2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

doi 10.26089/NumMet.v27r107

Optimal investment in S&P 500 using SDDP and the implied-calibrated ARMA–GARCH model

Petr A. Arbuzov
Lomonosov Moscow State University,
Faculty of Computational Mathematics and Cybernetics,
Moscow, Russian Federation
ORCID: 0009-0005-7864-509X, e-mail: arbuzov.parb@gmail.com

Dmitry Yu. Golembiovsky
Lomonosov Moscow State University,
Faculty of Computational Mathematics and Cybernetics,
Moscow, Russian Federation
ORCID: 0000-0002-1848-5988, e-mail: dgolembiovskiy@yandex.ru

Abstract: We study a dynamic portfolio optimization problem in which probabilistic forecasts of the S&P 500 index are derived from option market prices. To overcome the limitations of classical approaches to reconstructing implied density from option prices which fail to generate conditional multi-period return distributions, we propose a method that involves calibrating the discrete-time ARMA–GARCH model from the observed call option prices in a risk-neutral measure with subsequent transition to a physical measure using a representative-agent framework. The calibrated model provides conditional multi-period distributions of asset returns, which are used to construct scenario lattices in multi-stage stochastic optimization. The resulting portfolio optimization problem is formulated as a multi-stage stochastic programming problem. At each stage, a weighted combination of the expected negative objective value for the next stage and the Conditional Value-at-Risk (CVaR) is minimized. The optimization is performed by means of the Stochastic Dual Dynamic Programming (SDDP) method. Historical simulations over the period 2019–2023 demonstrate that the proposed option-calibrated ARMA–GARCH–SDDP method consistently outperforms benchmark approaches based on static implied densities, equal-probability scenarios, and buy-and-hold investment. The results underscore the economic value of using option-implied information in portfolio management.

Keywords: portfolio optimization, conditional Value-at-Risk (CVaR), stochastic dual dynamic programming (SDDP), ARMA–GARCH models, option-implied forecasting, scenario generation.

For citation: P. A. Arbuzov, D. Yu. Golembiovsky, “Optimal investment in S&P 500 using SDDP and the implied-calibrated ARMA–GARCH Model,” Numerical Methods and Programming. **27** (1), 90–100 (2026). doi 10.26089/NumMet.v27r107.

2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

Оптимальные инвестиции в S&P 500 с использованием SDDP и откалиброванной по подразумеваемым параметрам модели ARMA–GARCH

П. А. Арбузов
Московский государственный университет имени М. В. Ломоносова,
факультет вычислительной математики и кибернетики,
Москва, Российская Федерация
ORCID: 0009-0005-7864-509X, e-mail: arbuzov.parb@gmail.com

Д. Ю. Голембиовский
Московский государственный университет имени М. В. Ломоносова,
факультет вычислительной математики и кибернетики,
Москва, Российская Федерация
ORCID: 0000-0002-1848-5988, e-mail: dgolembiovskiy@yandex.ru

Аннотация: Работа посвящена изучению задачи динамической оптимизации портфеля, в которой вероятностные прогнозы индекса S&P 500 получаются на основе рыночных цен опционов. Чтобы преодолеть ограничения классических подходов к восстановлению подразумеваемой плотности по ценам опционов, которые не позволяют генерировать условные многопериодные распределения доходностей, мы предлагаем метод, включающий калибровку дискретно-временной модели ARMA–GARCH по наблюдаемым ценам на колл-опционы в риск-нейтральной вероятностной мере и последующий переход к физической мере с использованием степенной функции полезности агента. Калиброванная модель предоставляет условные многопериодные распределения доходностей активов, которые используются для построения сценарных решеток в многоэтапной стохастической оптимизации. Полученная задача оптимизации портфеля формулируется как многоэтапная задача стохастического программирования. На каждом этапе минимизируется взвешенная комбинация ожидаемого отрицательного значения целевой функции для следующего этапа и условное значение Value-at-Risk (CVaR). Оптимизация выполняется с помощью метода стохастического двойственного динамического программирования (SDDP). Исторические симуляции за период 2019–2023 демонстрируют, что предлагаемый метод ARMA–GARCH–SDDP, калиброванный по опционам, последовательно превосходит бенчмарк-подходы, основанные на статических подразумеваемых плотностях, сценариях с равными вероятностями и инвестициях buy-and-hold. Результаты подчеркивают экономическую ценность учета информации, извлекаемой из опционов при управлении портфелем.

Ключевые слова: портфельная оптимизация, условные ожидаемые потери (conditional Value-at-Risk, CVaR), двойственное стохастическое динамическое программирование, модели ARMA–GARCH, прогнозирование на основе цен опционов, генерация сценариев.

Для цитирования: Арбузов П.А., Голембиовский Д.Ю. Оптимальные инвестиции в S&P 500 с использованием SDDP и откалиброванной по подразумеваемым параметрам модели ARMA–GARCH // Вычислительные методы и программирование. 2026. **27**, № 1. 90–100. doi 10.26089/NumMet.v27r107.

1. Introduction. Forecasting financial asset prices is a core task in financial mathematics, as it provides the foundation for risk management, portfolio optimization, and derivatives pricing. However, market option prices frequently deviate from those predicted by theoretical models such as the Black-Scholes model [1]. This discrepancy suggests that market participants adhere to their own implicit view of the future distribution of the underlying asset's value, incorporating expectations about volatility and tail risks [2, 3].

Research on forecasting underlying asset prices using option prices originated with the seminal work of Breeden and Litzenberger [4]. They demonstrated that the risk-neutral probability density function for the option's expiration time can be extracted directly from option prices. This finding revealed that option prices contain crucial information about the future asset distribution that sets the market's expectations, providing a foundation for subsequent research.

Market option prices can be used to derive the implied volatility of the underlying asset, which often differs from historical volatility. The set of implied volatilities for different strikes forms the volatility smile [2]. To construct a smooth risk-neutral density of the future asset price using the Breeden–Litzenberger method, the implied volatility smile is typically interpolated (e.g., using cubic splines or polynomials). These interpolated volatilities are then converted into option prices by means of the Black-Scholes formula for untraded strikes. This approach ensures effective interpolation, leading to smoother and more accurate estimated densities than those obtained from direct interpolation of option prices [3, 5]. Finally, assuming a power-law utility function for the representative agent, the real-world (physical) distribution of the future asset value can be derived from the corresponding risk-neutral distribution [3, 6, 7].

However, these approaches suffer from two major limitations. First, they can estimate the probability density of the underlying asset's price only at specific future dates corresponding to option expirations, leaving other time points unaddressed. Second, and more critically, even for these expiration dates, the methods do not provide conditional distributions. In other words, if for example the asset price becomes known at some intermediate point, the model cannot describe how this information alters the future price distribution at the next expirations.

An alternative that resolves both issues is to model the underlying asset with a discrete-time ARMA–GARCH process, calibrated according to the observed option prices. This model explicitly defines the evolution of the price path over time. Accordingly, the model can produce probability distributions not only for expiration dates but for any intermediate date as well. It naturally provides conditional distributions given the price at any prior point, thereby incorporating the forward-looking information from the options [8].

A natural alternative is the Heston stochastic volatility model, which is widely used for option pricing and implied volatility modeling [9]. However, empirical evidence shows that Heston-type diffusions often struggle to accurately fit short-maturity options and deep out-of-the-money contracts, where implied volatility smiles exhibit pronounced skewness and kurtosis [10, 11]. In contrast, discrete-time ARMA–GARCH models provide greater flexibility in capturing short-term return dynamics and volatility clustering, making them particularly suitable for short-dated option calibration and conditional multi-period forecasting.

Probabilistic forecasts from option-calibrated ARMA–GARCH models serve as inputs for stochastic optimization in portfolio management. We formulate the problem of dynamic management of the S&P 500 index fund position as a multi-stage stochastic programming problem, aimed at optimizing sequential decisions under risk constraints while accounting for return uncertainty. This problem is solved using the stochastic dual dynamic programming (SDDP) [12]. The method efficiently approximates the value functions through forward and backward passes over a scenario lattice, making it particularly suited to handle the high-dimensional uncertainty inherent in our multi-period forecasting setting [13].

In the second section, we address the calibration of the risk-neutral ARMA–GARCH model. We consider this calibration as a stochastic optimization problem, which we solve using the Randomized Stochastic Projected Gradient-Free (RSPGF) algorithm [14]. This section also describes in detail the subsequent transition to the real-world probability measure.

In the third section, we formulate the management of the S&P 500 index fund position as the SDDP problem. We describe the theoretical setup and describe details of the algorithm for constructing the scenario lattice, and conduct historical simulations using relevant market data to compare the performance of various portfolio strategies.

In conclusion, we discuss the empirical results and explain why our approach outperforms other portfolio strategies.

2. Calibration of the ARMA–GARCH model for the underlying asset using option market prices. The ARMA–GARCH model eliminates the key limitations of the Breeden–Litzenberger approach, which recovers risk-neutral densities only at specific expiration dates and cannot provide conditional distributions at intermediate price levels. In contrast, ARMA–GARCH captures the dynamic evolution of asset prices by modeling autocorrelation in returns and volatility clustering. This capability enables probabilistic forecasts on any future horizon, depending on intermediate market states. Consequently, it is well-suited for dynamic portfolio management and risk assessment.

2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

2.1. Calibration using the RSPGF algorithm. The ARMA(1,1)–GARCH(1,1) model for the logarithmic returns $Y_t = \ln(S_t/S_{t-1})$, where S_t is the asset price at time t, is defined by the following system of equations:

$$Y_t = m_t + \epsilon_t, \quad \epsilon_t = \sqrt{h_t}\varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1),$$
$$m_t = \varphi_0 + \varphi_1 Y_{t-1} + \theta_1 \epsilon_{t-1},$$
$$h_t = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta_1 h_{t-1}.$$

Here m_t is the conditional mean, ϵ_t is the innovation term, h_t is the conditional variance, and $\varepsilon_t \sim \mathcal{N}(0,1)$ denotes a standard normal random variable with stationarity conditions $|\varphi_1 + \theta_1| < 1$ and $\alpha_1 + \beta_1 < 1$.

To ensure no-arbitrage option pricing, the ARMA–GARCH model must be adapted to the risk-neutral measure $\mathbb{Q}$. According to the approach of Danilishin [15], we employ a linear approximation of the asset's log-returns. According to this approximation, the risk-neutral dynamics of the ARMA(1,1)–GARCH(1,1) model are specified as:

$$\widetilde{Y}_t = e^{r\tau} - 1 + \sqrt{\widetilde{h}_t}\left(\frac{e^{r\tau}}{1+\widetilde{m}_t}\right)\varepsilon_t, \quad \varepsilon_t \sim \text{i.i.d. } \mathcal{N}(0,1),$$

where $\widetilde{Y}_t$ is the linearized log-return, r denotes the risk-free rate, τ is the time step (in years), $\widetilde{m}_t$ and $\widetilde{h}_t$ are the conditional mean and variance under $\mathbb{Q}$ respectively, and i.i.d. stands for independent and identically distributed. This enables Monte Carlo simulation for pricing European-style options within the ARMA–GARCH framework. For the purposes of this study, we focus exclusively on call options.

The model calibration is formulated as a stochastic optimization problem that minimizes a loss function measuring the discrepancy between model-generated and market-observed option prices. The chosen loss function is the expected relative pricing error:

$$f(\boldsymbol{x}) = \mathbb{E}_\xi \left[\sum_{t\in T}\sum_{i\in\mathcal{C}_t} \frac{|e^{-rt}\max(S_t - X_{t,i}, 0) - C_{t,i}|}{C_{t,i}}\right]. \tag{1}$$

In this formulation, $\boldsymbol{x} = (\varphi_0, \varphi_1, \theta_1, \alpha_0, \alpha_1, \beta_1)$ is the vector of model parameters, the expectation $\mathbb{E}_\xi$ is calculated over random trajectories ξ of the underlying asset, T is the set of expiration dates for options, S_t denotes the simulated asset price at the option's expiration t, $C_{t,i}$ is the market price of the call option i expiring at the time t, $X_{t,i}$ is the strike price of this option, $\mathcal{C}_t$ is the set of call options expiring at the time t.

Thus, the calibration task is to find

$$\min_{\boldsymbol{x}\in\Omega} f(\boldsymbol{x}),$$

where Ω is a convex compact set specifying model stationarity and practical parameter bounds. It is defined by the following constraints:

$$|\varphi_1 + \theta_1| \leqslant 1 - \delta_A, \qquad \alpha_1 + \beta_1 \leqslant 1 - \delta_G,$$
$$\alpha_0 \geqslant \delta_\alpha, \qquad \alpha_1 \geqslant 0, \qquad \beta_1 \geqslant 0,$$
$$|\varphi_0| \leqslant C, \qquad |\varphi_1| \leqslant C, \qquad |\theta_1| \leqslant C,$$

with small constants $\delta_A, \delta_G, \delta_\alpha > 0$ and $C > 0$ ensuring the compactness of Ω.

Due to the complex structure of the model, which excludes the use of analytical gradients, the calibration procedure is based on the RSPGF algorithm. In this method, the loss function at step k is approximated by means of Monte Carlo simulation with N simulated return paths $\xi_{k,i}$:

$$\hat{f}_k(\boldsymbol{x}) = \frac{1}{N}\sum_{i=1}^{N} f(\boldsymbol{x}, \xi_{k,i}).$$

This estimator is unbiased, and its variance remains bounded provided the underlying return process is stationary [14].

The gradients required for optimization are approximated using a random smoothing technique with Gaussian perturbations. This method constructs a smoothed version of the objective function:

$$f_\mu(\boldsymbol{x}) = \mathbb{E}_v[f(\boldsymbol{x} + \mu v)],$$

2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

where the smoothing parameter $\mu > 0$ controls the perturbation magnitude, and $v \sim \mathcal{N}(0, I_n)$ is an n-dimensional standard Gaussian vector (I_n denotes the $n \times n$ identity matrix). A stochastic gradient at the point $\boldsymbol{x}_k$ is then computed using a finite-difference-like estimator:

$$G_\mu(\boldsymbol{x}_k, \xi_k, v) = \frac{\hat{f}(\boldsymbol{x}_k + \mu v, \xi_k) - \hat{f}(\boldsymbol{x}_k, \xi_k)}{\mu}\, v,$$

where $\hat{f}(\cdot, \xi_k)$ is the Monte Carlo estimator of the loss from the previous step. To improve the convergence properties of the algorithm, we average m_k independent gradient estimates at each iteration k:

$$G_{\mu,k} = \frac{1}{m_k} \sum_{j=1}^{m_k} G_\mu(\boldsymbol{x}_k, \xi_{k,j}, v_{k,j}),$$

where $\xi_{k,j}$ and $v_{k,j}$ are independent samples of the asset trajectories and Gaussian perturbations, respectively. This averaging reduces the variance of the gradient direction, which is a key factor for stable convergence as shown in [16].

The RSPGF algorithm proceeds iteratively. Starting from an initial point $\boldsymbol{x}_1 \in \Omega$, the parameters are updated according to the projected gradient step:

$$\boldsymbol{x}_{k+1} = \pi_\Omega(\boldsymbol{x}_k - \lambda_k G_{\mu,k}),$$

where π_Ω denotes the Euclidean projection onto the feasible set Ω, and $\lambda_k > 0$ is a step size. Under standard assumptions, the algorithm is guaranteed to converge to an approximate stationary point. According to the definitions from [14], a random iterate $\boldsymbol{x}$ is called a (ϵ, Λ)-stationary point if it satisfies the condition

$$\mathbb{P}\big(\|g_\Omega(\boldsymbol{x})\|^2 \leqslant \epsilon\big) \geqslant 1 - \Lambda,$$

where $\mathbb{P}$ denotes probability, $g_\Omega(\boldsymbol{x})$ is the projected gradient mapping onto the feasible set Ω defined by the stationarity and parameter bound constraints above, $\epsilon > 0$ is the accuracy parameter, and $\Lambda \in (0, 1)$ is the confidence parameter. Let σ^2 be an upper bound for the variance of the stochastic gradient approximation $G_{\mu,k}$. Then, the RSPGF algorithm converges to an (ϵ, Λ)-stationary point at a rate of [14]

$$\mathcal{O}\left(\frac{n \ln(1/\Lambda)\sigma^2}{\epsilon}\left(\frac{1}{\epsilon} + \frac{\ln(1/\Lambda)}{\Lambda}\right)\right).$$

Sufficient conditions for the convergence of the algorithm were established in [14], and later in [17] it was shown that these conditions are fulfilled in the context of calibrating an ARMA(1,1)–GARCH(1,1) model based on market option prices.

Under the assumption of a representative agent with a power-law utility function characterized by constant relative risk aversion, the transformation from the risk-neutral measure $\mathbb{Q}$ back to the physical (real-world) measure $\mathbb{P}$ can be explicitly derived [18]. Maintaining the specification of standard normal innovations ($\varepsilon_t \sim \mathcal{N}(0, 1)$), the dynamics under $\mathbb{P}$ corresponding to the risk-neutral model (1) are given by [8]:

$$\widetilde{Y}_t = e^{r\tau} - 1 + \gamma \widetilde{h}_t \left(\frac{e^{r\tau}}{1 + \widetilde{m}_t}\right)^2 + \sqrt{\widetilde{h}_t}\left(\frac{e^{r\tau}}{1 + \widetilde{m}_t}\right)\varepsilon_t. \tag{2}$$

2.2. Empirical calibration results. The model was calibrated using historical data for weekly S&P 500 index options traded on the Chicago Board Options Exchange (CBOE) from January 2019 to December 2023. The dataset comprised 86 trading dates, with a total of 14460 liquid call options spanning a wide range of strike prices and maturities (from 7 to 60 days). The risk-free rate was proxied by the 3-month LIBOR.

The calibration achieved an average relative pricing error (the absolute difference between model and market prices, divided by the market price) of 3.14%, with a maximum observed error of 10.8%.

Figure 1 illustrates market option prices compared to those computed using the Monte Carlo method with the calibrated ARMA–GARCH model on October 30, 2023. Legend on the right lists the expiration dates of the options.

The risk-neutral densities extracted from the calibrated ARMA–GARCH model reproduced those derived by means of the standard Breeden–Litzenberger method for corresponding expiration dates. Figure 2 shows an example of probability densities calculated on June 10, 2019 for the future underlying price on June 28, 2019.

2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

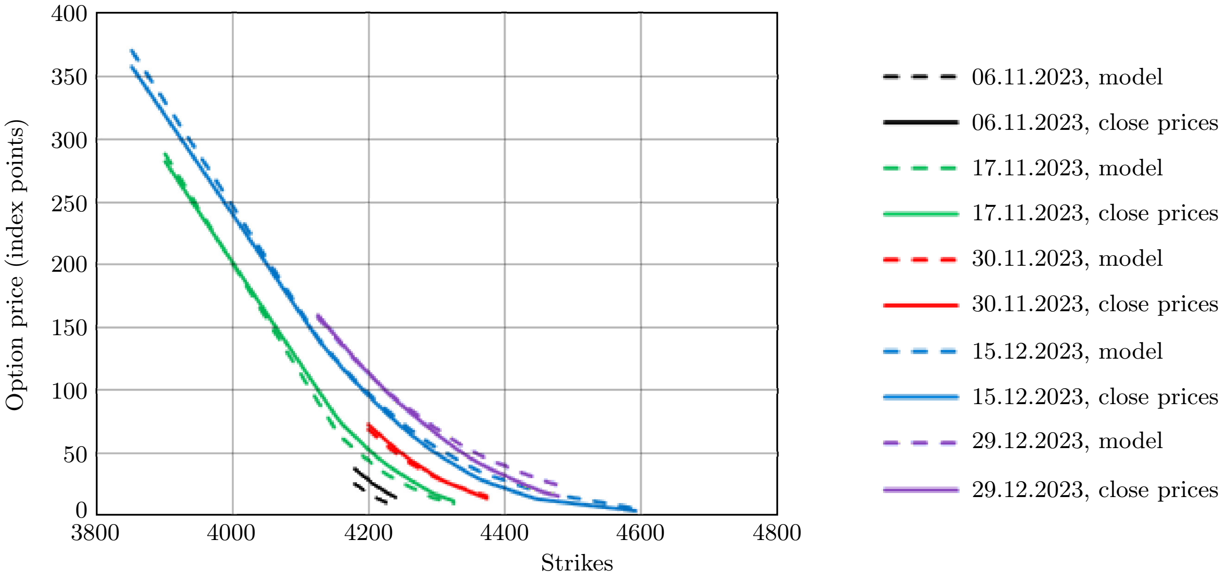

Figure 1. Market prices compared to model option prices as of October 30, 2023

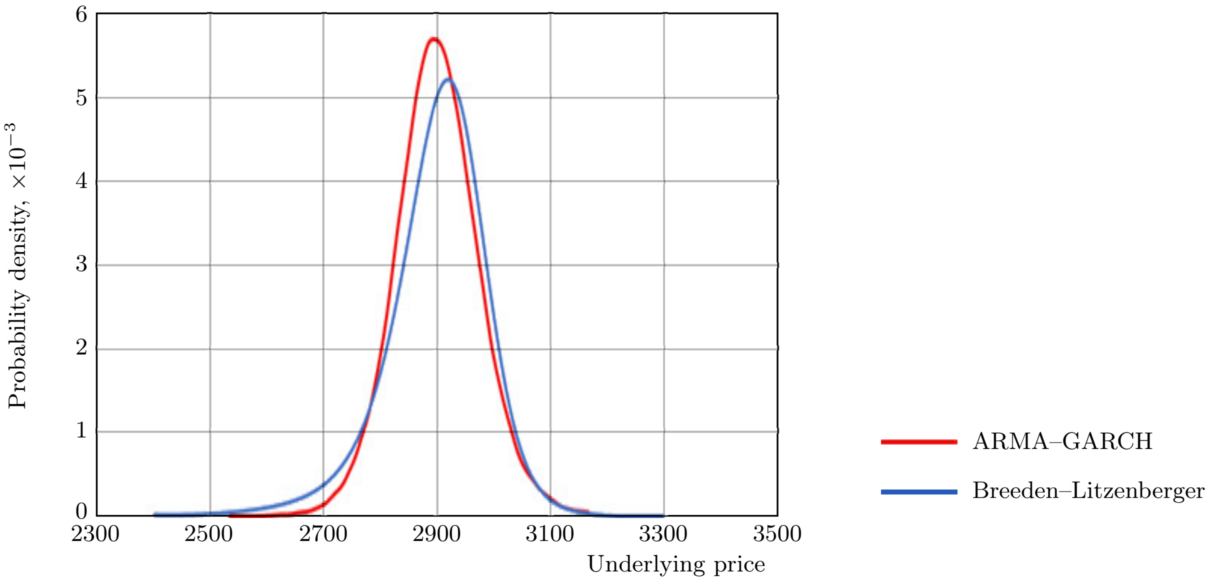

Figure 2. The Breeden–Litzenberger and ARMA–GARCH probability densities for the underlying asset price on June 28, 2019 estimated as of June 10, 2019

Finally, we applied the calibrated model to the real-world measure, incorporating risk aversion at different levels γ. The resulting simulated asset price trajectories reproduced the return characteristics and volatility dynamics of the historical behavior of the S&P 500 index. Figure 3 illustrates the historical dynamics of the S&P 500 index and scenarios generated using formula (2) with calibrated model parameters and $\gamma = 10$ starting from October 30, 2023.

Model forecast accuracy was evaluated using the Crnkovic–Drachman formula [19]. This approach treats each realized asset price as a quantile within the distribution predicted by the model for that date. For each forecast horizon (7, 10, 14, 18, 32, and 60 days), we computed these quantile levels using 1 million simulated price paths.

The resulting series of quantiles was then checked for uniformity using the Kolmogorov–Smirnov test and for independence using the Brock–Dechert–Scheinkman test [20]. Across all six horizons and for risk aversion

2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

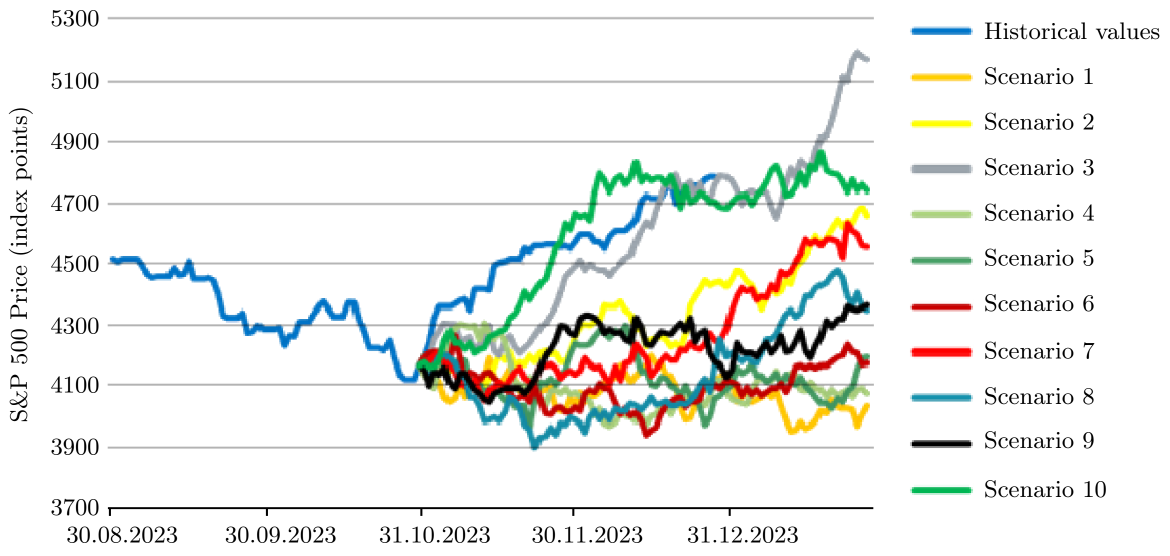

Figure 3. S&P 500 dynamics and generated scenarios

coefficient γ between 10 and 20, the Kolmogorov–Smirnov test yielded p-values well above the conventional 0.05 significance level (indicating no evidence against the null hypothesis of uniformity). This indicates that the realized price quantiles are statistically consistent with a uniform distribution, which does not allow rejection of the null hypothesis that the model's predictive distributions are well-calibrated [8].

3. Managing a position on the S&P 500 index fund. The calibrated ARMA–GARCH models provide a mechanism for generating realistic multi-period scenarios of the S&P 500 index dynamics, capturing essential stylized facts such as volatility clustering and asymmetry in returns. These scenarios can be used as direct inputs for solving practical stochastic optimization problems in portfolio management.

This section examines one such application: the dynamic management of a position on a hypothetical Exchange-Traded Fund (ETF) tracking the S&P 500 index. We formulate this problem as a multi-stage stochastic programming problem and solve it using the SDDP framework.

3.1. Stochastic dual dynamic programming formulation for index fund management. We formulate a position management problem for a hypothetical ETF that tracks the S&P 500 index, with an initial share price of \$2582 as of January 19, 2019. The portfolio consists of cash and ETF shares. Decisions on restructuring (purchase/sale) are made at separate stages. The aim is to minimize the weighted sum of the negative expected objective function value of the following stage and the CVaR term, thereby modeling a trade-off between return and risk [21]. Transaction costs are explicitly included in the formulation.

Our implementation of the SDDP algorithm [22] begins with the construction of a scenario lattice. In accordance with the methodology described in [23], this lattice is generated using simulations from the calibrated real-world ARMA–GARCH model.

It is assumed that ETF shares are infinitely divisible. The model is formulated in discrete time at stages $t = 0, \dots, T$. The stochastic price process is represented on a scenario lattice, where N_t denotes the set of nodes at stage t. For a given node $i \in N_t$, ω_i represents the price of the ETF. The state of the portfolio before a trading decision at node i is characterized by the pair (a_i, q_i), where a_i is the cash balance and q_i is the number of ETF shares held. The key control variable is q'_i, which denotes the number of shares held after rebalancing. This action leads to a post-decision state (a'_i, q'_i). The absolute trade size in number of shares is defined as $v_i = |q'_i - q_i|$.

For each node $i \in N_{T-1}$ (the stage immediately preceding the final horizon T), the following optimization criterion is applied:

$$V_i(a_i, q_i, \omega_i) = \min_{q'_i, v_i, \beta_i, y_j} \left[-(1-\lambda)\Big(q'_i \sum_{j \in i^+} p_{ij}\omega_j + a'_i\Big) + \lambda\Big(\beta_i + \frac{1}{1-\alpha}\sum_{j \in i^+} p_{ij} y_j\Big)\right],$$

2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

subject to:

$$v_i \geqslant q_i' - q_i,\ v_i \geqslant q_i - q_i', \tag{3}$$

$$q_i' \leqslant q_i + \frac{a_i}{\omega_i} - v_i c,\ -q_i' \leqslant q_i + \frac{a_i}{\omega_i} - v_i c, \tag{4}$$

$$a_i' = a_i - (q_i' - q_i)\,\omega_i - v_i \omega_i c, \tag{5}$$

$$a_i' + q_i' \omega_j \geqslant 0, \quad \forall j \in i^+, \tag{6}$$

$$y_j \geqslant 0,\ y_j \geqslant -\left(q_i' \omega_j + a_i'\right) - \beta_i, \quad \forall j \in i^+. \tag{7}$$

For nodes $i \in N_t,\ t = 0, \dots, T-2$:

$$V_i(a_i, q_i, \omega_i) = \min_{q_i', v_i, \beta_i, y_j} \left[-(1-\lambda)\Big(q_i' \sum_{j \in i^+} p_{ij}\omega_j + a_i' + \sum_{j \in i^+} p_{ij} V_j\Big) + \lambda\Big(\beta_i + \frac{1}{1-\alpha}\sum_{j \in i^+} p_{ij} y_j\Big)\right],$$

subject to the same constraints as above, except for:

$$y_j \geqslant -\left(q_i' \omega_j + a_i'\right) + V_j - \beta_i,\ \forall j \in i^+.$$

Here $\lambda \in (0, 1)$ is the risk aversion parameter, α is the confidence level for CVaR, and c is the proportional transaction cost, i^+ denotes the set of child nodes of node i in the scenario lattice, i.e. the nodes at stage $t+1$ reachable from node i at stage t, p_{ij} — transition probability between node i and its successor node j, β_i is a variable that at the optimum represents Value-at-Risk (VaR), and y_j are auxiliary variables used to linearize the CVaR computation. The second term in the square brackets corresponds to CVaR [24].

Inequalities (3) are the standard linear programming reformulation of the absolute value: they enforce $v_i \geqslant |q_i' - q_i|$, and since v_i incurs transaction costs in the objective, the optimum drives $v_i = |q_i' - q_i|$. Thus, v_i quantifies the absolute trade size. Condition (4) imposes budget constraints by restricting the post-trade position q_i' based on the available cash. Expression (5) defines the update of the cash balance a_i' accounting for the capital allocated to the trade and the incurred transaction costs. Condition (6) is a solvency constraint ensuring the portfolio's net value remains non-negative in every possible successor state. Finally, inequalities (7) provide the linear formulation necessary to integrate the CVaR measure into the optimization problem.

The formulated multi-stage stochastic optimization problem is solved using the SDDP.jl package, a Julia-based framework for stochastic dual dynamic programming [25].

3.2. Historical simulation of the S&P 500 index fund position management. To evaluate the practical effectiveness of the proposed approach, we conduct a historical trading simulation using S&P 500 data from January 14, 2019 to December 9, 2023. We compare the performance of the following 4 portfolio strategies:

1. ARMA–GARCH portfolio. The SDDP strategy is applied to scenario lattices constructed from price paths simulated by the ARMA–GARCH model, which is calibrated to market option prices. Stages in the lattice correspond to all consecutive business days in the planning horizon.
2. ARMA–GARCH–eq portfolio. This strategy uses the same lattice structure and price scenarios as the ARMA–GARCH portfolio, but assumes equal transition probabilities between all successor nodes at each stage.
3. Breeden–Litzenberger portfolio. This benchmark is based on the classic Breeden–Litzenberger method with smoothed volatility smiles, where risk-neutral densities are extracted directly from option prices without modeling the underlying time-series dynamics. Here each decision stage coincides with an option expiration date and all transition probabilities are equal.
4. Buy-and-hold portfolio. A passive benchmark, where all capital is invested in the index fund at the inception of the period and held without any subsequent rebalancing.

Scenario lattices were constructed sequentially with the lattice lengths that varied depending on the strategies. For the ARMA–GARCH and ARMA–GARCH–eq portfolios, the average lattice length was 32 stages (maximum — 60, minimum — 8). For the Breeden–Litzenberger portfolio, the average length was 7 stages (maximum — 17, minimum — 4), which corresponds to the intervals between option expiration dates. Each non-root stage in every lattice contained 20 nodes. The lattice for a given trading day was built by clustering 1 million price paths simulated by the Monte Carlo method.

2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

It is important to note that each lattice was generated using a model calibrated solely on the option prices observed on that specific trading day, with one calibrated model corresponding to one lattice. This design explicitly prevents look-ahead bias and ensures that future price realizations are never incorporated into past decision stages.

The simulation for all active strategies was configured with the following parameters: initial capital of $100,000 and risk-aversion coefficient $\gamma = 10$, which was selected based on the strongest calibration confidence across most horizons [8]. The remaining parameters were set to $\lambda = 0.5$ and CVaR confidence level $\alpha = 0.95$ (which are standard choices in risk-averse portfolio optimization), and a proportional transaction cost of $c = 0.009$. Preliminary experiments indicated that the model's performance is relatively insensitive to these hyperparameters, but highly dependent on the quality of model calibration.

The ARMA–GARCH and ARMA–GARCH–eq portfolios were rebalanced daily, while the Breeden–Litzenberger portfolio was adjusted only on option expiration days. The buy-and-hold portfolio, by design, remained static. The resulting wealth dynamics and the corresponding open ETF position values are illustrated in Figures 4 and 5.

4. Conclusion. The empirical analysis demonstrates that the ARMA–GARCH portfolio brings the highest average annual return (approximately 18% per year), which outperforms the results of the ARMA–GARCH–eq, Breeden–Litzenberger, and buy-and-hold strategies. These results underscore the value of incorporating option-implied information into the ARMA–GARCH framework, which enhances the model's ability to capture conditional return dynamics and volatility clustering.

A key factor contributing to the increased performance of the ARMA–GARCH portfolio is its ability to take short positions during rare high-volatility, downward-trending markets, notably in March 2020 and Septem-

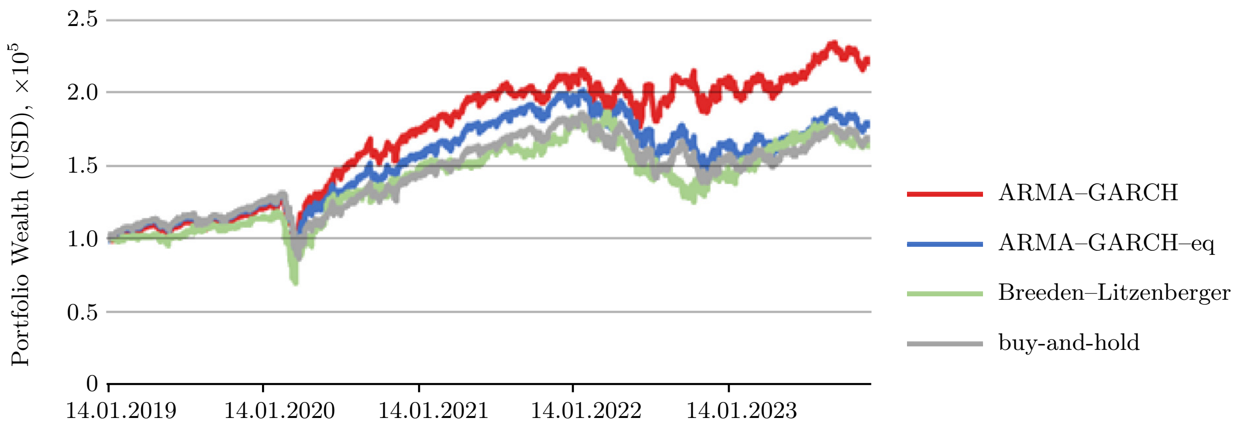

Figure 4. The dynamics of portfolio wealth

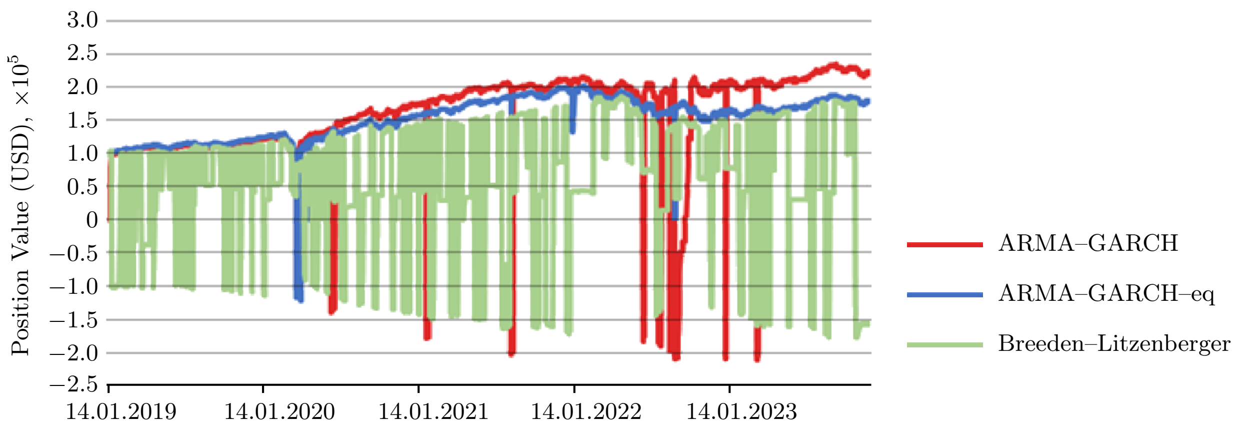

Figure 5. ETF open position values

2026, **27** (1), 90–100. doi 10.26089/NumMet.v27r107

ber 2022. However, long positions prevailed in the portfolio, reflecting typical behavior of investors that believe in the upward trend of the S&P 500 index. Moreover, the strategy does not provide for frequent adjustments to positions, which indicates that the performance is not driven by overfitting to short-term fluctuations or noise.

In contrast, the ARMA–GARCH–eq portfolio, which employed the same scenario lattice but assumed equal transition probabilities, deliver lower returns. The Breeden–Litzenberger scenario lattices with varying numbers of days between stages are noticeably shorter and do not incorporate conditional transition probabilities, causing the portfolio to frequently adjust its positions. This demonstrates why transition probabilities derived from the calibrated model dynamics are crucial. Unlike the approach that employs such calibrated dynamics, both Breeden–Litzenberger (which relies on static expiration-date densities) and passive buy-and-hold portfolios produce suboptimal results.

In general, the obtained results show that incorporation of the option-calibrated ARMA–GARCH models into SDDP-based multi-stage optimization can significantly improve portfolio performance compared to methods that ignore the conditional and path-dependent nature of price dynamics. The results also raise questions about the efficiency of the underlying asset's market in the presence of a liquid options market.

References

1. F. Black and M. Scholes, "The Pricing of Options and Corporate Liabilities," Journal of Political Economy **81** (3), 637–654 (1973).
2. J. C. Hull, *Options, Futures and Other Derivatives* (Pearson, 2021).
3. P. A. Arbuzov and D. Yu. Golembiovsky, "Calibration of the S_U-Johnson Distribution of the Future Price of the Underlying Asset Based on Option Prices," Problemy Analiza Riska **21** (2), 78–93 (2024). https://www.risk-journal.com/jour/article/view/808. Cited March 4, 2026.
4. D. T. Breeden and R. H. Litzenberger, "Prices of State-Contingent Claims Implicit in Option Prices," The Journal of Business **51** (4), 621–651 (1978). doi 10.1086/296025.
5. D. Shimko, "Bounds of Probability," Risk **6** (4), 33–37 (1993). https://www.researchgate.net/publication/306151578_Bounds_of_probability. Cited March 4, 2026.
6. Y. Ait-Sahalia and A. W. Lo, "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices," The Journal of Finance **53** (2), 499–547 (1998). doi 10.1111/0022-1082.215228.
7. R. R. Bliss and N. Panigirtzoglou, "Option-Implied Risk Aversion Estimates," The Journal of Finance **59** (1), 407–446 (2004). doi 10.1111/j.1540-6261.2004.00637.x.
8. P. A. Arbuzov and D. Yu. Golembiovsky, "Calibration of the ARIMA–GARCH Model of the Underlying Asset Based on Market Option Prices," Economika i Matematicheskie Metody **61** (3), 104–115 (2025). https://journals.eco-vector.com/0424-7388/article/view/691409. Cited March 4, 2026.
9. 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).
10. D. Guterding and W. Boenkost, "The Heston stochastic volatility model with piecewise constant parameters — efficient calibration and pricing of window barrier options," J. Comput. Appl. Math. **343**, 353–362 (2018). doi 10.1016/j.cam.2018.04.054.
11. G. Liu and W. Xu, "Application of Heston's Model to the Chinese Stock Market," Emerging Markets Finance and Trade **53** (8), 1749–1763 (2017). doi 10.1080/1540496X.2016.1219849.
12. M. V. F. Pereira and L. M. V. G. Pinto, "Multi-stage stochastic optimization applied to energy planning," Mathematical Programming **52** (1–3), 359–375 (1991). https://www.sci-hub.ru/10.1007/bf01582895?ysclid=mmbna1fjoq894065792. Cited March 4, 2026.
13. A. Shapiro, "Analysis of Stochastic Dual Dynamic Programming Method," European Journal of Operational Research **209** (1), 63–72 (2011). https://doi.org/10.1016/j.ejor.2010.08.007. Cited March 4, 2026.
14. S. Ghadimi, G. Lan, and H. Zhang, "Mini-Batch Stochastic Approximation Methods for Nonconvex Stochastic Composite Optimization," Mathematical Programming **155** (1), 267–305 (2016). doi 10.1007/s10107-014-0846-1.
15. A. R. Danilishin, "Approximation of the Girsanov Measure with Logarithmic Returns in the Case of Heavy-Tailed Distributions," Trudy ISA RAN. **73** (3), 21–30 (2023). http://www.isa.ru/proceedings/images/documents/2023-73-3/21-30.pdf. Cited March 5, 2026.
16. Y. Nesterov and V. Spokoiny, "Random Gradient-Free Minimization of Convex Functions," Foundations of Computational Mathematics **17** (2), 527–566 (2017). doi 10.1007/s10208-015-9296-2.

17. P. Arbuzov, "Convergence of the ARMA–GARCH Implied Calibration Algorithm," International Journal of Open Information Technologies **13** (9), 60–65 (2025). http://injoit.org/index.php/j1/article/view/2180. Cited March 5, 2026.

18. X. Liu, M. B. Shackleton, S. J. Taylor, and X. Xu, "Closed-Form Transformations from Risk-Neutral to Real-World Distributions," Journal of Banking & Finance **31** (5), 1501–1520 (2007). doi 10.1016/j.jbankfin.2006.09.005.

19. C. Crnkovic and J. Drachman, "Quality Control," Risk **9** (9), 138–143 (1996).

20. W. A. Broock, J. A. Scheinkman, W. D. Dechert, and B. LeBaron, "A Test for Independence Based on the Correlation Dimension," Econometric Reviews **15** (3), 197–235 (1996). doi 10.1080/07474939608800353.

21. A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory* (SIAM, Philadelphia, 2009).

22. A. Downward, O. Dowson, and R. Baucke, "Stochastic dual dynamic programming with stagewise dependent objective uncertainty," Optimization Online (2018). https://optimization-online.org/2018/02/6454/. Cited March 5, 2026.

23. D. Golembiovsky, A. Pavlov, and D. Smetanin, "Experimental Study of Methods of Scenario Lattice Construction for Stochastic Dual Dynamic Programming," Open Journal of Optimization **10** (2), 47–60 (2021). doi 10.4236/ojop.2021.102004.

24. R. T. Rockafellar and S. Uryasev, "Optimization of Conditional Value-at-Risk," Journal of Risk **2** (3), 21–41 (2000). https://ideas.repec.org/a/rsk/journ4/2161159.html. Cited March 5, 2026.

25. O. Dowson and L. Kapelevich, "SDDP.jl: A Julia Package for Stochastic Dual Dynamic Programming," INFORMS Journal on Computing **33** (1), 27–33 (2021). doi 10.1287/ijoc.2020.0987.

Received *January 26, 2026* | *Accepted* *March 3, 2026* | *Published* *March 17, 2026*

Information about the authors

Petr A. Arbuzov — Ph.D. student, Lomonosov Moscow State University, Faculty of Computational Mathematics and Cybernetics, Leninskie Gory, 1, building 52, 119991, Moscow, Russian Federation.

Dmitry Yu. Golembiovsky — Dr. Sci., Professor; Lomonosov Moscow State University, Faculty of Computational Mathematics and Cybernetics, Leninskie Gory, 1, building 52, 119991, Moscow, Russian Federation.