Optimizing with Transaction Costs

Filip Bašić

Department of Accounting and Finance

Lancaster University

Submitted in fulfillment of the requirements

for the degree of

Doctor of Philosophy in Finance

August 2024

Abstract

This thesis explores the integration of transaction costs into portfolio optimization models, addressing a significant gap in traditional financial theory where such costs are often overlooked or simplified. By treating transaction costs as dynamic and volatile elements, this research enhances the realism and practical applicability of portfolio construction, leading to more efficient and effective investment strategies. The work is structured across three main chapters. Chapter 1 introduces the concept of implementation shortfall variance and integrates transaction cost covariance into the mean-variance utility function. Through statistical modeling and simulations using S&P 500 data, it demonstrates how accounting for transaction cost variance can improve portfolio performance. Chapter 2 extends the analysis to equity factor portfolios, a critical component in asset management. It presents a novel optimization framework that reduces transaction costs while preserving the core characteristics of these portfolios. By applying this methodology to both single and multi-factor portfolios across global markets, the research shows that transaction-cost-optimized portfolios can achieve superior net returns, particularly in less liquid markets. Chapter 3 continues this approach to develop factor-enhanced market portfolios optimized for transaction costs. By integrating these factor strategies within a benchmark-relative context, the research provides actionable insights for asset managers, suggesting that traditional models which overlook or simplify transaction costs may be suboptimal.

Overall, this thesis makes significant contributions to the field of finance by emphasizing the importance of transaction costs in portfolio optimization and offering practical tools and models that can enhance portfolio performance in a realistic and cost-effective manner.

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and acknowledgments.

Filip Bašić
August 2024

Acknowledgements

Having reached the final steps of completing the daunting task a PhD is, I would have to describe this journey as challenging in variety of ways but also rewarding. Throughout this process, I have been fortunate to be surrounded by incredible individuals who have provided unwavering support and assistance.

To begin with, I would like to thank my wonderful group of supervisors, Prof. Alberto Martín-Utrera (Iowa State University), Dr. Harald Lohre (Robeco), Prof. Sandra Nolte and Prof. Ingmar Nolte (Lancaster University) for their unwavering support and belief in me, continuous feedback and insightful comments. Beginning with a modest foundation in Finance, both practically and academically, their expertise and patience created an environment where I was able to learn and develop.

I would also like to thank Invesco, Robeco and Lancaster University Management School for their generous financial support over the years, which funded my PhD. I would also like to express my gratitude for providing access to your facilities, which have greatly supported my PhD research. I am also very grateful to all the colleagues I have met here. There are too many wonderful people to mention, but without their help and selflessness, none of this would be possible.

I would also like to thank all friends I have met along the way for all the interesting conversations and ideas as well as fun times that made the whole process a lot more enjoyable.

Finally, I would like to thank my friends and family I left behind for their unconditional love and support throughout the years.

Table of Contents

List of Tables

List of Figures

Introduction

In the domain of portfolio management and financial optimization, the inclusion and impact of transaction costs have long been acknowledged, see Magill and Constantinides (1976); Davis and Norman (1990), yet they are often treated as peripheral concerns in classical optimization models. This thesis addresses this gap by integrating transaction costs directly into the optimization process, thereby enhancing the realism and applicability of portfolio construction models. By addressing transaction costs as dynamic and volatile components rather than fixed or negligible quantities, this research seeks to optimize portfolios in a manner that is both theoretically sound and practically relevant.

This thesis is structured into three chapters, each building upon the last to develop a comprehensive understanding of how transaction costs influence portfolio optimization and how these costs can be effectively managed across different financial strategies and market conditions.

Chapter 1: The Risk of Falling Short - Implementation Shortfall Variance in Portfolio Construction

The first chapter delves into the variance of transaction costs, particularly focusing on the implementation shortfall, a key metric for evaluating the real cost of executing trades defined in Perold (1988). Building on the foundational works of Markowitz (1952) on portfolio selection and risk diversification, this chapter broadens the mean-variance utility function with transaction costs seen in the works of Frazzini, Israel, and Moskowitz (2018) and Rossi, Hoch, and Steliaros (2022) to incorporate transaction cost volatility. This chapter estimates the variance and covariance of transaction costs, thereby providing a more comprehensive utility function that better captures the true risk-adjusted performance of portfolios. The practical implications are demonstrated through a multivariate simulation study, utilizing data from S&P 500 stocks, which shows that considering transaction cost variance can significantly improve portfolio outcomes.

Chapter 2: Transaction Cost-Optimized Equity Factors Around the World

Expanding on the insights gained in the first chapter, the second chapter investigates the impact of transaction costs on equity factor portfolios, which are widely used in asset management to capture specific risk premiums. This chapter presents a novel optimization framework that integrates transaction cost considerations into the construction of factor portfolios, thereby reducing costs while maintaining the desired factor exposures. Drawing on the methodologies of Novy-Marx and Velikov (2016), Frazzini, Israel, and Moskowitz (2018) and Rossi, Hoch, and Steliaros (2022), this chapter demonstrates how transaction-cost-optimized portfolios can achieve superior net returns without compromising their integrity. The analysis spans both single and multi-factor portfolios across global markets, illustrating that transaction cost considerations are critical for maintaining profitability, especially in less liquid markets such as those in Europe and Emerging markets.

Chapter 3: Transaction Cost-Optimal Factor-Enhanced Market Portfolios

The final chapter expands on the second chapter by developing factor-enhanced market portfolios optimized for transaction costs. Building on the factor models proposed by Fama and French (1993) and the subsequent enhancements by Asness, Frazzini, and Pedersen (2019), this chapter integrates these factor strategies within a market portfolio context, evaluating their performance when transaction costs are explicitly accounted for. By applying the transaction cost models developed in the earlier chapters, this chapter demonstrates that factor-enhanced portfolios can achieve a more efficient balance between cost and return. The practical implications for asset managers are profound, suggesting that traditional benchmark-relative models which overlook or simplify transaction costs may be suboptimal. The research thus provides actionable insights into constructing portfolios that are both cost-effective and return-efficient.

Through these chapters, this thesis makes significant contributions to the field of finance by not only highlighting the importance of transaction costs in portfolio optimization but also by offering practical tools and models that asset managers can use to enhance portfolio performance. The findings challenge the traditional approaches to portfolio construction, advocating for a more nuanced consideration of transaction costs as an integral component of the investment decision-making process.

Chapter 1

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

This project is joint work with my supervisors Alberto Martín-Utrera, Harald Lohre, Sandra Nolte and Ingmar Nolte. We thank 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.

1.1 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 et al., 2006 further develop square root price impact, basing their transaction cost model on Torre and Ferrari, 2000. This provides the theoretical framework for Almgren et al., 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, 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, Maglaras, and Moallemi, 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, Israel, and Moskowitz, 2018 and Rossi, Hoch, and Steliaros, 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 1.2 describes the data used for both the transaction cost modelling and time-series model estimation. Section 1.3 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 1.3 concludes with portfolio optimization and evaluation of the performance of the portfolios. Section 1.4 contains the estimation results for the transaction cost model, transaction cost variance model, as well as portfolio performance. Our conclusions are presented in Section 1.5.

1.2 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 1.1 reports the summary statistics on returns, trading volume and market capitalization across all 123 stocks.

Table 1.1: 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:

$$r_t^{HL} = \ln(H_t) - \ln(L_t), \tag{1.1}$$

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

Table 1.2 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 1.3.1. The median trade size observed is 0.16% of median daily volume (MDV), similar to what Frazzini, Israel, and Moskowitz, 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 1.2: 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 1.6, 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

1.3 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:

$$R^p_{t+1} = w_t^\top R_{t+1}, \tag{1.2}$$

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

$$\max_{w_t} \mathbb{E}_t[w_t^\top R_{t+1}] - \frac{\gamma}{2} Var_t[w_t^\top R_{t+1}], \tag{1.3}$$

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:

$$\max_{w_t} \mathbb{E}_t[w_t^\top R_{t+1}] - \frac{\gamma}{2} Var_t[w_t^\top R_{t+1}] - \mathbb{E}_t[\Delta w_t^\top TC_t(\Delta w_t, \theta)], \tag{1.4}$$

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 (1.4) to include a penalty term for the variance of transaction costs

$$\max_{w_t} \mathbb{E}_t[w_t^\top R_{t+1}] - \frac{\gamma}{2} Var_t[w_t^\top R_{t+1}] - \mathbb{E}_t[\Delta w_t^\top TC_t(\Delta w_t, \theta)] - \frac{\gamma}{2} Var_t[\Delta w_t^\top TC_t(\Delta w_t, \theta)], \tag{1.5}$$

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.

1.3.1 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),

$$IS = MI + OC = \sum_{k=1}^{T} q_k(p_k - p_{start}) + (n - m)(p_{end} - p_{start}), \tag{1.6}$$

where

$$m = \sum_{k=1}^{T} q_k. \tag{1.7}$$

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.

1.3.2 Modelling transaction costs

Transaction cost modelling has gained traction in recent times, ranging from (effective) bid-ask spreads, see 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, Israel, and Moskowitz, 2018 and Rossi, Hoch, and Steliaros, 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:

$$TC_t\left(\Delta w_{i,t}\right) = a_1 \sigma_{i,t} \frac{AuM \times \Delta w_{i,t}}{MDV_{i,t}} + \epsilon_{i,t}, \tag{1.8}$$

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.

1.3.3 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:

$$Var_t\left[TC_t\left(\Delta w_{i,t}\right)\right] = b_1 \sigma_{i,t} \frac{AuM \times \Delta w_{i,t}}{MDV_{i,t}} + \eta_{i,t}. \quad (1.9)$$

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.

1.3.4 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.

$$\hat{\varepsilon}_{i,w} = \frac{\sum_{t \in w} \hat{\varepsilon}_{i,t}}{\mathbf{card}(\{\hat{\varepsilon}_{i,t} : t \in w\})}. \tag{1.10}$$

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

$$u_{i,w} = \frac{\hat{\varepsilon}_{i,w}}{\sqrt{Var(\hat{\varepsilon}_{i,w} : w \in \{1, ..., N_w\})}}. \tag{1.11}$$

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

$$R_t^{TC} = U_t U_t^T. \tag{1.12}$$

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 $\hat{R}_t^{TC}$:

$$\hat{\Sigma}_t^{TC} = \hat{D}_t^{TC} \hat{R}_t^{TC} \hat{D}_t^{TC}, \tag{1.13}$$

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 1.9.

1.3.5 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:

$$\hat{\mu}_d = \hat{c}_d + \hat{\phi}_d r_{d-1}, \tag{1.14}$$

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.

$$\hat{\varepsilon}'_0 = r_{d-250} - \hat{\mu}_{d-250}, \tag{1.15}$$

$$\hat{\sigma}'^2_0 = \frac{\hat{\omega}_d}{1 - \hat{\alpha}_d - \hat{\beta}_d}, \tag{1.16}$$

$$\hat{\varepsilon}'_k = r_{d-250+k} - \hat{\mu}_{d-250+k}, \tag{1.17}$$

$$\hat{\sigma}'_k = \left(\hat{\omega}_d + \hat{\alpha}_d \left(\hat{\varepsilon}'_{k-1}\right)^2 + \hat{\beta}_d \left(\hat{\sigma}'_{k-1}\right)^2\right)^{\frac{1}{2}}, \tag{1.18}$$

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

$$\hat{\sigma}_d = \hat{\sigma}'_{250}. \tag{1.19}$$

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:

$$\hat{\Sigma}_d = \hat{D}\hat{R}\hat{D}, \tag{1.20}$$

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 $\hat{\Sigma}_d$,

$$r_d \sim \mathcal{N}\left(\hat{\mu}_d, \hat{\Sigma}_d\right). \tag{1.21}$$

We use Equation 1.21 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 1.1. 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 1.1: 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)**.

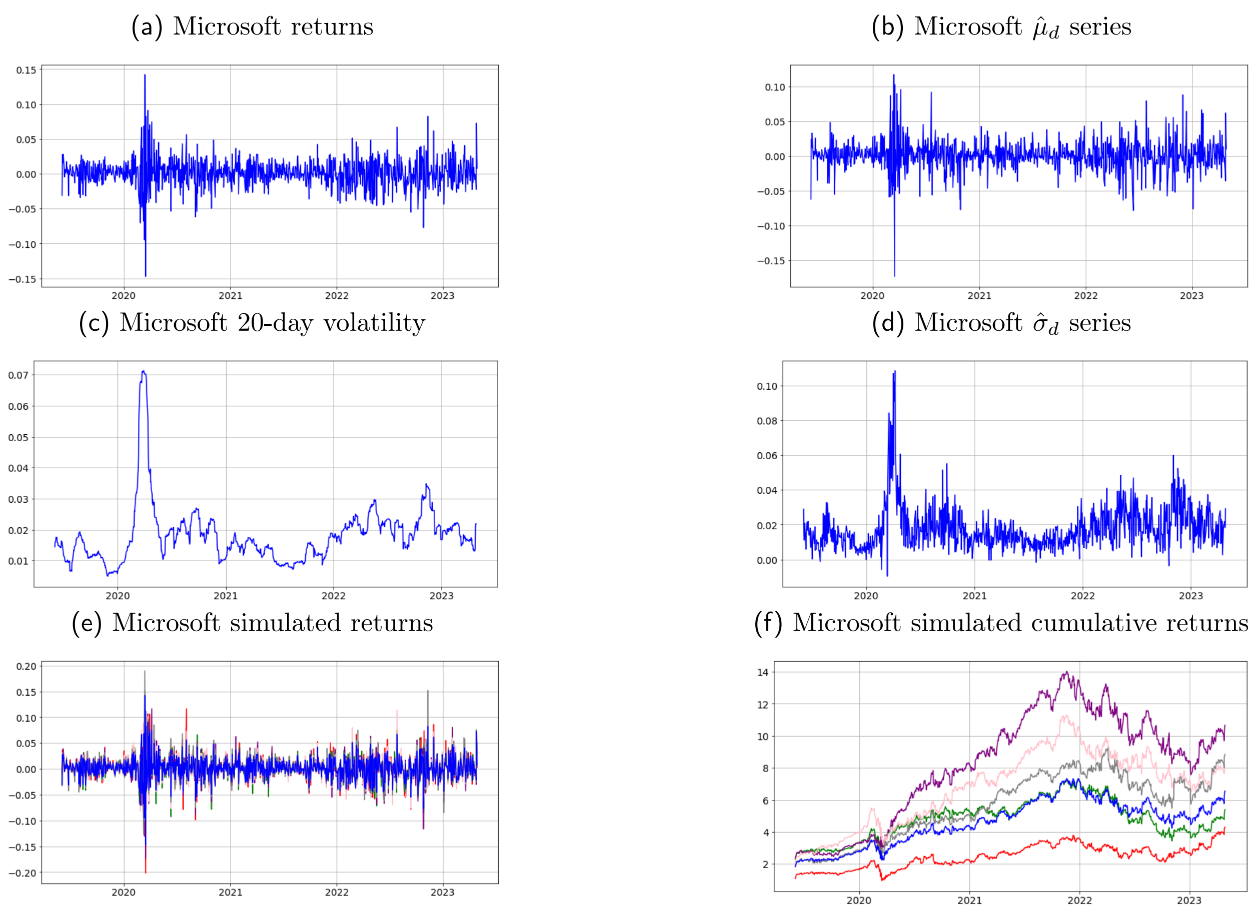

1.3.6 Expected variance

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.

1.3.7 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 1.3, 1.4 and 1.5 respectively.

1.3.8 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.

1.3.9 Performance evaluation

To compare the performance of our portfolios generated by different optimization problems, see Equations 1.3, 1.4 and 1.5, 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:

$$TC_t = \mathcal{N}\left(\widehat{TC}_t(\Delta w_t, \theta), \hat{\Sigma}_t^{TC}\right). \tag{1.22}$$

The mean will be given by the estimated transaction cost function following Equation 2.8, whereas $\hat{\Sigma}_t^{TC}$ is estimated following Equation 1.13 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.

1.4 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.

1.4.1 Transaction cost function parameters

In Figure 1.2, 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.

Figure 1.2: 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.

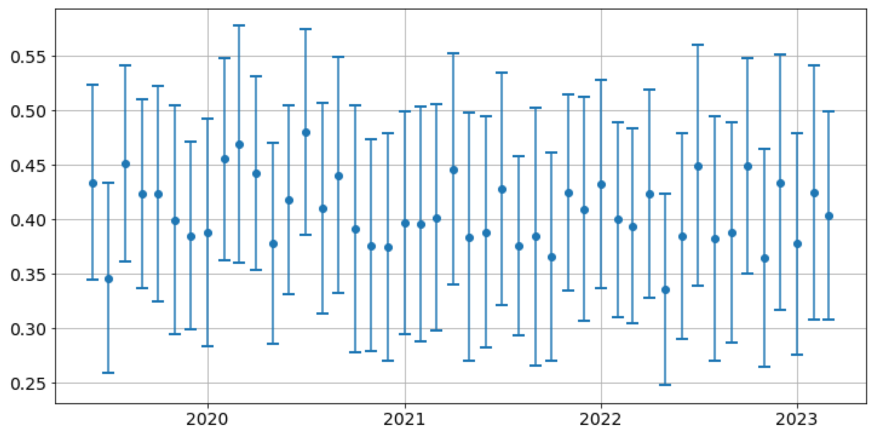

As shown in Table 1.3, the parameter is highly significant with an average t-stat of 8.24. The mean R-squared is 3.6% which is expected as transaction cost data is very volatile, see Frazzini, Israel, and Moskowitz, 2018. This further reinforces the importance of considering transaction cost variance.

Table 1.3: 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				

1.4.2 Transaction cost variance parameters

In Figure 1.3, 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 is stable across time, ensuring consistency of optimization and transaction cost simulations across periods.

Figure 1.3: 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.

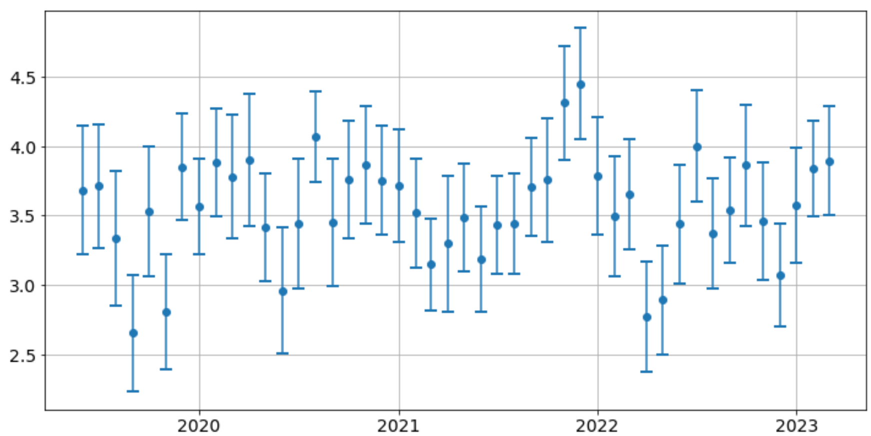

As shown in Table 1.4, b_1 has a mean of 3.55. Since the models we use are identical, we see that the variance of transaction costs are about an order of magnitude larger than transaction costs themselves. This ensures that adding a transaction cost variance covariance matrix penalty in our portfolio 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, Israel, and Moskowitz, 2018, meaning transaction cost variance estimates are more attainable than transaction costs estimates themselves.

Table 1.4: 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				

1.4.3 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 1.4. 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 1.4 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.

Figure 1.4: 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)

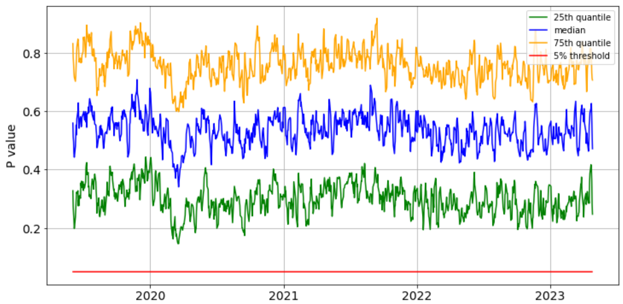

(b) GARCH(1,1)

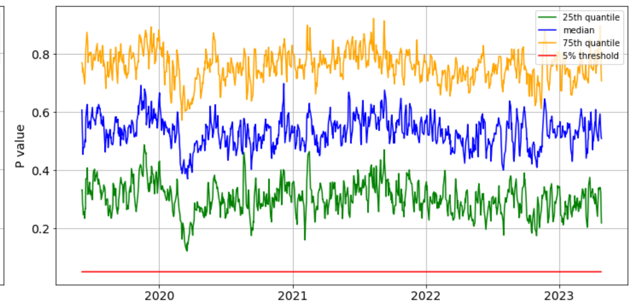

(c) AR(1)-DCC(1,1)-GARCH(1,1)

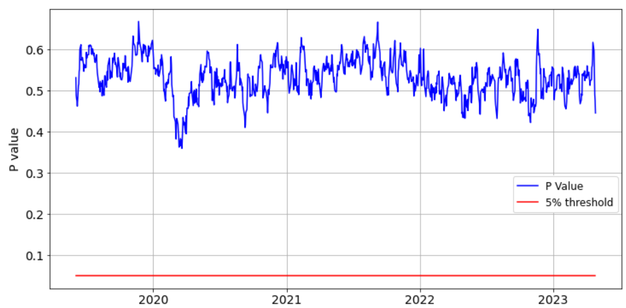

1.4.4 Portfolio performances

In this section, we present the performances of our three portfolios. Let P^1, P^2 and P^3 denote portfolios that optimize (1.3), (1.4) and (1.5), 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 1.5.

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.

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.

1.5 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, Maglaras, and Moallemi, 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, Israel, and Moskowitz, 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.

Table 1.5: 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

Chapter 2

Transaction Cost-Optimized Equity Factors Around the World

This project is joint work with my supervisors Alberto Martín-Utrera, Harald Lohre, Sandra Nolte and Ingmar Nolte. It is published in the Journal of Portfolio Management (*Journal of Portfolio Management*, 50 (6), 40-73). We thank David Blitz, Lars ter Braak, Khuram Chaudhry, Bart van der Grient, Matthias Hanauer, Iman Honarvar, Tobias Hoogteijling, Clint Howard, Maarten Jansen, Stefan Mittnik, Nick Mutsaers, Martin van der Schans, Dmitri Tikhonov, Weili Zhou, Machiel Zwanenburg, participants at the 13th CEQURA Conference on Advances in Financial and Insurance Risk Management in Munich, the J.P. Morgan Macro Quant and Derivatives Conference in London, and the Robeco Quant Research Seminar.

2.1 Introduction

We explore the impact of transaction costs on single and multi-factor equity portfolios, consisting of the six factors momentum, value, quality, size, low volatility and short-term reversal. Our aim is not only to learn about the net profitability of these strategies but to present a portfolio optimization framework for constructing transaction-cost-optimized factor portfolios. The extant literature has mostly focused on producing factors with lower transaction costs, but little has been done to verify if such portfolios retain the desired factor properties. Hence, our framework is centred around the concept of a target factor portfolio and allows tuning how much transaction cost optimization can deviate from the target. The resulting factor portfolios have a better transaction cost profile (how large our trading costs are, are they a result of inefficient trading) and preserve their factor integrity, which we define using a set of measures to monitor the presence of the original factor after TC augmentations. These measures are standard academic measures, including tracking error, active share, transfer coefficient. Novy-Marx and Velikov (2016) argue that many high-turnover factors deliver a return spread net of transaction costs that is not statistically significant. However, by properly taking transaction costs into account when building factor portfolios, transaction costs can be substantially reduced. Although value-weighted portfolios mitigate some of the transaction cost concerns faced by equally weighted portfolios, this route is sub-optimal as market capitalization is not a precise measure of liquidity. Moreover, value-weighted factor portfolios come with high portfolio concentration, which works against harvesting the factor's premium in a diversified manner.

In this light, this paper makes five contributions to the factor investing literature. First, we use factor portfolio trading data from a large institutional asset manager to measure transaction costs. The data used closely resembles similar publications, see for example, Frazzini, Israel, and Moskowitz (2018). Like Frazzini, Israel, and Moskowitz (2018), we show that one can thus more accurately estimate realistic transaction costs, as measured by a trade's implementation shortfall (Perold, 1988). In addition, we consider shorting costs as well as other fixed fees such as tax, foreign exchange and brokerage fees. Indeed, based on these transaction cost estimates we note that the gross benefits of most academic factors are eroded by transaction costs.

Second, we utilize a simple model to optimize the net performance of factor strategies. Our approach to constructing cost-efficient factor portfolios focuses on mean-variance factor portfolio optimization augmented by transaction cost penalties. We thus avoid the expensive trades and excessive turnover that academic factors typically require. Overall, our methodology is flexible with respect to the target portfolios, in terms of both the underlying characteristic and weighting scheme. After optimizing for transaction costs, a high turnover strategy like momentum increases from a net return of −1.52% per annum to 2.18% per annum. For a low turnover strategy like low volatility, the net returns increase from 4.48% to 6.31% per annum. This shows that our approach is not only useful for high turnover strategies, but also for more conservative ones, and one can thus improve net performance across the full spectrum of factor strategies.

Third, the transaction-cost-optimized factor portfolios preserve their factor integrity. For this purpose, we employ a set of measures to monitor how (dis)similar the optimized factor portfolios are relative to the original factors. We find that optimized factors preserve around 80% of their initial factor exposure. Also, the correlation of the optimized factors' portfolio weights with respect to the original factors are above 90% in most cases and the corresponding tracking errors are typically below 2%. Fourth, we explore the factor portfolio performance for different fund size assumptions and document factor portfolios to retain their profitability and factor integrity at different fund sizes.

Lastly, we show that these US results carry over to other regions. In particular, we find that transaction-cost-optimized factors in Europe and Emerging Markets deliver risk-adjusted returns net of transaction costs that are larger than that of the original factors, which are less liquid than those in the US. Considering transaction costs in portfolio construction is even more important for these markets.

The transaction costs literature branches into two segments (Patton and Weller, 2020). The first group resorts to estimating transaction costs using market-wide data; see, for instance, Korajczyk and Sadka (2004) and Novy-Marx and Velikov (2016). The authors use Trade and Quote (TAQ) data to estimate transaction cost model parameters and determine the after-cost performance of relevant asset pricing factors. The second group estimates transaction costs using proprietary data, as performed in the work of Engle, Ferstenberg, and Russell (2012), Frazzini, Israel, and Moskowitz (2018) and Rossi, Hoch, and Steliaros (2022). The goal of our paper coincides with that of Frazzini, Israel, and Moskowitz (2018), however, there are important differences. Our methodology allows us to explore how portfolios are impacted by transaction cost considerations. Exploring a mean-variance optimization setting with a changing transaction cost penalty presents different portfolios which we can analyze to observe the impact of transaction cost considerations on different strategies. In contrast, Frazzini, Israel, and Moskowitz (2018) minimize transaction costs given a tracking error constraint. We find that using a single metric like tracking error for gauging factor integrity can be misleading. Even though tracking error relative to the original factor is an important consideration, we find it to be too dependent on the factor in question. For example, low volatility can retain its factor score as tracking error increases, whereas the momentum factor sees a large decline in factor score as tracking error increases with higher transaction cost penalties. Hence, our methodology allows us to achieve portfolios with varying transaction cost profiles without restrictions on any given factor integrity metric.

Aiming to improve net performance of factor strategies, our work addresses the concerns raised by Chen and Velikov (2023) and Novy-Marx and Velikov (2016). Chen and Velikov (2023) consider transaction costs in single-factor portfolios and show their impact on 204 stock market anomalies. They argue that the average investor should expect tiny net profits from investing in any individual factor. They attribute the lack of net performance to two key factors, the post-publication performance decline and transaction costs. However, as argued in Frazzini, Israel, and Moskowitz (2018) and observed in our data, actual transaction costs are different from what one might infer when using TAQ-based (or similar) estimates.

Furthermore, Novy-Marx and Velikov (2016) focus on mitigating costs of a single-factor at a time. In contrast, we focus on a standard set of factors and propose a methodology to mitigate the cost of both single and multi-factor portfolios. Instead of defining a set of mechanical transaction cost mitigating procedures, we opt to reduce transaction costs by applying a transaction cost penalty in our mean-variance utility. Combining multiple factors helps to mitigate the impact of transaction costs due to trading diversification across factors. In particular, DeMiguel et al. (2020) show that combining characteristics reduces transaction costs since trades in the underlying stocks required to rebalance different factor portfolios can cancel out. As a result, they observe an increase in the number of relevant factor characteristics that explain the cross-section. The multi-factor portfolio framework of DeMiguel et al. (2020) optimizes the factor weights to minimize the impact of transaction costs. Similar to DeMiguel et al. (2020), we also construct transaction-cost optimal multi-factor portfolios, however, we propose a methodology to optimize the stock weights rather than factor weights.

The importance of optimizing factors with respect to transaction costs is further highlighted in the work of Brière et al. (2019) and Groot, Huij, and Zhou (2012). In Brière et al. (2019), single-factor portfolio performances are calculated net of transaction costs. They use institutional trading data to construct both a parametric and a non-parametric transaction cost model. In the case of the parametric transaction cost model, they build on the work of Frazzini, Israel, and Moskowitz (2018), which closely relates to our approach of modeling transaction costs. While we focus on constructing more efficient factor portfolios, Brière et al. (2019) instead focus on calculating the capacities of factor strategies. Groot, Huij, and Zhou (2012) show that short-term reversal strategies' profitability can largely be attributed to excessive trading in small capitalization stocks. These trades generate large transaction costs, diminishing the net profitability of such strategies. They further show that these costs can be reduced by limiting the stock universe to large caps as well as applying a more sophisticated portfolio construction algorithm to lower turnover.

Notably, while the literature agrees that transaction costs are an important consideration, there is no consensus on the profitability of factor portfolios once we account for transaction costs. On the one hand, we have Chen and Velikov (2023) positing that most factor portfolios have very little returns after publication and transaction costs. For instance, Lesmond, Schill, and Zhou (2004) argue that the momentum factor trades disproportionately in expensive stocks, concluding that the magnitude of the abnormal returns associated with momentum are illusory. On the other hand, Korajczyk and Sadka (2004) use non-proportional transaction costs to show that momentum retains profitability after transaction costs. Also, Groot, Huij, and Zhou (2012) manage to construct profitable reversal strategies, and Novy-Marx and Velikov (2016) obtain similar results for different factor portfolios. More recently, Rossi, Hoch, and Steliaros (2022) show that, for realistic portfolio sizes, the typical costs of rebalancing a single-factor portfolio are unlikely to erode the factor premium. Taking all these findings into account, one cannot help but wonder what is causing this apparent dissonance in the literature. To shed light on the matter, we show that transaction-cost-optimized factors deliver important economic gains relative to the original factor without jeopardizing the factor integrity. To achieve this, we develop a methodology of constructing transaction-cost efficient portfolios that is independent of the notion of factor integrity. Combining this with a set of factor integrity measures such as factor score, tracking error, or active share, we can better determine what constitutes a portfolio that is transaction-cost efficient as well as consistent with its factor integrity.

The remainder of this paper is organized as follows. Section 2.2 describes the stock and transaction cost data used to construct our factor portfolios and transaction cost model. Section 2.3 presents the methodology used in creating factor portfolios, estimating a transaction cost model and the optimization procedure. Section 2.4 presents our main results comparing the performances of optimized factor portfolios to their unoptimized counterpart. Section 2.5 compares our findings to Novy-Marx and Velikov (2016). Section 2.6 concludes.

2.2 Factor and transaction cost data

2.2.1 Factor investment universes

Our analysis covers three international equity markets: the United States (US), Europe (EU) and Emerging Markets (EM). Our universe of stocks for the US and Europe are constituents of the Standard and Poor's (S&P) Global Broad Market Index (BMI). This is a broad market index designed to capture the global investable opportunity set by including all listed companies with float-adjusted market capitalization greater than $100 million. This data spans over more than 50 developed and emerging markets, tracking over 99% of each constituent country's available market capitalization. The S&P Global BMI is regionally split into several broad indices. The one used to define our US and European universe is the S&P Developed BMI. The resulting universe consists of 9,834 stocks in the US with a median cross-section of 2,999 stocks dating from January 1980 to March 2023. For Europe, the resulting universe consists of 7,671 stocks with a median cross-section of 2,594 stocks also dating from January 1980 to March 2023.

For emerging markets, we use the constituents of the MSCI Emerging Markets Index. The MSCI Emerging Markets Index captures large and mid capitalization representation across 24 emerging markets countries. With 1,379 constituents as of March 2023, the index covers approximately 85% of the free float-adjusted market capitalization in each country. This results in a total of 3,438 stocks with a median cross-section of 715 stocks dating January 1995 to March 2023. We obtain characteristics, return and volume data from Compustat (for the US) and from Worldscope (for Europe and EM).

2.2.2 Transaction cost data

We utilize transaction cost data from a large institutional asset manager, covering 269,836 factor investment trades over the period from July 2017 to March 2023. The underlying assets are equities traded on various markets. Each trade is executed by the trading desk in multiple executions and the relevant information is then aggregated on all the executions done for the trade. This includes stock identifier and 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 execution is done and the trade is completed.

To measure the cost of executing a certain trade, we refer to the work of Perold (1988) where implementation shortfall (IS) is defined as the sum of execution cost and opportunity cost. Since all trades are fully executed, we can omit the opportunity cost and define implementation shortfall of a single trade as:

$$IS\left(q\right) = \sum_{i=1}^{m} q_i \left(p_i - p_b\right), \tag{2.1}$$

where the total numbers of shares traded, q, is assumed to be executed over m different executions. Each execution trades q_i shares at price p_i. p_b is the initial price of the underlying stock when the trade was submitted. Finally, we calculate the ratio between this cost and the total value of the trade, qp_b, to obtain relative implementation shortfall:

$$IS_{rel} = \frac{IS(q)}{qp_b} \tag{2.2}$$

where qp_b represents the initial value (size) of the trade in USD. We use the term implementation shortfall throughout the text when referring to the relative implementation shortfall.

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:

$$r_t^{HL} = \ln(H_t) - \ln(L_t), \tag{2.3}$$

where H_t and L_t are the highest and lowest prices on day t. We also include spreads in our transaction cost considerations which are sourced from FactSet. Missing spread observations are replaced by the cross-sectional median value.

In Table 2.1, we present the summary statistics for the transaction cost data split across regions. Note that the median trade size is similar to the one reported in Frazzini, Israel, and Moskowitz (2018), with the median traded size for the US being 0.17% of MDV. We also observe that trade size differs substantially across regions, with Europe having larger median MDV of a trade at 0.37%, and EM having the highest at 1.24%. The median transaction cost is also largest for the EM region with a median of 4, 3 and 9 basis points for the US, Europe and EM respectively.

2.3 Factor portfolio construction and transaction costs

In this section, we introduce a mean-variance approach to constructing optimized equity factor portfolios, explaining each of the components of the utility function in detail. Furthermore, we define the transaction cost and shorting cost models used and report the obtained parameters.

Table 2.1: Transaction cost data summary statistics across regions

We present the summary statistics of our trading data across regions. The data consists of 269,836 trades from a large institutional asset manager, covering the period from July 2017 to March 2023. The trading data is fairly evenly distributed across regions, with 99,420 trades for the US, 46,905 trades for Europe and 123,511 for EM. 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) and is reported as a fraction of the original trade value and capped at 10%.

	United States			Europe			Emerging Markets		
	% MDV	Vol (%)	IS (%)	% MDV	Vola (%)	IS (%)	% MDV	Vola (%)	IS (%)
min	0.00	5.15	-9.94	0.00	5.00	-9.51	0.00	5.01	-9.79
q25	0.04	19.50	-0.69	0.09	18.25	-0.36	0.31	20.02	-0.56
Median	0.17	25.75	0.04	0.37	23.64	0.03	1.24	27.43	0.09
q75	0.61	34.40	0.78	1.33	31.25	0.46	4.84	37.73	0.82
max	296.86	100.00	9.86	470.18	99.99	9.13	476.52	99.98	9.99
Mean	0.90	29.04	0.05	1.91	26.52	0.05	4.76	30.56	0.15
SD	3.07	14.06	1.48	6.78	12.35	0.88	10.82	14.58	1.46
# of obs.	99,420			46,905			123,511		

2.3.1 Academic factor portfolios

Assume a market with N stocks and K factors. A single-factor portfolio is a 100% long and 100% short portfolio that is typically obtained in a two-step procedure. First, all stocks considered are ranked by the underlying factor characteristic to identify the top and bottom stocks to populate the long and short leg of the factor portfolio. Second, a weighting scheme is assigned, usually with little regard to more than the ranks and size of the stocks chosen to be included in the portfolio. Common choices would be equal-, value- and characteristic rank-weighted portfolios. As we aim to exploit the entire cross-section, we opt for characteristic rank-weighted single factor portfolios. Characteristic rank-weighted portfolios are typically obtained by sorting the stock universe on the characteristic, assigning a rank to each security, and applying a linear weighting scheme across ranks. Then, define the return of the k-th single factor long-short portfolio at time t, denoted by r_t^k as:

$$r_t^k = w_t^{k\top} r_t = w_{1,t}^k r_{1,t} + w_{2,t}^k r_{2,t} + \ldots + w_{N,t}^k r_{N,t}, \tag{2.4}$$

with

$$\forall k \ \forall t \quad \sum_{i=1}^{N} w_{i,t}^{k} = 0, \ \sum_{i=1}^{N} |w_{i,t}^{k}| = 2, \tag{2.5}$$

where $w_{i,t}^k$ is the weight of stock i in the k-th single-factor long-short portfolio and $r_{i,t}$ the return on stock i at time t. w_t^k is a vector of the k-th single-factor long-short portfolio weights $w_{i,t}^k$, and r_t is the vector of stock returns $r_{i,t}$ at time t.

2.3.2 Mean-variance framework with transaction costs

Having constructed the standard academic single factor portfolios, we introduce a mean-variance optimization framework for constructing the weights of the optimized single-factor long-short portfolios, $w_{i,t}^k$. Define a mean-variance utility maximization problem with transaction costs on a single-factor long-short portfolio as:

$$\underset{w_{i,t}^{k}}{\arg\max} \ \mathbb{E}_{t-1}\left[w_t^{k\top} r_t\right] - \frac{\gamma}{2} Var_{t-1}\left[w_t^{k\top} r_t\right] - \delta \mathbb{E}_{t-1}\left[\Delta w_t^{k\top} TC\left(\Delta w_t^k\right)\right], \tag{2.6}$$

subject to

$$\forall k \ \forall t \quad \sum_{i=1}^{N} w_{i,t}^{k} = 0, \ \sum_{i=1}^{N} |w_{i,t}^{k}| = 2, \tag{2.7}$$

where $\mathbb{E}_{t-1}\left[w_t^{k\top} r_t\right]$ is the expected return of the i-th single-factor portfolio, $Var_{t-1}\left[w_t^{k\top} r_t\right]$ its variance, $\mathbb{E}_{t-1}\left[\Delta w_t^{k\top} TC(\Delta w_t^k)\right]$ the expected transaction costs, Δw_t^k is the difference in weights defined as $\Delta w_t^k = w_t^k - w_{t-1}^k$, γ the risk aversion coefficient, and δ is the transaction cost parameter.

For a multi-factor portfolio, we can define the same standard mean-variance utility maximization problem with transaction costs, but use expected returns, risk and transaction costs for multi-factor portfolios instead. Equations (2.6) and (2.7) defines the entire portfolio construction setup, leaving us with defining the elements of the utility function.

Estimating transaction costs

Transaction cost modeling

Transaction cost modeling has evolved significantly in recent history. Moving away from estimating effective bid-ask spreads and Kyle lambdas, most transaction cost models imply a price impact curve between constant and linear. The price impact curve denotes the change in relative price as a function of quantity bought, usually as a percentage of median daily volume (MDV). We define a transaction cost model following the work best described in Kissell (2014), and the associated I-Star market impact model is widely used in practice today. As a measure of transaction costs we thus consider relative implementation shortfall, that is, implementation shortfall as a fraction of total value of a given trade. Specifically, we estimate a transaction cost model of the following form:

$$TC\left(\Delta w_{i,t}\right) = a_1 \sigma_{i,t} \left(\frac{AuM \times \Delta w_{i,t}}{MDV_{i,t}}\right)^{a_2} + \frac{1}{2} s_{i,t} + \epsilon_{i,t} \tag{2.8}$$

where TC is the relative implementation shortfall, $s_{i,t}$ is the spread, $\sigma_{i,t}$ is the range volatility of stock i at time t, and AuM is the assumption of the assets under management, which we will keep constant and investigate different values. The product $AuM \times \Delta w_{i,t}$ gives us the total amount traded in USD, which we scale by MDV (based on the last 25 days).

Another important aspect of long-short portfolio construction are shorting fees. However, data regarding shorting fees is not publicly available. Cohen, Diether, and Malloy (2007) use proprietary data on shorting fees from a large institutional investor to examine the link between the shorting market and stock prices. For their sample period of September 1999 to August 2003, they show that the mean shorting fee was 2.60% per annum with a median of 1.82%. More recently, Muravyev, Pearson, and Pollet (2022) report that the mean borrow fee is 1.67% per annum with a median of 0.38%, first percentile of 0.25%, 10^{th} percentile of 0.28%, 90^{th} percentile of 3% and 99^{th} percentile of 30% for the sample period of July 2006 to December 2020. Since the sample period of shorting fee data is substantially larger in Muravyev, Pearson, and Pollet (2022), we will use their reported percentiles to fit a model for shorting costs, see the appendix for estimation details.

Transaction cost function in optimization

In order to make the optimization problem computationally more feasible, we estimate both the model given in Equation (2.8) together with a simplified transaction cost model, where we will fix $a_2 = 1$. Hence, the only parameter left to be estimated is a_1 which we denote as b_1 in the case of the simplified transaction cost model. The model in Equation (2.8) will be used to compute the realised transaction costs whereas the simplified model will be used in optimization. For the transaction cost model in Equation (2.8), a_2 is estimated to be 0.67, see Table 2.2. This suggests price impact (the rate at which we move the price with respect to traded amount) to be between square root and linear, often found in the literature. For example, Frazzini, Israel, and Moskowitz (2018) use both a square root term and a linear term in their model. Comparing a_1 to b_1, we see that b_1 is higher which would imply higher costs with a linear price impact, meaning lower transaction costs for small trades, and higher for large trades compared to the model in Equation (2.8). Looking at the R^2, we see a marginal increase from 3.8% to 4.5% in favour of the model in Equation (2.8). Similar to Rossi, Hoch, and Steliaros (2022), we impose a cap on transaction costs by assuming that transaction costs are maximized at MDV= 500%; in other words, trading more than 500% of MDV will not further increase the associated transaction costs.

Expected risk and return inputs

Factor portfolio risk model

As the cross-section of our universe is very large, constructing a variance-covariance matrix for each element is unfeasible. Therefore, we reduce the dimension of our problem using a simple linear factor model. Let

$$r_{i,t} = X_{i,t}^{\top} r_t^f + u_{i,t}, \tag{2.9}$$

describe the factor model where r_t^f is the vector of factor portfolio returns at time t, $r_t^f = \left[r_t^1, r_t^2, \ldots, r_t^K\right]$, $X_{i,t}$ denotes the sensitivity of stock i to factors, and $u_{i,t}$ is the stock specific excess return. Stock variances can then be estimated as

$$Var_{t-1}\left[r_{i,t}\right] = X_{i,t}^{\top} \Sigma_t^f X_{i,t} + U_{i,t}, \tag{2.10}$$

where Σ_t^f is the covariance matrix of factors and $U_{i,t} = Var_{t-1}\left(u_{i,t}\right)$ the specific risk variance

Table 2.2: Transaction cost model estimation

We present the parameters of the transaction cost model in Equation (2.8) and the parameter of the transaction cost model used in optimization. The data used consists of 269,836 trades from a large institutional asset manager, covering the period from July 2017 to March 2023. We estimate the model using nonlinear least squares, and report the coefficients obtained. Panel A gives results for the unrestricted model, and Panel B gives the results for the restricted model with a_1 set to 1.

Panel A: Transaction cost model

Parameter	Value	St. Dev.	t-stat	P-value	95% CI	
a_1	0.35	0.01	23.83	0.00	0.32	0.38
a_2	0.67	0.02	32.73	0.00	0.63	0.72

$R^2(\%)$	χ^2	Red. χ^2
4.5	38.57	0.00

Panel B: Simplified transaction cost model with $a_2 = 1$

Parameter	Value	St. Dev.	t-stat	P-value	95% CI	
a_1	0.43	0.02	18.11	0.00	0.38	0.48

$R^2(\%)$	Adj. $R^2(\%)$	$F-stat$	$p(F-stat)$	Durbin-Watson	Jarque-Bera	$p(JB)$
3.8	3.8	245.30	0.00	2.00	182,799.62	1.00

of stock i at time t. Let X_t denote the matrix of all stock sensitivities $X_{i,t}$ and U_t the diagonal matrix of $U_{i,t}$. Then, the factor portfolio variance becomes

$$Var_{t-1}\left[r_t^k\right] = Var_{t-1}\left[w_t^{k\top} r_t\right] = w_t^{k\top} X_t^\top \Sigma_t^f X_t w_t^k + w_t^{k\top} U_t w_t^k. \quad (2.11)$$

Expected factor portfolio returns

Since we aim to construct optimal factor portfolios with respect to transaction costs, our utility function should produce weights equal to the ones of the long-short standard factor portfolio when transaction costs are ignored. To this end, we simply extract the implied expected returns from a given factor portfolio which would return this very factor portfolio as the optimal solution in an unconstrained mean-variance optimization. Adding a transaction cost penalty as an additional friction then helps isolating the corresponding ramifications to the factor profile. By properly managing the impact of the transaction cost penalty, we can thus improve the factors' risk-return profile whilst still targeting the underlying exposures as best as possible. Since we have defined the risk term of utility, and assuming that $w^k_{t,target}$ are the factor weights of a given factor portfolio, we can directly calculate the implied expected returns by considering the following reverse optimization problem. Since $w^k_{t,target}$ are the solutions to Equation (2.12) below, we have

$$w^k_{t,target} = \arg\max_{w^k_t} \mathbb{E}_{t-1}\left[w^{k\top}_t r_t\right] - \frac{\gamma}{2} Var_{t-1}\left[w^{k\top}_t r_t\right], \tag{2.12}$$

where the variance term is estimated using Equation (2.11). The first order conditions of the above optimization problem are

$$\mathbb{E}^k_{t-1}\left[r_t\right] - \gamma w^{k\top}_{t,target} X^\top_t \Sigma^f_t X_t = 0, \tag{2.13}$$

which implies that the expected return vector for factor k is given by

$$\mathbb{E}^k_{t-1}\left[r_t\right] = \gamma w^{k\top}_{t,target} X^\top_t \Sigma^f_t X_t, \tag{2.14}$$

where $\mathbb{E}^k_{t-1}\left[r_t\right]$ denotes the implied expected return vector of the k-th factor portfolio.

2.4 Transaction-cost-optimized factors

In this section, we introduce factor portfolios and their optimized counterparts for all three regions, US, Europe, and Emerging Markets. We showcase the efficacy of our methodology in mitigating transaction costs while preserving the integrity of factors.

2.4.1 Single factor portfolios

We first explore optimizing single-factor long-short portfolios. Starting from defining a target portfolio, we look into the performance of optimized portfolios and show how our methodology impacts transaction costs.

Academic target factor portfolios

As argued by Chordia, Goyal, and Shanken (2017), firm characteristics are deemed good proxies of expected returns, further substantiated in Kagkadis et al. (2023) at the portfolio level. With this in mind, we follow a similar approach to that of Kozak, Nagel, and Santosh (2020) and apply a linear rank-weighting scheme across two size groups, large and small. Large stocks are those in the top 90% of the market capitalization for the region, and small stocks are those in the bottom 10%. To generate market-beta-neutral long-short portfolios for each factor we lever one of the two legs based on their respective market betas. Then, we rescale both legs to retain the total leverage of 2. The resulting portfolio is market neutral but does not have a net investment of zero, i.e, the weights do not sum to zero. This procedure is reiterated each month. The assumed net asset value of the position of our portfolios is $1 billion USD, i.e., $500 million long and short before beta adjustment. An exemplary weighting scheme is depicted in Figure 2.1.

Figure 2.1: Target factor portfolio weights

We illustrate the weighting scheme for constructing the target long-short factor portfolios based on a given factor characteristic for the US universe. To depict the weighting scheme, we show the momentum weights for May 2001. The weights are representative of what one would expect for an average portfolio of any factor.

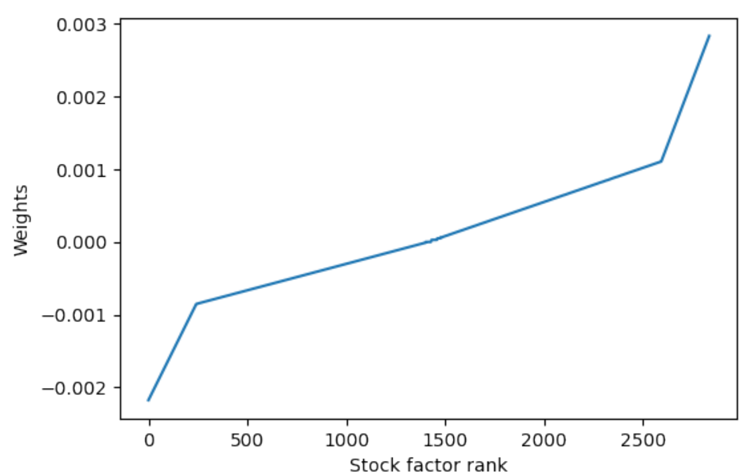

In Table 2.3, we collect the summary statistics for six salient US long-short factor portfolios following the above construction. We find that momentum is the highest gross yielding factor portfolio, with an annualized gross return of 7.71%. The return of short-term reversal

(STR) is second highest at 7.66%, and low volatility is next with 7.22%. Conversely, size is the lowest yielding portfolio, with a return of 0.81%. Value is next in line, with an annualized gross return of 2.49%, about half of quality with 6.25%. Turning from gross to net factor performance, the associated transaction costs follow intuition in many cases. For example, momentum profits are erased, seemingly confirming the notions of Lesmond, Schill, and Zhou (2004) that trading costs eliminate the profits to momentum portfolios at small fund sizes. Momentum and STR entail the highest turnover due to their fast-paced nature, with turnover figures of 534.6% and 1,659.2%, respectively. In turn, this turnover results in high annualized transaction costs of 7.19% and 31.58%. Although size has the lowest turnover at 134.8%, it still comes with considerable transaction costs at 2.08%, clearly showcasing that overweighting the smaller side of the universe comes at a hefty cost. The cheapest factors to trade were quality and low volatility, both due to their low turnovers and tendency to trade relatively larger stocks. Similar to Muravyev, Pearson, and Pollet (2022), we find momentum and quality to have the highest shorting costs at 1.89% and 1.76%, respectively, whereas value and size observe the lowest shorting costs at 0.88% and 0.18%. We find fixed fees (which include taxes, brokerage fees and foreign exchange fees) to mainly depend on turnover, with little difference between factor portfolios with similar turnover.

Our transaction cost decomposition in Panel B of Table 2.3 showcases how significant the two components, market impact and spread, are in estimating transaction costs. The ratio between spread and market impact clearly has a trade size dependency, further emphasizing how expensive trading illiquid stocks is. For the very liquid market portfolio, the market impact becomes small, showing that the spread component of transaction costs is not to be ignored when considering portfolios investing in liquid assets. Since the target factor portfolios are constructed with no regard to transaction costs, this naive portfolio construction proves detrimental to the resulting net performance. Apart from low volatility and quality, each academic factor portfolio observes a negative net return. STR reverses its profits to an average loss of −25.39% and the other factors, apart from low volatility, end up around zero. Low volatility becomes the best performing factor portfolio with a net Sharpe ratio of 0.39.

Table 2.3: Long-short factor portfolio summary statistics: US

We present the gross and net performance of the US market portfolio as well as of each long-short US factor portfolio. The considered sample period is January 1985 to March 2023. For each portfolio we present the gross and net return, Sharpe ratio, volatility, turnover, shorting costs, fixed fees and transaction costs decomposed into two components, spread and market impact. All numbers are annualized.

	Market	Momentum	Value	Quality	Size	Lowvol	STR
Panel A: Gross performance							
Gross return %	11.78	7.71	2.49	6.25	0.81	7.22	7.66
Volatility %	15.36	11.35	10.85	7.66	10.79	11.59	11.34
Gross Sharpe	0.77	0.68	0.23	0.82	0.08	0.62	0.68
Panel B: Transaction costs							
Turnover %	8.30	534.6	222.8	200.2	134.8	143.8	1,659.2
Transaction costs %	0.04	7.19	2.95	2.50	2.08	1.48	31.58
TC Spread %	0.03	2.17	0.95	0.82	0.59	0.59	6.69
TC MI %	0.01	5.02	2.00	1.68	1.49	0.89	24.89
Shorting Costs %		1.89	0.88	1.76	0.18	1.22	0.99
Fixed Fees %	0.01	0.15	0.06	0.06	0.04	0.04	0.48
Panel C: Net performance							
Net return %	11.73	-1.52	-1.40	1.93	-1.49	4.48	-25.39
Net Sharpe	0.76	-0.13	-0.12	0.26	-0.13	0.39	-2.24

How to monitor factor integrity

Many mutual funds and other investment vehicles are explicitly benchmarked against specific style or factor-based indices. For these funds, factor integrity becomes pivotal: investors expect them to maintain the distinctive characteristics—such as low volatility, momentum, or quality—that define their style. If an optimization process to reduce transaction costs alters holdings too extensively, the fund risks drifting from its stated factor exposure, potentially undermining its mandate and investor expectations. By monitoring the measures described below—factor scores, active share, transfer coefficients, and tracking error—portfolio managers can ensure they are preserving the essence of the targeted factor while still exercising prudent cost management. Hence, before optimizing factors, we elaborate on the notion of factor integrity and our approach to measuring it. Portfolio optimization with respect to transaction cost typically focuses on reducing turnover and the associated transaction costs to ultimately improve net (risk-adjusted) returns. While this is a reasonable goal, it might result in failing to capture the very factors' nature. In order to measure the closeness to the original target factor portfolio, we employ several measures. First, a key measure is the *factor score*, computed as the weighted average of the percentile characteristic rank of all stocks, normalized by the weighted average of the percentile characteristic rank of all stocks for the target portfolio. Second, the deviation from the target factor portfolio weights can be gauged by the *active share*, which is the absolute difference of weights from the target portfolio; it ranges from zero to two, with zero being an identical portfolio and two a completely different portfolio with no overlap in positions. Third, a related measure is the *transfer coefficient*, which is the correlation of optimized portfolio weights with target portfolio weights. We also report *tracking error* and *number of effective names* to track important portfolio characteristics. Obviously, the higher the transaction cost penalty, the harder it is to stay true to a given factor, and these metrics enable monitoring the degree of factors' integrity. It is important to recognize the notion of transaction cost and factor integrity trade-off. An intuitive way to express it is to maximize a measure of factor alignment minus a penalty for turnover—for example, weighting the portfolio's exposure to the target factor against the expected transaction costs from rebalancing. Concretely, this might take the form of maximizing factor score while imposing a transaction cost penalty that rises with portfolio turnover. The relative weighting of factor integrity versus cost can vary by investor type: a benchmark-aware mutual fund might place heavier emphasis on maintaining style consistency (e.g., value or low volatility) to avoid style drift, whereas a pure alpha-seeking hedge fund might tolerate more turnover (and thus higher costs) if it strengthens factor exposure. Hence, the objective function inherently reflects the investor's mandates, constraints, and tolerance for deviations from the benchmark, leading to different optimal trade-offs across fund types.

Optimized factor portfolios

In Table 2.4, we investigate the optimization of factors with respect to different transaction cost penalty parameters, denoted δ in Equation (2.6). There are several notable observations that apply across factors. The constructed optimal factor portfolios show a decline in gross returns when δ increases. This outcome is in line with Frazzini, Israel, and Moskowitz (2014) and can be rationalized by the factor score that is dropping with increasing transaction cost penalties. In an ideal scenario, the reduction in gross returns is made up for by the gains in net returns, see, e.g., the momentum factor (with exceptions in the higher δ values). After a certain δ threshold, the gross return simply deteriorates beyond the improvement on the transaction costs side. This results in a concave-shaped net return across deltas, facilitating the choice of the appropriate transaction cost parameter. For most factors, we observe close to constant volatility across different deltas, with a slight decline for $\delta > 0.0005$. As a result, Sharpe ratios yield to similar conclusions as returns, with a constant drop for gross portfolios, and a concave shape for the net portfolios. We experience a sharp decline in turnover and transaction costs across all factors, which comes with an increase in net returns as we increase the transaction cost parameter, see Figure 2.2.

Judging by net performance alone, one might be tempted to keep increasing the transaction cost penalty, yet factor integrity might be compromised as a result. Indeed, we observe a sharp decline in factor score and transfer coefficient for $\delta > 0.0005$ for most factors. Furthermore, the number of effective names drops (Table 2.4). These are important guideposts to gauging whether one has drifted too far from the target factor portfolios. Next, we will discuss the specific outcome of transaction cost optimization for each of the six factors.

Momentum

For momentum, we observe a rapid decline in gross returns as well as a decay in the factor score which is relatively large compared to other factors. This is consistent with a high-turnover factor like momentum, for which transaction cost penalties are expected to have more of an impact. Frazzini, Israel, and Moskowitz (2014) also find that momentum observes the largest drop in gross return after optimization and the highest net returns out of the three factors (momentum, value, size) they consider. Notably, momentum observes a peak net return and Sharpe ratio for $\delta = 0.0005$, while its factor integrity is still intact; the optimized portfolio is also well diversified, the factor score is above 90%, tracking error and active share are still reasonably low and the transfer coefficient is at 0.91. The top performing portfolio gains 3.70% in net returns over the original target factor portfolio, showing that sacrificing some momentum exposure for reducing costs can yield much better net outcomes. We see that the transaction costs are drastically reduced from 7.19% to 2.70% for the optimized

Figure 2.2: Key US factor characteristics by transaction cost penalty

We plot the net return, volatility, net Sharpe, factor score, transfer coefficient and turnover across different deltas, transaction cost penalty parameter from Equation (2.6), for the US factor portfolios. The period considered is January 1985 to March 2023. Net return, volatility and turnover are reported as percentages and annualized Net Sharpe ratio is also annualized. Net return and Sharpe ratios are calculated using the gross return and estimated transaction costs from our transaction cost model. Factor score is calculated as the weighted average of individual stock scores and given as a fraction of the respective value for the target portfolio. Transfer coefficient is calculated as the correlation between optimal weights for a given δ and target factor weights. Turnover for a given month ranges from 0 to 100%, which equates to a maximum of 1,200% per annum. Since short-term reversal exhibits very low net returns, we plot it on a separate right-hand-side y-axis in Figure 2.2a.

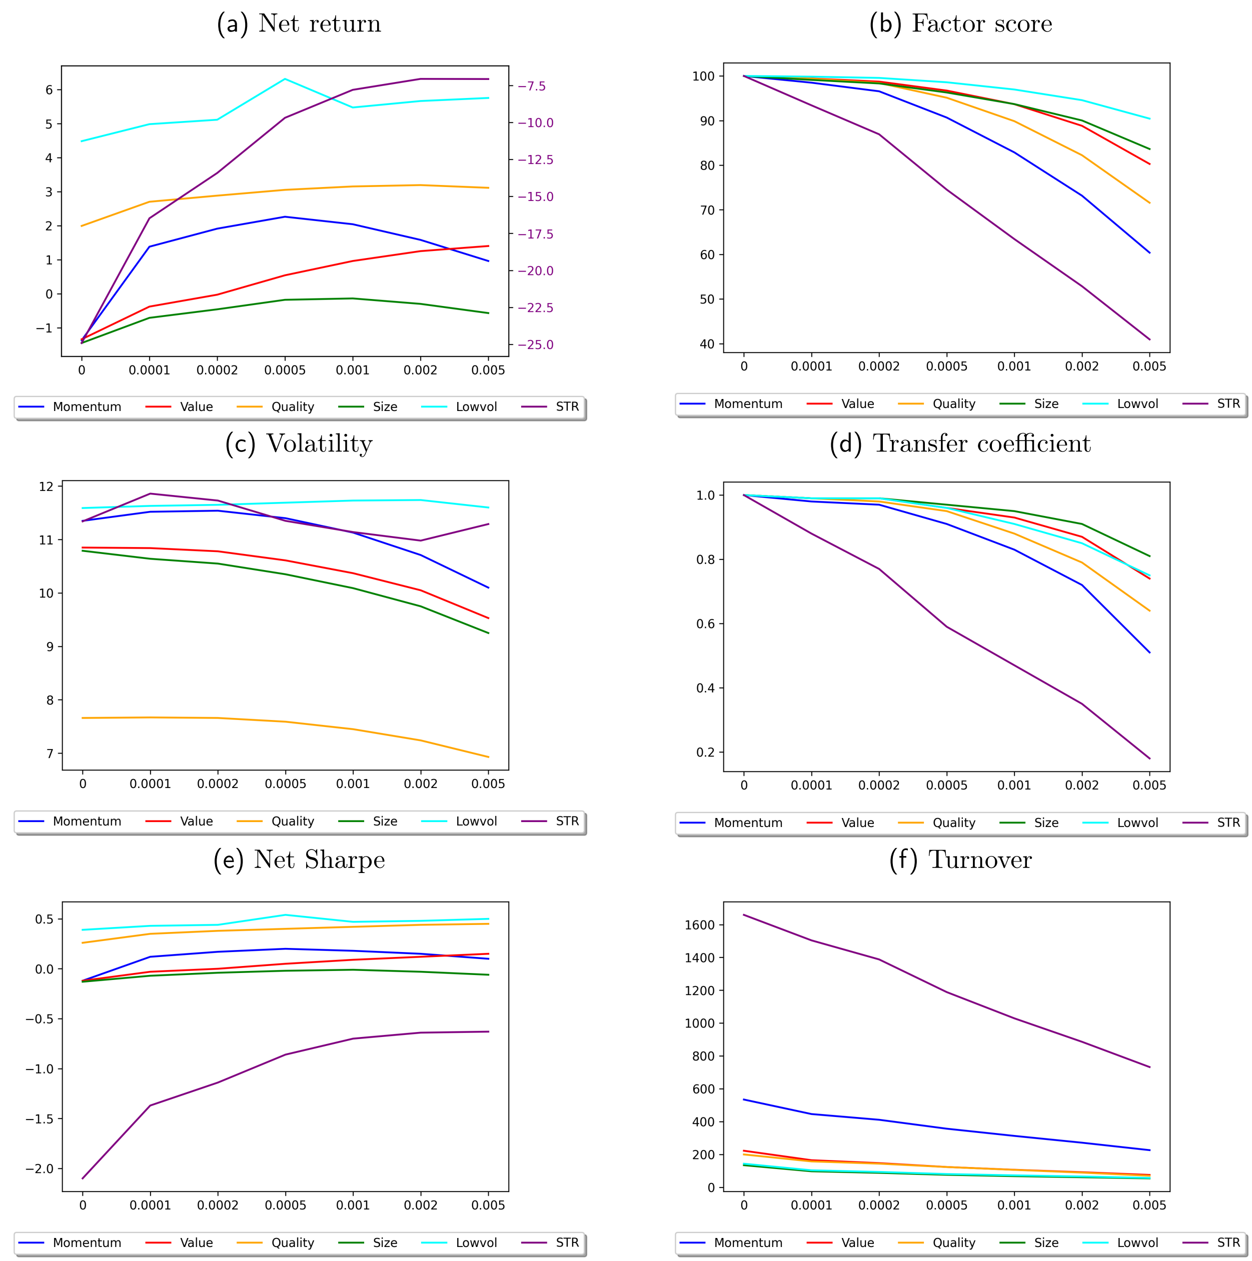

momentum factor. Furthermore, momentum observes the highest shorting costs among all factors, starting from 1.89% for the target portfolio and going down to 1.54% for the optimal portfolio.

Value

The optimized value factors display very different behaviour when compared to momentum. Their gross returns and factor scores decay much slower, but start from a lower point of 2.49% going down to 2.30% in the optimal case of $\delta = 0.001$. Transaction costs are though significantly reduced from 2.95% for the target portfolio down to 0.67% for the optimized value factor portfolio. We also observe improvements in net Sharpe ratio, going from -0.12 to 0.09 in the optimized case. Both tracking error and active share are comparable to that of the momentum factors. Shorting fees are relatively low at 0.88% and 0.67% for the target and optimized portfolios, respectively. Increasing δ further, net performance increases as well. Yet, for larger values of δ, we observe too low factor scores and too high tracking errors, making value a great example of how overly reducing transaction costs can create a portfolio that no longer represents the targeted factor portfolio.

Quality

Quality is one of the lower turnover strategies and offers the least room for improvement as a result. Notwithstanding, we can reduce the required transaction costs from 2.50% to 0.63% in the optimized case ($\delta = 0.001$). The gross performance is reduced by 1.25%, but the net return is still significantly increased, going from 1.93% to 3.11%. The optimized quality portfolio retains a low tracking error of 1.49% and a high transfer coefficient of 0.88. The associated shorting fees are very high (second only to the ones of momentum), standing at 1.76% and 1.22% for the target and optimal factor portfolios, respectively.

Size

The size factor observes the lowest target factor gross return out of all factors considered (0.81%). Its turnover is the lowest of all factors, starting at 134.8% and going down to 68.0% in the optimized case ($\delta = 0.001$). Consistent with Brière et al. (2019), we find size to

Table 2.4: Optimized US factor portfolios

We present the summary statistics of different optimized US factor portfolios. The optimized portfolios differ by the transaction cost parameter used in the utility function: $\delta = 0$ corresponds to the long-short target portfolios. Where applicable, all measures are annualized and quoted in percentages. All relative measures such as tracking error and active share are benchmarked to the $\delta = 0$ (target) portfolio. The considered sample period is January 1985 to March 2023.

Momentum	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$	$\delta = 0.005$
Panel A: Performance							
Gross return %	7.71	7.57	7.27	6.51	5.52	4.40	3.17
Net return %	-1.52	1.28	1.83	2.18	1.97	1.52	0.91
Volatility %	11.35	11.52	11.54	11.40	11.13	10.71	10.10
Gross Sharpe	0.68	0.66	0.63	0.57	0.50	0.41	0.31
Net Sharpe	-0.13	0.12	0.17	0.20	0.18	0.15	0.10
Panel B: Transaction costs							
Turnover %	534.6	446.2	411.4	357.0	313.2	271.8	226.4
Transaction costs %	7.19	4.43	3.67	2.70	2.10	1.66	1.29
TC Spread %	2.17	1.64	1.44	1.15	0.94	0.77	0.62
TC MI %	5.02	2.79	2.23	1.56	1.16	0.88	0.67
Shorting Costs %	1.89	1.76	1.69	1.54	1.38	1.17	0.91
Fixed Fees %	0.15	0.10	0.08	0.08	0.07	0.06	0.05
Panel C: Factor integrity							
Factor score %	100.00	98.49	96.56	90.67	82.86	73.16	60.38
Tracking error %	0.00	0.56	0.84	1.38	1.97	2.71	3.93
Active Share	0.00	0.25	0.39	0.68	0.97	1.30	1.73
Transfer coefficient	1.00	0.98	0.97	0.91	0.83	0.72	0.51
# of effective names	1489	1448	1426	1376	1279	1119	782

Size	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$	$\delta = 0.005$
Panel A: Performance							
Gross return %	0.81	0.57	0.60	0.65	0.55	0.28	-0.08
Net return %	-1.49	-0.75	-0.50	-0.21	-0.17	-0.33	-0.60
Volatility %	10.79	10.64	10.55	10.35	10.09	9.75	9.25
Gross Sharpe	0.08	0.05	0.06	0.06	0.05	0.03	-0.01
Net Sharpe	-0.13	-0.07	-0.04	-0.02	-0.01	-0.03	-0.06
Panel B: Transaction costs							
Turnover %	134.8	97.6	88.2	76.2	68.0	61.4	54.4
Transaction costs %	2.08	1.09	0.88	0.64	0.50	0.40	0.31
TC Spread %	0.59	0.37	0.33	0.27	0.23	0.20	0.17
TC MI %	1.49	0.72	0.56	0.38	0.27	0.20	0.14
Shorting Costs%	0.18	0.18	0.18	0.18	0.18	0.18	0.19
Fixed Fees %	0.04	0.04	0.04	0.03	0.03	0.03	0.03
Panel C: Factor integrity							
Factor score %	100.00	99.11	98.32	96.29	93.69	90.02	83.62
Tracking error %	0.00	0.48	0.73	1.20	1.72	2.37	3.36
Active Share	0.00	0.14	0.22	0.38	0.55	0.77	1.13
Transfer coefficient	1.00	0.99	0.99	0.97	0.95	0.91	0.81
# of effective names	1489	1479	1469	1434	1379	1283	1077

Value	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$	$\delta = 0.005$
Panel A: Performance							
Gross return %	2.49	2.02	2.02	2.15	2.30	2.35	2.28
Net return %	-1.40	-0.43	-0.07	0.50	0.92	1.22	1.37
Volatility %	10.85	10.84	10.78	10.61	10.37	10.05	9.53
Gross Sharpe	0.23	0.19	0.19	0.20	0.22	0.23	0.24
Net Sharpe	-0.12	-0.03	-0.00	0.05	0.09	0.12	0.15
Panel B: Transaction costs							
Turnover %	222.8	165.2	147.4	123.8	107.0	92.0	75.6
Transaction costs %	2.95	1.57	1.24	0.88	0.67	0.51	0.37
TC Spread %	0.95	0.62	0.52	0.40	0.33	0.27	0.21
TC MI %	2.00	0.95	0.72	0.48	0.34	0.24	0.16
Shorting Costs %	0.88	0.83	0.80	0.73	0.67	0.60	0.51
Fixed Fees %	0.06	0.05	0.04	0.04	0.04	0.03	0.03
Panel C: Factor integrity							
Factor score %	100.00	99.42	98.73	96.72	93.67	88.82	80.27
Tracking error %	0.00	0.49	0.73	1.21	1.77	2.49	3.48
Active Share	0.00	0.17	0.26	0.46	0.67	0.94	1.34
Transfer coefficient	1.00	0.99	0.99	0.96	0.93	0.87	0.74
# of effective names	1527	1510	1495	1452	1384	1268	1046

Lowvol	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$	$\delta = 0.005$
Panel A: Performance							
Gross return %	7.22	6.92	6.83	6.76	6.72	6.72	6.59
Net return %	4.48	4.98	5.11	6.31	5.47	5.66	5.75
Volatility %	11.59	11.63	11.65	11.69	11.73	11.74	11.60
Gross Sharpe	0.62	0.60	0.59	0.58	0.57	0.57	0.57
Net Sharpe	0.39	0.43	0.44	0.54	0.47	0.48	0.50
Panel B: Transaction costs							
Turnover %	143.8	103.8	93.8	81.0	73.1	66.0	57.2
Transaction costs %	1.48	0.76	0.61	0.44	0.36	0.30	0.24
TC Spread %	0.59	0.35	0.30	0.24	0.20	0.17	0.15
TC MI %	0.89	0.41	0.31	0.21	0.16	0.12	0.09
Shorting Costs %	1.22	1.14	1.08	0.97	0.86	0.73	0.57
Fixed Fees %	0.04	0.04	0.03	0.03	0.03	0.03	0.03
Panel C: Factor integrity							
Factor score %	100.00	99.84	99.56	98.58	96.96	94.57	90.43
Tracking error %	0.00	0.33	0.53	0.96	1.39	1.88	2.55
Active Share	0.00	0.16	0.26	0.48	0.71	0.97	1.29
Transfer coefficient	1.00	0.99	0.99	0.96	0.91	0.85	0.75
# of effective names	1,332	1,313	1,293	1,231	1,136	1,016	847

Quality	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$	$\delta = 0.005$
Panel A: Performance							
Gross return %	6.25	5.84	5.65	5.30	5.00	4.69	4.25
Net return %	1.93	2.65	2.83	3.01	3.11	3.15	3.08
Volatility %	7.66	7.67	7.66	7.59	7.45	7.24	6.93
Gross Sharpe	0.82	0.76	0.74	0.70	0.67	0.65	0.61
Net Sharpe	0.26	0.35	0.38	0.40	0.42	0.44	0.45
Panel B: Transaction costs							
Turnover %	200.2	157.8	144.2	123.6	106.4	89.6	70.6
Transaction costs %	2.50	1.50	1.22	0.86	0.63	0.45	0.30
TC Spread %	0.82	0.57	0.49	0.38	0.30	0.22	0.16
TC MI %	1.68	0.93	0.73	0.48	0.33	0.22	0.14
Shorting Costs %	1.76	1.64	1.55	1.39	1.22	1.05	0.85
Fixed Fees %	0.06	0.05	0.05	0.04	0.04	0.04	0.03
Panel C: Factor integrity							
Factor score %	100.00	99.37	98.42	95.12	89.86	82.23	71.57
Tracking error %	0.00	0.37	0.58	1.02	1.49	2.05	2.91
Active Share	0.00	0.18	0.30	0.56	0.84	1.16	1.56
Transfer coefficient	1.00	0.99	0.98	0.95	0.88	0.79	0.64
# of effective names	1522	1497	1476	1419	1336	1207	972

STR	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$	$\delta = 0.005$
Panel A: Performance							
Gross return %	7.66	5.61	4.46	3.15	2.20	1.08	-0.28
Net return %	-25.39	-16.88	-13.77	-10.02	-8.08	-7.31	-7.29
Volatility %	11.34	11.86	11.73	11.35	11.14	10.98	11.29
Gross Sharpe	0.68	0.47	0.38	0.28	0.20	0.10	-0.02
Net Sharpe	-2.10	-1.37	-1.14	-0.86	-0.70	-0.64	-0.63
Panel B: Transaction costs							
Turnover %	1,659.2	1,503.8	1,387.8	1,188.8	1,029	886.2	732.4
Transaction costs %	31.58	21.23	16.99	12.02	9.26	7.38	5.81
TC Spread %	6.69	5.69	4.98	3.89	3.15	2.58	2.09
TC MI %	24.89	15.54	12.02	8.14	6.12	4.80	3.72
Shorting Costs %	0.99	0.85	0.89	0.82	0.73	0.76	0.98
Fixed Fees %	0.48	0.40	0.35	0.33	0.28	0.25	0.22
Panel C: Factor integrity							
Factor score %	100.00	93.41	86.92	74.51	63.43	52.86	40.94
Tracking error %	0.00	1.50	1.93	2.83	3.66	4.40	6.69
Active Share	0.00	0.51	0.75	1.17	1.53	1.87	2.29
Transfer coefficient	1.00	0.88	0.77	0.59	0.47	0.35	0.18
# of effective names	1454	1047	705	406	326	292	143

incur lower transaction costs, second only to low volatility. However, when comparing size to other factors of similar turnover, its transaction costs to turnover ratio is high which is a consequence of frequently trading in small stocks. Its net performance increases from the initial −1.49% to −0.17% for the optimized case, rendering the size factor unattractive after transaction costs. Shorting costs are tiny compared to other factors, as size only short-sells large stocks with low shorting costs. We also see that the shorting costs are almost identical across all deltas. This is expected, as the transaction cost penalties apply mostly on the long (and less liquid) leg of the portfolio.

Low volatility

Low volatility has the second lowest turnover, with an annual turnover of the target portfolio at 143.8%. It also exhibits the lowest transaction costs of all factors, which is in line with Rossi, Hoch, and Steliaros (2022). When optimizing transaction costs its net performance is increased by more than one percent for the $\delta = 0.002$ case. Its volatility, however, slightly increases as well, moderating the improvement in Sharpe ratio. Interestingly, factor score is still above 90% for the last case of $\delta = 0.005$, yet active share is 1.29 and tracking error is at 2.55%. This highlights the importance of using different metrics when considering how close one tracks the target portfolio. Transaction costs are reduced from 1.48% to 0.30%, making the low volatility factor the cheapest optimized factor portfolio to trade.

Short-term reversal

The short-term reversal factor exhibits a huge turnover of $1,659.2\%$, leading to a vast transaction cost figure of 31.58% for the target factor portfolio. The corresponding gross portfolio return of 7.66% is thus reduced to a net return of -25.39%, highlighting the need to optimize the cost of implementing such fast-paced strategies. For the case of $\delta = 0.0005$, we observe a net performance increase of about 15% points, with a tracking error of 2.83%. Yet, despite cutting all these costs, the factor remains deeply in the negative as a standalone long-short portfolio. Similarly, Rossi, Hoch, and Steliaros (2022) also find that short-term reversal observes negative net returns even in the case of the smallest portfolio size. This is further substantiated in Blitz, Grient, and Honarvar (2023), where they argue using a single signal to construct short-term reversal results in strategies with negative net returns even under a low transaction cost assumption.

Decomposing transaction costs

Notably, our framework does not require to penalize transaction costs in factor portfolios with equal transaction cost penalties. For a high turnover signal like momentum, it can be more reasonable to penalize transaction costs less, as imposing too strict of a transaction cost penalty would result in a significant loss in factor score and thus exposure. Indeed, momentum shows the second fastest decline in factor score. Conversely, one can be more strict on some factors in terms of penalizing transaction costs but still conserving the integrity of factors.

Clearly, optimizing single-factor portfolios comes with sacrificing some turnover, and thus some factor exposure, in favor of lowering the cost of trading. What makes this optimization setting using a transaction cost model so beneficial is how effective it is in reducing expensive turnover, see Figure 2.3 which depicts the transaction cost decomposition for all target and optimized factor portfolios. We note the effect of reducing large relative trade sizes as the market impact component experiences a minimum threefold reduction in all cases. Naturally, the spread component can also be reduced by reducing turnover in stocks with higher spread, but not as effectively as the market impact component. This is a consequence of market impact growing with a larger exponent of amount traded than spread, since spread is a fixed cost per share.

2.4.2 International evidence

The performance improvement we find for US factor portfolios is only amplified for Europe and Emerging Markets. To illustrate, Figure 2.4 visualizes the net performance improvements that come from using the optimal transaction cost penalties for all six factors. Each bar starts from the theoretical maximum, i.e., the gross return of the target factor portfolio. Accounting for transaction costs, the net performance of the target portfolio is then contrasted with the net performance improvement of the optimal target portfolio. Across regions, we can successfully enhance net returns over the market returns for most factors, with some factors going from negative to positive territory.

Figure 2.3: Transaction cost decomposition for the US factor portfolios

We decompose the transaction cost of all single factors into the spread and market impact components for the target and optimized portfolios. Transaction costs of the market portfolio and each single factor portfolio is first split into two columns, spread (blue) and market impact (orange). Then, each column presents the transaction costs of a target factor portfolio and the optimized factor portfolio. This is illustrated by a lighter shade of the respective color for the optimized factor portfolio, i.e., light-blue and light-orange correspond to the spread and market impact of the optimal portfolio, respectively. In the case of short-term reversal, the value of market impact (orange) presented is lower than the actual value for better readability. All transaction cost numbers are quoted in percentages. The underlying sample period is January 1985 to March 2023.

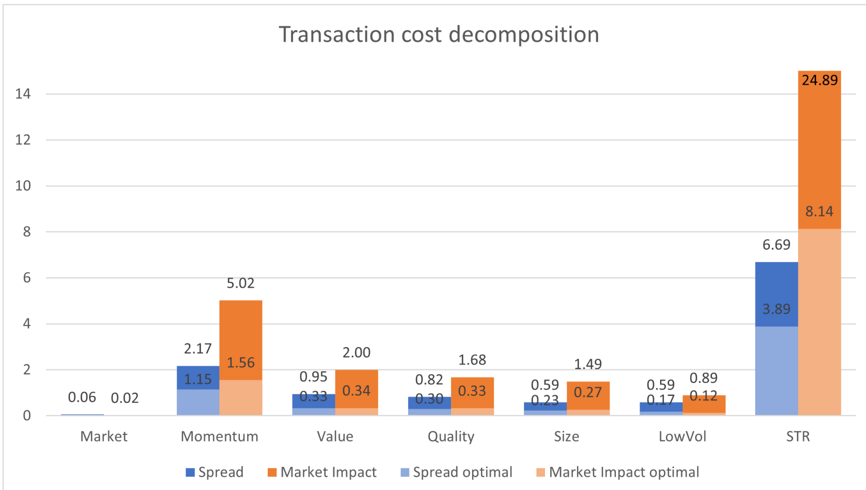

In Europe, the results are similar to that of the US but on a larger scale. The gross and net returns of the momentum target factor portfolios are 11.36% and −0.10%, respectively. After optimization, the net return of the optimized momentum portfolio is 4.25%. For EM, the gross return of the momentum target factor portfolio is close to the US one at 6.95% with a negative net return of −0.26%. After optimization, the net return goes up to 0.87% which is significantly lower than that of the US momentum factor, showcasing how expensive EM stocks are to trade frequently compared to the US. Consistent with Rossi, Hoch, and Steliaros (2022), the value factor in Europe observes a negative gross return in the target factor portfolio. Conversely, value in EM has a 2.99% gross return. However, the associated net returns stand at −1.67% which can be pushed to 0.02% only (after optimization). Therefore, unlike in the US, we cannot retain the value factor premium by avoiding expensive trades in the EM universe.

Quality factor performance is quite comparable across markets, with Europe observing the highest gross target factor portfolio returns. Net returns of the optimized quality factor portfolio go from 1.10% to 2.68% and 1.02% to 2.38% for the European and EM case, respectively. In contrast to the US, the size factor observes larger gross returns in European and Emerging markets. In both regions, we manage to retain about two thirds of the gross returns after optimization. Despite low volatility being one of the highest performing factors across all markets, this is where the US outperforms the other regions. Looking at the EM results, we again observe the increase of gross returns in the optimized case compared to the target factor portfolio, seemingly resulting in almost very low transaction costs. This again shows how important optimization is in the EM universe. Lastly, short-term reversal leaves a lot to be desired for across markets. For both Europe and EM, the costs of trading this turnover-heavy strategy are just too high to overcome. Such finding is consistent with Blitz, Grient, and Honarvar (2023) who argue that the classic short-term reversal factor needs refinement to become profitable. Specifically, the authors suggest controlling for industry and factor momentum effects. Also, the combination of different short-term signals is found an effective means to improving signal efficacy after transaction costs, see Blitz et al. (2023).

Figure 2.4: Target vs. optimized factor performance

We plot the gross and net return of target factor portfolios with the net return increase of the optimal portfolio across regions. For each factor portfolio, we show the gross performance of the target factor portfolio (gray), the net performance of the target factor portfolio (blue) and the net performance increase from the target factor portfolio to the optimal factor portfolio (orange). Overlapping areas are colored in dual colors. The underlying sample period is January 1985 to March 2023 for the US and Europe, and January 2000 to March 2023 for Emerging markets.

(a) US

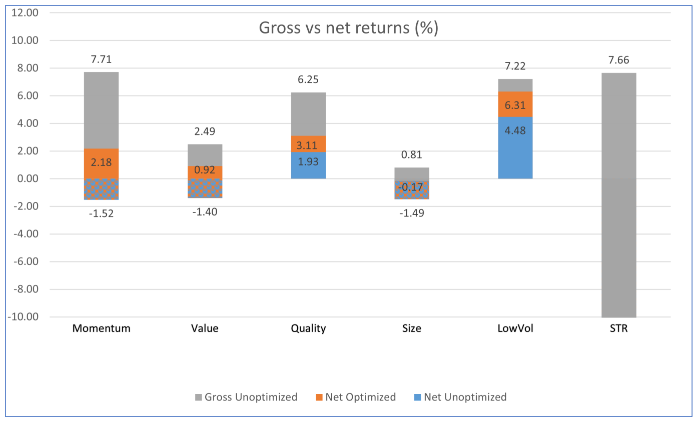

(b) Europe

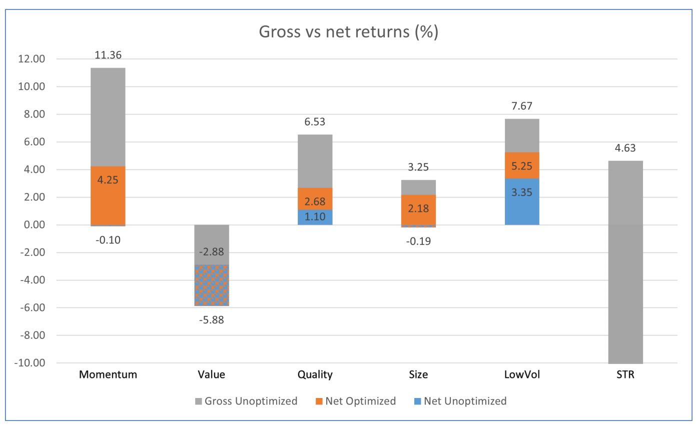

(c) Emerging Markets

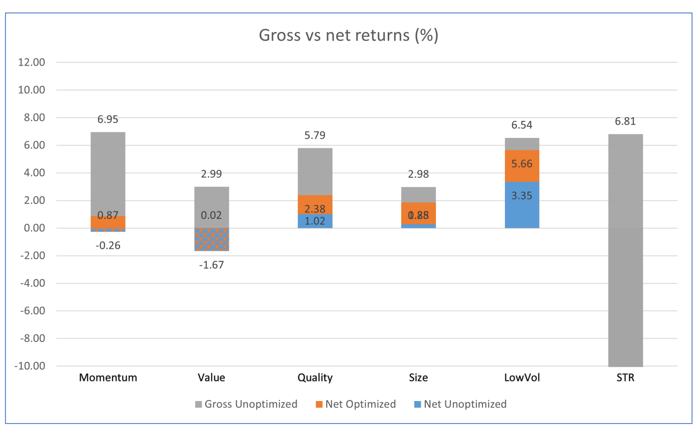

2.4.3 Robustness with respect to fund size

While our main results are based on an assumed fund size of $1 billion USD net asset value, Table 2.5 next showcases our framework's efficacy using two alternative NAV assumptions, $500 million USD and $2 billion USD. Apart from short-term reversal and value in emerging markets, we see positive net returns across factors and markets for the smaller NAV case of $500 million USD. For the $2 billion USD case, we see some factors becoming unprofitable even in the optimized case, for example momentum and value in emerging markets. This is mainly a result of change in market impact, as spread, shorting costs and fixed fees remain mostly unchanged.

Table 2.5: Long-short single-factor portfolios under different NAV assumptions

The table presents gross and net performance of optimized single-factor portfolios under three different net asset value assumptions: $500 million USD, $1 billion USD and $2 billion USD. The table is split into three regions: US, Europe and Emerging markets. The underlying sample period is January 1985 to March 2023 for the US and Europe, and January 2000 to March 2023 for Emerging markets. All numbers are quoted in percentages and annualized.

US	Momentum	Value	Quality	Size	Lowvol	STR
Gross return %	7.71	2.49	6.25	0.81	7.22	7.66
Net return % ($ 500 mil.)	3.88	1.22	3.45	0.09	6.53	-4.23
Net return % ($ 1 bil.)	2.18	0.92	3.11	-0.17	6.31	-10.02
Net return % ($ 2 bil.)	0.82	0.35	2.63	-0.48	6.12	-13.35
Europe	Momentum	Value	Quality	Size	Lowvol	STR
Gross return %	11.36	-1.00	6.53	3.25	7.67	4.63
Net return % ($ 500 mil.)	6.22	-2.10	3.63	2.84	5.82	-6.34
Net return % ($ 1 bil.)	4.25	-2.88	2.68	2.18	5.25	-10.79
Net return % ($ 2 bil.)	1.12	-3.81	1.32	1.29	4.52	-17.35
Emerging markets	Momentum	Value	Quality	Size	Lowvol	STR
Gross return %	6.95	2.99	5.79	2.98	6.54	6.81
Net return % ($ 500 mil.)	2.63	0.71	2.80	2.19	5.89	-4.98
Net return % ($ 1 bil.)	0.87	0.02	2.38	1.88	5.66	-11.43
Net return % ($ 2 bil.)	-1.86	-0.64	1.99	1.67	5.31	-18.49

2.4.4 Multi-factor portfolios

After constructing optimized single-factor portfolios, we next turn to the multi-factor case. It is natural to expect multi-factor portfolios to see some reduction in turnover (relative to the average turnover of considered factors) because of trading cost diversification. We show that our methodology naturally carries over to this case such that one can preserve the balance across the contributing factors. Similar to the single-factor portfolio case, we start by defining a multi-factor target portfolio. For simplicity, we choose the multi-factor target portfolio to be an equally weighted average of all six single-factor target portfolios. As a result, the multi-factor portfolio is a long-short portfolio based on characteristic-rank weighting of each factor. Not surprisingly, we find that combining factor portfolios increases the gross performance compared to individual factor portfolios, and we will dive into the regional multi-factor evidence next.

US

The results for the optimized US multi-factor portfolios across different transaction cost parameter δ are shown in Table 2.6. The multi-factor target portfolio has a gross performance of 14.90% against that of the market of 11.78%. However, the net performance of the multi-factor portfolio is -1.05%, as the massive turnover (938.8%) and subsequent transaction costs erode all of the gross performance. The volatility for the US multi-factor target portfolio is 10.23%, lower than that of the market at 15.36%. This results in a slightly negative net Sharpe ratio of -0.10. As in the single-factor portfolios, after optimization, we observe a decline in gross returns, going from 14.90% to 9.13% when increasing δ. Volatility remains similar across all tested portfolios, resulting in gross Sharpe ratios declining similarly to gross returns. Looking at net performance, we find that the $\delta = 0.0005$ portfolio comes with almost the best net return of 6.83%, an improvement of almost 8% over the target portfolio. In terms of staying true to the target portfolio, the tracking error stands at 1.69% with an active share of 0.75. In addition with a high transfer coefficient of 0.88, we conclude that factor integrity of the multi-factor target portfolio is kept.

Table 2.6: Long-short optimized multi-factor portfolios: US

We present the summary statistics of different optimized multi-factor portfolios. The optimized portfolios differ by the transaction cost parameter used in the utility function: $\delta = 0$ corresponds to the long-short target portfolio. Where applicable, all measures are annualized and quoted in percentages. All relative measures such as tracking error and active share are benchmarked to the $\delta = 0$ (target) portfolio. For each factor, we consider one target and seven optimized portfolios and present gross return, net return, transaction costs, turnover, volatility, gross Sharpe ratio, net Sharpe ratio, factor score, tracking error, active share, transfer coefficient and number of effective names. The sample period ranges from January 1985 to March 2023.

	Market	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$
Panel A: Performance							
Gross return %	11.78	14.90	13.35	12.75	11.92	11.16	10.19
Net return %	11.73	-1.05	4.46	5.65	6.83	7.13	6.88
Volatility %	15.36	10.23	10.33	10.36	10.40	10.38	10.30
Gross Sharpe	0.77	1.46	1.29	1.23	1.15	1.08	0.99
Net Sharpe	0.76	-0.10	0.43	0.55	0.66	0.69	0.67
Panel B: Transaction costs							
Turnover %	8.30	938.8	761.4	687.6	579.6	503.6	438.0
Transaction costs %	0.04	15.15	8.20	6.43	4.49	3.47	2.80
TC Spread %	0.03	3.78	2.73	2.32	1.76	1.42	1.18
TC MI %	0.01	11.37	5.46	4.11	2.73	2.05	1.62
Shorting Costs %		0.71	0.61	0.58	0.54	0.49	0.45
Fixed Fees %	0.01	0.09	0.08	0.08	0.07	0.07	0.07
Panel C: Factor integrity							
Factor score %		100.00	96.70	95.14	91.96	88.12	83.15
Tracking error %		0.00	0.81	1.14	1.69	2.11	2.60
Active Share		0.00	0.30	0.46	0.75	1.02	1.33
Transfer coefficient		1.00	0.98	0.95	0.88	0.80	0.68
# of effective names		1,511	1,484	1,447	1,320	1,160	921

Europe

The gross performance of the European target portfolio is better than that of the US, with a gross Sharpe ratio of 1.69. However, the transaction costs are considerably higher at 22.13%, resulting in a negative net Sharpe ratio of -0.95. As we increase δ, we again observe the loss of gross returns, which gradually declines from 14.62% to 9.60%. Volatility remains flat for most cases, resulting in Sharpe ratio patterns resembling net returns. The largest net return is 6.52% in the case of $\delta = 0.001$. Looking at tracking error, we see that the optimized portfolios quickly drift apart from the target multi-factor portfolio as we reduce transaction costs. Considering the rapidly increasing tracking error and active share, we conclude that using a higher transaction cost parameter leads to optimized portfolios too different from the target portfolio. We will further evidence this when analyzing the exposure to individual factors.

Table 2.7: Long-short optimized multi-factor portfolios: Europe

We present the summary statistics of different optimized multi-factor portfolios for Europe. The optimized portfolios differ by the transaction cost parameter used in the utility function: $\delta = 0$ corresponds to the long-short target portfolio. Where applicable, all measures are annualized and quoted in percentages. All relative measures such as tracking error and active share are benchmarked to the $\delta = 0$ (target) portfolio. For each factor, we consider one target and seven optimized portfolios and present gross return, net return, transaction costs, turnover, volatility, gross Sharpe ratio, net Sharpe ratio, factor score, tracking error, active share, transfer coefficient and number of effective names. The sample period ranges from January 1985 to March 2023.

	Market	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$
Panel A: Performance							
Gross return %	10.64	14.62	13.82	13.41	12.77	11.97	11.05
Net return %	10.38	-8.25	2.17	4.13	6.00	6.52	6.49
Volatility %	17.33	8.67	8.71	8.63	8.60	8.59	8.65
Gross Sharpe	0.61	1.69	1.59	1.55	1.49	1.39	1.28
Net Sharpe	0.60	-0.90	0.25	0.48	0.70	0.76	0.75
Panel B: Transaction costs							
Turnover %	16.27	980.0	768.2	694.2	588.6	512.8	448.2
Transaction costs %	0.23	22.13	11.06	8.72	6.26	4.97	4.11
TC Spread %	0.06	3.97	2.65	2.24	1.71	1.36	1.09
TC MI %	0.17	18.16	8.41	6.48	4.55	3.61	3.03
Shorting Costs %		0.60	0.48	0.45	0.42	0.39	0.36
Fixed Fees %	0.02	0.14	0.11	0.11	0.10	0.10	0.09
Panel C: Factor integrity							
Factor score %		100.00	95.26	92.91	87.97	81.76	74.14
Tracking error %		0.00	1.44	2.33	3.20	4.01	4.59
Active Share		0.00	0.44	0.64	1.00	1.34	1.69
Transfer coefficient		1.00	0.96	0.92	0.84	0.73	0.58
# of effective names		572	509	386	262	193	175

Emerging markets

For Emerging markets, the results are somewhere in between the US and Europe, starting with a gross Sharpe ratio of 1.58. The turnover is the highest $(1,035.6\%)$ resulting in transaction costs of 16.07%. The associated net Sharpe ratio is -0.15. On the other hand, the transaction cost-optimized portfolio construction yields even better outcomes in Emerging markets. We see familiar patterns for gross returns, volatility and net returns with one important difference, the pace of decay. For larger deltas, portfolios observe lower tracking errors and active shares, and higher transfer coefficients. Taking all of this into account, we find that the $\delta = 0.001$ portfolio gives the best return profile while retaining factor integrity to a great extent with the exception of a higher tracking error.

Table 2.8: Long-short optimized multi-factor portfolios: EM

We present the summary statistics of different optimized multi-factor portfolios for the EM region. The optimized portfolios differ by the transaction cost parameter used in the utility function: $\delta = 0$ corresponds to the long-short target portfolios. Where applicable, all measures are annualized and quoted in percentages. All relative measures such as tracking error and active share are benchmarked to the $\delta = 0$ (target) portfolio. For each factor, we consider one target and seven optimized portfolios and present gross return, net return, transaction costs, turnover, volatility, gross Sharpe ratio, net Sharpe ratio, factor score, tracking error, active share, transfer coefficient and number of effective names. The sample period ranges from January 2000 to March 2023.

	Market	$\delta = 0$	$\delta = 0.0001$	$\delta = 0.0002$	$\delta = 0.0005$	$\delta = 0.001$	$\delta = 0.002$
Panel A: Performance							
Gross return %	8.15	15.42	15.53	15.56	15.44	15.08	14.46
Net return %	7.70	-1.42	1.16	2.48	4.61	6.20	7.52
Volatility %	22.36	9.73	9.72	9.73	9.75	9.72	9.63
Gross Sharpe	0.36	1.58	1.60	1.60	1.58	1.55	1.50
Net Sharpe	0.35	-0.15	0.12	0.26	0.48	0.65	0.79
Panel B: Transaction costs							
Turnover %	30.30	1,035.6	982.2	947.5	873.4	791.4	689.6
Transaction costs %	0.43	16.07	13.64	12.37	10.15	8.22	6.32
TC Spread %	0.09	2.28	2.08	1.96	1.72	1.48	1.20
TC MI %	0.34	13.79	11.56	10.41	8.43	6.74	5.12
Shorting Costs%		0.62	0.60	0.58	0.56	0.54	0.51
Fixed Fees %	0.02	0.16	0.14	0.13	0.12	0.12	0.10
Panel C: Factor integrity							
Factor score %		100.00	99.27	98.74	97.39	95.54	92.50
Tracking error %		0.00	0.89	1.18	1.72	2.28	3.00
Active Share		0.00	0.11	0.18	0.32	0.49	0.71
Transfer coefficient		1.00	0.99	0.99	0.97	0.95	0.90
# of effective names		483	482	481	478	469	441

Factor balance

While it is comforting to see that the overall transfer of the multi-factor target is successful, we are also mindful of investigating as to how much of the underlying single factor exposures the multi-factor portfolios are able to transfer. Given the different single factor dynamics, we are eager to learn if the targeted equal factor balance can actually be maintained in a transaction-cost-optimized multi-factor portfolio. We perform such factor loading attribution by regressing the optimized multi-factor portfolio weights on the underlying target single-factor portfolio weights. Figure 2.5 collects the ensuing factor loadings and the overall decomposition by factor. Naturally, the multi-factor target portfolio can be fully replicated in absence of transaction costs ($\delta = 0$, leftmost stacked bars), hence we observe a fully balanced (i.e., equally weighted) factor loadings decomposition. As the transaction cost penalty δ increases we see the decay of short-term reversal and momentum exposures which we expect to suffer the most as expensive turnover is reduced. Conversely, low volatility is the most persistent factor with little loss in factor exposure across all deltas.

Focusing on US evidence, we see the preservation of factor balance even at higher deltas, reinforcing our choice of $\delta = 0.0005$. We find that the optimized multi-factor portfolio only loses exposure to the short-term reversal and momentum factors, as expected. Looking at higher deltas, we find that value and quality are reduced but retain significant presence in the portfolio. Size and low volatility remain almost unaffected. Still, the regression fit quickly deteriorates, seeing R^2 below 50% for δ higher than 0.01.

In contrast, the European multi-factor portfolios experience a more rapid decay in factor exposure as we increase the transaction cost penalty. Momentum, quality, value and short-term reversal gradually lose exposure with quality decaying the quickest. Size loses a lot of exposure in the higher deltas. Low volatility increases in exposure significantly in the middle range of deltas. The rapid decay in factor exposure is further evidenced by the faster decline of R^2. This outcome can be further rationalized when considering the fact that the European multi-factor portfolio is considerably less diversified than the US or EM portfolios in terms of number of effective names, making it relatively harder to sustain factor exposure at higher transaction cost penalties (and thus complicating the substitution of costly high factor exposure names with cheaper ones).

Against this backdrop, we see very persistent EM multi-factor loadings with little decay, similar to the US evidence. Naturally, we also observe the faster decay of short-term reversal and retention of the low volatility factor. The remaining four factors decay gradually and in a similar manner. This is further supported by a high R^2 across most deltas. The observed factor balance reinforces our choice of $\delta = 0.001$ as the best multi-factor portfolio for the EM region.

2.4.5 VMQL portfolios

When constructing a multi-factor portfolio, the approach was to equal-weight all six single-factor portfolios. However, as argued in Blitz and Hanauer (2021), the size factor is weak as a stand-alone factor (even after controlling for quality-versus-junk exposures) and is hardly explicitly targeted in factor investing propositions. Similarly, the short-term reversal factor has weakened over time when the underlying characteristic used is last month return (Blitz, Grient, and Honarvar, 2023). Indeed, our single factor results confirm the poor performance of short-term reversal and size, and thus we investigate a more common choice of multi-factor portfolio that focuses on the four factors value, momentum, quality and low volatility. We showcase the performance of these VMQL portfolios in Table 2.9. As expected, while removing short-term reversal results in a decrease in gross Sharpe, we observe a sharp increase in net Sharpe. The decrease in gross Sharpe becomes less pronounced in the optimized case, as the exposure to short-term reversal is usually considerably reduced as shown in Figure 2.5. Looking at factor integrity, VMQL portfolios generally observe larger factor scores, substantially lower tracking errors and active shares, when compared to the 6-factor portfolios.

Importantly, the removal of size and short-term reversal make the multi-factor performance more resilient with respect to larger fund sizes as well, see Panel D of Table 2.9 that shows multi-factor portfolio performance in two more cases, $500 million USD and $2 billion USD.

2.5 Benchmarking transaction cost-optimized factor portfolios

In order to mitigate transaction costs, Novy-Marx and Velikov (2016) suggest several turnover-reducing methods, concluding that the buy/hold spread rule is a highly effective cost mitiga-

Figure 2.5: Factor balance across optimal multi-factor portfolios

We regress the multi-factor portfolio weights on the individual factor weights for each optimal portfolio. Each bar shown represents the average factor exposure per factor. Starting from a balanced 1/6 target portfolio, we observe how multi-factor portfolios reduce exposure in some factors quickly while other factors remain unaffected or decay slowly. The underlying sample period is January 1985 to March 2023 for the US and Europe, and January 2000 to March 2023 for Emerging markets.

(a) US

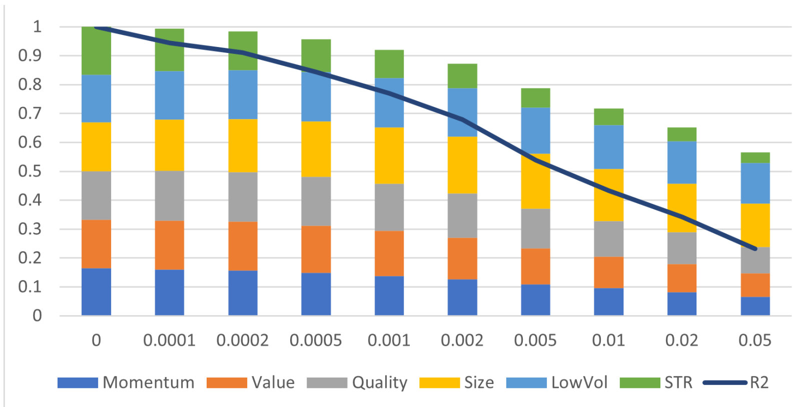

(b) Europe

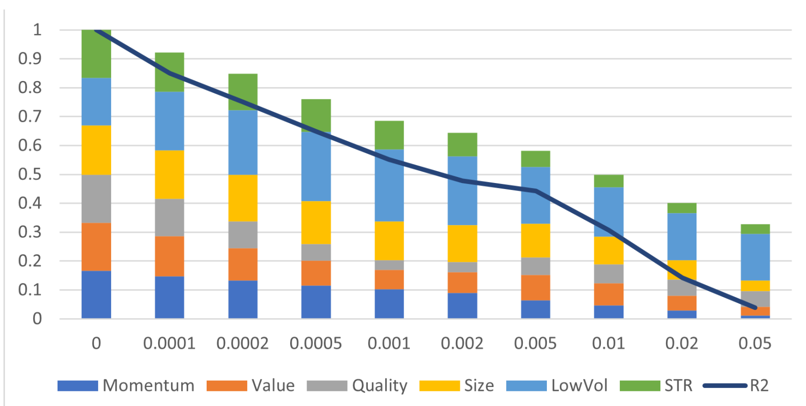

(c) EM

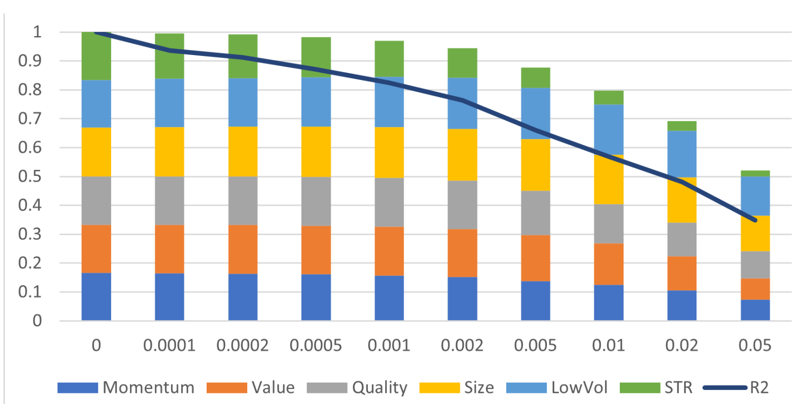

Table 2.9: Multi-factor vs. VMQL performance

The table presents the gross and net performance of optimized multi-factor portfolios under two different constructions. The first is the multi-factor portfolio using all six factors we investigate. VMQL is the multi-factor portfolio restricted to only four factors: momentum, value, quality and low volatility. All relative metrics such as tracking error and factor score are benchmarked to the respective unoptimized portfolios and all numbers are quoted in percentages and annualized. Panel D presents the gross and net performance of optimized multi-factor and VMQL portfolios under three different net asset value assumptions: $500 million USD, $1 billion USD and $2 billion USD. The underlying sample period is January 1985 to March 2023 for the US and Europe, and January 2000 to March 2023 for Emerging markets. All numbers are quoted in percentages and annualized.

	US		EU		EM	
	6-factor	*VMQL*	*6-factor*	*VMQL*	*6-factor*	*VMQL*
Panel A: Performance						
Gross return %	11.92	12.27	11.97	12.51	15.08	13.58
Net return %	6.83	10.73	6.52	8.82	6.20	10.87
Volatility %	10.40	12.24	8.59	9.31	9.72	9.22
Gross Sharpe	1.15	1.00	1.39	1.34	1.55	1.47
Net Sharpe	0.66	0.88	0.76	0.95	0.65	1.18
Panel B: Transaction costs						
Turnover %	579.6	500.9	512.8	551.6	791.4	529.92
Transaction costs %	4.49	1.01	4.97	3.13	8.22	2.17
TC Spread %	1.76	0.44	1.36	0.73	1.48	1.09
TC MI %	2.73	0.57	3.61	2.40	6.74	1.08
Shorting Costs%	0.54	0.48	0.39	0.35	0.54	0.43
Fixed Fees %	0.07	0.05	0.10	0.09	0.12	0.10
Panel C: Factor integrity						
Factor Score %	91.96	94.56	81.76	92.56	95.54	94.68
Tracking error %	1.69	1.06	4.01	1.82	2.28	1.49
Active Share	0.75	0.54	1.34	0.50	0.49	0.50
Transfer coefficient	0.88	0.89	0.73	0.94	0.95	0.93
# of effective names	1,320	1,075	193	370	469	424
Panel D: Performance under different NAV assumptions						
Gross return %	11.92	12.27	11.97	12.51	15.08	13.58
Net ret. % ($ 500 mil.)	9.15	11.15	8.47	9.87	9.73	11.34
Net ret. % ($ 1 bil.)	6.83	10.73	6.52	8.82	6.20	10.87
Net ret. % ($ 2 bil.)	4.44	10.21	4.02	7.45	2.47	9.90

tion technique. For example, their 10%/20% buy/hold spread rule for the long side implies buying a stock when the stock enters the top 10% based on the underlying characteristic, but only selling it when it falls below 20% and thus avoiding flip-flopping of a given name. They also argue that only factor portfolios with low to moderate turnover generate significant net returns, which deviates from our conclusions. To rationalize, we replicate the buy/hold methodology employed in Novy-Marx and Velikov (2016) based on our factor data and transaction cost model, see Table 2.10. Instead of using their value-weighting scheme, we use equally-weighted market-beta-neutral factor portfolios, rendering both factor construction methods more comparable. While the resulting portfolios yield higher gross returns than our factor portfolios, we note that this is driven by their considerably higher volatility, resulting in relatively lower Sharpe ratios. Furthermore, the net returns of these portfolios are much lower than those of our portfolios due to the large transaction costs they incur. The turnover is often very high for the equally-weighted starting portfolios, and is substantially reduced in the more restrictive h = 0.5 case. While their approach is effective at reducing turnover, the resulting reduction in transaction costs is much less rewarding when compared to our approach. Take momentum as an example: our approach reduces turnover from 534.6% to 357.0% in the optimal case and thus cuts transaction costs from 7.19% down to 2.70%. On the other hand, the buy/hold approach reduces turnover from 761.5% down to 313.6% in the h = 0.5 case. Despite having similar turnover to our approach, the buy/hold spread rule sees more than twice the transaction costs (5.65%), clearly showing that we can achieve much better cost mitigation at similar turnover values.

To illustrate, consider a scenario with a 10%/50% buy/hold spread portfolio where a stock enters the top decile in the first month, drops at the 30^{th} percentile next month, and then again drops at the 50^{th} the month after. In the buy/hold spread portfolio, we enter into a long position in the stock the first month, hold it for two months and finally sell it as it hits the 50^{th} percentile. In both our unoptimized and optimized portfolios, the stock is gradually traded off as the underlying characteristic reduces, where the trades are done even more gradually in the optimal case due to the transaction cost penalty. This results in very similar turnovers but the market impact incurred is substantially reduced. Importantly, in stark contrast to our approach, the buy/spread rule leads to a considerable divation from the starting factor portfolio, as reflected in high active share and tracking error figures. As a result, we observe lower transfer coefficents, and thus the buy/hold spread rule struggles to stay true to the original factor given relatively low factor scores either. To summarize, the buy/spread rule is too ad hoc and leaves ample room for further optimizing a given factor's transaction cost-factor integrity tradeoff.

Table 2.10: Comparison to Novy-Marx and Velikov (2016)

We present the summary statistics of factor portfolios created using the buy/hold spread approach of Novy-Marx and Velikov (2016). The portfolios differ by the hold threshold used. Equal-weighted corresponds to the equal-weighted long-short decile market-beta-neutral portfolios and $h = 0.2$ and 0.5 correspond to a hold threshold of 20 and 50%. Non-optimized and Optimized columns correspond to the $\delta = 0$ and optimal δ cases from Table 2.4. The considered sample period is January 1985 to March 2023. Where applicable, all measures are annualized and quoted in percentages. All relative measures such as tracking error and active share are benchmarked to the equal-weighted portfolio.

Momentum	Equal-weighted	h = 0.2	h = 0.5	Non-optimized	Optimized
Panel A: Performance					
Gross return %	11.30	10.11	7.21	7.71	6.51
Net return %	-10.58	-3.14	-0.94	-1.52	2.18
Volatility %	24.35	21.86	17.94	11.35	11.40
Gross Sharpe	0.46	0.46	0.40	0.68	0.57
Net Sharpe	-0.43	-0.15	-0.06	-0.12	0.20
Panel B: Transaction costs					
Turnover %	761.5	483.1	313.6	534.6	357.0
Transaction costs %	18.82	10.46	5.65	7.19	2.70
TC Spread %	3.12	1.97	1.27	2.17	1.15
TC MI %	15.70	8.49	4.38	5.02	1.56
Shorting Costs %	2.87	2.66	2.42	1.89	1.54
Fixed Fees %	0.20	0.14	0.09	0.15	0.08
Panel C: Factor integrity					
Factor score %	100.00	94.75	79.30	100.00	90.67
Tracking error %	0.00	3.67	8.21	0.00	1.38
Active Share	0.00	0.96	1.93	0.00	0.68
Transfer coefficient	1.00	0.87	0.72	1.00	0.91
# of effective names	414	544	809	1489	1376

Size	Equal-weighted	h = 0.2	h = 0.5	Non-optimized	Optimized
Panel A: Performance					
Gross return %	3.44	4.03	2.85	0.81	0.55
Net return %	-8.31	-1.31	0.22	-1.49	-0.17
Volatility %	15.65	13.42	10.78	10.79	10.09
Gross Sharpe	0.22	0.30	0.26	0.08	0.05
Net Sharpe	-0.53	-0.09	0.02	-0.13	-0.01
Panel B: Transaction costs					
Turnover %	262.6	146.5	112.5	134.8	68.0
Transaction costs %	11.49	5.09	2.42	2.08	0.50
TC Spread %	1.24	0.68	0.51	0.59	0.23
TC MI %	10.25	4.41	1.91	1.49	0.27
Shorting Costs %	0.18	0.18	0.18	0.18	0.18
Fixed Fees %	0.08	0.07	0.04	0.04	0.03
Panel C: Factor integrity					
Factor score %	100.00	95.14	82.66	100.00	93.69
Tracking error %	0.00	4.16	8.40	0.00	1.72
Active Share	0.00	1.05	1.90	0.00	0.55
Transfer coefficient	1.00	0.86	0.72	1.00	0.95
# of effective names	542	731	1,014	1,489	1,379

Value	Equal-weighted	h = 0.2	h = 0.5	Non-optimized	Optimized
Panel A: Performance					
Gross return %	1.51	-0.39	1.56	2.49	2.30
Net return %	-14.61	-10.91	-6.34	-1.40	0.92
Volatility %	23.80	22.03	18.48	10.85	10.37
Gross Sharpe	0.06	-0.02	0.08	0.23	0.22
Net Sharpe	-0.62	-0.49	-0.33	-0.12	0.09
Panel B: Transaction costs					
Turnover %	579.8	407.9	346.1	222.8	107.0
Transaction costs %	14.85	9.29	6.74	2.95	0.67
TC Spread %	2.34	1.65	1.48	0.95	0.33
TC MI %	12.51	7.63	5.26	2.00	0.34
Shorting Costs %	1.18	1.14	1.11	0.88	0.67
Fixed Fees %	0.10	0.08	0.05	0.06	0.04
Panel C: Factor integrity					
Factor score %	100.00	94.29	79.10	100.00	93.67
Tracking error %	0.00	8.49	14.21	0.00	1.77
Active Share	0.00	1.13	2.13	0.00	0.67
Transfer coefficient	1.00	0.84	0.66	1.00	0.93
# of effective names	348	479	743	1527	1384

Lowvol	Equal-weighted	h = 0.2	h = 0.5	Non-optimized	Optimized
Panel A: Performance					
Gross return %	14.22	12.61	11.56	7.22	6.76
Net return %	1.76	3.50	4.05	4.48	6.31
Volatility %	24.06	22.51	20.33	11.59	11.69
Gross Sharpe	0.59	0.56	0.57	0.62	0.58
Net Sharpe	0.08	0.16	0.20	0.39	0.54
Panel B: Transaction costs					
Turnover %	448.4	364.7	329.5	143.8	66.0
Transaction costs %	10.65	7.41	5.92	1.48	0.30
TC Spread %	1.76	1.44	1.35	0.59	0.17
TC MI %	8.89	5.97	4.57	0.89	0.12
Shorting Costs %	1.72	1.63	1.54	1.22	0.73
Fixed Fees %	0.08	0.07	0.05	0.04	0.03
Panel C: Factor integrity					
Factor score %	100.00	94.29	80.51	100.00	94.57
Tracking error %	0.00	6.92	11.03	0.00	1.88
Active Share	0.00	1.07	1.94	0.00	0.97
Transfer coefficient	1.00	0.85	0.70	1.00	0.85
# of effective names	306	415	598	1,332	1,016

Quality	Equal-weighted	h = 0.2	h = 0.5	Non-optimized	Optimized
Panel A: Performance					
Gross return %	14.40	12.02	10.54	6.25	5.00
Net return %	0.39	1.13	0.98	1.93	3.11
Volatility %	21.50	18.25	17.65	7.66	7.45
Gross Sharpe	0.67	0.66	0.60	0.82	0.67
Net Sharpe	0.02	0.06	0.06	0.26	0.42
Panel B: Transaction costs					
Turnover %	502.1	413.5	382.7	200.2	106.4
Transaction costs %	11.70	8.69	7.51	2.50	0.63
TC Spread %	1.99	1.65	1.53	0.82	0.30
TC MI %	9.71	7.04	5.97	1.68	0.33
Shorting Costs %	2.22	2.12	1.99	1.76	1.22
Fixed Fees %	0.09	0.08	0.06	0.06	0.04
Panel C: Factor integrity					
Factor score %	100.00	95.15	82.00	100.00	89.86
Tracking error %	0.00	8.05	12.94	0.00	1.49
Active Share	0.00	0.96	1.78	0.00	0.84
Transfer coefficient	1.00	0.87	0.73	1.00	0.88
# of effective names	347	457	633	1,522	1,336

STR	Equal-weighted	h = 0.2	h = 0.5	Non-optimized	Optimized
Panel A: Performance					
Gross return %	9.34	9.33	10.68	7.66	3.15
Net return %	-53.78	-47.46	-29.69	-25.39	-10.02
Volatility %	26.30	26.17	25.29	11.34	11.35
Gross Sharpe	0.36	0.36	0.42	0.68	0.28
Net Sharpe	-2.04	-1.81	-1.17	-2.10	-0.86
Panel B: Transaction costs					
Turnover %	2,180.7	2,016.5	1,562.9	1,659.2	1,188.8
Transaction costs %	61.42	55.19	39.00	31.58	12.02
TC Spread %	8.84	8.18	6.38	6.69	3.89
TC MI %	52.57	47.00	32.62	24.89	8.14
Shorting Costs %	1.06	1.04	0.95	0.99	0.82
Fixed Fees %	0.65	0.57	0.42	0.48	0.33
Panel C: Factor integrity					
Factor score %	100.00	97.85	81.45	100.00	74.51
Tracking error %	0.00	7.45	15.36	0.00	2.83
Active Share	0.00	0.46	1.45	0.00	1.17
Transfer coefficient	1.00	0.94	0.78	1.00	0.59
# of effective names	332	362	502	1,454	406

2.6 Conclusion

We present a framework to optimize the transactions cost of implementing salient long-short equity factors. The framework is independent of factor choice and can be applied to any weighting scheme or variance-covariance matrix. Using real trading data, we construct a tractable parametric transaction cost model to both evaluate and optimize factor portfolios. While many academic factor portfolios are characterized by underwhelming net performances, these factor portfolios can be optimized such that their net performance becomes economically relevant. Specifically, sacrificing less than half of the turnover in most cases, we can greatly reduce the incurred transaction costs. This is more noticeable in the market impact component, as it grows with a higher exponent than the spread component. We further show that these results can be achieved for various fund size assumptions. In addition to the standard US universe, we explore the performance of our framework in Europe and Emerging Markets. We find our results to remain consistent across those regions, resulting in larger improvements in the performance of the optimized portfolios when compared to the US.

Importantly, we document the relevance of monitoring the trade-off of reducing factor turnover and transaction costs versus preserving factor integrity. Specifically, some factors exhibit high net returns when constructed with a very high transaction cost penalty, yet bear little to no similarity with the original factors. This puts an emphasis on monitoring factor exposure as turnover is reduced.

We optimize a diversified multi-factor portfolio, where one does not only care about preserving the underlying factor's integrity but also about preserving the targeted factor balance. Indeed, the constructed optimized multi-factor portfolio observes increased net performance while retaining a well-balanced factor exposure. We show that our factor portfolio construction is robust to fund size assumption and choice of multi-factor portfolio constituents.

Chapter 3

Transaction Cost-Optimal Factor-Enhanced Market Portfolios

This project is joint work with my supervisors Alberto Martín-Utrera, Harald Lohre, Sandra Nolte and Ingmar Nolte. We thank David Blitz, Lars ter Braak, Khuram Chaudhry, Bart van der Grient, Matthias Hanauer, Iman Honarvar, Tobias Hoogteijling, Clint Howard, Maarten Jansen, Stefan Mittnik, Nick Mutsaers, Martin van der Schans, Dmitri Tikhonov, Weili Zhou, Machiel Zwanenburg, and the Robeco Quant Research Seminar.

3.1 Introduction

In the earlier chapters, we first introduced the role of transaction cost volatility in a utility-based framework, illustrating how higher-order features of transaction costs can influence optimal portfolio selection. Subsequently, we examined the construction of transaction cost-optimized factor portfolios, showing that explicitly accounting for transaction costs can lead to material improvements in net performance for standalone factor strategies. Building on these insights, this chapter extends that methodology to benchmark-relative portfolios, incorporating factors into a market (benchmark) portfolio subject to tracking error constraints and explicit transaction cost considerations. By addressing the reality that many institutional investors operate under benchmark comparisons, we demonstrate how optimizing factor tilts against a reference index can yield superior net outcomes while respecting both cost and tracking error requirements. In the realm of portfolio management, the concept of benchmark relative portfolios has garnered significant attention due to its pivotal role in evaluating portfolio performance. We look into benchmark relative portfolios inclusive of factor models, with a particular focus on tracking error and transaction costs; two critical aspects that significantly impact portfolio performance. Utilizing a simple yet versatile method of incorporating both factor model and transaction cost views into our utility based portfolio optimization procedure, we construct factor enhanced market portfolios optimal with respect to transaction costs. Beginning with simple factor portfolios and the difficulties of implementing standalone long-short single-factor portfolios, we construct market factors tilted towards factor portfolios and bound by tracking error. We conclude by showing that such portfolios' net performances are greatly improved when transaction costs are also taken into account.

Numerous studies have investigated the determinants and implications of tracking error. Roll (1992) highlighted the importance of tracking error in active portfolio management, emphasizing that minimizing tracking error is crucial for managers who aim to replicate benchmark performance while implementing their investment strategies. Furthermore, Jorion (2003) provided a comprehensive analysis of the sources of tracking error, identifying factors such as market timing, security selection, and portfolio rebalancing as primary contributors. Transaction costs, encompassing brokerage fees, bid-ask spreads, and market impact costs, play a vital role in the construction and management of benchmark relative portfolios. These costs can erode portfolio returns, making it imperative for portfolio managers to carefully consider them when making investment decisions. Research by Treynor and Black (1973) underscored the significance of transaction costs in active portfolio management, demonstrating that high transaction costs can negate the benefits of superior security selection. Additionally, Constantinides (1986) explored the impact of transaction costs on portfolio rebalancing, showing that frequent rebalancing can lead to substantial transaction costs, thereby reducing overall portfolio performance.

The interplay between tracking error and transaction costs presents a complex challenge for portfolio managers. While minimizing tracking error is essential for closely replicating benchmark performance, doing so often incentivizes frequent trading, which in turn increases transaction costs. Therefore, an optimal balance must be struck to ensure that the portfolio remains aligned with the benchmark while keeping transaction costs at a minimum. Grinold and Kahn (2000) addressed this trade-off, proposing a framework for managing portfolios that considers both tracking error and transaction costs. Their research highlighted the importance of strategic rebalancing and cost-effective trading practices in achieving desirable portfolio outcomes.

In this light, this paper makes three contributions to the factor investing literature. First, we use factor portfolio trading data from a large institutional asset manager to measure transaction costs. Like Frazzini, Israel, and Moskowitz (2018), we show that one can thus more accurately estimate realistic transaction costs, as measured by a trade's implementation shortfall (Perold, 1988). Second, we construct a simple way of incorporating factor exposure to benchmark portfolios while considering transaction costs. Our methodology is flexible with respect to benchmark portfolio choice, factor choice and factor construction. Finally, we show the consistence of our approach across different fund size assumptions.

The paper is structured as follows. Section 3.2 describes the data used for both the transaction cost modelling, factor portfolio and benchmark portfolio construction. Section 3.3 defines the utility functions we use, and discusses implementation shortfall, transaction cost models, tracking error, factor portfolios, benchmark portfolios as well as the corresponding estimation and construction procedures. Section 3.3 concludes with portfolio optimization and evaluation of the performance of the portfolios. Section 3.4 contains the estimation results for the transaction cost model, factor and benchmark portfolio performance as well as the benchmark relative factor portfolio performance. Our conclusions are presented in Section 3.5.

3.2 Transaction cost and stock data

For the stock data, we use the Center for Research in Security Prices (CRSP) daily data files ranging from January 2000 to December 2023. We filter the bottom 10% of stocks with respect to market capitalization, as that removes the small illiquid stocks. We further omit all stocks that have any missing observations, which makes up 3% of the stocks. Table 3.1 reports the summary statistics on number of stocks, returns, trading volume and market capitalization across all stocks.

Table 3.1: Stock data summary statistics

We present the summary statistics of CRSP monthly data. The data consists of monthly number of stocks, returns, volume traded and market capitalization covering the period from January 2000 to December 2023.

	# of stocks	Return (%)	Trading Volume ($\times 10^6$ shares)	Market Capitalization ($ bil)
min	3,558	-98.39	0.00	0.01
q25	3,777	-5.95	1.11	0.16
Median	3,962	0.58	5.03	0.61
q75	4,502	7.13	17.55	2.53
max	5,697	1,988.36	2,012.43	3,071.34
Mean	4,177	1.31	24.38	5.58
SD	521	16.96	107.66	32.25

Our trading data consists of 99,420 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. 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. The main parameter used in our transaction cost analysis is the trade size as a percentage of mean daily volume (MDV) which we calculate over the last month trading volume.

Table 3.2 reports summary statistics of the transaction costs data. As mentioned, the measure we use for transaction costs is implementation shortfall (IS), defined in Perold (1988). The median trade size observed is 0.17%, similar to what Frazzini, Israel, and Moskowitz (2018) report. 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.

Table 3.2: Transaction cost data summary statistics across regions

We present the summary statistics of our trading data. The data consists of 99,420 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 mean daily volume (MDV) and implementation shortfall. Mean daily volume is approximated as the last month trading mean of daily volume. Implementation shortfall is calculated in accordance to Perold (1988) and is reported as a fraction of the original trade value and capped at 10%.

	% MDV	IS (%)
min	0.00	-9.94
q25	0.04	-0.69
Median	0.17	0.04
q75	0.61	0.78
max	296.86	9.86
Mean	0.90	0.05
SD	3.07	1.48

3.3 Factor and market portfolio construction and transaction costs

We begin by considering a value-weighted market portfolio with weights w_t^M constructed out of N_t stocks. Then, we construct K single-factor portfolios as 100% long and 100% short portfolios that we obtain in a two-step procedure. First, all stocks considered are ranked by the underlying factor characteristic to identify the top and bottom stocks to populate the long and short leg of the factor portfolio. Second, a weighting scheme is assigned and as we aim to exploit the entire cross-section, we opt for characteristic rank-weighted single factor portfolios. Characteristic rank-weighted portfolios are typically obtained by sorting the stock universe on the characteristic, assigning a rank to each security, and applying a linear weighting scheme across ranks. Then, define the return of the k-th single factor long-short portfolio at time t, denoted by r_t^k as:

$$r_t^k = w_t^{k\top} r_t = w_{1,t}^k r_{1,t} + w_{2,t}^k r_{2,t} + \ldots + w_{N,t}^k r_{N_t,t}, \tag{3.1}$$

with

$$\forall k \ \forall t \quad \sum_{i=1}^{N_t} w_{i,t}^k = 0, \ \sum_{i=1}^{N_t} |w_{i,t}^k| = 2, \tag{3.2}$$

where $w_{i,t}^k$ is the weight of stock i in the k-th single-factor long-short portfolio and $r_{i,t}$ the return on stock i at time t. w_t^k is a vector of the k-th single-factor long-short portfolio weights $w_{i,t}^k$, and r_t is the vector of stock returns $r_{i,t}$ at time t.

3.3.1 Benchmark-relative mean-variance framework with transaction costs

Long-short factor portfolios are often attractive in theory but difficult to implement in practice, given constraints on short selling, regulatory limits, and overall capacity. As a result, most asset managers integrate factor models within long-only portfolios, using simple, liquid benchmarks—such as a market portfolio—as the base. This approach ensures both ease of investment and alignment with typical mandates that preclude extensive short exposure, while still striving to capture the potential performance benefits of factor tilts.

For many institutional investors, portfolio performance is evaluated relative to a designated benchmark. As such, it is not only absolute returns that matter but also consistency with the benchmark's behavior. Tracking error is a key risk metric in this setting, quantifying how much the portfolio's returns deviate from those of the benchmark. By incorporating a tracking error constraint or penalty, managers can control how "close" they stay to the benchmark—an especially important consideration when investors or regulators expect limited divergence from a reference index. Thus, using tracking error as a guiding metric enables the construction of benchmark-relative portfolios that strike a balance between active factor tilts and the need to avoid excessive deviation from the benchmark.

Having constructed the market and single-factor portfolios, we introduce a mean-variance optimization framework for constructing the weights of the optimized value-weighted market portfolio enhanced with single-factor long-short portfolios, which we will denote with w_t. Define a mean-variance utility maximization problem with transaction costs on a single-factor long-short portfolio as:

$$\underset{w_{i,t}}{\arg\max}\ \mathbb{E}_{t-1}\left[w_t^\top r_t\right] - \frac{\gamma}{2} Var_{t-1}\left[w_t^\top r_t\right] - \delta \mathbb{E}_{t-1}\left[\Delta w_t^\top TC\left(\Delta w_t\right)\right], \tag{3.3}$$

subject to

$$\forall k\ \forall t \quad \sum_{i=1}^{N_t} w_{i,t} = 1,\ w_{i,t} \geq 0,\ \sqrt{Var_{t-1}\left[\left(w_{i,t}^M - w_{i,t}^m\right) r_t\right]} \leq \gamma_{TE}, \tag{3.4}$$

where $\mathbb{E}_{t-1}\left[w_t^\top r_t\right]$ is the expected return of the enhanced market portfolio, $Var_{t-1}\left[w_t^\top r_t\right]$ its variance, $\mathbb{E}_{t-1}\left[\Delta w_t^\top TC(\Delta w_t)\right]$ the expected transaction costs, Δw_t is the difference in weights defined as $\Delta w_t = w_t - w_{t-1}$, γ the risk aversion coefficient, and δ is the transaction cost parameter. $\sqrt{Var_{t-1}\left[\left(w_{i,t}^M - w_{i,t}^m\right) r_t\right]}$ is the variance of the active portfolio return and γ_{TE} is the tracking error constraint.

3.3.2 Modelling transaction costs

Since our transaction cost data is based on realized trades of an asset manager, we can follow the more recent work of Rossi, Hoch, and Steliaros (2022) and Frazzini, Israel, and Moskowitz (2018), who both estimate a market impact model using real trading data. The market impact model is then used in portfolio construction. 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:

$$TC_t\left(\Delta w_{i,t}\right) = a_1 \frac{AuM \times \Delta w_{i,t}}{MDV_{i,t}} + \epsilon_{i,t}, \tag{3.5}$$

where TC is implementation shortfall as a fraction of trade size 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 mean daily volume (MDV). Using MDV also somewhat alleviates the concerns of under-penalizing transaction costs in the earlier period of our analysis, as trading costs have gone down significantly over time.

Inputs for Expected Risk and Return

Given the extensive cross-section of our universe, directly constructing a variance-covariance matrix for each individual element is impractical. To address this, we simplify the problem by adopting a linear factor model. The model is represented by:

$$r_{i,t} = X_{i,t}^{\top} r_t^f + u_{i,t}, \tag{3.6}$$

where r_t^f denotes the vector of factor portfolio returns at time t, with $r_t^f = \left[r_t^1, r_t^2, \ldots, r_t^K\right]$. Here, $X_{i,t}$ represents the sensitivities of stock i to the factors, and $u_{i,t}$ captures the stock-specific excess return. The variance of stock returns can then be estimated as:

$$Var_{t-1}\left[r_{i,t}\right] = X_{i,t}^{\top}\Sigma_t^f X_{i,t} + U_{i,t}, \tag{3.7}$$

where Σ_t^f is the covariance matrix of the factors, and $U_{i,t} = Var_{t-1}\left(u_{i,t}\right)$ represents the idiosyncratic risk variance of stock i at time t. Let X_t be the matrix of sensitivities for all stocks, and U_t the diagonal matrix of $U_{i,t}$. The variance of a portfolio can then be expressed as:

$$Var_{t-1}\left[w_t^{\top}r_t\right] = w_t^{\top}X_t^{\top}\Sigma_t^f X_t w_t + w_t^{\top}U_t w_t. \tag{3.8}$$

Since our goal is to construct optimal factor-enhanced market portfolios given a tracking error budget while accounting for transaction costs, the utility function should produce weights that match those of the long-short factor portfolio when weight and tracking error restrictions as well as transaction costs are disregarded. To achieve this, we extract the implied expected returns from a given factor portfolio that would yield this factor portfolio as the optimal solution under an unconstrained mean-variance optimization. By introducing a transaction cost penalty as an additional constraint, we can isolate its effects on the portfolio profile. Effectively managing the impact of the transaction cost penalty and the tracking error enables us to enhance the risk-return profile of the factor-enhanced market portfolio while still closely aligning with the original market portfolio.

Given that we have already defined the risk term within the utility function, and assuming that w_t^k represents the factor weights of a given factor portfolio, we can compute the implied expected returns by solving the following reverse optimization problem. As w_t^k solves Equation (3.9) below, we have:

$$w_t^k = \underset{w_t}{\arg\max}\, \mathbb{E}_{t-1}\left[w_t^{\top}r_t\right] - \frac{\gamma}{2}Var_{t-1}\left[w_t^{\top}r_t\right], \tag{3.9}$$

where the variance term is estimated using Equation (3.8). The first-order conditions for the optimization problem are:

$$\mathbb{E}_{t-1}^k\left[r_t\right] - \gamma w_t^{k\top}X_t^{\top}\Sigma_t^f X_t = 0, \tag{3.10}$$

which implies that the expected return vector for factor k is:

$$\mathbb{E}_{t-1}^{k}\left[r_t\right] = \gamma w_t^{k\top} X_t^{\top} \Sigma_t^f X_t, \tag{3.11}$$

where $\mathbb{E}_{t-1}^{k}\left[r_t\right]$ represents the implied expected return vector of the k-th factor portfolio.

3.3.3 Portfolio construction

The factor-enhanced market portfolios are constructed under a set of constraints to ensure appropriate allocation and risk management. First, a non-negativity constraint is imposed on the weights, such that the portfolio weights satisfy $w_i \geq 0$, ensuring a long-only strategy. Additionally, the weights are constrained to sum to 1, represented as $\sum_i w_i = 1$, which guarantees that the entire portfolio is fully invested. Furthermore, a tracking error constraint is included to limit the deviation of the portfolio from the benchmark, helping to manage risk more effectively.

These portfolios are the outcome of a utility optimization process. The utility function incorporates three key components: expected returns, risk, and transaction costs. Expected returns used are factor-specific and defined in Equation 3.11. Risk is measured through the variance of factor portfolio's returns and exposure of each stock to each factor, see Equation 3.8. Lastly, transaction costs are penalized using a transaction cost model given in Equation 3.5 and estimated using trading data. By balancing these elements, the optimization seeks to maximize the investor's utility while adhering to the specified constraints, see Equations 3.3 and 3.4.

The optimization process will be conducted under specific initial conditions and constraints. We assume that portfolios maintain a constant level of tracking error (γ_{TE}) and assets under management (AuM) throughout the observed period, though different γ_{TE} and AuM values will also be considered. The transaction cost penalty parameter δ and γ_{TE} are not mutually independent, as increasing δ reduces the turnover of the optimal portfolio in an optimization. Hence, the resulting optimal portfolio's tracking error will be smaller and the tracking error constraint γ_{TE} won't apply. Hence, we will fix δ to a value of 0.001, which we found to be reasonable for most portfolios in the previous chapter. The portfolios will be long only, subject to a tracking error constraint with respect to the value-weighted market portfolio benchmark. Turnover constraints will be implemented in the form of not allowing more than 100% of mean daily volume (MDV) to be traded in either direction in a given rebalancing period. The rebalancing will be performed monthly, over a period of 288 months resulting in 288 rebalances for each factor, starting in January 2000 and ending in December 2023. We assume the initially invested position to be the market portfolio at the first period.

3.4 Factor-enhanced market portfolios

In this section, we showcase the benefits of incorporating factor portfolios into market portfolios while considering transaction costs.

3.4.1 Single-factor portfolios

We consider four different single-factor portfolios starting with the momentum factor portfolio, which construct using the last 12 month return. Following the three factor model of Fama and French (1993), we construct the size factor portfolio using market capitalization and the value factor portfolio using the book equity to market capitalization ratio. Finally, we use operating profitability to construct what we will call the quality factor portfolio, following the work of Asness, Frazzini, and Pedersen (2019). In Table 3.3, we present the summary statistics of these portfolios. We can see that the value-weighted market portfolio has yielded and average of 8.25% over the last 24 years. Momentum and Size performed poorly, yielding a mere 2.98% and 1.32% per annum with a gross Sharpe of 0.15 and 0.12, respectively. Momentum experienced a significant decline in performance during the market crashes of 2008 and 2020, which we depict in Figure 3.1. Value and Quality were the winners, with sizable returns of 17.71% and 18.20%. We observe high returns in the early 2001-2009 period, as well as the post-2020 era, which are the main contributors to these returns.

Table 3.3: Long-short factor portfolio summary statistics

We present the gross performance of the value-weighted market portfolio as well as of each long-short single-factor portfolio. The considered sample period is January 2000 to December 2023. For each portfolio we present the gross return, Sharpe ratio, volatility and turnover. All numbers are annualized.

	Market	Momentum	Value	Quality	Size
Gross return %	8.25	2.98	17.71	18.20	1.32
Volatility %	16.11	20.24	12.94	16.25	11.22
Gross Sharpe	0.51	0.15	1.37	1.12	0.12
Turnover %	6.16	274.3	186.4	160.3	69.9

Figure 3.1: Market and single-factor returns

We plot the market and single-factor portfolios' cumulative returns from January 2000 to December 2023. The y axis is logarithmically scaled for readability.

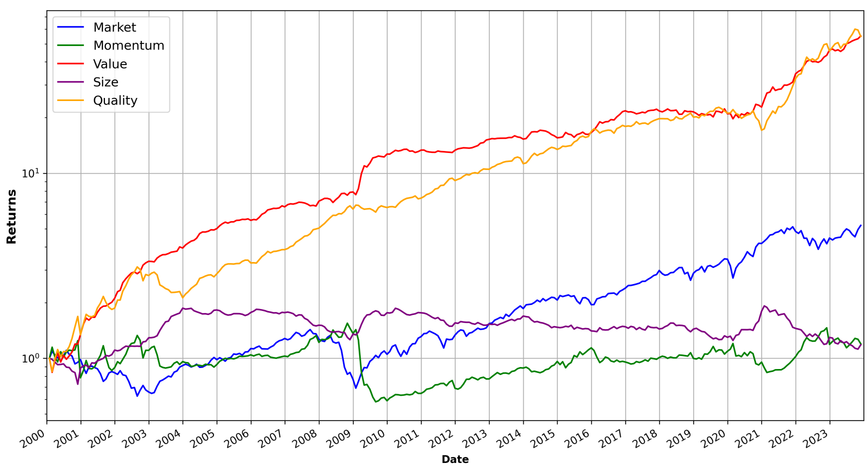

3.4.2 Portfolio performances

Having displayed the performance of standalone factors, we move onto the optimized market portfolios. Table 3.4 provides a comprehensive overview of the performance of momentum, value, quality, and size-enhanced market factor portfolios spanning from January 2000 to December 2023. This analysis is conducted across three different levels of AuM, \$500M, \$1B, and \$2B, and two tracking error constraints of 5% and 10%. The gross returns across portfolios range from 8.17% to 9.82%, with a notable decrease in returns as AuM increases, suggesting diminishing returns at higher asset levels due to increased transaction costs. Volatility is relatively stable, ranging from 15.08% to 16.68%, indicating consistent risk levels across different portfolio sizes. The gross Sharpe ratios, which measure the risk-adjusted return, peak at 0.650 for smaller AuM but decline slightly with larger AuM, indicating a potential trade-off between gross performance and transaction costs.

Turnover, reflecting the frequency of trading, is generally lower for portfolios with greater AuM, ranging from 8.55% to 25.32%. However, since the larger portfolios are more expensive to trade, they still observe higher transaction costs, which range from 0.07% to 0.54% depending on AuM size. These costs impact the net returns, which, after accounting for transaction costs, fall between 7.90% and 9.71%. As a result, the net Sharpe ratios, which incorporate transaction costs, vary from 0.475 to 0.642, showing a decreasing trend with larger AuM. Finally, the non-optimal net Sharpe ratios, calculated without considering transaction costs, are lower than their optimal counterparts, showcasing the substantial impact of transaction costs on overall portfolio performance as well as the benefits of transaction cost consideration in portfolio construction.

3.5 Conclusion

We present a comprehensive exploration of factor-enhanced market portfolios, focusing on the integration of transaction cost considerations into portfolio optimization. By leveraging real-world trading data and constructing a transaction cost model, we demonstrated that properly accounting for transaction costs significantly improves the net performance of factor-enhanced portfolios.

Our analysis reveals that while traditional factor portfolios such as momentum, value, and quality offer potential for enhanced returns, transaction costs can erode these gains if not managed effectively. The results show that our transaction cost-optimized portfolios consistently outperformed non-optimized portfolios in net returns across various asset sizes and tracking error constraints, confirming the critical importance of transaction cost consideration in portfolio construction.

Additionally, we highlight that our methodology is adaptable across different benchmark portfolios, factor exposures, and market conditions, making it versatile for a range of investment strategies. Our findings support the use of factor-enhanced market portfolios for achieving superior risk-adjusted performance, particularly when transaction costs are incor-

Table 3.4: Portfolio performances

We show the performance of momentum-, value- and quality-enhanced market factor portfolios spanning from January 2000 to December 2023. Three AuM assumptions, \$500M, \$1B, \$2B, and three tracking error constraints $\gamma_{TE} = 5, 10$ are considered. We report the annualized gross and net return, volatility, Sharpe ratio as well as turnover and associated transaction costs. Finally, we report the non-optimal Net Sharpe, which corresponds to the Net Sharpe where we applied no transaction cost penalty, $\delta = 0$.

AuM \$500M	$\gamma_{TE} = 5\%$				$\gamma_{TE} = 10\%$			
	Mom	***Value***	***Qual***	***Size***	***Mom***	***Value***	***Qual***	***Size***
Gross return (%)	8.41	9.03	8.90	8.31	8.57	9.82	9.61	8.37
Volatility (%)	16.27	15.67	16.36	16.11	16.50	15.11	16.02	15.08
Gross Sharpe	0.517	0.576	0.544	0.516	0.519	0.650	0.600	0.555
Turnover (%)	18.12	15.31	13.25	8.93	25.32	18.59	16.60	10.08
Transaction costs (%)	0.11	0.09	0.08	0.07	0.15	0.11	0.10	0.08
Net return (%)	8.30	8.94	8.82	8.24	8.42	9.71	9.51	8.29
Net Sharpe	0.511	0.571	0.539	0.511	0.510	0.642	0.594	0.550
Non-opt. Net Sharpe	0.499	0.560	0.535	0.508	0.499	0.637	0.590	0.547
AuM \$1B	$\gamma_{TE} = 5\%$				$\gamma_{TE} = 10\%$			
	Mom	***Value***	***Qual***	***Size***	***Mom***	***Value***	***Qual***	***Size***
Gross return (%)	8.38	9.00	8.85	8.24	8.55	9.70	9.55	8.33
Volatility (%)	16.31	15.61	16.41	16.06	16.45	15.17	16.07	15.02
Gross Sharpe	0.514	0.576	0.539	0.513	0.520	0.640	0.595	0.555
Turnover (%)	17.20	14.55	12.82	8.55	22.89	16.19	15.29	9.22
Transaction costs (%)	0.24	0.20	0.17	0.14	0.34	0.25	0.21	0.15
Net return (%)	8.14	8.80	8.68	8.10	8.21	9.45	9.34	8.18
Net Sharpe	0.499	0.564	0.529	0.505	0.499	0.623	0.581	0.545
Non-opt. Net Sharpe	0.478	0.548	0.517	0.497	0.481	0.608	0.567	0.528
AuM \$2B	$\gamma_{TE} = 5\%$				$\gamma_{TE} = 10\%$			
	Mom	***Value***	***Qual***	***Size***	***Mom***	***Value***	***Qual***	***Size***
Gross return (%)	8.34	8.89	8.74	8.18	8.47	9.52	9.45	8.17
Volatility (%)	16.45	15.79	16.24	16.29	16.68	15.20	16.15	15.27
Gross Sharpe	0.507	0.563	0.539	0.503	0.508	0.627	0.585	0.535
Turnover (%)	15.15	12.21	11.44	8.17	20.17	14.69	14.11	8.83
Transaction costs (%)	0.41	0.32	0.29	0.25	0.54	0.39	0.35	0.27
Net return (%)	7.93	8.57	8.45	7.93	7.93	9.13	9.10	7.90
Net Sharpe	0.482	0.542	0.521	0.487	0.475	0.601	0.564	0.518
Non-opt. Net Sharpe	0.441	0.511	0.491	0.477	0.434	0.588	0.542	0.502

porated into the portfolio optimization process. Future research could further explore the role of multi-factor strategies in enhancing market portfolios and maximizing their returns.

Appendix A

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.

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 $\varepsilon_t \overset{i.i.d}{\sim} N\left(0, \sigma^2\right)$, the joint density function for ε is:

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

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:

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

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.

Maximum likelihood

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

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

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

$$\mathbf{r_t} = \boldsymbol{\mu}_\mathbf{t} + \boldsymbol{\epsilon}_\mathbf{t}, \tag{A12}$$

$$\boldsymbol{\epsilon}_\mathbf{t} = \mathbf{H_t^{\frac{1}{2}} z_t}, \tag{A13}$$

$$\mathbf{H_t} = \mathbf{D_t R_t D_t}, \tag{A14}$$

where $\mathbf{r_t}$ is the $n \times 1$ return vector, $\boldsymbol{\mu}_\mathbf{t}$ the $n \times 1$ expected return vector, $\boldsymbol{\epsilon}_\mathbf{t}$ 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

$$\mathbf{R_t} = \mathbf{Q_t^{*-1} Q_t Q_t^{*-1}}, \tag{A17}$$

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

where

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

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

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

and u_t are the standardized errors

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

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

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

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.

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

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

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

A.4 Transaction cost data cleaning

The data is cleaned in several stages. First, we only retain trades that were executed on the market by brokers. This ensures that market impact is reflected on the underlying share price and can be measured using implementation shortfall. Additionally, if the time delay between the trade submission and the start of the execution of the trade is longer than one day for US and Europe and two days for Emerging markets, the trades are filtered out. Omitting this step would result in additional noise in implementation shortfall because of potentially large differences in price at the start of the execution and the prevailing benchmark price when the trade was submitted. Next, if there are multiple trades being executed within the same period, their transaction costs should be measured as a single trade. For this purpose, we merge trades that have overlapping arrival to end of trade dates. The merged trades' columns are then recalculated appropriately. For example, the trade size of the merged trade as a percentage of MDV is calculated as the sum of the individual trade sizes. IS is the value weighted average of the single trade IS, with volatility using the same approach.

A.5 Estimating shorting fees

This appendix explains how we compute reasonable shorting fee cost estimates based on the reported shorting fee information in Muravyev, Pearson, and Pollet (2022). In the sample period of July 2006 to December 2020, they document a mean borrow fee of 1.67% per annum with a median of 0.38%. The first percentile is 0.25%, the 10^{th} percentile is 0.28%, the 90^{th} percentile is 3% and the 99^{th} percentile is 30%. We estimate a quantile function based on the four observations given above, and we define the shorting fee quantile function as

$$\ln SF(q_i) - 0.1 = \ln a + bq_i^5 + \epsilon_i \tag{A26}$$

where q is a quantile ranging from 0 to 1 and 0.1% the smallest shorting fee assumed. Using log-linear least squares, we estimate the parameters a and b shown in Table A.1.

Table A.1: Shorting fees quantile function

We present the parameter of the shorting cost quantile function. Parameters are obtained based on quantile reports in Muravyev, Pearson, and Pollet (2022). The underlying sample period is July 2006 to December 2020. These values will be used as estimates for shorting fees based on a size rank.

Parameter	Value	St. Dev.	t-stat	P-value	95% CI	
$\ln(a)$	-1.73	0.24	-7.35	0.02	-2.74	-0.72
b	5.21	0.42	12.42	0.01	3.41	7.02

$R^2(\%)$	Adj. $R^2(\%)$	$F-stat$	$p(F-stat)$	Jarque-Bera	$p(JB)$	# of obs.
98.7	98.1	154.2	0.01	0.52	0.77	4

This functional form is motivated by the relatively low mean compared to function values in the higher quantiles, implying high skewness as mentioned in Muravyev, Pearson, and Pollet (2022). To further substantiate the validity of our quantile model in Equation (A26), we estimate the mean shorting fee using inverse transform sampling and drawing 10 million samples from our distribution. The obtained mean of the sample is 1.97%. We plot the fitted model against the data in Figure A.1. Upon obtaining the quantile function, we can then rank stocks based on size and apply the shorting fee quantile function to obtain an estimate for shorting fees, since size is shown to have a considerable impact on shorting fees as shown in Cohen, Diether, and Malloy (2007) Table I.

Figure A.1: Shorting costs model

We plot the shorting costs model estimates against the quantiles reported in Muravyev, Pearson, and Pollet (2022).

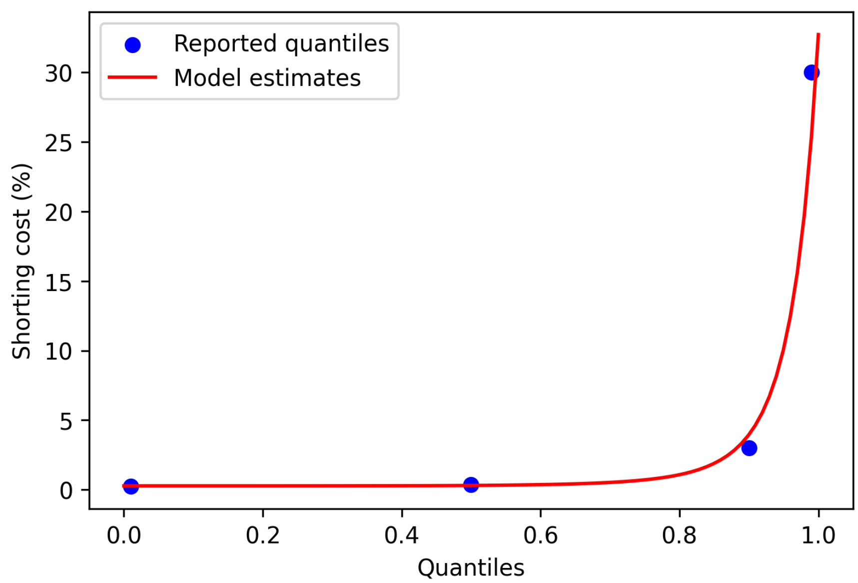

References

Almgren, Robert, Chee Thum, Emmanuel Hauptmann, and Hong Li (2005). "Direct Estimation of Equity Market Impact." *Journal of Risk* 3 (1), 5–39.

Asness, Clifford, Andrea Frazzini, and Lasse Pedersen (Mar. 2019). "Quality minus junk." *Review of Accounting Studies* 24 (2), 264–297.

Blitz, David, Bart van der Grient, and Iman Honarvar (2023). "Reversing the Trend of Short-Term Reversal." *Journal of Portfolio Management*, forthcoming.

Blitz, David and Matthias X. Hanauer (2021). "Settling the Size Matter." *Journal of Portfolio Management* 47 (2), 99–112.

Blitz, David, Matthias X. Hanauer, Iman Honarvar, Rob Huisman, and Pim van Vliet (2023). "Beyond Fama-French Factors: Alpha from Short-Term Signals." *Financial Analysts Journal* 79 (4), 96–117.

Breen, William J., Laurie Simon Hodrick, and Robert A. Korajczyk (2002). "Predicting equity liquidity." *Management Science* 48 (4), 470–483.

Brière, Marie, Charles-Albert Lehalle, Tamara Nefedova, and Amine Raboun (2019). "Stock Market Liquidity and the Trading Costs of Asset Pricing Anomalies,." Working paper.

Chen, Andrew and Mihail Velikov (2023). "Zeroing in on the Expected Returns of Anomalies." *Journal of Financial and Quantitative Analysis* 58 (3), 968–1004.

Chordia, Tarun, Amit Goyal, and Jay A. Shanken (2017). "Cross-Sectional Asset Pricing with Individual Stocks: Betas versus Characteristics,." Working paper.

Cohen, Lauren, Karl B. Diether, and Christopher J. Malloy (2007). "Supply and Demand Shifts in the Shorting Market." *Journal of Finance* 62 (5), 2061–2096.

Constantinides, George M. (1986). "Capital Market Equilibrium with Transaction Costs." *Journal of Political Economy* 94 (4), 842–862.

Davis, Mark H. A. and Andrew R. Norman (1990). "Portfolio Selection with Transaction Costs." *Mathematics of Operations Research* 15 (4), 676–713.

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 (6), 1443–1471.

DeMiguel, Victor, Alberto Martín-Utrera, Francisco J. Nogales, and Raman Uppal (2020). "A Transaction-Cost Perspective on the Multitude of Firm Characteristics." *Review of Financial Studies* 33 (5), 2180–2222.

Engle, Robert, Robert Ferstenberg, and Jeffrey Russell (2012). "Measuring and Modeling Execution Cost and Risk." *Journal of Portfolio Management* 38 (2), 14–28.

Fama, Eugene F. and Kenneth R. French (1993). "Common risk factors in the returns on stocks and bonds." *Journal of Financial Economics* 33 (1), 3–56.

Frazzini, Andrea, Ronen Israel, and Tobias J. Moskowitz (2014). "Trading Costs of Asset Pricing Anomalies,." Working paper.

— (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 (2), 461–504.

Glosten, Lawrence R. and Lawrence E. Harris (1988). "Estimating the components of the bid/ask spread." *Journal of Financial Economics* 21 (1), 123–142.

Grinold, Richard C. and Ronald N. Kahn (2000). "Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk." McGraw-Hill.

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

Jorion, Philippe (2003). "Financial Risk Manager Handbook." John Wiley & Sons.

Kagkadis, Anastasios, Ingmar Nolte, Sandra Nolte, and Nikolaos Vasilas (2023). "Factor Timing with Portfolio Characteristics." *Review of Asset Pricing Studies,* Forthcoming.

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

Korajczyk, Robert A. and Ronnie Sadka (2004). "Are momentum profits robust to trading costs?." *Journal of Finance* 59 (3), 1039–1082.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh (2020). "Shrinking the Cross-section." *Journal of Financial Economics* 135, 271–292.

Kyle, Albert S. (1985). "Continuous Auctions and Insider Trading." *Econometrica* 53 (6), 1315–1335.

Lesmond, David A., Michael J. Schill, and Chunsheng Zhou (2004). "The illusory nature of momentum profits." *Journal of Financial Economics*. Journal of Financial Economics 71 (2), 349–380.

Loeb, Thomas F. (1983). "Trading Cost: The Critical Link between Investment Information and Results." *Financial Analysts Journal* 39 (3), 39–44.

Magill, Michael J.P. and George M. Constantinides (1976). "Portfolio selection with transactions costs." *Journal of economic theory* 13 (2), 245–263.

Markowitz, Harry (1952). "Portfolio selection." *The Journal of Finance (New York)* 7 (1), 77–91.

— (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 (2), 830–846.

Muravyev, Dmitriy, Neil D. Pearson, and Joshua Matthew Pollet (2022). "Anomalies and Their Short-Sale Costs,." Working paper.

Novy-Marx, Robert and Mihail Velikov (2016). "A Taxonomy of Anomalies and Their Trading Costs." *Review of Financial Studies* 29 (1), 104–147.

Patton, Andrew J. and Brian M. Weller (2020). "What you see is not what you get: The costs of trading market anomalies." *Journal of Financial Economics* 137 (2), 515–549.

Perold, André F. (1988). "The implementation shortfall." *Journal of Portfolio Management* 14 (3), 4–9.

Roll, Richard (1992). "A Mean/Variance Analysis of Tracking Error." *The Journal of Portfolio Management* 18 (4), 13–22.

Rossi, Giuliano de, Eliad Hoch, and Michael Steliaros (2022). "The Cost of Trading Factor Strategies." *Journal of Systematic Investing* 2, 70–87.

Torre, Nicolo G. and M. Ferrari (2000). "The Market Impact Model." BARRA Inc, Berkeley.

Treynor, Jack L. and Fischer Black (1973). "How to Use Security Analysis to Improve Portfolio Selection." *The Journal of Business* 46 (1), 66–86.