# Portfolio selection with adaptive conformal prediction

Yusen Jia[1] and Bingyan Han[1(✉)][0000−0003−0680−7967]

Thrust of Financial Technology, The Hong Kong University of Science and Technology (Guangzhou)
yjia256@connect.hkust-gz.edu.cn
bingyanhan@hkust-gz.edu.cn

**Abstract.** We propose a model-free portfolio selection framework in which investment risk is estimated using conformal prediction. This approach accommodates distribution shifts, integrates with any regression method, and provides coverage guarantees. The value at risk (VaR) is derived from the lower bound of the prediction sets constructed via conformal prediction. We employ a projected gradient descent algorithm to optimize the portfolio weights under investor-specified constraints. Numerical results show that conformalized strategies with short-selling constraints consistently outperform both the equal-weighted portfolio and non-conformal counterparts across multiple performance metrics.

**Keywords:** Conformal Prediction · Portfolio Selection · Value at Risk

## 1 Introduction

Risk is a central concept in finance, playing a key role in investment decisions, risk management, and regulation. Markowitz [19] pioneered the use of mean and variance as proxies for the reward and risk of a portfolio, establishing the foundation of modern portfolio theory and inspiring a vast body of research in portfolio selection. A common criticism of this framework is that variance treats returns above the expected level as part of the risk. To address this limitation, alternative risk measures have been proposed in the literature, including semi-variance, value at risk (VaR), conditional VaR (also known as expected shortfall), mean absolute deviation, and maximal loss [3, 22, 12].

Estimating risk measures typically requires knowledge of the distribution of stock returns. In practice, agents often assume a parametric model and estimate its parameters using historical data and statistical techniques. However, accurate estimation generally demands a large number of observations, which are often unavailable in financial markets. As a result, even correctly specified models may yield unreliable estimates. In many situations, even identifying the correct model itself is challenging [18, 11]. Compounding the issue, optimal portfolio weights are usually sensitive to return parameters. An “optimal” portfolio may perform poorly in out-of-sample tests, rendering it practically ineffective. This raises a natural question: can we design distribution-free estimation methods that require minimal assumptions on the underlying distribution while still providing guarantees on estimation errors?

Inspired by recent developments in machine learning and uncertainty quantification, we adopt conformal inference (prediction) for estimating investment risk measured by VaR. The conformal inference framework was originally introduced by [24, 25] to construct prediction intervals in a sequential manner. [15] extended this to a general framework for distribution-free prediction intervals with finite-sample coverage guarantees, under the key assumption of exchangeability on data distributions. However, this assumption is often violated in real-world applications, including finance. Recent work has generalized conformal inference to handle a broader range of distribution shifts and dependency structures. Examples include methods tailored for cross-sectional time series [16], label shift [21], and covariate shift [23], among others. A comprehensive review is provided in [1].

Financial data often exhibit distinct statistical properties in bear versus bull markets, motivating the use of conformal inference techniques that adapt to distributional shifts. [8] introduces adaptive conformal inference by employing gradient descent to tune the width of prediction sets. While this method shows strong theoretical and empirical performance, it is sensitive to the step-size parameter, which is typically unknown in practice. To address this limitation, [9] proposes a re-weighting scheme based on the historical performance of a set of candidate step sizes, similar to the approach in [10]. [6] demonstrates that the method in [9] outperforms those in [26, 2]. Therefore, we adopt the algorithm in [9], with suitable modifications, to estimate portfolio investment risks.

We focus on portfolio risk measured by VaR, though the proposed method can be extended to other quantile-based risk measures. Given an investment strategy $\mathbf{w}$, VaR can be directly estimated using the endpoint of the prediction set generated by an algorithm adapted from [9], with the score function defined in (7). The covariate variables are derived from "virtual portfolios" [7], which use portfolio returns rather than individual stock returns. We then optimize the investment weight $\mathbf{w}$ using a projected gradient descent method, subject to the investor's specified constraints.

In the numerical study, we examine the investment problem in the U.S. stock market. In the benchmark setting, AR and GARCH models are used for VaR estimation without conformal prediction, and short selling is permitted. Both AR and GARCH strategies yield lower Sharpe ratios than the naive equal-weighted (EW) strategy, which requires no estimation. After incorporating conformal prediction, Table 2 shows that the GARCH model outperforms the EW strategy in terms of the Sharpe ratio. However, the AR model and most machine learning algorithms still underperform the EW portfolio. To address this, we impose a short-selling constraint on the conformalized strategies. Table 3 demonstrates that, under this constraint, most methods outperform the EW strategy. This result aligns with the findings in [13], which shows that imposing appropriate constraints can enhance the out-of-sample performance of investment strategies.

## 2 Conformal prediction

Let $\{(X_i, R_i)\}_{i=1,\dots,n}$ denote a set of $n$ observed training samples, where $X_i \in \mathcal{X}$ represents the covariate (or feature) and $R_i \in \mathcal{R}$ is the associated response variable. Consider a new data point $(X_{n+1}, R_{n+1})$, for which only the covariate $X_{n+1}$ is observed. The goal of conformal prediction (inference) is to construct a set $\hat{C}_{n+1}(X_{n+1})$ such that

$$\mathbb{P}(R_{n+1} \in \hat{C}_{n+1}(X_{n+1})) \geq 1 - \alpha, \tag{1}$$

where $\alpha$ is a nominal error level and the probability is taken over the joint distribution of all $n+1$ data points. The set $\hat{C}_{n+1}(X_{n+1})$ is referred to as a prediction set or conformal set.

A trivial choice of prediction set is

$$\hat{C}_{n+1}(X_{n+1}) := \begin{cases} \mathcal{R}, & \text{with probability } 1-\alpha, \\ \emptyset, & \text{with probability } \alpha. \end{cases}$$

However, this set is too large to be practically useful. Therefore, conformal prediction aims for a stronger goal: when $R_{n+1}$ is easier to predict from $X_{n+1}$, the prediction set $\hat{C}_{n+1}(X_{n+1})$ should be smaller. Remarkably, this is achievable in broad settings. Conformal prediction begins by introducing a score function, often based on a regression model trained using a chosen algorithm. Prediction sets become narrower when the target is easier to predict or when the model is more effective. Next, ranks of these scores are used to construct (adjusted) quantiles. Classical conformal prediction assumes data exchangeability, so that the score for $(X_{n+1}, R_{n+1})$ is uniformly distributed among prior scores. The method is distribution-free, providing exact finite-sample guarantees under any distribution, as long as exchangeability holds. A comprehensive review of recent developments can be found in [1].

First, we assume that the data points $(X_1, R_1), \dots, (X_{n+1}, R_{n+1})$ are exchangeable, a weaker condition than the standard independent and identically distributed (i.i.d.) assumption.

The split conformal prediction method proposed in [15] proceeds as follows. First, the training data $\{(X_i, R_i)\}_{1 \leq i \leq n}$ are randomly divided into two disjoint subsets: the training set $\mathcal{D}_{train} := \{(X_i, R_i) : i \in \mathcal{I}_1\}$ and the calibration set $\mathcal{D}_{cal} := \{(X_i, R_i) : i \in \mathcal{I}_2\}$, where $\mathcal{I}_1 \cap \mathcal{I}_2 = \emptyset$ and $\mathcal{I}_1 \cup \mathcal{I}_2$ includes all indices. Define $n_{train} = |\mathcal{D}_{train}|$ and $n_{cal} = |\mathcal{D}_{cal}|$. A typical choice is an equal split: $n_{train} = n_{cal} = n/2$, assuming without loss of generality that $n$ is even.

With the training set $\mathcal{D}_{train}$, we fit a regression model $\hat{f} : \mathcal{X} \to \mathcal{R}$ with a learning algorithm $\mathcal{A}$:

$$\hat{f} \leftarrow \mathcal{A}(\mathcal{D}_{train}).$$

With the fitted model $\hat{f}$, define a score function $S(\cdot, \cdot; \hat{f}) : \mathcal{X} \times \mathcal{R} \to \mathbb{R}$ and compute the scores for the calibration data:

$$S_i := S(X_i, R_i; \hat{f}),\ (X_i, R_i) \in \mathcal{D}_{cal}.$$

Common choices for the score function include

$$S(X,R;\hat{f}) := |R - \hat{f}(X)| \quad \text{and} \quad S(X,R;\hat{f}) := \hat{f}(X) - R.$$

The prediction set for a new input $X_{n+1}$ is given by

$$\begin{aligned}\hat{C}_{n+1}(X_{n+1}) := \big\{ y : S(X_{n+1}, y; \hat{f}) \le \text{the } \lceil (1-\alpha)(n_{cal}+1) \rceil\text{-th} \\ \text{smallest value of } \{S_i\}_{i\in\mathcal{D}_{cal}} \big\}.\end{aligned} \tag{2}$$

Since the data are exchangeable and we can break ties uniformly at random, the rank of $S(X_{n+1}, R_{n+1}; \hat{f})$ among the points $\{S_i\}_{i\in\mathcal{D}_{cal}}$ is uniformly distributed. Theorem 2.2 in [15] guarantees the desired coverage property (1).

The performance of prediction set $\hat{C}_{n+1}(X_{n+1})$ is typically evaluated based on two criteria: coverage guarantee and set width. The coverage guarantee depends on the underlying data distribution, while the width is influenced by the choice of score function, which reflects the predictive performance of the model $\hat{f}$ on the calibration set $\mathcal{D}_{cal}$. A more accurate model $\hat{f}$ generally yields narrower prediction sets.

In time series or other non-exchangeable settings, the coverage guarantee in (1) may no longer hold. This limitation has motivated efforts to extend conformal prediction methods beyond the exchangeable data framework. [9] introduced the dynamically-tuned adaptive conformal inference (DtACI) algorithm. DtACI constructs prediction sets in an online manner by dynamically updating both the error level and the training dataset. Without assuming any specific data distribution, the algorithm provides valid conformal coverage over long time horizons.

**Theorem 1 (Theorem 6 in [9]).** *Given positive weight parameters $\eta_t$ and $\sigma_t$, let $\gamma_{min} := \min_i \gamma_i$ and $\gamma_{max} := \max_i \gamma_i$. Define $err_t$ as the indicator of the error at time $t$, i.e. $R_t$ is not in the conformal set. Then the DtACI algorithm satisfies*

$$\begin{aligned}\left| \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}[err_t] - \alpha \right| \le & \frac{1+2\gamma_{max}}{T\gamma_{min}} + \frac{(1+2\gamma_{max})^2}{\gamma_{min}} \frac{1}{T}\sum_{t=1}^{T} \eta_t e^{\eta_t(1+2\gamma_{max})} \\ & + 2\frac{1+\gamma_{max}}{\gamma_{min}} \frac{1}{T}\sum_{t=1}^{T} \sigma_t.\end{aligned}$$

*If $\lim_{t\to\infty} \eta_t = \lim_{t\to\infty} \sigma_t = 0$, then $\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} err_t \overset{a.s.}{=} \alpha$.*

## 3 Portfolio selection

### 3.1 Formulation

Consider a financial market with $n$ risky assets (stocks) available for investment. Let $Z_{i,t}$ denote the random return of stock $i$ at period $t$. Assume that returns for the past $T$ periods have been observed and are given by $Z_{i,t} = z_{i,t}$.

We study a portfolio selection problem for period $T+1$. Let $w_i$ represent the proportion of wealth allocated to asset $i$, where $\sum_{i=1}^{n} w_i = 1$. The investment weight vector is denoted by $\mathbf{w} := (w_1, \ldots, w_n) \in \mathbb{R}^n$. Short selling is allowed in our formulation, while a leverage constraint is imposed to limit excessive risk-taking in asset investments. Specifically, the total absolute weight of the stocks is constrained to not exceed a given positive threshold $L$. Consequently, the space of admissible investment weights is given by

$$\mathcal{W} := \left\{ \mathbf{w} \in \mathbb{R}^n : \mathbf{1}^T \mathbf{w} = 1, \quad \|\mathbf{w}\|_1 \leq L \right\}. \tag{3}$$

Let $\mathbf{Z}_{T+1} := (Z_{1,T+1}, \ldots, Z_{n,T+1})$ denote the random return vector at period $T+1$. Then the portfolio return with the weight $\mathbf{w}$ is given by

$$R(\mathbf{w}; \mathbf{Z}_{T+1}) = \sum_{i=1}^{n} w_i Z_{i,T+1}. \tag{4}$$

The classical mean-variance theory uses variance as a risk measure. However, it penalizes both downside and upside deviations from the mean, even though returns above the mean are typically desirable. To overcome this limitation, VaR can be employed as an alternative risk measure:

$$\mathrm{VaR}_\alpha(R(\mathbf{w}; \mathbf{Z}_{T+1})) := \inf \{x \in \mathbb{R} : \mathbb{P}(-R(\mathbf{w}; \mathbf{Z}_{T+1}) \leq x) \geq 1 - \alpha\}. \tag{5}$$

where $\alpha$ denotes the specified level and the probability is taken over the random asset return vector $\mathbf{Z}_{T+1}$. The portfolio selection problem then minimizes VaR over all admissible investment weights:

$$\min_{\mathbf{w} \in \mathcal{W}} \mathrm{VaR}_\alpha(R(\mathbf{w}; \mathbf{Z}_{T+1})). \tag{6}$$

Although various risk measures have been studied in the literature, this paper focuses exclusively on VaR to maintain clarity of exposition, with extensions left for future work.

The estimation of VaR has been extensively explored [14, 17]. Given the well-known challenges in modeling stock returns, we adopt a distribution-free approach using conformal prediction to estimate VaR.

### 3.2 Portfolio selection with conformal prediction

We begin by estimating $\mathrm{VaR}_\alpha(R(\mathbf{w}; \mathbf{Z}_{T+1}))$ for a fixed investment weight vector $\mathbf{w}$. Recall that the previous $T$ periods of stock returns, denoted by $z_{i,t}$, are assumed to be observed.

As the number of stocks can be large, [7] suggests estimating the portfolio returns directly rather than modeling individual stock returns. High-dimensional time series modeling is often challenging, whereas portfolio returns can be forecasted using univariate models, making the approach more practical and easier to implement.

Given the weight vector $\mathbf{w}$, the portfolio return for each of the previous $T$ periods is computed as

$$R_t = \sum_{i=1}^{n} w_i z_{i,t}.$$

Following [7], we refer to this as a "virtual portfolio", since $\mathbf{w}$ may not correspond to the actual investment weights used during those periods. For notational simplicity, we omit the dependence of $R_t$ on $\mathbf{w}$.

In this paper, we use past portfolio returns as covariates to predict the return in the next period. We define $X_t := (R_{t-l}, \ldots, R_{t-1})$, where $l$ denotes the window size. Additional features, such as macroeconomic indicators, can also be included in $X_t$ if available.

After fitting the regression model $\hat{f}_t$, we define the score function as

$$S(X_t, R_t; \hat{f}_t) = \hat{f}_t(X_t) - R_t, \tag{7}$$

which is one-sided and referred to as the conformalized mean regression score metric [4]. The corresponding prediction set $\hat{C}_t(X_t; \alpha_t)$ is constructed as

$$\begin{aligned} \hat{C}_t(X_t; \alpha_t) = \big\{ r : r \geq \hat{f}_t(X_t) - \text{the } \lceil (1-\alpha_t)(n_{cal}+1) \rceil\text{-th} \\ \text{smallest value of } \{S_u\}_{t-n_{cal} \leq u \leq t-1} \big\}, \end{aligned} \tag{8}$$

where $\alpha_t$ is adaptively updated via DtACI.

Under suitable assumptions, the long-run average error rate of the event $R_t \notin \hat{C}_t(X_t; \alpha_t)$ converges to the specified level $\alpha$. This result motivates the estimation of $\mathrm{VaR}_\alpha(R(\mathbf{w}; \mathbf{Z}_{T+1}))$ using the negative lower bound of the prediction set $\hat{C}_{T+1}(X_{T+1}; \alpha_{T+1})$, defined as

$$\begin{aligned} &\mathrm{VaR}^{\mathrm{DtACI}}_{\alpha_{T+1}}(R(\mathbf{w}; \mathbf{Z}_{T+1})) \\ &\quad := \text{ the } \lceil (1-\alpha_{T+1})(n_{cal}+1) \rceil\text{-th smallest} \\ &\qquad \text{value of } \{S_u\}_{T+1-n_{cal} \leq u \leq T} - \hat{f}_{T+1}(X_{T+1}). \end{aligned}$$

Algorithm 2 then solves the optimization problem (6) using projected gradient descent, where $\mathrm{VaR}_\alpha(R(\mathbf{w}; \mathbf{Z}_{T+1}))$ is computed via Algorithm 1. While other optimization methods could be used, we do not pursue a comparative analysis here.

Based on DtACI, Algorithm 1 provides a distribution-free VaR forecast, denoted as $\mathrm{VaR}_\alpha(R(\mathbf{w}; \mathbf{Z}_{T+1}))$, for the portfolio return distribution at time $T+1$. The model $\hat{f}_u$ is recalibrated at each time step to incorporate the most recent data. To reduce computational cost, however, one may opt to fix a common regression model that does not depend on future data.

## 4 Empirical analysis

### 4.1 Conformalized VaR forecast with synthetic data

First, we demonstrate the effectiveness of our methods with simulated asset returns from a known RS-GARCH(1,1) process, without considering portfolio selection. The error level is set to $\alpha = 0.05$.

**Algorithm 1** Conformalized VaR forecast with a given weight $\mathbf{w}$

**Input:** Data $\{(X_t, R_t)\}_{t=1,\dots,T}$, $X_{T+1}$, training window size $n_{train}$, calibration window size $n_{cal}$, with $T - 2n_{cal} - n_{train} + 1 \geq 1$, algorithm $\mathcal{A}$
**for** $u = T - 2n_{cal} + 1, \dots, T - 1$ **do**
  Fit model $\hat{f}_u$ on $\{(X_h, R_h)\}_{u-n_{train} \leq h \leq u-1}$ using the algorithm $\mathcal{A}$
  Compute conformity score $S_u = S(X_u, R_u; \hat{f}_u)$
**end for**
**for** $s = T - n_{cal} + 1, \dots, T$ **do**
  Use scores $\{S_u\}_{s-n_{cal} \leq u \leq s-1}$ to compute

$$\beta_s := \sup\{\beta : R_s \in \hat{C}_s(X_s; \beta)\},$$

  where $\hat{C}_s(X_s; \beta)$ is defined in (8)
**end for**
Run DtACI over $T - n_{cal} + 1 \leq s \leq T$ with $\{\beta_s\}_{T-n_{cal}+1 \leq s \leq T}$
**Output:** The estimate $\mathrm{VaR}^{\mathrm{DtACI}}_{\alpha_{T+1}}(R(\mathbf{w}; \mathbf{Z}_{T+1}))$

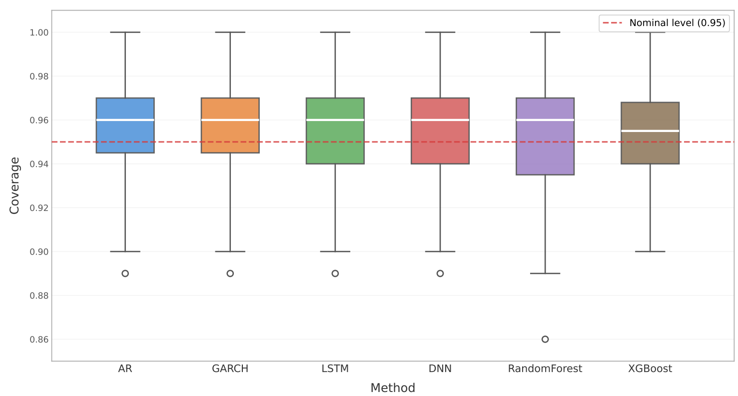


Fig. 1: Distribution of coverage rates for conformalized methods

The conformalized VaR estimation method can be integrated with any regression model. In Figure 1, we estimate $\hat{f}_t$ using AR, GARCH, Random Forest (RF), XGBoost, Multilayer Perceptron (denoted DNN in Figure 1), and Long Short-Term Memory (LSTM) networks. The calibration set size is fixed at 50. Figure 1 shows the coverage rate distributions under conformal prediction. These results underscore the robustness of conformal prediction under model misspecification.

### 4.2 Portfolio performance

This section compares our conformalized portfolios against several benchmarks. A natural baseline is the equal-weighted (EW) strategy, which allocates $1/n$ of total wealth to each of the $n$ available assets at each rebalancing date. [5]

**Algorithm 2** Portfolio selection with conformal prediction

**Input**: Stock returns $\{z_{i,t}\}_{1\leq i\leq n, 1\leq t\leq T}$, initial allocation $\mathbf{w} = (\frac{1}{n}, \ldots, \frac{1}{n})$, training window size $n_{train}$, calibration window size $n_{cal}$, learning rate $\zeta$
Initialize $\mathbf{w}^{(0)} \leftarrow \mathbf{w}$ and set iteration counter $j \leftarrow 0$
**repeat**
    Compute virtual portfolio returns $R_t = \sum_{i=1}^{n} w_i^{(j)} z_{i,t}$
    Define features $X_t := (R_{t-l}, \ldots, R_{t-1})$ and form dataset $\{(X_t, R_t)\}_{1\leq t\leq T}$
    Run Algorithm 1 to compute $\mathrm{VaR}^{\mathrm{DtACI}}_{\alpha_{T+1}}(R(\mathbf{w}^{(j)}; \mathbf{Z}_{T+1}))$ with the previous dataset
    Update portfolio weights via projected gradient descent:

$$\mathbf{w}^{(j+1)} \leftarrow \mathrm{proj}_{\mathcal{W}} \left( \mathbf{w}^{(j)} - \zeta \nabla_{\mathbf{w}} \mathrm{VaR}^{\mathrm{DtACI}}_{\alpha_{T+1}}(R(\mathbf{w}^{(j)}; \mathbf{Z}_{T+1})) \right)$$

    Set $j \leftarrow j + 1$
**until** convergence criterion is met
**Output:** Portfolio weights $\mathbf{w}^{(j)}$

finds that many sophisticated strategies fail to outperform the EW portfolio in terms of out-of-sample Sharpe ratios. Moreover, [20] shows that in a single-period setting, the EW strategy is optimal when model uncertainty is sufficiently high. Therefore, we include the EW strategy as a key benchmark. Additionally, several non-conformal strategies from Table 1 are also used for comparison.

All experiments involve investing in several component stocks from the S&P 500 index. Portfolio performance is evaluated from the start of 2019 to the end of 2023, a period marked by elevated market volatility due to the global pandemic and related economic disruptions. We apply the one-period model in (6) across multiple periods, with portfolio rebalancing at the beginning of each month. The VaR measure is estimated using monthly portfolio returns, with the error level set to $\alpha = 0.05$.

**Benchmark cases: Strategies without conformal prediction and short-selling constraint.** We first consider portfolio strategies based on VaR estimates from AR and GARCH models. In this setting, conformal prediction is not incorporated, and short selling is permitted. An additional leverage constraint is imposed by setting $L = 1.5$ in (3).

Table 1 shows that both the AR and GARCH strategies yield lower Sharpe ratios than the EW portfolio, while achieving higher Sortino ratios. The AR portfolio demonstrates a substantially lower maximum drawdown, which is approximately 33.5% less than that of the EW strategy. In contrast, the GARCH strategy exhibits a higher turnover, suggesting more volatile allocation adjustments.

Figure 2 illustrates that from early 2019 to the end of 2023, the EW strategy nearly doubles its initial wealth, consistent with its higher Sharpe ratio. The AR and GARCH portfolios display more stable trajectories over the same period.

Table 1: Performance without conformal prediction. Short selling is allowed

| Metric | AR | GARCH | EW |
|---|---|---|---|
| Monthly Return Mean | 0.0119 | 0.0115 | 0.0151 |
| Monthly Return Std | 0.0479 | 0.0500 | 0.0598 |
| Maximum Drawdown | 0.2225 | 0.2772 | 0.3343 |
| Turnover | 0.1898 | 0.5777 | 0.0449 |
| Annualized Sharpe | 0.8584 | 0.7936 | 0.8736 |
| Annualized Sortino | 1.1379 | 1.0900 | 1.0396 |

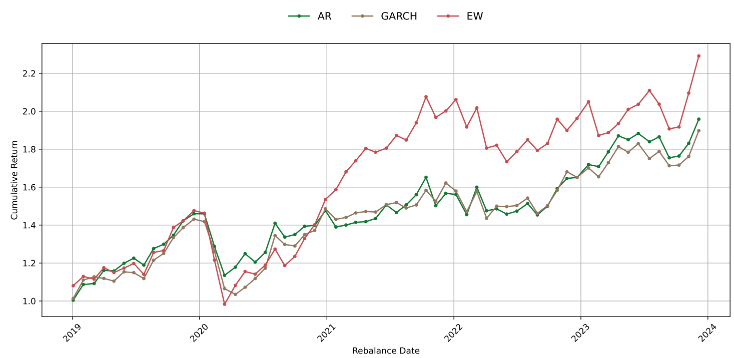


Fig. 2: Cumulative returns of strategies without conformal prediction. Short selling is allowed

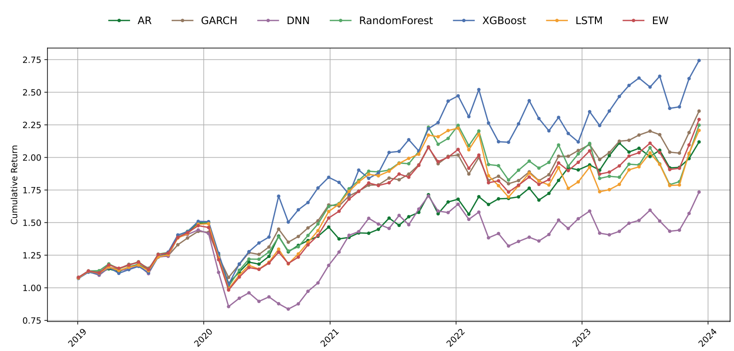


Fig. 3: Cumulative returns of conformalized strategies, when short selling is allowed. The EW strategy is the same as in Figure 2

Table 2: Performance of conformalized strategies when short selling is allowed

| Metric | AR | GARCH | LTSM | DNN | RF | XGBoost |
|---|---|---|---|---|---|---|
| Monthly Return Mean | 0.0136 | 0.0149 | 0.0149 | 0.0115 | 0.0152 | 0.0187 |
| Monthly Return Std | 0.0565 | 0.0499 | 0.0663 | 0.0720 | 0.0662 | 0.0718 |
| Maximum Drawdown | 0.3102 | 0.2469 | 0.3353 | 0.4203 | 0.3282 | 0.3199 |
| Turnover | 0.1120 | 0.0890 | 0.1107 | 0.0944 | 0.0925 | 0.1337 |
| Annualized Sharpe | 0.8345 | 1.0375 | 0.7780 | 0.5515 | 0.7936 | 0.9021 |
| Annualized Sortino | 1.0064 | 1.2998 | 0.9496 | 0.6983 | 0.9631 | 1.3892 |

**Performance of conformalized strategies when short selling is allowed.** After applying conformal prediction, the AR and GARCH strategies respond differently. As shown in Table 2, the GARCH strategy achieves higher Sharpe and Sortino ratios, while maintaining relatively low maximum drawdown and return volatility. Portfolio stability is notably improved, as indicated by a substantial decrease in turnover. In contrast, the AR strategy performs worse than in the unconformalized case presented in Table 1.

Following the simulation study, conformal prediction is combined with several machine learning algorithms to generate VaR estimates for portfolio selection. As shown in Figure 3, the XGBoost portfolio delivers the highest cumulative return, although its return volatility exceeds that of the EW benchmark. Other machine learning-based portfolios do not outperform the EW benchmark in terms of Sharpe and Sortino ratios.

In the following subsection, we show that the underperformance of the AR and machine learning strategies can be mitigated by imposing a short-selling constraint. All models use default hyperparameters without fine-tuning.

Table 3: Performance of conformalized strategies with the short-selling constraint

| Metric | AR | GARCH | LSTM | DNN | RF | XGBoost |
|---|---|---|---|---|---|---|
| Monthly Return Mean | 0.0144 | 0.0151 | 0.0165 | 0.0149 | 0.0161 | 0.0171 |
| Monthly Return Std | 0.0569 | 0.0507 | 0.0613 | 0.0610 | 0.0532 | 0.0629 |
| Maximum Drawdown | 0.3102 | 0.2625 | 0.3351 | 0.3362 | 0.2840 | 0.3614 |
| Turnover | 0.0941 | 0.0776 | 0.1038 | 0.0971 | 0.0983 | 0.0968 |
| Annualized Sharpe | 0.8747 | 1.0285 | 0.9339 | 0.8452 | 1.0507 | 0.9391 |
| Annualized Sortino | 1.0356 | 1.2516 | 1.1356 | 1.0661 | 1.2824 | 1.0679 |

**Performance of conformalized strategies with the short-selling constraint.** Table 3 shows that, under the short-selling constraint, the GARCH

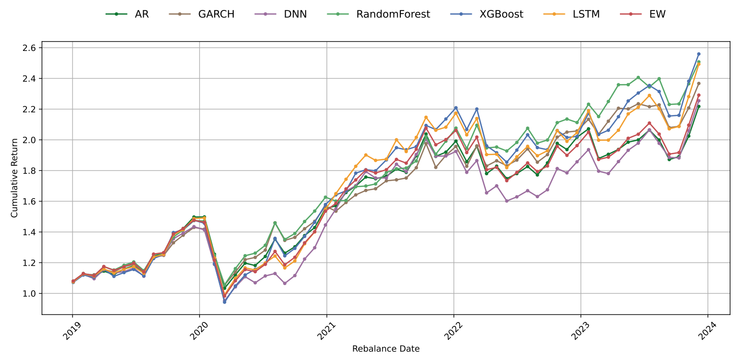


Fig. 4: Cumulative returns of conformalized strategies with the short-selling constraint. The EW strategy is the same as in Figure 2

portfolio continues to achieve higher Sharpe and Sortino ratios than its counterpart without conformal prediction, as seen in Table 1. These performance metrics are only slightly lower than those in Table 2, where short selling is permitted. The AR portfolio yields results comparable to the EW benchmark.

For portfolios using machine learning methods, Table 3 indicates improvements in both Sharpe and Sortino ratios, outperforming the EW benchmarks even with the short-selling constraint. Most portfolios also exhibit reduced monthly return volatility compared to those in Table 2, resulting in more stable performance as given by Figure 4.

This result may seem counterintuitive from an optimization perspective, as allowing short selling expands the feasible set and should, in principle, enhance performance. In fact, our finding is consistent with [13]. Empirical studies indicate that removing investment constraints can degrade out-of-sample performance. The short-selling constraint serves as an implicit form of regularization. By limiting overly aggressive positions, it reduces sensitivity to estimation errors and thereby improves out-of-sample outcomes.

## References

1. Angelopoulos, A.N., Bates, S., et al.: Conformal prediction: A gentle introduction. Foundations and Trends® in Machine Learning **16**(4), 494–591 (2023)
2. Bhatnagar, A., Wang, H., Xiong, C., Bai, Y.: Improved online conformal prediction via strongly adaptive online learning. In: International Conference on Machine Learning. pp. 2337–2363. PMLR (2023)
3. Campbell, R., Huisman, R., Koedijk, K.: Optimal portfolio selection in a value-at-risk framework. Journal of Banking & Finance **25**(9), 1789–1804 (2001)
4. Candès, E., Lei, L., Ren, Z.: Conformalized survival analysis. Journal of the Royal Statistical Society Series B: Statistical Methodology **85**(1), 24–45 (2023)
5. DeMiguel, V., Garlappi, L., Uppal, R.: Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? The Review of Financial Studies **22**(5), 1915–1953 (2009)

6. Fantazzini, D.: Adaptive conformal inference for computing market risk measures: An analysis with four thousand crypto-assets. Journal of Risk and Financial Management **17**(6), 248 (2024)
7. Francq, C., Zakoian, J.M.: Virtual historical simulation for estimating the conditional var of large portfolios. Journal of Econometrics **217**(2), 356–380 (2020)
8. Gibbs, I., Candes, E.: Adaptive conformal inference under distribution shift. Advances in Neural Information Processing Systems **34**, 1660–1672 (2021)
9. Gibbs, I., Candès, E.J.: Conformal inference for online prediction with arbitrary distribution shifts. Journal of Machine Learning Research **25**(162), 1–36 (2024)
10. Gradu, P., Hazan, E., Minasyan, E.: Adaptive regret for control of time-varying dynamics. In: Learning for Dynamics and Control Conference. pp. 560–572. PMLR (2023)
11. Hansen, L.P.: Nobel lecture: Uncertainty outside and inside economic models. Journal of Political Economy **122**(5), 945–987 (2014)
12. He, X.D., Jin, H., Zhou, X.Y.: Dynamic portfolio choice when risk is measured by weighted var. Mathematics of Operations Research **40**(3), 773–796 (2015)
13. Jagannathan, R., Ma, T.: Risk reduction in large portfolios: Why imposing the wrong constraints helps. The Journal of Finance **58**(4), 1651–1683 (2003)
14. Kuester, K., Mittnik, S., Paolella, M.S.: Value-at-risk prediction: A comparison of alternative strategies. Journal of Financial Econometrics **4**(1), 53–89 (2006)
15. Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R.J., Wasserman, L.: Distribution-free predictive inference for regression. Journal of the American Statistical Association **113**(523), 1094–1111 (2018)
16. Lin, Z., Trivedi, S., Sun, J.: Conformal prediction with temporal quantile adjustments. Advances in Neural Information Processing Systems **35**, 31017–31030 (2022)
17. Longin, F.: Extreme Events in Finance: A Handbook of Extreme Value Theory and its Applications. John Wiley & Sons (2016)
18. Maenhout, P.J.: Robust portfolio rules and detection-error probabilities for a mean-reverting risk premium. Journal of Economic Theory **128**(1), 136–163 (2006)
19. Markowitz, H.: Portfolio selection. The Journal of Finance **7**(1), 77–91 (1952)
20. Pflug, G.C., Pichler, A., Wozabal, D.: The $1/N$ investment strategy is optimal under high model ambiguity. Journal of Banking & Finance **36**(2), 410–417 (2012)
21. Podkopaev, A., Ramdas, A.: Distribution-free uncertainty quantification for classification under label shift. In: Uncertainty in Artificial Intelligence. pp. 844–853. PMLR (2021)
22. Rockafellar, R.T., Uryasev, S.: Conditional value-at-risk for general loss distributions. Journal of Banking & Finance **26**(7), 1443–1471 (2002)
23. Tibshirani, R.J., Foygel Barber, R., Candes, E., Ramdas, A.: Conformal prediction under covariate shift. Advances in Neural Information Processing Systems **32** (2019)
24. Vovk, V., Gammerman, A., Shafer, G.: Algorithmic Learning in a Random World, vol. 29. Springer (2005)
25. Vovk, V., Nouretdinov, I., Gammerman, A.: On-line predictive linear regression. The Annals of Statistics pp. 1566–1590 (2009)
26. Zaffran, M., Féron, O., Goude, Y., Josse, J., Dieuleveut, A.: Adaptive conformal predictions for time series. In: International Conference on Machine Learning. pp. 25834–25866. PMLR (2022)