Performance-Based Regularization in Portfolio Optimization Problems

Shijie Shao, Di Wu, and Tongshu Zhang
College of Engineering
Cornell University
{ss4339, dw668, tz368}@cornell.edu

Abstract

This study reproduces and critically evaluates the framework of Ban et al. (2018) [1] for incorporating machine-learning-based regularization into portfolio optimization. We implement both mean-variance and CVaR performance-based regularization methods and assess their out-of-sample performance on industry indices. To calibrate the regularization parameter, we introduce a logarithmically spaced grid search combined with k-fold cross-validation. The models are tested on a dataset spanning 2010 to 2020. Our code can be accessed through GitHub[1].

1 Introduction

Portfolio optimization is a cornerstone of quantitative finance: given p assets with (unknown) return distribution, an investor seeks weights $w \in \mathbb{R}^p$ that balance expected return against risk. The classical Mean–Variance [4] and its Mean–CVaR [5] problems both replace the true moments by their empirical counterparts and solve the resulting Sample Average Approximation (SAA). While SAA is consistent as the number of observations $n \to \infty$, in realistic, limited-data settings often yields highly unstable optima and disappointing out-of-sample performance [3].

In machine learning, regularization is routinely used to curb overfitting and improve robustness on unseen data. Inspired by this, *Performance-Based Regularization* (PBR) is introduced for portfolio optimization: adding constraints that limit the sample variance of the estimated portfolio risk (and return) to guard against estimation noise.

Because the optimization problems of PBR models are nonconvex, Ban et al. (2018) propose several convex relaxation approaches to make the model solvable. These relaxations preserve asymptotic optimality and are closely connected to robust optimization frameworks.

In this project, we reproduce the key methods and results from their work, including implementing PBR models, calibrating regularization thresholds, backtesting the out-of-sample performance-based cross-validation scheme. We identify key failings of the original OOS-PBSD calibration and design a log-spaced grid search algorithm(OOS-PBGS). Through extensive backtesting on daily industry-index data, we validate the empirical improvements of PBR over standard sample-average approximation (SAA) methods.

The rest of the paper is organized as follows. Section 2 presents the PBR models and their convex approximations. Section 3 describes regularization calibration procedure. Section 4 introduces backesting setup and performance metrics. Section 5 provides backtesting results and discussion. Section 6 gives a general summary of this study.

[1]`https://github.com/Jim-Shao/MLPO`

2 Portfolio Optimization with PBR

2.1 Classical Formulations

We consider two canonical portfolio optimization objectives: mean-variance (MV) optimization in the spirit of Markowitz, and mean–Conditional Value-at-Risk (CVaR) optimization. The formulations below assume full knowledge of population-level moments or distribution.

The population mean-variance problem is given by:

$$\begin{aligned} w_{\mathrm{MV}} = \arg\min_{w\in\mathbb{R}^p} \quad & w^\top \Sigma w \\ \text{s.t.} \quad & w^\top 1_p = 1, \\ & w^\top \mu = R, \qquad \text{(MV-population)} \end{aligned}$$

where μ and Σ denote the true mean vector and covariance matrix of excess returns X, and R is the investor's target return (which may be optionally enforced).

Similarly, the population mean–CVaR portfolio problem reads:

$$\begin{aligned} w_{\mathrm{CV}} = \arg\min_{w\in\mathbb{R}^p} \quad & \mathrm{CVaR}\big(-w^\top X;\, \beta\big) \\ \text{s.t.} \quad & w^\top 1_p = 1, \\ & w^\top \mu = R, \qquad \text{(CV-population)} \end{aligned}$$

with

$$\mathrm{CVaR}\big(-w^\top X;\, \beta\big) := \min_{\alpha\in\mathbb{R}} \left\{ \alpha + \frac{1}{1-\beta}\, \mathbb{E}\big[-w^\top X - \alpha\big]^+ \right\},$$

as introduced by Rockafellar and Uryasev [5], where $\beta \in (0.5, 1)$ is a confidence level.

2.2 Limitations of Sample Average Approximation (SAA)

In practice, population parameters μ, Σ, or the full distribution of returns are unknown. A common workaround is the Sample Average Approximation (SAA), where these quantities are replaced by their empirical estimates based on historical data $\{X_i\}_{i=1}^n$. This yields the following empirical optimization problems:

$$\begin{aligned} \widehat{w}_{n,\mathrm{MV}} = \arg\min_{w\in\mathbb{R}^p} \; & w^\top \widehat{\Sigma}_n\, w \\ \text{s.t.} \quad & w^\top 1_p = 1, \\ & w^\top \widehat{\mu}_n = R, \end{aligned} \qquad \text{(MV-SAA)}$$

$$\begin{aligned} \widehat{w}_{n,\mathrm{CV}} = \arg\min_{w\in\mathbb{R}^p} \; & \widehat{\mathrm{CVaR}}_n(-w^\top X; \beta) \\ \text{s.t.} \quad & w^\top 1_p = 1, \\ & w^\top \widehat{\mu}_n = R, \end{aligned} \qquad \text{(CV-SAA)}$$

where

$$\widehat{\mathrm{CVaR}}_n(-w^\top X; \beta) := \min_{\alpha\in\mathbb{R}} \left\{ \alpha + \frac{1}{n(1-\beta)} \sum_{i=1}^n \big(-w^\top X_i - \alpha\big)^+ \right\}.$$

While both $\widehat{w}_{n,\mathrm{MV}} \to w_{\mathrm{MV}}$ and $\widehat{w}_{n,\mathrm{CV}} \to w_{\mathrm{CV}}$ hold in probability as $n \to \infty$, the SAA approach is known to be statistically fragile under small sample sizes. In particular, prior studies (e.g., [3]) show that limited observations can lead to highly unstable and overfitted solutions, especially in high dimensions.

This motivates the development of Performance-Based Regularization (PBR), which introduces a data-dependent penalty to control instability and improve out-of-sample generalization.

2.3 PBR for Mean–Variance Optimization

In the PBR approach, one starts from the SAA problem and adds an upper-bound constraint on the sample variance of the portfolio variance:

$$
\begin{aligned}
\widehat{w}_{n,\mathrm{MV}} = \arg\min_{w\in\mathbb{R}^p}\ & w^\top \widehat{\Sigma}_n\, w,\\
\text{s.t.}\quad & w^\top 1_p = 1,\\
& \big(w^\top \widehat{\mu}_n = R\big),\\
& \mathrm{SVar}\big(w^\top \widehat{\Sigma}_n\, w\big) \le U, \qquad \text{(MV-PBR)}
\end{aligned}
$$

where $\mathrm{SVar}(\cdot)$ is the sample-variance operator, and U_1 controls the degree of regularization.

Proposition 1 (Sample variance of the sample variance) *The sample variance of the sample variance of the portfolio,* $\mathrm{SVar}\big(w^\top \widehat{\Sigma}_n\, w\big)$*, is given by*

$$
\mathrm{SVar}\big(w^\top \widehat{\Sigma}_n w\big) = \sum_{i=1}^{p}\sum_{j=1}^{p}\sum_{k=1}^{p}\sum_{l=1}^{p} w_i\, w_j\, w_k\, w_l\ \widehat{Q}_{ijkl},
$$

where

$$
\widehat{Q}_{ijkl} = \frac{1}{n}\big(\widehat{\mu}_{4,ijkl} - \widehat{\sigma}^2_{ij}\,\widehat{\sigma}^2_{kl}\big) \ +\ \frac{1}{n(n-1)}\big(\widehat{\sigma}^2_{ik}\,\widehat{\sigma}^2_{jl} + \widehat{\sigma}^2_{il}\,\widehat{\sigma}^2_{jk}\big).
$$

Here $\widehat{\mu}_{4,ijkl}$ *is the sample-average estimator of the fourth central moment*

$$
\mu_{4,ijkl} = \mathbb{E}\big[(X_i-\mu_i)(X_j-\mu_j)(X_k-\mu_k)(X_l-\mu_l)\big],
$$

and $\widehat{\sigma}^2_{ij}$ *is the sample-average estimator of the covariance*

$$
\sigma^2_{ij} = \mathbb{E}\big[(X_i-\mu_i)(X_j-\mu_j)\big].
$$

2.3.1 Rank-1 Approximation of MV-PBR

Here, we make a rank-1 approximation of the quartic polynomial constraint:

$$
(w^\top \hat{\alpha})^4 \ \approx\ \sum_{i,j,k,l=1}^{p} w_i\, w_j\, w_k\, w_l\ \widehat{Q}_{ijkl},
$$

By matching the diagonal terms, we get $\hat{\alpha}_i$ as

$$
\hat{\alpha}_i = \sqrt[4]{\widehat{Q}_{iiii}} = \sqrt[4]{\frac{1}{n}\,\hat{\mu}_{4,\,iiii} - \frac{n-3}{n(n-1)}\big(\hat{\sigma}^2_{ii}\big)^2}\,.
$$

We thus obtain the following convex approximation of (MV-PBR-1):

$$
\begin{aligned}
\widehat{w}_{n,\mathrm{PBR1}} = \arg\min_{w\in\mathbb{R}^p} & w^\top \widehat{\Sigma}_n\, w\\
\text{s.t.}\quad & w^\top 1_p = 1,\\
& (w^\top \widehat{\mu}_n = R),\\
& w^\top \hat{\alpha} \le \sqrt[4]{U}.
\end{aligned}
\qquad \text{(MV-PBR-1)}
$$

Proposition 2 (Solution to MV-PBR-1) *The solution to (MV-PBR-1) with the mean constraint* $w^\top \widehat{\mu}_n = R$ *is*

$$
\widehat{w}_{n,\mathrm{PBR1}} = \widehat{w}_{n,\mathrm{MV}} - \tfrac{1}{2}\,\lambda^*\,\widehat{\Sigma}_n^{-1}\big(\beta_1\, 1_p + \beta_2\,\widehat{\mu}_n + \hat{\alpha}\big).
$$

The coefficients β_1 *and* β_2 *are given by*

$$
\beta_1 = \frac{\hat{\alpha}^\top \widehat{\Sigma}_n^{-1}\widehat{\mu}_n\ \widehat{\mu}_n^\top \widehat{\Sigma}_n^{-1} 1_p - \hat{\alpha}^\top \widehat{\Sigma}_n^{-1} 1_p\ \widehat{\mu}_n^\top \widehat{\Sigma}_n^{-1}\widehat{\mu}_n}{1_p^\top \widehat{\Sigma}_n^{-1} 1_p\ \widehat{\mu}_n^\top \widehat{\Sigma}_n^{-1}\widehat{\mu}_n - (\widehat{\mu}_n^\top \widehat{\Sigma}_n^{-1} 1_p)^2},
$$

$$\beta_2 = \frac{\hat{\alpha}^\top \widehat{\Sigma}_n^{-1} 1_p \; 1_p^\top \widehat{\Sigma}_n^{-1} \widehat{\mu}_n - \hat{\alpha}^\top \widehat{\Sigma}_n^{-1} 1_p \; 1_p^\top \widehat{\Sigma}_n^{-1} 1_p}{1_p^\top \widehat{\Sigma}_n^{-1} 1_p \; \widehat{\mu}_n^\top \widehat{\Sigma}_n^{-1} \widehat{\mu}_n - (\widehat{\mu}_n^\top \widehat{\Sigma}_n^{-1} 1_p)^2}.$$

Moreover, the solution to (mv-PBR-1) without the mean constraint is

$$\widehat{w}_{n,\mathrm{PBR1}} = \widehat{w}_{n,\mathrm{MV}} - \tfrac{1}{2} \lambda^* \widehat{\Sigma}_n^{-1} \big(\beta_1 1_p + \hat{\alpha} \big).$$

Here $\widehat{w}_{n,\mathrm{MV}}$ is the SAA solution, λ^* is the optimal Lagrange multiplier for $w^\top \hat{\alpha} \le \sqrt[4]{U}$, and

$$\beta = -\frac{1_p^\top \widehat{\Sigma}_n^{-1} \hat{\alpha}}{1_p^\top \widehat{\Sigma}_n^{-1} 1_p}.$$

2.3.2 Best Convex Quadratic Approximation of MV-PBR

We also consider a convex quadratic approximation of the quartic polynomial constraint:

$$(w^\top A w)^2 \;\approx\; \sum_{i,j,k,l=1}^{p} w_i w_j w_k w_l \, \widehat{Q}_{ijkl},$$

where A is a positive semidefinite (PSD) matrix. Expanding the left-hand side gives

$$\sum_{i,j,k,l=1}^{p} w_i w_j w_k w_l \, A_{ij} A_{kl}.$$

To match the diagonal terms $A_{ij}^2 \approx \widehat{Q}_{ijij}$, we solve the semidefinite program

$$A^* \;=\; \arg\min_{A \succeq 0} \|A - Q_2\|_F,$$

where $\|\cdot\|_F$ is the Frobenius norm and $(Q_2)_{ij} = \widehat{Q}_{ijij}$. This leads to the quadratic approximation of (MV-PBR):

$$\begin{aligned} \widehat{w}_{n,\mathrm{PBR2}} = \arg\min_{w \in \mathbb{R}^p} \; & w^\top \widehat{\Sigma}_n w \\ \text{s.t.} \quad & w^\top 1_p = 1, \\ & (w^\top \widehat{\mu}_n = R), \\ & w^\top A^* w \le \sqrt{U}, \end{aligned} \tag{MV-PBR-2}$$

Proposition 3 (Solution to MV-PBR-2) *With the mean constraint $w^\top \widehat{\mu}_n = R$, the solution to* (MV-PBR-2) *is*

$$\widehat{w}_{n,\mathrm{PBR2}} = -\tfrac{1}{2} \widetilde{\Sigma}_n(\lambda^*)^{-1} \big(v_1^*(\lambda^*) \, 1_p + v_2^*(\lambda^*) \, \widehat{\mu}_n \big),$$

where

$$\widetilde{\Sigma}_n(\lambda^*) = \widehat{\Sigma}_n + \lambda^* A^*, \quad \lambda^* \text{ is the optimal multiplier for } w^\top A^* w \le \sqrt{U},$$

and the weights

$$v_1^*(\lambda) = 2 \frac{R \, \widehat{\mu}_n^\top \widetilde{\Sigma}_n(\lambda)^{-1} 1_p - \widehat{\mu}_n^\top \widetilde{\Sigma}_n(\lambda)^{-1} \widehat{\mu}_n}{1_p^\top \widetilde{\Sigma}_n(\lambda)^{-1} 1_p \; \widehat{\mu}_n^\top \widetilde{\Sigma}_n(\lambda)^{-1} \widehat{\mu}_n - (\widehat{\mu}_n^\top \widetilde{\Sigma}_n(\lambda)^{-1} 1_p)^2},$$

$$v_2^*(\lambda) = 2 \frac{-R \, 1_p^\top \widetilde{\Sigma}_n(\lambda)^{-1} 1_p + \widehat{\mu}_n^\top \widetilde{\Sigma}_n(\lambda)^{-1} 1_p}{1_p^\top \widetilde{\Sigma}_n(\lambda)^{-1} 1_p \; \widehat{\mu}_n^\top \widetilde{\Sigma}_n(\lambda)^{-1} \widehat{\mu}_n - (\widehat{\mu}_n^\top \widetilde{\Sigma}_n(\lambda)^{-1} 1_p)^2}.$$

Without the mean constraint, the solution is

$$\widehat{w}_{n,\mathrm{PBR2}} = \frac{\widetilde{\Sigma}_n(\lambda^*)^{-1} 1_p}{1_p^\top \widetilde{\Sigma}_n(\lambda^*)^{-1} 1_p}.$$

2.4 PBR for Mean–CVaR Optimization

The PBR formulation for the mean–CVaR problem adds upper-bounds on the variances of both the CVaR estimator and the sample mean to the SAA model:

$$
\begin{aligned}
\widehat{w}_{n,\mathrm{PBR}} = \arg\min_{w\in\mathbb{R}^p}\ & \widehat{\mathrm{CVaR}}_n(-w^\top X;\beta) \\
\text{s.t.}\quad & w^\top 1_p = 1, \\
& (w^\top \widehat{\mu}_n = R), \\
& \mathrm{SVar}\big(\widehat{\mathrm{CVaR}}_n(-w^\top X;\beta)\big) \le U_1, \\
& \mathrm{SVar}(w^\top \widehat{\mu}_n) \le U_2,
\end{aligned}
\tag{CV-PBR}
$$

Since

$$
\mathrm{Var}\big(w^\top \widehat{\mu}_n\big) = \frac{1}{n^2}\sum_{i=1}^{n} \mathrm{Var}(w^\top X_i) = \frac{1}{n}\, w^\top \Sigma\, w,
$$

we have

$$
\mathrm{SVar}(w^\top \widehat{\mu}_n) = n^{-1} w^\top \widehat{\Sigma}_n\, w.
$$

Proposition 4 *Under mild smoothness assumptions on the distribution of X, the variance of the CVaR estimator satisfies*

$$
\mathrm{Var}\big[\widehat{\mathrm{CVaR}}_n(-w^\top X;\beta)\big] = \frac{1}{n(1-\beta)^2}\,\mathrm{Var}\big[\big(-w^\top X - \alpha_\beta(w)\big)^+\big] + O(n^{-2}),
$$

where

$$
\alpha_\beta(w) = \inf\{\alpha :\ P(-w^\top X \ge \alpha) \le 1-\beta\}
$$

is the value-at-risk (VaR) of the portfolio at level β.

Thus, to first order,

$$
\mathrm{SVar}\big(\widehat{\mathrm{CVaR}}_n(-w^\top X;\beta)\big) = \frac{1}{n(1-\beta)^2}\, z^\top \Omega_n\, z,
$$

where $z_i = \max(0, -w^\top X_i - \alpha)$ and $\Omega_n = (1/(n-1))(I_n - n^{-1} 1_n 1_n^\top)$.

Accordingly, one can write the finite-sample PBR problem with explicit sample-variance constraints as

$$
\begin{aligned}
\min_{\alpha,w,z}\quad & \alpha + \frac{1}{n(1-\beta)}\sum_{i=1}^{n} z_i \\
\text{s.t.}\quad & w^\top 1_p = 1, \quad (w^\top \widehat{\mu}_n = R), \\
& \frac{1}{n(1-\beta)^2}\, z^\top \Omega_n\, z \le U_1, \\
& z_i = \max\big(0,\, -w^\top X_i - \alpha\big),\ i = 1,\dots,n, \\
& \frac{1}{n}\, w^\top \widehat{\Sigma}_n\, w \le U_2,
\end{aligned}
\tag{CV-PBR$'$}
$$

3 Calibration of Regularization Level

This section introduces the methodologies and algorithms authors used to search for the optimal regularization level U^*. Due to some limitations of their method, we put forward an alternative approach for selecting for U^*.

3.1 Practical Failure of Original OOS-PBSD Algorithm

Ban et al. [1] proceed as follows:

1. Compute feasible lower and upper bounds, $\underline{U}$ and $\bar{U}$, to define the search interval.

2. Perform performance-based k-fold cross-validation, using the out-of-sample performance-based steepest-descent (PBSD) Algorithm 4 to find the fold-wise estimates U^*_{-b}.
3. Aggregate by taking the arithmetic mean:

$$U^* = \frac{1}{k} \sum_{b=1}^{k} U^*_{-b}.$$

However, based on the practical experiments, we found that the OOS-PBSD procedure often gets stuck in an infinite loop or generates a U^* that is trivial to regularize on the model, resulting portfolio remains unchanged. We conclude the reasons that OOS-PBSD breaks down as follows:

- Even initialize at the theoretical lower bound $\underline{U}_{-b}$, the PBR penalty is still inactive, so $\hat{w}(U)$ (and hence $Sharpe(\hat{w}(U))$) remains constant over a non-trivial interval.
- Because the algorithm approximates $\frac{d\hat{w}(U)}{dU}$ by a finite difference $\frac{w(U)-w((1-bit)U)}{bit \cdot U}$, the resulting $\frac{Sharpe(\hat{w}(U))}{dU}$ is neither smooth nor monotonic, violating the smoothness assumptions of backtracking line search.
- the Armijo condition $Sharpe(U - t\Delta U) \geq Sharpe(U) + \alpha t \Delta U \frac{dSharpe}{dU}$ is immediately satisfied at initial t=1. In this case, the lowest U^* only allows to be $(1 - \frac{1}{div})\bar{U}_{-b}$ without searching for the $[\underline{U}_{-b}, (1 - \frac{1}{div})\bar{U}_{-b}]$ ranges.

3.2 Alternative Approaches for Selecting the Regularization Level U^*

Since our primary objective is to compare the PBR models with the SAA baseline, we aim to obtain *different* portfolios under the two approaches. These observations drive the motivation for the following grid-search approaches.

Geometric-Mean Aggregation In lieu of the arithmetic average, we aggregate the fold-wise estimates by their geometric mean:

$$U^* \;=\; \Bigl(\prod_{b=1}^{k} U^*_{-b}\Bigr)^{1/k}.$$

This approach mitigates the influence of extreme U^*_{-b} values, yielding a more balanced and stable regularization level.

Algorithm 1: Out-of-Sample Performance-Based k-Fold Cross-Validation (OOS-PBCV)

Input: Full training set X_{train}, number of folds k, target return R_{target}
Output: Regularization level U^*
Randomly split X_{train} into k equal folds $\{X^b\}_{b=1}^k$;
for $b \leftarrow 1$ **to** k **do**
  Define training sets and validation sets

$$X_{\text{train}}^{-b} \;=\; X_{\text{train}} \setminus X^b \quad , \quad X_{\text{val}}^b \;=\; X^b$$

  Get the best U on this fold according to OOS-Sharpe

$$U^*_{-b} \;\leftarrow\; \text{OOS-PBGS}\bigl(D_{\text{train}}^{-b},\, D_{\text{val}}^b,\, R_{\text{target}}\bigr)$$

Compute the geometric mean of $\bigl\{U^*_{-b}\bigr\}_{b=1}^k$

$$U^* \;\leftarrow\; \left(\prod_{b=1}^{k} U^*_{-b}\right)^{\frac{1}{k}}$$

return U^*;

Grid Search over $\mathcal{U}$ To impose a stronger and more global regularization, we replace the local descent search with a grid search over an approximately logarithmic grid $\{U_i\}_{i=1}^m \subset \mathcal{U}$. For each candidate U_i:

1. Solve the PBR problem to obtain $\hat{w}(U_i)$.
2. Compute the deviation $\|\hat{w}(U_i) - \hat{w}(\bar{U})\|$ from the baseline solution $\hat{w}(\bar{U})$.

We retain only those $\hat{w}(U_i)$ whose deviation lies in the top percentile, then evaluate their out-of-sample Sharpe ratios. The U^* corresponding to the highest Sharpe among this shortlist is selected. This procedure is detailed in Algorithm 1 and Algorithm 2.

Algorithm 2: Out-of-Sample Performance-Based Grid Search (OOS-PBGS)

Input: Training set X_{train}, validation set X_{valid}, method type `pbr`, target return R_{target}, fraction of candidates ρ, regularization set $\mathcal{U}$, (confidence level β for CVaR-based methods)

Output: Selected regularization level U^*

Estimate sample statistics:

$$\hat{\mu} = \frac{1}{n}\sum_{i=1}^{n} x_i, \quad \hat{\Sigma} = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \hat{\mu})(x_i - \hat{\mu})^\top$$

Compute the baseline portfolio:

$$w_0 \leftarrow \begin{cases} \text{MV\_opt}(\hat{\mu}, \hat{\Sigma}, R_{\text{target}}), & \text{if } \texttt{pbr} = \texttt{mv\_pbr1} \\ \text{CVaR\_opt}(X_{\text{train}}, \beta, R_{\text{target}}), & \text{if } \texttt{pbr} = \texttt{cvar\_pbr} \end{cases}$$

Initialize storage:

$$\mathcal{N} \leftarrow \emptyset, \qquad \mathcal{S} \leftarrow \emptyset$$

foreach $U \in \mathcal{U}$ **do**

Compute the regularized portfolio:

$$w \leftarrow \begin{cases} \text{mv\_pbr\_opt}(X_{\text{train}}, \hat{\mu}, \hat{\Sigma}, U, R_{\text{target}}), & \text{if } \texttt{pbr} = \texttt{mv\_pbr1} \\ \text{cvar\_pbr\_opt}(X_{\text{train}}, \beta, U, R_{\text{target}}), & \text{if } \texttt{pbr} = \texttt{cvar\_pbr} \end{cases}$$

Compute change in weights and validation Sharpe:

$$\Delta w \leftarrow w - w_0, \qquad \mathcal{N} \leftarrow \mathcal{N} \cup \{\|\Delta w\|_2\}, \qquad \mathcal{S} \leftarrow \mathcal{S} \cup \{\mathit{Sharpe}(X_{\text{valid}}, w)\}$$

Select top-ρ candidates by norm change:

$$m \leftarrow \lceil \rho \times |\mathcal{N}| \rceil$$

$$\mathcal{I} \leftarrow \text{indices of top-}m \text{ elements in } \mathcal{N}$$

Select best candidate based on Sharpe:

$$i^* \leftarrow \arg\max_{i \in \mathcal{I}} \mathcal{S}[i] \quad \Rightarrow \quad U^* \leftarrow \mathcal{U}[i^*]$$

return U^*

---

4 Backtesting

4.1 Data Sources and Description

This study uses daily value-weighted returns of 5- and 10-industry portfolios, constructed based on the CRSP industry classification from the CRSP database. The dataset is publicly available from the Kenneth French Data Library[2] [2].

[2] `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`

Table 1: Industry-index Fields in CRSP Database

Dataset	Industry-index Fields
FF5 Industries	Cnsmr, Manuf, HiTec, Hlth, Other
FF10 Industries	NoDur, Durbl, Manuf, Enrgy, HiTec, Telcm, Shops, Hlth, Utils, Other

4.2 Backtesting Setup

The backtesting period spans from 01/01/2010 to 12/31/2019, covering ten years of daily data. We adopt a rolling window framework to evaluate out-of-sample portfolio performance.

At each iteration, the model is trained using the most recent T_{train} days of data, and the resulting optimal weights are applied to the following T_{test} days. The window is then rolled forward by T_{test} days, and the process repeats until the end of the test period.

To evaluate robustness, we conduct experiments across different datasets and parameter configurations. The grid of experimental settings is as follows:

- **Dataset:** FF5 (5-industry portfolios), FF10 (10-industry portfolios)
- **Training window length** (T_{train})**:** 60, 120 days
- **Rebalancing frequency** (T_{test})**:** 60, 20, 10 days
- **Target return** (R_{target})**:** None (unconstrained), 0, 0.001, 0.002

4.3 Portfolio Strategies

We evaluate the out-of-sample performance of the following five portfolio allocation rules:

- **Simply equal weighted**
- **Sample average approximation (MV-SAA)** [Eq. MV-SAA]
- **Performance-based regularization–rank1 (MV-PBR-1)** [Eq. MV-PBR-1]
- **Sample average approximation (CV-SAA)** [Eq. CV-SAA]
- **Performance-based regularization (CV-PBR) with U2=∞** [Eq. CV-PBR]

In each PBR variant, we tune the remaining U-parameter to maximize the average out-of-sample Sharpe ratio over the k folds. A full description of our OOS-PBCV procedure appears in Section 3.2.

4.4 Evaluation Methodology

4.4.1 Out-of-Sample Sharpe Ratio

The out-of-sample daily return at time t is computed as the portfolio-weighted sum of individual asset returns:

$$r_t = \sum_{i=1}^{N} w_i \cdot r_t^{(i)}, \quad t = 1, \ldots, T$$

where w_i is the weight allocated to asset i, and $r_t^{(i)}$ is the return of asset i on day t.

The Sharpe ratio over the testing period is then calculated as:

$$\textit{Sharpe} = \frac{\hat{\mu}}{\hat{\sigma}} = \frac{\frac{1}{T}\sum_{t=1}^{T} r_t}{\sqrt{\frac{1}{T-1}\sum_{t=1}^{T}(r_t - \bar{r})^2}}, \quad \text{where } \bar{r} = \frac{1}{T}\sum_{t=1}^{T} r_t$$

4.4.2 Out-of-Sample Annualized Return

The out-of-sample cumulative return (CR) at time t is computed as the cumulative compounded return of the portfolio up to time t. Specifically, we define:

$$\text{CR}_t = \prod_{s=1}^{t}(1 + r_s) - 1, \quad t = 1, \ldots, T,$$

where r_s is the portfolio return at time s, computed as in the previous section. Convert cumulative return into annualized return given by:

$$\text{Annualized Return} = (1 + CR_T)^{\frac{252}{T}} - 1.$$

This measure captures the average growth of the portfolio value over the out-of-sample period.

5 Result & Discussion

This section summarizes FF5 and FF10 performance across PBR models, lookback windows, rebalancing frequencies, and return targets.

5.1 Back-Test Results

Our back-test results comprise three components:

- **Sharpe ratios** — summarised in Tables 2 and 3.
- **Annualized returns** — summary values are reported in Appendix B.1;
- **NAV back-testing graphs** — visual performance comparisons of net asset values. We analyze one representative graph in the main textin Figure 1, while the complete set is provided in Appendix B.2, B.3, B.4, and B.5.

5.2 FF5 Sharpe Ratios

Table 2: Sharpe Ratios for FF5 under Different Parameter Configurations

R_target	Lookback	Frequency	Equal	MV-SAA	MV-PBR-1	CV-SAA	CV-PBR
None	60	10	0.933	1.080	0.582	0.734	0.674
None	60	20	0.933	1.108	0.667	0.789	0.604
None	60	60	0.933	0.854	0.341	0.702	0.590
None	120	10	0.933	1.043	0.419	0.855	1.021
None	120	20	0.933	1.024	0.452	0.874	1.000
None	120	60	0.933	0.949	0.349	0.862	0.689
0	60	10	0.933	0.816	0.702	0.660	0.946
0	60	20	0.933	0.909	0.752	0.730	0.916
0	60	60	0.933	0.821	0.166	0.554	0.734
0	120	10	0.933	0.845	0.387	0.760	0.776
0	120	20	0.933	0.585	0.206	0.761	0.604
0	120	60	0.933	0.665	0.339	0.633	0.551
0.001	60	10	0.933	0.878	0.790	0.817	0.862
0.001	60	20	0.933	0.951	0.862	0.895	0.912
0.001	60	60	0.933	0.685	0.154	0.534	0.664
0.001	120	10	0.933	0.979	0.402	1.033	0.854
0.001	120	20	0.933	0.880	0.227	0.933	0.547
0.001	120	60	0.933	0.731	0.384	0.839	0.398
0.002	60	10	0.933	0.725	0.795	0.520	0.427
0.002	60	20	0.933	0.784	0.884	0.650	0.653
0.002	60	60	0.933	0.430	0.145	0.215	0.231
0.002	120	10	0.933	0.695	0.117	0.580	0.472
0.002	120	20	0.933	0.564	0.205	0.398	0.266
0.002	120	60	0.933	0.511	0.353	0.596	0.251

Notes. Red entries indicate cases where the PBR-regularized strategy outperforms the unregularized one.

- **Mild target helps, aggressive target hurts:** Adding a modest return target ($R_{\text{target}} = 0.001$) boosts Sharpe across most methods (e.g., = 1.033), but pushing it further to 0.002 leads to sharp declines, especially for MV-PBR-1 (as low as 0.117).

- **CVaR models are more robust:** CV-PBR consistently outperforms MV-PBR-1, especially under return constraints. Downside-focused risk regularization appears more resilient in finite-sample settings.
- **Lookback = 60 is often enough:** Longer lookbacks (120) offer limited gain and sometimes degrade performance due to slower adaptation (e.g., MV-SAA drops from 1.108 to 1.024 at freq = 20).
- **Medium rebalancing frequency performs best:** Frequency = 20 strikes a good balance—too frequent (10) leads to noisy weights; too infrequent (60) causes lag (e.g., MV-PBR-1 peaks at 0.884 for freq = 20).
- **PBR can backfire without careful tuning:** MV-PBR-1 shows large variability—high when conditions are favorable, disastrous otherwise. Rank-1 regularization needs caution.
- **Equal weight is a strong baseline:** Despite zero optimization, equal-weight achieves a steady 0.933 across all configs. It often beats poorly tuned optimizers and sets a high bar.

5.3 FF10 Sharpe Ratios

Table 3: Sharpe Ratios for FF10 under Different Parameter Configurations

R_target	Lookback	Frequency	Equal	MV-SAA	MV-PBR-1	CV-SAA	CV-PBR
None	60	10	0.891	1.143	0.773	1.081	1.134
None	60	20	0.891	1.126	0.795	1.034	1.141
None	60	60	0.891	1.010	0.846	0.933	0.942
None	120	10	0.891	1.229	0.713	0.977	0.921
None	120	20	0.891	1.226	0.673	1.173	1.130
None	120	60	0.891	1.148	0.771	1.153	1.108
0	60	10	0.891	1.060	0.653	1.052	1.088
0	60	20	0.891	1.138	0.834	1.248	1.356
0	60	60	0.891	0.821	0.166	0.554	0.734
0	120	10	0.891	1.114	0.739	1.211	1.278
0	120	20	0.891	1.185	0.807	1.312	1.288
0	120	60	0.891	0.964	0.687	0.845	0.996
0.001	60	10	0.891	1.372	0.895	1.323	1.412
0.001	60	20	0.891	1.356	0.996	1.379	1.431
0.001	60	60	0.891	0.685	0.154	0.534	0.664
0.001	120	10	0.891	1.368	0.891	1.344	1.187
0.001	120	20	0.891	1.353	0.928	1.423	1.316
0.001	120	60	0.891	1.098	0.806	1.039	1.148
0.002	60	10	0.891	1.413	1.040	1.183	1.202
0.002	60	20	0.891	1.327	1.063	1.086	1.158
0.002	60	60	0.891	0.756	0.629	0.672	0.654
0.002	120	10	0.891	1.187	0.897	1.034	1.132
0.002	120	20	0.891	1.123	0.903	0.958	0.970
0.002	120	60	0.891	0.911	0.808	0.770	0.899

Notes. Red entries indicate cases where the PBR-regularized strategy outperforms the unregularized one.

- **Mild target boosts, aggressive target moderates:** Introducing a modest return target $R_{\text{target}} = 0.001$) lifts Sharpe ratios for *all* methods—most notably CV-PBR, which reaches 1.431 (look-back 60, freq.=20). Pushing the target to 0.002 tempers performance; only MV-SAA maintains a Sharpe above 1.40, while CV-PBR slips to 1.202.
- **CVaR models remain the most robust:** Across nearly every configuration with a return constraint, CV-PBR outperforms MV-PBR-1 (red entries are concentrated in the CV-PBR column). Downside-risk regularisation is therefore more resilient to factor-model expansion from FF5 to FF10.

- **Lookback = 60 is still sufficient:** Extending the window to 120 months yields marginal (sometimes negative) gains. For example, MV-SAA drops from 1.143 (look-back 60, freq.=10) to 1.229 (look-back 120) under no target, but loses relative ground once a target is imposed.
- **Medium rebalancing frequency performs best:** A 20-day schedule balances responsiveness and noise. Too frequent (10 days) makes MV-PBR-1 volatile (e.g. Sharpe 0.773); too infrequent (60 days) causes lag (CV-PBR falls from 1.288 at freq.=20 to 0.996 at freq.=60 for $R_{\text{target}} = 0$).
- **PBR can still backfire without careful tuning:** MV-PBR-1 exhibits large dispersion: highs of 1.063 (target 0.002, freq.=20), but lows of 0.166 (target 0, freq.=60). Rank-1 regularisation therefore requires cautious calibration.
- **Equal weight remains a tough benchmark:** The naïve portfolio again posts a stable Sharpe (0.891) and is beaten consistently only by well-tuned MV-SAA and CV-PBR configurations, underscoring its value as a baseline.

5.4 Annualized Return of MV-PBR-1

Although our earlier Sharpe ratio analysis suggested that MV-PBR-1 under-performs the MV-SAA strategy on a risk-adjusted basis, its net asset value curve tells a different story.

An example is shown in Figure 1. The MV-PBR-1 portfolio delivers far greater cumulative NAV growth, reaching almost 20x of initial capital by early 2018, whereas the standard MV-SAA strategy only grows to about 4× over the same period. This example illustrates that, despite a slightly lower Sharpe, MV-PBR-1 can produce substantially higher absolute returns in favorable market regimes.

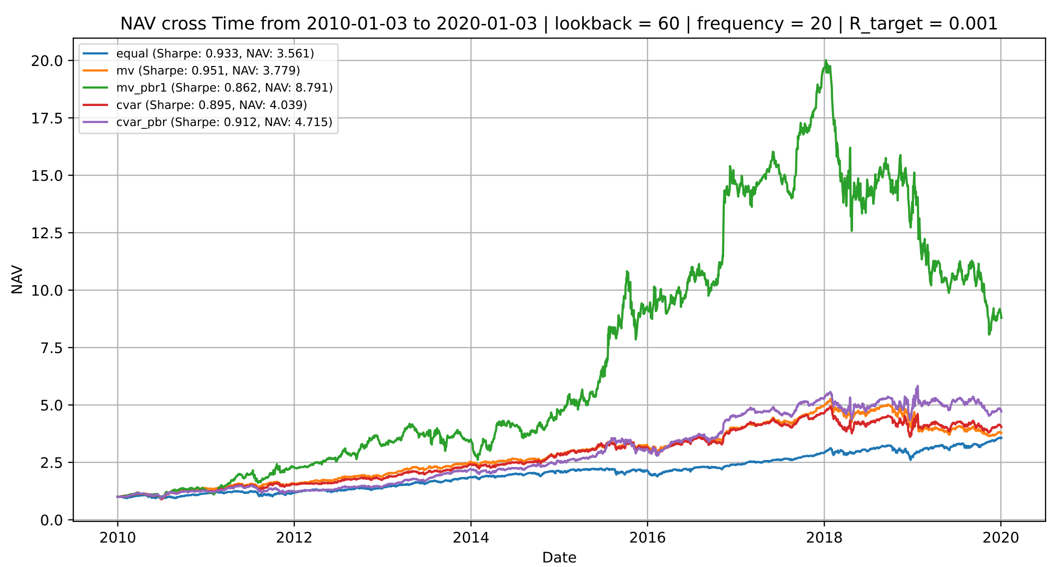

Figure 1: FF5 Performance with Lookback=60, Frequency=20, Target Return=1e-3

This observation is further supported by Table 4 and Table 5, where the MV-PBR-1 group consistently achieves the highest annualized return across several configurations—for instance, reaching 23.64% in one setting compared to only 14.50% for the standard MV-SAA strategy. Such results demonstrate that, in terms of pure return generation, MV-PBR-1 can significantly outperform traditional approaches despite exhibiting slightly lower Sharpe ratios.

From Table 4 and Table 5, we observe that the MV-PBR-1 strategy delivers significant improvement in annualized return over the vanilla MV-SAA strategy when the lookback window is short (e.g., 60). The performance gain is especially notable under more aggressive return targets, where MV-PBR-1 outperforms MV-SAA by a wide margin. However, the advantage diminishes as the lookback window extends, suggesting that MV-PBR-1 is most effective in more reactive, shorter-horizon settings.

6 Conclusion

In this paper, we analyze and implement the Performance-Based Regularization (PBR) model for both mean-variance (MV) and Conditional Value-at-Risk (CVaR) portfolio optimization. Our empirical results show that the PBR model performs particularly well when applied to CVaR optimization. Although it underperforms in the MV setting in terms of Sharpe ratio, it still delivers strong annualized returns overall, making it a promising framework worthy of further investigation.

References

[1] Gah-Yi Ban, Noureddine El Karoui, and Andrew EB Lim. Machine learning and portfolio optimization. *Management Science*, 64(3):1136–1154, 2018.

[2] Eugene F Fama and Kenneth R French. Production of us rm-rf, smb, and hml in the fama-french data library. *Chicago Booth Research Paper*, (23-22), 2023.

[3] Andrew EB Lim, J George Shanthikumar, and Gah-Yi Vahn. Conditional value-at-risk in portfolio optimization: Coherent but fragile. *Operations Research Letters*, 39(3):163–171, 2011.

[4] Harry M Markowits. Portfolio selection. *Journal of finance*, 7(1):71–91, 1952.

[5] R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

A Appendix - Algorithms

A.1 Original OOS-PBCV

Algorithm 3: Out-of-Sample Performance-Based k-Fold Cross-Validation (OOS-PBCV) [1]

Input: Full training set D_{train}, number of folds k
Output: Regularization level U^*
Solve PBR(U) on D_{train} to get $\hat{w}_{\text{train}}$
Compute $\bar{U} = \widehat{Risk_n}(-\hat{w_{train}}^T X)$
Solve U-min(U) on D_{train} to get $\hat{w}_{U_{\min}}$
Compute $\underline{U} = \widehat{Risk_n}(-\hat{w_{U_{min}}}^T X)$
Randomly split D_{train} into k equal folds $\{D^b_{\text{train}}\}^k_{b=1}$;
for $b \leftarrow 1$ **to** k **do**
  Define $D^{-b}_{\text{train}} = D_{\text{train}} \setminus D^b_{\text{train}}$
  Solve PBR(U) on D^{-b}_{train} to get $\hat{w}_{-b}$ Compute $\bar{U}_{-b} = \widehat{Risk_n}(-\hat{w_{-b}}^T X)$
  if $\bar{U}_{-b} < \underline{U}$ **then**
    Set $U^*_{-b} = \underline{U}$ and continue;
  else
    Solve U-min(U) on D^{-b}_{train} to get $\hat{w}^{\min}_{-b}$ Compute $\underline{U}_{-b} = \widehat{Risk_n}(-\hat{w_{-b}}^T X)$
    if $\underline{U}_{-b} > \bar{U}$ **then**
      Set $U^*_{-b} = \bar{U}$ and terminate
    else
      Update boundaries: $\bar{U}_{-b} = \min(\bar{U}, \bar{U}_{-b})$, $\underline{U}_{-b} = \max(\underline{U}, \underline{U}_{-b})$
      Run `OOS-PBSD` with boundaries $[\underline{U}_{-b}, \bar{U}_{-b}]$ to find U^*_{-b}

Compute final output: $U^* = \frac{1}{k} \sum_{b=1}^k U^*_{-b}$;

Where **U-min-MV1** is defined as:

$$\begin{aligned} \underline{U} = \min_{w} \quad & w^\top \alpha \\ \text{s.t.} \quad & w^\top \mathbf{1}_p = 1, \\ & w^\top \hat{\mu}_n = R, \end{aligned} \tag{U-min-MV1}$$

U-min-MV2 is defined as:

$$\begin{aligned} \underline{U} = \min_{w} \quad & w^\top A^\star w \\ \text{s.t.} \quad & w^\top \mathbf{1}_p = 1, \\ & w^\top \hat{\mu}_n = R, \end{aligned} \tag{U-min-MV2}$$

U-min-CV is defined as:

$$\begin{aligned} \underline{U} = \min_{w,z} \quad & z^\top \Omega_n z \\ \text{s.t.} \quad & w^\top \hat{\mu}_n = R, \\ & w^\top \mathbf{1}_p = 1, \\ & z_i \geq -w^\top X_i - \alpha, \quad i = 1, \dots, n, \\ & z_i \geq 0, \quad i = 1, \dots, n, \end{aligned} \tag{U-min-CV}$$

A.2 Original OOS-PBSD Algorithm

Algorithm 4: Out-of-Sample Performance-Based Steepest Descent (OOS-PBSD) [1]

Input: Training set D^{-b}_{train}, Valid set D^{b}_{valid}, boundaries $[\underline{U}_{-b}, \bar{U}_{-b}]$
Output: Optimal regularization level U^*_{-b}
Choose backtracking parameters $\alpha \in (0, 0.5)$, $\beta \in (0, 1)$;
Choose stepsize Div and perturbation size $bit \in (0, 0.5)$;
Initialize: $U = \bar{U}_{-b}$, $\Delta U = (\bar{U}_{-b} - \underline{U}_{-b})/Div$, $t = 1$;
Compute:

$$\frac{d\text{Sharpe}(U)}{dU} = \nabla_w \text{Sharpe}(\hat{w}_{-b}(U))^\top \cdot \left[\frac{d\hat{w}_{-b}(U)}{dU}\right]$$

where

$$\nabla_w \text{Sharpe}(\hat{w}_{-b}(U)) = \frac{(\hat{w}^\top_{-b}\Sigma\hat{w}_{-b})\mu - (\hat{w}^\top_{-b}\mu)\Sigma\hat{w}_{-b}}{(\hat{w}^\top_{-b}\Sigma\hat{w}_{-b})^{3/2}}$$

$$\frac{d\hat{w}_{-b}(U)}{dU} = \frac{\hat{w}_{-b}(U) - \hat{w}_{-b}((1-bit)U)}{bit \cdot U}$$

while $\text{Sharpe}(U - t\Delta U) < \text{Sharpe}(U) + \alpha t \Delta U \cdot \frac{d\text{Sharpe}(U)}{dU}$ **do**
  Update $t = \beta t$;
Return $U^*_{-b} = U - t\Delta U$;

B Appendix - Backtesting Performance

B.1 Annualized Returns of Backtest

Table 4: FF5 Annualized Returns

R_target	Lookback	Frequency	Equal	MV-SAA	MV-PBR-1	CV-SAA	CV-PBR
None	60	10	13.54%	13.61%	12.14%	10.41%	9.97%
None	60	20	13.54%	13.98%	14.16%	11.21%	8.81%
None	60	60	13.54%	10.71%	5.46%	10.26%	8.71%
None	120	10	13.54%	12.59%	7.59%	10.40%	14.60%
None	120	20	13.54%	12.50%	8.32%	11.48%	15.11%
None	120	60	13.54%	11.29%	5.46%	11.48%	9.86%
0	60	10	13.54%	12.05%	19.11%	9.99%	16.92%
0	60	20	13.54%	13.43%	20.34%	10.97%	16.94%
0	60	60	13.54%	11.66%	1.24%	7.75%	11.87%
0	120	10	13.54%	11.77%	7.99%	11.05%	12.89%
0	120	20	13.54%	11.96%	–1.71%	11.09%	9.68%
0	120	60	13.54%	9.87%	5.91%	9.63%	9.00%
0.001	60	10	13.54%	12.60%	22.67%	11.55%	15.09%
0.001	60	20	13.54%	14.50%	23.64%	13.03%	15.73%
0.001	60	60	13.54%	11.38%	0.13%	7.55%	10.77%
0.001	120	10	13.54%	13.39%	8.47%	15.07%	13.12%
0.001	120	20	13.54%	12.56%	1.20%	14.12%	8.54%
0.001	120	60	13.54%	10.35%	7.54%	12.49%	5.51%
0.002	60	10	13.54%	12.80%	24.89%	9.26%	8.01%
0.002	60	20	13.54%	14.62%	26.67%	11.83%	13.45%
0.002	60	60	13.54%	6.60%	–0.86%	2.31%	2.66%
0.002	120	10	13.54%	12.74%	–4.07%	11.77%	9.93%
0.002	120	20	13.54%	9.98%	–1.37%	7.02%	6.42%
0.002	120	60	13.54%	9.68%	6.54%	11.90%	5.62%

Table 5: FF10 Annualized Returns

R_target	Lookback	Frequency	Equal	MV-SAA	MV-PBR-1	CV-SAA	CV-PBR
None	60	10	13.54%	13.61%	12.14%	10.41%	9.97%
None	60	20	13.54%	13.98%	14.16%	11.21%	8.81%
None	60	60	13.54%	10.71%	5.46%	10.26%	8.71%
None	120	10	13.54%	12.59%	7.59%	10.40%	14.60%
None	120	20	13.54%	12.50%	8.32%	11.48%	15.11%
None	120	60	13.54%	11.29%	5.46%	11.48%	9.86%
0	60	10	13.54%	12.05%	19.11%	9.99%	16.92%
0	60	20	13.54%	13.43%	20.34%	10.97%	16.94%
0	60	60	13.54%	11.66%	1.24%	7.75%	11.87%
0	120	10	13.54%	11.77%	7.99%	11.05%	12.89%
0	120	20	13.54%	11.96%	–1.71%	11.09%	9.68%
0	120	60	13.54%	9.87%	5.91%	9.63%	9.00%
0.001	60	10	13.54%	12.60%	22.67%	11.55%	15.09%
0.001	60	20	13.54%	14.50%	23.64%	13.03%	15.73%
0.001	60	60	13.54%	11.38%	0.13%	7.55%	10.77%
0.001	120	10	13.54%	13.39%	8.47%	15.07%	13.12%
0.001	120	20	13.54%	12.56%	1.20%	14.12%	8.54%
0.001	120	60	13.54%	10.35%	7.54%	12.49%	5.51%
0.002	60	10	13.54%	12.80%	24.89%	9.26%	8.01%
0.002	60	20	13.54%	14.62%	26.67%	11.83%	13.45%
0.002	60	60	13.54%	6.60%	–0.86%	2.31%	2.66%
0.002	120	10	13.54%	12.74%	–4.07%	11.77%	9.93%
0.002	120	20	13.54%	9.98%	–1.37%	7.02%	6.42%
0.002	120	60	13.54%	9.68%	6.54%	11.90%	5.62%

B.2 FF5 (Lookback = 60) Backtest Net Asset Value (NAV) Figures

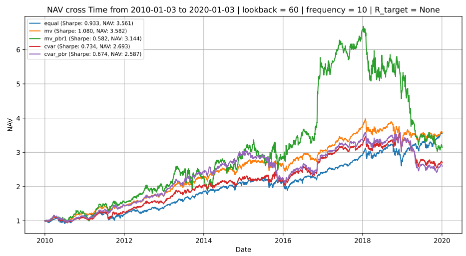

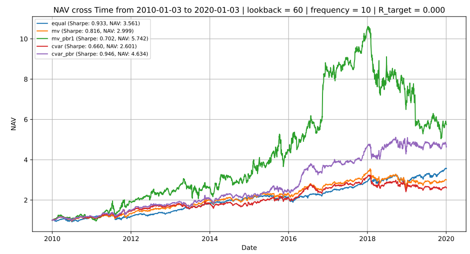

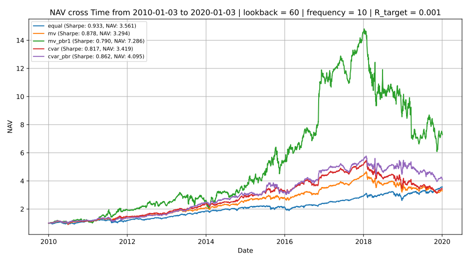

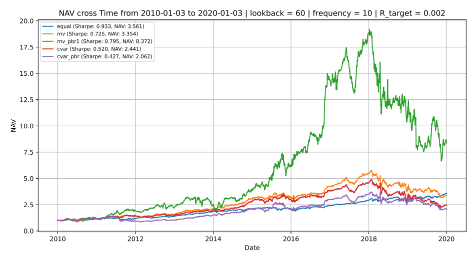

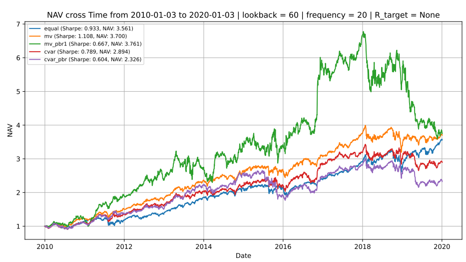

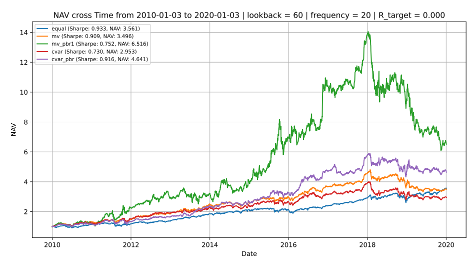

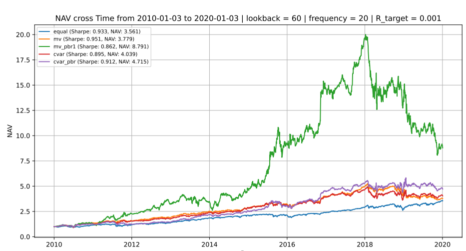

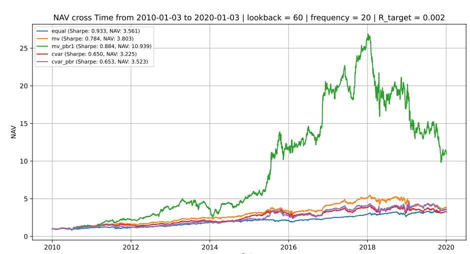

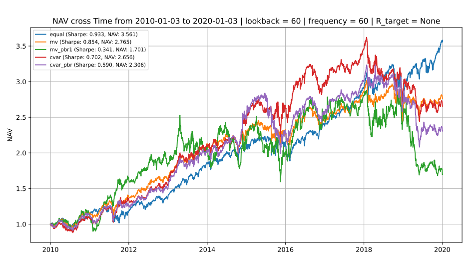

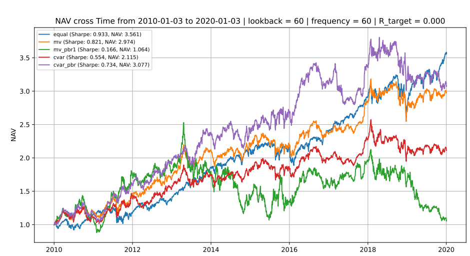

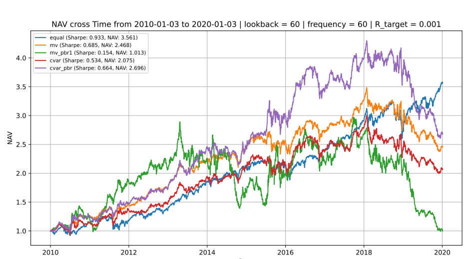

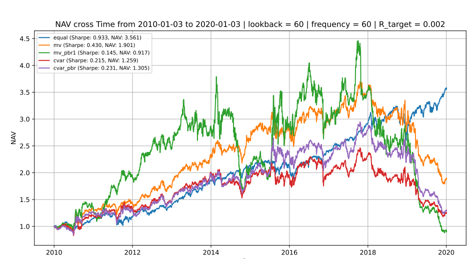

B.3 FF5 (Lookback = 120) Backtest Net Asset Value (NAV) Figures

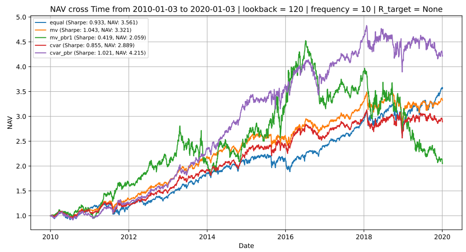

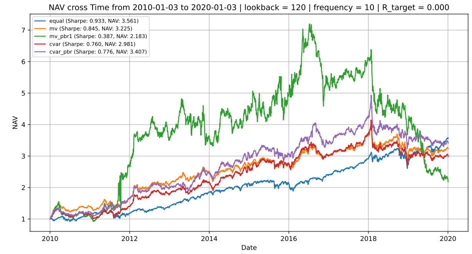

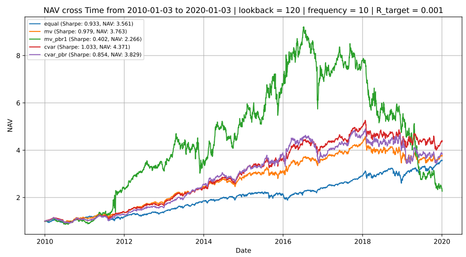

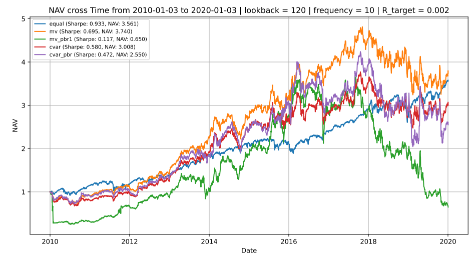

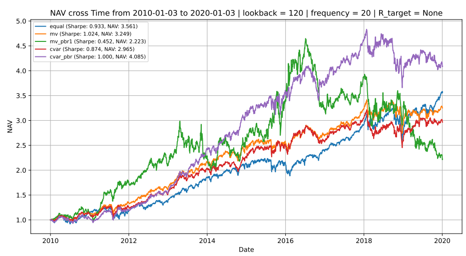

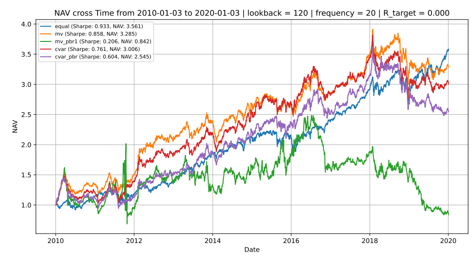

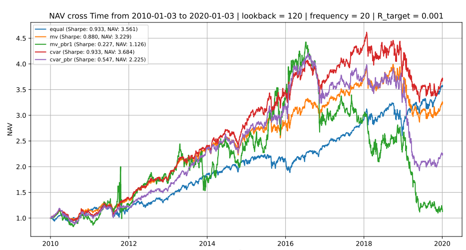

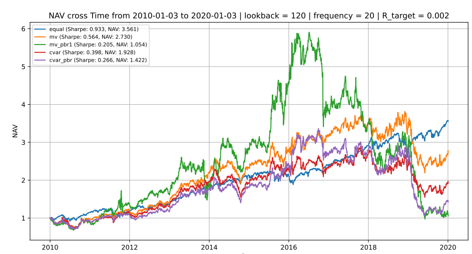

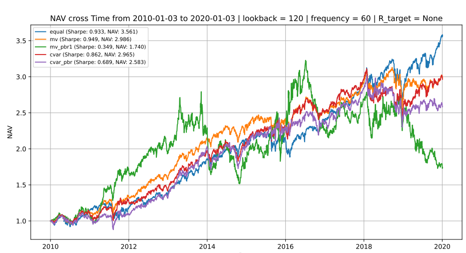

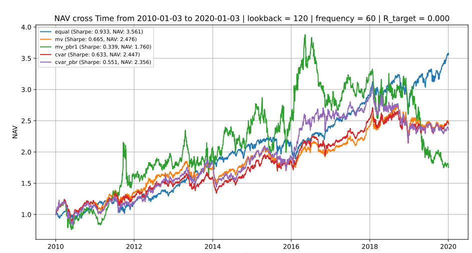

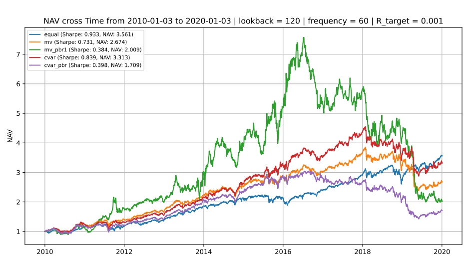

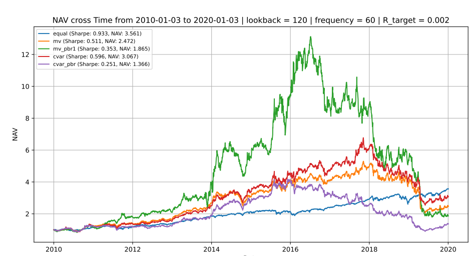

B.4 FF10 (Lookback = 60) Backtest Net Asset Value (NAV) Figures

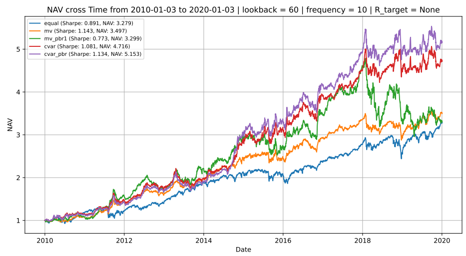

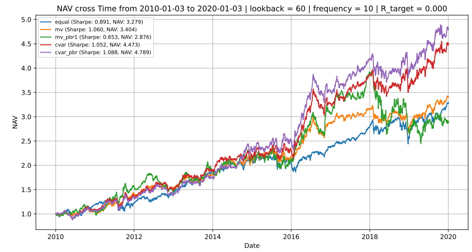

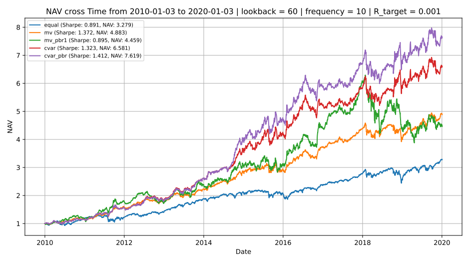

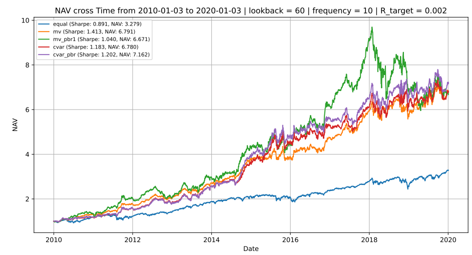

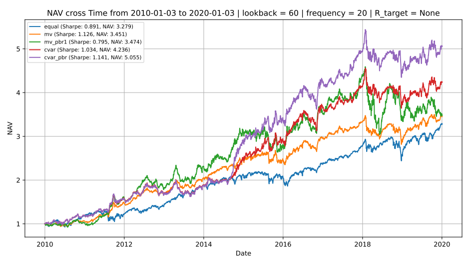

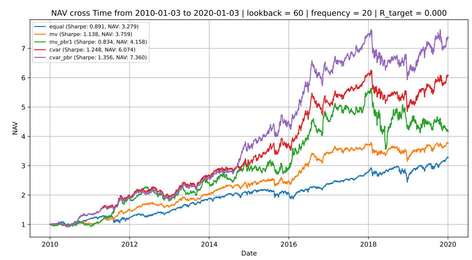

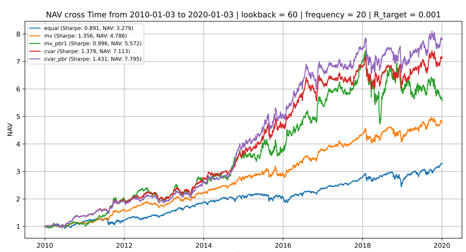

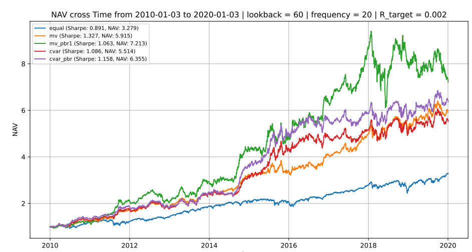

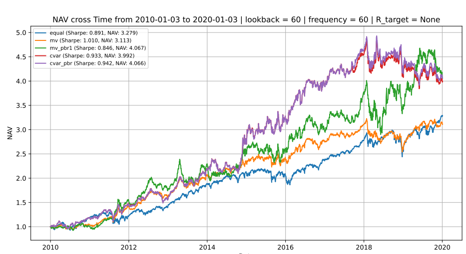

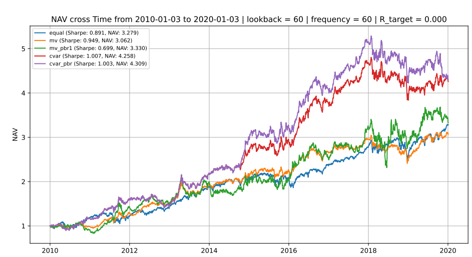

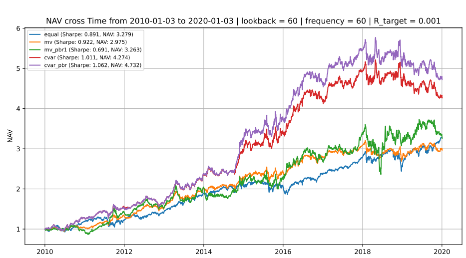

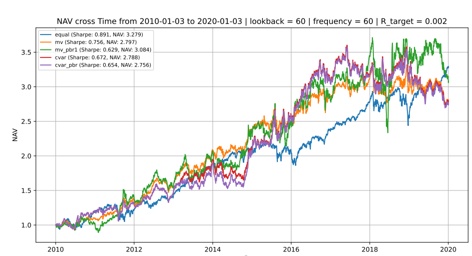

B.5 FF10 (Lookback = 120) Backtest Net Asset Value (NAV) Figures