

Regularizing Mean–Variance Portfolios with an ℓ_1 Total-Exposure Constraint: Evidence From S&P 100 (OEX) Monthly Data

Wentian Wang *

Department of Accountancy & Data Science, Gies College of Business, University of Illinois Urbana–Champaign (UIUC), Champaign, IL 61820, United States.

* Corresponding Author Email: wentian3@illinois.edu

Abstract. This study shows that a total exposure cap can improve the implementability of a mean–variance portfolio. The dataset contains monthly S&P 100 (OEX) returns and a matching nominal risk-free rate from the same workbook. The analysis compares four strategies: an unconstrained portfolio, portfolios with exposure caps of 2.0 and 1.8, and an equal-weighted portfolio. All strategies use the same inputs and the same rebalancing frequency, only the exposure budget changes. The results show that both constrained strategies outperform the unconstrained one in Sharpe ratio and drawdown control. The 1.8 cap performs best in most sub-periods in the later years. For the constrained portfolios, the gross exposure hits the cap each month. For the unconstrained portfolio, the average gross exposure is about 1.4. A normalized HHI based on absolute weights indicates lower concentration under caps, while the relative ranking across strategies remains unchanged. These findings suggest that exposure caps enhance portfolio stability and improve risk-adjusted returns. In practice, a 1.8 cap is a reasonable starting point, with later adjustments guided by edge-hitting behavior and concentration changes.

Keywords: Mean–variance; gross exposure; ℓ_1 constraint; portfolio concentration; OEX.

1. Introduction

Mean-Variance (MV) portfolio selection has always been an important tool for risk allocation. However, it is difficult to implement without constraints. The weight frequently switches between positive and negative, and the leverage drifts over time. If the input data contains noise, the turnover rate will also increase significantly. A simple and economically meaningful approach is to set an upper limit on the ℓ_1 norm of the weights. This can set a budget for the total exposure of "long + short" positions and suppress extreme positions [1]. Related weight-regularization and heuristic approaches for large portfolios include [2,3], and broader robust selection surveys are provided in [4]. In terms of estimating the input, stabilizing the covariance matrix (e.g., sparse precision/inverse-covariance or graphical-model approaches) can also improve results. However, this is logically distinct from weight constraints [5]. On the estimation side, sparse precision and graphical-model methods are surveyed in and instantiated in [6,7].

This article focuses on a specific issue with direct implementation value: In a setting that allows short selling and investable monthly rebalancing, can the Reg-T style ℓ_1 total exposure cap be enhanced in terms of implementability without sacrificing the risk-return ratio? The analysis is entirely based on the OEX stock pool and NRFR sequence within the workbook, comparing four types of strategies calculated in the same file: unconstrained MV ("Free"), with ℓ_1 constraints ($c = 2.0$ and tighter $c = 1.8$), and 1/N benchmark (EQW). All covariance matrices and other settings remain unchanged, only the c value is changed. The differences in the results obtained can be attributed to the weight constraints, rather than the estimation method [6]. The scoring method is based on the arithmetic annualization of monthly excess returns: Annualized return $\text{AnnRet} = 12 \times$ average monthly return, Annualized volatility $\text{AnnVol} = \sqrt{12} \times$ standard deviation, Sharpe ratio $\text{Sharpe} = \text{AnnRet}/\text{AnnVol}$.

Using Excel/Solver with identical μ and Σ across strategies, the analysis finds that imposing an ℓ_1 cap improves risk-adjusted performance and drawdown versus the unconstrained baseline, with $c =$

1.8 strongest in this sample. EQW achieves a high Sharpe due to low volatility yet lags in cumulative NAV. Looking at the HHI after normalization of absolute weights, during the sample period, $\text{Free} \approx 0.248$, $c=2.0 \approx 0.101$, $c=1.8 \approx 0.147$. Due to the geometric theoretical characteristics of ℓ_1 , the total exposure of the restricted portfolio operates close to the edge each month, which makes the leverage path more compact, and makes it easier to audit and operate.

To ensure the transparency and reproducibility of the data results in the paper, the data range is limited to the monthly panel of OEX and the sequence of NRFR. The risk-free interest rate is calculated based on the month-on-month change of NRFR. The portfolio excess returns are aligned by the same monthly timestamp. At the implementation level, Excel/Solver is used, and the constraint $\sum |w| \leq c$ is achieved through the ABS auxiliary column. This method does not require additional parameters and performs stably under the generalized reduced gradient (GRG) nonlinear solver [1]. The baseline does not introduce a transaction cost model, ensuring that all charts can be directly reproduced from the workbook.

The framework of the experimental study and the two research paths are related but deliberately separated: In the weight side, the ℓ_1 method can make the weights more stable, thereby reducing transaction friction and motivating the formation of the total exposure budget; while in the estimation side, the sparse precision (inverse-covariance) and graphical-model methods can improve the input stability in large asset pools. The calculation of the fixed covariance matrix Σ and the expected return μ is to identify the incremental effect of the total exposure limit on the feasibility (Riccobello et al., 2025). These two types of methods are complementary in practice, but logically they can be separated.

The contribution of this paper is practical and empirical. The data in the workbook demonstrates how the ℓ_1 constraint can change the efficient set in an investable manner. Moderate constraints ($c=1.8$) in the sample increased the Sharpe ratio and improved the drawdown, while fixing the leverage path. Looser constraints ($c=2.0$) are also superior to unconstrained MV, but the risk-return ratio is slightly inferior to $c=1.8$, which is also in line with the realistic theory. For large universes, ML-based heuristics and signal-combination schemes provide complementary allocation tools [8].

2. Data and Preprocessing

This study only uses the OEX monthly panel and NRFR sequence in the workbook. The sample period is from June 2009 to June 2024. Timestamps and components are exactly as described in the file; no external data is introduced, and the component changes are not altered. All four strategies are formed and scored on the same monthly grid: unconstrained MV (Free), two groups of MV with ℓ_1 total exposure limits ($c=2.0$ and $c=1.8$), and equal-weight benchmark (EQW) for the same component set. Risk-free returns come from the month-on-month changes in NRFR. Monthly portfolio return is computed as $(r_t = \text{NAV}_t / \text{NAV}_{t-1} - 1)$; excess return equals portfolio returns minus the NRFR change. All indicators are arithmetic annualized based on monthly excess returns: $\text{AnnRet} = 12 \times \text{mean}$, $\text{AnnVol} = \sqrt{12} \times \text{standard deviation}$, $\text{Sharpe} = \text{AnnRet} / \text{AnnVol}$. The maximum drawdown comes from the peak of the net asset value and the subsequent maximum decline. To ensure reproducibility of the results, the sequence is not filled in forward and backward, and the dates are aligned before calculating the returns. If $\text{NAV} = 0$, that month will be excluded. Outliers are not trimmed; robustness is demonstrated through horizontal comparisons of the budget and sub-samples.

The weights are directly taken from Excel/Solver. The net budget $\mathbf{1}^T \mathbf{w} = 1$ is for all dates. The total exposure of Free is $\sum |w|$ which is endogenous and is achieved through ABS column splitting of the restricted portfolio $\sum |w| \leq c$. Let $w_i = w_i^+ - w_i^-$, and $w_i^+, w_i^- > 0$. Calculate the absolute values in adjacent columns and aggregate at the top. GRG Nonlinear directly solves in the Sharpe cell without binary variables or additional parameters. Because the sample-Sharpe objective is non-convex, stability was verified via multiple random restarts (same solution within rounding) and by solving the equivalent target-return minimum-variance quadratic program (QP); both delivered identical weights under the same constraints; see also for related formulations in sparse minimum-variance settings. Related estimation-side regularization includes sparse-group lasso for portfolio

construction and sector selection [9]. The objective function is the tangent formula, directly maximizing the sample Sharpe. This is equivalent to tracing the efficient frontier via a target-return minimum-variance program; both yield the tangency (maximum-Sharpe) portfolio under identical constraints [6]. The covariance matrix Σ is consistent with the window, preprocessing, and possible contraction method of the mean μ , only changing the budget c . This can attribute the differences to the constraints rather than the estimation method.

ℓ_1 budget has a clear geometric meaning. The feasible region $\{w: \sum |w| \leq c\}$ is a convex set. When the unconstrained optimal point is outside the domain, the optimal solution will touch the edge. This edge-hitting behavior follows from convex optimization geometry. Therefore, the total exposure of the restricted portfolio is numerically equal to the upper limit. The monthly results in the document precisely confirm this: for $c = 2.0$, $\sum |w|$ is equal to 2.000 every month, with an average long/short ratio of approximately 1.5/0.5; for $c = 1.8$, $\sum |w|$ is equal to 1.800 every month, with an average long/short ratio of approximately 1.4/0.4; for Free, the average $\sum |w|$ is 1.463. This leads to three tests: the total exposure of the restricted portfolio should be long-term close to the edge; the NAV paths for different c should be significantly different; Under the concentration metric, the restricted portfolio should be lower than free. Concentration is calculated using normalized HHI: $HHI_t = \sum_i \left(\frac{|w_{i,t}|}{\sum_j |w_{j,t}|} \right)^2$. The results show that the average concentration for $c = 2.0$ is the lowest, $c = 1.8$ is in the middle, and Free is the highest. The "HHI = 1" in the workbook is only a definition artifact, not "fully invested in a single stock" [10].

The empirical procedure is clear: Solver outputs the current weight w_t each month, uses the next month's returns to form excess returns, and compounds them monthly to form NAV. The four strategies repeat this process in both the full sample and the sub-sample. Performance is scored by AnnRet, AnnVol, Sharpe, and MaxDD; implementability is explained by "whether the total exposure is close to the boundary" and "the level and variability of the normalized HHI". The charts come from the workbook: Cumulative NAV shows $c = 1.8 \approx 35.66$, $c = 2.0 \approx 24.89$, Free ≈ 6.20 , EQW ≈ 7.33 . Based on the average annualized excess performance: Free: 0.0517/0.2923/0.1769/-0.1979; ($c = 2.0$) = 0.1482/0.3043/0.4872/-0.1494; ($c = 1.8$): 0.1721/0.3070/0.5608/-0.1122; EQW = 0.1228/0.1205/1.0189/-0.1926. The results show that under the same Σ and μ , a single parameter c can systematically change the corner solution. Moderate budget ($c = 1.8$ or 2.0) can improve the risk-return ratio and maintain an auditable leverage path, while being too loose or too tight will reduce utility [2].

Notation and conventions. μ and Σ are fixed over the entire sample to isolate the effect of the weight-side budget c . Net budget is $(1^T w = 1)$. Gross exposure is $(G_t = \sum_i |w_{i,t}|)$. Annualization uses arithmetic monthly excess returns: (AnnRet = $12\bar{r}$, AnnVol = $\sqrt{12} \sigma(r)$, Sharpe = AnnRet/AnnVol.)

3. Experimental Results

This section only uses the time series and summary table of the workbook. Excel-based experiment reports the results under a unified scoring standard and date format. What to evaluate is the performance after adding the ℓ_1 total exposure constraint.

3.1. Cumulative Net Asset Value Path

Figure 1 (the summary page "Excess Performance") shows the cumulative net value of the four strategies over the common monthly grid (2009 - 2024). The end-of-period net value of the purple line ℓ_1 ($c = 1.8$) is the highest (approximately 35.66). Green line ℓ_1 ($c = 2.0$) is the second highest (approximately 24.89). The blue line EQW stays close to the orange line Free for a long time but remains lower than the two restricted combinations at the end (approximately 7.33 vs 6.20). The gap widens significantly in the latter part of the sample. After 2018, the relative advantage of the restricted combinations becomes clearer. This phenomenon supports an intuitive conclusion: when volatility

and dispersion increase, weight contraction can stabilize the spread between long and short positions and improve the compound interest path as displayed on the workbook's "Excess Performance" sheet.




Fig. 1 Cumulative Net Asset Value Paths of Portfolio Strategies (2009–2024)

3.2. Summary of Excess Profits

Table 1 (Sheet 4, $N \approx 183$ months) reports AnnRet, AnnVol, Sharpe, and MaxDD (NAV) under the "average excess annualized" metric: Free: 0.0517 / 0.2923 / 0.1769 / -0.1979, ℓ_1 ($c=2.0$): 0.1482 / 0.3043 / 0.4872 / -0.1494, EQW: 0.1228 / 0.1205 / 1.0189 / -0.1926, ℓ_1 ($c=1.8$): 0.1721 / 0.3070 / 0.5608 / -0.112, The Sharpe ratios of the two ℓ_1 .

constrained combinations are significantly higher than that of Free. ℓ_1 ($c=1.8$) is the strongest in this sample. It has the highest annualized return and Sharpe ratio, and the shallowest maximum drawdown. The volatility of EQW is very low, so its annualized Sharpe ratio is very high. However, its long-term compound return still lags the two constrained combinations. This is consistent with the net value ranking shown in Figure 1.

Table 1. Annualized Performance Metrics of Four Portfolio Strategies

Strategy	AnnRet_ex	AnnVol_ex	Sharpe	MaxDD_NAV
Free	0.0517	0.2923	0.1769	-0.1979
ℓ_1 ($c=2.0$)	0.1482	0.3043	0.4872	-0.1494
EQW	0.1228	0.1205	1.0189	-0.1926
ℓ_1 ($c=1.8$)	0.1722	0.3070	0.5608	-0.1122

3.3. Overall Exposure Margin and Implementability

The time series of the aggregated exposure directly validates "edge effect". For ℓ_1 ($c=2.0$), $\sum |w|$ equals 2.000 every month. For ℓ_1 ($c=1.8$), $\sum |w|$ equals 1.800 every month. The average of Free's $\sum |w|$ is 1.463. The corresponding average long position is approximately 1.2315, and the average short position is approximately 0.2315 (all from the workbook summary, with $L_t = \frac{G_t+1}{2}$ and $S_t = \frac{G_t-1}{2}$ implied by $1^T w = 1$ and $G_t = L_t + S_t$). Therefore, the constrained solution always operates on the ℓ_1 sphere; Free naturally drifts outside the constraints. This result, without changing the estimation module, provides clear and auditable boundaries for leverage and short selling.

3.4. Concentration Degree (Herfindahl)

The paper measures the concentration using the HHI metric that is normalized by absolute weights:

$$HHI_t = \sum_i \left(\frac{|w_{it}|}{\sum_j |w_{jt}|} \right)^2 \in [0,1] \quad (1)$$

At this scale, the sample average shows: Free ≈ 0.248 ; ℓ_1 ($c = 1.8$) ≈ 0.147 ; ℓ_1 ($c = 2.0$) is the lowest (≈ 0.101). The "unnormalized HHI ($\sum_i w_i^2$)" in the workbook is almost a constant under $c = 1.8$. This is defining a pseudo-effect, rather than "single stock full position". After normalization, the relative order of the three is consistent with the explanation of "restricted portfolio reducing concentration".

3.5. Robustness of Sub-Samples

The paper divided the samples into two periods: 2009-2016 and 2017-2024. The ranking of net values and the "edge-hitting" diagnosis remained consistent in both periods. The two restricted combinations consistently outperformed Free. The advantage was more pronounced in the latter period. This result is consistent with the role of weight regularization in suppressing unstable positions during high-discrepancy or high-noise phases. Because Σ and μ remain constant across the full sample and the sub-samples, the differences can be attributed to the change in a single knob, c , rather than the estimated input.

4. Robustness and Implementation Suggestions

The objective of this section is to test the robustness of the conclusion. The testing dimensions include the sampling window, the total exposure budget, and the concentration measure. The OEX investable set, the mean μ within the workbook, the covariance Σ , and the date grid remain unchanged. The statistical caliber is standardized as the average volume over the sample period. The indicators include the average annualized return, average annualized volatility, Sharpe, average total exposure, and average concentration based on the monthly sequence. This approach can avoid reliance on single-month readings.

4.1. Subsample and Stress Conditions

The samples were divided into two periods: 2009-2016 and 2017-2024. The ranking of net values and the average indicators were consistent in both periods. The overall performance of the two ℓ_1 constrained strategies were superior to the unconstrained benchmark. The average Sharpe and final net value of ℓ_1 ($c = 1.8$) were at the top. ℓ_1 ($c = 2.0$) followed. EQW lagged the two constrained strategies. Free performed the weakest. The stress episodes (e.g., the pandemic and the 2022 drawdown) did not change the above average ranking. The advantage in the latter period was more obvious. This result is consistent with the intuitive geometric meaning that the fixed total exposure budget suppresses unstable long-short weights and accelerates the recovery of drawdown in the high dispersion stage.

4.2. Edge-hitting Behavior and Budget Gradient

The total exposure path shows that the constraints have a substantive effect. For ℓ_1 ($c = 2.0$), $\sum |w|$ is equal to 2.000 month by month. For ℓ_1 ($c = 1.8$), $\sum |w|$ is equal to 1.800 month by month. The average sample $\sum |w|$ of Free is approximately 1.463, corresponding to an average Long/Short of approximately 1.2315/0.2315. The two restricted solutions always lie on the respective ℓ_1 ball boundaries. This property transparently limits leverage and short selling without changing the estimation module. The budget sensitivity shows a clear gradient. In this panel, relaxing the budget from 1.8 to 2.0 does not compromise feasibility, but the average risk return slightly weakens. Using the "average annualized" metric, ℓ_1 ($c = 1.8$) has a higher Sharpe ratio and a shallower maximum drawdown (approximately -0.112) compared to ℓ_1 ($c = 2.0$) (approximately -0.149) and Free (approximately -0.198). This result suggests that moderate tightening can weaken extreme hedging and improve risk calibration.

4.3. Selection and Interpretation of Concentration Measures

The concentration ratio is recommended to be calculated using the Herfindahl index normalized by absolute weights:

$$HHI_t = \sum_i \left(\frac{|w_{i,t}|}{\sum_j |w_{j,t}|} \right)^2 \in [0,1] \quad (2)$$

This metric can directly control the difference in the total volume of long and short positions and is suitable for comparing multi-long-short portfolios where $\sum |w| \neq 1$. The sample average shows that the normalized HHI is significantly lower than Free at both ℓ_1 ($c = 1.8$) and ℓ_1 ($c = 2.0$). The normalized HHI at ℓ_1 ($c = 2.0$) is even lower, indicating a higher degree of dispersion. This result is consistent with the intuition that "a wider budget allows for more weight dispersion across more targets". The workbook also provides the un-normalized $\sum_i w_i^2$. The ranking under this metric is typically $\text{Free} > c = 2.0 > c = 1.8$, with the difference arising from the influence of the total exposure size. The graphical suggestion is to use the time series or the box plot of the average value on a regular basis (in Excel:

`=SUMSQ(ABS(w_range))/POWER(SUM(ABS(w_range)),2)`). The un-normalized version can be placed in the appendix for verification.

4.4. The Purity of Attribution and Cost Extrapolation

The mean μ and covariance Σ remain consistent across strategies and sub-samples. Therefore, the average performance difference can be attributed to the weight budget c rather than changes in the estimation module. The drawdown statistics obtained from the NAV trajectories of each strategy provide the same information. The ℓ_1 budget reduces extreme long-short bets, resulting in a smaller average peak-to-trough decline and a faster recovery. The baseline does not incorporate explicit transaction costs to ensure one-click reproducibility. Ex post-testing can record the weight vector w_t , total exposure, and concentration for each rebalancing, and calculate the turnover rate based on.

$$TO_t = 0.5 \sum_i |w_{t,i} - w_{t-1,i}| \quad (3)$$

Then, replay the cost (κ as the unit transaction cost, in basis points) using.

$$r_t^* = r_t^P - \kappa * TO_t \quad (4)$$

and recalculate AnnRet/AnnVol/Sharpe/MaxDD using the "average annualized" metric. Dynamic allocation frameworks with asset-specific regimes provide a complementary route for implementability assessments [11]. The existing evidence and workbook behavior both show that the ℓ_1 budget usually reduces turnover, so a moderate budget still has advantages after considering costs.

4.5. Operational Suggestions

On this OEX panel, ℓ_1 ($c = 1.8$) can be set as the default budget. The total exposure of the capped portfolios equals the cap each month (edge-hitting), and its Sharpe ratio and drawdown are superior to ℓ_1 ($c = 2.0$) and Free. A preference for a wider leverage limit can be achieved by using ℓ_1 ($c = 2.0$). The operation process only requires adjusting the budget c , while the remaining settings (estimation window, μ , Σ , date) remain unchanged. After completion, two diagnostics should be conducted: one is whether $\sum |w|$ is close to the cap; the other is whether the normalized HHI is consistently lower than Free. If the mean or covariance engine is changed, the comparison between Free and the restricted solution should be re-established under the new engine before discussing the budget adjustment. This process retains the traceability of evidence and implementation and is consistent with the division of labor in the literature regarding "weight regularization improves realizability and estimation regularization stabilizes input". Our monthly setting abstracts from dynamic covariance and weight-predictability effects; recent evidence on large dynamic covariance (HAR) models and realized-weight dynamics can complement the present approach [12,13].

5. Conclusion

The total exposure budget is an economic constraint on leverage and short selling, rather than a statistical patch to the mean–variance framework. Based on monthly data of the S&P 100 (OEX) from June 2009 to June 2024 (approximately 183 periods), and under identical input and rebalancing rules, four strategies were evaluated and compared: an unconstrained portfolio, two restricted portfolios with total exposure limits of 1.8 and 2.0, and an equal-weighted portfolio. The restricted portfolios were generally more stable. The gross exposure of the restricted portfolios reached the cap every month, while the average gross exposure of the unconstrained portfolio was approximately 1.4. The 1.8 cap performed best in terms of Sharpe ratio and drawdown control, and its advantage became more pronounced after 2018. The 2.0 cap also improved performance, but to a lesser extent. The unconstrained portfolio showed the weakest performance throughout the sample period.

Concentration results support these findings. The normalized Herfindahl index based on absolute weights shows that the restricted portfolios had lower concentration than the unconstrained portfolio. A 2.0 cap typically led to a more dispersed portfolio, while a 1.8 cap struck a balance between stability and efficiency. Importantly, under this concentration metric, the relative ranking of the strategies remained unchanged.

The evidence indicates that a single transparent parameter can regulate the leverage path while maintaining risk–return characteristics. With fixed expected returns, covariances, and rebalancing frequency, adjusting only the exposure budget can alter portfolio behavior in an auditable way. A cap of 1.8 is a practical starting point. After implementation, two aspects should be checked: whether the exposure consistently reaches the cap, and whether concentration is lower than that of the unconstrained benchmark. If the unconstrained portfolio's average exposure is much lower than the candidate cap, or if simple transaction-cost adjustments alter the strategy rankings, the budget should be re-evaluated. Overall, the exposure cap enhances implementability, reduces concentration, and maintains a stable risk–return trade-off, providing a clear tool for portfolio design.

References

- [1] Dhingra, V., Gupta, S. K., & Sharma, A. (2023). Norm constrained minimum variance portfolios with short selling. *Computational Management Science*, 20, 6. <https://doi.org/10.1007/s10287-023-00438-2>
- [2] Kremer, P. J., Lee, S., Bogdan, M., & Paterlini, S. (2020). Sparse portfolio selection via the sorted ℓ_1 -norm. *Journal of Banking & Finance*, 110, 105687. <https://doi.org/10.1016/j.jbankfin.2019.105687>
- [3] Costola, M., Maillet, B., Yuan, Z., & Zhang, X. (2024). Mean–variance efficient large portfolios: A simple machine learning heuristic technique based on the two-fund separation theorem. *Annals of Operations Research*, 334, 133–155. <https://doi.org/10.1007/s10479-022-04881-3>
- [4] Ghahtarani, A., Saif, A., & Ghasemi, A. (2022). Robust portfolio selection problems: A comprehensive review. *Operational Research*, 22, 3203–3264. <https://doi.org/10.1007/s12351-022-00690-5>
- [5] Riccobello, R., Bonaccolto, G., Kremer, P. J., Sobczyk, P., Bogdan, M., & Paterlini, S. (2025). Sparse graphical modelling for global minimum variance portfolio. *Computational Management Science*, 22, 8. <https://doi.org/10.1007/s10287-025-00535-4>
- [6] Torri, G., Giacometti, R., & Paterlini, S. (2019). Sparse precision matrices for minimum variance portfolios. *Computational Management Science*, 16(3), 375–400. <https://doi.org/10.1007/s10287-019-00344-6>
- [7] Torri, G., Giacometti, R., & Paterlini, S. (2024). Penalized enhanced portfolio replication with asymmetric deviation measures. *Annals of Operations Research*, 332, 481–531. <https://doi.org/10.1007/s10479-023-05576-z>
- [8] Firoozye, N., Tan, V., & Zohren, S. (2023). Canonical portfolios—Optimal asset and signal combination. *Journal of Banking & Finance*, 148, 106851. <https://doi.org/10.1016/j.jbankfin.2023.106851>
- [9] Chen, J., Dai, G., & Zhang, N. (2020). An application of sparse-group lasso regularization to equity portfolio optimization and sector selection. *Annals of Operations Research*, 284, 243–262. <https://doi.org/10.1007/s10479-019-03429-2>

- [10] Lassance, N., & Vrms, F. (2021). Portfolio selection with parsimonious higher comoments estimation. *Journal of Banking & Finance*, 126, 106115. <https://doi.org/10.1016/j.jbankfin.2021.106115>
- [11] Shu, Y., Yu, C., & Mulvey, J. M. (2025). Dynamic asset allocation with asset-specific regime forecasts. *Annals of Operations Research*, 346, 285–318. <https://doi.org/10.1007/s10479-024-06266-0>
- [12] Honig, I., & Kircher, F. (2025). Large dynamic covariance matrices and portfolio selection with a heterogeneous autoregressive model. *Journal of Banking & Finance*, 178, 107505. <https://doi.org/10.1016/j.jbankfin.2025.107505>
- [13] Golosnoy, V., & Gribisch, B. (2022). Modeling and forecasting realized portfolio weights. *Journal of Banking & Finance*, 138, 106404. <https://doi.org/10.1016/j.jbankfin.2022.106404>