R}^{N_l \times (N_{l-1} + 1)}$; the bias $b_l \in \mathbb{R}^{N_l}$;

The (scalar) **activation function** $a_l: \mathbb{R} \rightarrow \mathbb{R}$ and the vector activation function $\mathbf{a}_l: \mathbb{R}^{N_l} \rightarrow \mathbb{R}^{N_l}$, which, for $x = (x_1; x_2; \dots; x_{N_l})$, is defined by

$$\mathbf{a}_l(x) = \begin{pmatrix} 0 \\ a_l(x_1) \\ \vdots \\ a_l(x_{N_l}) \end{pmatrix} + \begin{pmatrix} 1 \\ \vdots \\ A \end{pmatrix} ;$$

-) The output of the network should obey the **trading constraints**, which are managed by choosing an **appropriate activation function** in the output layer

Neural networks representing the trading strategy

The trading strategy consists of three parts: i) the static amount invested in each option, ii) the static strike prices of the options K , and iii) the dynamic amount invested in each stock

and K are **decided at $t = 0$** and with a deterministic initial wealth x_0^{IC} , we have a deterministic representation for α and K .

α may **depend on previous performance**, which is affected by randomness through the stock process (a dynamic strategy).

For the dynamic trading strategy, we use a deep neural network taking the current wealth as input and outputs the stock allocation.

The admissible trading strategies are $\alpha^{NN} = f(\cdot; K; x_0; \dots; x_{N-1})g$, where $\alpha; \dots; x_{N-1}$, may depend on the stock.

Optimization problem with neural networks

8

maximize $\frac{1}{2} \sum_{NN} U^M(\cdot)$; where M i.i.d. random variables are distributed according to

$$R(S; \cdot) = R_{SB}(S; \cdot) + R_O(S; \cdot)$$

$$R_{SB}(S; \cdot) = \sum_{k=1}^{N^{stocks}} \lambda_{t_N}^k x_0^k TC^k; \quad R_O(S; \cdot) = \sum_{i=1}^{N^{options}} \gamma_{t_N}^i \gamma_0^i$$

$$\lambda_{t_N}^k = \lambda_0^k + \sum_{n=0}^{N-1} \lambda_{t_n}^k \int_{x>0} \lambda_{t_n}^k \hat{\Lambda}_n^0(B_{t_{n+1}}^k - B_{t_n}^k) + \sum_{k=1}^{N^{stocks}} \hat{\Lambda}_n^k(S_{t_{n+1}}^k - S_{t_n}^k); \quad \lambda_0^k = \lambda_0^{IC} \gamma_0^k$$

$$\gamma_{t_N}^i = \sum_{i=1}^{N^{options}} \lambda_i V^i(T; S_{t_N}; K^i); \quad \gamma_0^i = \sum_{i=1}^{N^{options}} \lambda_i V^i(0; S_0; K^i);$$

$$\hat{\Lambda}_0^k = \lambda_0^k \sum_{k=1}^{N^{stocks}} \hat{\Lambda}_0^k; \quad (\hat{\Lambda}_0^1, \dots, \hat{\Lambda}_0^{N^{stocks}})^\top = a^0(\lambda_0^0); \quad \hat{\Lambda} = a(\cdot); \quad K = a^K(\lambda^K);$$

$$\hat{\Lambda}_n^k = \frac{1}{B_{t_n}^k} \sum_{k=1}^{N^{stocks}} \lambda_n^k S_{t_n}^k; \quad (\hat{\Lambda}_n^1, \dots, \hat{\Lambda}_n^{N^{stocks}})^\top = (\lambda_{t_n}^k; \lambda_n^0):$$

(11)

General neural network settings

We use a **sequence of neural networks**, as tools to solve the problem.

The **number of training samples** $M_{\text{train}} = 2^{22}$, the batch size $M_{\text{batch}} = 2^{12}$, the number of epochs $M_{\text{epoch}} = 10$ and the number of layers $L = 4$.

For the **interior layers**, i.e., $2 \leq l \leq 3$, set the number of nodes $n_l = 20$ and the activation function $\sigma_l(\cdot) = \text{ReLU}(\cdot)$.

-) $D_{\text{input}} = 1$ and D_{output} , as well as the **activation function in the output layer**, depend on the trading constraints and are specified for each specific problem.

Initial learning rate is 0.01. After two batches, it decreases by a factor $\exp(-0.5)$ for each new batch.

Classical continuous mean-variance optimization

Classical MV problem: asset process is geometric Brownian motion. Trading is carried out **without transaction costs**, i.e., setting $C = 0$. There are **no constraints** and trading in the options is not allowed. The objective function is given by

$$U(x) = E[x_T] - \frac{\gamma}{2} \text{Var}[x_T]$$

where $\gamma > 0$ controls the risk aversion.

Closed-form expression for the optimal allocation as well as an optimal mean and variance of the terminal wealth $T = 2$; $N = 20$; $r = 0.06$; $\sigma = 0.17$

$$a = \begin{pmatrix} 0 \\ 0.08 \\ 0.07 \\ 0.06 \\ 0.05 \\ 0.04 \end{pmatrix}; \quad = \begin{pmatrix} 0 \\ 0.23 \\ 0.05 \\ 0.05 \\ 0.05 \\ 0.05 \end{pmatrix} \begin{pmatrix} 0.05 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.215 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.2 & 0.05 & 0.05 & 0.185 & 0.05 \\ 0.05 & 0.05 & 0.05 & 0.05 & 0.17 \end{pmatrix} \quad (12)$$

Stochastic control problem

The optimal value of the objective function is approximately 1637.

Figure: Upper: Convergence of the loss to the analytic counterpart with respect to number of training epochs. Comparison with reference solution. Lower: Comparison of the empirical pdfs and reference. Comparison of the empirical CDFs and the reference.

Beyond MV, with market frictions and jumps

Consider the **full generality** of the asset model, as well as transaction costs, no bankruptcy constraint and trading in European call and put options.

The parameter values **are reused** and $\sigma = 0:05$, $\mu = (0; \dots; 0)^T$,
 $\Sigma = \text{diag}(0:2; \dots; 0:2)$, $NB = 1$ and $C = 0:005$.

An interpretation of C is as a penalizing term for **too heavy reallocation** (which is something that for instance pension funds want to avoid).

$$U(\omega) = E[R] \omega_1 + \omega_2 \text{Var}[R] + \omega_3 \text{ES}_{p_1}^-(R) + \omega_4 \text{ES}_{p_2}^+(R):$$

$p_1 = 0:01$ i.e., we **penalize low values** of the expected return of the worst 1% performance of the portfolio. For the **upper tail**, we maximize the expected return of the 5% best outcomes $p_2 = 0:95$.

The weights are set to $\omega_1 = 0:552$, $\omega_2 = 0:276$ and $\omega_3 = 0:110$.

Problem specific neural network settings

We set $x_{\max} = 1$ (the **maximum amount of allocation** into the options is 100% of the initial wealth).

We use a slight modification of the **activation function** for this.

Then, the **option allocation range** is $[0, x_{\max}]$, while keeping the sum of the allocations into each option to $[0, x_{\max}]$.

For the strike prices, we use $K^{\text{low}} = (0.75; 0.75; \dots; 0.75)$ and $K^{\text{high}} = (1.25; 1.25; \dots; 1.25)$, i.e., setting the **range for strike prices** between 75% to 125% of the stock price at the initial time.

We set $x_n^{\text{low}} = 2x_n$ and $x_n^{\text{high}} = 2x_n$ implying that we can allocate into each stock **between 200% and 200%** of the total value of the stocks and bond.

Evaluation of the results

The algorithm returns a **dynamic strategy** for the bond and stocks, the static strategies for the allocations into the options and a strike for each option.

Figure: Left: Average allocation to stocks, bond and options over time. Right: Average allocation to stocks, bond and options over time for each stock. Asterisks and bull/bear symbols represent call and put option holdings, respectively.

Stochastic control problem

The **strike prices** are optimized by the neural networks to 0.75 for all call options and 1.25 for all put options, e., deep in the money.

For the best performing outcomes, the main option contribution comes from call options; for the worst performing outcomes from put options.

Figure: Contribution to the portfolios for terminal wealth less than 1.03 (33% of the outcomes), between 1.03 and 1.12 (41%), and above 1.12 (26%).

Stochastic control problem

With options, we observe **i)** a thinner left tail, **ii)** a higher density around the expected terminal wealth, and **iii)** a fatter right tail.

The first and last items are **beneficial** since the objective function aims to prevent large losses (by the lower expected shortfall term) and encourages large gains (by the upper expected shortfall term).

Regarding a comparison with the MV-strategy, in contrast to our strategies we encounter a **fatter left than the right tail**, which is non-desirable.

Stochastic control problem

For all measures, but the variance, the **portfolio with options** performs best. By the strategy with options, **transaction costs decrease** by 60% compared to the strategy without options and by 90% compared to the MV-strategy. This is beneficial since **less aggressive** re-allocation is desirable for a fund.

	$E[R]$	$\text{Var}[R]$	$ES_{p_1}^+(R)$	$ES_{p_2}(R)$	$U(\cdot)$	Tr. cost
With options	1.146	0.081	0.971	2.18	1.61	0.386%
Without options	1.140	0.045	0.931	1.93	1.58	1.01%
MV strategy	1.146	0.077	-0.208	1.48	1.21	3.98%

Table: For the MV-strategy, λ is set to make the mean coincide with the mean obtained from the strategy with options. The trading cost is a percentage of the initial wealth.

Stochastic control problem

The market **does not** behave exactly as the model.

Test the algorithm's **robustness for model miss-specification**, applying the strategies with higher and lower volatility of the underlying asset process.

-) In the high volatility case, we multiply by two and in the low volatility case we divide by two.

Stochastic control problem

Most notable is the **lower expected shortfall** which expresses a loss of 172.8% and 98.4% and the trading costs at 10.1% and 2.74% for the higher and lower volatilities, respectively.

Options in the portfolio are **beneficial** when the volatility increases and less beneficial when the volatility decreases.

Variance is larger with options, due to the fatter right tail of the distribution.

	$E[R]$	$\text{Var}[R]$	$ES_{\rho_1}^+(R)$	$ES_{\rho_2}(R)$	$U(\cdot)$	Trading cost
Evaluation with increased volatility for the underlying assets ($\sqrt{2}$).						
With options	1.318	0.660	0.763	3.268	1.540	0.838%
Without options	1.350	0.175	0.734	2.635	1.532	1.23%
MV strategy	1.081	0.674	-0.728	1.460	0.668	10.1%
Evaluation with decreased volatility for the underlying assets ($\sqrt{0.5}$).						
With options	1.074	0.0262	0.969	1.620	1.506	0.261%
Without options	1.143	0.0160	0.957	1.548	1.569	0.636%
MV strategy	1.163	0.0569	0.0160	1.487	1.301	2.74%

Table: Comparison of three strategies. For the MV-strategy, ρ is set to make the mean coincide with the mean obtained from the strategy with options. The trading cost, as percentage of initial wealth reflects the volatility of portfolio re-allocations.

Conclusions

The choice of objective function should reflect the **true incentives** of a rational trader

Adding options makes shaping of the distribution of the terminal wealth more flexible due to the asymmetric distribution of option returns.

Options significantly **reduce re-allocation** and in turn the trading cost;

A **sequence of neural networks** produces high quality allocation strategies in high dimensions (many assets, options and strike prices for each option).

Extension to trading options in a dynamic setting is straightforward if we have access to an efficient option valuation along stochastic asset trajectories.

