

EFZG WORKING PAPER SERIES

EFZG SERIJA ČLANAKA U NASTAJANJU

ISSN 1849-6857

UDC 33:65

No. 2026-05

Silvija Vlah Jerić
Margareta Gardijan Kedžo
Zrinka Lovretin Golubić

Ridge Regularization or Weight Constraints? Portfolio Optimization in a Small Frontier Market

Ridge Regularization or Weight Constraints? Portfolio Optimization in a Small Frontier Market

Silvija Vlah Jerić

Faculty of Economics & Business, University of Zagreb

svlah@net.efzg.hr

Margareta Gardijan Kedžo

Faculty of Economics & Business, University of Zagreb

mgardijan@net.efzg.hr

Zrinka Lovretin Golubić

Faculty of Economics & Business, University of Zagreb

zlovretin@net.efzg.hr

The views expressed in this working paper are those of the author(s) and not necessarily represent those of the Faculty of Economics and Business – Zagreb. The paper has not undergone formal review or approval. The paper is published to bring forth comments on research in progress before it appears in final form in an academic journal or elsewhere.

Copyright September 2026 by Silvija Vlah Jerić, Margareta Gardijan Kedžo & Zrinka Lovretin Golubić. All rights reserved.

Sections of text may be quoted provided that full credit is given to the source.

Abstract

This paper examines whether ridge regularization improves portfolio performance in the Croatian equity market and whether it adds value when individual portfolio weights are already capped. We compare long-only global minimum variance (GMV) and minimum conditional value-at-risk (CVaR) portfolios using CROBEX constituent returns from 2004 onward, covering 43 index revisions and six estimation and holding-window schemes. The regularization parameter is selected adaptively using information from previous holding periods. Performance is assessed against corresponding unregularized portfolios, an equally weighted portfolio, and the CROBEX index. We also examine turnover, transaction costs, and crisis performance. Under the baseline 30% maximum-weight constraint, adaptive ridge regularization does not produce a statistically significant improvement in Sharpe ratios over the corresponding unregularized portfolios. For CVaR portfolios, the weight constraint accounts for a substantial part of the observed Sharpe-ratio improvement, while the additional benefit of ridge regularization is limited and varies across specifications. Both mechanisms have only small effects on GMV performance. Although selected CVaR specifications outperform CROBEX in cumulative returns, GMV generally provides better realized downside protection. These results suggest that ridge regularization and maximum-weight constraints are partly substitutable in this small frontier market.

Key words

portfolio optimization; ridge regularization; maximum-weight constraints; conditional value-at-risk; global minimum variance; frontier markets; CROBEX.

JEL classification

G11; C61; C58.

1. Introduction

Because the theoretical market portfolio is unobservable, capitalization-weighted equity indices have become the standard proxy for the market portfolio in both empirical research and investment practice. Their widespread use is justified by the fact that they follow the same market-value weighting principle as the theoretical CAPM market portfolio while remaining readily observable, easily tradable, and inexpensive to replicate. However, empirical studies dating back to the early 1990s demonstrated that capitalization-weighted equity indices are not mean-variance efficient (Grinold, 1992; Haugen & Baker, 1991). This creates a fundamental gap between financial theory and investment practice, providing strong motivation for developing alternative portfolio construction methods (Amenc et al., 2013).

By construction, the Global Minimum Variance (GMV) and Maximum Sharpe Ratio (MSR) portfolios lie on the estimated efficient frontier, whereas capitalization-weighted indices need not do so. Empirical evidence consistently shows that capitalization-weighted indices fall below the efficient frontier of their constituent securities and, in many cases, even equal-weighted portfolios are located closer to the efficient frontier (DeMiguel et al., 2009b). Nevertheless, despite this apparent theoretical disadvantage, empirical studies based on CROBEX index constituents, the main stock market index of the Croatian capital market, have shown that the CROBEX index is remarkably difficult to outperform (Dolinar et al., 2017; Zoričić et al., 2018; Lovretin Golubić, 2025). This apparent contradiction reflects the characteristics of small frontier markets rather than the superiority of capitalization weighting. The Croatian equity market consists of a relatively small number of listed stocks, limiting diversification opportunities and reducing the statistical reliability of estimated portfolio parameters. As a result, the sample covariance matrix is often ill-conditioned, making mean-variance optimization highly sensitive to estimation error. Low market liquidity further amplifies parameter uncertainty, frequently producing unstable portfolio weights and disappointing out-of-sample performance. These challenges provide strong motivation for applying regularization techniques, such as ridge regularization, to improve the robustness of portfolio optimization (Ledoit & Wolf, 2004; DeMiguel et al., 2009a).

The purpose of this research is to examine whether ℓ_2 ridge regularization improves the out-of-sample performance of GMV and minimum-CVaR portfolios in a small frontier market and, importantly, whether it provides additional benefits once portfolio concentration is already controlled through maximum-weight (ℓ_∞) constraints. The empirical analysis begins in 2004 and relies on stock returns of all CROBEX index constituents, including a total of 43 CROBEX revisions. The combination of a limited cross-section, relatively short estimation windows, correlated asset returns, and noisy market data can make portfolio inputs unstable in a small frontier market. Ridge regularization addresses this instability by penalizing the Euclidean norm of the portfolio-weight vector, thereby discouraging extreme or highly concentrated allocations and reducing sensitivity to estimation noise.

This distinction is practically relevant because the two mechanisms control concentration in different ways: the ℓ_2 penalty discourages large weights continuously, whereas a maximum-weight constraint directly limits the largest individual position. The empirical analysis therefore evaluates not only whether ridge regularization improves portfolio performance in isolation, but also whether its benefits remain when concentration is already restricted.

2. Portfolio models and evaluation framework

Ridge regularization is introduced directly into the portfolio optimization objective through an ℓ_2 penalty on portfolio weights (DeMiguel et al., 2009a). For the GMV strategy, the regularized problem is:

$$\min w^T \hat{\Sigma} w + \lambda \|w\|_2^2 \quad \text{subject to} \quad 1^T w = 1, w \geq 0$$

For the minimum CVaR strategy (Rockafellar & Uryasev, 2000), ridge regularization is incorporated analogously by adding the same ℓ_2 penalty directly to the risk objective, following the regularized Expected Shortfall framework of Still and Kondor (2010):

$$\min \text{CVaR}\alpha(w) + \lambda \|w\|_2^2 \quad \text{subject to} \quad 1^T w = 1, w \geq 0$$


Here, Σ denotes the covariance matrix estimated from the estimation window, while $\text{CVaR}\alpha(w)$ denotes the sample CVaR objective evaluated from portfolio losses in the same window. The parameter $\lambda \geq 0$ controls the strength of regularization. In the GMV problem, the ℓ_2 penalty is algebraically equivalent to replacing Σ by $\Sigma + \lambda I$ in the quadratic objective, but the penalty formulation makes the regularization mechanism explicit and provides a formulation that extends naturally to the CVaR objective. When $\lambda = 0$, the corresponding unregularized GMV or CVaR problem is recovered. As λ increases, concentrated weight vectors are penalized more strongly; under the long-only full-investment constraint, the solution is increasingly pushed toward the equally weighted portfolio. The predefined λ grid ranged from 0 (the unregularized case) to 1000 (very strong regularization), with denser coverage at low and moderate penalty values. Because portfolio variance and CVaR are measured on different numerical scales, the same numerical value of λ does not represent the same effective degree of regularization in the two optimization problems. Consequently, λ values should be interpreted within each strategy rather than compared directly between GMV and CVaR; the common grid is used as a practical exploratory grid, and λ is selected separately within each strategy.

Six rolling-window schemes were used for different estimation and holding periods: T60_H20, T60_H60, T90_H90, T180_H90, REV (revision-aligned), and T756_HREV. For each window, the GMV strategy used the covariance matrix estimated over the estimation period, whereas the CVaR strategy used the empirical loss distribution from the same period; the corresponding ridge-regularized optimization problem was then solved over the λ grid to obtain portfolio weights. The selected portfolio was then held over the corresponding holding period, with performance evaluated exclusively out of sample. For the adaptive strategy, λ was selected using out-of-sample performance in previously completed holding windows, as described below. Performance was compared against three benchmarks: EW portfolio, the unregularized portfolio ($\lambda=0$), and the capitalization-weighted CROBEX index.

3. Empirical analysis

We first examine the out-of-sample performance of GMV and CVaR portfolios across fixed values of the regularization parameter λ . Figure 1 presents the corresponding mean Sharpe ratios for the different estimation and holding-window schemes during crisis and tranquil periods.

Figure 1. Out-of-sample mean Sharpe ratio and different lambda grids for GMV and CvaR strategy and all T/H labels in crisis and tranquil period



Source: Authors' calculation


Figure 1 shows a hump-shaped Sharpe-ratio pattern for the CVaR strategy, particularly during crisis periods: moderate regularization can improve performance, whereas stronger penalization eventually reduces it. For GMV, the Sharpe ratio generally declines as λ increases, and the differences are not statistically significant. Because λ is not directly comparable across the GMV and CVaR objectives, these patterns should be interpreted within each strategy rather than as evidence that a given numerical λ has the same strength across both models.

The fixed- λ analysis illustrates the effect of regularization but does not provide a rule for selecting λ in advance. We therefore use an adaptive procedure that selects the regularization parameter using only previously available information. For each holding window t , the optimal λ was determined exclusively from historical information, using the previous L holding windows ($v < t$) and never information from the current evaluation window. Within each past window, the best-performing λ from the predefined grid was identified according to either the modified Sharpe ratio (Israelson, 2005) or the return/CVaR ratio. These historical winners were then aggregated using either the mode or the median to obtain the λ applied in period t . The lookback horizon L was set to 3, 5, 10, and 12 windows, as well as an expanding

window including all available historical observations. In addition, to assess implementation robustness, portfolio weight constraints (maximum individual asset weights of 10%, 20%, 30%, 50%, and 100%) were examined. In the long-only setting, these maximum-weight constraints are equivalent to an ℓ_∞ -norm constraint, $\|w\|_\infty \leq c$, where c denotes the maximum permitted individual weight, and therefore provide an additional form of concentration control alongside the ℓ_2 ridge penalty. The use of portfolio-weight constraints as a means of limiting extreme allocations and improving estimation robustness is well established in the portfolio optimization literature (Jagannathan & Ma, 2003; Behr et al., 2013). The 30% maximum-weight constraint is used as the baseline specification, while the alternative caps are treated as robustness checks; a 100% cap is effectively the no-additional-cap benchmark in the long-only full-investment setting. Proportional transaction costs, modelled as bid-ask spreads of 10, 20, 30, 50, 100, and 200 basis points, were also examined.

Before focusing on the baseline specification, we first examine the cumulative returns of the adaptive GMV and CVaR strategies across different lookback horizons and estimation/holding label combinations.

Figure 2. Cumulative return of adaptive lambda¹ per strategy and T/H label across different lookback horizons




Source: Authors' calculation

As depicted in Figure 2, GMV cumulative returns are negative or close to zero in most specifications, although a positive return is observed for T90_H90 with $L=3$. CVaR cumulative returns remain positive under the T90_H90 and T180_H90 specifications across all lookback horizons examined. While the extent of the outperformance differs across lookback periods, cumulative returns remain positive throughout and often surpass the CROBEX benchmark. However, performance alone does not determine whether a strategy is practically attractive. We therefore examine portfolio turnover and transaction costs.

Figure 3. Turnover² of adaptive lambda per strategy and T/H label across different lookback horizons

¹ Mode aggregation rule for adaptive λ selection.

² Turnover is calculated using formula $12j=1Nw_{j,t}-w_{j,t-1}$ where $w_{j,t}$ is the portfolio weight of asset j after rebalancing at time t and N is the number of assets in the portfolio.




Source: Authors' calculation

Although the CVaR strategy generates substantially higher cumulative returns than the GMV strategy (as shown in Figure 2), this advantage is accompanied by considerably higher portfolio turnover as depicted in Figure 3. Therefore, any comparison of the two strategies should account for transaction costs before concluding that CVaR is the superior approach.

Performance varies with the lookback horizon, although CVaR cumulative returns remain positive under the T90_H90 and T180_H90 specifications across all horizons examined. The baseline specification uses a five-window lookback ($L=5$), mode aggregation, the modified Sharpe ratio as the selection criterion, and a 30% maximum asset weight.

Figure 4. Walk-forward (adaptive) lambda per strategy and T/H label




Source: Authors' calculation

To better understand the behavior of the adaptive procedure, Figure 4 presents the movement of the selected regularization parameter λ over time. The GMV strategy generally selects lower values of λ , with occasional spikes toward very strong regularization (and therefore near-EW allocations) during the Global Financial Crisis (GFC), whereas the CVaR strategy switches λ more frequently across the full sample, not only during crisis periods.

Comparing the adaptive GMV and CVaR strategies with the EW portfolio and the CROBEX index reveals notable differences in their relative performance.

Figure 5. Adaptive lambda annualized Sharpe ratio – GMV and CVaR compared to CROBEX and EW




Source: Authors' calculation

Figure 5 shows that the adaptive GMV strategy outperforms the EW portfolio across all window specifications but fails to outperform the CROBEX index. The adaptive CVaR strategy also consistently outperforms the EW portfolio and, under the T180_H90 specification, exceeds the CROBEX index, achieving the highest annualized Sharpe ratio among all strategies.

Further on, cumulative returns of the adaptive strategies are compared with the corresponding unregularized portfolios ($\lambda = 0$), EW and the CROBEX index.

Figure 6. Cumulative return of adaptive lambda per strategy and T/H label compared to EW, CROBEX and unregularized portfolio



Source: Authors' calculation


The adaptive GMV strategy does not outperform the CROBEX index under any window specification, although it generally improves upon the EW portfolio. In contrast, the adaptive

CVaR strategy generates substantial cumulative returns under the T90_H90 and T180_H90 specifications, outperforming both the EW portfolio and the CROBEX index. However, in both cases the unregularized CVaR portfolio achieves even higher cumulative returns, suggesting that adaptive regularization does not improve upon the best fixed- λ specification. A Wilcoxon signed-rank test finds no statistically significant difference in Sharpe ratios between the adaptive strategy and the corresponding unregularized strategy. Hence, adaptive walk-forward λ selection does not deliver a statistically significant improvement in risk-adjusted performance. Since 30% maximum weight constraint is already applied it provides diversification pressure, reducing the marginal benefit of ridge regularization.

In addition, to examine the impact of using the median as the aggregation rule for adaptive λ selection and the Return/CVaR ratio as the performance criterion, the corresponding cumulative returns are presented in Figure 7.

Figure 7. Cumulative return of adaptive lambda (median and mode) per strategy and T/H label compared to EW, CROBEX and unregularized portfolio with additional Return/CVaR ratio

Cumulative return start-to-end — adaptive λ^* median vs mode (Sharpe*) vs mod (Return/CVaR), EW, CROBEX (lookback=5 | median/mode criterion=Sharpe* | mod criterion=Return/CVaR ratio)



Source: Authors' calculation

The effect of using the median rather than the mode varies across window specifications. Replacing the modified Sharpe ratio with the Return/CVaR ratio as the selection criterion improves cumulative returns for both GMV and CVaR under T90_H90.

We next examine whether the observed Sharpe-ratio improvement is driven primarily by adaptive ℓ_2 regularization or by the 30% maximum-weight (ℓ_∞) constraint. To disentangle the


two mechanisms, four specifications are considered for each optimization criterion: the unregularized portfolio without an additional maximum-weight cap, ℓ_2 regularization only, the 30% ℓ_∞ constraint only, and the combination of adaptive ℓ_2 regularization with the 30% ℓ_∞ constraint. This comparison separates the standalone effect of ridge regularization from the effect of directly limiting portfolio concentration and allows us to assess whether the two mechanisms are complementary or largely substitutable.

Table 1. Overview of portfolio optimization specifications

	No additional maximum-weight cap	30% maximum-weight cap (ℓ_∞)
$\lambda = 0$	Unregularized	ℓ_∞ constraint only
Adaptive λ	ℓ_2 regularization only	$\ell_2 + \ell_\infty$

Source: Authors' own elaboration

Figure 8. Incremental Sharpe-ratio improvement from the 30% maximum-weight constraint and adaptive ℓ_2 regularization




Source: Authors' calculation

Figure 8 indicates that, for GMV portfolios, both adaptive ℓ_2 regularization and the 30% maximum-weight constraint have only small effects on the Sharpe ratio, with no statistically significant differences. For the CVaR strategy, the 30% maximum-weight constraint produces larger Sharpe-ratio improvements under T90_H90, T180_H90, and REV, whereas adaptive regularization without the cap produces larger improvements under T60_H20, T60_H60, and T756_HREV. Overall, the results suggest that ℓ_2 regularization and ℓ_∞ concentration constraints may act partly as substitutes rather than as fully complementary stabilization mechanisms in this small frontier market. The marginal value of ridge regularization therefore appears to depend on how strongly portfolio concentration is already controlled.

Because the strategies differ substantially in portfolio turnover, cumulative returns were also evaluated net of transaction costs of 10, 20, 30, 50, 100, and 200 bps to assess whether the observed performance persists after accounting for implementation costs. Increasing transaction costs from 10 to 200 bps has only a limited impact on the overall conclusions. The relative ranking of the strategies remains unchanged, with the T90_H90 and T180_H90 CVaR specifications continuing to deliver the strongest performance. In both cases, the unregularized portfolio ($\lambda = 0$) continues to outperform the CROBEX index even after accounting for transaction costs, whereas the adaptive strategy remains more conservative and generates lower cumulative returns.

Figure 9. Crisis protection for T90_H90 label with adaptive lambda compared to EW, CROBEX and unregularized portfolio



Source: Authors' calculation

The results from Figure 9 provide little evidence that the CVaR portfolio offers superior realized downside protection during market crises. Across the three crisis episodes and all together, the GMV portfolio generally achieves lower realized CVaR and lower maximum drawdown than the CVaR portfolio. Minimizing estimated in-sample CVaR does not necessarily translate into lower realized out-of-sample tail risk, particularly when tail-risk estimates are based on relatively short and noisy samples. This challenge may be especially relevant in a small and illiquid market, where estimation error can offset the theoretical advantages of a tail-risk optimization criterion.

Despite its weaker downside protection, CVaR portfolio achieves a higher Sharpe ratio, suggesting that it generates better risk-adjusted returns, although at the cost of larger losses during periods of severe market stress.

4. Conclusion

Under the baseline 30% maximum-weight constraint, adaptive ridge regularization does not produce a statistically significant improvement in Sharpe ratios over the corresponding unregularized portfolios. Although the T90_H90 and T180_H90 CVaR specifications outperform CROBEX in cumulative returns, their unregularized versions achieve higher returns than the adaptive portfolios. In the crisis analysis for T90_H90, GMV generally provides better realized downside protection than CVaR, highlighting that stronger return performance does not necessarily imply better protection during market stress.

Taken together, the results suggest that the effectiveness of ridge regularization depends on the portfolio objective and on whether concentration is already controlled directly. In particular, for the CVaR strategy the 30% maximum-weight constraint appears to capture a substantial part of the performance improvement, while the incremental contribution of

adaptive ℓ_2 regularization is smaller and specification-dependent. For GMV, both effects are generally limited. These findings point to partial substitutability between ℓ_2 regularization and ℓ_∞ concentration control. Because λ is not on a common numerical scale across the GMV and CVaR objectives, the conclusions concern the within-strategy effect of regularization rather than direct comparisons of penalty magnitudes across the two models.

Acknowledgement: This work was funded by the European Union – NextGenerationEU under Call IIP-581-2025.

Declaration of AI assistance: The authors used ChatGPT to assist with language editing and refinement. The authors reviewed and revised the resulting text and take full responsibility for the content of this paper.

Literature:

1. Amenc, N., Goltz, F. & Martellini, L. (2013). Smart Beta 2.0., *EDHEC-Risk Institute*, <https://risk.edhec.edu/publications/smart-beta-2-0>
2. Behr, P., Guettler, A., & Miebs, F. (2013). On portfolio optimization: Imposing the right constraints. *Journal of Banking & Finance*, 37(4), 1232–1242.
3. DeMiguel, V., Garlappi, L., Nogales, F. J., & Uppal, R. (2009a). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, 55(5), 798–812.
4. DeMiguel, V., Garlappi, L., & Uppal, R. (2009b). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The Review of Financial Studies*, 22(5), 1915–1953.
5. Dolinar, D., Zoričić, D., & Kožul, A. (2017). Towards the Estimation of an Efficient Benchmark Portfolio: The Case of Croatian Emerging Market. *Zagreb international review of economics & business*, 20(SCI), 13-23.
6. Grinold, R. C. (1992). Are benchmark portfolios efficient?. *Journal of Portfolio Management*, 19, 34-34.
7. Haugen, R. A., & Baker, N. L. (1991). The efficient market inefficiency of capitalization-weighted stock portfolios. *Journal of portfolio management*, 17(3), 35.
8. Israelsen, C. L. (2005.). A refinement to the Sharpe ratio and information ratio, *Journal of Asset Management*, 5(6), 423-427.
9. Jagannathan, R., & Ma, T. (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance*, 58(4), 1651–1683.
10. Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411.
11. Lovretin Golubić, Z., Dolinar, D., & Zoričić, D. (2025). A heuristic approach to the estimation of an efficient benchmark in the Croatian stock market. *Ekonomski preglad*, 76(1), 3-14.
12. Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of risk*, 2, 21-42.
13. Still, S., & Kondor, I. (2010). Regularizing portfolio optimization. *New Journal of Physics*, 12(7), 075034.
14. Zoričić, D., Dolinar, D., & Lovretin Golubić, Z. (2018). A test of global minimum variance portfolio in the Croatian capital market. In *7th International Scientific Symposium ECONOMY OF EASTERN CROATIA-VISION AND GROWTH* (pp.

1165-1173). Osijek: Ekonomski fakultet Sveučilišta Josipa Jurja Strossmayera u Osijeku.