Mean-Semivariance Optimization Under Transaction Costs

Andrea Rigamonti
Masaryk University
andrea.rigamonti@econ.muni.cz

Matúš Horváth
Masaryk University
matus.horvath@econ.muni.cz

Abstract

We propose a novel portfolio optimization algorithm that accounts for transaction costs without requiring the selection of a turnover penalization parameter. Specifically, we consider the problem of minimizing portfolio risk subject to a target return net of proportional transaction costs measured using a quadratic turnover. An important advantage of the proposed formulation is that it can be directly integrated into a mean-semivariance optimization procedure that employs a re-weighting scheme. We evaluate the approach using empirical datasets of weekly asset returns and find that it significantly improves out-of-sample performance net of transaction costs. The gains are particularly pronounced in mean-semivariance optimization, whose standard implementation can generate high portfolio turnover. Moreover, the regularization induced by the quadratic turnover in the constraint leads to a higher performance even before transaction costs, especially with larger portfolios.

Keywords: semivariance, portfolio optimization, transaction costs
JEL Classification Code: C61, G11

1 Introduction

Targeting the portfolio semivariance has long been suggested as preferable to variance minimization. Markowitz (1959) defines the semivariance as

$$\sigma_B^2 = \frac{1}{T}\sum_{t=1}^{T}[\text{Min}(R_t - B, 0)]^2 \tag{1}$$

where T is the sample size used for estimation, B is a benchmark set by the investor, and R_t is the return at time t. The downside deviation, σ_B, is the square root of σ_B^2. Contrary to the variance, which considers all variability as equally undesirable, σ_B^2 only focuses on volatility below B. Minimizing the variance σ^2 is equivalent to targeting σ_B^2 only if the return distribution is symmetric and B is set equal to the sample mean $\bar{R}$ (the latter condition is not necessary if a target return is specified, see Klebaner et al., 2016). The Sharpe ratio is arguably the most used measure of risk-adjusted return when measuring risk using the standard deviation σ. When σ_B is used, a more appropriate performance metric is the Sortino ratio from Sortino and Prince (1994), defined as:

$$\text{Sortino} = \frac{\bar{R} - B}{\sigma_B} \tag{2}$$

While B can formally take any value, it is advisable to set it equal to the risk-free rate, as this has desirable properties both in terms of objective function (Hoechner et al., 2017) and of performance evaluation using the Sortino ratio (Pedersen and Satchell, 2002).

Despite its theoretical attractiveness, targeting the semivariance has not become as widespread as targeting the variance. This is partly because the corresponding optimization procedure requires as input a semicovariance matrix which is difficult to provide due to its endogenous nature. Specifically, its parameters change with the portfolio weights, which in turn are what the optimization is supposed to provide. This challenging problem has only recently been fully solved (see Ferrari et al., 2026). While this allowed the implementation of standard semivariance-minimizing optimization procedures, expanding the basic algorithms presents additional difficulties. In this paper, we address the inclusion of transaction costs when minimizing the portfolio semivariance given a target return.

Ignored in the original Markowitz (1952) setting, transaction costs were later recognized as important to determine the optimal allocation in a realistic setting. Their first formal comprehensive integration into the optimization algorithm dates back to Pogue (1970). Because introducing transaction costs makes the problem significantly harder to solve, some proposed solutions rely on simplifying assumptions, like costs of fixed amount instead of proportional to the turnover (e.g., Patel and Subrahmanyam, 1982). A computationally efficient solution for large-scale portfolios with realistic variable costs is then proposed by Perold (1984). An arguably more intuitive approach is taken by Yoshimoto (1996), who considers the expected return net of proportional linear transaction costs when maximizing mean-variance utility. The resulting optimal portfolio is shown to be different (and preferable) to the one obtained from traditional utility maximization. More recent works have considered the link between transaction costs and weight instability due to estimation error. Since transaction costs increase with turnover, accounting for them can also reduce large and frequent weights rebalancing driven by unreliable inputs. Notable examples include Olivares-Nadal and deMiguel (2018) and Hautsch and Voigt (2019). The latter shows that using quadratic transaction costs is equivalent to performing a covariance matrix shrinkage, while linear costs induce a LASSO-type regularization. Despite this variety of proposed solutions, Ledoit and Wolf (2025) argue that ignoring transaction costs when computing the optimal weights remains the dominant approach, and they also criticize the simplifying assumption of equal costs across assets. They use instead asset-specific costs, and propose to add a penalization term for the turnover. They show that this results in a more efficient allocation. As in Olivares-Nadal and deMiguel (2018), their solution requires the selection of a constant to specify the penalization intensity, which in practice needs to be selected empirically (usually via cross-validation).

The literature on the analogous problem targeting the semivariance is much more limited. This is likely because portfolio semivariance minimization is itself a challenging problem. The earliest, and arguably still most notable to date, work is from Hamza and Janssen (1998). Instead of estimating a semicovariance matrix, they directly target portfolio semivariance using auxiliary variables that separate deviations above and below the mean. This allows to set up the problem as a mixed-integer quadratic program. Transaction costs can then be accounted for by subtracting them from the available budget. However, their approach has important limitations. Firstly, it requires B to be equal to the mean. Secondly, the problem dimension grows with the length of the estimation window, as it introduces $2T$ auxiliary variables and T equality constraints, plus additional constraints to account for transaction costs. Later works exist, and they address more comprehensive problems (e.g., multiperiod optimization). However, they abandon exact optimization in favour of metaheuristic algorithms (see Zhang et al., 2012; Najafi and Mushakhian, 2015; Chen et al., 2019).

In this paper, we propose a novel optimization framework that accounts for transaction costs without requiring restrictive assumptions, complex metaheuristics, or the selection of a penalization parameter. This allows for its direct integration into the original Ferrari et al. (2026) algorithm, and therefore its application to mean-semivariance optimization problems. However, it is also an attractive solution for the mean-variance framework.

The optimization problem for both settings is described in Section 2. In Section 3 we carry out an empirical application of our solution, and compare it with the standard optimization procedures that ignore transaction costs. Section 4 concludes.

2 Portfolio optimization with transaction costs

In the case of variance minimization with a target return, adding to the objective function a term that penalizes for the turnover is arguably among the most convenient ways to account for trading costs. This solution preserves most of the original problem's structure, and although a closed form solution is typically not available, the resulting problem remains convex and easily solvable numerically. However, this formulation requires specifying a constant term that determines the penalization intensity, whose value needs to be selected empirically. Similar tasks are generally performed via cross-validation, see e.g. Olivares-Nadal and deMiguel (2018). This is not feasible when employing the algorithm introduced by Ferrari et al. (2026), which is based on a re-weighting scheme that simultaneously estimates the optimal weights and the associated semicovariance matrix. In practice, this involves a loop that re-estimates the weights and the semicovariance matrix until reaching convergence. Nesting a cross-validation procedure within this algorithm is not a viable solution. Therefore, in order to account for transaction costs when computing the optimal mean-semivariance weights, a different approach that does not require selecting a penalization parameter is needed. The solution proposed by Yoshimoto (1996) avoids this step by subtracting transaction costs from the mean return in the objective function, but it is tailored to mean-variance utility maximization.

We propose a new formulation that uses the volatility minimization framework, and incorporates the transaction costs into the target return constraint. In other words, the required expected return of the portfolio is specified net of transaction costs. We first describe this approach in the familiar mean-variance setting. Consider a set of N risky assets and one risk-free asset, with vector of mean return μ and covariance matrix Σ. Moreover, assume that trading is associated with proportional transaction costs c, and that the investor has a target portfolio return Re. The one-period optimization problem is:

$$\begin{aligned} \min_{w_t} \quad & w_t'\Sigma w_t \\ \text{s.t.} \quad & w_t'\mu - c(w_t - w_{t-1})'(w_t - w_{t-1}) \geq Re \end{aligned} \tag{3}$$

Hence, the transaction costs are obtained using the quadratic turnover computed as $\sum_{i=1}^{N} (w_{i,t} - w_{i,t-1})^2$ rather than the actual turnover given by $\sum_{i=1}^{N} |w_{i,t} - w_{i,t-1}|$. This allows us to still have a convex and differentiable problem, although without a closed form solution, making problem (3) particularly easy and fast to solve. Moreover, it indirectly performs portfolio regularization by imposing a higher penalty on large weight swings, therefore limiting extreme positions that are often driven by estimation errors.

We now consider the mean-semivariance optimization problem

$$\begin{aligned} \min_{w_t} \quad & w_t' \Sigma_s w_t \\ \text{s.t.} \quad & w_t'\mu - c(w_t - w_{t-1})'(w_t - w_{t-1}) \geq Re \end{aligned} \tag{4}$$

where Σ_s is the semicovariance matrix given by the smooth estimator introduced by Ferrari et al. (2026). In this case, the implementation requires a re-weighting algorithm that, starting from an equally weighted portfolio, simultaneously estimates a new semicovariance matrix and the associated portfolio weights at each iteration, until reaching convergence. Since our setting does not require a data-driven selection of a penalization value, no changes to this algorithm are necessary when applied to solve problem (4).

3 Empirical application

We solve problems (3) and (4) using the solver from the R package *CVXR*. We test our strategy on the weekly returns of the 17, 30 and 48 industry portfolios from the Kenneth R. French Data Library, computed from daily returns.[1] The risk-free rates are downloaded from the same source and, while weekly rates are available, for consistency we compute them by compounding the daily rates. Since data before July 1, 1969, are not available for some assets in the 48 industry dataset, we do not consider periods before that date. Hence, all datasets span from July 11, 1969, to May 29, 2026, for a total of 2969 observations. For the estimation we use a 520-period rolling window, yielding 2449 out-of-sample portfolio returns starting on June 29, 1979. The target return is set equal to 0.1%. The benchmark for the semivariance and the Sortino ratio is set equal to the risk-free rate. Portfolio weights are updated at every new period using inputs estimated with a rolling window. For simplicity, and because we are working with research portfolios, c is set equal to 10 basis points following Frazzini et al. (2012). However, the algorithm can also accommodate asset-specific costs based on the bid-ask spread as in Ledoit and Wolf (2025). In each period t, given a vector of weights w, we compute the turnover

$$TO_t = \sum_{i=1}^{N} \left| w_{i,t} - w_{i,t-1}^{+} \right| \tag{5}$$

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

where $w_{i,t-1}^{+}$ is the weight for asset i at the end of period $t-1$, hence accounting for the effect of its realized return $R_{i,t-1}$, and given by

$$w_{i,t-1}^{+} = \frac{w_{i,t-1} + w_{i,t-1} R_{i,t-1}}{1 + \sum_{i=1}^{N} w_{i,t-1} R_{i,t-1}} \tag{6}$$

The turnover is then used to compute returns net of transaction costs.

We expect our solution to increase the mean portfolio return net of transaction costs, thanks to a reduction in the turnover. However, since we are indirectly performing a regularization by penalizing the quadratic turnover, some performance improvement can be expected also before transaction costs, especially with larger portfolios. We evaluate the results by computing the mean, standard deviation (SD), downside deviation (DD), Sharpe ratio (SR), Sortino ratio and average turnover (TO). For each measure, we also report the difference between the basic strategy and the one that accounts for transaction costs. For the mean, we test the statistical significance of this difference using a two-sided t-test, with heteroskedasticity and autocorrelation consistent (HAC) standard errors from Newey and West (1987). For the Sharpe ratio, we use the test proposed by Ledoit and Wolf (2008), and provided in the R package *PeerPerformance* from Ardia and Seguin (2025). Since we work with time-series data, we employ the studentized circular bootstrap, with 10000 replications.

Table 1 reports the annualized (except for TO) results for the standard mean-variance portfolio (MV), those obtained by solving (3) (MV-TC), and the difference between the two (Change). With all three datasets, MV-TC has a much lower turnover, which results in higher mean and Sharpe ratio, with the difference significant at the 1% for both. While no formal test is available for the Sortino ratio, the difference compared to the one achieved by MV is even larger than for the Sharpe ratio. Though not as large, significant performance improvements of MV-TC over MV are also registered for the results before transaction costs, except for the 17 industry dataset. We interpret this as the effect of the regularization indirectly performed by setting a target return constraint net of quadratic trading costs. As expected, the benefits get larger with a higher N, since larger dimension implies greater estimation error in the inputs. The same applies to results after transaction costs, which explains why bigger performance gains are achieved with larger datasets despite a smaller decrease (both in absolute and relative terms) in the turnover.

The results for mean-semivariance optimization performed with the original Ferrari et al. (2026) algorithm (MSV) and by solving problem (4) (MSV-TC) are reported in Table 2. Compared to the mean-variance setting, we find an even larger increase in mean, Sharpe ratio and Sortino ratio as a result of applying our strategy, both before and after transaction costs. As observed in Table 1, bigger gains are recorded with larger portfolios, even

Strategy	Mean	SD	DD	Sharpe	Sortino	TO
	Panel A: 17 industry – no transaction costs					
MV	0.141	0.237	0.164	0.390	0.564	0.461
MV-TC	0.143	0.239	0.165	0.395	0.572	0.221
Change	0.002	0.002	0.001	0.005	0.008	-0.240
	Panel B: 17 industry – 10 bp transaction costs					
MV	0.114	0.237	0.166	0.289	0.414	0.461
MV-TC	0.130	0.239	0.166	0.347	0.500	0.221
Change	0.016***	0.002	0.000	0.058***	0.086	-0.240
	Panel C: 30 industry – no transaction costs					
MV	0.094	0.199	0.142	0.253	0.356	0.458
MV-TC	0.098	0.200	0.142	0.270	0.380	0.267
Change	0.004***	0.001	0.000	0.017**	0.024	-0.192
	Panel D: 30 industry – 10 bp transaction costs					
MV	0.069	0.199	0.144	0.134	0.186	0.458
MV-TC	0.083	0.200	0.143	0.201	0.281	0.267
Change	0.015***	0.001	0.000	0.067***	0.095	-0.192
	Panel E: 48 industry – no transaction costs					
MV	0.081	0.165	0.116	0.231	0.329	0.464
MV-TC	0.085	0.165	0.116	0.252	0.360	0.309
Change	0.004***	0.000	0.000	0.021***	0.031	-0.154
	Panel F: 48 industry – 10 bp transaction costs					
MV	0.055	0.165	0.118	0.085	0.119	0.464
MV-TC	0.068	0.165	0.117	0.155	0.219	0.309
Change	0.012***	0.000	-0.001	0.070***	0.100	-0.154

Table 1: Annualized performance of mean-variance optimization. Weekly out-of-sample returns from June 29, 1979, to May 29, 2026. Significance level for the differences of mean (using a two-sided t-test with HAC standard errors) and Sharpe ratio (using Ledoit and Wolf, 2008): ***1%, **5%, *10%.

after transaction costs, despite a lower decrease in turnover compared to what is observed with smaller datasets. Overall, our methodology produces the same effects as in the mean-variance framework, but on a larger scale.

Strategy	Mean	SD	DD	Sharpe	Sortino	TO
Panel A: 17 industry – no transaction costs						
MSV	0.145	0.244	0.168	0.394	0.572	0.546
MSV-TC	0.148	0.245	0.169	0.401	0.582	0.252
Change	0.003	0.001	0.001	0.007	0.010	-0.294
Panel B: 17 industry – 10 bp transaction costs						
MSV	0.113	0.244	0.170	0.277	0.397	0.546
MSV-TC	0.133	0.245	0.170	0.347	0.502	0.252
Change	0.020***	0.001	0.000	0.071***	0.104	-0.294
Panel C: 30 industry – no transaction costs						
MSV	0.086	0.205	0.148	0.208	0.289	0.548
MSV-TC	0.091	0.206	0.147	0.231	0.322	0.312
Change	0.005***	0.000	0.000	0.022**	0.032	-0.236
Panel D: 30 industry – 10 bp transaction costs						
MSV	0.056	0.205	0.150	0.069	0.095	0.548
MSV-TC	0.074	0.206	0.149	0.152	0.210	0.312
Change	0.018***	0.000	-0.001	0.082***	0.115	-0.236
Panel E: 48 industry – no transaction costs						
MSV	0.065	0.172	0.123	0.136	0.191	0.575
MSV-TC	0.072	0.173	0.122	0.170	0.240	0.374
Change	0.006***	0.000	-0.001	0.034***	0.048	-0.201
Panel F: 48 industry – 10 bp transaction costs						
MSV	0.034	0.173	0.125	-0.037	-0.051	0.575
MSV-TC	0.051	0.173	0.124	0.057	0.080	0.374
Change	0.017***	0.000	-0.001	0.094***	0.131	-0.201

Table 2: Annualized performance of mean-semivariance optimization. Weekly out-of-sample returns from June 29, 1979, to May 29, 2026. Significance level for the differences of mean (using a two-sided t-test with HAC standard errors) and Sharpe ratio (using Ledoit and Wolf, 2008): ***1%, **5%, *10%.

4 Conclusions

In this paper we proposed a portfolio optimization framework that incorporates proportional transaction costs directly into the target-return constraint using a quadratic turnover measure. The proposed formulation does not require selecting a penalization parameter, and can easily be solved with off-the-shelf numerical solvers. These features allow it to be used in the re-weighting scheme from Ferrari et al. (2026) to compute optimal mean-semivariance weights accounting for transaction costs. It is also an attractive approach in the traditional mean-variance framework. Empirical results on industry portfolios show that our solution consistently reduces portfolio turnover and improves out-of-sample performance after transaction costs. Moreover, the regularization induced by the quadratic turnover in the constraint also leads to performance gains before transaction costs, especially in larger portfolios. While the benefits are substantial even for mean-variance optimization, the improvements are particularly pronounced when targeting the semivariance.

References

Ardia, D. and Seguin, B. (2025). Peerperformance: Luck-corrected peer performance analysis in R. Available at SSRN: https://ssrn.com/abstract=5418256.

Chen, W., Li, D., Lu, S., and Liu, W. (2019). Multi-period mean-semivariance portfolio optimization based on uncertain measure. *Soft Computing*, 23:6231–6247.

Ferrari, D., Paterlini, S., Rigamonti, A., and Weissensteiner, A. (2026). Smoothed semicovariance estimation for portfolio selection. *Annals of Operations Research*, 357:565–604.

Frazzini, A., Israel, R., and Moskowitz, T. (2012). Trading costs of asset pricing anomalies. Available at SSRN: https://ssrn.com/abstract=2294498.

Hamza, F. and Janssen, J. (1998). The mean-semivariances approach to realistic portfolio optimization subject to transaction costs. *Applied Stochastics Models and Data Analysis*, 14(4):275–283.

Hautsch, N. and Voigt, S. (2019). Large-scale portfolio allocation under transaction costs and model uncertainty. *Journal of Econometrics*, 212(1):221–240.

Hoechner, B., Reichling, P., and Schulze, G. (2017). Pitfalls of downside performance measures with arbitrary targets. *International Review of Finance*, 17(4):597–610.

Klebaner, F., Landsman, Z., Makov, U., and Yao, J. (2016). Optimal portfolios with downside risk. *Quantitative Finance*, 17(3):315–325.

Ledoit, A. and Wolf, M. (2008). Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance*, 15(5):850–859.

Ledoit, A. and Wolf, M. (2025). Markowitz portfolios under transaction costs. *Journal of Multivariate Analysis*, 100.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.

Markowitz, H. (1959). *Portfolio selection*. John Wiley and Sons.

Najafi, A. A. and Mushakhian, S. (2015). Multi-stage stochastic mean-semivariance-cvar portfolio optimization under transaction costs. *Applied Mathematics and Computation*, 256:445–458.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.

Olivares-Nadal, A. and deMiguel, V. (2018). Technical note – a robust perspective on transaction costs in portfolio optimization. *Operations Research*, 66(3).

Patel, N. R. and Subrahmanyam, M. G. (1982). A simple algorithm for optimal portfolio selection with fixed transaction costs. *Management Science*, 28(3):303–314.

Pedersen, C. S. and Satchell, S. E. (2002). On the foundation of performance measures under asymmetric returns. *Quantitative Finance*, 2:217–223.

Perold, A. F. (1984). Large-scale portfolio optimization. *Management Science*, 30(10):1143–1160.

Pogue, G. A. (1970). An extension of the markowitz portfolio selection model to include variable transactions' costs, short sales, leverage policies and taxes. *The Journal of Finance*, 25(5):1005–1027.

Sortino, F. A. and Prince, L. N. (1994). Performance measurement in a downside risk framework. *The Journal of Investing*, 3(3):59–64.

Yoshimoto, A. (1996). The mean-variance approach to portfolio optimization subject to transaction costs. *Journal of the Operations Research Society of Japan*, 39:99–117.

Zhang, W.-G., Liu, Y., and Xu, W.-J. (2012). A possibilistic mean-semivariance-entropy model for multi-period portfolio selection with transaction costs. *European Journal of Operational Research*, 222(2):341–349.