

CEIS Tor Vergata

RESEARCH PAPER SERIES

Vol. 21, Issue 1, No. 552 – February 2023

A Reinforcement Learning Algorithm for Trading Commodities

Federico Giorgi, Stefano Herzel and Paolo Pigato

A Reinforcement Learning Algorithm for Trading Commodities

FEDERICO GIORGI STEFANO HERZEL
PAOLO PIGATO

Department of Economics and Finance
University of Rome Tor Vergata

29/05/2023

Abstract. We propose a Reinforcement Learning (RL) algorithm for generating a trading strategy in a realistic setting, that includes transaction costs and factors driving the asset dynamics. We benchmark our algorithm against the analytical optimal solution, available when factors are linear and transaction costs are quadratic, showing that RL is able to mimic the optimal strategy. Then we consider a more realistic setting, including non-linear dynamics, that better describes the WTI spot prices time series. For these more general dynamics, an optimal strategy is not known and RL becomes a viable alternative. We show that on synthetic data generated from WTI spot prices, the RL agent outperforms a trader that linearizes the model to apply the theoretical optimal strategy.

Keywords: Portfolio Optimization, Reinforcement Learning, SARSA, Commodities, Threshold Models.

1 Introduction

We propose an algorithm, based on Reinforcement Learning (RL), to determine an effective trading strategy for a dynamic model that includes transaction costs and where the asset price dynamics is driven by factors that can be either internal, like the average of previous returns, or external, like the level of a market index or its volatility. We tackle the classical problem of optimal portfolio allocation, initiated with the seminal work by Markowitz, who formulated the problem in a static setting under mean-variance utility as a quadratic programming. Markowitz result was extended to a dynamic setting by Merton [17], who considered a model where prices follow a geometric Brownian motion and solved the associated dynamic optimization problem through the Bellman Equation. Later on, Gârleanu and Pedersen [12] (henceforth GP), provided an analytical solution to a problem where prices follow a diffusion and are linearly predicted by mean reverting factors, in a market where transactions costs are quadratic in the number of shares traded. In their paper, GP also apply their solution to determine the optimal investment strategy on a portfolio of futures on commodities, showing that, by calibrating their model in-sample and by assuming quadratic transaction costs, their strategy significantly outperform, in terms of Sharpe ratio, the Markowitz-Merton one.

The GP result is a significant theoretical breakthrough but their model assumptions are not always realistic. For instance, predicting factors, albeit present, may not be observable, their relation to price changes may be non-linear, transaction costs may not be quadratic, price dynamics may not be well represented by a diffusion. To deal with more realistic settings, where an analytical solution is not known, we propose to use RL, as it is often effective under general assumptions (see Sutton and Barto [28] for a general discussion). In RL, an agent is interacting with the environment over time. At each time, the agent observes the state of the environment and chooses an action; at the following time, the agent observes a new state of the environment and receives a reward for its action. RL algorithms provide a way to train the agent by letting it learn through positive reinforcement with the goal of maximizing its reward. The agent does not know the dynamics of the environment and learns how to improve its strategy from the observation of a large number of trials.

Our algorithm combines the standard SARSA (State-Action-Reward-State-Action) with a neural network to estimate the value function on a continuous set of states (current shares, market realization) and actions (shares to trade). The value function is estimated by generating consecutive batches of episodes: within each batch, the agent trained on previous batches is initially used, and then its choices are gradually improved via simulations of new episodes. RL algorithms, even if effective in most cases, are not guaranteed to work on any possible instance of the problem at hand. For this, reason we perform a comparative test in a situation where the optimal solution is known. The application of the algorithm to crude oil West Texas Intermediate (WTI) shows that RL may bring some statistically significant advantages over the current best theoretical choice represented by the GP model. The comparison of several models shows that the best fit is attained for a specification of the factor dynamics that is not linear and hence does not fit GP modeling assumptions. This is a clear example of an instance where the RL approach may present some advantage over the theoretical model.

We adopt a value-based algorithm, where an approximation to the optimal value function

is computed on simulations. Value-based algorithms have been originally developed in a tabular fashion for discrete and finite state-action spaces (refer to [28] for more insights). In finance, several strategies have been adopted in order to generalize to the continuous case. In [2], tabular methods have been extended to the continuous case by means of kernel approximation methods. Value-based algorithms have been used for example by [15] and [6] for hedging derivatives and by [23] for optimal trading, as they are easy to implement for arbitrarily complex reward functions. On the negative side they are sensitive to hyperparameters and they need a large number of simulations to converge (see [19]).

A different approach, called “policy search method”, optimizes directly on the policies space. These methods have been applied to finance in [3], to compare the performances of an agent using RL to one that adopts the GP formula, in [18] and [33] to produce a trading algorithm and in [4] and [30] for hedging derivatives. In [10], RL is used to construct real-time trading strategies in the presence of trading signals. They are considered to be robust on noisy datasets and fast converging. Furthermore, it is possible to build algorithms with safety guarantees that ensure an improvement of the policy at each update (see [13, 21, 25, 26]). However, policy search methods are often complex to implement and very objective-specific (see [9] for a survey).

In Section 2 we revisit the RL theory and describe the modified SARSA algorithm. In Section 3 we compare the trained algorithm to the GP optimal strategy for trading with transaction costs. In Section 4 we calibrate several non-linear models to the WTI time-series and combine the RL algorithm with the most realistic one. We prove by simulation that, in this case, the RL outperforms GP. In Section 5 we conclude the paper and propose possible developments.

2 The Reinforcement Learning algorithm

Let us introduce the basic ideas and definitions on RL. The RL agent learns how to interact with an environment by receiving some rewards for its actions: it observes the state of the environment s_t , selects an action a_t , receives a reward R_{t+1} and observes a new state s_{t+1} as a consequence, possibly random, of s_t and a_t . The procedure is then repeated until a final time T (possibly infinite). The sequence of states, action and rewards is stochastic, and the mathematical framework used for modeling such decision making problem is provided by a Markov decision process (MDP). A MDP is a Markov process that is embedded with the possibility for an agent to interact with its evolution through its actions. In a MDP, the law of the time $t + 1$ state and reward does not depend on the full path of state, action and rewards up to t , but only on the most recent state and action

$$\mathcal{L}(s_{t+1}, R_{t+1} | s_0, a_0, R_1, s_1, a_1, \dots, R_t, s_t, a_t) = \mathcal{L}(s_{t+1}, R_{t+1} | s_t, a_t) \quad (1)$$

where $\mathcal{L}(X|y)$ denotes the law of the random variable X conditional on the observation of $Y = y$. The policy, or strategy, followed by the agent of selecting an action after observing a state is expressed as a function $\pi(s) = a$. For any initial probability distribution on the states s_0 , the transition probability $\mathcal{L}(s_{t+1}, R_{t+1} | s_t, a_t)$ (1), together with the policy π completely determines the progress of the system through a *homogeneous Markov chain*. Refer to [8] for details on how to generalize the setting to time-changing environments. The agent’s goal is to maximize the

expectation of the cumulative (discounted) rewards, that is to determine the optimal policy

$$\pi^* = \operatorname{argmax}_{\pi} \mathbb{E}^{\pi} \left[\sum_{k \geq 0} \gamma^k R_{t+k+1} \right] \quad (2)$$

where \mathbb{E}^{π} is the expectation operator based on the MDP distribution (1) and following the policy π , while γ is a discount factor.

The state-action value function of a policy π is defined as the expected value of cumulative rewards when the initial state is s , the selected action is a and the policy π is followed,

$$q_{\pi}(s, a) = \mathbb{E}^{\pi} \left[\sum_{k \geq 0} \gamma^k R_{t+k+1} | s_t = s, a_t = a \right]. \quad (3)$$

Observe that, due to the stationarity of the policy, both the problem definition (2) and the action-value function (3) are time independent.

All optimal policies share the same optimal action-value function $q_*(s, a) = \max_{\pi} q_{\pi}(s, a)$, which satisfies the Bellman optimality equations (see [28] for a proof)

$$q_*(s, a) = \mathbb{E} \left[R_{t+1} + \gamma \max_{a_{t+1} \in \mathcal{A}(s_{t+1})} q_*(s_{t+1}, a_{t+1}) | s_t = s, a_t = a \right] \quad (4)$$

where $\mathcal{A}(s_{t+1})$ determines the state-dependent action space. To each action-value function $(s, a) \rightarrow q(s, a)$ it is possible to associate its q -greedy policy. Such policy consists in choosing the action a that, in a given state s , maximizes the action-value function (3)

$$\pi_{greedy}^q(s) = \operatorname{argmax}_{a \in \mathcal{A}(s)} q(s, a) \quad (5)$$

The solution to (5) does not need to exist or to be unique, depending on the nature of the RL problem¹. From the Bellman equation (4), the greedy policy (5) associated to the optimal value function q_* is a solution to the RL problem (2), or $\pi^* = \pi_{greedy}^{q^*}$. In turn, the Bellman equation (4) can be expressed as

$$q_*(s, a) = \mathbb{E} \left[R_{t+1} + \gamma q_*(s_{t+1}, \pi_{greedy}^{q^*}(s_{t+1})) | s_t = s, a_t = a \right] \quad (6)$$

The main idea behind value-based algorithms is to approximate the solution to (6) by an iterative procedure

$$q_{k+1}(s, a) = \mathbb{E} \left[R_{t+1} + \gamma q_k(s_{t+1}, \pi_{greedy}^{q_k}(s_{t+1})) | s_t = s, a_t = a \right], \quad k = 0, 1, \dots \quad (7)$$

This is called ‘‘iterative policy evaluation’’ and its convergence to the value function is a consequence of the Fixed-Value Theorem. Therefore, an approximation to the optimal policy can be obtained

¹Consider, for instance, the the action-value function $q(s, a) = a^2$ on the unbounded action space $\mathcal{A}(s) = \mathbb{R}$, in which case the maximum does not exist, or on the symmetric action space $\mathcal{A}(s) = [-A, A]$, in which case the maximum is not unique

by a sequence of greedy policies $\pi_{greedy}^{q_k}$, a procedure called “iterative policy improvement”. Refer to [28] and references therein.

The estimation of the optimal value function is typically done by simulations, where the MDP assumption is used to generate a large number of paths, or “episodes”, which are then used to approximate the solution of the Bellman equation (4).

Indeed, ideas and definitions introduced so far are well known in dynamic programming. A relevant difference between RL and dynamic programming is that the RL agent does not need to know the distribution of the MDP and learns the optimal strategy after repeated interactions with the environment. In RL, the process of learning about the environment, and how it reacts to different actions, is called exploration, while the process of following the greedy policy associated to the current estimate of the state-value function is called exploitation. RL algorithms alternate between exploration and exploitation.

In particular the RL algorithm adopted in this paper belongs to the class of SARSA methods, proposed by [24], and refined by [28]. The name SARSA derives from the fact that the value function is estimated by successive simulations of the sequence State-Action-Reward-State-Action. To select the action after observing the state, the SARSA algorithm balances exploration and exploitation adopting the ϵ -greedy policy (of a value function q) defined as

$$\pi_{\epsilon\text{-greedy}}^q(s) = \begin{cases} \tilde{a}, & \text{if } u < \epsilon \\ \operatorname{argmax}_a q(s, a), & u \geq \epsilon \end{cases} \quad (8)$$

where $\epsilon \leq 1$ is a positive constant, u is uniformly sampled on $(0, 1)$ and \tilde{a} is randomly chosen according to the uniform distribution over $\mathcal{A}(s)$. The $\pi_{\epsilon\text{-greedy}}^q$ policy selects a random action with probability ϵ (exploration) and coincides with the π_{greedy}^q policy with probability $1 - \epsilon$ (exploitation). At the first step of the algorithm the value of ϵ is large, so that more weight is assigned to exploration, and after each step of the algorithm it is decreased towards zero so that the policy becomes more exploitative (closer to the greedy one).

The name SARSA derives from the fact that the value function is estimated by successive simulations of the sequence State-Action-Reward-State-Action $s_0^{(j)}, a_0^{(j)}, r_1^{(j)}, s_1^{(j)}, a_1^{(j)}, \dots$ as produced by the MDP (1), which are then used to estimate the optimal value function q^* by iteratively solving the Bellman equation (4) by means of the SARSA updating formula

$$q_{k+1}(s_{t-1}^{(j)}, a_{t-1}^{(j)}) = q_k(s_{t-1}^{(j)}, a_{t-1}^{(j)}) + \alpha \left(R_t^{(j)} + \gamma q_k(s_t^{(j)}, a_t^{(j)}) - q_k(s_{t-1}^{(j)}, a_{t-1}^{(j)}) \right) \quad (9)$$

where $\alpha \in [0, 1]$ is a “learning rate” that allows to fine tune the learning speed: setting it equal to 0 means that the value function is never updated, whereas setting it equal to 1 means that learning can occur quickly. The reader may notice that the full-strength SARSA updating formula (9) resembles the term in the expectation on the right-hand-side of the iteration formula (7). Originally, the SARSA algorithm was developed under the assumption of finite state space and action space \mathcal{A} . In such situation, it can be shown that, provided that enough episodes are simulated, then all the possible state-action couples (s, a) are updated often enough to reach convergence of the value function q_k on the entire grid. Under suitable conditions, the algorithm produces a sequence of value functions converging to the optimal one (see [28] and references therein.)

Since in our case the action space is a continuous set, we propose a modification of the standard SARSA algorithm that uses a supervised regressor to estimate the value function. An extensive

literature has been developed to extend the application of SARSA to continuous state-action spaces. One possibility, is to discretize the state-action spaces and then to apply the standard SARSA. With this approach, convergence is ensured from the results obtained for the finite case (see [28]). However, providing a discretization that is dense enough to capture the continuous features of states and actions and yet coarse enough to ensure reasonable computational effort is not an easy task. A more refined extension of SARSA to the continuous case has been studied in [20] by means of Newton’s Method for direct optimization of the state-action value function for action selection, in order to allow continuous action reinforcement learning without a separate policy function. In this paper we propose to use Neural Networks.

Let us now describe the algorithm, whose pseudo-code is shown in Table 1 . The training process, is performed on consecutive “batches” of episodes, obtained by simulating the MDP (1). At the beginning of the algorithm, an estimate for the value function $\hat{q}^{(0)}(\cdot, \cdot)$ must be initialized. Such initialization imposes a bias on the algorithm and can be tackled in different ways. The most agnostic approach is the zero initialization $\hat{q}^{(0)}(s, a) = 0$, reflecting the fact that no episode has yet been generated. A somewhat opposite approach is the “optimistic” initialization. In this case, the initial estimate is set to a constant, large value $\hat{q}^{(0)}(s, a) = Q \gg 0$; then, as the training proceeds, the estimate on the state-action pairs that are actually visited tends to decrease to its true value, leaving the large value for un-visited state-action pairs. Then, by following the greedy policy (5), the agent tends to select actions that are yet-to-be visited, thus encouraging exploration. In this setting, exploration is tuned by means of the ϵ -greedy policy (8), therefore we do not use the optimistic initialization approach. However, the zero initialization approach presents some technical difficulties in the definition of greedy policy (5): in fact, if the value function is constant, the action providing the maximum value is not well defined. Therefore, we choose an approach whereby the initial value function is set $\hat{q}^{(0)}(s, a) = 0$ and the ϵ parameter in the ϵ -greedy policy is set to $\epsilon_0 = 1$.

In the exploration phase occurring in the ϵ -greedy policy (8), the agent can explore the action space in different fashions by choosing different strategies to randomize the exploring action $\tilde{a} \in \mathcal{A}(s)$. The most agnostic way is to draw \tilde{a} uniformly on the action space $\mathcal{A}(s)$. Alternatively more probability weight can be assigned to areas where one believes better rewarding actions are located.

In the exploitation phase occurring in the ϵ -greedy policy (8), an optimization of the current estimate of the state-action value function $\hat{q}^{(n)}(s, a)$ must be performed on the action space $\mathcal{A}(s)$. In order to do so, we use the SHGO optimization algorithm, which has been proven to be appropriate for solving blackbox low-dimensional global optimization problems (refer to [11] for details).

The simulations of the episodes are used to obtain simulations of the SARSA updating formula scheme (9)

$$q^{(j)}(s_t^{(j)}, a_t^{(j)}) = \hat{q}^{(n-1)}(s_t^{(j)}, a_t^{(j)}) + \alpha \left(R_{t+1}^{(j)} + \gamma \hat{q}^{(n-1)}(s_{t+1}^{(j)}, a_{t+1}^{(j)}) - \hat{q}^{(n-1)}(s_t^{(j)}, a_t^{(j)}) \right) \quad (10)$$

At the end of each batch, a Neural Network interpolation $N(s, a; \boldsymbol{\theta}^*)$ is fitted of the relation $(s_t^{(j)}, a_t^{(j)}) \rightarrow q^{(j)}(s_t^{(j)}, a_t^{(j)})$ by using the grid of points generated within the batch. Then, a new state value function is produced by model averaging

$$\hat{q}^{(n)}(s, a) = \eta N(s, a; \boldsymbol{\theta}^*) + (1 - \eta) \hat{q}^{(n-1)}(s, a) \quad (11)$$

Set $n = 0$, the initial value function estimate $\hat{q}^{(0)}(\cdot, \cdot)$, the value of $\epsilon = \epsilon_0$, the value of the decreasing factor k , the model averaging parameter η , and the number of episodes J in each batch.

Repeat

1 Set $n = n + 1$

2 For $j = 1, \dots, J$

a Simulate state $s_0^{(j)}$, compute $a_0^{(j)} = \pi_{\epsilon\text{-greedy}}^{\hat{q}^{(n-1)}}(s_0^{(j)})$

b For $t = 1, \dots, T$

i Simulate the new state $s_t^{(j)}$ and compute the reward $R_t^{(j)}$

ii Compute the new action $a_t^{(j)} = \pi_{\epsilon\text{-greedy}}^{\hat{q}^{(n-1)}}(s_t^{(j)})$

iii Compute the estimate of the value function for the previous state and action

$$q^{(j)}(s_{t-1}^{(j)}, a_{t-1}^{(j)}) = \hat{q}^{(n-1)}(s_{t-1}^{(j)}, a_{t-1}^{(j)}) + \alpha \left(R_t^{(j)} + \gamma \hat{q}^{(n-1)}(s_t^{(j)}, a_t^{(j)}) - \hat{q}^{(n-1)}(s_{t-1}^{(j)}, a_{t-1}^{(j)}) \right)$$

3 Fit the Neural Network $N(s, a; \theta)$ over all the episodes of the batch

$$\theta^* = \operatorname{argmin}_{\theta} \sum_{t,j} \left(q^{(j)}(s_t^{(j)}, a_t^{(j)}) - N(s_t^{(j)}, a_t^{(j)}; \theta) \right)^2$$

4 Update the value function for the next batch

$$\hat{q}^{(n)}(s, a) = \eta N(s, a; \theta^*) + (1 - \eta) \hat{q}^{(n-1)}(s, a)$$

5 Set $\epsilon = k\epsilon$

Until convergence. Return $\hat{q}^{(n)}(\cdot, \cdot)$

Table 1: The pseudo-code for the SARSA algorithm.

where we set the parameter $\eta \in [0, 1]$ to assign the proper weight to previous estimates of the state value function. Finally, the exploration parameter ϵ is reduced at the end of each step to get convergence to the greedy policy. In our implementation², a Neural Network $N(s, a; \theta^*)$ with only 3 layers, with 64 hidden neurons on the first layer, 32 neurons on the second layer and 8 neurons on the last proved to be sufficient. The activation function defining each hidden neuron is the standard rectified linear unit, or ReLU (see [22]). The solver for the weights optimization is the standard Adam algorithm by [14].

²Source code available on GitHub at <https://github.com/dynamic-trading-RL/dynamic-trading>

3 The Garleanu-Pedersen model: Reinforcement Learning vs Dynamic Programming

To validate the procedure explained in the previous section, we benchmark our RL algorithm against the model proposed by GP in [12] for which the optimal strategy can be computed in closed form using dynamic programming. GP consider the prices p_t of S securities and the changes in prices in excess of the risk-free return

$$x_{t+1} = p_{t+1} - (1 + r^f)p_t.$$

The price changes are driven by a K -dimensional factor f as

$$x_{t+1} = \mu_r + Bf_t + u_{t+1} \tag{12}$$

$$f_{t+1} - f_t = \mu_f - \Phi f_t + \varepsilon_{t+1} \tag{13}$$

where Φ is a constant matrix (subject to conditions to make the process stationary) and the noises are Gaussian and uncorrelated, $u_t \sim N(0, \Sigma)$ i.i.d., $\varepsilon_t \sim N(0, \Omega)$ i.i.d., $u_t \perp \varepsilon_t$. The goal is to determine the strategy π , defined by the number of shares n_t hold in the portfolio at time t , maximizing the risk-cost penalized present value of the future expected gains

$$\pi^* = \operatorname{argmax}_{\pi} \mathbb{E}^{\pi} \left\{ \sum_{t=1}^{\infty} \gamma^{t-1} \left(\gamma(n'_{t-1}x_t - \frac{\kappa}{2}n'_{t-1}\Sigma n_{t-1}) - c(\Delta n_{t-1}) \right) \right\} \tag{14}$$

where γ and κ are constants representing the discount factor and the risk-aversion, $c(\Delta n)$ are the transaction costs paid when trading $\Delta n_{t-1} = n_t - n_{t-1}$ shares, and $(\cdot)'$ denotes the transpose.

The optimal strategy for quadratic trading costs $c(a) = \frac{\lambda}{2}a'\Sigma a^3$ satisfies the relation

$$n_t = (1 - \eta)n_{t-1} + \eta \times aim_t \tag{15}$$

where η is a constant depending on the parameters of the model and aim_t is a weighted combination of the expected values of the future Markowitz portfolios (see [16]) defined as

$$Markowitz_t := (\kappa\Sigma)^{-1}(\mu_f + Bf_t) \tag{16}$$

When trading costs are very high, that is when λ grows to infinite, the solution converges to the static strategy $n_t = n_{t-1}$; when there are no trading costs, or $\lambda \rightarrow 0$, the solution converges to the Markowitz solution.

For the RL algorithm we need to define state and actions spaces and reward function. The definition of the state space is very important, as it should contain all the relevant information needed by the agent to select an action, but to keep the computational workload feasible it should not be too large. The state variables should include the current time t , the price change x_t , the

³More precisely, Gârleanu and Pedersen solve the problem with the more general cost formulation $c(a) = \frac{1}{2}a'\Lambda a$, where Λ is a symmetric positive-definite matrix. Then, they specialize the discussion on the case where transaction costs are proportional to the P&L covariance Σ in terms of a cost friction parameter λ . In fact, such formulation provides intuitive results (as discussed below) as well as having micro-foundations.

$\mu_r = 0.007$	$B = -0.083$	$u_t \sim N(0, 1.349)$
$\mu_f = 0.001$	$\Phi = 0.228$	$\varepsilon_t \sim N(0, 0.100)$

Table 2: Maximum likelihood estimated parameters for the GP model (12)-(13).

observed factor f_t and the current number of shares n_{t-1} . The action variable is given by the number of shares traded

$$a_t = \Delta n_t = n_t - n_{t-1}. \quad (17)$$

The reward is obtained from the objective function (14) as

$$R_{t+1} = \gamma \left(n'_t x_{t+1} - \frac{\kappa}{2} n'_t \Sigma n_t \right) - c(\Delta n_t) \quad (18)$$

In order to get a bounded action space we introduce an upper bound on the absolute value of the shares

$$\sup_t \|n_t\|_\infty < M. \quad (19)$$

where $M \in \mathbb{R}_{>0}^S$ is a vector of positive bounds for each of the assets positions. Therefore we impose that, for all times t ,

$$|a_t^i| < M^i - n_{t+1}^i$$

for all $i = 1, \dots, S$, thereby considering a state-dependent action space. We remark that, by appropriately defining a lower and an upper bound on the trades, one can also include short sale or budget constraints in the RL formulation, generalizing the unconstrained GP solution (15).

To compare the performances of the RL algorithm to that of the theoretical optimum by GP, we consider the West Texas Intermediate (WTI) daily spot prices from Jan-1986 to Jul-2019⁴. We assume that the factor depends on previous observations of the process x , as could be believed by an investor following a momentum strategy

$$f_t = \frac{1}{5}(x_t + x_{t-1} + x_{t-2} + x_{t-3} + x_{t-4}). \quad (20)$$

In Figure 1a we display the WTI time series, in Figure 1b the corresponding factor time series computed as in (20). Maximum likelihood estimation of the GP model (12)-(13) returns the parameters in Table 2⁵. We consider a trading period of $T = 50$ days. Following GP in [12], we assume the costs parameter to be $\lambda = 0.015$ (which yields a cost for unit trade of $c(1) = \$0.01$), a risk aversion parameter $\kappa = 10^{-3}$ and a 2% continuously compounded annual rate, providing a daily discount factor $\gamma = e^{-0.02/252}$.

⁴Source: FRED (Federal Reserve Economic Data) - Economic Research Division - Federal Reserve Bank of St. Louis

⁵We fit all the models presented in this paper on the time period from 1988-05-17 to 2018-10-29; then, we test our strategies on the time period that was held out from the estimation, i.e. from 2018-10-30 to 2019-01-07.




Figure 1: WTI observed prices (1a) and increments of the corresponding factor (1b) computed as in (20).

We train the agent on consecutive batches; in each batch, $J = 15,000$ episodes are generated by using the model (12)-(13). The ϵ -greedy policy starts with a parameter $\epsilon = 1\%$, which is then decreased as $\epsilon \leftarrow \epsilon/3$ at each batch iteration.

Model averaging is also applied with $\eta = 0.5$ in (11), in order to avoid overfitting and get a more robust convergence, less dependent on the specific simulations of the current batch. The algorithm stops when the update of the value function is negligible, in the sense that $\|\hat{q}^{(n)}(\cdot, \cdot) - \hat{q}^{(n-1)}(\cdot, \cdot)\|$ is smaller than a pre-determined level. Table 3 displays the values $\bar{q}^{(n)}$ of the average of the value function at step n for a random sample of starting points extracted from the stationary distribution of the process in (12)-(13). Table 3 shows that the average of the value function is improving at each step of the algorithm.

To set the bound M in (19) so that the action space is sufficiently large to include the optimal strategy, we compute the trades of the Markowitz strategy (16) on simulated paths and set M as the 99.5 upper percentile of their absolute values. This is motivated by the fact that we expect the trades given by the Markowitz strategy (which does not take into account trading costs) to be an upper bound for the optimal strategy (15).

The maximization over the action space required by the ϵ -greedy policy (8) is performed by using the SHGO global optimization algorithm [11].

The left (right) panel of Figure 2 reports the shares hold at time t by the GP (RL) trader as a function of the values of the factor f_t and of the previous number of shares. We see how the strategy of the RL agent mimics the analytical optimal one. If we focus on the vertical line $n_{t-1} = 0$, we can see that, when the factor f_t is positive, the next step shares are negative, whereas when the factor f_t is negative, the next step shares are positive. This behavior is expected, since the estimated parameter $B = -0.083$ implies that positive factors predict negative price changes.

To backtest the strategies, we consider a trading period of T days on the historical time series

n	ϵ	$\bar{q}^{(n)}$
0	1.000%	-1,238.259
1	0.333%	5.445
2	0.111%	7.084
3	0.037%	7.790
4	0.012%	8.176
5	0.004%	8.425

Table 3: Progress of the RL algorithm: n is the number of batches, ϵ is the exploration parameter, $\bar{q}^{(n)}$ is the average of the value function at step n .

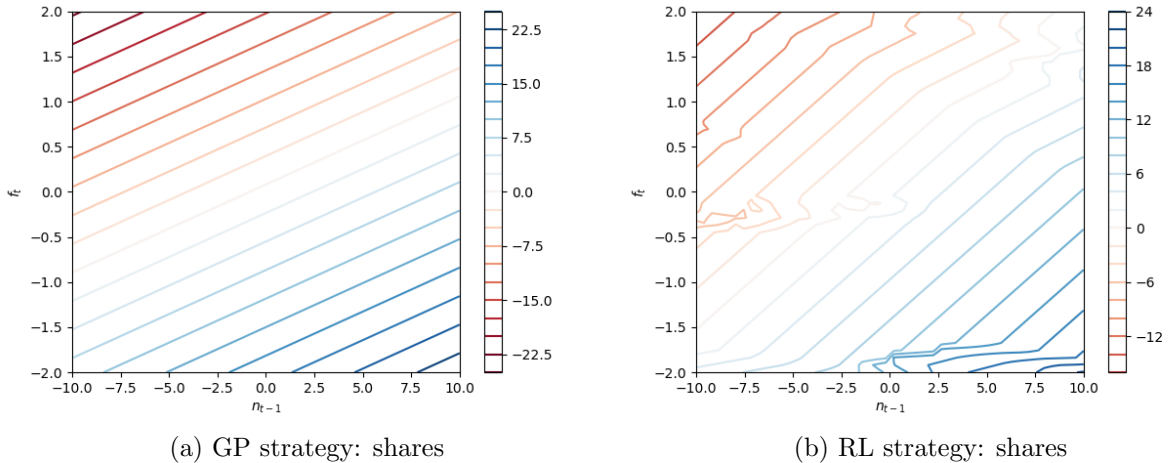


Figure 2: The number of shares n_t given the factor f_t and previous number of shares n_{t-1} according to GP (2a) and RL (2b).

and compare the three strategies GP, Markowitz and RL. For each strategy, we report in Figure 3b the realized cumulative risk-cost penalized present value of the future expected gains, which, following [23] or [15], we refer to as “wealth” (even if it should be noted that this is not a monetary quantity, but rather a total utility indicator)

$$w_t = \sum_{s=1}^t \gamma^{s-1} \left(\gamma(n'_{s-1}x_s - \frac{\kappa}{2}n'_{s-1}\Sigma n_{s-1}) - c(\Delta n_{s-1}) \right) \quad (21)$$

In Figure 3a we see that the Markowitz strategy shows wider oscillations in the number of traded shares than the optimal and RL strategies. This depends on the fact that Markowitz does not consider transaction costs, resulting in a lower cumulative wealth. We also see that the RL strategy follows closely the optimal strategy.

For a statistical comparison we simulate 10,000 paths and store the final wealths (21) obtained




Figure 3: Evolution of portfolio position (3a) and cumulative wealth (3b) according to Markowitz, GP and RL .

on each path. We get the following estimates for the sample mean and standard deviation

$$\begin{aligned} \mathbb{E}\{w_T^{GP}\} &= 11.24, & \mathbb{S}d\{w_T^{GP}\} &= 100.25, \\ \mathbb{E}\{w_T^{RL}\} &= 8.55, & \mathbb{S}d\{w_T^{RL}\} &= 101.88. \end{aligned}$$

To test the hypothesis that the expected final wealth of the optimal and RL agents are equal, we use a two-sided Welch's t -test

$$\begin{aligned} H_0 &: \mathbb{E}\{w_T^{RL}\} = \mathbb{E}\{w_T^{GP}\} \\ H_1 &: \mathbb{E}\{w_T^{RL}\} \neq \mathbb{E}\{w_T^{GP}\} \\ t\text{-statistic} &= -1.879, p = 0.060 \end{aligned}$$

The null hypothesis of equality between the final wealth attained by RL and GP is not rejected.

The histograms of the final wealths in Figure 4 visually confirm the statistical results, with the distributions of GP and RL being almost indistinguishable, while the Markowitz one exhibiting a larger variance and a smaller mean.

4 Reinforcement learning with nonlinear models

The recent empirical study [7] uses data on commodity futures transactions in the 2004-2013 period, including energy commodities, to uncover several stylized facts on commodity futures prices and volatilities, that clearly do not fit the GP framework. With this in mind, in what follows we determine a non-linear model that better fits the WTI and factors time series x_t and f_t (see Figure




Figure 4: Distribution of final wealths produced by Markowitz, GP and RL on 10,000 simulations

$\mu_r^{(0)} = 0.025$	$B^{(0)} = 0.014$	$u_t^{(0)} \sim N(0, 1.370)$
$\mu_r^{(1)} = 0.081$	$B^{(1)} = -0.276$	$u_t^{(1)} \sim N(0, 1.325)$

Table 4: Parameters of model (22), with $c = 0$, fitted to WTI time series.

1), and then compare the out-of-sample performances of a trading strategy produced by the RL algorithm trained on the non-linear model to those produced by GP.

We consider a non-linear model, where two different regimes are in force, depending on the value taken by the factor f_t

$$x_{t+1} = \begin{cases} \mu_r^{(0)} + B^{(0)} f_t + u_{t+1}^{(0)} & \text{if } f_t < c \\ \mu_r^{(1)} + B^{(1)} f_t + u_{t+1}^{(1)} & \text{if } f_t \geq c \end{cases}, \quad u_{t+1}^{(i)} \sim N(0, \Sigma^{(i)}) \quad (22)$$

where c is a constant threshold (that we set equal to 0) and $u_t^{(0)}, u_t^{(1)}$ are mutually independent. Calibrating the model (22) to the WTI time series and using the same factors as in (20) we get the parameters displayed in Table 4. Figure 5 represents the expected price change as a function of the observed factor for the linear and the non-linear model. The estimated slope of the linear model is $B = -0.083$, therefore a negative realization of the factor predicts a positive realization of x . Differently, the nonlinear model predicts negative values for all the occurrences of the factor, both

positive and negative, provided that their absolute value is sufficiently large. In particular, large positive factors predict price changes that are more negative than those predicted by the linear model.




Figure 5: Expected value $E[x_{t+1}|f_t]$ as a function of f_t , for the linear (13) and non-linear (22) model.

Let us turn our attention to the forecasting factors. Figure 1b shows that the factor time series is not homoskedastic and displays volatility clustering. To model it we propose five alternatives:

1. AR

$$f_{t+1} - f_t = \mu_f - \Phi f_t + \varepsilon_{t+1}$$

This is the dynamics assumed by GP. As it is well known, it does not support the homoskedasticity and volatility clustering that we observe in the data.

2. SETAR

$$f_{t+1} - f_t = \begin{cases} \mu_f^{(0)} - \Phi^{(0)} f_t + \varepsilon_{t+1}^{(0)} & \text{if } f_t < c \\ \mu_f^{(1)} - \Phi^{(1)} f_t + \varepsilon_{t+1}^{(1)} & \text{if } f_t \geq c \end{cases}, \quad \varepsilon_{t+1}^{(i)} \sim N(0, \Omega^{(i)}) \quad (23)$$

The SETAR (Self-Exciting Threshold AutoRegressive) model was proposed by [29]. Similarly to the AR model, it is mean-reverting. However, unlike the AR model, it allows different autoregression and volatility intensities, according to the position of the process at the previous step, above or below a given threshold c (that we set equal to zero in our calibration).

3. AR-TARCH

$$f_{t+1} - f_t = \mu_f - \Phi f_t + \varepsilon_{t+1} \quad (24)$$

$$\varepsilon_{t+1} = \sigma_{t+1} e_{t+1}, \quad e_{t+1} \sim N(0, 1) \quad (25)$$

$$\sigma_{t+1}^2 = \omega + \alpha \varepsilon_t^2 + \gamma \varepsilon_t^2 1_{\varepsilon_t < 0} + \beta \sigma_t^2 \quad (26)$$

The factor f follows the AR model with TARCH (Threshold ARCH) volatility (see [32]) In the TARCH model, the volatility process includes a threshold (that we set to 0) with an

Model	Log-likelihood	AIC	BIC
AR-TARCH	848.31	-1,684.61	-1,642.73
TARCH	419.78	-829.55	-794.65
GARCH	389.13	-770.27	-742.34
SETAR	-88.29	-182.58	-201.62
AR	-2,136.98	4,279.97	4,300.91

Table 5: In this table we report Log-likelihood, Akaike information criterion (AIC) and Bayesian information criterion (BIC) for the models considered. We can see how the best model in terms of likelihood is given by the AR model with TARCH residuals (26). In order to verify that the higher Log-likelihood is not a consequence of overfitting, we verify that the AIC and BIC for such model is the lowest among the models considered.

$\mu_f = 0.001$	$\Phi = 0.228$	$\omega = 0.002$
$\alpha = 0.200$	$\gamma = 0.010$	$\beta = 0.775$

Table 6: Parameters of model (26), with $c = 0$, fitted to WTI factors.

extra term that is activated only when the previous realization of the factor innovation is below the threshold. This model should be able to fit both the mean reversion on the factor process and the (threshold) volatility clustering on the volatility process.

4. TARCH. This is a sub-case of the $AR - TARCH$ model, with the parameter Φ set to zero, that is, the process f is not mean-reverting.
5. GARCH. This is also a sub-case of the $AR - TARCH$ model, with the parameters Φ and γ both set to zero. The GARCH model is ubiquitous in finance and is used to capture volatility clustering of time series by modeling the volatility process in an autoregressive fashion.

In Figure 6 we show the residuals obtained after fitting each model. The time series are clearly non-stationary for AR and SETAR models while look more stationary for TARCH, GARCH and AR-TARCH. Figure 7 shows the autocorrelation functions for the AR model and for the AR-TARCH model; as we can see, the AR provides residuals with stronger autocorrelation than the more refined AR-TARCH, which is able to capture volatility clustering. To chose between these models we apply the Akaike information criterion (AIC) [1], and the Bayesian information criterion (BIC) [27]. BIC and AIC are model selection criteria that rank the explanatory power of the models taking into account their likelihood as well as the number of parameters. We report in Table 5 the Log-likelihood, AIC and BIC for the models under consideration. The model to be preferred is the AR-TARCH, since its BIC and AIC are the lowest: therefore, in what follows we use AR-TARCH as the data generating process to train the RL algorithm. The best fit of model (26) with $c = 0$ on the factors time series (20) gives the parameters in Table 6.

Let us now compare the strategies of two traders. The first trader (GP) fits the model (12)-

(13), that is linear on the price changes and autoregressive on the factors, and then applies the GP solution (15). The second trader (RL) fits the non-linear model on the price changes (22) and the AR-TARCH (26) model on the factors and applies the RL algorithm.

In Figure 8 we show the policies defining the shares owned by the GP investor and by the RL-trained investor. Focusing on the vertical line $n_{t-1} = 0$, we see that the monotonic behavior shown in the GP solution (Figure 8a) is not preserved by the RL solution (Figure 8b). In fact, since $B^{(0)} = 0.014, B^{(1)} = -0.276$, both positive and negative factors predict negative price changes. For this reason it appears that the allocation strategy of the RL trader takes on less extreme values than that of the GP trader.

In Figure 9 we see the results of backtesting obtained by applying the policies of the three agents on the WTI time series from November 2018 to January 2019. In Figure 9a we see that allocation policy of the RL trader is more static than the one of the GP trader, in Figure 9b that it results in a larger cumulative wealth (21).

For a statistical comparison we generate data from the calibrated model. In Figure 10 we visualize the distribution of the final cumulative wealth on 10,000 simulated AR - TARCH paths, according to the three strategies. The resulting statistics is reported below

$$\begin{aligned} \mathbb{E}\{w_T^{GP}\} &= 6.37, & \mathbb{Sd}\{w_T^{GP}\} &= 77.51, \\ \mathbb{E}\{w_T^{RL}\} &= 11.52, & \mathbb{Sd}\{w_T^{RL}\} &= 94.69. \end{aligned}$$

We see that RL trader outperforms the GP trader in terms of wealth; this is confirmed by the one-sided Welch's t -test, which results in rejecting the null hypothesis of underperformance.

$$\begin{aligned} H_0 &: \mathbb{E}\{w_T^{RL}\} \leq \mathbb{E}\{w_T^{GP}\} \\ H_1 &: \mathbb{E}\{w_T^{RL}\} > \mathbb{E}\{w_T^{GP}\} \\ t\text{-statistic} &= 4.206, p < 10^{-3} \end{aligned}$$

5 Conclusion

We constructed an RL algorithm, combining the SARSA approach with neural networks, applicable to a continuous state-action space, to determine a trading strategy for a dynamic model that includes transaction costs.

We first considered a setting with linear factors and quadratic transaction costs, where an analytical optimal solution is known from [12], finding that the RL approach is able to recover the optimal strategy. Then, we considered a setting with more realistic, non-linear dynamics, where a theoretical optimal solution is not known and the RL approach can still be applied, since it only requires to be able to simulate the model. We provided statistical evidence that, on synthetic data generated from WTI spot prices, the RL agent outperforms a trader that linearizes the model to apply the theoretical optimal strategy.

The main advantage of using RL is its flexibility, as it works under any model dynamics. The main disadvantage is that, to train it in a satisfactory way, it needs a significant number of trials. The problem with simulating these trials is not only the computational burden, but also that, in this setting, the user needs to specify a parametric model to produce the data. We think that an

important improvement on the methodology proposed here will be achieved when new simulated data can be produced in a non-parametric way, so that the approach becomes really “model-free”. The use of signatures [5] and/or Generating Adversarial Network [31] are promising tools in this direction.

Our algorithm produces an approximation to the optimal value function by employing a modification of the SARSA method to consider a continuous state-action space. A promising different approach, used in a similar setting by [3], is based on policy approximation methods.

The selection of the predicting factor is key for a successful use of our methodology. In this respect, the empirical analysis [7] points out that commodity volatility is strongly related with volatility in other markets, and in particular to stock market volatility. This may be taken into account by including an implied or realized volatility index as a factor.


Acknowledgements. The authors gratefully acknowledge financial support from University of Rome Tor Vergata via the Project AIF - E83C22002610005. We thank Riccardo Cogo, Giovanni Nappi, Andrea Cosentini and Ranieri Dugo for valuable inputs, and the participants of the XXIII Workshop on Quantitative Finance, the EFMA Meeting 2022, the AMASES Conference 2022 for their feedback. We are very grateful to an anonymous reviewer of “Applied Stochastic Models in Business and Industry” for his/her careful reading and valuable suggestions.

References


- [1] Hirotogu Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike*, pages 199–213. Springer, 1998.
- [2] Francesco Bertoluzzo and Marco Corazza. Testing different reinforcement learning configurations for financial trading: Introduction and applications. *Procedia Economics and Finance*, 3:68–77, 2012.
- [3] Alessio Brini and Daniele Tantari. Deep Reinforcement Trading with Predictable Returns. *arXiv e-prints*, page arXiv:2104.14683, April 2021.
- [4] Hans Buehler, Lukas Gonon, Josef Teichmann, and Ben Wood. Deep hedging. *Quantitative Finance*, 19(8):1271–1291, 2019.
- [5] Hans Buehler, Blanka Horvath, Terry Lyons, Imanol Perez Arribas, and Ben Wood. A Data-driven Market Simulator for Small Data Environments. *arXiv e-prints*, page arXiv:2006.14498, June 2020.
- [6] Jay Cao, Jacky Chen, John Hull, and Zissis Poulos. Deep hedging of derivatives using reinforcement learning. *The Journal of Financial Data Science*, 3(1):10–27, 2021.
- [7] Peter Christoffersen, Asger Lunde, and Kasper V. Olesen. Factor structure in commodity futures return and volatility. *Journal of Financial and Quantitative Analysis*, 54(3):1083–1115, 2019.

- [8] Balázs Csanád Csáji and László Monostori. Value function based reinforcement learning in changing markovian environments. *Journal of Machine Learning Research*, 9(8), 2008.
- [9] Marc Peter Deisenroth, Gerhard Neumann, Jan Peters, et al. A survey on policy search for robotics. *Foundations and trends in Robotics*, 2(1-2):388–403, 2013.
- [10] Yue Deng, Feng Bao, Youyong Kong, Zhiquan Ren, and Qionghai Dai. Deep direct reinforcement learning for financial signal representation and trading. *IEEE transactions on neural networks and learning systems*, 28(3):653–664, 2016.
- [11] Stefan C Endres, Carl Sandrock, and Walter W Focke. A simplicial homology algorithm for lipschitz optimisation. *Journal of Global Optimization*, 72(2):181–217, 2018.
- [12] N. Gârleanu and L. H. Pedersen. Dynamic trading with predictable returns and transaction costs. *The Journal of Finance*, 68(6):2309–2340, 2013.
- [13] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer, 2002.
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [15] Petter N Kolm and Gordon Ritter. Dynamic replication and hedging: A reinforcement learning approach. *The Journal of Financial Data Science*, 1(1):159–171, 2019.
- [16] Harry Markowitz. Portfolio selection in the journal of finance vol. 7. 1952.
- [17] Robert C Merton. An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, pages 867–887, 1973.
- [18] John Moody and Matthew Saffell. Learning to trade via direct reinforcement. *IEEE transactions on neural Networks*, 12(4):875–889, 2001.
- [19] Remi Munos, Leemon C Baird, and Andrew W Moore. Gradient descent approaches to neural-net-based solutions of the hamilton-jacobi-bellman equation. In *IJCNN'99. International Joint Conference on Neural Networks. Proceedings (Cat. No. 99CH36339)*, volume 3, pages 2152–2157. IEEE, 1999.
- [20] Barry D Nichols and Dimitris C Dracopoulos. Application of newton’s method to action selection in continuous state-and action-space reinforcement learning. ESANN, 2014.
- [21] Matteo Papini, Matteo Pirotta, and Marcello Restelli. Adaptive batch size for safe policy gradients. In *The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [22] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- [23] Gordon Ritter. Machine learning for trading. *Available at SSRN 3015609*, 2017.


- [24] Gavin A Rummery and Mahesan Niranjan. *On-line Q-learning using connectionist systems*, volume 37. University of Cambridge, Department of Engineering Cambridge, UK, 1994.
- [25] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- [26] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [27] Gideon Schwarz. Estimating the dimension of a model. *The annals of statistics*, pages 461–464, 1978.
- [28] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 2018.
- [29] Howell Tong. *Threshold models in non-linear time series analysis*, volume 21. Springer Science & Business Media, 2012.
- [30] Edoardo Vittori, Michele Trapletti, and Marcello Restelli. Option hedging with risk averse reinforcement learning. *arXiv preprint arXiv:2010.12245*, 2020.
- [31] Tianlin Xu, Li K Wenliang, Michael Munn, and Beatrice Acciaio. Cot-gan: Generating sequential data via causal optimal transport. *arXiv preprint arXiv:2006.08571*, 2020.
- [32] J. M. Zakoian. Threshold heteroskedastic models. *Unpublished paper, Institut National de la Statistique et des Etudes Economiques, Paris*, 1991.
- [33] Zihao Zhang, Stefan Zohren, and Stephen Roberts. Deep reinforcement learning for trading. *The Journal of Financial Data Science*, 2(2):25–40, 2020.




(a) AR residuals




(b) AR-TARCH residuals



(c) TARCH residuals




(d) GARCH residuals




(e) SETAR residuals

Figure 6: Time series of residuals ε_t for different non-linear models (6a) AR; (6b) AR-TARCH; (6c) TARCH; (6d) GARCH; (6e) SETAR




(a) AR(1) residuals autocorrelation




(b) AR(1)-TARCH(1, 1, 1) residuals autocorrelation

Figure 7: Autocorrelation functions of the time series of residuals ε_t for different non-linear models (6a) AR; (6b) AR-TARCH



(a) GP strategy: shares



(b) RL strategy: shares

Figure 8: Policy (number of shares) of the GP trader (8a) and RL trader (8b) as a function of the current values of factor and shares.

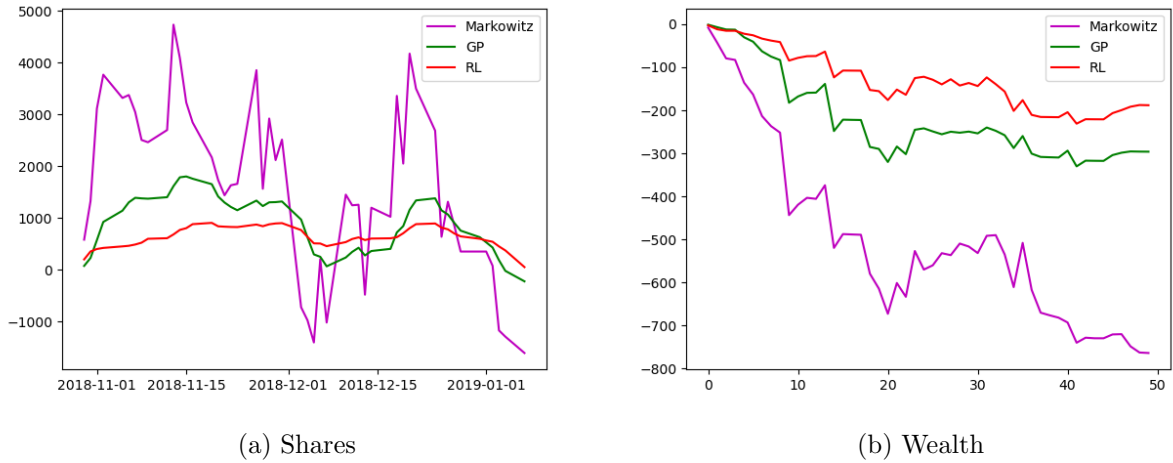


Figure 9: Evolution of the portfolio position according to Markowitz, GP and RL (9a) and evolution of the corresponding wealth 9b.

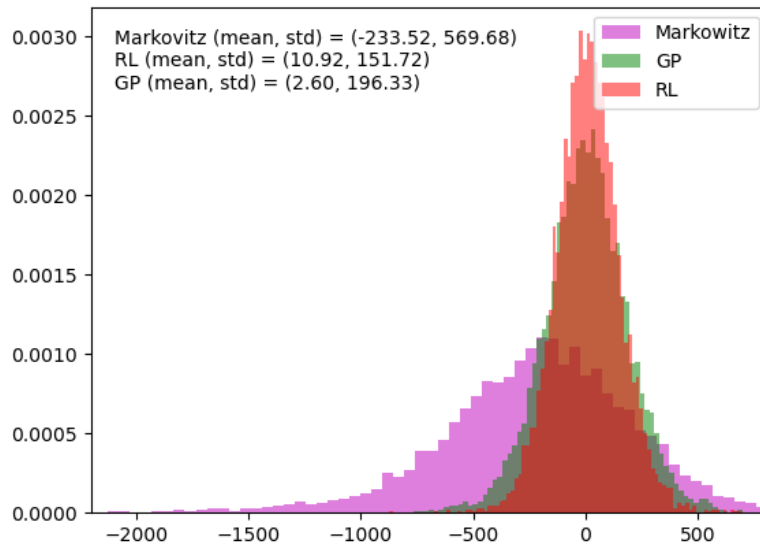


Figure 10: Distribution of the final wealth according to Markowitz, GP and RL on 10,000 simulations generated by the AR(1) - TARCH model, calibrated to the WTI time series.

RECENT PUBLICATIONS BY *CEIS Tor Vergata*

On the Role of Bargaining Power in Nash-in-Nash Bargaining: When More is Less

Bassam Fattouh, Beniamino Pisicoli and Pasquale Scaramozzino

CEIS Research Paper, 551 February 2023

On the Role of Bargaining Power in Nash-in-Nash Bargaining: When More is Less

Marc ESCRIHUELA-VILLAR, Walter Ferrarese and Alberto Iozzi

CEIS Research Paper, 550 December 2022

Traumatic Experiences Adversely Affect Life Cycle Labor Market Outcomes of the Next Generation - Evidence from WWII Nazi Raids.

Vincenzo Atella, Edoardo Di Porto, Joanna Kopinska and Maarten Lindeboom

CEIS Research Paper, 549 December 2022

Inter-municipal Cooperation in Public Procurement

Giampaolo Arachi, Debora Assisi, Berardino Cesi, Michele G. Giuranno and Felice Russo

CEIS Research Paper, 548, December 2022

Testing for Endogeneity of Irregular Sampling Schemes

Aleksey Kolokolov, Giulia Livieri and Davide Pirino

CEIS Research Paper, 547, December 2022

Aggregation Trees

Riccardo Di Francesco

CEIS Research Paper, 546, December 2022

Food Security during the COVID-19 Pandemic: the Impact of a Rural Development Program and Neighbourhood Spillover Effect in the Solomon Islands

Leonardo Becchetti, Sara Mancini and Sara Savastano

CEIS Research Paper, 545, December 2022

Attitude Towards Financial Planning of Italian Households

Marianna Brunetti, Rocco Ciciretti, Monica Gentile, Nadia Linciano and Paola Soccorso

CEIS Research Paper, 544, December 2022

Climate Policies, Macroprudential Regulation, and the Welfare Cost of Business Cycles

Barbara Annicchiarico, Marco Carli and Francesca Diluiso

CEIS Research Paper, 543, October 2022

The Social Value of Health: Amenable Deaths and Estimated the Gap with the Life Expectancy Frontier

Leonardo Becchetti, Gianluigi Conzo and Giovanni Trovato

CEIS Research Paper, 542, August 2022

DISTRIBUTION

Our publications are available online at www.ceistorvergata.it

DISCLAIMER

The opinions expressed in these publications are the authors' alone and therefore do not necessarily reflect the opinions of the supporters, staff, or boards of CEIS Tor Vergata.

COPYRIGHT

Copyright © 2023 by authors. All rights reserved. No part of this publication may be reproduced in any manner whatsoever without written permission except in the case of brief passages quoted in critical articles and reviews.

MEDIA INQUIRIES AND INFORMATION

For media inquiries, please contact Barbara Piazzi at +39 06 72595652/01 or by e-mail at piazzi@ceis.uniroma2.it. Our web site, www.ceistorvergata.it, contains more information about Center's events, publications, and staff.

DEVELOPMENT AND SUPPORT

For information about contributing to CEIS Tor Vergata, please contact at +39 06 72595601 or by e-mail at segr.ceis@economia.uniroma2.it