The risk of falling short: Implementation Shortfall variance in portfolio construction*

Filip Bašić[a] Alberto Martín-Utrera[b]
Ingmar Nolte[a] Sandra Nolte[a]

[a]Lancaster University Management School, Lancaster, UK
[b]Iowa State University, Iowa, USA

September 2, 2025

Abstract

Transaction cost variance introduces a risk often neglected in portfolio optimization. Adopting a mean-variance portfolio optimization problem, we show that including a transaction cost variance term can significantly impact the associated portfolios' performances. Transaction cost variance is estimated based on a transaction cost model constructed using proprietary data from a large institutional investment company. In addition to variance, we estimate transaction cost covariances and construct a transaction cost covariance matrix. Using a standard time-series model setup for returns, we show that considering transaction cost covariance leads to improved net risk-adjusted performance.

Keywords: Portfolio optimization, Implementation Shortfall, Time-series models.
JEL Classification: C01, C32, C51, C52, C61, G11, G12.

*We wish to thank Chris Adcock (editor), the associate editor, two anonymous referees, Lars ter Braak, Iman Honarvar, Edward Leung, Harald Lohre, Martin van der Schans, Jerry Sun, participants of the 2022 Financial Engineering and Banking Society (FEBS) Conference in Pourtsmouth, Invesco Quantitative Strategies, and the Robeco Quant Research Seminar for very helpful comments and suggestions.

1 Introduction

Portfolio optimization is a common means to construct and investigate portfolios in both theoretical and empirical asset management, as highlighted in the seminal work by Markowitz (1952) on portfolio selection and risk diversification. Markowitz's contributions laid the foundation for modern portfolio theory, emphasizing the importance of optimizing the trade-off between risk and return in portfolio construction, see Markowitz (1959). The portfolio optimization setups used in the more recent literature vary significantly, but usually consist of two parts, one based on expected risk and return of the underlying assets, and one based on transaction costs. The latter part usually builds on a model that estimates the expected transaction costs for a given transaction. The simple question we seek to answer is, what if we include transaction costs for what they are, negative returns, and include an additional risk penalty for their uncertainty? In other words, if considering expected transaction costs improves net portfolio performance, would adding a risk term for transaction costs help too?

In this paper, we broaden the traditional focus on return variance by explicitly considering the variance of transaction costs, which captures the uncertainty associated with executing trades under real-world market conditions. In contrast to the broader market-driven risks embedded in portfolio returns, transaction cost variance is driven by microstructural and liquidity-related factors such as bid-ask spreads, market depth, price impact, and order book volatility. These elements can fluctuate substantially over short horizons, especially when trading involves multiple securities that share liquidity pools or overlapping market participants. Consequently, transaction costs often exhibit cross-security dependencies, highlighting the relevance of recent insights on cross-impact effects, see Min, Maglaras, and Moallemi (2022). While the risk of prices moving away from a target position introduces some degree of correlation between returns and transaction costs, the relatively short execution window in our setting—where all trades are typically completed within a few days—mitigates this overlap in practice. By modeling the covariance structure of transaction costs alongside return variance, our approach aims to better capture the interplay of liquidity shocks, market microstructure dynamics, and cross-security effects, ultimately providing a more comprehensive framework for mean-variance portfolio selection.

The first examples of studies incorporating transaction costs into a portfolio optimization go back to Magill and Constantinides (1976) and Davis and Norman (1990). The key source of development in this area is the transaction cost modelling literature, starting with Loeb (1983) specifying a proportional transaction cost model, followed by Kyle (1985) introducing linear price impact in transaction cost modelling, which is further developed in Glosten and Harris (1988). This theoretical work has been further empirically developed in Breen, Hodrick, and Korajczyk (2002), which allowed DeMiguel, Martín-Utrera, and Nogales (2015) to construct portfolios using linear price impact and Korajczyk and Sadka (2004) to study the performance of momentum portfolios under different transaction cost measures. Following this, Gabaix, Gopikrishnan, Plerou, and Stanley (2006) further develop square root price impact, basing their transaction cost model on Torre and Ferrari (2000). This provides the theoretical framework for Almgren, Thum, Hauptmann, and Li (2005), who estimate these non-linear price impact models based on real trading data. More recently, Frazzini, Israel, and Moskowitz (2018) use real trading data to estimate a linear transaction cost model using both linear and square root price impact to investigate how transaction costs vary across trade types, stock characteristics, trade size, time, and exchanges globally. Their transaction cost model is the motivation behind the transaction cost model developed in the present paper, as their model is in agreement with the theoretical and empirical findings to date. The data we use is also of similar format to what they use, and they show their model to work well with this type of transaction cost data. The measure we use to measure transaction costs is implementation shortfall (IS), which is the most common measure used in literature to quantify transaction costs, see Perold (1988). In more recent work, de Rossi, Hoch, and Steliaros (2022) use implementation shortfall as a measure for transaction costs estimated using real trading data to compute the net performance of various factor strategies. The inclusion of transaction costs as an explicit penalty term in portfolio optimization has become a common practice in both academic research and industry applications. Commercial portfolio management software, such as Axioma, incorporates transaction costs into the optimization process to reflect real-world constraints and improve the implementability of optimized portfolios.

To showcase the benefits of modelling transaction costs in such a manner, we carry out a multivariate simulation study for 123 stocks, which are constituents of the S&P 500 index. Leveraging a proprietary dataset of real trading data for these 123 stocks from a large institutional asset manager we estimate a transaction cost model, with the aim of constructing a variance-covariance matrix of transaction costs and then including it in portfolio optimization. The importance of considering transaction cost covariance is further highlighted in Min et al. (2022), who show when working with highly correlated assets, trades in one security can influence the prices of other securities. Their work extends the understanding of market impact beyond individual securities by analyzing the interdependencies between assets, providing insights into how trading activity in one asset affects the broader portfolio. Specifically, we construct implementation shortfall residuals for every given trade using the estimated transaction cost. We propose a new transaction cost volatility model specified as a linear model with squared residuals as the dependent variable, enabling us to estimate the variance of transaction costs for a given trade. We obtain the correlation matrix by constructing weekly residual estimates for 123 stocks considered. Combining the correlation matrix and the transaction cost volatility model, we construct a variance-covariance matrix of transaction costs that is entering portfolio optimization as an additional term. The latter is estimated using an AR(1)-DCC(1,1)-GARCH(1,1) model estimated on return data of the 123 stocks considered. Upon constructing every element of all the portfolio optimization setups considered, we construct these portfolios. Portfolios are path dependant and optimized monthly based on simulated stock returns, the stock return covariance matrix, transaction costs estimates and the transaction cost covariance matrix estimates. Finally, we calculate the net performance of each portfolio and compare them across different parameter assumptions.

This paper makes two key contributions. First, since we find transaction costs to be volatile, we model their covariance. Although volatility of transaction costs has been evidenced in for example Frazzini et al. (2018) and de Rossi et al. (2022), no attempts were made to model the covariance matrix of transaction costs in the literature. We show that modelling transaction cost covariance leads to a better overall fit. Second, we broaden the way of incorporating transaction costs into a mean-variance portfolio optimization seeing this addition improves portfolio performance in most cases. More specifically, we show that Sharpe ratios are increased, implying that incorporating transaction costs in portfolio optimization is more efficient when transaction cost volatility is also included.

The paper is structured as follows. Section 2 describes the data used for both the transaction cost modelling and time-series model estimation. Section 3 defines the portfolio optimization setups we use, and discusses implementation shortfall, transaction cost models, transaction cost variance and covariance as well as the corresponding estimation procedures. Further, it describes the return simulation procedure and estimation of expected variance-covariance matrix of the returns. Section 3 concludes with portfolio optimization and evaluation of the performance of the portfolios. Section 4 contains the estimation results for the transaction cost model, transaction cost variance model, as well as portfolio performance. Our conclusions are presented in Section 5.

2 Transaction cost and stock data

For the return data, we use the Center for Research in Security Prices (CRSP) daily data files ranging from January 2015 to December 2023 for the 123 stocks we consider. All of the stocks considered are large market capitalization companies and constituents of the S&P 500. We increase the period compared to the transaction cost data for the purpose of model fitting. Table 1 reports the summary statistics on returns, trading volume and market capitalization across all 123 stocks.

Table 1: Stock data summary statistics

We present the summary statistics of our stock data. The data consists of 123 stocks' daily returns, volume traded and market capitalization covering the period from January 2015 to March 2023. All values reported are averages across 123 stocks and 2,075 days.

	Return (%)	Volume ($\times 10^6$ shares)	Volume ($ bil)	Market Cap ($ bil)
min	-14.35	2.29	0.26	54.59
q25	-0.75	8.58	1.08	80.82
Median	0.08	11.30	1.51	114.74
q75	0.92	15.23	1.82	172.77
max	14.62	99.78	11.51	254.64
Mean	0.06	12.94	1.70	128.44
SD	1.77	7.10	0.99	54.52

Our trading data consists of 38,250 trades from a large institutional asset manager, covering the period from July 2017 to March 2023. The underlying assets of the trades are US equities. The database is compiled by the trading desk and covers all US trades executed subject to a frequency requirement. Each of the 123 stocks included in the database has at least one occurrence of being traded every single week throughout the entire period considered. Every trade is executed by the trading desk in multiple smaller executions and the relevant information is then aggregated on all the executions done for the trade. This includes trade identifier, stock identifier, timestamps of the beginning and the end of the trade, where the beginning of the trade is its arrival time to the trading desk and the end of the trade timestamp is created once the last trade is executed and the trade is completed.

The size of the trade is given in number of shares, which is the initially intended number of shares to be traded; as all of the trades we consider are fully executed, it equates to the number of shares traded. The total value of the executed position is given as the product of the number of shares traded and average execution price, quoted in USD. The trading data also includes the share price at the start of the trade, which we can use to calculate the average price impact exerted by the trade. Each trade is executed within 5 days from order creation, with an average execution time of 1.9 days.

The main parameter used in our transaction cost analysis is the trade size as a percentage of median daily volume (MDV) which we calculate over the last 25 trading days. Another important variable is the range volatility, measured as the variance of r_t^{HL} over the last 15 days, where:

$$r_t^{HL} = \ln(H_t) - \ln(L_t), \tag{1}$$

where H_t and L_t are the highest and lowest prices on day t.

Table 2 reports summary statistics of the transaction costs data. As mentioned, the measure we use for transaction costs is implementation shortfall (IS), which we define in Section 3.1. The median trade size observed is 0.16% of median daily volume (MDV), similar to what Frazzini et al. (2018) report. We also observe the volatile nature of implementation shortfall, as the mean observed is 0.06% and standard deviation is 1.39%. As almost half of our trades are negative in implementation shortfall, we can expect our transaction cost model, which will (and should) always estimate transaction costs to be positive, to have a large prediction error. This is why we argue penalizing the variance of transaction costs will help the risk-adjusted performance of our portfolios.

Table 2: Transaction cost data summary statistics across regions

We present the summary statistics of our trading data. The data consists of 38,250 trades from a large institutional asset manager, covering the period from July 2017 to March 2023. We include the trade size as a percentage of median daily volume (MDV), volatility and implementation shortfall. Median daily volume is calculated as the last 25 trading days median of daily volume. Range volatility is calculated using the past 15 days, and we only consider observations with range volatility between 5 and 100%. Implementation shortfall is calculated in accordance to Perold (1988), see Equation 6, and is reported as a fraction of the original trade value and capped at 10%.

	% MDV	Vola (%)	IS (%)
min	0.00	5.17	-9.37
q25	0.05	18.89	-0.79
Median	0.16	24.92	0.06
q75	0.63	33.55	0.69
max	172.19	100.00	9.15
Mean	0.85	28.83	0.06
SD	2.88	13.75	1.39

3 Portfolio construction and transaction costs

Consider a market with N stocks and no risk-free asset available to trade. Let $R_t \in \mathbb{R}^N$ be the corresponding return vector at time t. The return of the portfolio at time $t+1$ is:

$$R^p_{t+1} = w_t^\top R_{t+1}, \tag{2}$$

where w_t is the weight vector at time t. We define a mean-variance portfolio as the set of weights w_t that satisfy

$$\max_{w_t} \mathbb{E}_t[w_t^\top R_{t+1}] - \frac{\gamma}{2} Var_t[w_t^\top R_{t+1}], \tag{3}$$

where $\gamma > 0$ is the risk aversion parameter and $\forall t \ \sum_{i=1}^N w_{i,t} = 1$. Following this procedure, we can construct monthly mean-variance portfolios at time t using information available to us up to time t. Now consider the following transaction cost function, $TC_t(w_t, \theta) = MI(w_t, \theta) + \varepsilon_t$, where MI denotes market impact and is a function of the change in weights, Δw_t, and other parameters contained in θ, and ε_t is the error term. Assuming an investor constructs a portfolio using mean-variance optimization, we have:

$$\max_{w_t} \mathbb{E}_t[w_t^\top R_{t+1}] - \frac{\gamma}{2} Var_t[w_t^\top R_{t+1}] - \mathbb{E}_t[\Delta w_t^\top TC_t(\Delta w_t, \theta)], \tag{4}$$

where $TC(w_t)$ are the corresponding transaction costs.

Assuming transaction costs of a portfolio are a stochastic process $\{TC_t\}_{t \in \mathbb{N}_T}$ with an associated variance-covariance matrix, we can define a new optimization problem by extending (4) to include a penalty term for the variance of transaction costs

$$\max_{w_t} \mathbb{E}_t[w_t^\top R_{t+1}] - \frac{\gamma}{2} Var_t[w_t^\top R_{t+1}] - \mathbb{E}_t[\Delta w_t^\top TC_t(\Delta w_t, \theta)] - \frac{\gamma}{2} Var_t[\Delta w_t^\top TC_t(\Delta w_t, \theta)], \tag{5}$$

Note that we will be using the same risk-aversion parameter for returns and transaction costs, as there is no inherent difference between the two. They are both reflective of returns, with opposing signs which does not affect the risk considerations.

3.1 Implementation Shortfall and Market Impact

Following Perold (1988), we consider the following simplified scenario. Suppose we have an order to buy n shares within a period and denote the price of one share at the start (end) of our trade net of fixed costs as p_{start} (p_{end}). Assume we finish trading having bought $m < n$ shares. Furthermore, assume the m shares were bought in T individual transactions. Let q_k denote the amount of shares bought in transaction k, and let p_k be the price of the share during our transaction k. Then we can define implementation shortfall (IS) as the sum of market impact (MI) and opportunity cost (OC),

$$IS = MI + OC = \sum_{k=1}^{T} q_k(p_k - p_{start}) + (n - m)(p_{end} - p_{start}), \tag{6}$$

where

$$m = \sum_{k=1}^{T} q_k. \tag{7}$$

Intuitively, market impact measures the additional cost of trading due to the price drifting away from its value at the start of the trade, and opportunity cost is the difference in values of the underlying asset that one failed to obtain by not executing the trade fully. Since our trade data is constructed using executed trades, we omit opportunity cost and use implementation shortfall directly as a proxy for market impact.

3.2 Modelling transaction costs

Transaction cost modelling has gained traction in recent times, ranging from (effective) bid-ask spreads, see de Groot, Huij, and Zhou (2012), to estimating market impact with the purpose of capturing the impact a given trade has on the market price. For example, Frazzini et al. (2018) and de Rossi et al. (2022) both estimate a market impact model which they go on to use in portfolio construction. The most relevant literature suggests market impact, when taken as a function of the trade size in median daily volume, behaves as a polynomial function with an exponent between 0.5 and 1. We model our transaction costs using the I-Star model of Kissell (2014), which is widely used in practice. In order to make the optimization problem computationally more feasible, we estimate a simplified version of the I-Star model of the following form:

$$TC_t(\Delta w_{i,t}) = a_1 \sigma_{i,t} \frac{AuM \times \Delta w_{i,t}}{MDV_{i,t}} + \epsilon_{i,t}, \tag{8}$$

where TC is implementation shortfall as a fraction of trade size, $\sigma_{i,t}$ is the range volatility of stock i at time t, and AuM is the assumption of the assets under management. In other words, AuM denotes the dollar value of the portfolio we trade with. We will keep this constant across time, as we wish for transaction costs to be comparable across the entire period and not be skewed by portfolio size, and investigate for different values. The product $AuM \times \Delta w_{i,t}$ gives us the total amount traded in USD, which we scale by median daily volume (MDV).

The parameter estimation is carried out in a rolling-window step procedure, with a base period of 2 years adding the next and subtracting the last month every step starting with July 2017, making the first period July 2017 – June 2019. For each period, we estimate the a_1 parameter using linear regression.

3.3 Variance of transaction costs

Having fitted the transaction cost model, we can compute the IS residuals $\hat{\varepsilon}_{i,t}$ for every trade within a window. We observe that squared residuals are increasing in both trade size and range volatility, with correlations of 29% and 12%, respectively. This leads us to believe that we can model our residuals in a similar manner as we modelled implementation shortfall. Specifically, we model the variance of our transaction costs as:

$$Var_t[TC_t(\Delta w_{i,t})] = b_1 \sigma_{i,t} \frac{AuM \times \Delta w_{i,t}}{MDV_{i,t}} + \eta_{i,t}. \tag{9}$$

Equivalently to implementation shortfall, the estimation is done in a rolling-window fashion, using a base period of 2 years adding the next and subtracting the last month every step.

In our framework, the transaction cost mean and variance function coefficients, a_1 and b_1, are used both for simulating transaction costs and in the optimization process. This modeling choice ensures consistency between the simulation and optimization frameworks, allowing us to isolate the effects of incorporating transaction costs into portfolio selection. While this approach assumes no estimation error in the transaction cost parameters, it is appropriate within the scope of this study, as the primary objective is to evaluate the impact of transaction cost-aware optimization rather than the estimation accuracy of transaction cost models.

3.4 Covariance of transaction costs

In order to use transaction cost variance in the same manner as we use variance of returns as a penalty function, we require the computation of a covariance matrix of transaction costs. However, unlike returns, transaction costs (and trades) are irregular discrete events, where we can observe many or no trades in a given period. To resolve this, we define a procedure of obtaining standardized residuals of transaction costs for each stock in a given window. We construct weekly residuals for each stock and use those to construct a covariance matrix.

Using the computed residuals $\hat{\varepsilon}_{i,t}$, we estimate an unconditional covariance matrix in the following way: for each window, we split our observations into weeks, yielding on average 105 weekly groups per window for every stock with an average of 1.3 observations. We calculate a weekly residual for stock i in week w, $\hat{\varepsilon}_{i,w}$, as the average residual in the corresponding week. Hence, we obtain a $N_{stocks} \times N_{weeks}$ matrix of residuals of equal number of observed residuals for every stock.

$$\hat{\varepsilon}_{i,w} = \frac{\sum_{t \in w} \hat{\varepsilon}_{i,t}}{\mathbf{card}(\{\hat{\varepsilon}_{i,t} : t \in w\})}. \tag{10}$$

To obtain the standardized residuals, $u_{i,w}$, we scale the residuals with the standard deviation of the residuals of the respective company,

$$u_{i,w} = \frac{\hat{\varepsilon}_{i,w}}{\sqrt{Var(\hat{\varepsilon}_{i,w} : w \in \{1, ..., N_w\})}}. \tag{11}$$

Denoting the matrix of these standardized residuals U_t, we compute the correlation matrix as

$$R_t^{TC} = U_t U_t^T. \tag{12}$$

Again, for consistency, this computation is carried out in a rolling-window step procedure, with a base period of 2 years adding the next and subtracting the last month every step. The reason behind choosing such a large base period is mainly due to the size of our covariance matrix, and how many observations needed to be taken into account for a stable estimate with full rank, avoiding any rank defective matrix issues. Finally, the resulting transaction cost covariance matrix is a product of the transaction cost variance function and the transaction cost correlation matrix $\hat{R}_t^{TC}$:

$$\hat{\Sigma}_t^{TC} = \hat{D}_t^{TC} \hat{R}_t^{TC} \hat{D}_t^{TC}, \tag{13}$$

where $\hat{D}_t^{TC}$ is a diagonal matrix of the transaction cost volatility estimates obtained by taking a square root of the estimated transaction cost variance, given by Equation 9.

3.5 Simulation study

In order to evaluate and compare the performance of the three portfolio optimizations, we simulate daily returns and the covariance matrix for the 123 stocks we consider over the given period. Our base assumption will be that returns follow an AR(1)-DCC(1,1)-GARCH(1,1) process. A detailed description of all of these models and their estimation procedures is given in the Appendix A. To ensure the robustness of our results, we employ a fully out-of-sample framework in our analysis. At each time t, all model parameters are estimated using only the data available up to time t. These estimates are then used in portfolio optimization for the subsequent period $t+1$. We further clarify that the GARCH parameters are estimated only once using the historical data available up to that point. Once estimated, we use these parameters and daily simulated return data to simulate conditional variances and covariances daily. The conditional means and variances are then drawn daily and are directly used as inputs for the portfolio optimizer. This ensures that the optimizer only relies on realistic, forward-looking estimates rather than true underlying parameters, maintaining consistency with an out-of-sample framework.

For our mean process, AR(1) was chosen over a more complex ARMA(p,q) process due to the quality of the parameters estimated, leading us to the conclusion that adding any moving average terms did not yield any additional explanatory power. Hence, we simplify to an AR(1) model with a reasonable ϕ_1 parameter.

For the GARCH model, we find the stability of our parameters to be the highest for a GARCH(1,1) model. For the correlation part, the DCC(1,1) model resulted in the dynamic parameter being equal to zero, implying a CCC (constant) model, which means our correlation does not observe a strong dynamic pattern.

To simulate returns, we first use the estimated AR(1) model to compute the daily mean series μ_d. Our time series modeling will be done on a daily basis to improve the quality of our estimates, and we will use the daily predictions to construct monthly ones. The period we consider begins in June 2017 and ends in April 2023. We calculate the μ_d series daily, with a burn-in period of one year used to estimate the AR(1) parameter ϕ_d. The AR(1) model used for the conditional mean is given by:

$$\hat{\mu}_d = \hat{c}_d + \hat{\phi}_d r_{d-1}, \tag{14}$$

where $\hat{c}_d$ and $\hat{\phi}_d$ denote the AR(1) model parameters estimated using data up to day d. Upon obtaining the estimates of the μ_d series, we move on to generating the estimates of the σ_d series. The estimation of σ_d is done in the following way for each given day d: we create two series of length 251, σ'_k and ε'_k which will represent the σ and ε series for the GARCH(1,1) process finishing at day d with parameters ω_d, α_d, β_d. We will set the starting values for σ'_k and ε'_k and use the parameters to calculate the series. The final observation of σ'_k will be σ_d.

$$\hat{\varepsilon}'_0 = r_{d-250} - \hat{\mu}_{d-250}, \tag{15}$$

$$\hat{\sigma}'^2_0 = \frac{\hat{\omega}_d}{1 - \hat{\alpha}_d - \hat{\beta}_d}, \tag{16}$$

$$\hat{\varepsilon}'_k = r_{d-250+k} - \hat{\mu}_{d-250+k}, \tag{17}$$

$$\hat{\sigma}'_k = \left(\hat{\omega}_d + \hat{\alpha}_d\left(\hat{\varepsilon}'_{k-1}\right)^2 + \hat{\beta}_d\left(\hat{\sigma}'_{k-1}\right)^2\right)^{\frac{1}{2}}, \tag{18}$$

with $k = 1, ..., 250$.
Finally,

$$\hat{\sigma}_d = \hat{\sigma}'_{250}. \tag{19}$$

Now that we have the μ_d and σ_d series estimates, we can simulate our cross-section of returns on a daily basis. Since our DCC(1,1) model resulted in a CCC model, we obtain the correlation matrix $\hat{R}$ by simply estimating the sample correlation. We construct the variance covariance matrix $\hat{\Sigma}_d$ using a diagonal matrix $\hat{D}$ of the estimated individual volatilities $\hat{\sigma}_d$ and a correlation matrix $\hat{R}$ using the following equation:

$$\hat{\Sigma}_d = \hat{D}\hat{R}\hat{D}, \tag{20}$$

Assuming a normal distribution, the simulation of returns on day d will be the realisation of a normal random variable with mean vector $\hat{\mu}_d$ and variance-covariance matrix $\hat{\Sigma}_d$,

$$r_d \sim \mathcal{N}\left(\hat{\mu}_d, \hat{\Sigma}_d\right). \tag{21}$$

We use Equation 21 to construct 100 simulations of returns for our cross-section of 123 stocks. Each simulation consists of 2,075 daily returns of 123 stocks for the period from June 2017 to April 2023. The results of one such simulation for one stock (Microsoft) are shown in Figure 1. We plot the simple and cumulative returns for the period starting June 2019, since that will be the starting point of our optimization. For comparison, in the same figures, we plot the historical return series and the cumulative return series denoted with a blue line. Furthermore, we plot the μ_d series, historic returns, σ_d series and 20-day volatility. We see that our mean (volatility) estimates closely resemble the returns (20-day volatility), suggesting our model fits the data properly. This means our simulations can be used to ensure the robustness of our results, assuming our results are consistent across simulations.

Figure 1: Microsoft simulations

We plot the simple returns, 20-day volatility, $\hat{\mu}_d$ and $\hat{\sigma}_d$ series, simulated simple and cumulative returns of Microsoft from June 2019 to April 2023. In **(e)**, we plot 5 return simulations alongside the historic returns which are depicted in blue. We plot the corresponding cumulative returns of both the simulated and historic returns, which are depicted in blue, in **(f)**.

(a) Microsoft returns

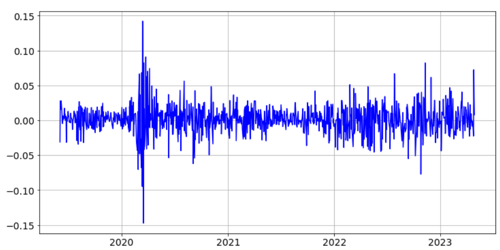

(b) Microsoft $\hat{\mu}_d$ series

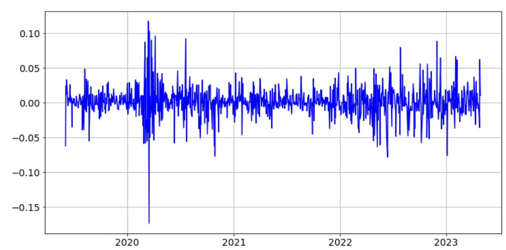

(c) Microsoft 20-day volatility

(d) Microsoft $\hat{\sigma}_d$ series

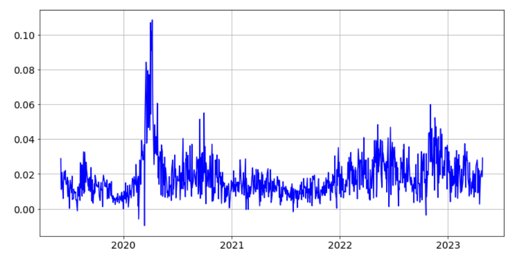

(e) Microsoft simulated returns

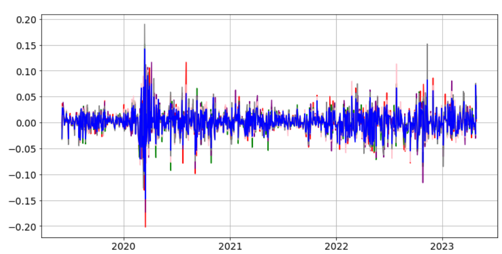

(f) Microsoft simulated cumulative returns

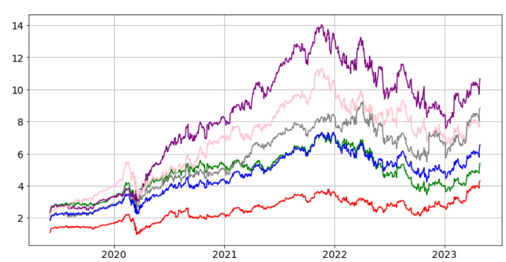

3.6 Expected variance

In order to finalize the optimization procedure, we have to define the return and covariance matrix estimates. Since the optimization procedure will be carried out monthly, assuming a monthly rebalanced portfolio, we have to construct monthly forecasts for the mean and variance. Our mean estimates will be forecasted using the AR(1) model estimated by simply forecasting the next month of daily means and summing them. To compute the 1-month forecasted volatility, we use the daily volatility estimate multiplied by the square root of the number of days in a given month, hence rescaling the daily volatilities into monthly ones. For the correlation, since we do not observe any dynamic properties, we simply use the correlation matrix of returns R estimated at the day d of the rebalance.

3.7 Optimizing portfolios

For notation purposes, let us define a monthly rebalanced portfolio P as a $N_{months} \times N_{stocks}$ matrix of weights with each row representing weights for a given month. As previously mentioned, we will evaluate the performance of three different portfolios, each characterized by the different objective functions they optimize. There are three key components to constructing our portfolios: 1. expected variance of returns, which we obtain from the daily estimates of the covariance matrix and rescale to monthly ones, 2. transaction costs and 3. transaction cost covariance matrices, both of which are estimated monthly, so every 3 months we update our expectations with the newly fitted models. We can now define portfolios P^1, P^2 and P^3, as the resulting portfolios of Equations 3, 4 and 5 respectively.

3.8 Optimization

The optimization we perform will be subject to certain initial conditions and constraints. We will assume that portfolios have a constant number of assets under management (AuM) during the observed period, but different values will be explored. The portfolios will be long only with a maximum holding constraint of 5% in a single stock. Turnover constraints will be implemented in the form of not allowing more than 100% of median daily volume (MDV) to be traded in either direction in a given rebalancing period. The rebalancing will be performed monthly, over a period of 47 months resulting in 46 rebalances, starting in June 2019 and ending in April 2023. The AR-GARCH as well as TC and TC covariance parameters used are based on data up until that month, avoiding any look-ahead bias. We assume a value-weighted starting portfolio. This bears little to no importance to our analysis, as our portfolios generate turnovers large enough to trade out of initial positions within a few months. That being said, we want to start from a liquid position that will be cheap to trade out of, hence why we set the initial position to be the value-weighted portfolio.

3.9 Performance evaluation

To compare the performance of our portfolios generated by different optimization problems, see Equations 3, 4 and 5, we calculate the out-of-sample performance net of transaction costs based on the transaction cost and the variance of transaction costs models. Returns are obtained from the return simulation and transaction costs are calculated by a draw from a multivariate normal random variable:

$$TC_t = \mathcal{N}\Big(\widehat{TC}_t(\Delta w_t, \theta), \hat{\Sigma}_t^{TC}\Big). \tag{22}$$

The mean will be given by the estimated transaction cost function following Equation 8, whereas $\hat{\Sigma}_t^{TC}$ is estimated following Equation 13 using the transaction cost variance function and the estimated transaction cost correlation matrix $\hat{R}_t^{TC}$. Since it can be done independently of the portfolio optimization process, we simulate 100 transaction costs per portfolio for every month.

4 The relevance of transaction cost variance

In this section, we showcase the benefits of including the transaction cost variance in portfolio optimization. Starting with the analysis of all parameters obtained and validating the return simulation, we then look into the net performances of our portfolios.

4.1 Transaction cost function parameters

In Figure 2, we plot the estimated transaction cost model parameter. The parameter is similar across the entire estimation period which is important for our analysis as we do not wish to penalize transaction costs differently across different periods. This is important for both the optimization and net return calculations.

As shown in Table 3, the parameter is highly significant with an average t-stat of 8.24. The mean R-squared is 3.6% which is expected as transaction cost data is very volatile, see Frazzini et al. (2018). This further reinforces the importance of considering transaction cost variance.

4.2 Transaction cost variance parameters

In Figure 3, we plot the transaction cost variance function parameters estimated across different periods. Similar to the transaction cost parameter a_1, the transaction cost variance parameter b_1

Figure 2: Transaction cost parameter estimates

We plot the transaction cost model parameter a_1 estimated from June 2019 to March 2023. The estimation is done every month using past two years' trading data. Every point is plotted with the corresponding 95% confidence interval.

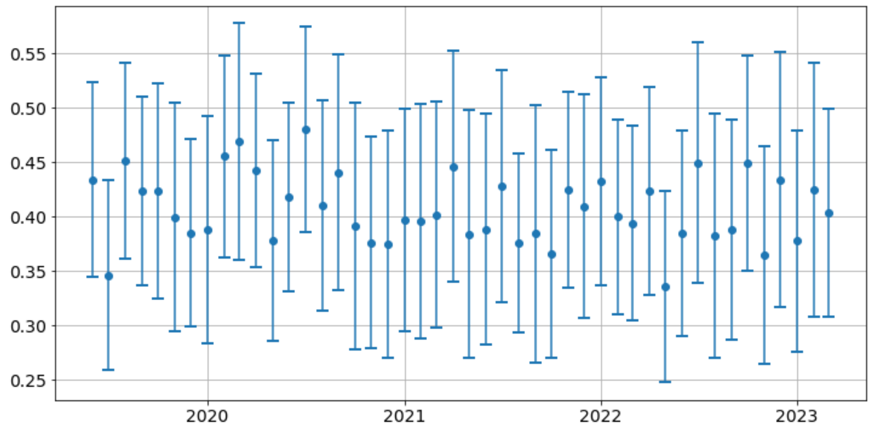

Table 3: Transaction cost estimation results

We show the mean, median, standard deviation minimum and maximum of parameter values, t-statistics and R-squared estimated from June 2019 to March 2023. The estimation is done every month using past two years' trading data.

	Mean	Median	Std	Min	Max
a_1	0.41	0.40	0.032	0.34	0.48
t-statistic	8.24	8.12	0.96	6.48	10.10
R^2 (%)	3.62	3.51	0.72	2.10	5.44
p-value	<0.001				

is stable across time, ensuring consistency of optimization and transaction cost simulations across periods.

As shown in Table 4, b_1 has a mean of 3.55. Since the models we use are identical, we see that the variance of transaction costs are about an order of magnitude larger than transaction costs themselves. This ensures that adding a transaction cost variance covariance matrix penalty in our portfolio optimization setup impacts the weights in a meaningful manner. b_1 is also highly statistically significant with a mean t-statistic of 17.61. The mean R-squared is 72.35 %; this far above any transaction cost model R-squared estimated on real trading data observed in the recent literature, see Frazzini et al. (2018), meaning transaction cost variance estimates are more attainable than transaction costs estimates themselves.

Figure 3: Transaction cost variance parameter estimates

We plot the transaction cost variance model parameter b_1 estimated from June 2019 to March 2023. The estimation is done every month using past two years' trading data. Every point is plotted with the corresponding 95% confidence interval.

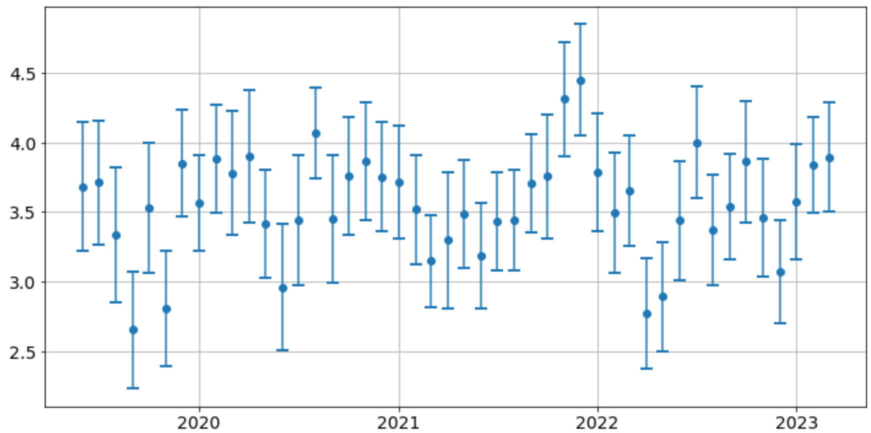

Table 4: Transaction cost variance estimation results

We show the mean, median, standard deviation minimum and maximum of parameter values, t-statistics and R-squared estimated from June 2019 to March 2023. The estimation is done every month using past two years' trading data.

	Mean	Median	Std	Min	Max
b_1	3.55	3.55	0.37	2.66	4.45
t-statistic	17.61	17.64	2.65	12.69	25.25
R^2 (%)	72.35	71.34	4.23	59.35	84.52
p-value	<0.001				

4.3 Simulation performances

A key component of our methodology is the return simulation, as we need to ensure our results are robust to the behaviour of the underlying assets. To strengthen the validity of our return simulations, we show our time series model to be a good fit by performing a series of Ljung-Box tests on our residuals. The number of lags used to conduct the tests is 20. Since our model is estimated in a rolling window (adding and subtracting one day at a time), we will have daily model fits for the entire period. With all this in mind, we present the results of our Ljung-Box tests in three plots: one for the AR(1) series residuals, one for the GARCH(1,1) residual series and one for the AR(1)-DCC(1,1)-GARCH(1,1) residual series in Figure 4. For the AR(1) and GARCH(1,1) residual series, since we have 123 stocks, meaning 123 tests per period, we only plot the 25th, 50th and 75th quantiles of the P-value of our Ljung-Box statistic for each period. For the GARCH(1,1) and AR(1)-DCC(1,1)-GARCH(1,1) series, we standardize residuals using the estimated variances and variance-covariance matrices respectively.

Figure 4: Ljung-Box test P-values

Plots for the P-values of the Ljung-Box statistic based on three models' residuals: AR(1) residual series, GARCH(1,1) residual series and the AR(1)-DCC(1,1)-GARCH(1,1) residual series. We plot the 25th, 50th and 75th quantiles of the P-value of our Ljung-Box statistic for each period for the AR(1) and GARCH(1,1) residual series. For the AR(1)-DCC(1,1)-GARCH(1,1) residual series, as we observe one value daily, we simply plot the P-value over time.

(a) AR(1)

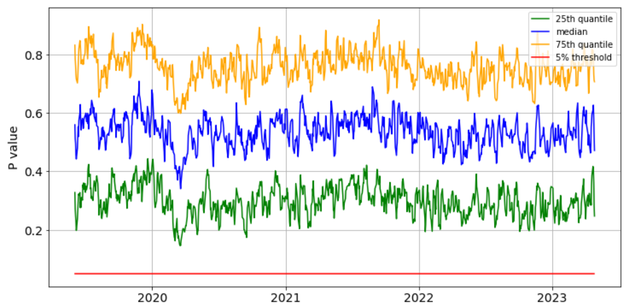

(b) GARCH(1,1)

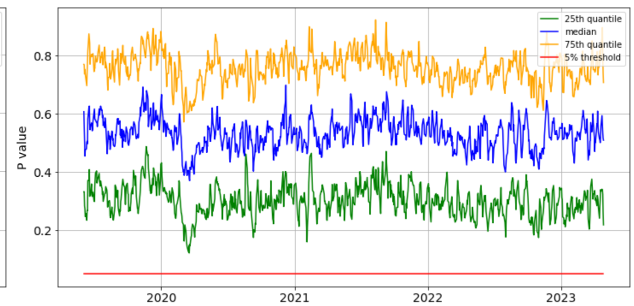

(c) AR(1)-DCC(1,1)-GARCH(1,1)

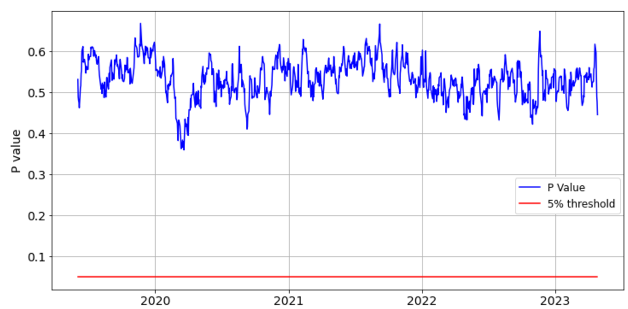

Figure 4 shows that the P-values for all periods are above the 5% threshold, meaning we can not reject the null hypothesis of absence of autocorrelation in our residuals. This gives us confirmation that our model has managed to capture the return data movements successfully.

4.4 Portfolio performances

In this section, we present the performances of our three portfolios. Let P^1, P^2 and P^3 denote portfolios that optimize (3), (4) and (5), respectively. P^1 is then the mean-variance optimal portfolio with no transaction cost consideration, P^2 is the mean-variance optimal portfolio with a transaction cost penalty, and P^3 is the mean-variance optimal portfolio with a transaction cost penalty and a transaction cost variance penalty. We construct these portfolios in three AuM assumptions of \$500M, \$1B and \$2B, as well as three risk aversion parameter values of $\gamma = 1$, $\gamma = 5$ and $\gamma = 10$, resulting in 9 different cases considered. In each case, we construct 100 different portfolios based on 100 simulations of returns. Additionally, each portfolio's net performance is calculated 100 times based on 100 simulations of transaction costs. The summary statistics of our results are reported in Table 5.

To properly understand the behaviour of our portfolios, let us first look into how AuM and risk parameter assumptions affect our portfolios. Changes in AuM impact the portfolios via transaction costs, as larger AuM implies larger transaction costs. Hence, AuM impacts gross performance as we shift our portfolios away from expensive trades and potentially from taking positions that are desirable from a purely return-oriented view. In addition, the effects of AuM are also observable in net performance, where larger AuM assumptions, ceteris paribus, result in higher transaction costs. As the P^1 portfolio does not consider transaction costs in optimization, increasing AuM does not impact the gross performance of the portfolio while increasing transaction costs. The risk aversion parameter γ impacts all the portfolios. Increasing γ results in portfolios with lower variance at the cost of lower returns. Additionally, it impacts the transaction cost variance penalty of portfolio P^3 as we observe a slight but consistent reduction in transaction costs as γ increases. This could be explained by the tendency of portfolios with higher risk aversion parameter to have less concentrated positions, making them less expensive to trade in and out of.

Table 5: Portfolio performances

We show the performance of all three portfolios considered spanning from June 2019 to March 2023. Three AuM assumptions, \$500M, \$1B, \$2B, and three risk aversion parameters $\gamma = 1, 2, 5$ are considered. We report the monthly gross and net return, net cumulative return, maximum drawdown, volatility, Sharpe ratio as well as turnover and associated transaction costs. Reported numbers are averages across all simulations of both returns and transaction costs. Best portfolio denotes the number of times a portfolio had the highest net Sharpe ratio in the 100 return simulations considered.

AuM \$500M	$\gamma = 1$			$\gamma = 5$			$\gamma = 10$		
	P^1	P^2	P^3	P^1	P^2	P^3	P^1	P^2	P^3
Gross return (%)	3.03	2.73	2.66	2.91	2.59	2.54	2.82	2.48	2.46
Volatility (%)	5.45	4.79	4.72	5.15	4.64	4.52	4.85	4.37	4.34
Gross Sharpe ratio	0.556	0.570	0.564	0.565	0.558	0.562	0.581	0.568	0.567
Turnover (%)	62.34	33.24	32.98	59.35	32.25	31.87	56.64	31.19	30.84
Transaction costs (%)	0.77	0.42	0.39	0.75	0.41	0.39	0.72	0.38	0.37
Net return (%)	2.26	2.31	2.27	2.16	2.18	2.15	2.10	2.10	2.09
Net cum. return (%)	176	160	177	178	215	174	146	179	158
Max. drawdown	80	77	77	80	79	76	79	76	75
Net Volatility (%)	5.76	4.93	4.77	5.37	4.87	4.63	5.14	4.55	4.45
Net Sharpe ratio	0.392	0.469	0.476	0.402	0.448	0.464	0.409	0.462	0.470
Best portfolio	2	16	82	1	11	88	0	9	91
AuM \$1B	$\gamma = 1$			$\gamma = 5$			$\gamma = 10$		
	P^1	P^2	P^3	P^1	P^2	P^3	P^1	P^2	P^3
Gross return (%)	3.03	2.36	2.30	2.91	2.25	2.21	2.82	2.23	2.19
Volatility (%)	5.45	4.57	4.53	5.15	4.51	4.46	4.85	4.09	4.04
Gross Sharpe ratio	0.556	0.516	0.508	0.565	0.499	0.496	0.581	0.545	0.542
Turnover (%)	62.34	27.71	27.35	59.35	26.67	26.32	56.64	25.04	24.44
Transaction costs (%)	1.54	0.78	0.74	1.49	0.76	0.73	1.44	0.73	0.71
Net return (%)	1.49	1.58	1.56	1.42	1.49	1.48	1.38	1.50	1.48
Net cum. return (%)	88	133	97	94	133	125	77	116	93
Max. drawdown	80	75	75	80	76	74	79	75	73
Net Volatility (%)	5.93	4.82	4.67	5.73	4.91	4.63	5.16	4.41	4.28
Net Sharpe ratio	0.251	0.328	0.334	0.248	0.303	0.320	0.267	0.340	0.346
Best portfolio	0	23	77	0	14	86	0	28	72
AuM \$2B	$\gamma = 1$			$\gamma = 5$			$\gamma = 10$		
	P^1	P^2	P^3	P^1	P^2	P^3	P^1	P^2	P^3
Gross return (%)	3.03	1.98	1.91	2.91	1.72	1.69	2.82	1.68	1.65
Volatility (%)	5.45	4.32	4.28	5.15	4.38	4.34	4.85	4.33	4.29
Gross Sharpe ratio	0.556	0.458	0.446	0.565	0.393	0.389	0.581	0.388	0.385
Turnover (%)	62.34	22.49	22.18	59.35	20.14	19.84	56.64	19.27	18.77
Transaction costs (%)	3.08	1.16	1.09	2.99	1.10	1.08	2.88	1.05	1.03
Net return (%)	-0.05	0.82	0.82	-0.08	0.62	0.61	-0.06	0.63	0.62
Net cum. return (%)	-5	35	56	-13	36	30	4	31	49
Max. drawdown	80	73	72	80	73	73	79	73	72
Net Volatility (%)	6.43	4.73	4.55	6.03	4.87	4.66	5.23	4.64	4.51
Net Sharpe ratio	-0.01	0.173	0.180	-0.01	0.127	0.131	-0.01	0.136	0.137
Best portfolio	0	13	87	0	18	82	0	33	67

Looking at the underlying portfolio optimization setups, portfolio P^1 optimizes the risk-adjusted return with no considerations to the net performance. This leads to portfolios with high gross return, Sharpe ratio, turnover and transaction costs and low net returns and net Sharpe ratio. We see that the P^1 portfolios observe the highest gross performance for each case considered, but are rarely the "Best portfolio", as it displays the number of times a portfolio has the highest net Sharpe ratio. This number ranges from 0 to 100, as we have 100 return simulations for which we construct portfolios P^1, P^2 and P^3. For instance, P^1 is the best portfolio in only two out of 100 simulations when the assumed AuM is \$500M and γ is 1. In most other instances it is never the best performing portfolio. Portfolio P^2 introduces a transaction cost penalty using our transaction cost model. This results in a reduced gross performance relative to P^1, but decreases turnover and transaction costs which substantially improves the net portfolio performance in most cases. Looking at the AuM \$1B and $\gamma = 5$ case, we see a reduction in gross Sharpe ratio from 0.565 to 0.499. On the other hand, turnover is significantly reduced from 59.35% to 26.67%, resulting in lower transaction costs from 1.49% to 0.76%. The final result is a large increase in net Sharpe ratio from 0.248 to 0.303 further evidenced by P^1 never being the best performing portfolio and P^2 being it 14 times. Portfolio P^3 introduces a transaction cost variance penalty that has a slight (and directionally inconsistent) impact on the gross performance of the portfolio. Transaction costs and turnover are slightly reduced as a result of an additional penalty on transaction costs, and the resulting net returns and net variance are also reduced. However, the resulting net Sharpe ratio is increased in most instances, best depicted by the "Best portfolio" measure. Looking at maximum drawdown, we can see it is somewhat stable across all portfolios, keeping within the 75-80 range. As expected, raising gamma, as well as reducing turnover, which, in our case reduces volatility going across the three different portfolios, slightly reduces drawdown. In the case where AuM is \$1B and $\gamma = 5$, we see a small decrease in gross Sharpe ratio from 0.499 to 0.496. Turnover is slightly lowered from 26.67% to 26.32% resulting in comparably lower transaction costs from 0.76% to 0.73%. The resulting net returns are slightly lower from 1.49% to 1.48%. Despite this, the resulting net Sharpe ratio observes a noticeable increase from 0.303 to 0.320 as a consequence of the drop in net volatility going from 4.91% to 4.63%. Maximum drawdown remains stable, reducing going from portfolio P^1 to P^3, and with higher gammas. Looking at net cumulative return, we can see it closely resembles the net return, dropping in magnitude across larger AuM and gammas, and performing best in portfolio P^2. There are a few subtle conclusions to be drawn from these results. Going from P^2 to P^3, the additional penalty to transaction cost variance reduces transaction costs in an unfavourable manner, which results in lower net returns. However, the resulting portfolio P^3 observes considerably less volatile transaction costs, which in turn reduce the net volatility of P^3. Finally, this results in an increase in net Sharpe ratio as we see that P^3 is the best performing portfolio in 86 cases. This means we manage to achieve superior net risk-adjusted performance by reducing the volatility of our trades, which is exactly what the additional penalty seeks to do.

5 Conclusion

This paper set out to refine mean-variance portfolio optimization by incorporating more realistic transaction cost considerations. We addressed both the expected level of transaction costs and their uncertainty across time, ultimately proposing an integrated framework that models the covariance of transaction costs. Our work was motivated by two main observations in the empirical literature. First, while the majority of portfolio optimization studies incorporate transaction costs in only a rudimentary or deterministic manner, real-world data consistently exhibit significant variability in these costs. Second, by leveraging a proprietary dataset of 38,250 trades from a large institutional asset manager, we observe that variance of transaction costs can be large enough to meaningfully alter trading decisions and portfolio allocations.

To evaluate our framework, we compared three portfolios that differ in how they handle transaction costs: a baseline mean-variance portfolio with no explicit transaction cost term, a mean-variance portfolio that penalizes the expected transaction costs, and a mean-variance portfolio that penalizes both the expected and variance of transaction costs. We used the proprietary trade dataset to build models for both the expected transaction costs and their variance-covariance structure. We show that transaction cost variance is significant and can be estimated accurately. In accordance with Min et al. (2022), we also constructed transaction cost covariates, finding that they are statistically significant and exert a meaningful impact on optimal portfolio weight selections.

We then conducted a simulation study based on a multivariate time-series model with parameters estimated using historical data. Through this process, we constructed the three portfolios and compared their performances out-of-sample. The findings reveal that the portfolios incorporating transaction cost information, particularly the one that includes the variance-covariance of transaction costs, outperform the basic mean-variance portfolio on a net-risk-adjusted basis in most cases. In other words, although explicitly penalizing transaction cost variance and covariance might lead to a moderate reduction in gross returns (due to slightly more constrained trading), this trade-off proves beneficial once the volatility of actual trading costs are accounted for. The higher net Sharpe ratios of these cost-aware portfolios highlight their superiority in balancing returns against overall risk, which now includes not just the covariance of asset returns but also that of transaction costs.

The main takeaway is twofold. First, transaction costs are clearly volatile, as argued by Frazzini et al. (2018) and much of the empirical transaction cost literature, and it is therefore a major simplification to treat them as a static, one-size-fits-all deduction from returns. Second, explicitly modeling this volatility in the portfolio construction process translates into more realistic and robust allocation decisions. Including the variance and covariance of transaction costs leads to improved net risk-adjusted performance, providing a compelling case for portfolio managers who seek to implement strategies that maintain a competitive edge in real-world settings.

References

Almgren, Robert, Chee Thum, Emmanuel Hauptmann, and Hong Li, 2005, Direct estimation of equity market impact, *Journal of Risk* 3, 5–39.

Breen, William J., Laurie Simon Hodrick, and Robert A. Korajczyk, 2002, Predicting equity liquidity, *Management Science* 48, 470–483.

Davis, Mark H. A., and Andrew R. Norman, 1990, Portfolio selection with transaction costs, *Mathematics of Operations Research* 15, 676–713.

de Groot, Wilma, Joop Huij, and Weili Zhou, 2012, Another look at trading costs and short-term reversal profits, *Journal of Banking & Finance* 36, 371–382.

de Rossi, Giuliano, Eliad Hoch, and Michael Steliaros, 2022, The cost of trading factor strategies, *Journal of Systematic Investing* 2, 70–87.

DeMiguel, Victor, Alberto Martín-Utrera, and Francisco J. Nogales, 2015, Parameter uncertainty in multiperiod portfolio optimization with transaction costs, *Journal of Financial and Quantitative Analysis* 50, 1443–1471.

Frazzini, Andrea, Ronen Israel, and Tobias J. Moskowitz, 2018, Trading costs, Working paper.

Gabaix, X., P Gopikrishnan, V Plerou, and HE Stanley, 2006, Institutional investors and stock market volatility, *Quarterly Journal of Economics* 121, 461–504.

Glosten, Lawrence R., and Lawrence E. Harris, 1988, Estimating the components of the bid/ask spread, *Journal of Financial Economics* 21, 123–142.

Kissell, Robert, 2014, The science of algorithmic trading and portfolio management, Chapter 5 - Estimating I-Star Model Parameters, *Academic Press* 163–191.

Korajczyk, Robert A., and Ronnie Sadka, 2004, Are momentum profits robust to trading costs?, *Journal of Finance* 59, 1039–1082.

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

Loeb, Thomas F., 1983, Trading cost: The critical link between investment information and results, *Financial Analysts Journal* 39, 39–44.

Magill, Michael J.P., and George M. Constantinides, 1976, Portfolio selection with transactions costs, *Journal of economic theory* 13, 245–263.

Markowitz, Harry, 1952, Portfolio selection, *The Journal of Finance (New York)* 7, 77–91.

Markowitz, Harry, 1959, *Portfolio selection: Efficient Diversification of Investments* (Yale University Press).

Min, Seungki, Costis Maglaras, and Ciamac C. Moallemi, 2022, Cross-sectional variation of intraday liquidity, cross-impact, and their effect on portfolio execution, *Operations Research* 70, 830–846.

Perold, André F., 1988, The implementation shortfall, *Journal of Portfolio Management* 14, 4–9.

Torre, Nicolo G., and M. Ferrari, 2000, The market impact model (BARRA Inc, Berkeley).

A Appendix

A.1 ARMA parameter estimation

The general $ARMA(p, q)$ model is given by

$$r_t = c + \sum_{i=1}^{p} \phi_i r_{t-i} + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j} + \varepsilon_t, \tag{A1}$$

where r_t is the return at time t, ε_t the error term at time t and c a constant. ϕ_i and θ_j are the parameters of the model. The goal here is to estimate the given parameters, including c, using maximum likelihood methods. We assume the error term to be normally distributed.

A.1.1 Maximum likelihood

Assuming r_t can be modelled using an $ARMA(p, q)$ model we have,

$$r_t = c + \sum_{i=1}^{p} \phi_i r_{t-i} + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j} + \varepsilon_t. \tag{A2}$$

Rearranging gives us:

$$\varepsilon_t = r_t - c - \sum_{i=1}^{p} \phi_i r_{t-i} - \sum_{j=1}^{q} \theta_j \varepsilon_{t-j}. \tag{A3}$$

Setting $\tilde{r}_t = r_t - \tilde{c}$ where

$$\tilde{c} = \frac{c}{1 - \sum_{i=1}^{p} \phi_i}, \tag{A4}$$

gives us:

$$\varepsilon_t = \tilde{r}_t - \sum_{i=1}^{p} \phi_i \tilde{r}_{t-i} - \sum_{j=1}^{q} \theta_j \varepsilon_{t-j}. \tag{A5}$$

Now we move on to the joint density function and calculating the likelihood function.
Given $\varepsilon_t \overset{i.i.d}{\sim} N(0, \sigma^2)$, the joint density function for ε is:

$$p(\varepsilon_1, \ldots, \varepsilon_n) = \frac{1}{(2\pi)^{n/2} \sigma^n} \exp\left\{ -\frac{1}{2\sigma^2} \sum_{t=1}^{n} \varepsilon_t^2 \right\}. \tag{A6}$$

Assuming a natural filtration $\mathcal{F}$ contains p observations of r_t before the series started, q noises of ε_t before the series started and all returns, we have:

$$L\Big(\phi, \theta, c, \sigma^2\Big) = p(r_1, \ldots, r_n|\mathcal{F}) = \frac{1}{(2\pi)^{n/2}\sigma^n} \exp\Bigg\{-\frac{1}{2\sigma^2}\sum_{t=1}^{n}\varepsilon_t(\phi, \theta, c|\mathcal{F})^2\Bigg\}, \tag{A7}$$

which we have to maximize with respect to the given parameters ϕ, θ, c in order to obtain their values.

A.2 GARCH parameter estimation

The general GARCH(p,q) model is given by

$$r_t = \mu_t + \epsilon_t, \tag{A8}$$

$$\sigma_t^2 = \omega + \sum_{j=1}^{p}\beta_j\sigma_{t-j}^2 + \sum_{i=1}^{q}\alpha_i\epsilon_{t-i}^2, \tag{A9}$$

$$\epsilon_t = \sigma_t e_t, \quad e_t \sim N(0,1), \tag{A10}$$

where r_t is the return at time t, ε_t the error term at time t and ω a constant. α_i and β_j are the parameters of the model. The goal here is to estimate the given parameters, including ω, using maximum likelihood methods. We assume the error term to be normally distributed.

A.2.1 Maximum likelihood

Using a similar procedure as in the ARMA process, we obtain the GARCH likelihood to have the form of

$$L(\alpha, \beta, \omega) = p(r_1, \ldots, r_n|\mathcal{F}) = \frac{1}{(2\pi)^{n/2}\prod_{t=1}^{n}\sigma_i} \exp\Bigg\{-\frac{1}{2}\sum_{t=1}^{n}\frac{\varepsilon_t^2}{\sigma_t^2}(\alpha, \beta, \omega)\Bigg\}, \tag{A11}$$

which we have to maximize with respect to the given parameters α, β, ω in order to obtain their values.

A.3 DCC-GARCH parameter estimation

After estimating the individual GARCH(1,1) parameters, we continue by estimating the covariance matrix. The general DCC(m,n)-GARCH(p,q) model is given by

$$\mathbf{r_t} = \mu_\mathbf{t} + \epsilon_\mathbf{t}, \tag{A12}$$

$$\epsilon_{\mathbf{t}} = \mathbf{H}_{\mathbf{t}}^{\frac{1}{2}} \mathbf{z}_{\mathbf{t}}, \tag{A13}$$

$$\mathbf{H}_{\mathbf{t}} = \mathbf{D}_{\mathbf{t}} \mathbf{R}_{\mathbf{t}} \mathbf{D}_{\mathbf{t}}, \tag{A14}$$

where $\mathbf{r_t}$ is the $n \times 1$ return vector, $\mu_{\mathbf{t}}$ the $n \times 1$ expected return vector, $\epsilon_{\mathbf{t}}$ the $n \times 1$ residual vector with a covariance matrix H_t and $\mathbf{z_t}$ a $n \times 1$ vector of standard iid error terms (in our case, standard normal random variables). D_t is a diagonal matrix of conditional standard deviations of ϵ_t obtained using the GARCH(p,q) model and R_t is its correlation matrix.

$$D_t = \begin{bmatrix} \sigma_{1t} & 0 & \cdots & 0 \\ 0 & \sigma_{2t} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sigma_{nt} \end{bmatrix}, \tag{A15}$$

$$R_t = \begin{bmatrix} 1 & \rho_{1,2,t} & \cdots & \rho_{1,n,t} \\ \rho_{1,2,t} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \rho_{n-1,n,t} \\ \rho_{1,n,t} & \cdots & \rho_{n-1,n,t} & 1 \end{bmatrix}. \tag{A16}$$

Furthermore, we define

$$\mathbf{R}_{\mathbf{t}} = \mathbf{Q}_{\mathbf{t}}^{*-\mathbf{1}} \mathbf{Q}_{\mathbf{t}} \mathbf{Q}_{\mathbf{t}}^{*-\mathbf{1}}, \tag{A17}$$

$$\mathbf{Q}_{\mathbf{t}} = \left(1 - \sum_{i=1}^{m} a_i - \sum_{j=1}^{n} b_j\right) \bar{\mathbf{Q}} + \sum_{i=1}^{m} a_i u_{t-i} u_{t-i}^T + \sum_{j=1}^{n} b_j \mathbf{Q}_{\mathbf{t}-\mathbf{j}}, \tag{A18}$$

where

$$\bar{\mathbf{Q}} = Cov\left[u_t u_t^T\right] = \mathbb{E}\left[u_t u_t^T\right] = \frac{1}{T} \sum_{t=1}^{T} u_t u_t^T, \tag{A19}$$

$\mathbf{Q}_{\mathbf{t}}^{*}$ is a diagonal matrix with the square root of the diagonal elements of $\mathbf{Q}_{\mathbf{t}}$ at the diagonal

$$\mathbf{Q}_{\mathbf{t}}^{*} = \begin{bmatrix} \sqrt{q_{11t}} & 0 & \cdots & 0 \\ 0 & \sqrt{q_{22t}} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sqrt{q_{nnt}} \end{bmatrix}, \tag{A20}$$

and u_t are the standardized errors

$$u_{it} = \frac{\epsilon_{it}}{\sigma_{it}}. \tag{A21}$$

A.3.1 Maximum likelihood

Assuming our error terms $\mathbf{z_t}$ follow a multivariate Gaussian distribution, the joint distribution of $\mathbf{z_1}, \ldots, \mathbf{z_T}$ becomes

$$f(\mathbf{z_t}) = \prod_{t=1}^{T} \frac{1}{(2\pi)^{\frac{n}{2}}} \exp\left\{-\frac{1}{2}\mathbf{z_t^T}\mathbf{z_t}\right\}. \tag{A22}$$

Using this, we can obtain the likelihood function for ϵ_t as

$$L(\theta) = \prod_{t=1}^{T} \frac{1}{(2\pi)^{n/2}|H_t|^{1/2}} \exp\left\{-\frac{1}{2}\epsilon_{\mathbf{t}}{}^T H_t^{-1}\epsilon_{\mathbf{t}}\right\}. \tag{A23}$$

where θ is the set of all parameters. Having already estimated the GARCH model parameters, we are left with estimating the DCC parameters hence $\theta = \{a_1, \ldots, a_m, b_1, \ldots, b_n\}$. We continue with the expression for the log likelihood.

$$\begin{aligned}
\ln(L(\theta)) &= -\frac{1}{2}\sum_{t=1}^{T}\Big(\mathrm{nln}(2\pi) + \ln(|H_t|) + \epsilon_{\mathbf{t}}{}^T H_t^{-1}\epsilon_{\mathbf{t}}\Big) \\
&= -\frac{1}{2}\sum_{t=1}^{T}\Big(\mathrm{mln}(2\pi) + \ln(|D_t R_t D_t|) + \epsilon_{\mathbf{t}}{}^T D_t^{-1} R_t^{-1} D_t^{-1}\epsilon_{\mathbf{t}}\Big) \\
&= -\frac{1}{2}\sum_{t=1}^{T}\Big(n\ln(2\pi) + 2\ln(|D_t|) + \ln(|R_t|) + \epsilon_{\mathbf{t}}{}^T D_t^{-1} R_t^{-1} D_t^{-1}\epsilon_{\mathbf{t}}\Big) \\
&= -\frac{1}{2}\sum_{t=1}^{T}\Big(n\ln(2\pi) + 2\ln(|D_t|) + \ln(|R_t|) + u_t^T R_t^{-1} u_t\Big).
\end{aligned} \tag{A24}$$

Since D_t is constant having estimated the GARCH parameters, our problem becomes equivalent to maximizing

$$\ln(L^*(\theta)) = -\frac{1}{2}\sum_{t=1}^{T}\Big(\ln(|R_t|) + u_t^T R_t^{-1} u_t\Big). \tag{A25}$$