# Importance of Transaction Costs for Asset Allocations in FX Markets*

Thomas A. Maurer[†] Luca Pezzo[‡] Mark P. Taylor[§]

December 12, 2019

**Abstract**

Taking transaction costs into account in a mean-variance portfolio optimization in FX markets significantly improves the achievable after costs Sharpe ratio out-of-sample. The optimization reduces trading costs while the performance before costs is unaffected. The price impact due to large buy and sell orders or market illiquidity can turn popular currency trading strategies unprofitable, while our optimized portfolios remain profitable. This is because our optimized strategy is trading less aggressively. Rules-of-thumb to tackle transaction costs – such as (i) construct equally weighted strategies instead of optimized portfolios, (ii) trade at a low frequency, (iii) restrict trading to low cost assets, (iv) only rebalance if the current position is too far from the desired position, or (v) use expected returns net of costs in the optimization – are inefficient because there are adverse effects on the (before cost) performance which dominate the savings in transaction costs.

JEL-Classification: F31, G11, G15. Keywords: Transaction Costs, Mean-Variance, Optimization, Asset Allocation, Foreign Exchange, Carry Trade.

---

*We are thankful to Jason Donaldson, Phil Dybvig, Radha Gopalan, Nishad Kapadia, Mina Lee, Hong Liu, Asaf Manela, George Panayotov, Ivelina Pavlova, Alessio Saretto, Ngoc-Khanh Tran, Michael Weber, Di Woody Wu, Guofu Zhou, seminar participants at Tulane U, U New Orleans, Washington U, and conference participants at the Eastern Finance Association, Hong Kong Joint Finance Research Workshop, Multinational Finance Society, Southwestern Finance Association, TsinghuaFinanceWorkshop for many helpful comments and suggestions.

[†]University of Hong Kong, +852 5236 2121, maurer@hku.hk.

[‡]University of New Orleans, +1 314 422 6083, lpezzo@uno.edu.

[§]Olin Business School, Washington University in St. Louis, mark.p.taylor@wustl.edu.

# 1 Introduction

There is mounting evidence in equity markets that transaction costs have a significant impact on the profitability of trading strategies. Korajczyk and Sadka (2004) estimate that the price impact of a fund with over \$5 billion under management implies trading costs which exceed the abnormal returns of momentum strategies. Lesmond et al. (2004), Novy-Marx and Velikov (2016) and Chen and Velikov (2019) document that most of the abnormal returns of 120 different stock anomalies are eroded by transaction costs (after accounting for the post-publication bias).[1]

We analyze the implications of transaction costs on the profitability of trading strategies in foreign exchange (FX) markets. We document that transaction costs can be large and significantly affect the profitability of popular currency trading strategies. We further show that a mean-variance optimization which accounts for trading costs in the optimization is able to efficiently mitigate costs without any adverse implications on the before cost performance. Our optimization ensures that the portfolio stays "close enough" to the optimal before cost balance between expected return and risk, while minimizing turnover and trading costs. Our optimization approach works well out-of-sample, to the extent that the impact of transaction costs on the after cost performance becomes almost negligible. In contrast, common rules-of-thumb to tackle transaction costs are inefficient. We document that rules-of-thumb lead to a significant deterioration in the before cost profitability and this adverse effect outweighs the savings in transaction costs.

Accordingly, the main insights of our research paper are twofold. First, despite the high trading volume in FX markets transaction costs cannot be ignored. Costs can have a first order effect on the profitability of trading strategies. Second, we provide an approach to efficiently reduce the impact of costs on the after cost performance to a negligible amount

[1] Exceptions are Frazzini et al. (2015) and DeMiguel et al. (2019). Frazzini et al. (2015) argue that large institutional investors face much smaller trading costs than what is estimated using publicly available data, and thus, they estimate a break-even fund size that is more than an order of magnitude larger than suggested by Korajczyk and Sadka (2004). DeMiguel et al. (2019) argue that in a portfolio of anomaly strategies trades across anomalies are often offsetting, and thus, transaction costs for a portfolio of anomalies are significantly lower than the sum of the costs of individual anomaly strategies.

without a loss in the before cost profitability. Hence, it is important for investors to optimally tackle transaction costs in the portfolio construction. However, once investors efficiently mitigate costs, then from a macro perspective trading costs appear unimportant. Given these insights we draw an analogy to the importance of idiosyncratic risks. Idiosyncratic risks can be a first order component of an asset's volatility. It is important for investors to tackle idiosyncratic risks and construct a diversified portfolio. However, once idiosyncratic risks are diversified away, then from a macro perspective they appear unimportant.[2]

Our paper further contributes to the theory literature on optimal portfolio choice subject to transaction costs. First, we extend the model of Dybvig and Pezzo (2019) to account for (i) directional costs, i.e., we allow costs for opening new positions to be different from those of closing/reducing existing positions, and (ii) a price impact specification which is less then linear in the size of the trades as in Frazzini et al. (2015). Second, we make a technical contribution. The solution approach of Dybvig and Pezzo (2019) requires a full rank covariance matrix. However, Maurer et al. (2018a,b) document that a robust estimate of the covariance matrix based on a principal component analysis significantly improves the out-of-sample performance of optimized currency portfolios. The principal component analysis reduces the rank of the covariance matrix, and thus, we cannot readily use the solution approach of Dybvig and Pezzo (2019). We introduce a transformation of the original problem to deal with this technical difficulty.

In the following we provide a detailed preview of our analysis and results. We investigate the profitability of a mean-variance optimized portfolio which ignores costs in its construction (denoted by $MV$), a mean-variance optimized portfolios which explicitly accounts for trading costs in the optimization (denoted by $MV_{TC}$), and characteristic sorted, equally weighted long-short currency trading strategies. In our sample the most profitable equally weighted strategy is the currency carry trade which equally borrows (lends) in currencies with low (high) interest rates. $MV$ is well-known to earn a significantly higher (and almost twice

[2] We acknowledge that idiosyncratic risk can be priced if perfect diversification is impossible (Merton, 1987). Moreover, we acknowledge that there is (puzzling) empirical evidence that idiosyncratic risk is negatively related to expected returns (Ang et al., 2006).

as high) Sharpe ratio before costs than the the carry trade (Baz et al., 2001; Della Corte et al., 2009; Daniel et al., 2017; Ackermann et al., 2016; Maurer et al., 2018a,b). For this reason most of our analysis focuses on the comparison between $MV_{TC}$ and $MV$, and to a lesser extend on inferior equally weighted currency trading strategies. The problem with $MV$ is that it has fine tuned portfolio weights which are sensitive to time-series variation in conditional expected returns and covariances which leads to a high turnover and trading costs. We document that $MV$'s turnover and costs are more than five times those of the carry trade.

In contrast to $MV$, $MV_{TC}$ reduces its trading activity in the presence of transaction costs. $MV_{TC}$ is trading off the costs of rebalancing against the desire to optimally balance the before cost expected return and risk of the portfolio. If trading costs are high, $MV_{TC}$ does not trade much and may end up holding a portfolio which is far away from the optimal combination of expected return and risk. Since $MV$ rebalances its portfolio weights at all costs (to achieve the optimal combination of expected return and risk), it is intuitive to expect that the before cost Sharpe ratio of $MV$ dominates $MV_{CT}$. Surprisingly, we do not find evidence in the data to support this intuition. Out-of-sample the before cost Sharpe ratio of $MV$ and $MV_{TC}$ are statistically indistinguishable. $MV$ has a before cost Sharp ratio of 1.26 (0.96) while it is 1.27 (1.04) for $MV_{TC}$ for our data of 29 developed and emerging (15 developed) currencies from 1976 to 2016.

As expected we find in the data that transaction costs are larger for $MV$ than for $MV_{TC}$. The difference is driven by a significant reduction in the trading activity and the turnover of $MV_{TC}$. Throughout most of our paper we use bid-ask spreads as a measure of transaction costs. Using this measure of costs, $MV_{TC}$ pays 1.25% (0.85%) of the portfolio value per year in transaction costs which is substantially less than the 5.13% (2.19%) paid by $MV$ for our set of 29 developed and emerging (15 developed) currencies. Combining this with the finding about the before cost performance we find that $MV_{TC}$ outperforms $MV$ after transaction costs. $MV_{TC}$ has an after cost Sharpe ratio of 1.16 (0.95), while the Sharpe ratio of $MV$ is 0.78 (0.75) for our set of 29 developed and emerging (15 developed) currencies. The differences are statistically significant and economically meaningful. Note that these Sharpe ratios are out-of-sample. Other moments of the return distributions are similar across the two strategies, and if anything, $MV_{TC}$ has a smaller crash risk exposure.

In order to achieve the superior performance of $MV_{TC}$, it is important to properly account for correlations between assets. Our model suggests that if assets are positively correlated, then the no trading region of $MV_{TC}$ is larger than the one of $MV_{TC\setminus Corr}$, a strategy which accounts for transaction costs in the optimization but assumes that assets are uncorrelated when constructing the no trading region. The construction of $MV_{TC\setminus Corr}$ is in the same spirit as Liu (2004) who derives an exact solution for the optimal portfolio in a continuous-time model with transaction costs for multiple uncorrelated assets. Indeed, we confirm our theoretical insight in the data that the no trading region of $MV_{TC}$ is larger than the one of $MV_{TC\setminus Corr}$ and we show that this is empirically relevant. $MV_{TC}$ incurs lower trading costs and achieves a significantly higher after cost Sharpe ratio than $MV_{TC\setminus Corr}$. This is an important contribution to the literature because it demonstrates that we should not make the assumption that assets are uncorrelated in order to simplify the optimization problem. This is unfortunate because the model of Liu (2004) is the only continuous-time (multi period) portfolio choice problem with many risky assets and subject to direct transaction costs for which we have an exact solution.[3]

We further investigate whether the issue of transaction costs can be mitigated if we (i) construct equally weighted long-short strategies instead of an optimized portfolio, (ii) trade at a low frequency, (iii) remove assets with relatively high transaction costs from the set of admissible assets, (iv) only rebalance if the current position is too far from the desired position, or (v) use expected returns net of costs in the optimization. First, if there is time-series variation in the investment opportunity set, it is reasonable to expect that an equally weighted long-short strategy has a lower turnover and trading costs than a mean-

[3] Garleanu and Pedersen (2013, 2016) solve dynamic problems in the presence of an arbitrary covariance structure but restrict costs to the price impact, i.e., they do not consider proportional costs from a bid-ask spread, for instance.

variance optimized portfolio with fine tuned weights.[4] Second, we expect that trading at a lower frequency will reduce trading activities and costs. Third, we expect that transaction costs decrease if we restrict trading to assets with relatively low costs. Fourth, if we only rebalance when the current position is too far from the desired position then we mechanically trade less frequently and reduce costs. Finally, if we use expected returns net of costs in the optimization we take less extreme positions, which should reduce turnover and trading costs.

We confirm in our data that these four intuitive rules-of-thumb reduce the turnover and transaction costs. Nevertheless, we find that they are inefficient because the performance significantly worsens if we deviate from the optimal expected return-risk tradeoff, decrease the trading frequency, or restrict the asset universe. That is, all four rules-of-thumb have large unintended adverse implications on the performance, which dominate the intended savings in trading costs. Overall, our findings advise against the use of intuitive rules-of-thumb to mitigate transaction costs. The out-of-sample performance is significantly better if we invest in a fully optimized portfolio which accounts for costs in the optimization. This is an interesting contribution because the empirical finance literature often relies on such rules-of-thumb to argue that costs are of second order importance. Our findings also provide important guidance for practitioners.

In the last section of the paper we explore the additional implications of a price impact. Large trading orders or low liquidity move the execution price outside the current bid-ask spread, and thus, potentially increase trading costs. We do not have data to accurately estimate the price impact of trading. Therefore, we investigate its implications in a sensitivity analysis. Our goal is to explore how large the price impact has to be for trading strategies to become unprofitable. We show that $MV_{TC}$ performs well after costs even if an investor

[4] To illustrate this point compare a mean-variance optimized portfolio to an equally weighted portfolio which sorts assets according to expected returns and buys the top and sells the bottom quintiles. A time-series variation in the covariance matrix will only affect the portfolio weights of the optimized portfolio while the weights in the equally weighted portfolio remain constant. Moreover, small changes in expected returns will likely leave the allocation of the equally weighted portfolio unchanged while the weights of the optimized portfolio may substantially change. In these scenarios the optimized portfolio has more turnover and higher costs than the equally weighted strategy.

has a large portfolio size and executes large buy and sell orders or FX markets are illiquid and the price impact of trading is severe. In contrast, the returns of $MV$ and many equally weighted currency trading strategies are quickly eroded by trading costs if an investor faces a severe price impact. Furthermore, we show that on average it is optimal to trade less in the presence of proportional costs and price impact than if there were only proportional costs. This is an important contribution in that the theory is silent on the size and direction of the optimal trades in the presence of a price impact (since they crucially depend on the price impact parameters). Hence, to the best of our knowledge, we are the first to empirically document the average behavior of the optimal mean-variance strategy in the presence of proportional costs and a price impact.

A limitation of our optimized portfolio $MV_{TC}$ is that it is derived in a single period model. Thus, it is myopic and in general it is suboptimal in a multi-period setting. In Appendix A.6 we discuss heuristic adjustments to account for this shortcoming. We find it difficult to improve the outstanding performance of $MV_{TC}$ even if we allow for adjustments which feature a look ahead bias. Therefore, we conclude that the myopic feature of $MV_{TC}$ does not seem to be a first order problem and $MV_{TC}$ is the best strategy that we know to address transaction costs.

Transaction costs are decreasing over time and larger among emerging currencies. Traders who specialize in more developed currencies may be tempted to ignore such costs when constructing mean-variance efficient portfolios. In Appendix A.2 we show how our cost-optimized strategies still remain considerably better even if we start the implementation after the introduction of the Euro. Moreover even when transaction costs are low during normal times, they can substantially increase during crises and become relevant (Mancini et al., 2013; Karnaukh et al., 2015). Many currency traders have shifted their focus to emerging and frontier markets because exchange rate forward discounts (and expected returns) among developed currencies are close to zero in the past decade (see Figure 1). Transaction costs in emerging and frontier markets are larger than the costs considered in our analysis. Hence, the implications of transaction costs on the optimal portfolio choice are even more important for these traders than what we report. More generally, FX markets are more liquid and have a higher trading volume than stock and other asset markets. Thus, we expect that our empirical findings set a lower bound for the importance to optimize over transaction costs for traders in other asset markets. Moreover, carry trade strategies are known to outperform stock markets over the past four decades. Therefore, FX markets do not only provide a useful environment to study mean-variance efficient portfolios but they are also among the most important markets to investors.

Besides the aforementioned literature our paper is related to the literature on portfolio optimization in the presence of transaction costs. Due to the complexity of the problem most of the literature solves frameworks with only two or three assets, either directly (Taksar et al., 1988; Davis and Norman, 1990; Dumas and Luciano, 1991; Shreve and Soner, 1994; Balduzzi and Lynch, 1999, 2000; Liu and Loewenstein, 2004; Abel et al., 2013; Campanale et al., 2015; Buss and Dumas, 2017), through the indirect martingale approach (Goodman and Ostrov (2010) and Schachermayer (2017) for a comprehensive summary of this approach), or approximately (Leland, 2000; Donohue and Yip, 2003; Muthuraman and Kumar, 2006; Irle and Prelle, 2008; Myers, 2009; Lynch and Tan, 2009). In the case of many assets Liu (2004) provides an exact solution in a continuous time model under the strong assumption that assets are uncorrelated, Garleanu and Pedersen (2013) and Garleanu and Pedersen (2016) find closed form solutions in discrete and continuous time settings in the presence of return predictability but only consider costs in terms of a price impact. Dybvig and Pezzo (2019) and Brandt et al. (2009) exactly solve one period models in the presence of an arbitrary covariance structure and number of risky assets.[5] DeMiguel et al. (2019), building on the setup of Brandt et al. (2009), also provide a myopic mean-variance setup. They solve for optimal deviations from a benchmark by investing in a small number of assets (portfolios of stock characteristics). We build on the setup of Dybvig and Pezzo (2019) since it does not require a sophisticated calibration in a pre-sample, directly chooses the optimal weights in

[5]The setup of Brandt et al. (2009) faces a “curse of dimensionality” problem if the set of assets to be used is too large.

the assets we want to trade in, adapts to the actual asset-specific cost structure of the FX markets, and scale up nicely with the number of assets.

# 2 Theory

Following Dybvig and Pezzo (2019) we work with a single period model. We ignore potentially interesting implications of a dynamic model to keep the solution exact and tractable. Therefore, our strategy is myopic which in general leads to a suboptimal outcome. However, we show that our myopic strategy has a strong out-of-sample performance in the data.

The investment opportunity set at time $t$ consists of one risk-free asset with risk-free rate of return $r_{f,t}$ and $N_t$ risky assets with conditional expected excess returns over the risk-free rate $\mu_{\mathbf{t}}^{\mathbf{e}}$ and conditional covariance matrix $\mathbf{V_t}$. If there are no transaction costs, then a myopic investor with mean-variance preferences with risk aversion $\lambda$ selects the $N_t$-vector of risky asset portfolio weights $\theta_{\mathbf{t}}^{\mathbf{MV}} = \arg\max_{\{\theta_{\mathbf{t}} \in \mathbb{R}^N\}} \left\{ \theta_{\mathbf{t}}{}' \mu_{\mathbf{t}}^{\mathbf{e}} - \frac{\lambda}{2} \theta_{\mathbf{t}}{}' \mathbf{V_t} \theta_{\mathbf{t}} \right\}$ to maximize her 1-period utility at time $t$. The optimal investment in the $N_t$ risky assets is $\theta_{\mathbf{t}}^{\mathbf{MV}} = \frac{1}{\lambda} \mathbf{V_t}^{-1} \mu_{\mathbf{t}}^{\mathbf{e}}$ and the investment in the risk-free asset is $\theta_{0,t}^{MV} = 1 - \mathbf{1}'_{\{N \times 1\}} \theta_{\mathbf{t}}^{\mathbf{MV}}$ where $\mathbf{1}_{\{N_t \times 1\}}$ is a $N_t$-vector with all elements equal to 1 (Markowitz, 1952). We denote this strategy by $MV$.

In what follows we describe the strategies that take into account transaction costs in the mean-variance optimization. We defer to Appendix B.1 for the required proofs of existence and uniqueness of solutions and to Appendix B.3 for the description of the algorithm to implement them.

## 2.1 Baseline Model

We denote the mean-variance strategy where the investor takes into account transaction costs by $MV_{TC}$. Let $\theta_{\mathbf{t}}^{\mathbf{0}}$ be the $N_t$-vector of initial weights before trading. The investor has to choose by how much to increase ($\mathbf{\Delta}_{\mathbf{t}}^{\mathbf{P}+} \geq 0$) or decrease ($\mathbf{\Delta}_{\mathbf{t}}^{\mathbf{S}+} \geq 0$) open long positions (positive weights in $\theta_{\mathbf{t}}^{\mathbf{0}}$), and by how much to increase ($\mathbf{\Delta}_{\mathbf{t}}^{\mathbf{S}-} \geq 0$) or decrease ($\mathbf{\Delta}_{\mathbf{t}}^{\mathbf{P}-} \geq 0$)

open short positions (negative weights in $\theta_\mathbf{t}^\mathbf{0}$). The allocation after trading at time $t$ is given by the weight vector $\theta_\mathbf{t} = \theta_\mathbf{t}^\mathbf{0} + \boldsymbol{\Delta}_\mathbf{t}^\mathbf{P+} + \boldsymbol{\Delta}_\mathbf{t}^\mathbf{P-} - \boldsymbol{\Delta}_\mathbf{t}^\mathbf{S+} - \boldsymbol{\Delta}_\mathbf{t}^\mathbf{S-}$. Increasing a long position in asset $i$ at time $t$ by $\boldsymbol{\Delta}_\mathbf{i,t}^\mathbf{P+}$ entails a cost of $\mathbf{C}_\mathbf{i,t}^\mathbf{P+}\boldsymbol{\Delta}_\mathbf{i,t}^\mathbf{P+}$, where $\mathbf{C}_\mathbf{i,t}^\mathbf{P+}$ is the per dollar or proportional cost associated to a trade of size $\boldsymbol{\Delta}_\mathbf{i,t}^\mathbf{P+}$. $\mathbf{C}_\mathbf{i,t}^\mathbf{P+}$ and $\boldsymbol{\Delta}_\mathbf{i,t}^\mathbf{P+}$ are the $i$-th elements of the $N_t$-vectors $\mathbf{C}_\mathbf{t}^\mathbf{P+}$ and $\boldsymbol{\Delta}_\mathbf{t}^\mathbf{P+}$. Thus, $\mathbf{C}_\mathbf{t}^{\mathbf{P+}\prime}\boldsymbol{\Delta}_\mathbf{t}^\mathbf{P+}$ captures the costs to open new long positions across all assets. Similarly, $\mathbf{C}_\mathbf{t}^{\mathbf{S+}\prime}\boldsymbol{\Delta}_\mathbf{t}^\mathbf{S+}$ describes the costs to close long positions, and $\mathbf{C}_\mathbf{t}^{\mathbf{S-}\prime}\boldsymbol{\Delta}_\mathbf{t}^\mathbf{S-}$ respectively $\mathbf{C}_\mathbf{t}^{\mathbf{P-}\prime}\boldsymbol{\Delta}_\mathbf{t}^\mathbf{P-}$ the costs to open respectively close short positions. Costs are asset-specific and proportional to the size and direction of the trades reducing the portfolio return by $\sum_{z\in\{P+,P-,S+,S-\}} \boldsymbol{\Delta}_\mathbf{t}^{\mathbf{z}\prime}\mathbf{C}_\mathbf{t}^\mathbf{z}$. All transaction costs are functions of the sizes of the directional trades ($\boldsymbol{\Delta}_\mathbf{t}^\mathbf{P+}, \boldsymbol{\Delta}_\mathbf{t}^\mathbf{P-}, \boldsymbol{\Delta}_\mathbf{t}^\mathbf{S+}, \boldsymbol{\Delta}_\mathbf{t}^\mathbf{S-}$) and we assume that there are no fixed costs.

Our baseline setting is a straightforward extension of the case studied by Dybvig and Pezzo (2019) where costs to adjust long and short positions are identical, that is, $\mathbf{C}_\mathbf{t}^\mathbf{P+} = \mathbf{C}_\mathbf{t}^\mathbf{P-}, \mathbf{C}_\mathbf{t}^\mathbf{S+} = \mathbf{C}_\mathbf{t}^\mathbf{S-}$. The optimization problem is:

**Problem 1 (Strategy $MV_{TC}$ baseline)**

$$\max_{\{\boldsymbol{\Delta}_\mathbf{t}^\mathbf{P+},\boldsymbol{\Delta}_\mathbf{t}^\mathbf{P-},\boldsymbol{\Delta}_\mathbf{t}^\mathbf{S+},\boldsymbol{\Delta}_\mathbf{t}^\mathbf{S-}\}} \left\{\theta_\mathbf{t}{}'\mu_\mathbf{t}^\mathbf{e} - \frac{\lambda}{2}\theta_\mathbf{t}{}'\mathbf{V}_\mathbf{t}\theta_\mathbf{t} - \sum_{z\in\{P+,P-,S+,S-\}} \boldsymbol{\Delta}_\mathbf{t}^{\mathbf{z}\prime}\mathbf{C}_\mathbf{t}^\mathbf{z}\right\}$$

$$\begin{aligned} s.t. \quad \theta_\mathbf{t} &= \theta_\mathbf{t}^\mathbf{0} + \boldsymbol{\Delta}_\mathbf{t}^\mathbf{P+} + \boldsymbol{\Delta}_\mathbf{t}^\mathbf{P-} - \boldsymbol{\Delta}_\mathbf{t}^\mathbf{S+} - \boldsymbol{\Delta}_\mathbf{t}^\mathbf{S-} \\ \mathbf{0} &\leq \boldsymbol{\Delta}_\mathbf{t}^\mathbf{P+}, \\ 0 &\leq \Delta_{i,t}^{P-} \leq -\min(\theta_{i,t}^0, 0) \textit{ for every asset } i, \\ 0 &\leq \Delta_{i,t}^{S+} \leq \max(\theta_{i,t}^0, 0) \textit{ for every asset } i, \\ \mathbf{0} &\leq \boldsymbol{\Delta}_\mathbf{t}^\mathbf{S-}. \end{aligned}$$

In our data $\mathbf{C}_\mathbf{t}^\mathbf{P-} \leq \mathbf{C}_\mathbf{t}^\mathbf{P+}$ and $\mathbf{C}_\mathbf{t}^\mathbf{S+} \leq \mathbf{C}_\mathbf{t}^\mathbf{S-}$ (element by element). This implies that whenever it is optimal to buy more of asset $i$ (either to close currently open short positions $\Delta_{i,t}^{P-} > 0$ or to open new long positions $\Delta_{i,t}^{P+} > 0$), the investor first closes the open short positions (if any) and then opens new long positions. Similarly, when it is optimal to sell more of the asset the investor first closes open long positions (if any) and then opens new short positions.

The mean-variance setup without transaction costs ($\sum_{z\in\{P+,P-,S+,S-\}} \mathbf{C_t^{z}}'\mathbf{\Delta_t^z} = 0$) is a special case of Problem 1. The portfolio of strategy $MV$ is independent of the initial position $\theta_{\mathbf{t}}^{\mathbf{0}}$, and it is always optimal to trade all the way to $\theta_{\mathbf{t}}^{\mathbf{MV}}$. In contrast, if there are transaction costs ($\sum_{z\in\{P+,P-,S+,S-\}} \mathbf{C_t^{z}}'\mathbf{\Delta_t^z} > 0$), then $\theta_{\mathbf{t}}^{\mathbf{MV_{TC}}}$ crucially depends on the origin $\theta_{\mathbf{t}}^{\mathbf{0}}$. Intuitively, there is a trade-off between paying transaction costs and utility gains when moving towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$. If the initial allocation $\theta_{\mathbf{t}}^{\mathbf{0}}$ is close enough to $\theta_{\mathbf{t}}^{\mathbf{MV}}$, it is optimal not to trade at all since the marginal cost required to move towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$ is higher than the marginal utility. Thus, there is a no trading region around $\theta_{\mathbf{t}}^{\mathbf{MV}}$. If the initial allocation $\theta_{\mathbf{t}}^{\mathbf{0}}$ is far enough from $\theta_{\mathbf{t}}^{\mathbf{MV}}$ (i.e., outside of the no trading region), then it is optimal to move towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$ but only until $\theta_{\mathbf{t}}^{\mathbf{MV_{TC}}}$. This is because the marginal utility of moving towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$ is diminishing, and at the boundary of the no trading region, where $\theta_{\mathbf{t}}^{\mathbf{MV_{TC}}}$ is located, the marginal utility is equal to the relevant entries of the marginal costs $\mathbf{C_t^{P+}}, \mathbf{C_t^{P-}}, \mathbf{C_t^{S+}}$ and $\mathbf{C_t^{S-}}$.

Without loss of generality, Figure 3 illustrates the optimal solution to Problem 1 in a simplified setting with two risky assets (and one risk-free asset) and $\mathbf{C_t^{P+}} = \mathbf{C_t^{P-}} = \mathbf{C_t^{S+}} = \mathbf{C_t^{S-}} > 0$.[6] The horizontal axis describes the weight placed on asset 1 and the vertical axis the weight on asset 2. The weight on the risk-free asset is 1 minus the sum of the weights on the two risky assets. The green rectangle represents the optimal portfolio $\theta_{\mathbf{t}}^{\mathbf{MV}}$ if there were no transaction costs. The blue parallelogram surrounding $\theta_{\mathbf{t}}^{\mathbf{MV}}$ defines the no trading region when the two assets are positively correlated. If the initial allocation $\theta_{\mathbf{t}}^{\mathbf{0}}$ is inside the no trading region (i.e., within the blue parallelogram), then there is no trade and $\theta_{\mathbf{t}}^{\mathbf{MV_{TC}}} = \theta_{\mathbf{t}}^{\mathbf{0}}$, because the marginal cost to trade towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$ exceeds the marginal utility.

If the initial portfolio $\theta_{\mathbf{t}}^{\mathbf{0}}$ lies outside of the no trading region, then the investor wants to move towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$ but stops trading once she reaches the boundary of the no trading

[6] In Appendix B.2 we refine this example by allowing directional proportional costs $\mathbf{C_t^{P+}} \neq \mathbf{C_t^{P-}}$ and $\mathbf{C_t^{S+}} \neq \mathbf{C_t^{S-}}$.

region. The arrows indicate the direction of trade and the arrow heads show how far to trade. Purple, brown or orange colors of the arrows indicate that only asset 1, 2 or both assets are traded.

If we start anywhere beyond the corners of the no trading region it is optimal to trade in both assets (red arrows) until we reach the closest corner. The optimal portfolio $\theta_{\mathbf{t}}^{\mathbf{MV_{TC}}}$ is exactly at one of the corners of the no trading region. For example, if $\theta_{\mathbf{t}}^{\mathbf{0}}$ lies beyond the bottom, right corner we are under-weighted in asset 2 and over-weighted in asset 1. It is optimal to buy some units of asset 2 ($\mathbf{\Delta_{2,t}^{P}} > 0$) and sell some units of asset 1 ($\mathbf{\Delta_{1,t}^{S}} > 0$) until we reach the bottom, right corner of the blue no trading region.[7]

If we start anywhere else outside the no trading region it is optimal to trade only one asset at a time. For example, if $\theta_{\mathbf{t}}^{\mathbf{0}}$ lies in the area labeled "sell 1" we are over-weighted in asset 1 and thus it is optimal to reduce our position buy selling units of that asset (along the dark red arrows) until we reach the closest boundary of the blue no trading region.

Finally, if we construct the no trading region around $MV$ ignoring the correlation between the assets, then it reduces to the yellow square. For instance, Liu (2004) assumes uncorrelated assets to derive an exact solution for the optimal dynamic trading strategy. We denote this approximate solution by $MV_{TC\setminus Corr}$ and provide details about the formal problem in Section 2.4.

If the two assets are positively correlated, then the no trading region of $MV_{TC}$ is larger along the $-45°$ line than the one of $MV_{TC\setminus Corr}$. This is because the two assets are substitutes if they are positively correlated, while they are not substitutable if they are uncorrelated. Note that if the two assets were perfect substitutes (i.e. a correlation equal to 1 and identical volatilities), then selling asset 1 and at the same time buying asset 2 would leave our risk exposure unaffected. In the same spirit, if the two assets are imperfect substitutes (i.e. correlation between 0 and 1), then there is less benefit in selling one and at the same time buying the other asset than if they are not substitutable at all (i.e. correlation equal to

[7]Since $\mathbf{C_t^{P+}} = \mathbf{C_t^{P-}}$ in our simplified example we only have one type of purchase cost, which we define as $\mathbf{C_t^{P}}$, and one type of buy trade, which we define as $\mathbf{\Delta_t^{P}}$. Similarly, since $\mathbf{C_t^{S+}} = \mathbf{C_t^{S-}}$ we only have one type of sale cost, which we define as $\mathbf{C_t^{S}}$, and one type of sell trade, which we define as $\mathbf{\Delta_t^{S}}$.

0). Since an initial position $\theta_\mathbf{t}^\mathbf{0}$ close to the $-45°$ line requires the investor to buy one and sell the other asset, the marginal utility from trading towards $\theta_\mathbf{t}^\mathbf{MV}$ is smaller and the no trading region larger if the two assets are positively correlated than if they are uncorrelated. Conversely, a similar argument can be applied to the case of a negative correlation, and the no trading region of $MV_{TC}$ is smaller along the $-45°$ line than the one of $MV_{TC\setminus Corr}$.

## 2.2 Price Impact

When trades are large enough to distort prices, which is possible for a large institutional trader like BlackRock, the optimization problem should take into account the price impact of the chosen trades. In order to capture this effect, for each type of trade $z \in \{P+, P-, S+, S-\}$ we add the price impact component $\frac{1}{2}PI_t(z, \mathbf{\Delta_t^z})\mathbf{\Delta_t^z}$ to the proportional cost $\mathbf{C_t^{z\prime}}\mathbf{\Delta_t^z}$.[8] Following the price impact literature we use two parametric functional forms: the linear price impact function

$$\frac{1}{2}PI_t(z, \mathbf{\Delta_t^z}) \equiv \frac{1}{2}PI_t^L(z, \mathbf{\Delta_t^z}) = \frac{1}{2}\mathbf{\Delta_t^{z\prime}}\mathbf{\Pi_t^{z,L}}$$

as in (Novy-Marx and Velikov, 2016; Garleanu and Pedersen, 2013), and the square-root price impact function

$$\frac{1}{2}PI_t(z, \mathbf{\Delta_t^z}) \equiv \frac{1}{2}\sqrt{PI}_t(z, \mathbf{\Delta_t^z}) = \frac{1}{2}\sqrt{\mathbf{\Delta_t^z}}'\mathbf{\Pi_t^{z,SR}}$$

as in (Frazzini et al., 2015). $\mathbf{\Pi_t^{z,L}}$ and $\mathbf{\Pi_t^{z,SR}}$ are positive definite diagonal matrices containing the asset-specific price impact parameters. $\frac{1}{2}PI_t(z, \mathbf{\Delta_t^z})\mathbf{\Delta_t^z}$ measures the extra cost to the investor due to price movements induced by the submitted order $\mathbf{\Delta_{i,t}^z}$. For example, an order to open $\mathbf{\Delta_{i,t}^{P+}}$ new long positions puts upward pressure on the ask price, and thus, the per dollar costs to open new long positions in asset $i$ increases by $\frac{1}{2}PI_t(P+, \mathbf{\Delta_t^{P+}})_i$, either

[8]Pre-multiplying by $\frac{1}{2}$ makes the units comparable with the other terms in the mean-variance utility function. Also the number 2 cancels out in the first order condition when the price impact is a linear function of the size of the trades as defined below.

linearly if $PI_t = PI_t^L$, or less than linearly if $PI_t = \sqrt{PI}_t$. The price impact parameter $\pi_{ii,t}^{P+}$ is inversely related to the liquidity and depth of the market and positively related to the investor's portfolio size and buy and sell order amounts.[9] A large (small) portfolio size implies that changes in the portfolio allocation entail large (small) buy and sell orders, which will have a large (small) effect on the execution prices.

Estimating the price impact is difficult and there is no definite answer on whether or not larger trades have less price impact per dollar traded. This is why in our analysis we use the two different functional forms to model the price impact. The new optimization problem is:

**Problem 2 (Strategy $MV_{TC}$ with price impact)**

$$\max_{\{\mathbf{\Delta_t^{P+}},\mathbf{\Delta_t^{P-}},\mathbf{\Delta_t^{S+}},\mathbf{\Delta_t^{S-}}\}} \left\{ \theta_\mathbf{t}'\mu_\mathbf{t}^\mathbf{e} - \frac{\lambda}{2}\theta_\mathbf{t}'\mathbf{V_t}\theta_\mathbf{t} - \sum_{z\in\{P+,P-,S+,S-\}} \left( \mathbf{\Delta_t^{z\prime}}\mathbf{C_t^z} + \frac{1}{2}PI_t(z,\mathbf{\Delta_t^z})\mathbf{\Delta_t^z} \right) \right\}$$

$$\begin{aligned} s.t. \quad \theta_\mathbf{t} &= \theta_\mathbf{t}^\mathbf{0} + \mathbf{\Delta_t^{P+}} + \mathbf{\Delta_t^{P-}} - \mathbf{\Delta_t^{S+}} - \mathbf{\Delta_t^{S-}} \\ \mathbf{0} &\leq \mathbf{\Delta_t^{P+}}, \\ 0 &\leq \Delta_{i,t}^{P-} \leq -\min(\theta_{i,t}^0, 0) \textit{ for every asset } i, \\ 0 &\leq \Delta_{i,t}^{S+} \leq \max(\theta_{i,t}^0, 0) \textit{ for every asset } i, \\ \mathbf{0} &\leq \mathbf{\Delta_t^{S-}}. \end{aligned}$$

Note that Problem 2 is a generalization of Problem 1. If we set $\mathbf{\Pi_t^{P+}} = \mathbf{\Pi_t^{P-}} = \mathbf{\Pi_t^{S+}} = \mathbf{\Pi_t^{S-}} = 0$ in Problem 2 then we recover Problem 1. The economics of the new model are similar to the one in the baseline model. There is still a no trading region around $\theta_\mathbf{t}^\mathbf{MV}$, which is the same as the one implied by Problem 1, but the optimal trades starting outside of it will never reach its borders (see Dybvig and Pezzo (2019) Theorem 5). This is because the marginal utility of trading has now an additional component which only cancels when it is optimal not to trade. Thus, at the border of the no trading region the marginal utility is the same as in Problem 1 and equals the relevant entries of the marginal costs $\mathbf{C_t^{P+}}, \mathbf{C_t^{P-}}, \mathbf{C_t^{S+}}$

[9] Remember that $\mathbf{\Delta_t^z}$ is the change in the asset weight as a percentage of the portfolio value.

and $\mathbf{C_t^{S-}}$. However, if $\theta_\mathbf{t}^\mathbf{0}$ lies outside of the no trading region the additional component due to the price impact does not cancel out, making the marginal costs a function of the price impact parameters. Thus, the starting point $\theta_\mathbf{t}^\mathbf{0}$ and the price impact parameters, $\mathbf{\Pi_t^{z,L}}$ or $\mathbf{\Pi_t^{z,SR}}$, will determine the size and the direction of the optimal trades.

## 2.3 A Robust Estimator of the Covariance Matrix

Maurer et al. (2018a,b) show that eliminating the principal components of $\mathbf{V_t}$ that explain less than 1% of the common variation of the currency returns substantially increases the out-of-sample performance of $MV$. They define the robust version of the covariance matrix as $\mathbf{\tilde{V}_t} = \mathbf{\tilde{W}_t}\mathbf{\tilde{\Lambda}_t}\mathbf{\tilde{W}_t'}$ where $\mathbf{\tilde{W}_t}$ is the matrix of eigenvectors and $\mathbf{\tilde{\Lambda}_t}$ is the diagonal matrix of eigenvalues after removing eigenvectors and eigenvalues which explain less than 1% of the common variation of the returns. The problem is that $\mathbf{\tilde{V}_t}$ is not full rank. Since the algorithm developed by Dybvig and Pezzo (2019) requires the covariance matrix $\mathbf{V_t}$ to be full rank, we need to re-write Problem 2 in a slightly more general form as:

**Problem 3 (Strategy $MV_{TC}$ with price impact and PCA)**

$$\max_{\{\mathbf{\Delta_t^{P+}},\mathbf{\Delta_t^{P-}},\mathbf{\Delta_t^{S+}},\mathbf{\Delta_t^{S-}}\}} \left\{ \begin{array}{c} \frac{1}{2}\mu_\mathbf{t}^{\mathbf{e}\prime}\tilde{\theta}_\mathbf{t}^\mathbf{MV} - \frac{\lambda}{2}(\theta_\mathbf{t} - \tilde{\theta}_\mathbf{t}^\mathbf{MV})'\mathbf{V_t}(\theta_\mathbf{t} - \tilde{\theta}_\mathbf{t}^\mathbf{MV}) \\ -\sum_{z\in\{P+,P-,S+,S-\}}\left(\mathbf{\Delta_t^{z\prime}}\mathbf{C_t^z} + \frac{1}{2}PI_t(z,\mathbf{\Delta_t^z})\mathbf{\Delta_t^z}\right) \end{array} \right\}$$

$$\begin{aligned} s.t. \quad \theta_\mathbf{t} &= \theta_\mathbf{t}^\mathbf{0} + \mathbf{\Delta_t^{P+}} + \mathbf{\Delta_t^{P-}} - \mathbf{\Delta_t^{S+}} - \mathbf{\Delta_t^{S-}} \\ \mathbf{0} &\leq \mathbf{\Delta_t^{P+}}, \\ 0 &\leq \Delta_{i,t}^{P-} \leq -\min(\theta_{i,t}^0, 0) \textit{ for every asset } i, \\ 0 &\leq \Delta_{i,t}^{S+} \leq \max(\theta_{i,t}^0, 0) \textit{ for every asset } i, \\ \mathbf{0} &\leq \mathbf{\Delta_t^{S-}}, \end{aligned}$$

where $\tilde{\theta}_\mathbf{t}^\mathbf{MV} = \frac{1}{\lambda}\mathbf{\tilde{V}_t^{-1}}\mu_\mathbf{t}^\mathbf{e}$ and $\mathbf{\tilde{V}_t^{-1}} = \mathbf{\tilde{W}_t}\mathbf{\tilde{\Lambda}_t^{-1}}\mathbf{\tilde{W}_t'}$. When $\mathbf{\tilde{V}_t} = \mathbf{V_t}$, then $\tilde{\theta}_\mathbf{t}^\mathbf{MV} = \theta_\mathbf{t}^\mathbf{MV}$ and the problem to solve is exactly Problem 2. This is because the objective function of Problem 3 in this case is just an algebraic re-arrangement of that of Problem 2 and we have not changed any constraint. More generally, by using $\tilde{\theta}_{\mathbf{t}}^{\mathbf{MV}}$ instead of $\theta_{\mathbf{t}}^{\mathbf{MV}}$, we are constructing the no trading region using the original full rank covariance $\mathbf{V_t}$ around $\tilde{\theta}_{\mathbf{t}}^{\mathbf{MV}}$, the optimal mean-variance strategy of Maurer et al. (2018a,b). Using this transformation we retain the feasibility of the Dybvig and Pezzo (2019) setup while exploiting the superior performance of $\tilde{\theta}_{\mathbf{t}}^{\mathbf{MV}}$. The $MV$ and $MV_{TC}$ strategies that we study and present in the rest of the paper use $\tilde{\theta}_{\mathbf{t}}^{\mathbf{MV}}$ and are the solution of Problem 3 (the baseline setup is thus considered to be Problem 3 with $\mathbf{\Pi_t^{z,L}} = \mathbf{\Pi_t^{z,SR}} = \mathbf{0}$).

## 2.4 Disregarding Asset Correlations when Forming the No Trading Region

Liu (2004) shows that in a continuous-time model the assumption of uncorrelated assets greatly reduces the complexity to dynamically optimize a portfolio subject to transaction costs. This is because with uncorrelated assets we can solve $N_t$ independent problems each one associated with only one asset. The optimal solution is a rectangular no trading region (or its multidimensional analog) surrounding the solution in the setting without transaction costs.

We re-create the same insights in our myopic framework to construct $MV_{TC\backslash Corr}$. That is, we build a rectangular no trading region (or its multidimensional analog) around $MV$.[10] As a consequence, the optimal portfolio $\theta_t^{MV\backslash Corr}$ is the solution to Problem 3 with diagonal matrix $\mathbf{V_t^d}$ in place of $\mathbf{V_t}$, where the diagonal elements are equal to the diagonal elements of $\mathbf{V_t}$.

## 2.5 Model Predictions and Limitations

Our models leaves us with four theoretical predictions. First, we expect $MV$ to outperform $MV_{TC}$ if the performance is measured in returns before transaction costs. This is because,

[10] As in the example in Figure 13 the actual no trading region might have some edges cut when $\mathbf{C_t^{P+}} \neq \mathbf{C_t^{P-}}$ or $\mathbf{C_t^{S+}} \neq \mathbf{C_t^{S-}}$ and the point $(0, 0)$ is inside the no trading region.

by definition, $MV$ is the optimal portfolio when evaluated before transaction costs (in the single period model). Second, we expect transaction costs to be larger for $MV$ than for $MV_{TC}$. Third, we expect $MV_{TC}$ to outperform $MV$ after transaction costs. Moreover, the size of these differences between $MV$ and $MV_{TC}$ are expected to depend on the size of the no trading region of $MV_{TC}$. In turn, the no trading region is expected to be increasing in the size of the transaction costs. Fourth, if assets are correlated, then we expect that the no trading region of $MV_{TC}$ is larger than that of $MV_{TC\backslash Corr}$. Thus, we expect $MV_{TC}$ to outperform $MV_{TC\backslash Corr}$. We quantify and assess the empirical importance of these four predictions and compare the performances of $MV$, $MV_{TC}$ and $MV_{TC\backslash Corr}$ for our baseline setup in Section 4.1 and conduct a separate analysis for the price impact in Section 4.3.

Finally, as explained earlier our theory model assumes a single investment period. A multi-period model is complex for two reasons. First, changes in the investment opportunity set lead to a hedging demand (Merton, 1971). Second, since the optimal portfolio crucially depends on the initial portfolio $\theta_t^0$, in a multi-period model an investor should take into account how the portfolio at time $t$ will affect the initial position at time $t+1$, which leads to a non-trivial adjustment of the no trading region. We are not aware of a tractable solution to solve the second problem, unless we make additional strong assumptions such as imposing that all assets are uncorrelated as in Liu (2004) or restricting costs to be purely quadratic as in Garleanu and Pedersen (2013, 2016). Our empirical results suggest that our myopic solution performs well in the data and in Appendix A.6 we argue that multi-period adjustments do not seem of first order importance.

# 3 FX Markets

The investment strategies $MV$ (ignoring transaction costs in the optimization), $MV_{TC}$ (taking into account costs in the optimization) and $MV_{TC\backslash Corr}$ (taking into account costs in the optimization but constructing the no trading region assuming uncorrelated assets) are based on the mean-variance optimization outlined in Section 2. In order to construct mean-variance efficient portfolios that perform well out-of-sample, we need sensible estimates of conditional expected returns and the covariance matrix. Estimation errors are a well-known problem in the portfolio optimization literature and can lead to a bad out-of-sample performance of optimized portfolios (Brandt, 2005). For instance, DeMiguel et al. (2009) show that in the US stock market an equally weighted portfolio outperforms mean-variance optimized portfolios out-of-sample due to estimation errors.

FX markets are special. First, forward discounts are good proxies for conditional expected excess returns of currency trades because exchange rate growths are well-described by a random walk (Meese and Rogoff, 1983). Second, there is a strong factor structure to describe the covariance matrix (Lustig et al., 2011). This is helpful to reduce estimation errors. These properties are exploited in several recent papers and mean-variance optimized portfolios in FX markets are shown to be very profitable in out-of-sample analyses (Baz et al., 2001; Della Corte et al., 2009; Daniel et al., 2017; Ackermann et al., 2016; Maurer et al., 2018a,b). We follow this literature and implement $MV$, $MV_{TC}$ and $MV_{TC\backslash Corr}$ in FX markets to quantify the importance of accounting for transaction costs in the construction of mean-variance optimized portfolios.

## 3.1 Investment Opportunity Set in FX Markets

We denote spot and 1-month forward exchange rates as USD (US-dollar) per unit of currency $i$ at time $t$ by $X_{i,t}$ and $F_{i,t}$. Following the literature, we define the 1-month realized currency return between currency $i$ and the USD (denominated in USD) by

$$r_{i,t+1} \equiv \ln\left(\frac{X_{i,t+1}}{F_{i,t}}\right) = fd_{i,t} + \Delta x_{i,t+1},$$

where $fd_{i,t} = \ln\left(\frac{X_{i,t}}{F_{i,t}}\right)$ (known at time $t$) is the forward discount, and $\Delta x_{i,t+1} = \ln\left(\frac{X_{i,t+1}}{X_{i,t}}\right)$ (realized at time $t+1$) is the exchange rate growth. $r_{i,t+1}$ is the excess return (over the risk-free rate in USD) of entering an uncovered long position in the 1-month forward exchange rate contract.[11]

We use currency returns for $N_t$ currencies (against the USD) as our universe of $N_t$ risky assets. Due to data availability the number of currencies $N_t$ changes through time. The excess returns from time $t$ to $t+1$ of strategies $MV$, $MV_{TC}$ and $MV_{TC\backslash Corr}$ are $r'_{t+1}\theta_{\mathbf{t}}^{\mathbf{MV}}$, $r'_{t+1}\theta_{\mathbf{t}}^{\mathbf{MV_{TC}}}$ and $r'_{t+1}\theta_{\mathbf{t}}^{\mathbf{MV_{TC\backslash Corr}}}$, where $r_{t+1}$ is the vector of excess returns of all $N_t$ currency returns.

We set the initial weight of asset $i$ at time $t$ equal to the optimal portfolio at time $t-1$ plus any changes due to realized returns. That is, for generic strategy $j$, $\theta_{\mathbf{i,t}}^{\mathbf{0,j}} = \theta_{\mathbf{t-1,i}}^{\mathbf{j}} + \theta_{\mathbf{t-1,i}}^{\mathbf{j}} r_{i,t}$, where $\theta_{t-1,i}^{j}$ is the optimal weight of asset $i$ in the portfolio at time $t-1$ according to strategy $j$ and $r_{i,t}$ is the realized return of asset $i$ from time $t-1$ to $t$.

The constructions of $MV$, $MV_{TC}$ and $MV_{TC\backslash Corr}$ require estimates of conditional expected excess returns $\mu_{\mathbf{t}}^{\mathbf{e}}$ and the covariance matrix $\mathbf{V_t}$. We follow the literature and use the current forward discount $fd_{i,t}$ as a proxy for conditional expected excess return $\mu_{\mathbf{i,t}}^{\mathbf{e}}$ (Baz et al., 2001; Della Corte et al., 2009; Daniel et al., 2017; Ackermann et al., 2016; Maurer et al., 2018a,b). This is motivated by the empirical finding that exchange rate changes are difficult to predict over a short horizon, i.e., $E_t\left[\Delta x_{i,t+1}\right] \approx 0$ (Meese and Rogoff, 1983).

To estimate the conditional covariance matrix $\mathbf{V_t}$ we follow Maurer et al. (2018a,b). We refer to their work for details about the estimation. First, we estimate $\mathbf{V_t}$ using an exponentially weighted moving average estimator for the covariance of daily returns. We use a decay factor of 0.97 and a 9 month window preceding month $t$ such that our estimate uses only information available prior to $t$ and the subsequent portfolio construction is out-of-sample. Second, as explained in Section 2.3, we apply a spectral decomposition to $\mathbf{V_t}$, and remove eigenvectors and eigenvalues which explain less than 1% of the common variation in returns.

[11] Under the premise of the covered interest rate parity (CIP), the forward discount is equal to the interest rate differential $fd_{i,t} = \ln\left(\frac{R_{i,t}}{R_{US,t}}\right)$ where $R_{US,t}(= e^{r_{f,t}})$ and $R_{i,t}$ are 1-month risk-free interest rates in the USD and currency $i$, and the carry trade return is equivalent to borrow $\frac{1}{R_{US,t}}$ USD and lend $\frac{1}{R_{US,t}X_{i,t}}$ units of currency $i$. Note that we do not require the CIP to hold for the construction of our portfolios or the out-of-sample performance analysis. We implement all carry trade returns using forward and spot exchange rates and do not need information about interest rates.

The constructions of $MV_{TC}$ and $MV_{TC\backslash Corr}$ further require estimates of transaction costs. We compute currency returns before and after transaction costs. We use mid exchange rate quotes for $X_{i,t}$ and $F_{i,t}$ to compute returns before transaction costs. To account for proportional transaction costs we use bid-ask quotes, indicated by superscripts $^b$ and $^a$. Since it is relatively cheap to roll a contract over from month to month, the literature typically assumes no roll-over fees and only accounts for transaction costs if there is a change in a position (Menkhoff et al., 2012; Della Corte et al., 2016; Maurer et al., 2018a). Our estimates of the per dollar transaction costs to open new long positions ($\mathbf{C_{i,t}^{P+}}$), close existing long positions ($\mathbf{C_{i,t}^{S+}}$), open new short positions ($\mathbf{C_{i,t}^{S-}}$) and close existing short positions ($\mathbf{C_{i,t}^{P-}}$) are

$$
\begin{aligned}
\mathbf{C_{i,t}^{P+}} &\equiv \ln\left(\frac{X_{i,t+t}}{F_{i,t}}\right) - \ln\left(\frac{X_{i,t+1}}{F_{i,t}^a}\right) = \ln\left(\frac{F_{i,t}^a}{F_{i,t}}\right) \\
\mathbf{C_{i,t}^{S+}} &\equiv \ln\left(\frac{X_{i,t}}{F_{i,t-1}}\right) - \ln\left(\frac{X_{i,t}^b}{F_{i,t-1}}\right) = \ln\left(\frac{X_{i,t}}{X_{i,t}^b}\right) \\
\mathbf{C_{i,t}^{S-}} &\equiv -\ln\left(\frac{X_{i,t+1}}{F_{i,t}}\right) + \ln\left(\frac{X_{i,t+1}}{F_{i,t}^b}\right) = \ln\left(\frac{F_{i,t}}{F_{i,t}^b}\right) \\
\mathbf{C_{i,t}^{P-}} &\equiv -\ln\left(\frac{X_{i,t}}{F_{i,t-1}}\right) + \ln\left(\frac{X_{i,t}^a}{F_{i,t-1}}\right) = \ln\left(\frac{X_{i,t}^a}{X_{i,t}}\right).
\end{aligned}
$$

In Appendix A.5 we implement our analysis using full round-trip costs (i.e. assume that a position is completely closed and re-opened every month) and show that this leads to substantially larger transaction costs and a quantitatively larger effect of our results. Full round-trip costs are considered too conservative and larger than the trading costs paid in practice. Thus, we consider the robustness results presented in Appendix A.5 less relevant.

It is more difficult to obtain an estimate for the price impact parameters $\mathbf{\Pi_t^{z,L}}$ and $\mathbf{\Pi_t^{z,SR}}$ for $z \in \{P+, P-, S+, S-\}$. Mancini et al. (2013) and Karnaukh et al. (2015) study liquidity in FX markets. They show that the linear price impact of trading is important and there is a substantial variation in the cross-section and the time-series. They have access to a proprietary dataset from Electronic Broking Services, which contains second-by-second bid and ask quotes, volume and information on the direction of trades for seven major currencies between January 2007 and December 2009. Unfortunately, this data is not available to us. Moreover, it covers only spot (but not forward) exchange rates and only a small subset of the currencies and the time horizon which we consider in our analysis. While we assume that there is no price impact in our baseline analysis, following the setup developed in Section 2.2, we investigate in Section 4.3 the implications of a wide range of values for $\mathbf{\Pi_t^{z,L}}$ and $\mathbf{\Pi_t^{z,SR}}$ on the performance of our trading strategies.

Figure 4 plots the time-series of the cross-sectional average of annualized costs (as a percentage of the portfolio value; not including price impact costs) for a set of 29 developed and emerging currencies (black solid line), a subsets of 14 emerging currencies (blue dashed line), and a subset of 15 developed currencies (red dotted line).[12] As expected, transaction costs to trade emerging currencies are substantially larger than developed currencies. Transaction costs generally decrease over time, except during FX market crises, which do not necessarily coincide with NBER recessions (grey shaded areas). The costs reach low levels between 0.04% and 0.06% of the portfolio value in the final year of our sample. In Appendix A.2 we show that $MV_{TC}$ continues to outperform $MV$ even in the period after the introduction of the Euro in January 1999.

Traders who specialize in developed currencies may be tempted to ignore transaction costs when constructing mean-variance efficient portfolios. However, even if transaction costs are low during normal times, they can substantially increase during crises and become relevant (Mancini et al., 2013; Karnaukh et al., 2015). Moreover, many currency traders have shifted their focus to emerging and frontier markets because exchange rate forward discounts (or expected returns) among developed currencies are close to zero in the past decade. Figure 1 illustrates the decline in the average (across currencies) of absolute forward discounts. Transaction costs in emerging and frontier markets are generally larger than the costs considered in our analysis, and thus, the implications of transaction costs on the optimal portfolio choice are even more important for these traders. In addition, FX markets

[12]The cross-sectional average of costs is computed as $\frac{1}{4\times N_t}\sum_{i=1}^{N_t}\left(\mathbf{\hat{C}_t^{P}}+\mathbf{\hat{C}_t^{P-}}+\mathbf{\hat{C}_t^{S+}}+\mathbf{\hat{C}_t^{S-}}\right)$, where $N_t$ is the number of exchange rates for which we have data available at time $t$ and the average is taken over the time $t$ cross sectional median of each type of cost $\mathbf{\hat{C}_t^{z}}$ for $z \in \{P+, P-S+, S-\}$.

are more liquid and have lower transaction costs than most other asset markets. Therefore, if we generalize our finding and consider the implications for other markets, we expect that our results can be viewed as a lower bound.

All baseline strategies ($MV$, $MV_{TC}$, $MV_{TC\setminus Corr}$) use information (i.e. estimates for $\mu_\mathbf{t}^\mathbf{e}$, $\mathbf{V_t}$, $\mathbf{C_{i,t}^z}$ $\forall z \in \{P+, S+, P-, S-\}$) available at the end of month $t$ to construct a portfolio which we then hold until the end of the subsequent month $t+1$. Thus, all returns are out-of-sample and none of the trading strategies suffers from a look-ahead bias.

Finally, we set the risk aversion coefficient $\lambda = 50$ in our portfolio optimizations. We follow Maurer et al. (2018a) who choose this value so that the mean-variance efficient portfolio has an unconditional volatility comparable to typical currency trading strategies. In Appendix A.4 we show that our results are empirically and theoretically unaffected by the choice of $\lambda$.

## 3.2 Data

We collect daily spot and 1-month forward bid, ask and mid exchange rates from Barclays Bank International and Reuters via Datastream. We use quotes of the last day of the month to compute monthly returns $r_{i,t+1}$. A concern with currencies of emerging countries is that there are capital controls and major trading frictions. Menkhoff et al. (2012) and Della Corte et al. (2016) suggest to exclude countries with a negative score on the capital account openness index of Chinn and Ito (2006).[13] Following this literature, we include currencies of 29 countries in our analysis. According to Lustig et al. (2011) 15 of them are classified as developed, while the remaining 14 are "emerging" countries. The 15 developed countries are: Australia, Belgium, Canada, Denmark, Euro Area, France, Germany, Italy, Japan, Netherlands, New Zealand, Norway, Sweden, Switzerland, United Kingdom. The 14

[13]We further exclude a currency at time $t$ if it is pegged to another currency, more than 20% of its daily exchange rate changes are missing over the past 9 months, or if the absolute value of the annualized forward discount $12 \times |fd_{i,t}|$ is larger than 25%. Forward discounts of more than 25% are rare and we believe that such large values likely indicate non-tradable outliers in the data, the presence of severe trading frictions, sizable sovereign default risk or an extraordinary large expected currency devaluation. Under these conditions, a currency trader is likely not able or willing to consider a currency as part of the investment opportunity set.

emerging countries are: Brazil, Czech Republic, Greece, Hungary, Iceland, Ireland, Mexico, Poland, Portugal, Singapore, South Africa, South Korea, Spain, Taiwan. The Euro was introduced in January 1999 and we exclude all countries which have joined the Euro after that date and only keep the Euro as a currency.

Exchange rates of all 29 currencies are quoted against the USD for the sample starting on October 11th, 1983 and ending on March 2nd, 2016. We are able to extend our sample further back to January 2nd, 1976 for the following subset of 14 countries with exchange rates quoted against the GBP (Great British Pound): Belgium, Canada, Denmark, France, Germany, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, USA. For the period from January 2nd, 1976 to October 11th, 1983 we convert all data to exchange rates quoted against the USD using mid exchange rate quotes of USD/GBP. Because the data quoted against the GBP is less reliable, in Appendix A.3 we restrict our analysis to a subsample starting in October 11th, 1983. We show that our results are robust in this subsample.

# 4 Results

In section 4.1 we show that the out-of-sample performance after transaction costs of mean-variance efficient portfolios improves if transaction costs are taken into account in the baseline optimization (i.e., Problem 3 without price impact). The outperformance after transaction costs of $MV_{TC}$ over $MV$ is due to a significant reduction in the trading activity, which reduces its implementation costs. In addition, we show that it is important to properly account for correlations between assets for the superior performance of $MV_{TC}$. This is an important contribution to the literature because it demonstrates that we should not make the assumption of Liu (2004) that assets are uncorrelated.

In section 4.2 we use the baseline setup to investigate whether the issue of transaction costs can be mitigated if we (i) construct equally weighted long-short strategies instead of an optimized portfolio, (ii) trade at a low frequency, (iii) remove assets with relatively high transaction costs from the set of admissible assets, (iv) only rebalance if the current position is too far from the desired position, or (v) use expected returns net of costs in the optimization.

While we confirm that these four intuitive rules-of-thumb reduce transaction costs, we find that they are not efficient because the (before and after cost) performance significantly worsens if we deviate from the optimal tradeoff between expected returns and risk, decrease the trading frequency, or restrict the asset universe. Overall, our findings advise against the use of an intuitive rule-of-thumb to mitigate transaction costs. The out-of-sample performance is significantly better if we invest in a fully optimized portfolio which accounts for costs in the optimization.

In section 4.3 we analyze how a price impact of trading affects the performance of our strategies in addition to effects of proportional costs. That is, we implement the version of Problem 3, letting $\mathbf{\Pi_t^{z,L}} = \mathbf{\Pi_t^{z,SR}} \neq \mathbf{0}$. $MV_{TC}$ still remain the best choice. No matter the functional form of the price impact or the set of currencies, minimizing over costs in the optimization is crucial if the price impact is not infinitesimal. $MV_{TC}$ and $MV_{TC\setminus Corr}$ are the *only* strategies that consistently perform well after costs. Moreover, the two strategies remain profitable even if an investor has a large portfolio size and executes large dollar amount buy and sell orders or FX markets are illiquid and the price impact of trading is extremely severe. In contrast, $MV$ and many popular equally weighted currency trading strategies in the literature significantly underperform, even if the price impact is moderate. Transaction costs quickly erode the returns if there is a price impact and several strategies turn unprofitable. Furthermore, we find that on average is optimal to trade less in the presence of proportional costs and price impact that what it would be optimal if costs were only proportional. This is interesting because the theory is silent on the size and direction of the optimal trades in the presence of price impact as they depend on the price impact parameters (see Section 2.2). Thus, to the best of our knowledge, we are the first to document the empirical behavior of the optimal mean-variance trades in the presence of proportional costs and price impact.

Appendix A contains additional results which illustrate the robustness of our findings. Sections A.1, A.2 and A.3 provide robustness results in subsamples considering NBER vs non-NBER time periods, samples before vs after the introduction of the Euro, and a shorter sample excluding (arguably less reliable) data before November 1983. Section A.4 provides robustness results showing that our results are indistinguishable across for various choices of the risk aversion coefficient $\lambda$. Section A.5 shows that our main results are quantitatively even more important if we consider full round-trip transaction costs. In section A.6 we discuss heuristic adjustments to $MV_{TC}$ to account for its shortcoming as a myopic strategy. We find it difficult to improve the performance of $MV_{TC}$ even if we allow for adjustments which feature a look ahead bias. We conclude that the myopic feature of $MV_{TC}$ does not seem to be a first order problem and $MV_{TC}$ is the best strategy that we know to address transaction costs.

## 4.1 Comparison of $MV_{TC}$, $MV_{TC\backslash Corr}$ and $MV$

### Performance Before Transaction Costs

From our model we get the following theoretical result in a single period model:

**Prediction 1:** *$MV$ is expected to outperform $MV_{TC}$ if the performance is measured in returns before transaction costs.*

Table 1 quantifies the difference between $MV$ and $MV_{TC}$ (computed solving Problem 3 in the absence of price impact) and summarizes the monthly out-of-sample excess returns of both strategies for our full set of 29 currencies (columns 1 and 2) and the subset of 15 developed currencies (columns 3 and 4) from January 1976 to February 2016. The first panel of Table 1 reports the Sharpe ratios (SR) and average excess returns (Mean) before transaction costs. Interestingly, the annualized Sharpe ratios of $MV$ and $MV_{TC}$ are almost identical for the set of 29 currencies, 1.26 and 1.27, while $MV_{TC}$, with a Sharpe ratio of 1.04, actually outperforms $MV$, with a Sharpe ratio of 0.96, for the set of 15 developed currencies. However, the differences in Sharpe ratios between $MV$ and $MV_{TC}$ are not significant (neither in the set of all 29 nor the 15 developed currencies). The average annual return before transaction costs of $MV_{TC}$ is 1.41% and 0.69% lower than the average return of $MV$ (denoted by $\Delta$Mean in Table 1) when measured with respect to the 29 and the 15 developed currencies respectively, but the volatility (Vol) of $MV_{TC}$ is proportionally lower only for the set of 29 currencies while much lower for the set of the 15 developed currencies, which implies almost identical Sharpe ratios before costs across the two strategies in the former case and an higher Sharpe ratio for $MV_{TC}$ in the latter case.

The top panel of Figure 5 displays the cumulative returns of $MV$ (black dashed line) and $MV_{TC}$ (red solid line) before transaction costs for our full set of 29 currencies (left panel) and for the set of 15 developed currencies (right panel). For the set of 29 currencies the two time-series closely track each other and the returns of the two strategies are almost identical at every point in time. For the case of the 15 developed currencies the cumulative returns of $MV_{TC}$ are almost always superior to the ones of $MV$ with the gap between the two becoming sensibly noticeable right after the 1981-1982 NBER recession.

To conclude, we do not find a significant difference in the performance before transaction costs between $MV$ and $MV_{TC}$. Although $MV_{TC}$ trades less actively than $MV$ due to the no trading region and theoretically holds an ex-ante "suboptimal" position[14], the out-of-sample performance before transaction costs is almost identical for the case of the 29 currencies, and economically, albeit not statistically, better for the case of the 15 developed currencies. Therefore, the theoretical prediction 1 is at best empirically irrelevant.

### Transaction Costs

Our second model implication (from the baseline setup) is:

**Prediction 2:** *The transaction costs of $MV$ are expected to be higher than $MV_{TC}$.*

We first investigate how much $MV_{TC}$ is trading compared to $MV$ to quantify the size of the no trading region. The theoretical prediction rests on the insight that if the investor

[14]Suboptimal in a single period model if there are no transaction costs.

optimizes over transaction costs (i.e. $MV_{TC}$), she trades from her initial position $\theta_{\mathbf{t}}^{\mathbf{0}}$ towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$ but stops at the boundary of the no trading region. The size of the no trading region of strategy $S \in \{MV, MV_{TC}\}$ is inversely related to the trade aggressiveness $TA_t^S$ defined as,

$$TA_t^S = \frac{\sum_i \|\theta_{\mathbf{i,t}}^{\mathbf{S}} - \theta_{\mathbf{i,t}}^{\mathbf{0,S}}\|}{\sum_i \|\theta_{\mathbf{i,t}}^{\mathbf{MV}} - \theta_{\mathbf{i,t}}^{\mathbf{0,S}}\|} \in [0,1]\,. \tag{1}$$

A large $TA_t^S$ indicates that the investor trades aggressively and chooses a position $\theta_{\mathbf{t}}^{\mathbf{S}}$ close to $\theta_{\mathbf{t}}^{\mathbf{MV}}$, which in turn implies that the no trading region is small. In the extreme case where $TA(\theta_{\mathbf{t}}^{\mathbf{S}}) = 1$, $\theta_{\mathbf{t}}^{\mathbf{S}} = \theta_{\mathbf{t}}^{\mathbf{MV}}$ and the no trading region is a singleton. In contrast, a small value indicates that the investor does not trade aggressively and $\theta_{\mathbf{t}}^{\mathbf{S}}$ is far away from $\theta_{\mathbf{t}}^{\mathbf{MV}}$, which in turn means that the no trading region is large. In the extreme case where $TA_t^S = 0$, strategy $S$ does not trade at all, $\theta_{\mathbf{t}}^{\mathbf{S}} = \theta_{\mathbf{t}}^{\mathbf{0,S}}$, and the initial position lies within the no trading region. Thus, $TA_t^S$ measures how aggressive an investor trades from the initial position $\theta_t^{0,S}$ towards the optimum without transaction costs $\theta_t^{MV}$, and its inverse quantifies the size of the no trading region of strategy $S$.

The last row of table 1 summarizes the average of the monthly $TA_t^S$ for $S \in \{MV, MV_{TC}\}$. By definition $TA_t^{MV} = 1$ and $TA_t^{MV_{TC}} \in [0,1]$. On average the trade aggressiveness of strategy $MV_{TC}$ is 0.41 for our set of 29 currencies and 0.35 for our set of 15 developed currencies. That is, the investor reduces the amount of trading by 59% or 65%. This reduction in the trading activity decreases the cost to implement $MV_{TC}$ compared to $MV$.

We further plot the time-series of the turnover $\sum_i \|\theta_{i,t}^S - \theta_{i,t}^{0,S}\|$ of $S = MV$ (black dashed line) and $S = MV_{TC}$ (red solid line) in the top panel in Figure 7 for our full set of 29 currencies (left panel) and for the set of 15 developed currencies (right panel). The turnover of $MV$ is on average almost two times and a half larger than the turnover of $MV_{TC}$. In a similar fashion, in the bottom panel of Figure 7, we report the average portfolio holdings and 1-standard deviation error bars of $MV$ (downward pointing triangles and black lines) and $MV_{TC}$ (upward pointing triangles and red lines). The average portfolio holdings are similar across the two strategies but the standard deviation is systematically larger for $MV$, which again indicates more trading activity.

Finally, we directly quantify the trading costs. The second panel in Table 1 reports the average transaction costs paid per year as a percentage of the portfolio value (or alternatively as a reduction in the portfolio return). The costs paid by $MV$ are substantial, i.e., 5.13% for the set of 29 currencies and 2.19% for the set of 15 developed currencies. That is, 39%-24% of $MV$'s expected return is lost to transaction costs. The costs paid by $MV_{TC}$ are less than one-forth the size of the costs paid by $MV$, i.e., 1.25% for the set of 29 currencies and a bit more than one-third, i.e. 0.85%, for the set of 15 developed currencies. These savings in transaction costs are economically meaningful. The difference in costs between $MV$ and $MV_{TC}$ is highly statistically significant for both the set of 29 currencies and the subset of 15 developed currencies.

Figure 6 visualizes this striking result by plotting the time-series of cumulative transaction costs (top panels) and the monthly costs (bottom panels) paid by $MV$ (black dashed line) and $MV_{TC}$ (red solid line) for our full set of 29 currencies (left panels) and for the set of 15 developed currencies (right panels). The spread between the cumulative costs of $MV$ and $MV_{TC}$ is steadily increasing, while the monthly costs incurred by $MV$ are without exception always larger than the costs of $MV_{TC}$. Note that transaction costs are decreasing over time which is consistent with the decline in average costs illustrated in figure 4. To sum up, our second theoretical prediction that $MV$ is subject to larger transaction costs than $MV_{TC}$ is empirically important.

### Performance After Transaction Costs

The main prediction of our single period model is:

**Prediction 3:** *$MV_{TC}$ is expected to outperform $MV$ after transaction costs. Moreover, the outperformance is expected to be more substantial if transaction costs are large.*

The third panel in Table 1 compares returns after transaction costs. The Sharpe ratios after transaction costs are highlighted in boldface. For the full set of 29 currencies, the annualized Sharpe ratio of $MV$ is 0.78 and the one of $MV_{TC}$ is 1.16. For the set of the 15 developed currencies , the annualized Sharpe ratio of $MV$ is 0.75 and the one of $MV_{TC}$ is 0.96. The differences of $\Delta\text{SR} = 0.38$ and 0.20 respectively are economically meaningful and almost half, one-forth respectively, of $MV$'s Sharpe ratio. $MV_{TC}$ is compensated by a 5.7%, 3.15% respectively, higher annual risk premium than $MV$ per 15% return volatility (which is roughly equal to the unconditional volatility of the value weighted US stock market index). We further find that the difference is statistically significant with a p-value of 0.01 and 0.03 respectively. We employ the test proposed by Ledoit and Wolf (2008), which uses block bootstrapping and is robust to heteroskedasticity and cross- and auto-correlation.[15]

The bottom panel in Figure 5 illustrates the striking dominance of $MV_{TC}$ by plotting cumulative returns after transaction costs. The spreads in cumulative returns after costs are steadily opening. Our results are not driven by outliers or a crisis. Our finding suggests that our third theoretical prediction is empirically important. Optimizing transaction costs when constructing mean-variance efficient portfolios substantially improves the out-of-sample performance.

As mentioned in section 2 $MV_{TC}$ is the solution in a single period model and does not take into account potentially interesting dynamics in a multi-period model. It is not obvious at the outset that $MV_{TC}$ performs well when implemented in the data with many trading dates. The result that $MV_{TC}$ performs significantly better than $MV$ and achieve a high Sharpe ratio is important.

In addition to the Sharpe ratio analysis, we investigate the (ex-post) utility when switching from $MV$ to $MV_{TC}$. The last four rows of Table 1 immediately above the trade aggressiveness measure report the annualized return or certainty equivalent $CE_\lambda$ a mean-variance investor with risk aversion $\lambda \in \{1, 5, 10, 50\}$ is willing to give up in order to switch from $MV$ to $MV_{TC}$. For our set of 29 currencies, a log investor ($\lambda = 1$) is willing to give up 0.81% to switch from $MV$ to $MV_{TC}$. For an investor with $\lambda$ equal to 5, 10 or 50, $CE_\lambda$ increases to 3.04%, 5.82% or 28.09%. For the set of 15 developed currencies, the certainty equivalents

[15]We choose a block size of 5 observations for the block bootstrapping.

are smaller, i.e., for $\lambda \in \{1, 5, 10, 50\}$, $CE_\lambda \in \{0.35\%, 1.49\%, 2.91\%, 14.27\%\}$. The monotonically increasing relation highlights that more risk averse investors have a stronger desire to manage transaction costs efficiently.

All other moments of returns after transaction costs are comparable across $MV$ and $MV_{TC}$. Table 1 lists the monthly return skewness (Skew), kurtosis (Kurt), the percentage of positive monthly returns (% Positive), the maximum draw down (MDD), which measures the maximum loss from peak to trough of the strategy in the entire sample, and the autocorrelation (AC). If anything the skewness and the MDD of $MV_{TC}$ are more favorable than the ones of $MV$, suggesting that $MV_{TC}$ has less crash risk exposure than $MV$.

For completeness, we plot the time-series of the notional values or total dollar exposures $\sum_i \|\theta^S_{i,t}\|$ of $S = MV$ (black dashed line) and $S = MV_{TC}$ (red solid line) for our full set of 29 currencies and the set of 15 developed currencies in the top, respectively bottom, panel of Figure 8. The notional value is slightly smaller for $MV_{TC}$ than for $MV$ and mostly below 10. The notional values spikes mostly during the volatile periods around the end of 70s and the first half of the 80s. Margin requirements in FX derivatives markets are low and implementing a strategy with a notional value of 20 or even 50 is unproblematic.

To sum up, we recall that the performance before transaction costs of $MV$ and $MV_{TC}$ are at most identical (if not better for $MV_{TC}$). $MV$ trades more aggressively and faces larger transaction costs than $MV_{TC}$. In other words, $MV_{TC}$ efficiently reduces unnecessary trading and allows the investor to save money. In turn, $MV_{TC}$ substantially outperforms $MV$ after transaction costs.

### Importance of Correlations between Assets

The last model prediction investigates the assumption of uncorrelated assets which is necessary in the model of Liu (2004) to derive an exact solution for the optimal dynamic trading strategy subject to transaction costs.

**Prediction 4:** *If assets are correlated, then the no trading region of $MV_{TC}$ is expected to be larger than the one of $MV_{TC\backslash Corr}$. Moreover, transaction costs of $MV_{TC}$ are expected to be lower than the costs of $MV_{TC\backslash Corr}$ (at least in the absence of price impact), and we expect $MV_{TC}$ to outperform $MV_{TC\backslash Corr}$ after costs.*

Recall that $MV_{TC\backslash Corr}$, which we introduced in Section 2.4, is the strategy which optimizes transaction costs similar to $MV_{TC}$ but assumes that assets are uncorrelated in the construction of the no trading region and obtain an approximate solution. Figure 9 shows the time-series of the average conditional correlation of each exchange rate growth $i$ with all other exchange rate growths, $\rho_{i,t} = \frac{1}{N-1}\sum_{j=1}^{N-1} Corr_t\left(\Delta x_{i,t}, \Delta x_{j,t}\right)$ for our full set of 29 currencies (top panel) and for the set of the 15 developed currencies (bottom panel). To estimate the conditional correlation $Corr_t\left(\Delta x_{i,t}, \Delta x_{j,t}\right)$ between exchange rate growths $i$ and $j$ in month $t$ we use daily exchange rate growths within the month. The bold black line is the average of all correlations $\rho_t = \frac{1}{N-1}\sum_{i=1}^{N} \rho_{i,t}$ in month $t$. Correlations $\rho_{i,t}$ are almost always positive and on average close to 0.5. The average correlation $\rho_t$ is always between 0.18 and 0.85, which is clearly different from zero.

Table 2 compares the monthly excess returns of $MV$, $MV_{TC\backslash Corr}$ and $MV_{TC}$ for our full set of 29 currencies from 1976 to 2016, while Table 3 shows the same analysis for the set of the 15 developed currencies. Note that $MV$ and $MV_{TC}$ are also described in Table 1. Consistent with the previous finding, the average returns and Sharpe ratios before transaction costs are similar across the three strategies for the set of 29 currencies. For the set of 15 developed currencies the Sharpe ratio of $MV_{TC}$ is at least 0.08 higher, due to the ability of the strategy to achieve a lower volatility without giving away too much expected return.

$MV_{TC\backslash Corr}$ has a trade aggressiveness of 0.70, 0.72 respectively, and transaction costs of 2.34%, 1.63% respectively, per year for the set of 29 currencies and the set of 15 developed currencies. $MV_{TC\backslash Corr}$ trades less aggressively than $MV$ and is able to save 2.79%, 0.56% respectively, in costs compared to $MV$. On the other hand, $MV_{TC\backslash Corr}$ trades considerably more aggressively than $MV_{TC}$, with $\Delta TA = 29\%$ and 37% respectively (both statistically significant at the 1% level). This means that the no trading region of $MV_{TC}$ is 29%(37%)

bigger, resulting in average annual cost savings of 0.96%, respectively 0.78%, over $MV$.

After transaction costs, the Sharpe ratio of $MV_{TC\setminus Corr}$ is 1.01, 0.79 respectively, which is 0.23, 0.04 respectively, higher than the ratio of $MV$ but 0.15, 0.17 respectively, lower than the ratio of $MV_{TC}$. The differences in Sharpe ratios between $MV$ and $MV_{TC\setminus Corr}$ and $MV_{TC\setminus Corr}$ and $MV_{TC}$ are statistically and economically significant (p-values of 0.03 and 0.04) for the set of 29 currencies. For the set of the 15 developed currencies only the difference in Sharpe ratio between $MV_{TC\setminus Corr}$ and $MV_{TC}$ is statically and economically significant (p-value of 0.04). Therefore, accounting for correlations in the optimization is important to significantly increase the Sharpe ratio when minimizing over transaction costs. So much that this is the only thing that matters in the superior performance of $MV_{TC}$ for the set of the 15 developed currencies. Employing the approximate solution $MV_{TC\setminus Corr}$ to optimize over transaction costs is suboptimal.

The certainty equivalent $CE_\lambda$ an investor with risk aversion $\lambda \in \{1, 5, 10, 50\}$ is willing to pay to switch from $MV_{TC\setminus Corr}$ to $MV$ is always negative (i.e., $MV_{TC\setminus Corr}$ is preferred to $MV$). The certainty equivalent for a switch from $MV_{TC\setminus Corr}$ to $MV_{TC}$ are positive, implying that investors prefer $MV_{TC}$. All higher order return moments are comparable across the three strategies and if anything $MV_{TC}$ is preferred especially with regards to crash risk exposure.

We conclude that the no trading regions of $MV_{TC\setminus Corr}$ and $MV_{TC}$ are not only theoretically but also quantitatively different. Accounting for correlations when optimizing transaction costs is empirically important and the out-of-sample out-performance of $MV_{TC}$ over $MV_{TC\setminus Corr}$ is economically and statistically significant. This empirical finding is an important contribution to the literature. We should not make the assumption of Liu (2004) that assets are uncorrelated in order to simplify the optimization problem. This assumption is a costly mistake (at least in the context of FX markets) and leads to a significant reduction in the out-of-sample performance.

## 4.2 Rules-of-Thumb to Address Transaction Costs

We investigate whether intuitive rules-of-thumb are helpful to mitigate the issue of transaction costs. First, if there is time-series variation in the investment opportunity set, it is reasonable to expect that an equally weighted long-short strategy has a lower turnover than a mean-variance optimized portfolio with fine tuned weights. The portfolio weights of a characteristic sorted, equally weighted long-short strategy are in general less sensitive to time-series variation in forward discounts and covariances than those of a mean-variance optimized portfolio.[16] Thus, it is reasonable to expect that the turnover and transaction costs are smaller for the equally weighted strategy than an optimized portfolio. We check whether we can reduce transaction costs and improve the performance after costs if we build equally weighted portfolios instead of $MV$ or $MV_{TC}$. We consider five long-short currency strategies which are well-known in the literature. Second, we expect that trading at a lower frequency will reduce transaction costs. We repeat our analysis but change the monthly holding period to either weekly or quarterly and study how the trading frequency affects the performance. Third, we expect that transaction costs decrease if we restrict trading to assets with relatively low costs. We repeat our analysis but remove assets with high costs from the set of admissible assets. Fourth, following the idea of Novy-Marx and Velikov (2016) we impose a rule that we only rebalance the weight in asset $i$ if $\|\theta_{i,t}^{MV} - \theta_{i,t}^{0}\|$ is larger than some ad hoc threshold value. This rule ensures that there is less frequent trading and a position is only rebalanced if the current position is too far away from the desired position, and thus, the utility gain from the rebalancing is arguably large. Finally, we use expected returns net of costs to construct $MV$, which arguably leads to less extreme positions and a smaller notional value. In turn, we expect that a smaller notional value implies a smaller turnover and lower transaction costs.

While we confirm that these four intuitive rules-of-thumb help to reduce transaction costs, we find that they are inefficient because the (before and after cost) performance significantly worsens. Our finding advises against the use of intuitive rules-of-thumb to

[16] See the illustrative example in footnote 4.

mitigate transaction costs because deviating from the optimal tradeoff between expected return and risk, decreasing the trading frequency, or restricting the asset universe have adverse consequences on the performance. The out-of-sample performance is significantly better if we invest in a fully optimized portfolio which accounts for costs in the optimization. This is one of the key contributions of our paper.

### Equally Weighted Portfolios

It is common in the literature to rank assets based on characteristics and construct equally weighted long-short portfolios. Intuitively, this helps to mitigate two problems: (i) estimation errors (which lead to an optimistic in-sample but bad out-of-sample performance) and (ii) transaction costs. With regard to the first problem DeMiguel et al. (2009) demonstrate that optimized portfolios in the US stock market perform poorly and yield a less attractive out-of-sample performance than equally weighted portfolios. However, as explained in section 3 estimation errors are not problematic in FX markets. A growing literature shows that mean-variance optimized portfolios achieve a significantly better out-of-sample performance (before costs) than equally weighted portfolios (Baz et al., 2001; Della Corte et al., 2009; Daniel et al., 2017; Ackermann et al., 2016; Maurer et al., 2018a,b).

We focus on the second problem, i.e., transaction costs. The portfolio weights of a characteristic sorted, equally weighted long-short strategy are in general less sensitive to time-series variation in forward discounts and covariances than an optimized portfolio with fine tuned weights. Thus, it is reasonable to expect that the turnover and transaction costs are smaller for the equally weighted strategy than an optimized portfolio. We construct five equally weighted currency strategies which are well-known to deliver high returns. The dollar strategy ($DOL$) borrows in the USD and equally invests in all other currencies. The dynamic dollar ($DDOL$) takes a long position in $DOL$ if the median exchange rate forward discount is positive, and a short position otherwise. The carry ($HML$) sorts currencies according to the forward discount into quintiles and borrows in the bottom and invests in the top quintile. The momentum ($MOM$) sorts currencies according to their past 12 month performance into quintiles and borrows in the bottom and invests in the top quintile. The value ($VAL$) sorts currencies according to the power purchase parity adjusted exchange rate into quintiles and borrows in the top quintile (overvalued currencies with high real exchange rates) and invests in the bottom quintile (undervalued currencies with low real exchange rates).

Figure 2 and Table 4 compare the out-of-sample Sharpe ratios of $MV_{TC}$, $MV_{TC\setminus Corr}$, $MV$, $DOL$, $DDOL$, $HML$, $MOM$, $VAL$. In Figure 2 the out-of-sample Sharpe ratios before (after) costs are illustrated by blue (red) bars and the difference between before and after cost ratios, which indirectly measures transaction costs, by yellow bars. The figure is generated using monthly returns of our full set of 29 currencies (top panel) and those of the set of the 15 developed currencies (bottom panel) from 1976 to 2016. For the set of 29 and 15 currencies columns 2 and 5 (3 and 6) in table 4 provide the Sharpe ratios before (after) costs. Columns 4 and 7 report the differences between the Sharpe ratios after costs of $MV_{TC}$ and all other strategies for the set of 29 and 15 currencies. Moreover, columns 4 and 5 in Table 7 and 8 report the average annual turnover and transaction costs of all the strategies for the set of 29 currencies and 15 currencies respectively.

We confirm our intuition that the costs for the equally weighted portfolios are significantly smaller than the costs of the optimized portfolios. Among the five equally weighted portfolios the $DOL$ has the lowest average annual cost of 0.10%, 0.03% respectively, and $MOM$ the highest costs of 1.14%, 1.04% respectively. In comparison $MV$ and $MV_{TC}$ have average annual costs of 5.13%, 2.19% respectively, and 1.25%, 0.85% respectively. However, the Sharpe ratios (before and after costs) of the equally weighted portfolios are economically and statistically significantly smaller than the Sharpe ratio of $MV_{TC}$. For the set fo 29 (15) currencies $HML$ yields the highest after cost Sharpe ratio of 0.60 (0.59) among the five equally weighted portfolios. $MV$ and $MV_{TC}$ achieve significantly higher after cost Sharpe ratios of 0.78 and 1.16 (0.75 and 0.96) for the set of 29 (15) currencies. Despite the savings in transaction costs the rule-of-thumb to construct equally weighted portfolios is not efficient because a deviation from the optimal tradeoff between expected returns and risk has a dominant negative effect on the performance. Therefore, we advise against the use of this rule-of-thumb.

**Trading Frequencies: Weekly, Monthly, Quarterly**

It is intuitive that frequent trading leads to more turnover and higher transaction costs. We investigate how the performances of our portfolios are affected if we change the trading frequency from monthly to either weekly or quarterly. We expect that the problem of transaction costs is more (less) severe at a weekly (quarterly) frequency. Therefore, a simple rule-of-thumb suggests that we should trade at a quarterly frequency so that we do not have to worry about transaction costs.

Table 5 and 6 provide the analogous results to table 4 but for weekly and quarterly trading frequencies. Moreover, table 7 and table 8 report the average annual turnover and transaction costs. We observe that the costs to implement $MV$ and the equally weighted portfolios are the highest at the weekly frequency and the lowest at the quarterly frequency. For the set of 29 currencies and the set of 15 respectively, $MV$ has an average annual cost of 11.18%, 5.23% respectively, at the weekly, 5.13%, 2.19% respectively, at the monthly and 2.46%, 1.04% respectively, at the quarterly frequency. Among the five equally weighted portfolios average annual costs also decrease, even if much less so in proportion, as we change the trading frequency from weekly to quarterly. This confirms our intuition that reducing the trading frequency decreases costs. Surprisingly, there is no monotonic relationship between the trading frequency and the turnover and costs in the case of $MV_{TC}$. The average annual costs of $MV_{TC}$ is 0.36%, 0.25% respectively, at the weekly, 1.25%, 0.85% respectively, at the monthly and 1.16%, 0.72% respectively, at the quarterly frequency. Therefore, the rule-of-thumb to trade at a lower frequency reduces trading costs (except in the case of $MV_{TC}$).

We further observe that the trading frequency do not always has a significant effect on the before cost Sharpe ratios of the optimized portfolios. While for the set of the 29 currencies, the Sharpe ratios of $MV_{TC}$, $MV_{TC\backslash Corr}$ and $MV$ are the highest at the monthly frequency, i.e., between 1.26 and 1.27, against an average level between 0.87 and 1.07 at the other frequencies. For the set of the 15 developed currencies there is no substantial difference between those at the monthly and weekly frequencies while the Sharpe ratios at the quarterly frequency appear lower. Finally, we observe that the after cost Sharpe ratio of $MV_{TC}$ is higher at the monthly trading frequency than either at the weekly or quarterly frequency. Moreover, $MV_{TC}$ dominates the after cost Sharpe ratio of $MV$ and $MV_{TC \setminus Corr}$ at any frequency.

The trading frequency does not appear to have a significant effect on the before cost Sharpe ratios of the five equally weighted strategies. While this is a desirable feature of the equally weighted strategies, we find that the $MV_{TC}$ dominates all other trading strategies in terms of Sharpe ratios after costs at any trading frequency.

To sum up, we find that trading at a lower frequency decreases transaction costs but changing the trading frequency has a first order effect on the performance. Overall, it is not advisable to trade at the quarterly frequency and the after cost Sharpe ratio of $MV_{TC}$ is higher at the monthly trading frequency than either at the weekly or quarterly frequency.

### Removing Assets with High Transaction Costs

Our third rule-of-thumb builds on the intuition that transaction costs decrease if we restrict trading to assets with low costs and exclude high costs assets from the set of admissible assets. We focus on the monthly trading frequency and report results for the set of 29 currencies, and the set of 15 developed currencies respectively. In table 9, table 13 respectively, we show Sharpe ratios before (after) costs of $MV$ and $MV_{TC}$ in columns 2 and 3 (4 and 5). The difference between the after cost Sharpe ratios of $MV_{TC}$ and $MV$ and the corresponding p-value are provided in columns 6 and 7. The first row reports the performance of $MV$ and $MV_{TC}$ constructed from the full set of 29(15) currencies. The second row reports the results when we reduce our set of admissible assets by one asset and drop the currency with the highest median transaction cost over the past 9 months.[17] Each subsequent row reduces the set of admissible assets by an additional asset and row $i$ removes the $i-1$ currencies

[17]Changing the window length does not affect our results.

with the highest median costs over the past 9 months. Note that the data availability changes through time and we do not have data for all 29(15) currencies in every month. The maximum number of currencies we can drop is 10(7) so that we have at least two currency to construct a portfolio in every month.

First, we confirm our earlier finding that before cost Sharpe ratios are very similar between $MV_{TC}$ and $MV$. Second, we observe that the difference between before and after cost Sharpe ratios are much larger for $MV$ than for $MV_{TC}$ independently of the number of currencies dropped. For instance the before and after cost difference for $MV$ is 4.36 times, 2.63 respectively, the one for $MV_{TC}$ for the set of 29 currencies, and the set of 15 currencies respectively, if we do not drop any currency. Such gap stabilizes around 2 times as we start removing currencies.

The differential between after cost Sharpe ratios of $MV$ and $MV_{TC}$ is narrowing as we drop currencies (column 6 in table 9 and table 13). The differential is 0.38(0.21) if we do not drop any currency, 0.21(0.20) if we drop the 4 most expensive currencies and 0.01(0.17) if we drop the 10(7) most expensive currencies for the set of 29(15) currencies. The differential is significant if we drop up to 6(3) currencies but turns insignificant thereafter. This is consistent with our intuition that removing high cost assets from the set of admissible assets decreases costs and an investor does not need to bother to optimize costs. That is, when we only trade assets with low costs then the costs and after cost performance of $MV$ and $MV_{TC}$ are indistinguishable.

We further document an interesting average decline in the before and after cost Sharpe ratios of $MV_{TC}$ as we drop more currencies. This effect can be clearly seen for the full set of 29 currencies where the before(after) cost Sharpe ratio monotonically drops from 1.27(1.16) to 1.06(0.93) if we remove the 5 most expensive currencies, and then averages around 1.03(0.93) if we remove the 10 most expensive currencies. But it is also noticeable by comparing the before and after cost Sharpe ratios starting from the set of 29 currencies with those starting from the set of 15 currencies. The before(after) cost average Sharpe ratio drops from 1.10(0.99) to 1.04(0.95), with the before(after) cost Sharpe ratios of $MV_{TC}$

starting with the full set of 29 currencies and removing up to two currencies being at least 0.12(0.08) higher than the best before(after) cost Sharpe ratios achievable starting from the set of the 15 developed currencies.

The decrease in the Sharpe ratios before costs is a first order effect and dominates the savings in transaction costs, leading to a decline in Sharpe ratios after costs. Therefore, it is not efficient to follow a rule-of-thumb to remove high cost currencies to mitigate transaction costs because shrinking the set of currencies has an adverse effect on the performance which outweighs the decrease in transaction costs. It is more profitable to work with the full set of 29 admissible assets and construct a fully optimized portfolio which takes into account transaction costs.

Thus, focusing on the full set of 29 currencies, in Tables 10 and 11 we provide more details and decompose the Sharpe ratios to investigate why removing high cost currencies significantly depresses the performance. First, as we drop high cost currencies (and keep the risk aversion coefficient $\lambda$ constant) we observe a decrease in the notional value for $MV$ and $MV_{TC}$. We further document a general decrease in average returns and volatilities of $MV$ and $MV_{TC}$ as we drop costly currencies. Average returns decrease faster than volatilities and thus Sharpe ratios are declining. This is consistent with the intuition that an investor reduces his position in the risky portfolio and invests more in the risk-free asset (in our case the risk-free 1-month USD bond) as the tradeoff between expected return and risk worsens. The skewness of $MV$ and $MV_{TC}$ does not display a regular pattern but the kurtosis in general increases. Maximum draw downs are less severe if we drop costly currencies due to the fact that the notional values decrease. More importantly, if we compare maximum draw downs per 1% expected return (i.e. we divide the column of Table 11labeled "After TC MDD" by that of Table 10 labeled "After TC Mean"), we observe that the crash risk of $MV$ deteriorates at a much faster rate than the one for $MV_{TC}$, which eventually converges from above as we shrink the set of admissible assets.[18] Similar to the observed decrease

[18]Note that maximum draw down per 1 % expected return is a measure in a similar spirit like the Sharpe ratio because it compares risk and expected return. It is a more sensible comparison because both expected returns and maximum draw down (and volatility) scale linearly with leverage.

in the Sharpe ratios of $MV_{TC}$, the increase in the maximum draw down per 1% expected returns suggests a worsening in the tradeoff between expected returns and (crash) risk as we drop costly currencies from the set of admissible assets. Finally, in Table 12 we investigate the return distributions of the currencies which we successively drop in Table 9. In the $i$th row we construct a "portfolio" which consists of only the currency with the $i$th highest average transaction costs over the past 9 months. The table reports the average transaction costs, Sharpe ratio, average return, volatility, average correlation with all other currencies, skewness and kurtosis. We do not observe a statistically significant correlation between transaction costs and any of the six characteristics. In other words, it does not appear to be the case that currencies with high costs deliver a more (or less) attractive tradeoff between expected returns than currencies with low costs. Therefore, our evidence suggests that as we shrink the set of admissible currencies the benefits of diversification decreases, which has an adverse effect on the portfolio performance (and decreases the Sharpe ratio and increases the maximum draw down per 1% expected return for $MV_{TC}$). On the flip side, our findings suggest that if we had a larger original set of assets, it is possible that the decrease in diversification and the adverse effects on the performance are less severe as we remove just the first few of the most expensive currencies. Our finding that it is not efficient to follow a rule-of-thumb to remove high cost currencies to mitigate costs may be due to the fact that our original set of currencies is relatively small. In this case we expect a bell shaped relation on the performance as we sequentially eliminate the most expensive currencies from the tradable universe.

We conclude that in our data it is not efficient to follow a rule-of-thumb to remove high cost currencies to mitigate transaction costs because shrinking the set of currencies has an adverse effect on the performance which outweighs the decrease in transaction costs. The adverse effects seem to be due to a deterioration in the portfolio diversification as the set of admissible assets becomes smaller. Thus, it is more profitable to work with the original full set of 29 admissible assets and construct a fully optimized portfolio which takes into account transaction costs.

### No Trading Rule based on the Trade Size

In our fourth rule-of-thumb we follow the idea of Novy-Marx and Velikov (2016) and only rebalance the weight in asset $i$ if $\|\theta_{i,t}^{MV} - \theta_{i,t}^{0}\|$ is larger than some ad hoc threshold value. Thus, a position is less frequently rebalanced and trading only occurs if the expected utility gain is large enough because our current position is relatively far away from the optimal balance between expected return and risk. This effect is stronger for larger threshold values. Since such strategies are in general known as "sS" rules we denote them as $MV_{sS}$.

Columns 2 to 5 in Tables 14 and 15 compare the before and after cost Sharpe ratios of the $MV_{sS}$ strategies to the $MV_{TC}$ for the various threshold values listed in the first column. Table 14 presents the results for the set of 29 and Table 15 for 15 currencies at the monthly frequency. The last two columns of the tables report the annualized costs and turnover of the $MV_{sS}$ strategies. Imposing a threshold of 0 is the same as implementing strategy $MV$. Thus, the first row of the tables report the earlier discussed comparison between $MV$ and $MV_{TC}$.

As expected a higher threshold implies that we rebalance positions less frequently, and the turnover and costs are lower. As we increase the threshold value from 0 to 1, the turnover drops form 40.47 (27.6) to 9.59 (8.1) and trading costs drop from 5% (2%) to 2% (1%) for our set of 29 developed and emerging (15developed) currencies. However, a higher threshold value is also associated with a drop in the before cost Sharpe ratio because our portfolio is not optimally trading off expected return and risk. The adverse effect on the before cost performance dominates the savings in trading costs, and thus, this rule-of-thumb is inefficient. We observe a significantly smaller after cost Sharpe ratio for $MV_{sS}$ compared to $MV_{TC}$ for any threshold value between 0 and 1.

### Conditional Expected Returns Net of Costs

To address transaction costs Burnside et al. (2008) and Burnside et al. (2011) sort currencies based on forward discounts adjusted for bid-ask spreads to construct $HML$. That is, they lend in currencies with high $\ln\left(\frac{X^b_{i,t}}{F^a_{i,t}}\right)$ and borrow in currencies with low $\ln\left(\frac{X^a_{i,t}}{F^b_{i,t}}\right)$. Notice that if we take a long position in asset $i$ and pay full round-trip costs to open and close it, then the expected return reduces to $\ln\left(\frac{X^b_{i,t}}{F^a_{i,t}}\right) = r_{i,t+1} - \mathbf{C^{P+}_{i,t}} - \mathbf{C^{S+}_{i,t}}$.[19] Because the expected return of an asset is the expected loss to a short seller, when we take a short position in asset $i$ and pay full round-trip costs to open and close it, then the expected loss of short selling increases to $\ln\left(\frac{X^a_{i,t}}{F^b_{i,t}}\right) = r_{i,t+1} + \mathbf{C^{S-}_{i,t}} + \mathbf{C^{P-}_{i,t}}$. Thus, the idea of Burnside et al. (2008) and Burnside et al. (2011) is to sort currencies based on expected returns net of costs to construct $HML$.

In our last rule-of-thumb we follow this idea and construct strategy $MV_{Net}$ which is the same as strategy $MV$ except that we use expected returns net of full round-trip costs $\hat{\mu}^{\mathbf{e}}_{\mathbf{t}}$ as an input instead of the before cost expected returns $\mu^{\mathbf{e}}_{\mathbf{t}}$. We define

$$\theta_{\mathbf{t}}^{\mathbf{MV_{Net}}} = \arg\max_{\theta_{\mathbf{t}}} \left\{ \theta_{\mathbf{t}}{}' \hat{\mu}^{\mathbf{e}}_{\mathbf{t}} - \frac{\lambda}{2} \theta_{\mathbf{t}}{}' \mathbf{V_t} \theta_{\mathbf{t}} \right\},$$

with $\hat{\mu}^{\mathbf{e}}_{\mathbf{i,t}} = \mu^{\mathbf{e}}_{\mathbf{i,t}} - \mathbf{1}_{\left\{\theta^{\mathbf{MV_{Net}}}_{\mathbf{i,t}} \geq \mathbf{0}\right\}} \left(\mathbf{C^{P+}_{i,t}} + \mathbf{C^{S+}_{i,t}}\right) + \mathbf{1}_{\left\{\theta^{\mathbf{MV_{Net}}}_{\mathbf{t}} < \mathbf{0}\right\}} \left(\mathbf{C^{S-}_{i,t}} + \mathbf{C^{P-}_{i,t}}\right)$.[20] The solution can be obtained by iteratively computing a sequence of standard mean-variance weights until $\hat{\mu}^{\mathbf{e}}_{\mathbf{t}}$ converges. This is complex and only feasible with few assets since it requires iterating on all possible combinations of the assets for which the weights change signs from one iteration to the next.

Hence, there are two ways to attempt a solution to this problem. We either (i) recognize that $MV_{Net}$ is the solution to Problem 1 (or Problem 3 if we center the no-trading region around $\tilde{\theta}^{MV}_t$ as we do in our baseline setup) under the unrealistic assumption of round-trip costs,[21] or (ii) approximate the solution by prematurely stopping the iterations before the curse of dimensionality becomes too costly (in terms of computing power).

[19] The costs to close a position are only known in the future when the position is closed. We use the current costs to close a position as an approximation for the expectation of the future costs.

[20] Adjusting the second moment of the return distribution do not materially affect the results. We do not report these results but they are available upon request.

[21] The mapping with Problem 1 is obtained by setting $\theta^{\mathbf{0}}_{\mathbf{t}} = \mathbf{0}$ and setting the cost to open $\mathbf{\Delta^{P+}_{i,t}}$ new long positions in asset $i$ equal to $\mathbf{\Delta^{P+}_{i,t}} \left(\mathbf{C^{P+}_{i,t}} + \mathbf{C^{S+}_{i,t}}\right)$ and the costs to open $\mathbf{\Delta^{S-}_{i,t}}$ new short position equal to $\mathbf{\Delta^{S-}_{i,t}} \left(\mathbf{C^{S-}_{i,t}} + \mathbf{C^{P-}_{i,t}}\right)$.

The purpose of this section is to give a simple rule-of-thumb to solve the problem from the perspective of an investor who does not have access to our optimization, otherwise there is no point in solving for $MV_{Net}$ under the wrong cost structure when an exact and efficient solution for the correct problem is available (namely $MV_{TC}$). Therefore we provide a simple algorithm to get an approximate solution for $\theta_{\mathbf{t}}^{\mathbf{MV_{Net}}}$. First, we use

$$\theta_{\mathbf{t}}^{\mathbf{MV_{Net}^{(0)}}} = \frac{1}{\lambda}\mathbf{V_t^{-1}}\mu_{\mathbf{t}}^{\mathbf{e}},$$

to construct $\hat{\mu}_{\mathbf{i,t}}^{\mathbf{e,(1)}} = \mu_{\mathbf{i,t}}^{\mathbf{e}} - \mathbf{1}_{\left\{\theta_{\mathbf{i,t}}^{\mathbf{MV_{Net}^{(0)}}} \geq \mathbf{0}\right\}}\left(\mathbf{C_{i,t}^{P+}} + \mathbf{C_{i,t}^{S+}}\right) + \mathbf{1}_{\left\{\theta_{\mathbf{t}}^{\mathbf{MV_{Net}^{(0)}}} < \mathbf{0}\right\}}\left(\mathbf{C_{i,t}^{S-}} + \mathbf{C_{i,t}^{P-}}\right)$. Second, we construct

$$\theta_{\mathbf{t}}^{\mathbf{MV_{Net}^{(1)}}} = \frac{1}{\lambda}\mathbf{V_t^{-1}}\hat{\mu}_{\mathbf{t}}^{\mathbf{e,(1)}}.$$

If $sign\left(\theta_{\mathbf{i,t}}^{\mathbf{MV_{Net}^{(1)}}}\right) \neq sign\left(\theta_{\mathbf{i,t}}^{\mathbf{MV_{Net}^{(0)}}}\right)$ for asset $i$, then we set the portfolio weight $\theta_{\mathbf{i,t}}^{\mathbf{MV_{Net}^{(2)}}} = 0$. Finally, we do one last mean-variance optimization using all assets $j$ for which $sign\left(\theta_{\mathbf{j,t}}^{\mathbf{MV_{Net}^{(1)}}}\right) = sign\left(\theta_{\mathbf{j,t}}^{\mathbf{MV_{Net}^{(0)}}}\right)$. That is,

$$\theta_{\{\mathbf{j}\},\mathbf{t}}^{\mathbf{MV_{Net}^{(2)}}} = \frac{1}{\lambda}\mathbf{V}_{\{\mathbf{j}\}\times\{\mathbf{j}\},\mathbf{t}}^{-\mathbf{1}}\hat{\mu}_{\{\mathbf{j}\},\mathbf{t}}^{\mathbf{e,(1)}},$$

where set $\{j\}$ includes all $j$ for which $sign\left(\theta_{\mathbf{j,t}}^{\mathbf{MV_{Net}^{(1)}}}\right) = sign\left(\theta_{\mathbf{j,t}}^{\mathbf{MV_{Net}^{(0)}}}\right)$, and vectors $\theta_{\{\mathbf{j}\},\mathbf{t}}^{\mathbf{MV_{Net}^{(2)}}}$ and $\hat{\mu}_{\{\mathbf{j}\},\mathbf{t}}^{\mathbf{e,(1)}}$ contain all elements of vectors $\theta_{\mathbf{t}}^{\mathbf{MV_{Net}^{(2)}}}$ and $\hat{\mu}_{\mathbf{t}}^{\mathbf{e,(1)}}$ with the locations specified by $\{j\}$, and matrix $\mathbf{V}_{\{\mathbf{j}\}\times\{\mathbf{j}\},\mathbf{t}}$ contains all elements of matrix $\mathbf{V_t}$ with the locations specified by $\{j\} \times \{j\}$ . We use $\theta_{\mathbf{t}}^{\mathbf{MV_{Net}^{(2)}}}$ as an approximation of $\theta_{\mathbf{t}}^{\mathbf{MV_{Net}}}$.

Intuitively, adjusting expected returns to incorporate trading costs reduces the attractiveness to buy or short sell assets, and thus, the notional value decreases. In turn, a smaller notional value is expected to imply less turnover and trading costs.

Tables 4, 5 and 6 compare the before and after cost performance of $MV_{Net}$ with the other aforementioned strategies. Moreover, tables 7 and 8 report the average annual turnover and transaction costs. We observe that $MV_{Net}$ has substantially lower costs and turnover than $MV$, achieving levels similar to those provided by $MV_{TC\setminus Corr}$. These findings are robust to (i) the set of 29 or 15 currencies and (ii) the trading frequency being weekly, monthly or quarterly.

Similarly to the other rules-of-thumb, the reduction in trading costs come at the expense of a suboptimal risk-return tradeoff. As shown in Table 4, 5 and 6 the before cost performance of $MV_{Net}$ is comparable to that of $MV_{TC}$ only at the quarterly frequency where costs do not matter much. $MV_{Net}$ underperform $MV_{TC}$ the monthly frequency by a Sharpe ratio gap of at least 0.12 (representing 11.5% of the $MV_{TC}$ performance). Moreover, it proofs not to be an adequate strategy for trading at higher frequencies, delivering a Sharpe ratio of no more than 0.26 when the asset universe encompasses the full set of 29 currencies.

Similarly, the after-cost performance of $MV_{Net}$ is comparable to that of $MV_{TC}$ only at the quarterly frequency. At the monthly frequency the Sharpe ratio of $MV_{Net}$ is 0.32 lower than that of $MV_{TC}$ which is statistically significant on the 1% level when trading involves the full set of 29 currencies. For the set of 15 developed currencies the difference is 0.14, which is not statistically significant but nevertheless it is economically meaningful since it represents approximately 15% of the performance of $MV_{TC}$. Finally, at the weekly frequency $MV_{Net}$ turns unprofitable, and thus, it is inadequate to tackle transaction costs at higher frequencies.

### 4.3 Price Impact

In the previous discussions we assume no price impact when an investor trades and trading costs are fully captured by current bid and ask quotes. As explained in sections 2.2 and 3.1 it is difficult to obtain sensible estimates for the price impact. In particular, there is no definite answer on whether larger trades have less price impact per dollar traded, leaving aside the issue of how to estimate it.

For this reason we employ the two mainstream functional forms for the price impact described in section 2.2 and a wide range of estimates. In particular, for each type of trade $z \in \{P+, P-, S+, S-\}$, the price impact $PI_t$ due to the trade $\mathbf{\Delta_t^z}$ is either linear and $PI_t = PI_t^L = \mathbf{\Delta_t^{z\prime} \Pi_t^{z,L}}$, or it grows at a rate of of 0.5 with the size of the trade and $PI_t = \sqrt{PI}_t = \sqrt{\mathbf{\Delta_t^{z}}'\mathbf{\Pi_t^{z,SR}}}$. The diagonal matrices $\mathbf{\Pi_t^{z,L}}$ and $\mathbf{\Pi_t^{z,SR}}$ contain the asset specific price impact parameters as in Novy-Marx and Velikov (2016) and Frazzini et al. (2015).

In order to get precise estimates of these parameters we need data on trade orders of investors in the FX markets. Since we do not have these data we perform a sensitivity analysis. We assume $\mathbf{\Pi_t^{z,L}} = \mathbf{\Pi_t^{z,SR}} = \pi \mathbf{I_t}$, where $I_t$ is the $N_t \times N_t$ identity matrix, and we vary $\pi$ between 0 and 100 basis points (bps), which includes a back-of-the-envelope average estimate from Mancini et al. (2013).[22] We therefore analyze the implications of a time and currency invariant, linear or square-root price impact on our trading strategies.

Mancini et al. (2013) provide linear price impact estimates for nine currency pairs over the period January 2007 to December 2009 based on data from the Electronic Broking Services, a leading global marketplace for spot inter-dealer FX trading.[23] In their framework $\pi_{ii,t} = \phi_t \times x_{i,t}$ where $x_{i,t}$ is the net trade in currency $i$ at time $t$ in millions of USD and $\phi_t$ is the sensitivity (in bps) at time $t$ of currency returns to $x_{i,t}$. As back-of-the-envelope estimates, we take the average of their estimates $\hat{\phi}_t$ and average $|x_{i,t}|$ for all the currency pairs involving the USD which are part of our traded universe. This yields an average price impact of $\bar{\pi} \approx 0.44 \times 41.42 \approx 18.22$ bps.

We therefore analyze the role of the time-currency invariant price impact $\pi$ on our strategies for many different values inside the basis point interval $[0, 100]$ for both a linear and a square-root functional form. There is no price impact when $\pi = 0$. This is a reasonable assumption for an investor with a relatively small portfolio size and small buy and sell orders. In contrast, if $\pi = 100$ bps, then an $x\%$ change in the portfolio allocation of asset $i$ changes the execution price and increases the per dollar cost to trade the asset by $x$ bps. A large $\pi$

[22]To our knowledge Mancini et al. (2013) is the only study that reports the price impact coefficients for FX rates.

[23]The nine currency pairs are the AUD/USD, EUR/CHF, EUR/GBP, EUR/JPY, EUR/USD, GBP/USD, USD/CAD, USD/CHF and USD/JPY.

is more relevant to an investor with a large portfolio size and large buy and sell orders.

The top graphs in Figure 10 show how annualized out-of-sample Sharpe ratios after transaction costs depend on the price impact parameter $\pi$ for various trading strategies constructed from our set of 29 developed and emerging (left panel) and 15 developed currencies (right panel) under the assumption of a linear functional form. We find that even a small value of $\pi$ has a considerable impact on the after cost Sharpe ratios. The solid black line displays the results for $MV_{TC}$. If there is no price impact then the after cost Sharpe ratio is 1.16 (0.96) for our set of 29 developed and emerging (15 developed) currencies. If $\pi$ is 1, 5 or 10 bps, then the Sharpe ratio drops to 1.11 (0.93), 1.03 (0.88) or 0.97 (0.86). While this decrease in the performance of $MV_{TC}$ due to the introduction of a price impact is considerable, we find the decrease in the after cost Sharpe ratio becomes less sensitive to changes in $\pi$ when $\pi$ is larger. For instance, the after cost Sharpe ratio experiences a proportionally moderate decrease from 0.89 (0.81) to 0.79 (0.74) or 0.72 (0.67) when we increase $\pi$ from 20 to 50 or 100 bps. Thus, even if we assume an investor has a large portfolio size with large buy and sell orders or FX markets are illiquid and the price impact is large, $MV_{TC}$ remains a profitable trading strategy.

The introduction of a price impact has a significant negative effect on the performance of $MV$ (bold black dashed line). The after cost Sharpe ratio sharply decreases, and if $\pi \geq 2$ bps, $MV$ turns unprofitable, i.e., the average returns after costs are negative. This is a stark result. Ignoring a price impact can have disastrous consequences for an investor. Hence, if $\pi \geq 2$ bps, taking transaction costs into account in the optimization does not only improves the out-of-sample performance but it is necessary to earn a positive average return after costs.

$MV_{TC\backslash Corr}$ is more robust to a price impact than $MV$. Its after cost Sharpe ratio remains positive and above 0.53 (0.34) for any level of $\pi$ for our set of 29 developed and emerging (15 developed) currencies.

The bottom plots of Figure 10 show the impact of $\pi$ on the trade aggressiveness as defined in equation (1) of $MV$ (black dashed line), $MV_{TC\backslash Corr}$ (black dashed-dotted line) and $MV_{TC}$

(black solid line). The left (right) figure shows the results for our set of 29 (15) currencies. By definition the trade aggressiveness of $MV$ is constant and equal to 1. We observe a rapid decline in the trade aggressiveness of $MV_{TC}$ and $MV_{TC\setminus Corr}$ as $\pi$ increases, and the decrease is faster for $MV_{TC}$. If there is no price impact the trade aggressiveness of $MV_{TC}$ is 0.39 (0.32), and it is only 0.29 (0.23), 0.19 (0.15), 0.12 (0.09) or even 0.03 (0.02) if $\pi$ is 1, 5, 10 or 100 bps for our set of 29 (15) currencies. The corresponding trade aggressiveness of $MV_{TC\setminus Corr}$ is 0.72 (0.75), 0.64 (0.67), 0.48 (0.51), 0.37 (0.40) or 0.08 (0.08). While the price impact causes transaction costs to surge and turns $MV$ unprofitable, $MV_{TC}$ and $MV_{TC\setminus Corr}$ decrease their trading activity to save costs and remain profitable. This result shed light on the empirical behavior of the optimal strategy in the presence of price impact, which theory is silent on. The fact that the trade aggressiveness is inversely related to the magnitude of the price impact parameter implies that on average in the presence of a price impact it is optimal to trade less.

Finally, the top panel in figure 10 also reports how $\pi$ affects the after cost Sharpe ratios of popular equally weighted currency trading strategies. In the case of $MOM$ (purple solid line) even a small price impact $\pi$ of 2 bps turns the strategy unprofitable, quickly making it the worst among all the discussed strategies. In the case of $HML$ (yellow solid line) and $VAL$ (green solid line) transaction costs exceed the returns and the after cost Sharpe ratio is negative if $\pi$ is in the neighborhood of 30 bps or larger. Thus, the optimized strategies $MV_{TC}$ and $MV_{TC\setminus Corr}$ are less sensitive to costs arising from a price impact than $HML$, $MOM$ or $VAL$. Although our setting assumes that the price impact is constant through time, we expect that $MV_{TC}$ and $MV_{TC\setminus Corr}$ are less susceptible to periods of high market illiquidity than some of the popular currency strategies.

The performance of $DOL$ (blue solid line) is almost unaffected by $\pi$, though the after cost Sharpe ratio of $DOL$ is small even before costs, i.e., 0.09. The after cost Sharpe ratio of $DDOL$ (red solid line) steadily decreases as $\pi$ increases. It is 0.43 (0.56) if there is no price impact and almost linearly decreases to 0.20 (0.23) if $\pi$ is 100 bps. $DDOL$ is less sensitive to changes in $\pi$ than $MV_{TC}$ and $MV_{TC\setminus Corr}$ if $\pi$ is small but the opposite is true for larger values of $\pi$. Moreover, for any value of $\pi$ the after cost Sharpe ratio is always substantially larger for $MV_{TC}$ than for $DDOL$.

As reported in Figure 11 the inference remains the same if we adopt a square-root functional form for the price impact. If anything, the relative out-performance of $MV_{TC}$ and $MV_{TC\backslash Corr}$ become even more evident. As a matter of fact, for values of $\pi$ higher than 50, while all but the low performing DOL strategy turn unprofitable, the Sharpe ratios of $MV_{TC}$ and $MV_{TC\backslash Corr}$ stabilize at around 0.42 (0.55) and 0.40 (0.30) for our set of 29 (15) currencies.

To sum up, independent of the functional form of the price impact or the set of currencies, minimizing over costs in the optimization is crucial if the price impact is not infinitesimal. $MV_{TC}$ and $MV_{TC\backslash Corr}$ are the *only* strategies that perform well after costs even if an investor has a large portfolio size and executes large USD amounts of buy and sell orders or FX markets are illiquid and the price impact of trading is severe. In contrast, transaction costs due to a price impact quickly erode the returns of $MV$ and many popular equally weighted currency trading strategies. These strategies significantly underperform $MV_{TC}$ even if the price impact is moderate and some strategies turn unprofitable. Moreover, while the theory is silent on the size and direction of the optimal trades in the presence of price impact (because they depend on the price impact parameters), we empirically demonstrate how in our setup, it is optimal to trade less.

# 5 Conclusion

Using monthly FX market returns of 29 developed and emerging currencies from 1976 to 2016, we show that taking transaction costs into account in a mean-variance portfolio optimization leads to an economically large and statistically significant improvement in the out-of-sample performance. For the set 29 (15) currencies the out-of-sample Sharpe ratio after costs increases from 0.78 (0.75) for $MV$, which ignores costs in the optimization, to 1.16 (0.96) for $MV_{TC}$, which optimizes costs. The outperformance after transaction costs of $MV_{TC}$ over $MV$ is due to a significant reduction in the trading activity which reduces trading costs. Moreover, we show that it is important to properly account for correlations between assets to achieve the superior performance of $MV_{TC}$. This is an important contribution to the transaction costs literature because it demonstrates that we should not make the assumption of Liu (2004) that assets are uncorrelated. This is unfortunate because the solution of the optimal portfolio in a dynamic setting with many assets and direct transaction costs is only known if the assets are uncorrelated.

We further investigate whether the issue of transaction costs can be mitigated if we (i) construct equally weighted long-short strategies instead of an optimized portfolio, (ii) trade at a low frequency, (iii) remove assets with high transaction costs from the set of admissible assets, (iv) only rebalance if the current position is too far from the desired position, or (v) use expected returns net of costs in the optimization. While we confirm that these intuitive rules-of-thumb reduce transaction costs, we find that they are inefficient because the performance significantly worsens if we deviate from the optimal tradeoff between expected returns and risk, decrease the trading frequency, or restrict the asset universe. Overall, our findings advise against the use of an intuitive rule-of-thumb to mitigate transaction costs. The out-of-sample performance is significantly better if we invest in a fully optimized portfolio which accounts for costs in the optimization.

Finally, we analyze how a price impact of trading affects the performance of our strategies. We show that $MV_{TC}$ performs well after costs even if an investor has a large portfolio size and executes buy and sell orders with a large USD amount or FX markets are illiquid and the price impact of trading is severe. In contrast, $MV$ and many popular equally weighted currency trading strategies in the literature significantly underperform $MV_{TC}$. Moreover, trading costs quickly erode returns and turn many of these strategies unprofitable if the price impact is large enough. Finally, while the theory is silent on the size and direction of the optimal trades in the presence of a price impact (because they depend on the price impact parameters), we empirically demonstrate how in our setup, it is optimal to trade less in the presence of proportional costs and price impact that what it would be optimal if costs were only proportional.

## References


Abel, Andrew B., Janice C. Eberly, and Stavros Panageas, 2013, Optimal Inattention To The Stock Market With Information Costs and Transaction Costs, *Econometrica* 81, 14551481.

Ackermann, Fabian, Walt Pohl, and Karl Schmedders, 2016, Optimal and Naive Diversification in Currency Markets, *Management Science* .

Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross section of volatility and expected returns, *Journal of Finance* 51, 259–299.

Balduzzi, Pierluigi, and Anthony W. Lynch, 1999, Transaction costs and predictability: some utility cost calculations, *Journal of Financial Economics* 52, 47–78.

Balduzzi, Pierluigi, and Anthony W. Lynch, 2000, Predictability and Transaction Costs: The Impact on Rebalancing Rules and Behavior, *Journal of Finance* 55, 2285–2309.

Baz, Jamil, Frances Breedon, Vasant Naik, and Joel Peress, 2001, Optimal Portfolios of Foreign Currencies, *The Journal of Portfolio Management* 28, 102–111.

Brandt, M. W., P. Santa-Clara, , and R. Valkanov, 2009, Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns, *Review of Financial Studies* 22, 3411–3447.

Brandt, Michael W., 2005, Portfolio choice problems, in Y. Ait-Shahalia, and L. P. Hansen, eds., *Handbook of Financial Econometrics*, chapter 5, 269–336 (Elsevier, Amsterdam).

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2011, Do Peso Problems Explain the Returns to the Carry Trade?, *Review of Financial Studies* 24, 853–891.

Burnside, Craig, Sergio Rebelo, and Martin Eichenbaum, 2008, Carry Trade: The Gains of Diversification, *Journal of the European Economic Association* 6, 581–588.

Buss, Adrian, and Bernard Dumas, 2017, The Dynamic Properties of Financial-Market Equilibrium with Trading Fees, Working paper, INSEAD.

Campanale, Claudio, Carolina Fugazza, and Francisco Gomez, 2015, Life-cycle portfolio choice with liquid and illiquid financial assets, *Journal of Monetary Economics* 71, 67–83.

Chen, Andrew Y., and Mihail Velikov, 2019, Accounting for the Anomaly Zoo: a Trading Cost Perspective, Working paper.

Chinn, Menzie D., and Hiro Ito, 2006, What Matters for Financial Development? Capital Controls, Institutions, and Interactions, *Journal of Development Economics* 81, 163–192.

Daniel, Kent, Robert J. Hodrick, and Zhongjin Lu, 2017, The Carry Trade: Risks and Drawdowns, *Critical Finance Review* 6, 211–262.

Davis, Mark H.A., and A. R. Norman, 1990, Portfolio Selection With Transaction Costs, *Mathematics of Operations Research* 15, 676–713.

Della Corte, Pasquale, Tarun Ramadorai, and Lucio Sarno, 2016, Volatility Risk Premia and Exchange Rate Predictability, *Journal of Financial Economics* 120, 21–40.

Della Corte, Pasquale, Lucio Sarno, and Ilias Tsiakas, 2009, An Economic Evaluation of Empirical Exchange Rate Models, *Review of Financial Studies* 22, 3491–3530.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal, 2009, Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?, *Review of Financial Studies* 22, 1915–1953.

DeMiguel, Victor, Alberto Martin-Utrera, Francisco J. Nogales, and Raman Uppal, 2019, A Transaction-Cost Perspective on the Multitude of Firm Characteristics, *Review of Financial Studies* .

Donohue, C., and K. Yip, 2003, Optimal Portfolio Rebalancing with Transaction Costs: Improving on Calendar- or Volatility-Based Strategies, *Journal of Portfolio Management* 29, 49–63.

Dumas, Bernard, and Elisa Luciano, 1991, An Exact Solution to a Dynamic Portfolio Choice Problem under Transactions Costs, *Journal of Finance* 46, 577–595.

Dybvig, Phillip H., and Luca Pezzo, 2019, Mean-Variance Porfolio Rebalancing with Transaction Costs, Working paper, Washington University.

Frazzini, Andrea, Ronen Israel, and Tobias J. Moskowitz, 2015, Trading Costs of Asset Pricing Anomalies, Working paper, AQR Capital Management, LLC and Yale University.

Garleanu, Nicolae, and Lasse H. Pedersen, 2013, Dynamic Trading with Predictable Returns and Transaction Costs, *Journal of Finance* .

Garleanu, Nicolae, and Lasse H. Pedersen, 2016, Dynamic portfolio choice with frictions, *Journal of Economic Theory* .

Goodman, Jonathan, and Daniel N. Ostrov, 2010, Balancing Small Transaction Costs with Loss of Optimal Allocation in Dynamic Stock Trading Strategies, *SIAM Journal on Applied Mathematics* 70, 19771998.

Irle, A., and C. Prelle, 2008, A renewal theoretic result in porfolio theory under transaction costs with multiple risky assets, Working paper, Kiel Institute for the World Economy.

Karnaukh, Nina, Angelo Ranaldo, and Paul Soederlind, 2015, Understanding FX Liquidity, *The Review of Financial Studies* 28, 3073–3108.

Korajczyk, Robert A., and Ronnie Sadka, 2004, Are Momentum Profits Robust to Trading Costs?, *Journal of Finance* 59, 1039–1082.

Ledoit, Olivier, and Michael Wolf, 2008, Robust Performance Hypothesis Testing with the Sharpe Ratio, *Journal of Empirical Finance* 15, 850–859.

Leland, Hayne E., 2000, Optimal porfolio implemenentation with transaction costs and capital gain taxes, Working paper, University of California at Berkeley.

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

Liu, Hong, 2004, Optimal Consumption and Investment with Transaction Costs and Multiple Risky Assets, *Journal of Finance* 59, 289–338.

Liu, Hong, and M. Loewenstein, 2004, Optimal Porfolio Selection with Transaction Costs and Finite Horizons, *Review of Financial Studies* 15, 805–835.

Liu, Hong A., and Jing Xu, 2018, Optimal Liquidation and Acquisition Strategy with Price Impact, Working paper, Washington University.

Lustig, Hanno, Nick Roussanov, and Adrien Verdelhan, 2011, Common Risk Factors in Currency Returns, *Review of Financial Studies* 24, 3731–3777.

Lynch, Anthony W., and Sinan Tan, 2009, Multiple Risky Assets, Transaction Costs, and Return Predictability: Allocation Rules and Implications for the U.S. Investors, *Journal of Financial and Quantitative Analysis* 45, 1015–1053.

Mancini, Loriano, Angelo Ranaldo, and Jan Wrampelmeyer, 2013, Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums, *Journal of Finance* 68, 1805–1841.

Markowitz, Harry, 1952, Portfolio Selection, *Journal of Finance* 7, 77–91.

Maurer, Thomas A., Thuy-Duong To, and Ngoc-Khanh Tran, 2018a, Optimal Factor Strategy in FX Markets, Working paper, Washington University.

Maurer, Thomas A., Thuy-Duong To, and Ngoc-Khanh Tran, 2018b, Pricing Implications of Covariances and Spreads in Currency Markets, Working paper, Washington University.

Meese, R., and K. Rogoff, 1983, Empirical Exchange Rate Models of the Seventies: Do They Fir Out of Sample?, *Journal of International Economics* 14, 3–24.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012, Currency Momentum Strategies, *Journal of Financial Economics* 106, 620–684.

Merton, Robert, 1971, Optimum Consumption and Portfolio Rules in a Continuous-Time Model, *Journal of Economic Theory* 3, 373–413.

Merton, Robert C., 1987, A simple model of capital market equilibrium with incomplete information, *Journal of Finance* 42, 483–510.

Muthuraman, K., and S. Kumar, 2006, Multidimensional Porfolio Optimization with Proportional Transaction Costs, *Mathematical Finance* 16, 301–335.

Myers, Jeremy D., 2009, Portfolio Optimization with Transaction Costs and Preconceived Portfolio Weights, Ph.d dissertation, Massachusetts Institute of Technology.

Newey, Whitney K., and Kenneth D. West, 1987, A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703708.

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

Schachermayer, Walter, 2017, Asymptotic Theory of Transaction Costs , in Jessica James, Ian W. Marsh, and Lucio Sarno, eds., *Asymptotic Theory of Transaction Costs* (European Mathematical Society, Hoboken, NJ).

Shreve, S. E., and H. M. Soner, 1994, Optimal investment and consumption with transaction costs, *Annals of Applied Probability* 4, 609–692.

Taksar, M., M. J. Klass, and D. Assaf, 1988, A Diffusion Model for Optimal Portfolio Selection in the Presence of Brokerage Fees, *Mathematics of Operations Research* 13, 277–294.

Table 1: **Mean-Variance Strategies:** $MV_{TC}$ **vs** $MV$

| | All 29 Currencies | | 15 Developed Currencies | |
|---|---|---|---|---|
| | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| **Before Transaction Costs:** | | | | |
| SR | 1.26 | 1.27 | 0.96 | 1.04 |
| Mean | 13.30 | 11.89 | 9.18 | 8.50 |
| ΔMean | - | -1.41 | - | -0.69 |
| **Transaction Costs:** | | | | |
| Mean Costs | 5.13 | 1.25 | 2.19 | 0.85 |
| ΔMean Costs | - | -3.88$^{***}$ | - | -1.34$^{***}$ |
| **After Transaction Costs:** | | | | |
| **SR** | **0.78** | **1.16** | **0.75** | **0.96** |
| **ΔSR** | **-** | **0.38**$^{***}$ | **-** | **0.20**$^{**}$ |
| **(p-value)** | - | **(0.01)** | **-** | **(0.03)** |
| Mean | 8.17 | 10.64 | 6.99 | 7.64 |
| Vol | 10.52 | 9.21 | 9.27 | 7.98 |
| Skew | 0.00 | 0.47 | -1.00 | -0.28 |
| Kurt | 7.60 | 6.59 | 15.31 | 12.94 |
| % Positive | 61.49 | 68.30 | 62.77 | 65.96 |
| MDD | -37.61 | -18.90 | -35.97 | -24.16 |
| AC | 0.11 | 0.15 | 0.10 | 0.17 |
| $CE_{\lambda=1}$ | - | 0.81 | - | 0.35 |
| $CE_{\lambda=5}$ | - | 3.04 | - | 1.49 |
| $CE_{\lambda=10}$ | - | 5.82 | - | 2.91 |
| $CE_{\lambda=50}$ | - | 28.09 | - | 14.27 |
| $TA$ | 1 | 0.41 | 1 | 0.35 |

*Notes*: Summary statistics of monthly excess returns of $MV$ and $MV_{TC}$. First two columns report results for all 29 currencies, last two columns for 15 developed currencies. The sample period is 1976-2016. SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation, $CE_\lambda$ the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV$ to strategy $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. $^{***}$, $^{**}$, $^{*}$ indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 2: $MV_{TC}$ vs $MV$: **Importance of Correlations, 29 Currencies**

| | $MV$ | $MV_{TC\setminus Corr}$ | $MV_{TC}$ |
|---|---|---|---|
| **Before Transaction Costs:** | | | |
| SR | 1.26 | 1.22 | 1.27 |
| Mean | 13.30 | 12.35 | 11.89 |
| ΔMean | 0.95 | - | -0.46 |
| **Transaction Costs:** | | | |
| Mean Costs | 5.13 | 2.34 | 1.25 |
| ΔMean Costs | 2.79*** | - | -1.10*** |
| **After Transaction Costs:** | | | |
| **SR** | **0.78** | **1.01** | **1.16** |
| **ΔSR** | **-0.23**** | **-** | **0.15**** |
| **(p-value)** | **(0.03)** | **-** | **(0.04)** |
| Mean | 8.17 | 10.01 | 10.64 |
| Vol | 10.52 | 9.93 | 9.21 |
| Skew | 0.00 | 0.79 | 0.47 |
| Kurt | 7.60 | 8.18 | 6.59 |
| % Positive | 61.49 | 65.11 | 68.30 |
| MDD | -37.61 | -18.92 | -18.90 |
| AC | 0.11 | 0.10 | 0.15 |
| $CE_{\lambda=1}$ | -0.56 | - | 0.25 |
| $CE_{\lambda=5}$ | -2.03 | - | 1.01 |
| $CE_{\lambda=10}$ | -3.87 | - | 1.95 |
| $CE_{\lambda=50}$ | -18.57 | - | 9.51 |
| $TA$ | 1 | 0.70 | 0.41 |
| ΔTA | - | - | -0.29*** |

*Notes*: Summary statistics of monthly excess returns of $MV$, $MV_{TC}$ and $MV_{TC\setminus Corr}$. for all 29 currencies from 1976 to 2016. SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation. $CE_\lambda$ is the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV_{TC\setminus Corr}$ to strategy $MV$ or $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. ΔTA is the difference in $TA$ between $MV_{TC\setminus Corr}$ and $MV_{TC}$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs and ΔTA are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 3: $MV_{TC}$ **vs** $MV$**: Importance of Correlations, 15 Developed Currencies**

| | $MV$ | $MV_{TC\backslash Corr}$ | $MV_{TC}$ |
|---|---|---|---|
| **Before Transaction Costs:** | | | |
| SR | 0.96 | 0.94 | 1.04 |
| Mean | 9.18 | 8.62 | 8.50 |
| ΔMean | 0.57 | - | -0.12 |
| **Transaction Costs:** | | | |
| Mean Costs | 2.19 | 1.63 | 0.85 |
| ΔMean Costs | 0.56*** | - | -0.78*** |
| **After Transaction Costs:** | | | |
| **SR** | **0.75** | **0.79** | **0.96** |
| **ΔSR** | **-0.04** | **-** | **0.17**** |
| **(p-value)** | **(0.13)** | **-** | **(0.04)** |
| Mean | 6.99 | 6.99 | 7.64 |
| Vol | 9.27 | 8.88 | 7.98 |
| Skew | -1.00 | -0.88 | -0.28 |
| Kurt | 15.31 | 14.73 | 12.94 |
| % Positive | 62.77 | 63.19 | 65.96 |
| MDD | -35.97 | -32.29 | -24.16 |
| AC | 0.10 | 0.11 | 0.17 |
| $CE_{\lambda=1}$ | -0.07 | - | 0.28 |
| $CE_{\lambda=5}$ | -0.37 | - | 1.11 |
| $CE_{\lambda=10}$ | -0.75 | - | 2.16 |
| $CE_{\lambda=50}$ | -3.74 | - | 10.53 |
| $TA$ | 1 | 0.72 | 0.35 |
| ΔTA | - | - | -0.37*** |

*Notes*: Summary statistics of monthly excess returns of $MV$, $MV_{TC}$ and $MV_{TC\backslash Corr}$. for 15 developed currencies from 1976 to 2016. SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation. $CE_\lambda$ is the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV_{TC\backslash Corr}$ to strategy $MV$ or $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. ΔTA is the difference in $TA$ between $MV_{TC\backslash Corr}$ and $MV_{TC}$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs and ΔTA are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 4: **Optimized vs Equally Weighted Portfolios: Monthly Frequency**

| | All 29 Currencies | | | 15 Developed Currencies | | |
|---|---|---|---|---|---|---|
| | Before TC | After TC | | Before TC | After TC | |
| Strategies | SR | SR | ΔSR | SR | SR | ΔSR |
| $\mathbf{MV_{TC}}$ | **1.27** | **1.16** | - | **1.04** | **0.96** | - |
| $MV_{TC\backslash Corr}$ | 1.22 | 1.01 | $0.15^{**}$ | 0.94 | 0.79 | $0.17^{**}$ |
| | | | (0.03) | | | (0.05) |
| $MV$ | 1.26 | 0.78 | $0.38^{**}$ | 0.96 | 0.75 | $0.20^{**}$ |
| | | | (0.01) | | | (0.02) |
| $MV_{Net}$ | 1.01 | 0.84 | $0.32^{***}$ | 0.92 | 0.82 | 0.14 |
| | | | (0.01) | | | (0.24) |
| $DOL$ | 0.10 | 0.09 | $1.07^{***}$ | 0.07 | 0.07 | $0.89^{***}$ |
| | | | (0.00) | | | (0.00) |
| $DDOL$ | 0.49 | 0.43 | $0.72^{***}$ | 0.59 | 0.56 | $0.40^{*}$ |
| | | | (0.00) | | | (0.07) |
| $HML$ | 0.74 | 0.60 | $0.55^{***}$ | 0.64 | 0.59 | $0.37^{*}$ |
| | | | (0.00) | | | (0.10) |
| $MOM$ | 0.37 | 0.28 | $0.87^{***}$ | 0.31 | 0.24 | $0.72^{***}$ |
| | | | (0.00) | | | (0.01) |
| $VAL$ | 0.47 | 0.38 | $0.78^{***}$ | 0.53 | 0.50 | $0.46^{*}$ |
| | | | (0.00) | | | (0.06) |

*Notes*: Columns 2 and 5 report Sharpe ratios before costs (Before TC SR). Columns 3 and 6 report Sharpe ratios after costs (After TC SR). Columns 4 and 7 report the difference between the Sharpe ratios after costs of $MV_{TC}$ and the strategy in the corresponding row (After TC ΔSR). $MV_{TC}$ is the mean-variance optimized portfolio which optimizes over transaction costs. $MV_{TC\backslash Corr}$ is the mean-variance optimized portfolio which optimizes over transaction costs but makes the simplifying assumption that assets are uncorrelated. $MV$ is the mean-variance optimized portfolio without taking into account transaction costs in the optimization. $MV_{Net}$ is analogous to $MV$ only it is based on returns net of costs. $DOL$ borrows in the USD and equally invests in all other currencies. $DDOL$ takes a long position in $DOL$ if the median exchange rate forward discount is positive, and a short position otherwise. $HML$ sorts currencies according to the forward discount into quintiles and borrows in the bottom and invests in the top quintile. $MOM$ sorts currencies according to their past 12 month performance into quintiles and borrows in the bottom and invests in the top quintile. $VAL$ sorts currencies according to the power purchase parity adjusted exchange rate into quintiles and borrows in the top quintile and invests in the bottom quintile. The data are monthly returns for our full set of 29 currencies (columns 2 to 4) and a subsample of 15 developed currencies (columns 5 to 7) from January 1976 to February 2016. Standard errors for ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008).

Table 5: **Optimized vs Equally Weighted Portfolios: Weekly Frequency**

| | All 29 Currencies | | | 15 Developed Currencies | | |
|---|---|---|---|---|---|---|
| | Before TC | After TC | | Before TC | After TC | |
| Strategies | SR | SR | ΔSR | SR | SR | ΔSR |
| $\mathbf{MV_{TC}}$ | **0.89** | **0.84** | - | **0.98** | **0.93** | - |
| $MV_{TC\setminus Corr}$ | 0.97 | 0.76 | 0.08 | 0.97 | 0.77 | $0.17^{*}$ |
| | | | (0.47) | | | (0.10) |
| $MV$ | 1.07 | -0.06 | $0.90^{***}$ | 1.15 | 0.54 | $0.40^{***}$ |
| | | | (0.00) | | | (0.00) |
| $MV_{Net}$ | 0.26 | -0.19 | $1.03^{***}$ | 0.11 | -0.24 | $1.17^{***}$ |
| | | | (0.00) | | | (0.00) |
| $DOL$ | 0.09 | 0.08 | $0.76^{***}$ | 0.08 | 0.07 | $0.86^{***}$ |
| | | | (0.00) | | | (0.00) |
| $DDOL$ | 0.46 | 0.40 | $0.43^{**}$ | 0.61 | 0.58 | $0.36^{*}$ |
| | | | (0.04) | | | (0.09) |
| $HML$ | 0.70 | 0.57 | 0.27 | 0.64 | 0.59 | $0.34^{**}$ |
| | | | (0.11) | | | (0.06) |
| $MOM$ | 0.36 | 0.26 | $0.58^{***}$ | 0.25 | 0.17 | $0.77^{***}$ |
| | | | (0.01) | | | (0.00) |
| $VAL$ | 0.41 | 0.31 | $0.53^{**}$ | 0.53 | 0.47 | $0.47^{**}$ |
| | | | (0.01) | | | (0.04) |

*Notes*: Columns 2 and 5 report Sharpe ratios before costs (Before TC SR). Columns 3 and 6 report Sharpe ratios after costs (After TC SR). Columns 4 and 7 report the difference between the Sharpe ratios after costs of $MV_{TC}$ and the strategy in the corresponding row (After TC ΔSR). $MV_{TC}$ is the mean-variance optimized portfolio which optimizes over transaction costs. $MV_{TC\setminus Corr}$ is the mean-variance optimized portfolio which optimizes over transaction costs but makes the simplifying assumption that assets are uncorrelated. $MV$ is the mean-variance optimized portfolio without taking into account transaction costs in the optimization. $MV_{Net}$ is analogous to $MV$ only it is based on returns net of costs. $DOL$ borrows in the USD and equally invests in all other currencies. $DDOL$ takes a long position in $DOL$ if the median exchange rate forward discount is positive, and a short position otherwise. $HML$ sorts currencies according to the forward discount into quintiles and borrows in the bottom and invests in the top quintile. $MOM$ sorts currencies according to their past 12 month performance into quintiles and borrows in the bottom and invests in the top quintile. $VAL$ sorts currencies according to the power purchase parity adjusted exchange rate into quintiles and borrows in the top quintile and invests in the bottom quintile. The data are weekly returns for our full set of 29 currencies (columns 2 to 4) and a subsample of 15 developed currencies (columns 5 to 7) from January 1976 to February 2016. Standard errors for ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008).

Table 6: **Optimized vs Equally Weighted Portfolios: Quarterly Frequency**

| | All 29 Currencies | | | 15 Developed Currencies | | |
|---|---|---|---|---|---|---|
| | Before TC | After TC | | Before TC | After TC | |
| Strategies | SR | SR | ΔSR | SR | SR | ΔSR |
| $\mathbf{MV_{TC}}$ | **0.91** | **0.83** | - | **0.89** | **0.83** | - |
| $MV_{TC\backslash Corr}$ | 0.87 | 0.73 | 0.09<br>(0.24) | 0.87 | 0.78 | 0.05**<br>(0.03) |
| $MV$ | 0.91 | 0.73 | 0.10<br>(0.15) | 0.87 | 0.78 | 0.05**<br>(0.03) |
| $MV_{Net}$ | 0.86 | 0.76 | 0.06<br>(0.43) | 0.91 | 0.85 | -0.02<br>(0.74) |
| $DOL$ | 0.08 | 0.08 | 0.75***<br>(0.01) | 0.06 | 0.06 | 0.77***<br>(0.01) |
| $DDOL$ | 0.41 | 0.39 | 0.44**<br>(0.05) | 0.64 | 0.63 | 0.20<br>(0.33) |
| $HML$ | 0.58 | 0.53 | 0.30**<br>(0.03) | 0.43 | 0.41 | 0.42**<br>(0.03) |
| $MOM$ | 0.29 | 0.24 | 0.59**<br>(0.02) | 0.17 | 0.13 | 0.70**<br>(0.01) |
| $VAL$ | 0.45 | 0.41 | 0.42**<br>(0.03) | 0.55 | 0.52 | 0.31<br>(0.13) |

*Notes*: Columns 2 and 5 report Sharpe ratios before costs (Before TC SR). Columns 3 and 6 report Sharpe ratios after costs (After TC SR). Columns 4 and 7 report the difference between the Sharpe ratios after costs of $MV_{TC}$ and the strategy in the corresponding row (After TC ΔSR). $MV_{TC}$ is the mean-variance optimized portfolio which optimizes over transaction costs. $MV_{TC\backslash Corr}$ is the mean-variance optimized portfolio which optimizes over transaction costs but makes the simplifying assumption that assets are uncorrelated. $MV$ is the mean-variance optimized portfolio without taking into account transaction costs in the optimization. $MV_{Net}$ is analogous to $MV$ only it is based on returns net of costs. $DOL$ borrows in the USD and equally invests in all other currencies. $DDOL$ takes a long position in $DOL$ if the median exchange rate forward discount is positive, and a short position otherwise. $HML$ sorts currencies according to the forward discount into quintiles and borrows in the bottom and invests in the top quintile. $MOM$ sorts currencies according to their past 12 month performance into quintiles and borrows in the bottom and invests in the top quintile. $VAL$ sorts currencies according to the power purchase parity adjusted exchange rate into quintiles and borrows in the top quintile and invests in the bottom quintile. The data are quarterly returns for our full set of 29 currencies (columns 2 to 4) and a subsample of 15 developed currencies (columns 5 to 7) from January 1976 to February 2016. Standard errors for ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008).

Table 7: **Average Turnover and Transaction Costs, 29 Currencies**

| | Weekly | | Monthly | | Quarterly | |
|---|---|---|---|---|---|---|
| Strategies | Turnover | Costs | Turnover | Costs | Turnover | Costs |
| $MV_{TC}$ | 9.33 | 0.36 | 18.49 | 1.25 | 12.64 | 1.16 |
| $MV_{TC\backslash Corr}$ | 31.76 | 1.94 | 28.74 | 2.34 | 15.45 | 1.83 |
| $MV$ | 101.15 | 11.18 | 40.47 | 5.13 | 17.77 | 2.46 |
| $MV_{Net}$ | 39.59 | 3.12 | 23.11 | 1.73 | 12.47 | 1.36 |
| $DOL$ | 0.82 | 0.12 | 0.50 | 0.10 | 0.27 | 0.05 |
| $DDOL$ | 3.78 | 0.45 | 3.41 | 0.44 | 1.39 | 0.19 |
| $HML$ | 9.39 | 1.01 | 8.64 | 0.98 | 3.79 | 0.45 |
| $MOM$ | 20.08 | 1.45 | 14.81 | 1.14 | 7.83 | 0.72 |
| $VAL$ | 6.00 | 0.71 | 3.46 | 0.55 | 1.71 | 0.25 |

*Notes*: Average annual turnover and transaction costs for diverse strategies. $MV_{TC}$ is the mean-variance optimized portfolio which optimizes over transaction costs. $MV_{TC\backslash Corr}$ is the mean-variance optimized portfolio which optimizes over transaction costs but makes the simplifying assumption that assets are uncorrelated. $MV$ is the mean-variance optimized portfolio without taking into account transaction costs in the optimization. $MV_{Net}$ is analogous to $MV$ only it is based on returns net of costs. $DOL$ borrows in the USD and equally invests in all other currencies. $DDOL$ takes a long position in $DOL$ if the median exchange rate forward discount is positive, and a short position otherwise. $HML$ sorts currencies according to the forward discount into quintiles and borrows in the bottom and invests in the top quintile. $MOM$ sorts currencies according to their past 12 month performance into quintiles and borrows in the bottom and invests in the top quintile. $VAL$ sorts currencies according to the power purchase parity adjusted exchange rate into quintiles and borrows in the top quintile and invests in the bottom quintile. The sample contains all 29 currencies from January 1976 to February 2016.

Table 8: **Average Turnover and Transaction Costs, 15 Developed Currencies**

| | Weekly | | Monthly | | Quarterly | |
|---|---|---|---|---|---|---|
| Strategies | Turnover | Costs | Turnover | Costs | Turnover | Costs |
| $MV_{TC}$ | 5.30 | 0.25 | 12.51 | 0.85 | 8.60 | 0.72 |
| $MV_{TC\setminus Corr}$ | 24.12 | 1.60 | 21.46 | 1.63 | 10.82 | 0.95 |
| $MV$ | 63.73 | 5.23 | 27.53 | 2.19 | 11.75 | 1.04 |
| $MV_{Net}$ | 19.28 | 1.80 | 18.25 | 1.31 | 9.18 | 0.84 |
| $DOL$ | 0.68 | 0.05 | 0.40 | 0.03 | 0.24 | 0.02 |
| $DDOL$ | 3.62 | 0.27 | 3.30 | 0.25 | 1.44 | 0.12 |
| $HML$ | 7.16 | 0.47 | 6.45 | 0.45 | 3.38 | 0.25 |
| $MOM$ | 18.72 | 1.09 | 15.80 | 1.04 | 7.77 | 0.52 |
| $VAL$ | 7.89 | 0.50 | 4.10 | 0.28 | 2.42 | 0.19 |

*Notes*: Average annual turnover and transaction costs for diverse strategies. $MV_{TC}$ is the mean-variance optimized portfolio which optimizes over transaction costs. $MV_{TC\setminus Corr}$ is the mean-variance optimized portfolio which optimizes over transaction costs but makes the simplifying assumption that assets are uncorrelated. $MV$ is the mean-variance optimized portfolio without taking into account transaction costs in the optimization. $MV_{Net}$ is analogous to $MV$ only it is based on returns net of costs. $DOL$ borrows in the USD and equally invests in all other currencies. $DDOL$ takes a long position in $DOL$ if the median exchange rate forward discount is positive, and a short position otherwise. $HML$ sorts currencies according to the forward discount into quintiles and borrows in the bottom and invests in the top quintile. $MOM$ sorts currencies according to their past 12 month performance into quintiles and borrows in the bottom and invests in the top quintile. $VAL$ sorts currencies according to the power purchase parity adjusted exchange rate into quintiles and borrows in the top quintile and invests in the bottom quintile. The sample contains 15 developed currencies from January 1976 to February 2016.

Table 9: **Restricting 29 Admissible Currencies to Low Transaction Cost ones: Sharpe Ratio**

| Drop # Top | Before TC Sharpe Ratios | | After TC Sharpe Ratios | | | |
|---|---|---|---|---|---|---|
| TC Currencies | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $\Delta$SR | (p-val) |
| Keep all | 1.26 | 1.27 | 0.78 | 1.16 | **0.38**$^{***}$ | (0.01) |
| Drop 1 | 1.21 | 1.24 | 0.93 | 1.12 | **0.20**$^{***}$ | (0.00) |
| Drop 2 | 1.14 | 1.19 | 0.83 | 1.06 | **0.23**$^{**}$ | (0.01) |
| Drop 3 | 1.09 | 1.12 | 0.80 | 1.00 | **0.19**$^{**}$ | (0.04) |
| Drop 4 | 1.01 | 1.07 | 0.73 | 0.94 | **0.21**$^{**}$ | (0.03) |
| Drop 5 | 1.00 | 1.06 | 0.74 | 0.93 | **0.19**$^{**}$ | (0.02) |
| Drop 6 | 1.07 | 1.07 | 0.83 | 0.95 | **0.13**$^{*}$ | (0.10) |
| Drop 7 | 1.02 | 1.05 | 0.78 | 0.94 | 0.15 | (0.22) |
| Drop 8 | 1.06 | 1.07 | 0.85 | 0.97 | 0.12 | (0.21) |
| Drop 9 | 1.11 | 1.05 | 0.94 | 0.96 | 0.02 | (0.64) |
| Drop 10 | 1.01 | 0.93 | 0.83 | 0.84 | 0.01 | (0.80) |

*Notes*: Out-of-sample Sharpe ratios of $MV$ and $MV_{TC}$. Column 2 and 3 provide Sharpe ratios before transaction costs and columns 4 and 5 Sharpe ratios after costs. Column 6 shows the difference between the Sharpe ratio of $MV_{TC}$ and $MV$ and column 7 indicates the p-value of the difference. Standard errors for $\Delta$SR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). The top row includes all 29 available currencies. In subsequent rows we remove one-by-one the currency with the highest median transaction cost over the previous 9 months. The data are monthly returns from January 1976 to February 2016.

Table 10: **Restricting 29 Admissible Currencies to Low Transaction Cost ones: Notional Value, Mean and Volatility**

| Drop # Top | Notional Value | | After TC Mean | | After TC Volatility | |
|---|---|---|---|---|---|---|
| TC Currencies | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| Keep all | 4.05 | 3.38 | 0.08 | 0.11 | 0.11 | 0.09 |
| Drop 1 | 4.00 | 3.39 | 0.10 | 0.10 | 0.11 | 0.09 |
| Drop 2 | 4.09 | 3.45 | 0.09 | 0.10 | 0.11 | 0.09 |
| Drop 3 | 4.08 | 3.39 | 0.09 | 0.09 | 0.11 | 0.09 |
| Drop 4 | 3.99 | 3.30 | 0.08 | 0.09 | 0.11 | 0.09 |
| Drop 5 | 3.87 | 3.14 | 0.08 | 0.08 | 0.10 | 0.09 |
| Drop 6 | 3.70 | 3.01 | 0.08 | 0.08 | 0.10 | 0.09 |
| Drop 7 | 3.62 | 2.97 | 0.08 | 0.08 | 0.10 | 0.09 |
| Drop 8 | 3.53 | 2.90 | 0.08 | 0.09 | 0.10 | 0.09 |
| Drop 9 | 3.34 | 2.82 | 0.10 | 0.09 | 0.10 | 0.09 |
| Drop 10 | 3.08 | 2.60 | 0.07 | 0.07 | 0.09 | 0.08 |

*Notes*: Columns 2 and 3 report the average notional values of $MV$ and $MV_{TC}$, column 4 and 5 report after cost average returns and columns 6 and 7 after costs return volatilities. The top row includes all 29 available currencies. In subsequent rows we remove one-by-one the currency with the highest median transaction cost over the previous 9 months. The data are monthly returns from January 1976 to February 2016.

Table 11: **Restricting 29 Admissible Currencies to Low Transaction Cost ones: Crash Risk**

| Drop # Top TC Currencies | After TC Skewness $MV$ | After TC Skewness $MV_{TC}$ | After TC Kurtosis $MV$ | After TC Kurtosis $MV_{TC}$ | After TC MDD $MV$ | After TC MDD $MV_{TC}$ |
|---|---|---|---|---|---|---|
| Keep all | 0.00 | 0.47 | 7.60 | 6.59 | -0.38 | -0.19 |
| Drop 1 | 0.19 | 0.38 | 6.25 | 6.02 | -0.25 | -0.18 |
| Drop 2 | -0.45 | -0.06 | 6.25 | 5.67 | -0.31 | -0.20 |
| Drop 3 | -0.62 | -0.12 | 5.99 | 5.55 | -0.29 | -0.20 |
| Drop 4 | -0.46 | -0.08 | 5.91 | 5.78 | -0.34 | -0.20 |
| Drop 5 | -0.47 | -0.04 | 7.84 | 6.37 | -0.27 | -0.18 |
| Drop 6 | -0.39 | 0.01 | 7.42 | 6.83 | -0.22 | -0.19 |
| Drop 7 | -0.82 | 0.14 | 10.03 | 6.71 | -0.27 | -0.24 |
| Drop 8 | 0.00 | 0.47 | 10.26 | 8.31 | -0.23 | -0.23 |
| Drop 9 | 1.57 | 1.30 | 17.88 | 16.02 | -0.23 | -0.23 |
| Drop 10 | 0.01 | 0.14 | 6.58 | 7.99 | -0.22 | -0.22 |

*Notes*: Columns 2 and 3 report the after cost skewness of $MV$ and $MV_{TC}$, olumn 4 and 5 report after cost kurtosis and columns 6 and 7 after costs maximum draw downs. The top row includes all 29 available currencies. In subsequent rows we remove one-by-one the currency with the highest median transaction cost over the previous 9 months. The data are monthly returns from January 1976 to February 2016.

Table 12: **Currencies Sorted by Transaction Costs (using all 29 currencies)**

| # Top TC Currency | TC | SR | Mean | Vol | Corr | Skew | Kurt |
|---|---|---|---|---|---|---|---|
| # 1 | 0.3700 | 0.0400 | 0.4400 | 11.7900 | 0.5400 | -0.6800 | 6.3400 |
| # 2 | 0.1500 | 0.1000 | 1.1800 | 12.3800 | 0.5700 | -0.0700 | 5.5800 |
| # 3 | 0.1200 | 0.0700 | 0.8200 | 12.0800 | 0.5400 | -0.6200 | 4.8600 |
| # 4 | 0.1000 | 0.3200 | 4.0800 | 12.6100 | 0.6300 | -0.5200 | 5.0100 |
| # 5 | 0.0900 | 0.1200 | 1.4400 | 12.3900 | 0.5700 | -0.4000 | 4.2400 |
| # 6 | 0.0800 | 0.0100 | 0.0700 | 11.5600 | 0.5500 | -0.3500 | 4.4800 |
| # 7 | 0.0800 | 0.2300 | 2.6000 | 11.4400 | 0.5700 | -0.1700 | 4.6000 |
| # 8 | 0.0700 | 0.1500 | 1.5900 | 10.9300 | 0.6000 | 0.0100 | 3.9100 |
| # 9 | 0.0600 | 0.1100 | 1.3600 | 11.8600 | 0.5900 | -0.8400 | 9.9300 |
| # 10 | 0.0600 | -0.1200 | -1.3800 | 11.3400 | 0.6100 | -0.3000 | 5.0400 |

*Notes*: Return distributions of individual currencies. Column 2 reports the transaction costs, column 3 Sharpe ratios after transaction costs, columns 4 and 5 after cost average returns and volatilities, column 6 average correlations between a currency's exchange rate and all other exchange rates, columns 7 and 8 return skewness and kurtosis. Row $i$ provides values for a "portfolio" consisting of the currency with the $i$th highest transaction costs over the previous 9 months. The data are monthly returns from January 1976 to February 2016.

Table 13: **Restricting 15 Admissible Currencies to Low Transaction Cost ones: Sharpe Ratio**

| Drop # Top | Before TC Sharpe Ratios | | After TC Sharpe Ratios | | | |
|---|---|---|---|---|---|---|
| TC Currencies | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $\Delta$SR | (p-val) |
| Keep all | 0.96 | 1.04 | 0.75 | 0.96 | **0.21**$^{**}$ | (0.02) |
| Drop 1 | 0.96 | 1.04 | 0.75 | 0.96 | **0.21**$^{**}$ | (0.02) |
| Drop 2 | 1.03 | 1.05 | 0.81 | 0.96 | **0.16**$^{*}$ | (0.08) |
| Drop 3 | 1.05 | 1.07 | 0.82 | 0.98 | **0.16**$^{*}$ | (0.07) |
| Drop 4 | 0.96 | 1.04 | 0.75 | 0.95 | 0.20 | (0.24) |
| Drop 5 | 0.97 | 1.01 | 0.76 | 0.91 | 0.14 | (0.43) |
| Drop 6 | 0.93 | 1.02 | 0.74 | 0.92 | 0.18 | (0.36) |
| Drop 7 | 1.00 | 1.06 | 0.80 | 0.97 | 0.17 | (0.16) |

*Notes*: Out-of-sample Sharpe ratios of $MV$ and $MV_{TC}$. Column 2 and 3 provide Sharpe ratios before transaction costs and columns 4 and 5 Sharpe ratios after costs. Column 6 shows the difference between the Sharpe ratio of $MV_{TC}$ and $MV$ and column 7 indicates the p-value of the difference. Standard errors for $\Delta$SR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). The top row includes all the 15 developed currencies. In subsequent rows we remove one-by-one the currency with the highest median transaction cost over the previous 9 months. The data are monthly returns from January 1976 to February 2016.

Table 14: **sS Mean-Variance strategies, 29 Currencies**

| | Before TC Sharpe Ratios | | After TC Sharpe Ratios | | | |
|---|---|---|---|---|---|---|
| Threshold | $MV_{TC}$ | $MV_{sS}$ | $MV_{sS}$ | ΔSR | Costs | Turnover |
| 0 | 1.27 | 1.26 | 0.78 | **-0.38**$^{***}$ | 0.05 | 40.47 |
| 0.01 | 1.27 | 1.26 | 0.78 | **-0.38**$^{***}$ | 0.05 | 40.40 |
| 0.05 | 1.27 | 1.27 | 0.79 | **-0.36**$^{***}$ | 0.05 | 39.08 |
| 0.10 | 1.27 | 1.26 | 0.80 | **-0.35**$^{***}$ | 0.05 | 36.30 |
| 0.15 | 1.27 | 1.18 | 0.75 | **-0.40**$^{***}$ | 0.05 | 33.26 |
| 0.25 | 1.27 | 1.19 | 0.83 | **-0.33**$^{***}$ | 0.04 | 27.93 |
| 0.50 | 1.27 | 0.92 | 0.66 | **-0.49**$^{***}$ | 0.03 | 18.97 |
| 0.75 | 1.27 | 0.84 | 0.66 | **-0.49**$^{***}$ | 0.02 | 13.41 |
| 1.00 | 1.27 | 0.54 | 0.40 | **-0.76**$^{***}$ | 0.02 | 9.59 |

*Notes*: Summary statistics for, $MV_{sS}$, mean-variance strategies following the "sS" rule. The strategies are at the monthly frequency over the period January 1976 to February 2016 and use all the 29 available currencies. Each strategy, defined by its associated threshold weight (constant across currencies), trades in period $t$ in currency $i$ only if the differential from the $i$-th mean-variance weight and the one inherited from the previous period is more than the threshold in absolute value. Column 2 through 4 compare the out-of-sample Sharpe ratios of $MV_{sS}$ and $MV_{TC}$. Column 2 and 3 provide Sharpe ratios before transaction costs. Column 4 shows the after cost Sharpe ratios of $MV_{sS}$, column 5 its difference, ΔSR, with respect to the after cost Sharpe ratio of $MV_{TC}$ (negative values indicate underperformance relative to $MV_{TC}$). Standard errors for ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Column 6 and column 7 display the annualized trading costs and the turnover for the different mean-variance "sS" rule strategies. $^{***}$, $^{**}$, $^{*}$ indicate a statistical significance at the 1%, 5%, 10% level of ΔSR.

Table 15: **sS Mean-Variance strategies, 15 Currencies**

| | Before TC Sharpe Ratios | | After TC Sharpe Ratios | | | |
|---|---|---|---|---|---|---|
| Threshold | $MV_{TC}$ | $MV_{sS}$ | $MV_{sS}$ | ΔSR | Costs | Turnover |
| 0 | 1.04 | 0.96 | 0.75 | **-0.21**$^{**}$ | 0.02 | 27.60 |
| 0.01 | 1.04 | 0.96 | 0.75 | **-0.21**$^{**}$ | 0.02 | 27.54 |
| 0.05 | 1.04 | 0.96 | 0.76 | **-0.20**$^{**}$ | 0.02 | 26.66 |
| 0.10 | 1.04 | 0.95 | 0.76 | **-0.20**$^{**}$ | 0.02 | 25.24 |
| 0.15 | 1.04 | 0.97 | 0.79 | **-0.17**$^{**}$ | 0.02 | 23.60 |
| 0.25 | 1.04 | 0.81 | 0.66 | **-0.30**$^{***}$ | 0.02 | 20.38 |
| 0.50 | 1.04 | 0.77 | 0.66 | **-0.30**$^{**}$ | 0.02 | 14.82 |
| 0.75 | 1.04 | 0.64 | 0.57 | **-0.39**$^{**}$ | 0.01 | 11.06 |
| 1.00 | 1.04 | 0.54 | 0.49 | **-0.47**$^{**}$ | 0.01 | 8.10 |

*Notes*: Summary statistics for, $MV_{sS}$, mean-variance strategies following the "sS" rule. The strategies are at the monthly frequency over the period January 1976 to February 2016 and use the 15 developed currencies. Each strategy, defined by its associated threshold weight (constant across currencies), trades in period $t$ in currency $i$ only if the differential from the $i$-th mean-variance weight and the one inherited from the previous period is more than the threshold in absolute value. Column 2 through 4 compare the out-of-sample Sharpe ratios of $MV_{sS}$ and $MV_{TC}$. Column 2 and 3 provide Sharpe ratios before transaction costs. Column 4 shows the after cost Sharpe ratios of $MV_{sS}$, column 5 its difference, ΔSR, with respect to the after cost Sharpe ratio of $MV_{TC}$ (negative values indicate underperformance relative to $MV_{TC}$). Standard errors for ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Column 6 and column 7 display the annualized trading costs and the turnover for the different mean-variance "sS" rule strategies. $^{***}$, $^{**}$, $^{*}$ indicate a statistical significance at the 1%, 5%, 10% level of ΔSR.

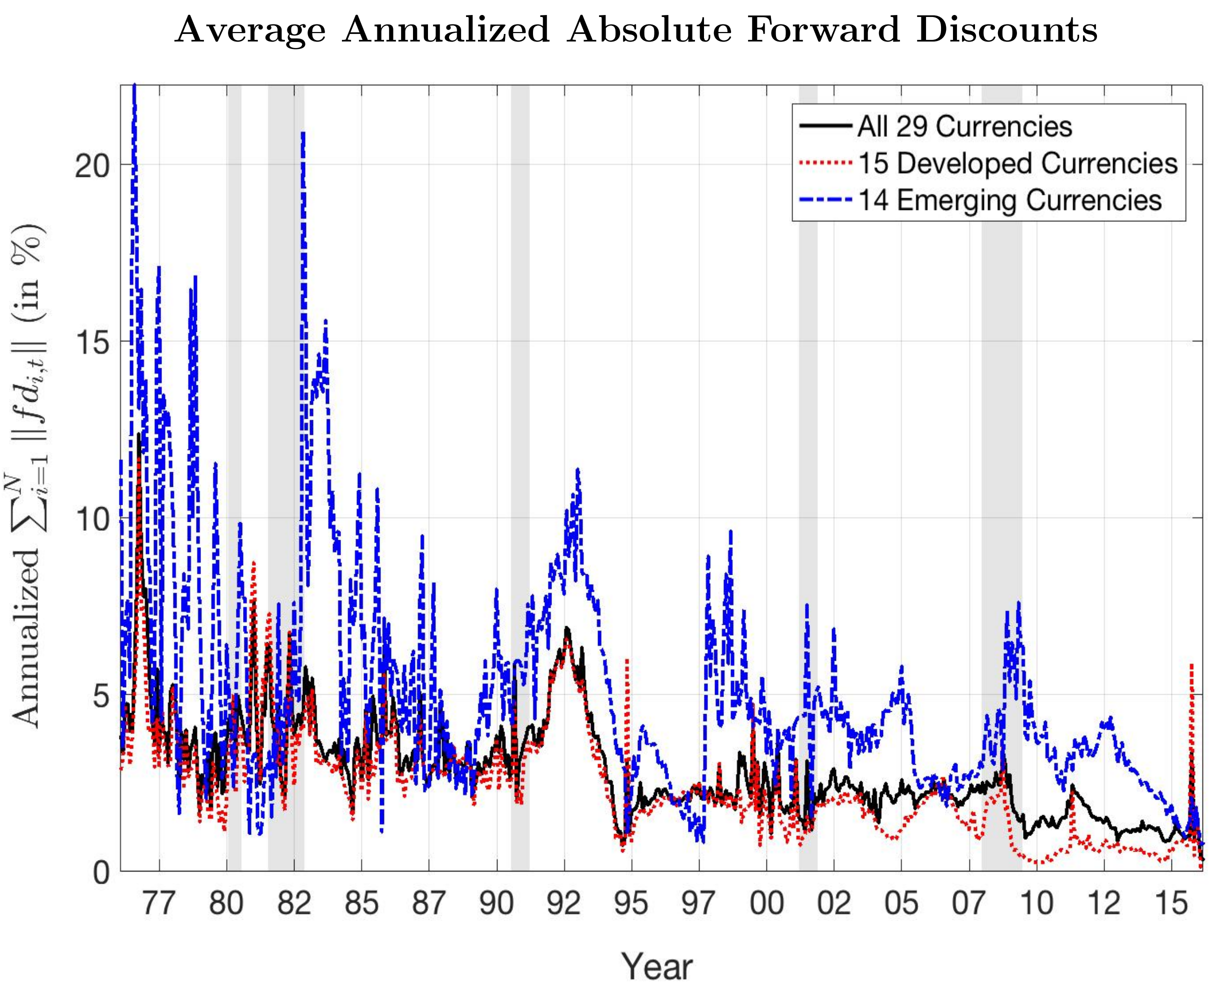


Figure 1: Average (across currencies) annualized absolute forward discounts (in percentage points) for the the full set of 29 currencies (black solid line), the subset of 15 developed currencies (red dotted line), and the subset of 14 emerging currencies (blue dashed line) from January 1976 to February 2016. Grey shaded areas indicate NBER recessions.

## Importance of Transaction Costs in FX Markets

**All 29 Currencies:**

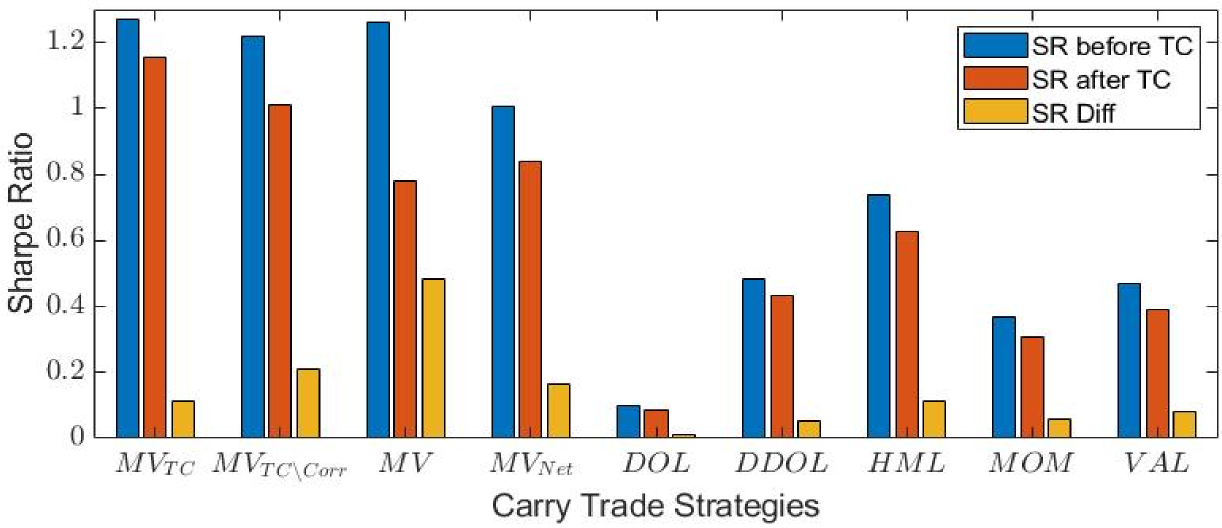


**15 Developed Currencies:**

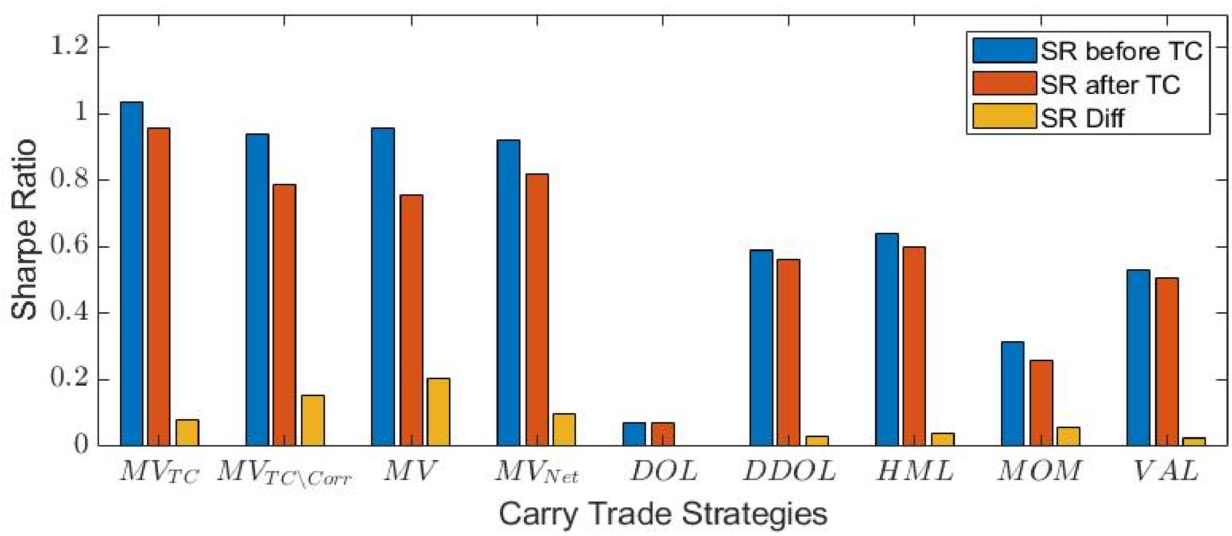


Figure 2: Annualized out-of-sample Sharpe ratios before (blue bars to the left) and after (green bars in the middle) transaction costs of various currency trading strategies and transaction costs (yellow bar to the right) paid by them. $MV_{TC}$ is the mean-variance optimized portfolio which optimizes over transaction costs. $MV_{TC\setminus Corr}$ is the mean-variance optimized portfolio which optimizes over transaction costs but makes the simplifying assumption that assets are uncorrelated. $MV$ is the mean-variance optimized portfolio without taking into account transaction costs in the optimization. $MV_{Net}$ is analogous to $MV$ only it is based on returns net of costs. $DOL$ invests equally in all bilateral carry trades. $DDOL$ takes a long position in $DOL$ if the median exchange rate forward discount is positive, and a short position otherwise. $HML$ sorts bilateral carry trades according to the forward discount into quintiles and shorts the bottom and invests in the top quintile. $MOM$ sorts bilateral carry trades according to their past 12 month performance into quintiles and shorts the bottom and invests in the top quintile. $VAL$ sorts bilateral carry trades according to the power purchase parity adjusted exchange rate into quintiles and shorts the top quintile (overvalued currencies with high real exchange rates) and invests in the bottom quintile (undervalued currencies with low real exchange rates). The data are monthly returns for our full set of 29 currencies from January 1976 to February 2016.

## Mean-Variance Problem with TC: Case of 2 Risky Assets

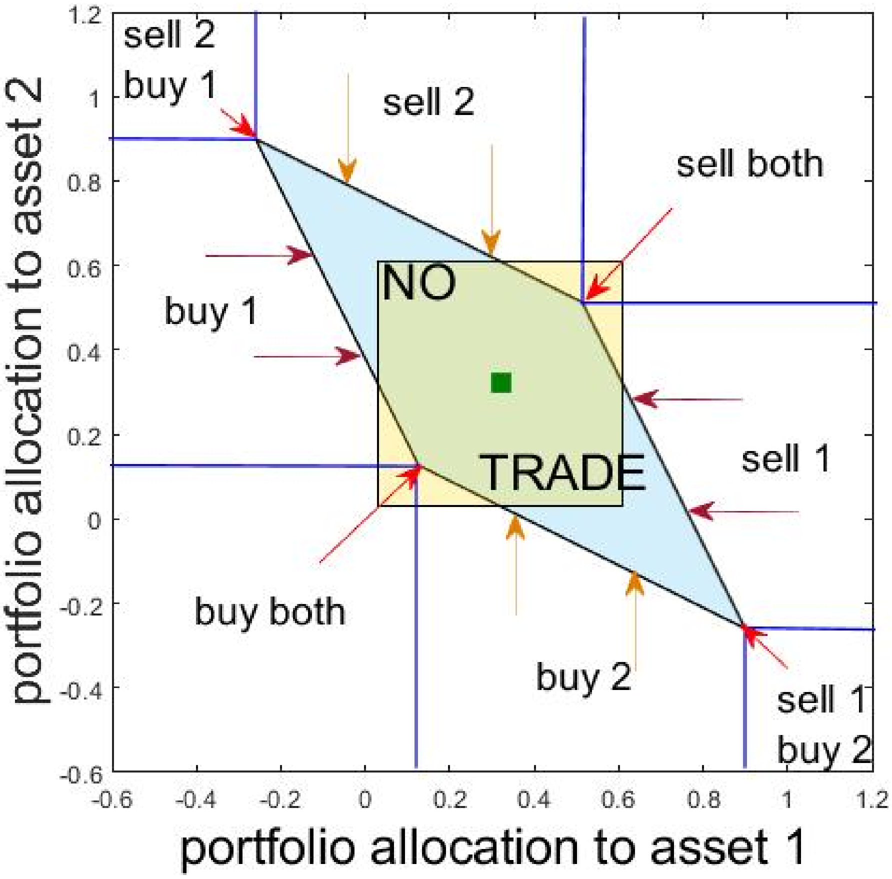


Figure 3: The investment opportunity set consists of two positively correlated risky assets. The horizontal axis measures the weight a portfolio places on asset 1 and the vertical axis the weight on asset 2. The green square is the optimal portfolio $\theta_{\mathbf{t}}^{\mathbf{MV}}$ if there are no transaction costs. The blue parallelogram illustrates the no trading region of $MV_{TC}$, which optimizes over transaction costs. The yellow square determines the no trading region of $MV_{TC\backslash Corr}$, which optimizes over transaction costs but assumes that the two assets are uncorrelated in the construction of the no trading region. If the initial position is within the no trading region, then the investor does not trade. If it is outside, then the investor trades towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$ until she reaches the boundary of the no trading region as indicated by the arrows. Purple, brown or orange colors of the arrows indicate that only asset 1, 2 or both assets are traded.

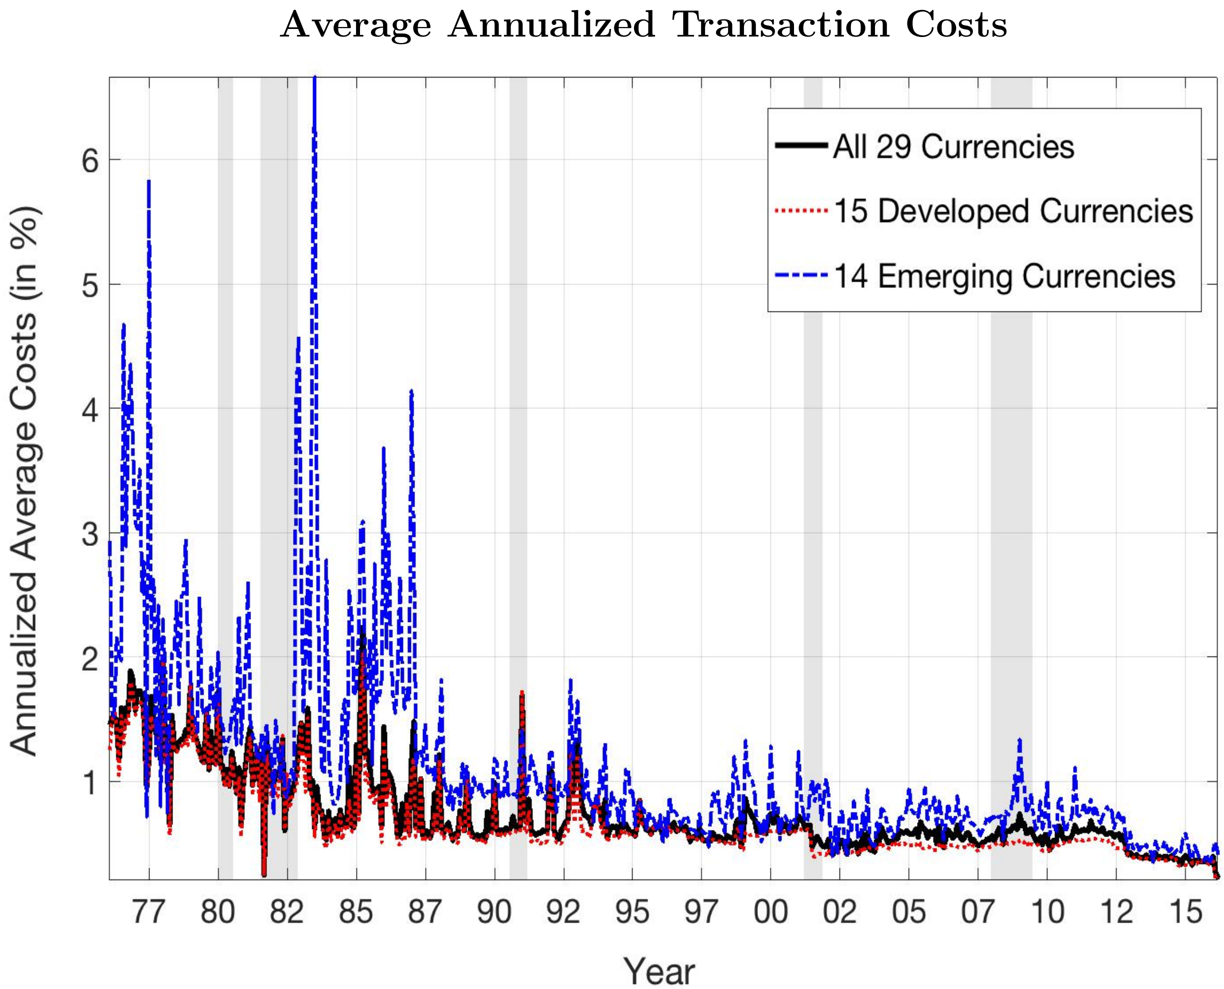


Figure 4: Average (across currencies) annualized costs (in percentage points) to trade a currency for the the full set of 29 currencies (black solid line), the subset of 15 developed currencies (red dotted line), and the subset of 14 emerging currencies (blue dashed line) from January 1976 to February 2016. Grey shaded areas indicate NBER recessions.

# Cumulative Returns of $MV$ and $MV_{TC}$

**Cumulative Returns Before Transaction Costs:**

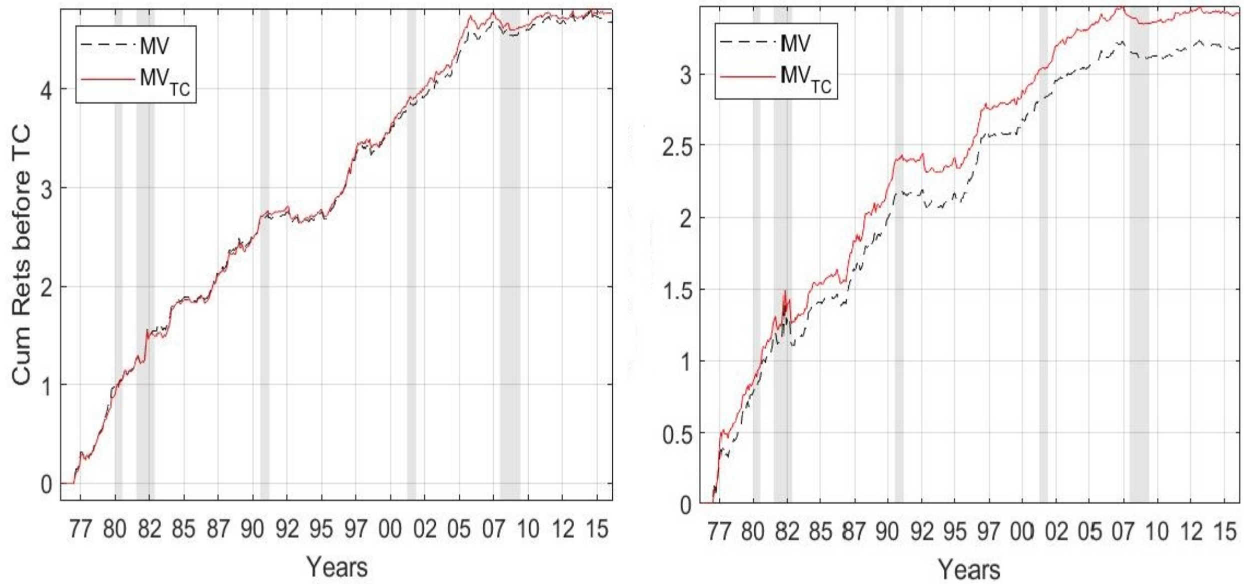


**Cumulative Returns After Transaction Costs:**

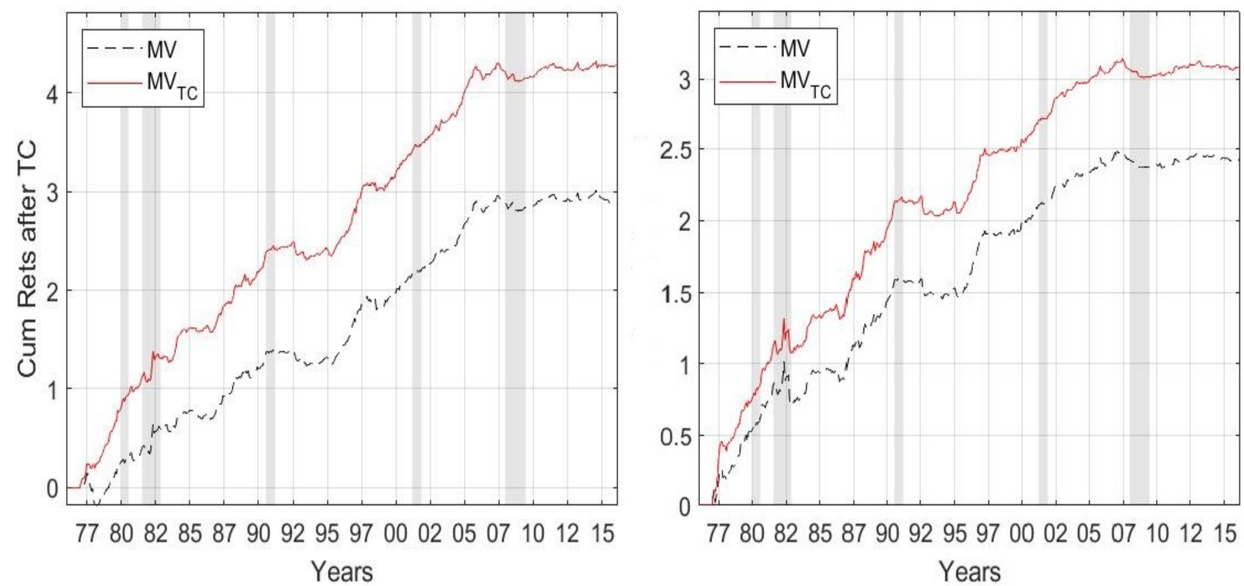


Figure 5: Time-series of cumulative returns of $MV$ (black dashed line) and $MV_{TC}$ (red solid line) for our set of 29 currencies (left panels) and for the set of 15 developed currencies (right panels) from January 1976 to February 2016. Returns before transaction costs are shown in the top panel, and returns after transaction costs in the bottom panel. Grey shaded areas indicate NBER recessions.

## Transaction Costs of $MV$ and $MV_{TC}$

**Cumulative Transaction Costs:**

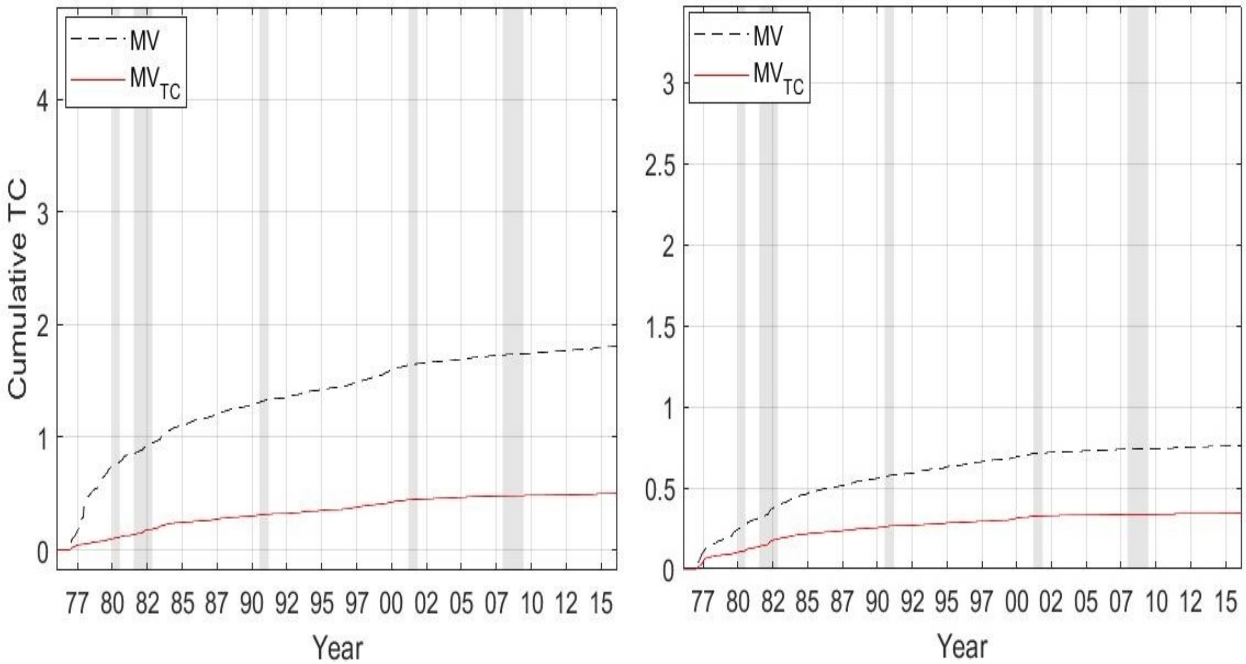


**Monthly Transaction Costs:**

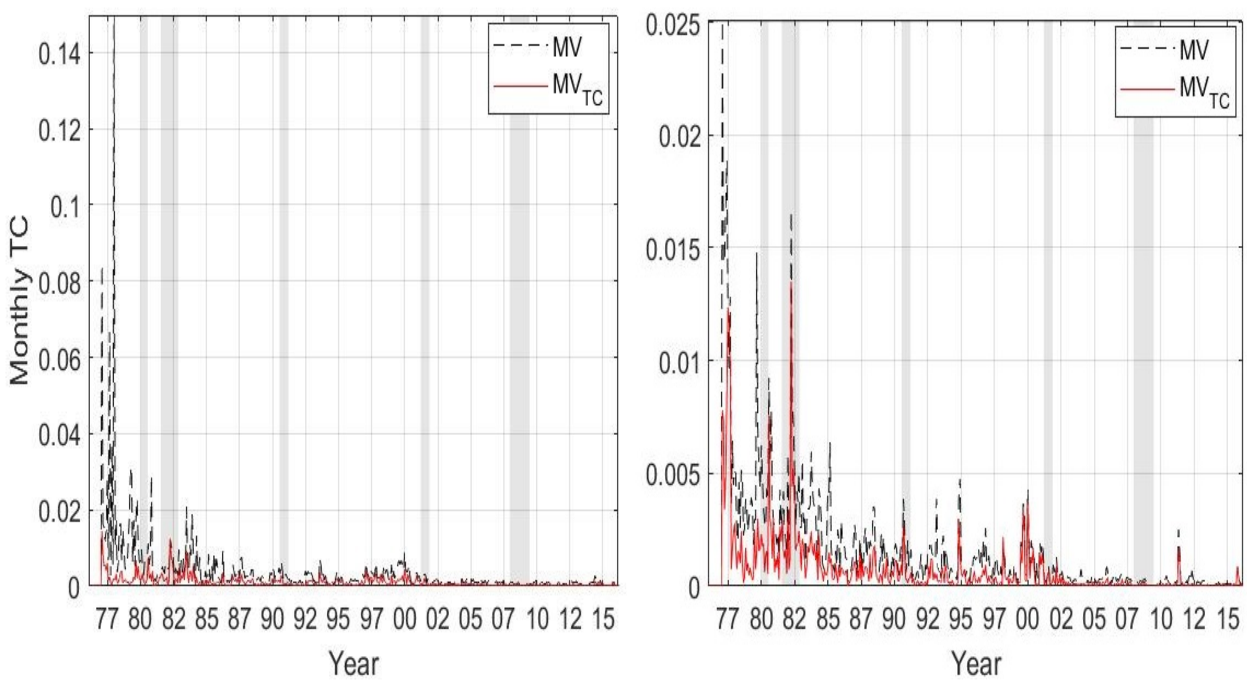


Figure 6: Time-series of transaction costs of $MV$ (black dashed line) and $MV_{TC}$ (red solid line) for our set of 29 currencies (left panels) and for the set of 15 developed currencies (right panels) from January 1976 to February 2016. Cumulative costs are shown in the top panel, and monthly costs in the bottom panel. Grey shaded areas indicate NBER recessions.

# Trading Activity of $MV$ and $MV_{TC}$

**Turnover:**

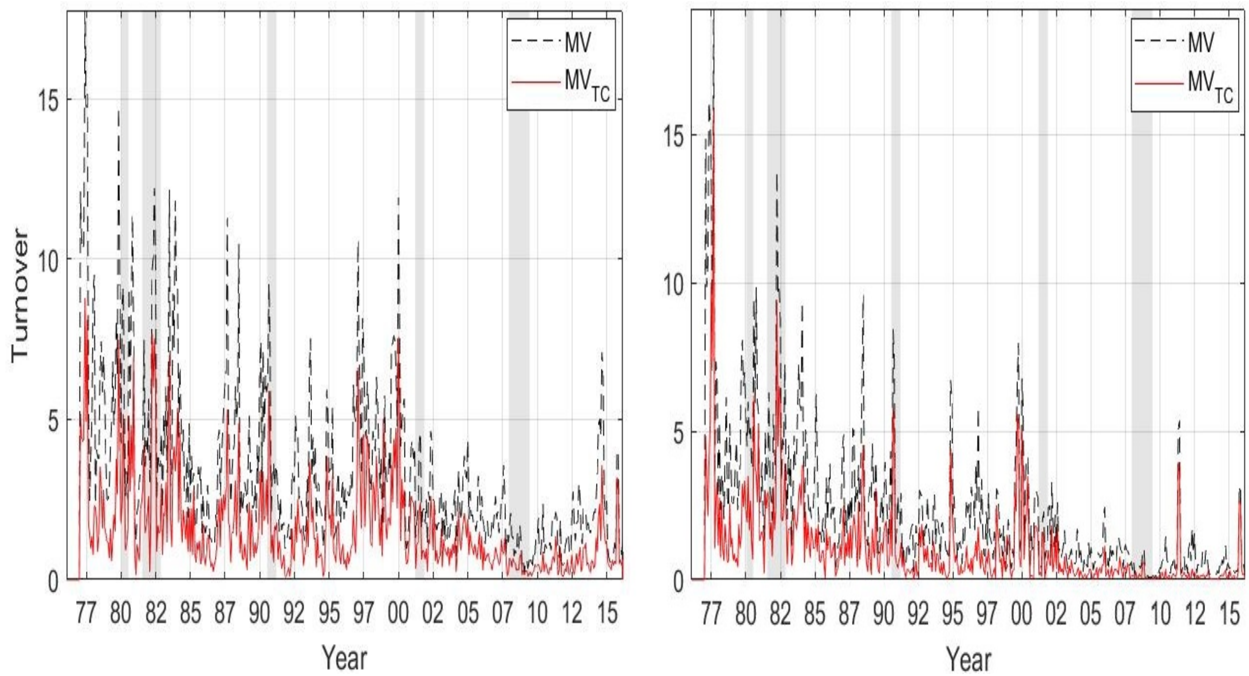


**Average Portfolio Weights and Standard Deviation Bars:**

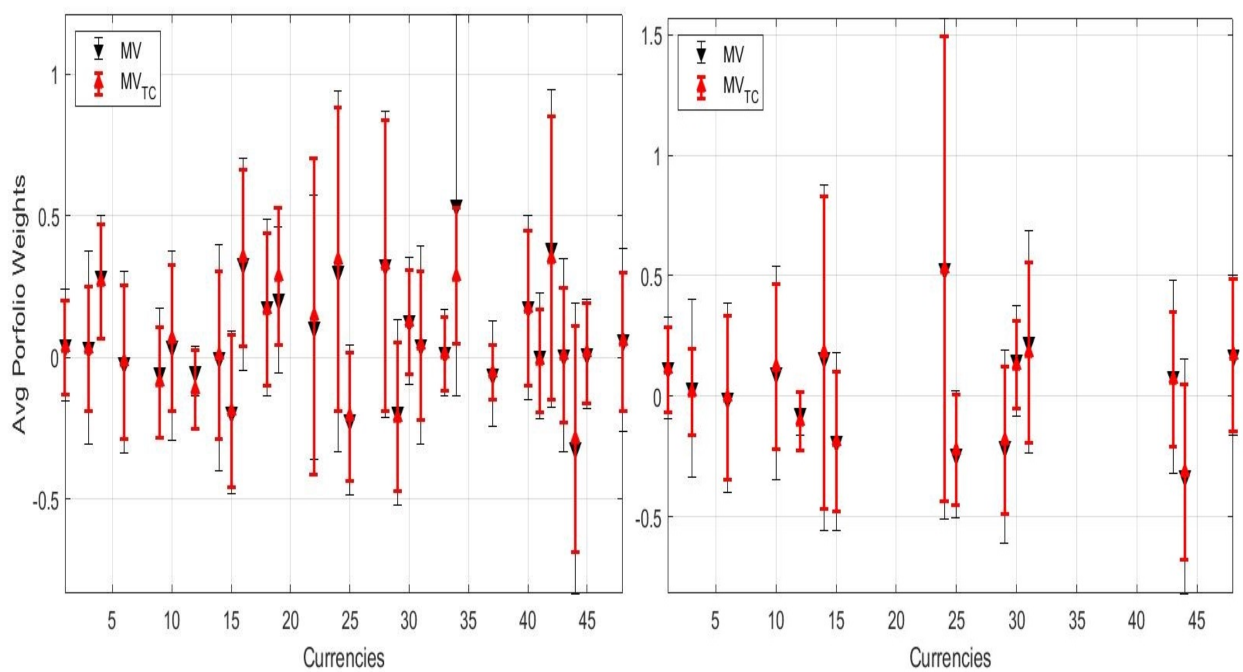


Figure 7: Top panel: Time-series of the turnover $\sum_i \|\theta_{i,t} - \theta_{i,t-1}\|$ of $MV$ (black dashed line) and $MV_{TC}$ (red solid line) for our set of 29 currencies (left panels) and for the set of 15 developed currencies (right panles) from January 1976 to February 2016. Grey shaded areas indicate NBER recessions. Bottom panel: Average portfolio weights and 1-standard deviation error bars of $MV$ (downward pointing triangle and thin black line) and $MV_{TC}$ (upward pointing triangle and thick red line).

## Notional Value of $MV$ and $MV_{TC}$

**All 29 Currencies:**

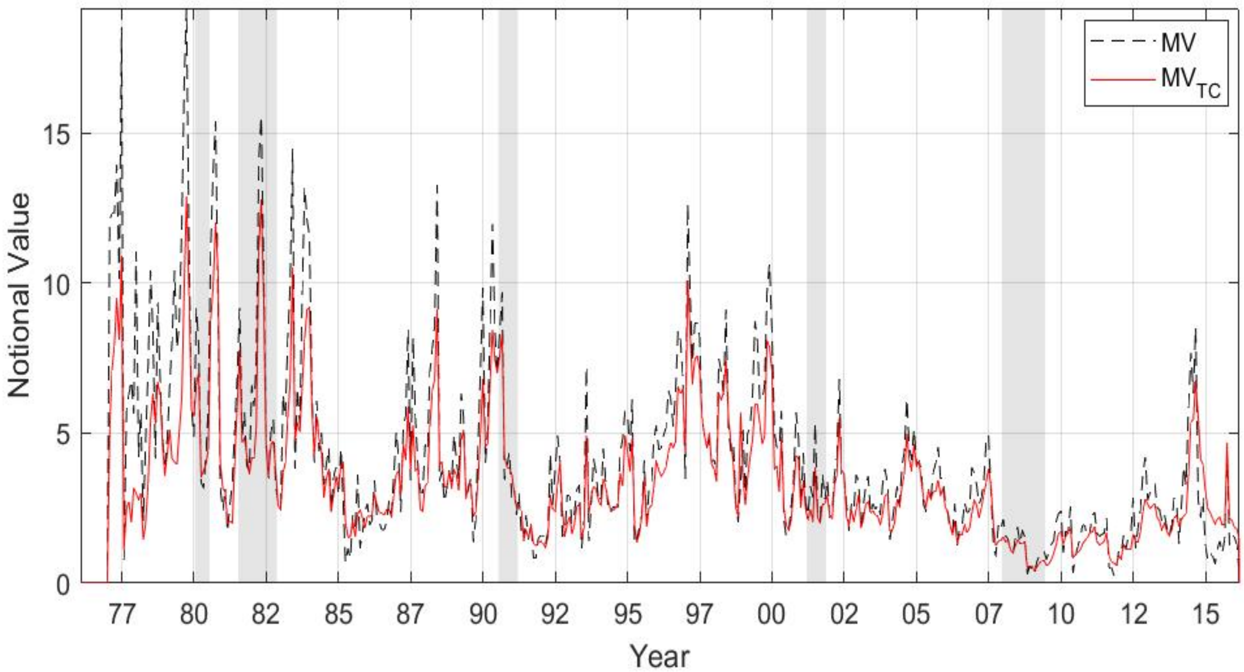


**15 Developed Currencies:**

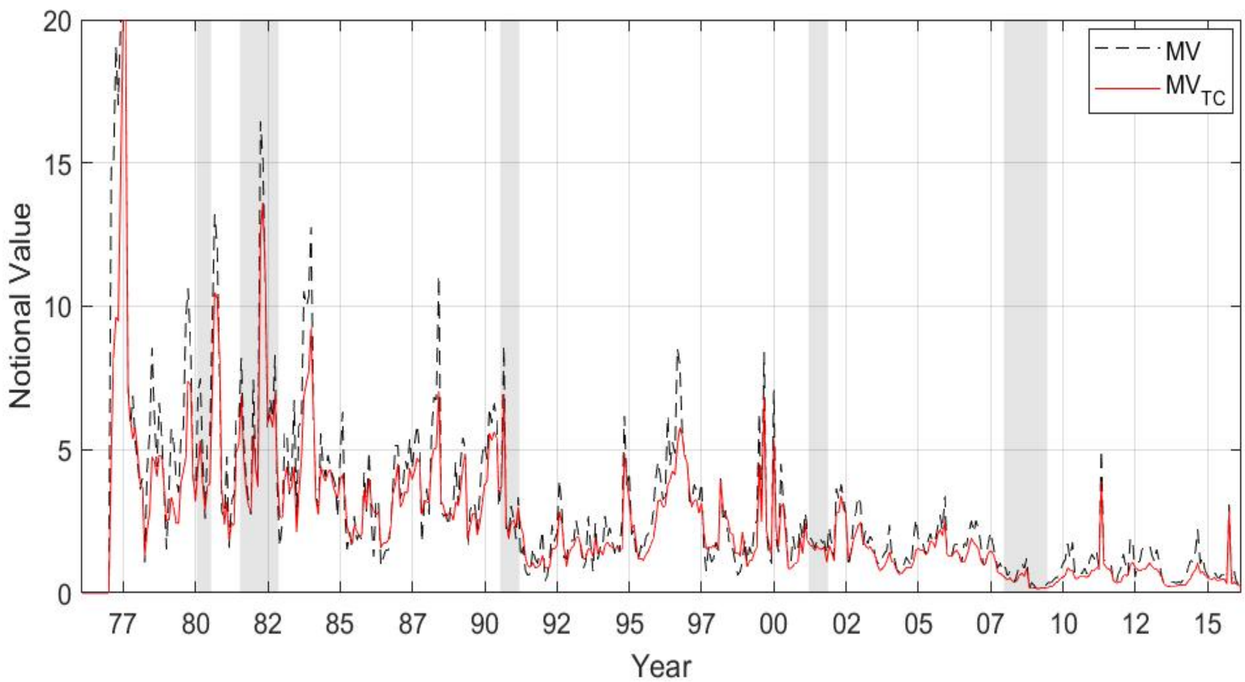


Figure 8: Time-series of the notional value or total dollar exposure $\sum_i \|\theta_{i,t}\|$ of $MV$ (black dashed line) and $MV_{TC}$ (red solid line) for our set of 29 currencies (top panel) and for the set of the 15 developed currencies (bottom panel) from January 1976 to February 2016. Grey shaded areas indicate NBER recessions.

# Average Correlations

**All 29 Currencies:**

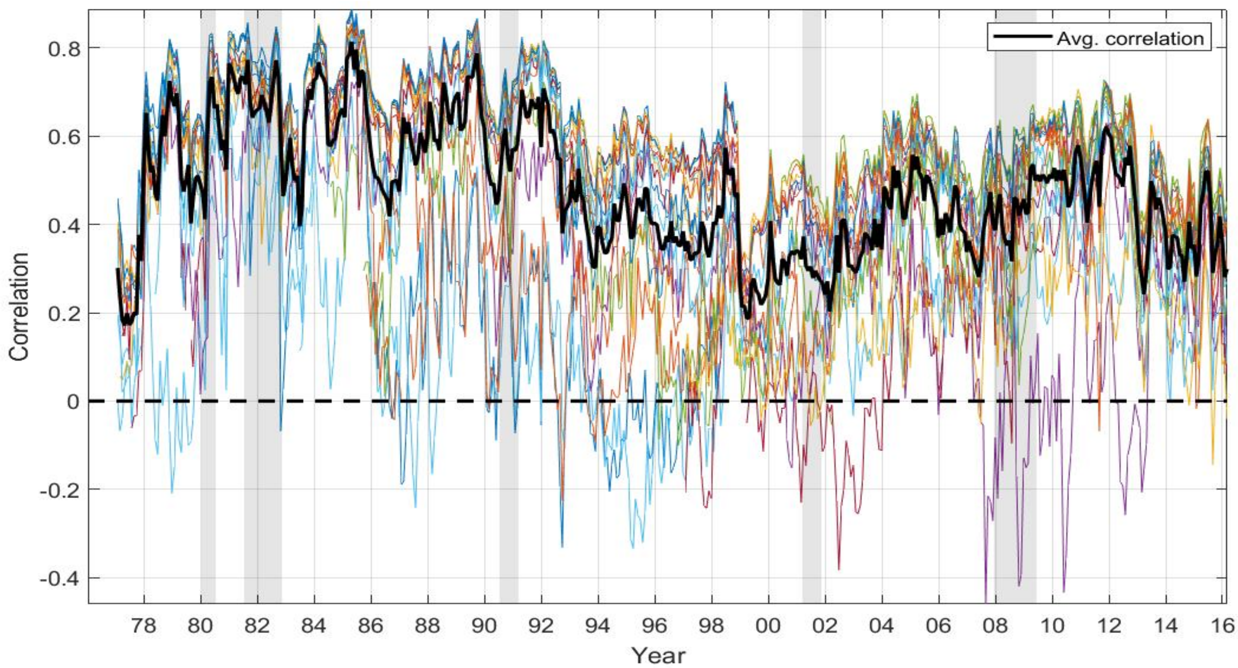


**15 Developed Currencies:**

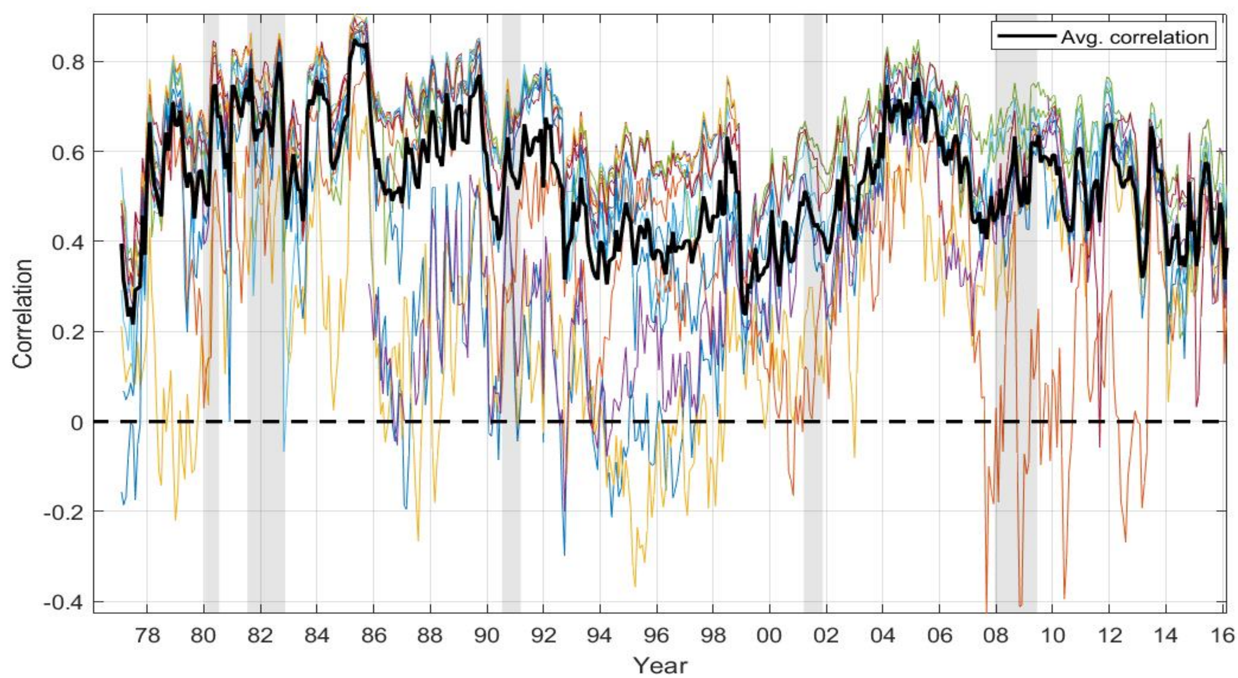


Figure 9: Time-series of the average conditional correlation of each exchange rate growth $i$ with all other exchange rate growths for our full set of $N_t = 29$ currencies (top panel) and for the set of 15 developed currencies (bottom panel), $\rho_{i,t} = \frac{1}{N-1}\sum_{j=1}^{N-1} Corr_t\left(\Delta x_{i,t}, \Delta x_{j,t}\right)$ estimated using daily data within each month from January 1976 to February 2016. The bold black line captures the time-series of the cross-sectional average across all correlations, $\rho_t = \frac{1}{N_t-1}\sum_{i=1}^{N_t} \rho_{i,t}$. Grey shaded areas indicate NBER recessions.

# Time-currency invariant Linear Price Impact

**After cost Sharpe ratios:**

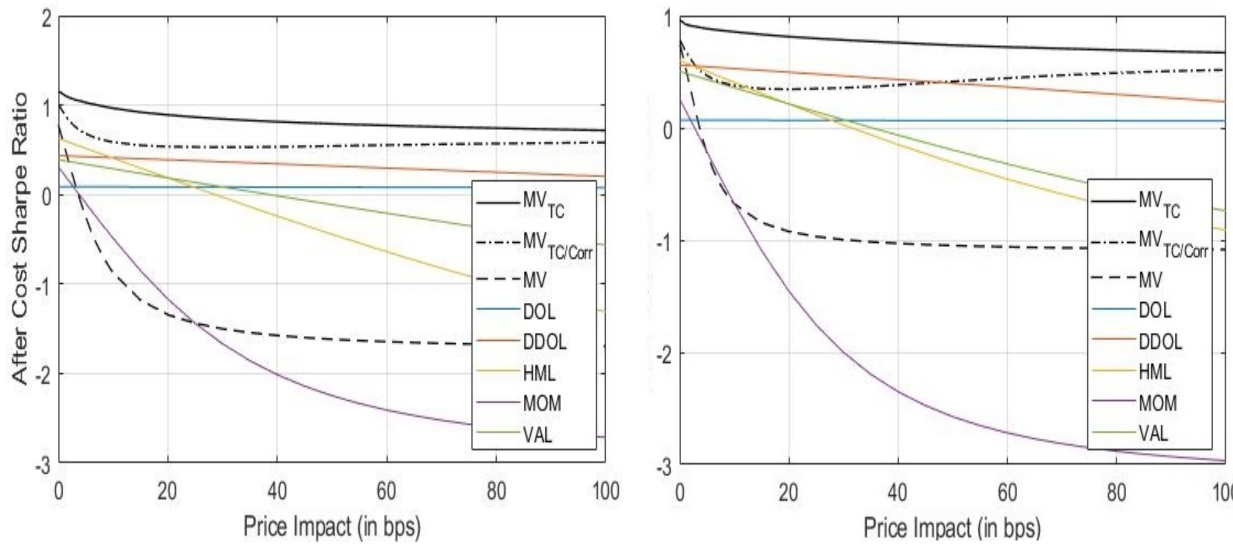


**Trade Aggressiveness:**

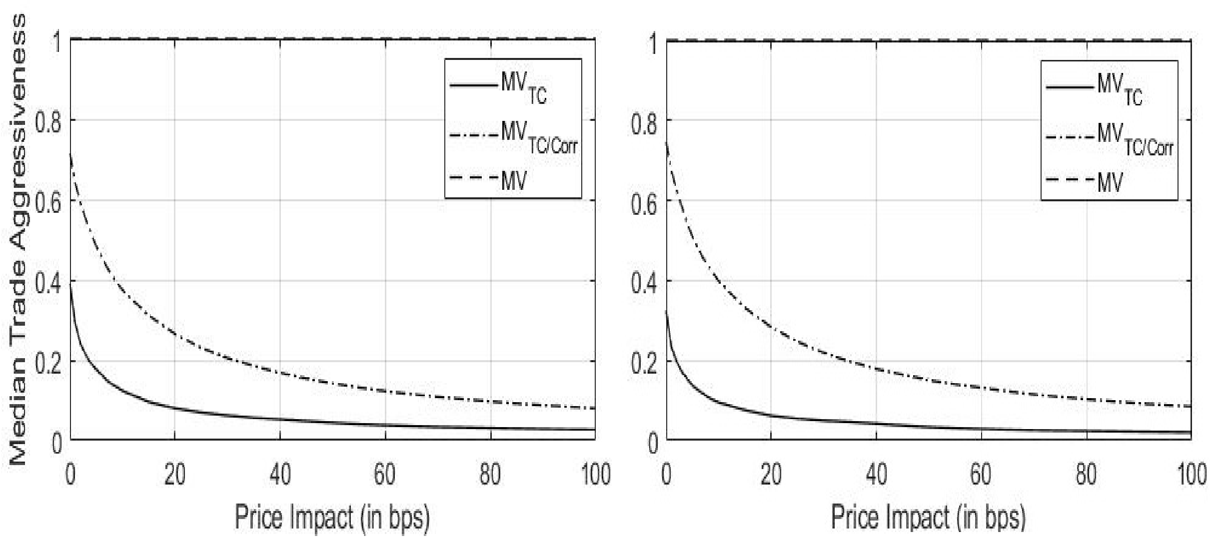


Figure 10: The top graphs show annualized out-of-sample Sharpe ratios after transaction costs of various trading strategies described in Figure 2, as a function of the price impact parameter $\pi \in [0, 100]$, which is measured in basis points. The price impact is a linear function of the size of the trades, constant across currencies $i$ and time $t$, i.e., $\pi = \pi^z_{i,t}$ $\forall i, t$ and $z \in \{P+, P-, S+, S-\}$, where $\pi^z_{i,t}$ are the diagonal elements of $\mathbf{\Pi_t}^{L,z}$. Similarly the bottom graphs plot the median trade aggressiveness $TA^S_t$ (defined in (1)) of the various trading strategies as a function of the price impact parameter. Left (right) plots report results for the set of 29 (15) currencies over the sample from January 1976 to February 2016.

# Time-currency invariant Square-root Price Impact

**After cost Sharpe ratios:**

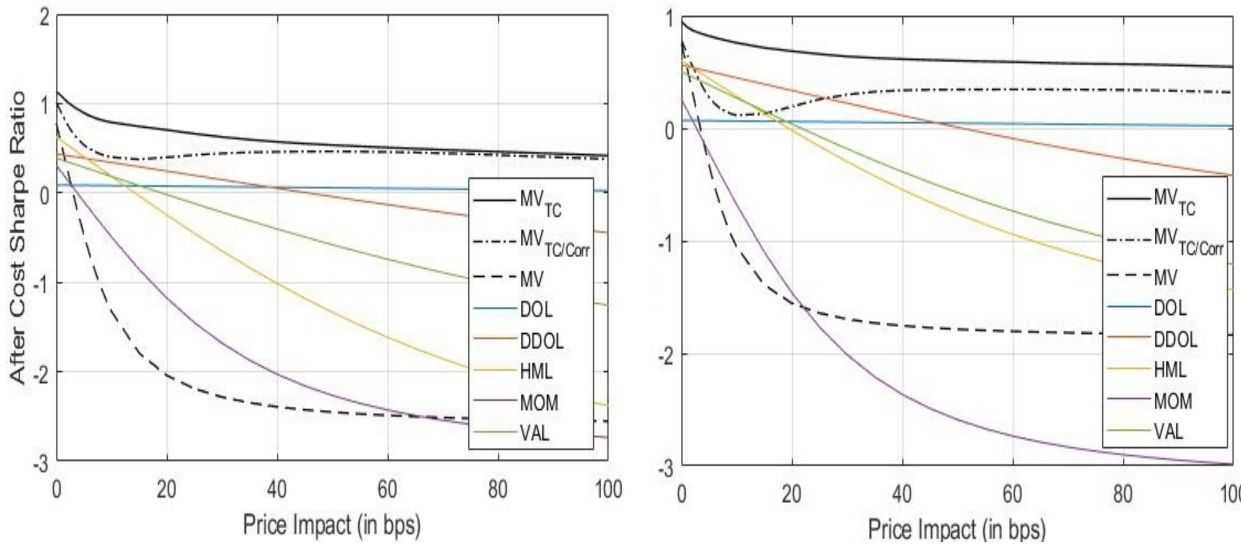


**Trade Aggressiveness:**

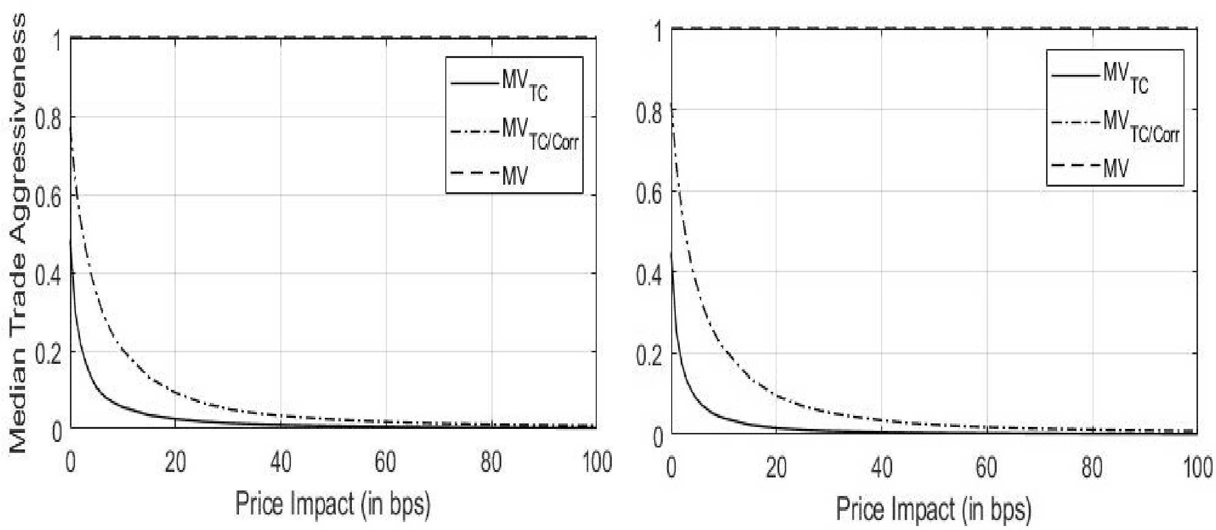


Figure 11: The top graphs show annualized out-of-sample Sharpe ratios after transaction costs of various trading strategies described in Figure 2, as a function of the price impact parameter $\pi \in [0, 100]$, which is measured in basis points. The price impact is a square-root function of the size of the trades, constant across currencies $i$ and time $t$, i.e., $\pi = \pi^z_{i,t}\ \forall i, t$ and $z \in \{P+, P-, S+, S-\}$, where $\pi^z_{i,t}$ are the diagonal elements of $\mathbf{\Pi_t}^{SR,z}$. Similarly the bottom graphs plot the median trade aggressiveness $TA^S_t$ (defined in (1)) of the various trading strategies as a function of the price impact parameter. Left (right) plots report results for the set of 29 (15) currencies over the sample from January 1976 to February 2016.

# Appendix

## A Robustness

Sections A.1, A.2 and A.3 provide robustness results of our main result in various subsamples. Section A.4 provides robustness results for various choice of risk aversion coefficients $\lambda$ backing the empirics with the theory. Section A.5 shows that our main results are quantitatively more important if we consider full round-trip transaction costs. In section A.6 we argue that the myopic feature of $MV_{TC}$ does not seem to be a first order problem and $MV_{TC}$ is the best strategy that we know to address transaction costs.

### A.1 NBER Recessions

We investigate the impact of recessions. We confirm our results in both NBER recessions and non-recession periods. Tables 22 and 23 summarizes the monthly excess returns of $MV$ and $MV_{TC}$ during NBER recessions (columns 1 and 2) and during non-recession periods (columns 3 and 4) for our the set of 29(15) currencies from 1976 to 2016. Sharpe ratios before transaction costs are more than twice(ten), in non-recession periods than during recessions. This difference is driven by both higher average returns and lower volatilities in non-recession periods. Transaction costs in recessions are at least 50% more for $MV_{TC}$ but almost identical for $MV$. The difference in costs between $MV$ and $MV_{TC}$ is statistically significant. $MV_{TC}$ outperforms $MV$ and the difference in after cost Sharpe ratios is 0.18(0.20) in recessions and 0.43(0.21) during non-recession periods. Such differences in Sharpe ratios are statistically significant at the 10% except for $MV_{TC}$ during recessions, where the p-value is 0.26. The p-values in recession periods are large because we only have 56 monthly observations during recessions and the power of the test is low. However, the economic magnitude of our result is meaningful in and out of recessions. It also become apparent how diversification helps in recessions, the after costs Sharpe ratios of $MV$ and $MV_{TC}$ are much higher when we are allowed to trade in the set of 29 currencies rather then constrained to the set of 15 developed ones.

We conclude, that our findings in Section 4 are present both in and out of recession periods and diversification helps in recessions.

## A.2 Subsamples before and after the Introduction of the Euro

The introduction of the Euro non-trivially affected the investment opportunity set in FX markets. Our results from Section 4 are stronger in the subsample before the introduction of the Euro, which is mostly due to the general decline in average transaction costs over time (Figure 4). Nonetheless we still find sizable improvements after the introduction of Euro, with statistical significance at the 10%(15%) level for the set of 29(15) currencies.

Tables 24 and 25 summarizes the monthly excess returns of $MV$ and $MV_{TC}$ for our set of 29(15) currencies before (columns 1 and 2) and after (columns 3 and 4) the introduction of the Euro on January 2nd, 1999. In both samples, there is no difference in Sharpe ratios before transaction costs between $MV$ and $MV_{TC}$ (except for noticeable improvement of 0.1 for $MV_{TC}$ in the pre-Euro sample for the set of 15 currencies). Transaction costs are substantially larger in the pre-Euro sample. Costs incurred by $MV$ and $MV_{TC}$ are 7.65%(3.40%) and 1.76%(1.30%) per year, a difference of 5.89%(2.10%), in the pre-Euro, and 1.74%(0.57%) and 0.56%(0.25%) per year, a difference of 1.18%(0.32%), in the post-Euro sample. The difference in costs between $MV$ and $MV_{TC}$ is statistically significant in both pre- and post-Euro samples. The Sharpe ratios after transaction costs of $MV$ and $MV_{TC}$ are 0.75(0.83) and 1.26(1.11) in the pre-Euro and 0.86(0.74) and 1.01(0.83) in the post-Euro sample. The difference of 0.50(0.26) in the pre-Euro sample is economically and statistically significant with a p-value of 0.011(0.03). The difference of 0.15(0.09) in the post-Euro sample is economically meaningful but statistically significant at the 10% level for the set of 29 currencies (it has a p-value of 0.148 for the set of 15 currencies). The decline in the difference in Sharpe ratios after costs between $MV_{TC}$ and $MV$ from the pre- to the post-Euro sample is mostly due to the strong decline in average transaction costs. However, this does not mean that optimizing transaction costs is irrelevant in the post-Euro era. The bottom graphs of Figure 5 show how the superior performance of $MV_{TC}$ over $MV$ is steady over the entire period. The cumulative returns after costs of $MV_{TC}$ are always above those of $MV$ and the spread is monotonically increasing. To conclude, the main results of Section 4 are confirmed in subsamples before and after the introduction of the Euro.

## A.3 Sample from November 1983 to February 2016

Our main analysis uses the sample from January 2nd, 1976 to March 2nd, 2016. The data before October 11th, 1983 is quoted against the Great British Pound (GBP), and we convert all data to exchange rates quoted against the USD (using mid quotes between the USD and GBP). The data quoted against the GBP are less reliable compared to the later sample quoted against the USD. Moreover, the bid and ask quotes after converting the 1976-1983 data to quotes against the USD do not exactly reflect the true bid and ask quotes against the USD, i.e., they are the bid and ask quotes against the GBP converted by the mid quote between USD and GBP. We show that our results are robust independent of whether we use the full sample from 1976 to 2016 or the shorter sample from 1983 to 2016.

Columns 3 and 4 of Tables 26 and 27 summarize the out-of-sample excess returns of $MV$ and $MV_{TC}$ for our set of 29(15) currencies from November 1983 to February 2016. The Sharpe ratios before transaction costs of $MV$ and $MV_{TC}$ are 1.17(1.04) and 1.20(1.04), which is 0.09(0.08) and 0.07(0.00) lower than in the full sample from 1976 to 2016. The costs paid by $MV$ and $MV_{TC}$ are 2.76%(1.19%) and 0.86%(0.44%) per year. These numbers are lower than in the full sample, which is consistent with the fact that average transaction costs are decreasing over time (Figure 4). $MV_{TC}$ costs less than half of $MV$ to implement. The difference in costs between $MV$ and $MV_{TC}$ is statistically significant. In contrast and more importantly, the Sharpe ratios after transaction costs are very similar in the sample starting in 1983 and in the full sample. The Sharpe ratio of $MV$ is 0.88(0.88) and the one of $MV_{TC}$

is 1.11(0.98). The difference between the strategies is 0.23(0.10) and statistically significant with a p-value of 0.007(0.03). Therefore, we confirm the results from from our main sample in the smaller sample starting in 1983.

## A.4 Risk Aversion Coefficient $\lambda$

We show, both empirically and theoretically, that the choice of the risk aversion coefficient $\lambda$ is irrelevant for our analysis and does not affect our results. We analyze the out-of-sample performance of $MV$ and $MV_{TC}$ for values $\lambda \in \{1, 5, 10, 25, 50, 100, 200\}$. Our baseline analysis sets $\lambda = 50$. This is a useful contribution for the one-period mean-variance setup with proportional costs since the independence of the Sharpe ratio measure to portfolios constructed with $\theta^{\mathbf{MV_{TC}}}$ was not known.

Tables 16 and 19 compare before and after cost Sharpe ratios of $MV$ and $MV_{TC}$ for the set of 29(15) currencies. For any value of $\lambda$ the before cost Sharpe ratios of $MV$ and $MV_{TC}$ are 1.26(0.96) and 1.27(1.04) while the after cost ratios are 0.78(0.75) and 1.16(0.96). As already shown in our main results, the difference of 0.38(0.20) in after cost Sharpe ratios between $MV_{TC}$ and $MV$ is highly statistically significant. Therefore, the Sharpe ratios and our conclusion are unaffected by the choice of $\lambda$.

While the independence of $\lambda$ from the Sharpe ratio of $MV$ is a well known straight-forward feature of the standard mean-variance setup, the same cannot be said for $MV_{TC}$. The following example with 2 excess risky returns (stacked in the vector $\mathbf{r^e}$) shows why. In the case of $MV$ $SR_{MV} \equiv \frac{E[\mathbf{r^e}'\theta^{\mathbf{MV}}]}{\sigma(\mathbf{r^e}'\theta^{\mathbf{MV}})} = \frac{\frac{1}{\lambda}\mu^{\mathbf{e}\prime}\mathbf{V}\mu^{\mathbf{e}}}{\frac{1}{\lambda}\sqrt{\mu^{\mathbf{e}\prime}\mathbf{V}\mu^{\mathbf{e}}}}$. The reason why the Sharpe ratio is independent from $\lambda$ is because $\theta^{\mathbf{MV}}$ is proportional to $\frac{1}{\lambda}$. To illustrate what is going on with $MV_{TC}$ further assume, without loss of generality, that there are no directional costs (i.e. $\mathbf{C^{P+}} = \mathbf{C^{P-}} \equiv \mathbf{C^{+}}$ and $\mathbf{C^{S+}} = \mathbf{C^{S-}} \equiv \mathbf{C^{-}}$). We need to show that $\theta^{\mathbf{MV_{TC}}}$ is proportional to $\frac{1}{\lambda}$. From the first order conditions for asset $i$ $\mathbf{V}_i\theta^{\mathbf{MV_{TC}}} \geq \frac{1}{\lambda}(\mu^{\mathbf{e}}{}_i - \mathbf{C_i^+})$ and $\mathbf{V}_i\theta^{\mathbf{MV_{TC}}} \leq \frac{1}{\lambda}(\mu^{\mathbf{e}}{}_i + \mathbf{C_i^-})$ where $\mathbf{V}_i$ represents the $i$-th row of $\mathbf{V}$. The intersections of these half-planes for all assets $i$ form the no trading region. Replacing the inequalities with equalities yields the equations for the borders of the no trading region; solving this system of equations yields the optimal weights for the corners of the region. From Dybvig and Pezzo (2019) we know that if the first(second) inequality is violated we need to buy(sell) asset $i$ until the inequality are restored with equality. For brevity we will pick two representative situations to show that $\theta^{\mathbf{MV_{TC}}}$ is indeed proportional to $\frac{1}{\lambda}$, the other situations are analogous. Suppose we are undervalued in asset 1 and overvalued in asset 2, then it is optimal to trade to the upper-left corner of the blue no trading region shown in Figure 3 and buy more of asset 1 while selling more of asset 2. Then $\theta^{\mathbf{MV_{TC}}\prime} = [\theta_{\mathbf{1}}^{\mathbf{MV}} - \frac{V_{12}}{V_{11}}\left(\frac{V_{12}\mathbf{C_1^+} + V_{11}\mathbf{C_2^-}}{\lambda(V_{22}V_{11} - V_{12}^2)}\right) - \frac{\mathbf{C_1^+}}{\lambda V_{11}}, \theta_{\mathbf{2}}^{\mathbf{MV}} + \frac{V_{12}\mathbf{C_1^+} + V_{11}\mathbf{C_2^-}}{\lambda(V_{22}V_{11} - V_{12}^2)}]'$ which is proportional to $\frac{1}{\lambda}$. Now consider a situation where we are only undervalued in asset 1, then it is optimal to only buy more of asset $i$. This violates the first inequality defining the no trading region, thus we need to trade until we restore it with an equality and leave the position in asset 2 unchanged at $\theta_2^0$. In this case $\theta^{\mathbf{MV_{TC}}\prime} = [\frac{1}{\lambda V_{11}}(\mu_{\mathbf{1}}^{\mathbf{e}} - \mathbf{C_1^+}) - V_{12}\theta_2^0, \theta_2^0]'$ with $\frac{1}{\lambda(V_{22} - V_{12}^2)}\left[\mu_{\mathbf{2}}^{\mathbf{e}} - \mathbf{C_2^+} - \frac{V_{12}}{V_{11}}(\mu_{\mathbf{1}}^{\mathbf{e}} - \mathbf{C_1^+})\right] \leq \theta_2^0 \leq \frac{1}{\lambda(V_{22} - V_{12}^2)}\left[\mu_{\mathbf{2}}^{\mathbf{e}} + \mathbf{C_2^+} - \frac{V_{12}}{V_{11}}(\mu_{\mathbf{1}}^{\mathbf{e}} - \mathbf{C_1^+})\right]$, hence $\theta^{\mathbf{MV_{TC}}}$ is proportional to $\frac{1}{\lambda}$.

Tables 17 and 20 show that the notional values, after cost average returns and after cost volatilities of $MV$ and $MV_{TC}$ linearly scale with the risk aversion coefficient $\lambda$. This linear relationship is well-known in the case of $MV$. It does not trivially hold in the case of $MV_{TC}$ but empirically we find that it holds up to a precision of four decimal places. Finally Tables 18 and 21 show that the choice of $\lambda$ has no effect on the skewness and kurtosis. In the case of $MV$ this holds mechanically. For $MV_{TC}$ this is need to be shown, nonetheless empirically skewness and kurtosis are virtually unaffected by the choice of $\lambda$. The maximum draw down linearly scales with $\lambda$. While this is again a trivial relationship in the case of $MV$, we find that it empirically holds for $MV_{TC}$ up to a precision of three decimal places.

To conclude, the choice of $\lambda$ has no effect on our result that taking costs into account in a mean-variance portfolio optimization in FX markets significantly improves the out-of-sample performance after costs, furthermore this fact can also be theoretically proven in a one-period model.

## A.5 Full Round-Trip Transaction Costs

As a robustness check we repeat our analysis and use full round-trip transaction costs, which assume that a position is fully closed and re-opened every month, instead of the costs specified in section 3.1. That is, instead of subtracting $\mathbf{C_t^{P+}}'\mathbf{\Delta_t^{P+}} + \mathbf{C_t^{P-}}'\mathbf{\Delta_t^{P-}} + \mathbf{C_t^{S+}}'\mathbf{\Delta_t^{S+}} + \mathbf{C_t^{S-}}'\mathbf{\Delta_t^{S-}}$ from the mid-quote realized returns we subtract the full round-trip costs $(\mathbf{C_t^{P+}} + \mathbf{C_t^{S+}})'\mathbf{\Delta_t^{P+}} + (\mathbf{C_t^{P-}} + \mathbf{C_t^{S-}})'\mathbf{\Delta_t^{S-}}$.

As expected, we find that trading in the presence of full round-trip costs yields higher costs. As we move from our baseline analysis with no roll-over fees to our robustness analysis with full round-trip costs the increase in transaction costs is least severe for $MV_{TC}$. While the average annualized costs of $MV_{TC}$, $MV_{TC\backslash Corr}$ and $MV$ in our baseline analysis for the set of 29(15) currencies are 1.25%(0.85%), 2.34%(1.63%) and 5.13%(2.19%), they are 2.76%(2.12%), 4.99%(3.90%) and 13.06%(5.61%) in our robustness analysis with full round-trip costs.

Table 28 to 29 summarize the strategies' performances. First, Table 28 confirms our main result that $MV_{TC}$ has a significantly higher Sharpe ratio than $MV$. Results exacerbate the insights we gathered from the analysis in the main text. While the baseline difference in after cost Sharpe ratios for the set of 29(15) currencies is 0.38(0.20) with a p-value of 0.01(0.03), with round-trip costs it becomes 0.78(0.35) with a p-value of 0.00(0.01). Therefore, we can view our results in the main text as conservative.

In the main analysis we documented how properly accounting for correlations while optimizing over costs is important. The first three rows of Table 29 demonstrates how this is still the case under full round-trip costs, especially for the set of the 15 developed currencies. Going from $MV$ to $MV_{TC}$ trhough $MV_{TC\backslash Corr}$ improves the after cost Sharpe ratios from 0.02(0.40) to 0.70(0.51) and 0.80(0.74) for the set of 29(15) currencies. With most of the improvement, a highly statistically significant 0.67(0.24), obtained switching from $MV(MV_{TC\backslash Corr})$ to $MV_{TC\backslash Corr}(MV_{TC})$. For the set of 29(15) currencies, the improvement 0.11(0.11) from switching from $MV_{TC\backslash Corr}(MV)$ to $MV_{TC}(MV_{TC\backslash Corr})$ is more modest and not significant at the 10% level with a p-value of 0.11(0.16).

Because a marginally significant Sharpe ratio increment of 0.11 from $MV_{TC\backslash Corr}$ to $MV_{TC}$ for the full set of 29 currencies is still economically important (it corresponds to an increment of 1.1% in risk premium given a 10% increment in volatility), we conclude that accounting for correlations in the optimization still captures a first order effect even in the presence of full round-trip costs.

Table 29 also illustrates that $MV_{TC}$ continuoes to outperforms all equally weighted strategies as well as $MV_{net}$.

## A.6 Heuristic Adjustments to Approximate the Optimal Portfolio in a Multi-Period Setting

As discussed in section 2, $MV_{TC}$ is the optimal solution in a single period model but in general it is suboptimal in a multi-period framework. Intuitively, there are two reasons for that. First, the time-series variation in the investment opportunity set introduces a hedging demand (Merton, 1971). Second, since the optimal portfolio and value function at time $t+1$ crucially depend on the initial position $\theta^{\mathbf{0}}_{\mathbf{t+1}}$ at time $t+1$, the portfolio choice at time $t$ should not only trade off current expected returns, risks and costs but also take into account the implications on $\theta^{\mathbf{0}}_{\mathbf{t+1}}$.

Suppose now we extend our model to a multi-period setting. The investor trades in every period $t$ and her utility at time $t$ is $U_t = E_t\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} u_\tau\right]$ where $E_t[.]$ is the conditional expectation operator and $\beta \in [0,1]$ is a subjective time discount factor of future period $\tau$ mean-variance utility $u_\tau = {\theta_\tau}'\mu^{\mathbf{e}}_\tau - \frac{\lambda}{2}{\theta_\tau}'\mathbf{V}_\tau\theta_\tau - \sum_{z\in\{P+,P-,S+,S-\}} \mathbf{\Delta^{z}_{t}}'\mathbf{C^{z}_{t}} + \frac{1}{2}PI_t(z,\mathbf{\Delta^{z}_{t}})\mathbf{\Delta^{z}_{t}}$. For simplicity, suppose the investment opportunity set is constant so that there is no hedging demand (Merton, 1971), i.e., $\mu^{\mathbf{e}}_\tau = \mu^{\mathbf{e}}$, $\mathbf{V}_\tau = \mathbf{V}$, $\mathbf{C^{z}_{\tau}} = \mathbf{C^{z}}, PI_t(z,\mathbf{\Delta^{z}_{t}}) = PI(z,\mathbf{\Delta^{z}_{t}})\ \forall z \in \{P+,S+,P-,S-\}$. If there are no transaction costs ($\mathbf{C^{P+}} = \mathbf{C^{S+}} = \mathbf{C^{P-}} = \mathbf{C^{S-}} = PI(P+,\mathbf{\Delta^{P+}_{t}}) = PI(P-,\mathbf{\Delta^{P-}_{t}}) = PI(S+,\mathbf{\Delta^{S+}_{t}}) = PI(S-,\mathbf{\Delta^{S-}_{t}}) = 0$), then it is well-known that the optimal solution in every period $t$ is the same as the solution in the single period model, $\theta_{\mathbf{t}}^{\mathbf{MV}^*} = \theta^{\mathbf{MV}} = \frac{1}{\lambda}\mathbf{V}^{-1}\mu^{\mathbf{e}}$, where the superscript $^*$ indicates that the portfolio is the solution to the multi-period setting. In contrast, if there are positive transaction costs then in general the optimal solution is not equal to the single period solution, $\theta_{\mathbf{t}}^{\mathbf{MV}^*_{\mathbf{TC}}} \neq \theta^{\mathbf{MV}_{\mathbf{TC}}}$.

Unfortunately, we do not have an algorithm to solve the multi-period model with many correlated assets. Instead, we use intuitive arguments to construct heuristic adjustments for $MV_{TC}$ to approximate $MV^*_{TC}$. If the investment opportunity set is constant, we expect that the no trading region of the multi-period strategy $MV^*_{TC}$ is smaller than for the single period strategy $MV_{TC}$. The marginal utility to move towards $\theta^{\mathbf{MV}}$ is larger in the multi-period setting than in the single period model because the benefit to be close to $\theta^{\mathbf{MV}}$ is reaped for multiple periods instead of only once.

A similar intuition applies to settings with stochastic changes in the investment opportunity set. We expect that the size of the no trading region depends inversely on the persistence in the state variables that determine the investment opportunity set. We illustrate the intuition using two extreme examples. First, if state variables are independently and identically distributed (i.i.d.), then $\theta_{\mathbf{t}}^{\mathbf{MV}}$ is also i.i.d. (because it is a function of the state variables). Intuitively, it is optimal to trade towards the unconditional average $\overline{\theta}^{\mathbf{MV}} = E\left[\theta_{\mathbf{t}}^{\mathbf{MV}}\right]$. Once we get close enough, it is optimal to choose $\theta_{\mathbf{t}}^{\mathbf{MV}^*}$ in the neighborhood of $\overline{\theta}^{\mathbf{MV}}$. This is because in expectation $\theta_{\mathbf{t+1}}^{\mathbf{MV}}$ is equal to $\overline{\theta}^{\mathbf{MV}}$, and thus, choosing $\theta_{\mathbf{t}}^{\mathbf{MV}^*}$ close to $\overline{\theta}^{\mathbf{MV}}$ means that $\theta_{\mathbf{t+1}}^{\mathbf{0}}$ is close to $\theta_{\mathbf{t+1}}^{\mathbf{MV}}$ and we expect low trading costs in period $t+1$. Therefore, once our portfolio is in the neighborhood of $\overline{\theta}^{\mathbf{MV}}$, the no trading region is large and we do not trade aggressively. This is why, on average, when the state variables are not persistent and mean-revert quickly we expect a large no trading region (potentially larger than the no trading region in the one period model). By the same logic, we expect that the no trading region is smaller (larger) when the investment opportunity set experiences a change which pushes $\theta_{\mathbf{t}}^{\mathbf{MV}}$ towards (away from) $\overline{\theta}^{\mathbf{MV}}$.

In our second extreme example state variables follow a random walk, i.e., changes are i.i.d., and accordingly, $\theta_{\mathbf{t}}^{\mathbf{MV}}$ follows a non-stationary process. It is then optimal to trade aggressively and set $\theta_{\mathbf{t}}^{\mathbf{MV}^*}$ close to $\theta_{\mathbf{t}}^{\mathbf{MV}}$. This is because it is our best guess that $\theta_{\mathbf{t+1}}^{\mathbf{MV}}$ is close to $\theta_{\mathbf{t}}^{\mathbf{MV}}$. Thus, choosing $\theta_{\mathbf{t}}^{\mathbf{MV}^*}$ close to $\theta_{\mathbf{t}}^{\mathbf{MV}}$ makes $\theta_{\mathbf{t+1}}^{\mathbf{0}}$ the best guess for $\theta_{\mathbf{t+1}}^{\mathbf{MV}}$ and we expect low trading costs in period $t+1$. Therefore, when the state variables are persistent we expect the no trading region is small (and potentially smaller than the no trading region in the one period model) and we trade aggressively.

Following our intuition, we propose the following heuristic solutions to approximate the unknown, true solution to the multi-period model. We define the cost multiplier

$$\begin{aligned}\mathbf{M_{i,t}}(c_1,a_1,c_2,a_2) \quad = \quad & c_1 + a_1 \times \rho\left(\frac{\mu_{\mathbf{i,t}}^{\mathbf{e}}}{\sigma_{i,t}^2}\right) + \left(1 - \left|\rho\left(\frac{\mu_{\mathbf{i,t}}^{\mathbf{e}}}{\sigma_{i,t}^2}\right)\right|\right) \times I_{\left\{\theta_{\mathbf{i,t}}^{\mathbf{MV_{TC}}} \in \left[\theta_{\mathbf{i,t}}^{\mathbf{0}}, \bar{\theta}_{\mathbf{i}}^{\mathbf{MV}}\right]\right\}} \\ & \times I_{\left\{\left(\theta_{\mathbf{i,t}}^{\mathbf{MV}} - \theta_{\mathbf{i,t}}^{\mathbf{0}}\right)\left(\bar{\theta}_{\mathbf{i}}^{\mathbf{MV}} - \theta_{\mathbf{i,t}}^{\mathbf{0}}\right) > 0\right\}} \times \left[c_2 + a_2 \times \rho\left(\frac{\mu_{\mathbf{i,t}}^{\mathbf{e}}}{\sigma_{i,t}^2}\right)\right], \end{aligned} \tag{2}$$

where $\rho(x_t)$ is the first auto-correlation operator of the time-series $x_t$, $\sigma_{i,t}^2$ is the $i$th diagonal element of $\mathbf{V_t}$, and the adjusted transaction costs for trading asset $i$ is

$$\mathbf{M_{i,t}}(c_1,c_2,a_1,a_2)\left(\mathbf{C_{i,t}^z}\mathbf{\Delta_{i,t}^z} + \pi_{i,t}^z(\mathbf{\Delta_{i,t}^z})^2\right), \quad \forall z \in \{P+, S+, P-, S-\}. \tag{3}$$

We conjecture that there exist a parameter configuration for $c_1, c_2, a_1, a_2$ such that the solution $\theta^{\mathbf{MV_{TC}^M}}$ of Problem 1 with the adjusted transaction costs (given in equation (3)) approximates the true solution $\theta_{\mathbf{t}}^{\mathbf{MV_{TC}^*}}$ in the multi-period model.

Notice that $\theta^{\mathbf{MV_{TC}^M}} = \theta^{\mathbf{MV_{TC}}}$ if $c_1 = 1, c_2 = a_1 = a_2 = 0$, and $\theta^{\mathbf{MV_{TC}^M}} = \theta^{\mathbf{MV}}$ if $c_1 = c_2 = a_1 = a_2 = 0$. Therefore, our heuristic adjustments nest the single period model solutions with and without transaction costs. More generally, via the multiplier equation (2), we make the perceived costs from equation (3): 1) on average higher or lower than the actual costs, 2) a function of the persistence of the state variables, 3) dependent on the relative path of $\theta_{\mathbf{t}}^{\mathbf{MV_{TC}}}$ with respect to its long run mean $\bar{\theta}^{\mathbf{MV}}$. Pre-multiplying the actual costs buy a higher (lower) than 1 currency-specific factor $\mathbf{M_{i,t}}$ is the heuristics adopted to induce, according to our intuitive insights, less (more) trading in our myopic strategy due to the temporal dynamics. Specifically, 1) $c_1$ captures the average level of the currency-specific factor, 2) $a_1$ captures the auto-correlation function which keeps track of the persistence of each currency price of risk $\rho\left(\frac{\mu^{\mathbf{e}}_{\mathbf{i,t}}}{\sigma^2_{i,t}}\right)$ (remember that $\theta^{\mathbf{MV_{TC}}}_{\mathbf{t}}$ is proportional to prices of risk), and 3) the second factor in equation (2) makes the relative path of $\theta^{\mathbf{MV^M_{TC}}}_{\mathbf{t}}$ dependent on the long run mean $\bar{\theta}^{\mathbf{MV}}$ (inducing more trades – a lower factor $\mathbf{M_{i,t}}$ - when the myopic weights are converging towards the long run mean and vice versa).[24]

We empirically assess the importance of our heuristic approximation without and with the time-currency invariant linear price impact analyzed in Section 4.3 for the set of 29 and 15 currencies. We can construct $\theta^{\mathbf{MV^M_{TC}}}$ and compute returns for any parameter configuration $c_1, c_2, a_1, a_2$. We then find the configuration which yields the maximum Sharpe ratio after costs.

Figure 12 plots the annualized after cost Sharpe ratios of the myopic portfolio $MV_{TC}$ (i.e., $c_1 = 1, c_2 = a_1 = a_2 = 0$) and the portfolio $MV^M_{TC}$ with the maximum after cost Sharpe ratio in the space $c_1 \in \mathbb{R}, c_2 = a_1 = a_2 = 0$, $(c_1, c_2) \in \mathbb{R}^2, a_1 = a_2 = 0$, $(c_1, c_2, a_1, a_2) \in \mathbb{R}^4$ as a function of the price impact parameter $\pi \in [0, 100]$ for the set of 29(15) currencies. The price impact, when present (i.e. for $\pi > 0$), is modeled as linear in the size of the trades $\mathbf{\Delta^z_t}$ for $z \in \{P+, P-, S+, S-\}$. [25]

In the absence of price impact, that is when $\pi = 0$, we observe only small and statistically insignificant differences between the Sharpe ratios of the myopic portfolio $MV_{TC}$ and any of the adjusted portfolios $MV^M_{TC}$. The strategies' after cost Sharpe ratios range between 1.156(0.962) and 1.185(0.980) with (unreported) very similar levels of trade aggressiveness. Therefore, in the absence of a price impact, our proposed heuristic approximations of the multi-period model solution do not improve the out-of-sample performance over the myopic solution $MV_{TC}$.

[24]We also multiply the term which characterizes to the reversion to the long run mean $\bar{\theta}^{\mathbf{MV}}$ by $\left(1 - \left|\rho\left(\frac{\mu^{\mathbf{e}}_{\mathbf{i,t}}}{\sigma^2_{i,t}}\right)\right|\right)$. This is motivated by the intuition that a higher persistence and longer expected time for the weights $\theta^{\mathbf{MV^M_{TC}}}_{\mathbf{t}}$ to revert to the mean implies that the long run mean is less important as a benchmark. In the extreme cases when the processes for $\frac{\mu^{\mathbf{e}}_{\mathbf{i,t}}}{\sigma^2_{i,t}}$ are unit roots, the long run mean does not exists, and thus, this term vanishes.

[25]This is without loss of generality given the similar results for the different price impact setups analyzed in Section **??**.

Price impacts can have a first order effects on the performance of our strategies. Those should be viewed as the average effects suffered over an interval of time in a more realistic dynamic model. However, in such frameworks it is feasible and optimal to smooth price impacts over time rather than bear their effects all at once (see for example Liu and Xu (2018)). For this reason we expect that applying our heuristics adjustments here should improve the performance. As the two plots in Figure 12 shows for $\pi > 0$, the heuristic adjustments improve the performance of $MV_{TC}$ in the presence of a price impact. The more flexible model (the dot-dashed green line) yields bigger and bigger Sharpe ratio improvements as we move from a price impact of 5 to 100 basis points, reaching improvements of approximately 0.07 at 100 basis points. Nonetheless, these improvements are second order and almost never significant. (except for the case of $MV_{TC}^{M3}$ at 100 basis points with associated p-value of 0.063)

Note that the construction of our heuristic adjustments and in particular the parameter search suffers from a severe look-ahead bias. Ex-ante we did not know which parameter configuration $c_1, c_2, a_1, a_2$ maximizes the after cost Sharpe ratio. It is in favor of $MV_{TC}$ that $MV_{TC}^{M}$, whether or not a price impact is considered, does not significantly outperform $MV_{TC}$ despite of the look-ahead bias.

To conclude, our heuristic analysis suggests that the myopic strategy $MV_{TC}$ earns a high out-of-sample Sharpe ratio and our intuitive adjustments to approximate the true solution in a dynamic multi-period setting do not substantially improve the performance.[26]

[26]Because it is impossible to gauge the approximation error of $MV_{TC}^{M}$, we cannot rule out the possibility that the true, unknown solution $MV_{TC}^{*}$ of the multi-period model significantly outperforms the myopic portfolio $MV_{TC}$ in out-of-sample tests.

# B Details on Portfolio Optimization Problem

## B.1 Existence and Uniqueness of Solutions

Theorem 4 below is the natural extension of Proposition 1 in Dybvig and Pezzo (2019) when there is no price impact as in Problem 1. The following two corollaries are the straightforward extensions that take price impact and PCA adjustments into account as in Problem 2 and 3 respectively.

**Theorem 4** *Given the assumption of a positive definite covariance matrix* $\mathbf{V_t}$*, a positive risk aversion coefficient* $\lambda > 0$*, and a cost structure of the form* $\mathbf{0} \leq \mathbf{C^{P-}_t} \leq \mathbf{C^{P+}_t}$*,* $\mathbf{0} \leq \mathbf{C^{S+}_t} \leq \mathbf{C^{S-}_t}$*, the optimal solution to Problem 1 exists and the optimal portfolio* $\theta_t^{MV_{TC}}$ *we trade to is unique. The optimal trades are such that* $\mathbf{\Delta^{P+}_{n,t}} + \mathbf{\Delta^{P-}_{n,t}} = \max(\theta_{n,t} - \theta^0_{n,t}, 0) + x_n$ *and* $\mathbf{\Delta^{S+}_{n,t}} + \mathbf{\Delta^{S-}_{n,t}} = -\min(\theta_{n,t} - \theta^0_{n,t}, 0) + x_{n,t}$ *where* $x_{n,t} \geq 0$*. If the round trip cost for the relevant combination* $(j,k) \in \{+,-\}$ *is positive,* $\mathbf{C^{Pj}_{n,t}} + \mathbf{C^{Sk}_{n,t}} > 0$*, then it is suboptimal to simultaneously buy and sell so that* $x_{n,t} = 0$*. In particular, if the round-trip trading cost is positive for all securities, the optimal trades are unique and Problem 1 is unique.*

**Proof.** Let $U(\mathbf{\Delta^{P+}_t}, \mathbf{\Delta^{P-}_t}, \mathbf{\Delta^{S+}_t}, \mathbf{\Delta^{S-}_t})$ be the objective function of Problem 1 and define $\hat{U}(\theta) \equiv \max\{U(\mathbf{\Delta^{P+}_t}, \mathbf{\Delta^{P-}_t}, \mathbf{\Delta^{S+}_t}, \mathbf{\Delta^{S-}_t})$ subject to the constraints of Problem 1$\}$. Define $\mathbf{\hat{\Delta}^{P+}_t}, \mathbf{\hat{\Delta}^{P-}_t}, \mathbf{\hat{\Delta}^{S+}_t}$ and $\mathbf{\hat{\Delta}^{S-}_t}$ such that $\mathbf{\hat{\Delta}^{P+}_{n,t}} + \mathbf{\hat{\Delta}^{P-}_{n,t}} = \max(\theta_{n,t} - \theta^0_{n,t}, 0) \geq 0$ and $\mathbf{\hat{\Delta}^{S+}_{n,t}} + \mathbf{\hat{\Delta}^{S-}_{n,t}} = -\min(\theta_{n,t} - \theta^0_{n,t}, 0) \geq 0$. Then $\hat{U}(\theta_t) = U(\mathbf{\hat{\Delta}^{P+}_t}, \mathbf{\hat{\Delta}^{P-}_t}, \mathbf{\hat{\Delta}^{S+}_t}, \mathbf{\hat{\Delta}^{S-}_t})$, because any other way of generating $\theta$ either has the same value (if trading in at least some securities with round-trip cost of zero) or smaller value. Therefore,

$$\hat{U}(\theta_t) = \theta_t'\mu^{\mathbf{e}} - \frac{\lambda}{2}\theta_t'\mathbf{V_t}\theta_t - \mathbf{\hat{\Delta}^{P+}_t}(\theta_t)'\mathbf{C^{P+}_t} - \mathbf{\hat{\Delta}^{P-}_t}(\theta_t)'\mathbf{C^{P-}_t} - \mathbf{\hat{\Delta}^{S+}_t}(\theta_t)'\mathbf{C^{S+}_t} - \mathbf{\hat{\Delta}^{S-}_t}(\theta_t)'\mathbf{C^{S-}_t} \quad (4)$$

Note that $-\mathbf{hat\Delta^{P+}_t}(\theta_t)'\mathbf{C^{P+}_t}, -\mathbf{\hat{\Delta}^{P-}_t}(\theta_t)'\mathbf{C^{P-}_t}, -\mathbf{\hat{\Delta}^{S+}_t}(\theta_t)'\mathbf{C^{S+}_t}$ and $-\mathbf{\hat{\Delta}^{S-}_t}(\theta_t)'\mathbf{C^{S-}_t}$ are concave (since $\mathbf{\hat{\Delta}^{P+}_t}(\theta_t)'\mathbf{C^{P+}_t}, \mathbf{\hat{\Delta}^{P-}_t}(\theta_t)'\mathbf{C^{P-}_t}, \mathbf{\hat{\Delta}^{S+}_t}(\theta_t)'\mathbf{C^{S+}_t}$ and $\mathbf{\hat{\Delta}^{S-}_t}(\theta_t)'\mathbf{C^{S-}_t}$ are convex) and the quadratic terms are strictly concave. Therefore $\hat{U}(\theta_t)$ is strictly concave and if an optimal

$\theta_t$ exists it is unique. Such an optimum does exists because $\hat{U}(\theta_t) < \hat{U}(\theta_t^0)$ outside of the compact set

$$\theta_t'\mu^{\mathbf{e}} - \frac{\lambda}{2}\theta_t'\mathbf{V_t}\theta_t \geq \theta_t^{0\prime}\mu^{\mathbf{e}} - \frac{\lambda}{2}\theta_t^{0\prime}\mathbf{V_t}\theta_t^0,$$

which is a multidimensional ellipse.

The final allocation $\theta_t$ can be achieved by any combination $(\mathbf{\Delta_t^{P+}}, \mathbf{\Delta_t^{P-}}, \mathbf{\Delta_t^{S+}}, \mathbf{\Delta_t^{S-}})$ such that $\mathbf{\Delta_t^{P+}} + \mathbf{\Delta_t^{P-}} = \hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{P+} + \hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{P-} + x_t = (\hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{P+} + \frac{1}{2}x_t) + (\hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{P-} + \frac{1}{2}x_t)$ and $\mathbf{\Delta_t^{S+}} + \mathbf{\Delta_t^{S-}} = \hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{S+} + \hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{S-} + x_t = (\hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{S+} + \frac{1}{2}x_t) + (\hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{S-} + \frac{1}{2}x_t)$, where $x_{n,t} \geq 0$. The utility function for these pairs is given by $\hat{U}(\theta_t) - \frac{1}{2}x_t'(\mathbf{C_t^{P+}} + \mathbf{C_t^{P-}} + \mathbf{C_t^{S+}} + \mathbf{C_t^{S-}})$. If round-trip trading costs are positive for all securities $n$, then $x_t = 0$ is the unique choice. Given the assumption of a cost structure of the form $\mathbf{C_t^{P-}} \leq \mathbf{C_t^{P+}}$, $\mathbf{C_t^{S+}} \leq \mathbf{C_t^{S-}}$, whenever it is optimal to buy more of some security $n$, the open short position (if any) will optimally be closed or reduced first (i.e. $\mathbf{\Delta_t^{P-}} \in [0, -\theta_{n,t}^0]$) and then a new long position (if any) will optimally be opened (i.e $\mathbf{\Delta_t^{P+}} \geq 0$.). Similarly, whenever it is optimal to sell more of some security $n$, the open long position (if any) will optimally be closed or reduced first (i.e. $\mathbf{\Delta_t^{S+}} \in [0, \theta_{n,t}^0]$) and then a new short position (if any) will optimally be opened (i.e $\mathbf{\Delta_t^{S-}} \geq 0$.). Thus, the optimal directional trades $(\hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{P+}, \hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{P-}, \hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{S+}, \hat{\mathbf{\Delta}}_\mathbf{t}^\mathbf{S-})$ are unique. ■

**Corollary 5** *If we include $\sum_{z\in\{P+,P-,S+,S-\}} \frac{1}{2}PI_t(z, \mathbf{\Delta_t^z})\mathbf{\Delta_t^z}$ with $PI_t(z, \mathbf{\Delta_t^z})$ as defined in Section 2.2 for $z \in \{P+, P-, S+, S-\}$ to take the price impact of trades into account, then the solution to the new problem, Problem 2, still exists and is unique.*

**Proof.** Problem 2 can more generally be rewritten as

$$\min_{\mathbf{x}} \quad \mathbf{q}'\mathbf{x} + \frac{1}{2}\mathbf{x}'\mathbf{H}\mathbf{x} + (1 - 1_Q)||\mathbf{\Pi}^{\frac{2}{3}}\mathbf{x}||_{1.5}^{1.5}$$

subject to

$$\mathbf{0} \leq \mathbf{x} \leq \overline{\mathbf{x}}$$

where the exact mapping is detailed in Section B.3. By Theorem 4 if $\mathbf{\Pi} = \mathbf{0}$ there exist a unique solution. This is because the objective function is continuous and strictly convex over the relevant domain. In the presence of price impact $\mathbf{\Pi} \neq \mathbf{0}$. If $1_Q = 1$ than the Hessian of the objective function is $\mathbf{H} = \lambda\bar{\mathbf{I}}'\mathbf{V}_t\bar{\mathbf{I}} + 2\mathbf{\Pi}$ which is positive definite since $\lambda\bar{\mathbf{I}}'\mathbf{V}_t\bar{\mathbf{I}}$ is positive definite as implied by Theorem 4 and $\mathbf{\Pi}$ is positive definite by assumption. If $1_Q = 0$ than $\mathbf{\Pi} \geq 0$ and diagonal by assumption and the Hessian of the objective function is $2\lambda\bar{\mathbf{I}}'\mathbf{V}_t\bar{\mathbf{I}} + \frac{3}{4}\mathbf{\Pi}diag(\mathbf{x}^{0.5})$ which is positive definite because it is the sum of a positive definite matrix $2\lambda\bar{\mathbf{I}}'\mathbf{V}_t\bar{\mathbf{I}}$ and a positive semi-definite matrix $\frac{3}{4}\mathbf{\Pi}diag(\mathbf{x}^{0.5})$. ■

**Corollary 6** *A solution to Problem 3 exists and it is unique.*

Since $\mathbf{V_t}$ is positive definite, from the previous corollary it follows that the objective function in Problem 3 is strictly concave.

## B.2 The No Trading Region with Directional Costs

In the main text we provide a graphical visualization of the no trading region of the optimal trading strategy for the case of 2 risky assets when $\mathbf{C_t^{P+}} = \mathbf{C_t^{P-}} = \mathbf{C_t^{S+}} = \mathbf{C_t^{S-}}$ and $\sum_{z\in\{P+,P-,S+,S-\}} \frac{1}{2} PI_t(z, \mathbf{\Delta_t^z})\mathbf{\Delta_t^z} = 0$ (Figure 3). Figure 13 generalizes the cost structure in this illustration without considering the price impact (i.e. $\sum_{z\in\{P+,P-,S+,S-\}} \frac{1}{2} PI_t(z, \mathbf{\Delta_t^z})\mathbf{\Delta_t^z} = 0$).[27]

Explicitly modeling a price impact would unnecessarily complicate the graph without yielding substantial new insights. With a price impact for any initial allocation from which it is worth trading it is optimal to stop trading before reaching the boundary of the no trading region. This is because costs are now function of the trades' sizes $(\mathbf{\Delta_t^{P+}}, \mathbf{\Delta_t^{P-}}, \mathbf{\Delta_t^{S+}}, \mathbf{\Delta_t^{S-}})$.

We choose the parameters of the investment opportunity set such that the 2 risky assets match the mean values of our full set of 29 currencies from 1976 to 2016. In particular, we set $\mu_\mathbf{t}^\mathbf{e} = 2.4\%$, $\rho = 0.5$, $\sigma_\mathbf{t} = 10\%$ (the diagonal elements of $\mathbf{V_t}$), $\mathbf{C_t^{P+}} = 1.45\%$ for the costs of increasing long positions, $\mathbf{C_t^{S+}} = 0.71\%$ for the costs of decreasing long positions,

[27]The no=trading region in the presence of price impact would be exactly the same. The optimal trades would be different. Depending on the starting point $\theta_\mathbf{t}^\mathbf{0}$ and the price impact parameters, $\mathbf{\Pi_t^{z,L}}$ or $\mathbf{\Pi_t^{z,L}}$, it might or might not be optimal to trade at all, and when it is optimal to trade, the trades move towards the border of the no trading region but they might stop before.

$\mathbf{C_t^{P-}} = 0.71\%$ for the costs of reducing short positions, and $\mathbf{C_t^{S-}} = 1.45\%$ for the costs of increasing short positions. We set the coefficient of risk aversion $\lambda = 5$.

The no trading region for $MV_{TC}$ is the blue polygon, the one for $MV_{TC\backslash Corr}$ is the yellow square, and the standard mean-variance optimum $\theta_{\mathbf{t}}^{\mathbf{MV}}$ (represented by the green square) lies inside both. There are two main implications stemming from directional costs: 1) the no trading region for both $MV_{TC}$ and $MV_{TC\backslash Corr}$ are no more centered around $\theta_{\mathbf{t}}^{\mathbf{MV}}$, 2) the no trading region for $MV_{TC}$ is still bigger than the one for $MV_{TC\backslash Corr}$ but is not a parallelogram anymore.

Because the costs for incrementing long positions for assets 1 and 2, ($\mathbf{C_t^{P+}} = 1.45\%$) are higher than those for decrementing them ($\mathbf{C_t^{S+}} = 0.71\%$), it is optimal to stay further away from $\theta_{\mathbf{t}}^{\mathbf{MV}}$ along the buy directions than the sell ones. So, both $MV_{TC}$ and $MV_{TC\backslash Corr}$ no trading regions are shifted downward (along the 45° line passing through $(-0.2, 0.2)$) with respect to locations where $\theta_{\mathbf{t}}^{\mathbf{MV}}$ would have been centered.

The no trading region for $MV_{TC}$, as in figure Figure 3, continues to be bigger than that of $\theta_{\mathbf{t}}^{\mathbf{MV}}$. The two would be equal only if assets were uncorrelated. Another way to see the same phenomenon is to compute the trade aggressiveness for the two strategies according to equation (1) and noticing that on average $MV_{TC\backslash Corr}$ trades more aggressively than $MV_{TC}$. This is because the only regions where $MV_{TC}$ gets us closer to $\theta_{\mathbf{t}}^{\mathbf{MV}}$ are the bright yellow corners of the $MV_{TC\backslash Corr}$ no trading region along the 45° line passing through $(-0.2, -0.2)$. These cover a total area less than that covered by the bright blue corners of the $MV_{TC}$ no trading region along the $-45°$ line (passing through the point (-0.2,0.75)), which represent the regions where $MV_{TC\backslash Corr}$ trades more aggressively.

The directional costs also shape the no trading region for $MV_{TC}$ in a peculiar way. If there were only two costs, $\mathbf{C_t^{P-}} = 0.71\%$ and $\mathbf{C_t^{S-}} = 1.45\%$ the no trading region for $MV_{TC}$ would be the dash-dotted parallelogram. In contrast, if the two costs were $\mathbf{C_t^{S+}} = 0.71\%$ and $\mathbf{C_t^{P+}} = 1.45\%$, then the no trading region for $MV_{TC}$ would be the dashed parallelogram. Notice that the latter is the actual no trading region for $MV_{TC}$ with the two extreme corners falling outside the positive orthant cut. This observations explains the shape (and the logic) of the actual $MV_{TC}$ no trading region. Whenever the starting position $\theta_{\mathbf{t}}^{\mathbf{0}}$ is far enough from $\theta_{\mathbf{t}}^{\mathbf{MV}}$ and at least one asset position is negative, then as we travel towards $\theta_{\mathbf{t}}^{\mathbf{MV}}$ and are in the negative territory, we are targeting the dash-dotted parallelogram along the prescribed optimal direction(s). However, as soon as we turn into the positive territory we switch our target to the the dashed parallelogram. Thus, the no trading region of $MV_{TC}$ is the dashed no trading region with the corners falling into negative territories cut. Once we cross the line dividing the negative from the positive territory we switch our target from the dashed-dotted to the dashed parallelogram, but at that point we are already inside the targeted no trading region and we optimally stop. This intuition also explains why the no trading region for $MV_{TC\setminus Corr}$ is a rectangle with no edges cut. This no trading region is all contained in the positive orthant and coincides with the target region for a problem where we neglect correlation and only have two costs, $\mathbf{C_t^{S+}} = 0.71\%$ and $\mathbf{C_t^{P+}} = 1.45\%$.

## B.3 Algorithms

Following the solution approach of Dybvig and Pezzo (2019), we can rewrite Problem 3 as a standard quadratic program of the form

$$\min_{\mathbf{x}} \quad \mathbf{q}'\mathbf{x} + \frac{1}{2}\mathbf{x}'\mathbf{H}\mathbf{x} + (1 - 1_Q)||\mathbf{\Pi}^{\frac{2}{3}}\mathbf{x}||_{1.5}^{1.5}$$

subject to

$$\mathbf{0} \leq \mathbf{x} \leq \overline{\mathbf{x}}$$

where $\mathbf{q}' \equiv \mathbf{C}' + \lambda(\theta_{\mathbf{t}}^{\mathbf{0}} - \tilde{\theta}_{\mathbf{t}}^{\mathbf{MV}})'\mathbf{V}_t\bar{\mathbf{I}}$, with $\tilde{\theta}_{\mathbf{t}}^{\mathbf{MV}} = \frac{1}{\lambda}\tilde{\mathbf{V}}_{\mathbf{t}}^{\mathbf{-1}}\mu_{\mathbf{t}}^{\mathbf{e}}$ (further details are provided in Section 2.3), $\mathbf{C}' \equiv [\mathbf{C}_t^{P+\prime}, \mathbf{C}_t^{P-\prime}, \mathbf{C}_t^{S+\prime}, \mathbf{C}_t^{S-\prime}]$, $\mathbf{H} = \lambda\bar{\mathbf{I}}'\mathbf{V}_t\bar{\mathbf{I}} + 1_Q 2\mathbf{\Pi}$, $\bar{\mathbf{I}} \equiv [I_N, I_N, -I_N, -I_N]$, $\mathbf{\Pi} =$
$$\begin{bmatrix} \mathbf{\Pi_t}^{P+} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Pi_t}^{P-} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{\Pi_t}^{S+} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{\Pi_t}^{S-} \end{bmatrix}, 1_L = \begin{cases} 1, & \text{if } \frac{1}{2}PI_t = \frac{1}{2}PI_t^L \text{ and } \mathbf{\Pi}_t^z = \mathbf{\Pi}_t^{z,L} \text{ for } z \in \{P+, P-, S+, S-\} \\ 0, & \text{if } PI_t = \sqrt{PI}_t \text{ and } \frac{1}{2}\mathbf{\Pi}_t^z = \frac{1}{2}\mathbf{\Pi}_t^{z,SR} \text{ for } z \in \{P+, P-, S+, S-\} \end{cases}$$
where $PI_t, \mathbf{\Pi}_t^{z,L}$ and $\mathbf{\Pi}_t^{z,SR}$ are defined in Section 2.2, $I_a$ is the $a \times a$ identity matrix, and $||\cdot||_p^p$

represent the vector $p$-norm operator to the power of $p$. The program returns the solution $\mathbf{x} \equiv [\mathbf{\Delta_t^{P+}}', \mathbf{\Delta_t^{P-}}', \mathbf{\Delta_t^{S+}}', \mathbf{\Delta_t^{S-}}']$ which is asset-wise bounded above for $\mathbf{\Delta_{i,t}^{P-}}$ by $-\min(\theta^0_{i,t}, 0)$, the $N_t + 1$-th through the $2N_t$-th element of $\overline{\mathbf{x}}$, and for $\mathbf{\Delta_{i,t}^{S+}}$ by $\max(\theta^0_{i,t}, 0)$, the $2N + 1$-th through the $3N$-th element of $\overline{\mathbf{x}}$, all the other elements in $\overline{\mathbf{x}}$ are set to $+\infty$. The optimal portfolio $\theta_t^{MV_{TC}}$ is finally obtained by

$$\theta_t^{MV_{TC}} = \theta_\mathbf{t}^\mathbf{0} + \overline{\mathbf{I}}\mathbf{x}.$$

Notice that:

- solving such program for every $t$ produces strategy $MV_{TC}$ (i) with no-price impact if $\mathbf{\Pi} = \mathbf{0}$, (ii) with linear price impact if $1_L = 1$ and $\mathbf{\Pi} \neq \mathbf{0}$, (iii) with square-root price impact if $1_L = 0$ and $\mathbf{\Pi} \neq \mathbf{0}$

- by setting $\mathbf{C} = \mathbf{\Pi}_t^{P+} = \mathbf{\Pi}_t^{P-} = \mathbf{\Pi}_t^{S+} = \mathbf{\Pi}_t^{S-} = \mathbf{0}$ we solve the standard mean-variance problem following the PCA approach proposed by Maurer et al. (2018a), doing so for every $t$ is the way to implement strategy $MV$

- by setting $\mathbf{\tilde{V}_t^{-1}} = \mathbf{V_t^{-1}}$ if $\mathbf{\Pi} = \mathbf{0}$ we would solve Problem 1, else we would solve Problem 2

- by replacing $\mathbf{V_t}$ with $\mathbf{V_t^d}$ (see Section 2.4 for more details) and solving the program for every $t$ delivers strategy $MV_{TC \backslash Corr}$

In any case this is a well-behaved convex program and we solve it using the Matlab Optimization ToolBox.

# C Data Sources: Spot and Forward Exchange Rates

In Table 30 we list the Datastream mnemonics for spot and forward exchange rate quotes against the USD, whereas those against the GBP are listed in Table 31. To obtain mid-, bid-and ask-exchange rates, the suffixes (ER), (EB) and (EO) are added to the corresponding mnemonics.

# D Tables and Figures of the Appendix

Table 16: **Sensitivity to the Risk Aversion Coefficient $\lambda$ (29 currencies): Sharpe Ratio**

| Risk Aversion | Before TC Sharpe Ratios | | After TC Sharpe Ratios | | | |
|---|---|---|---|---|---|---|
| Coefficient $\lambda$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $\Delta$SR | (p-val) |
| 1 | 1.26 | 1.27 | 0.78 | 1.16 | 0.38 | 0.00 |
| 5 | 1.26 | 1.27 | 0.78 | 1.16 | 0.38 | 0.00 |
| 10 | 1.26 | 1.27 | 0.78 | 1.16 | 0.38 | 0.01 |
| 25 | 1.26 | 1.27 | 0.78 | 1.16 | 0.38 | 0.01 |
| 50 | 1.26 | 1.27 | 0.78 | 1.16 | 0.38 | 0.01 |
| 100 | 1.26 | 1.27 | 0.78 | 1.16 | 0.38 | 0.01 |
| 200 | 1.26 | 1.27 | 0.78 | 1.16 | 0.38 | 0.01 |

*Notes*: Out-of-sample Sharpe ratios of $MV$ and $MV_{TC}$. Column 2 and 3 provide Sharpe ratios before transaction costs and columns 4 and 5 Sharpe ratios after costs. Column 6 shows the difference between the Sharpe ratio of $MV_{TC}$ and $MV$ and column 7 indicates the p-value of the difference. Standard errors for $\Delta$SR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). The rows report results for portfolios constructed with risk aversion coefficients $\lambda \in \{10, 25, 50, 100, 200\}$. The data are monthly returns from January 1976 to February 2016.

Table 17: **Sensitivity to the Risk Aversion Coefficient $\lambda$ (29 currencies): Notional Value, Mean and Volatility**

| Risk Aversion | Notional Value | | After TC Mean | | After TC Volatility | |
|---|---|---|---|---|---|---|
| Coefficient $\lambda$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| 1 | 202.75 | 168.88 | 4.10 | 5.32 | 5.26 | 4.61 |
| 5 | 40.55 | 33.77 | 0.82 | 1.06 | 1.05 | 0.92 |
| 10 | 20.27 | 16.88 | 0.41 | 0.53 | 0.53 | 0.46 |
| 25 | 8.11 | 6.75 | 0.16 | 0.21 | 0.21 | 0.18 |
| 50 | 4.05 | 3.38 | 0.08 | 0.11 | 0.11 | 0.09 |
| 100 | 2.03 | 1.69 | 0.04 | 0.05 | 0.05 | 0.05 |
| 200 | 1.01 | 0.84 | 0.02 | 0.03 | 0.03 | 0.02 |

*Notes*: Columns 2 and 3 report the average notional values of $MV$ and $MV_{TC}$, column 4 and 5 report after cost average returns and columns 6 and 7 after costs return volatilities. The rows report results for portfolios constructed with risk aversion coefficients $\lambda \in \{10, 25, 50, 100, 200\}$. The data are monthly returns from January 1976 to February 2016.

Table 18: **Sensitivity to the Risk Aversion Coefficient $\lambda$ (29 currencies): Crash Risk**

| Risk Aversion | After TC Skewness | | After TC Kurtosis | | After TC MDD | |
|---|---|---|---|---|---|---|
| Coefficient $\lambda$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| 1 | 0.00 | 0.46 | 7.58 | 6.58 | -18.78 | -9.46 |
| 5 | 0.00 | 0.46 | 7.58 | 6.58 | -3.76 | -1.89 |
| 10 | 0.00 | 0.46 | 7.58 | 6.58 | -1.88 | -0.95 |
| 25 | 0.00 | 0.46 | 7.58 | 6.58 | -0.75 | -0.38 |
| 50 | 0.00 | 0.47 | 7.58 | 6.59 | -0.38 | -0.19 |
| 100 | 0.00 | 0.47 | 7.58 | 6.59 | -0.19 | -0.09 |
| 200 | 0.00 | 0.47 | 7.58 | 6.59 | -0.09 | -0.05 |

*Notes*: Columns 2 and 3 report the after cost skewness of $MV$ and $MV_{TC}$, column 4 and 5 report after cost kurtosis and columns 6 and 7 after costs maximum draw downs. The rows report results for portfolios constructed with risk aversion coefficients $\lambda \in \{10, 25, 50, 100, 200\}$. The data are monthly returns from January 1976 to February 2016.

Table 19: **Sensitivity to the Risk Aversion Coefficient $\lambda$ (15 currencies): Sharpe Ratio**

| Risk Aversion | Before TC Sharpe Ratios | | After TC Sharpe Ratios | | | |
|---|---|---|---|---|---|---|
| Coefficient $\lambda$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $\Delta$SR | (p-val) |
| 1 | 0.96 | 1.04 | 0.75 | 0.96 | 0.20 | 0.03 |
| 5 | 0.96 | 1.04 | 0.75 | 0.96 | 0.20 | 0.03 |
| 10 | 0.96 | 1.04 | 0.75 | 0.96 | 0.20 | 0.03 |
| 25 | 0.96 | 1.04 | 0.75 | 0.96 | 0.20 | 0.03 |
| 50 | 0.96 | 1.04 | 0.75 | 0.96 | 0.20 | 0.03 |
| 100 | 0.96 | 1.04 | 0.75 | 0.96 | 0.20 | 0.03 |
| 200 | 0.96 | 1.04 | 0.75 | 0.96 | 0.20 | 0.03 |

*Notes*: Out-of-sample Sharpe ratios of $MV$ and $MV_{TC}$. Column 2 and 3 provide Sharpe ratios before transaction costs and columns 4 and 5 Sharpe ratios after costs. Column 6 shows the difference between the Sharpe ratio of $MV_{TC}$ and $MV$ and column 7 indicates the p-value of the difference. Standard errors for $\Delta$SR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). The rows report results for portfolios constructed with risk aversion coefficients $\lambda \in \{10, 25, 50, 100, 200\}$. The data are monthly returns from January 1976 to February 2016.

Table 20: **Sensitivity to the Risk Aversion Coefficient $\lambda$ (15 currencies): Notional Value, Mean and Volatility**

| Risk Aversion | Notional Value | | After TC Mean | | After TC Volatility | |
|---|---|---|---|---|---|---|
| Coefficient $\lambda$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| 1 | 154.45 | 125.06 | 3.71 | 3.89 | 4.93 | 4.06 |
| 5 | 30.89 | 25.00 | 0.74 | 0.78 | 0.99 | 0.81 |
| 10 | 15.44 | 12.50 | 0.37 | 0.39 | 0.49 | 0.41 |
| 25 | 6.18 | 5.00 | 0.15 | 0.16 | 0.20 | 0.16 |
| 50 | 3.09 | 2.50 | 0.07 | 0.08 | 0.10 | 0.08 |
| 100 | 1.54 | 1.25 | 0.04 | 0.04 | 0.05 | 0.04 |
| 200 | 0.77 | 0.62 | 0.02 | 0.02 | 0.02 | 0.02 |

*Notes*: Columns 2 and 3 report the average notional values of $MV$ and $MV_{TC}$, column 4 and 5 report after cost average returns and columns 6 and 7 after costs return volatilities. The rows report results for portfolios constructed with risk aversion coefficients $\lambda \in \{10, 25, 50, 100, 200\}$. The data are monthly returns from January 1976 to February 2016.

Table 21: **Sensitivity to the Risk Aversion Coefficient $\lambda$ (15 currencies): Crash Risk**

| Risk Aversion | After TC Skewness | | After TC Kurtosis | | After TC MDD | |
|---|---|---|---|---|---|---|
| Coefficient $\lambda$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| 1 | -0.12 | -0.08 | 16.64 | 13.36 | -17.98 | -12.08 |
| 5 | -0.12 | -0.08 | 16.64 | 13.36 | -3.60 | -2.42 |
| 10 | -0.12 | -0.08 | 16.64 | 13.36 | -1.80 | -1.21 |
| 25 | -0.12 | -0.08 | 16.64 | 13.36 | -0.72 | -0.48 |
| 50 | -0.12 | -0.08 | 16.64 | 13.36 | -0.36 | -0.24 |
| 100 | -0.12 | -0.08 | 16.64 | 13.37 | -0.18 | -0.12 |
| 200 | -0.12 | -0.07 | 16.64 | 13.38 | -0.09 | -0.06 |

*Notes*: Columns 2 and 3 report the after cost skewness of $MV$ and $MV_{TC}$, column 4 and 5 report after cost kurtosis and columns 6 and 7 after costs maximum draw downs. The rows report results for portfolios constructed with risk aversion coefficients $\lambda \in \{10, 25, 50, 100, 200\}$. The data are monthly returns from January 1976 to February 2016.

Table 22: $MV_{TC}$ **vs** $MV$**: NBER Recessions, 29 Currencies**

| | NBER Recessions | | non-NBER Recessions | |
|---|---|---|---|---|
| | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| **Before Transaction Costs:** | | | | |
| SR | 0.63 | 0.63 | 1.39 | 1.41 |
| Mean | 8.79 | 8.05 | 13.89 | 12.39 |
| ΔMean | - | -0.74 | - | -1.50 |
| **Transaction Costs:** | | | | |
| Mean Costs | 4.46 | 1.78 | 5.22 | 1.17 |
| ΔMean Costs | - | -2.67*** | - | -4.04*** |
| **After Transaction Costs:** | | | | |
| **SR** | **0.32** | **0.51** | **0.86** | **1.29** |
| **ΔSR** | **-** | **0.18** | **-** | **0.43*** ** |
| **(p-value)** | - | **(0.26)** | **-** | **(0.01)** |
| Mean | 4.33 | 6.26 | 8.67 | 11.22 |
| Vol | 13.43 | 12.38 | 10.09 | 8.71 |
| Skew | 1.45 | 1.25 | -0.42 | 0.21 |
| Kurt | 9.64 | 8.41 | 6.40 | 5.03 |
| Positive | 55.36 | 60.71 | 62.32 | 69.32 |
| MDD | -10.69 | -9.23 | -37.61 | -17.96 |
| AC | 0.15 | 0.10 | 0.09 | 0.15 |
| $CE_{\lambda=1}$ | - | 0.60 | - | 0.84 |
| $CE_{\lambda=5}$ | - | 2.22 | - | 3.15 |
| $CE_{\lambda=10}$ | - | 4.24 | - | 6.03 |
| $CE_{\lambda=50}$ | - | 20.44 | - | 29.12 |
| $TA$ | 1 | 0.46 | 1 | 0.40 |

*Notes*: Summary statistics of monthly excess returns of $MV$ and $MV_{TC}$ for all the 29 currencies for the sample 1976-2016. First two columns report results for the NBER recession periods, last two columns report results for non-recession periods. SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation, $CE_\lambda$ the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV$ to strategy $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 23: $MV_{TC}$ vs $MV$: **NBER Recessions, 15 Currencies**

| | NBER Recessions | | non-NBER Recessions | |
|---|---|---|---|---|
| | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| **Before Transaction Costs:** | | | | |
| SR | 0.03 | 0.14 | 1.23 | 1.29 |
| Mean | 0.44 | 1.91 | 10.33 | 9.36 |
| ΔMean | - | 1.47 | - | -0.97 |
| **Transaction Costs:** | | | | |
| Mean Costs | 3.03 | 1.49 | 2.08 | 0.77 |
| ΔMean Costs | - | -1.54*** | - | -1.31*** |
| **After Transaction Costs:** | | | | |
| **SR** | **-0.17** | **0.03** | **1.02** | **1.22** |
| **ΔSR** | **-** | **0.20*** | **-** | **0.21*** |
| **(p-value)** | - | **(0.10)** | **-** | **(0.06)** |
| Mean | -2.59 | 0.42 | 8.25 | 8.59 |
| Vol | 15.35 | 13.11 | 8.11 | 7.02 |
| Skew | -1.82 | -1.11 | 0.27 | 0.77 |
| Kurt | 12.09 | 9.58 | 6.67 | 7.69 |
| Positive | 48.21 | 53.57 | 64.73 | 67.63 |
| MDD | -39.17 | -25.37 | -16.99 | -13.89 |
| AC | -0.03 | -0.14 | 0.13 | 0.28 |
| $CE_{\lambda=1}$ | - | 1.21 | - | 0.24 |
| $CE_{\lambda=5}$ | - | 4.82 | - | 1.04 |
| $CE_{\lambda=10}$ | - | 9.35 | - | 2.04 |
| $CE_{\lambda=50}$ | - | 45.54 | - | 10.04 |
| $TA$ | 1 | 0.44 | 1 | 0.34 |

*Notes*: Summary statistics of monthly excess returns of $MV$ and $MV_{TC}$ for all the 15 developed currencies for the sample 1976-2016. First two columns report results for the NBER recession periods, last two columns report results for non-recession periods. SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation, $CE_\lambda$ the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV$ to strategy $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and autocorrelation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 24: $MV_{TC}$ **vs** $MV$**: Pre- vs Post-EURO, 29 Currencies**

| | Pre-EURO | | Post-EURO | |
|---|---|---|---|---|
| | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| **Before Transaction Costs:** | | | | |
| SR | 1.40 | 1.41 | 1.07 | 1.08 |
| Mean | 16.74 | 14.81 | 8.67 | 7.96 |
| ΔMean | - | -1.93 | - | -0.71 |
| **Transaction Costs:** | | | | |
| Mean Costs | 7.65 | 1.76 | 1.74 | 0.56 |
| ΔMean Costs | - | -5.89*** | - | -1.18*** |
| **After Transaction Costs:** | | | | |
| **SR** | **0.75** | **1.26** | **0.86** | **1.01** |
| **ΔSR** | **-** | **0.50**** | **-** | **0.15*** |
| **(p-value)** | - | **(0.01)** | **-** | **(0.06)** |
| Mean | 9.09 | 13.05 | 6.93 | 7.40 |
| Vol | 12.04 | 10.36 | 8.05 | 7.30 |
| Skew | -0.02 | 0.48 | -0.08 | -0.04 |
| Kurt | 6.93 | 6.14 | 5.32 | 4.64 |
| Positive | 63.40 | 71.70 | 59.02 | 63.90 |
| MDD | -37.61 | -17.96 | -18.19 | -18.90 |
| AC | 0.13 | 0.15 | 0.04 | 0.11 |
| $CE_{\lambda=1}$ | - | 0.22 | - | 1.27 |
| $CE_{\lambda=5}$ | - | 0.90 | - | 4.69 |
| $CE_{\lambda=10}$ | - | 1.76 | - | 8.96 |
| $CE_{\lambda=50}$ | - | 8.60 | - | 43.16 |
| $TA$ | 1 | 0.39 | 1 | 0.42 |

*Notes*: Summary statistics of monthly excess returns of $MV$ and $MV_{TC}$ for all the 29 currencies for the sample 1976-2016. First two columns report results for the pre-Euro period (1976-1999), the last two columns report results for the post-Euro period (1999-2016). SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation, $CE_\lambda$ the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV$ to strategy $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 25: $MV_{TC}$ **vs** $MV$**: Pre- vs Post-EURO, 15 Currencies**

| | Pre-EURO | | Post-EURO | |
|---|---|---|---|---|
| | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| **Before Transaction Costs:** | | | | |
| SR | 1.10 | 1.20 | 0.85 | 0.88 |
| Mean | 12.95 | 12.08 | 4.11 | 3.68 |
| ΔMean | - | -0.88 | - | -0.44 |
| **Transaction Costs:** | | | | |
| Mean Costs | 3.40 | 1.30 | 0.57 | 0.25 |
| ΔMean Costs | - | -2.10*** | - | -0.32*** |
| **After Transaction Costs:** | | | | |
| **SR** | **0.83** | **1.11** | **0.74** | **0.83** |
| **ΔSR** | **-** | **0.26**** | **-** | **0.09** |
| **(p-value)** | - | **(0.03)** | **-** | **(0.15)** |
| Mean | 9.56 | 10.78 | 3.55 | 3.43 |
| Vol | 11.49 | 9.83 | 4.77 | 4.13 |
| Skew | -1.12 | -0.53 | 0.77 | 0.81 |
| Kurt | 11.46 | 9.79 | 6.75 | 6.88 |
| Positive | 66.42 | 71.32 | 58.05 | 59.02 |
| MDD | -35.97 | -24.16 | -14.66 | -12.81 |
| AC | 0.08 | 0.15 | 0.12 | 0.17 |
| $CE_{\lambda=1}$ | - | 0.04 | - | 0.59 |
| $CE_{\lambda=5}$ | - | 0.27 | - | 2.43 |
| $CE_{\lambda=10}$ | - | 0.55 | - | 4.73 |
| $CE_{\lambda=50}$ | - | 2.77 | - | 26.16 |
| $TA$ | 1 | 0.30 | 1 | 0.40 |

*Notes*: Summary statistics of monthly excess returns of $MV$ and $MV_{TC}$ for all the 15 developed currencies for the sample 1976-2016. First two columns report results for the pre-Euro period (1976-1999), the last two columns report results for the post-Euro period (1999-2016). SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation, $CE_\lambda$ the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV$ to strategy $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 26: $MV_{TC}$ **vs** $MV$**: 1976-2016 vs 1983-2016, 29 Currencies**

| | 1976-2016 | | 1983-2016 | |
|---|---|---|---|---|
| | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| **Before Transaction Costs:** | | | | |
| SR | 1.26 | 1.27 | 1.17 | 1.20 |
| Mean | 13.30 | 11.89 | 10.85 | 10.03 |
| $\Delta$Mean | - | -1.41 | - | -0.82 |
| **Transaction Costs:** | | | | |
| Mean Costs | 5.13 | 1.25 | 2.76 | 0.86 |
| $\Delta$Mean Costs | - | -3.88*** | - | -1.90*** |
| **After Transaction Costs:** | | | | |
| **SR** | **0.78** | **1.16** | **0.88** | **1.11** |
| **$\Delta$SR** | **-** | **0.38*** ** | **-** | **0.23*** ** |
| **(p-value)** | - | **(0.01)** | **-** | **(0.00)** |
| Mean | 8.17 | 10.64 | 8.09 | 9.17 |
| Vol | 10.52 | 9.21 | 9.20 | 8.28 |
| Skew | 0.00 | 0.47 | -0.22 | -0.06 |
| Kurt | 7.60 | 6.59 | 5.33 | 4.82 |
| Positive | 61.49 | 68.30 | 61.70 | 67.35 |
| MDD | -37.61 | -18.90 | -19.17 | -18.90 |
| AC | 0.11 | 0.15 | 0.06 | 0.16 |
| $CE_{\lambda=1}$ | - | 0.88 | - | 0.33 |
| $CE_{\lambda=5}$ | - | 3.04 | - | 1.19 |
| $CE_{\lambda=10}$ | - | 5.82 | - | 2.28 |
| $CE_{\lambda=50}$ | - | 28.09 | - | 10.96 |
| $TA$ | 1 | 0.41 | 1 | 0.41 |

*Notes*: Summary statistics of monthly excess returns of $MV$ and $MV_{TC}$ for all the 29 currencies for the full sample 1976-2016 versus the sample starting in 1983. First two columns report results for the pre-Euro period (1976-1999), the last two columns report results for the post-Euro period (1999-2016). SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation, $CE_\lambda$ the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV$ to strategy $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). $\Delta$Mean, $\Delta$Mean Costs, $\Delta$SR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. Standard errors of $\Delta$SR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of $\Delta$Mean Costs are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of $\Delta$SR and $\Delta$Mean Costs. We only report the p-value for $\Delta$SR after costs.

Table 27: $MV_{TC}$ **vs** $MV$**: 1976-2016 vs 1983-2016, 15 Currencies**

| | 1976-2016 | | 1983-2016 | |
|---|---|---|---|---|
| | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| **Before Transaction Costs:** | | | | |
| SR | 0.96 | 1.04 | 1.04 | 1.04 |
| Mean | 9.18 | 8.50 | 7.13 | 6.42 |
| ΔMean | - | -0.69 | - | -0.71 |
| **Transaction Costs:** | | | | |
| Mean Costs | 2.19 | 0.85 | 1.19 | 0.44 |
| ΔMean Costs | - | -1.34*** | - | -0.75*** |
| **After Transaction Costs:** | | | | |
| **SR** | **0.75** | **0.96** | **0.88** | **0.98** |
| **ΔSR** | **-** | **0.20**** | **-** | **0.10**** |
| **(p-value)** | - | **(0.03)** | **-** | **(0.03)** |
| Mean | 6.99 | 7.64 | 5.94 | 5.99 |
| Vol | 9.27 | 7.98 | 6.73 | 6.12 |
| Skew | -1.00 | -0.28 | 0.32 | 0.15 |
| Kurt | 15.31 | 12.94 | 6.02 | 6.60 |
| Positive | 62.77 | 65.96 | 61.70 | 64.52 |
| MDD | -35.97 | -24.16 | -18.54 | -13.89 |
| AC | 0.10 | 0.17 | 0.21 | 0.22 |
| $CE_{\lambda=1}$ | - | 0.35 | - | 0.08 |
| $CE_{\lambda=5}$ | - | 1.49 | - | 0.38 |
| $CE_{\lambda=10}$ | - | 2.91 | - | 0.75 |
| $CE_{\lambda=50}$ | - | 14.27 | - | 3.74 |
| $TA$ | 1 | 0.35 | 1 | 0.33 |

*Notes*: Summary statistics of monthly excess returns of $MV$ and $MV_{TC}$ for all the 15 developed currencies for the full sample 1976-2016 versus the sample starting in 1983. First two columns report results for the pre-Euro period (1976-1999), the last two columns report results for the post-Euro period (1999-2016). SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation, $CE_\lambda$ the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV$ to strategy $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 28: **Mean-Variance Strategies with Full Round-Trip Costs: $MV_{TC}$ vs $MV$**

| | All 29 Currencies | | 15 Developed Currencies | |
|---|---|---|---|---|
| | $MV$ | $MV_{TC}$ | $MV$ | $MV_{TC}$ |
| **Before Transaction Costs:** | | | | |
| SR | 1.26 | 1.14 | 0.96 | 1.00 |
| Mean | 13.30 | 8.54 | 9.18 | 6.59 |
| ΔMean | - | -4.75 | - | -2.59 |
| **Transaction Costs:** | | | | |
| Mean Costs | 13.06 | 2.77 | 5.61 | 2.12 |
| ΔMean Costs | - | -10.29*** | - | -3.49*** |
| **After Transaction Costs:** | | | | |
| **SR** | **0.02** | **0.80** | **0.40** | **0.74** |
| **ΔSR** | **-** | **0.78***** | **-** | **0.35**** |
| **(p-value)** | - | **(0.00)** | **-** | **(0.01)** |
| Mean | 0.24 | 5.78 | 3.58 | 4.47 |
| Vol | 10.98 | 7.19 | 9.03 | 6.03 |
| Skew | -0.70 | 0.43 | -1.92 | -1.01 |
| Kurt | 7.34 | 8.79 | 19.35 | 18.05 |
| Positive | 53.62 | 62.55 | 57.87 | 64.47 |
| MDD | -147.35 | -17.24 | -41.84 | -23.67 |
| AC | 0.18 | 0.08 | 0.05 | 0.13 |
| $CE_{\lambda=1}$ | - | 1.55 | - | 0.68 |
| $CE_{\lambda=5}$ | - | 5.49 | - | 3.03 |
| $CE_{\lambda=10}$ | - | 10.40 | - | 5.97 |
| $CE_{\lambda=50}$ | - | 49.79 | - | 29.49 |
| $TA$ | 1 | 0.43 | 1 | 0.40 |

*Notes*: Summary statistics of monthly excess returns of $MV$ and $MV_{TC}$ using round-trip costs. First two columns report results for all 29 currencies, last two columns for 15 developed currencies. The sample period is 1976-2016. SR is the annualized Sharpe ratio, Mean the annualized average return (in percentage points), Mean Costs the average annualized transaction costs measured in percentage of the portfolio value, Vol the annualized standard deviation (in percentage points), Skew the skewness, Kurt the kurtosis, % Positive the percentage of positive monthly returns, MDD the Maximum Draw Down, AC the autocorrelation, $CE_\lambda$ the annualized rate of return (Certainty Equivalent) an investor with mean-variance preferences and risk aversion $\lambda$ is willing to give up in order to switch from strategy $MV$ to strategy $MV_{TC}$. $TA$ measures the average of trade aggressiveness defined in equation (1). ΔMean, ΔMean Costs, ΔSR are the differences in the Mean, Mean Costs, SR between $MV_{TC}$ and $MV$. Standard errors of ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). Standard errors of ΔMean Costs are estimated using Newey and West (1987) to account for heteroskedasticity and auto-correlation. ***, **, * indicate a statistical significance at the 1%, 5%, 10% level of ΔSR and ΔMean Costs. We only report the p-value for ΔSR after costs.

Table 29: **Optimized vs Equally Weighted Portfolios with Full Round-Trip Costs, Monthly Frequency**

| | All 29 Currencies | | | 15 Developed Currencies | | |
|---|---|---|---|---|---|---|
| | Before TC | After TC | | Before TC | After TC | |
| Strategies | SR | SR | ΔSR | SR | SR | ΔSR |
| $\mathbf{MV_{TC}}$ | **1.14** | **0.80** | - | **1.00** | **0.74** | - |
| $MV_{TC\setminus Corr}$ | 1.20 | 0.70 | 0.11 | 0.91 | 0.51 | $0.24^{***}$ |
| | | | (0.17) | | | (0.00) |
| $MV$ | 1.26 | 0.02 | $0.78^{***}$ | 0.96 | 0.40 | $0.35^{***}$ |
| | | | (0.00) | | | (0.01) |
| $MV_{Net}$ | 1.17 | 0.36 | $0.44^{***}$ | 0.91 | 0.38 | $0.36^{***}$ |
| | | | (0.00) | | | (0.00) |
| $DOL$ | 0.10 | -0.14 | $0.95^{***}$ | 0.07 | -0.10 | $0.84^{***}$ |
| | | | (0.00) | | | (0.00) |
| $DDOL$ | 0.49 | 0.24 | $0.56^{***}$ | 0.59 | 0.42 | $0.32^{**}$ |
| | | | (0.01) | | | (0.01) |
| $HML$ | 0.74 | 0.06 | $0.74^{***}$ | 0.64 | 0.30 | $0.44^{*}$ |
| | | | (0.00) | | | (0.07) |
| $MOM$ | 0.37 | 0.12 | $0.68^{***}$ | 0.31 | 0.10 | $0.64^{**}$ |
| | | | (0.00) | | | (0.02) |
| $VAL$ | 0.47 | -0.39 | $1.20^{***}$ | 0.53 | 0.17 | $0.57^{**}$ |
| | | | (0.00) | | | (0.04) |

*Notes*: Columns 2 and 5 report Sharpe ratios before costs (Before TC SR). Columns 3 and 6 report Sharpe ratios after costs (After TC SR). Columns 4 and 7 report the difference between the Sharpe ratios after costs of $MV_{TC}$ and the strategy in the corresponding row (After TC ΔSR). $MV_{TC}$ is the mean-variance optimized portfolio which optimizes over transaction costs. $MV_{TC\setminus Corr}$ is the mean-variance optimized portfolio which optimizes over transaction costs but makes the simplifying assumption that assets are uncorrelated. $MV$ is the mean-variance optimized portfolio without taking into account transaction costs in the optimization. $DOL$ borrows in the USD and equally invests in all other currencies. $DDOL$ takes a long position in $DOL$ if the median exchange rate forward discount is positive, and a short position otherwise. $HML$ sorts currencies according to the forward discount into quintiles and borrows in the bottom and invests in the top quintile. $MOM$ sorts currencies according to their past 12 month performance into quintiles and borrows in the bottom and invests in the top quintile. $VAL$ sorts currencies according to the power purchase parity adjusted exchange rate into quintiles and borrows in the top quintile and invests in the bottom quintile. The data are monthly returns for our full set of 29 currencies (columns 2 to 4) and a subsample of 15 developed currencies (columns 5 to 7) from January 1976 to February 2016. Standard errors for ΔSR are estimated using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008).

Table 30: Datastream mnemonics for currency quotes against the U.S. dollar

| Currency | Spot rate | Forward rate | Quote convention |
|---|---|---|---|
| Australian dollar | BBAUDSP | BBAUD1F | FCU/USD |
| Belgian franc | BELGLU$ | USBEF1F | FCU/USD |
| Brazilian real | BRACRU$ | USBRL1F | FCU/USD |
| British pound | BBGBPSP | BBGBP1F | USD/FCU |
| Canadian dollar | BBCADSP | BBCAD1F | FCU/USD |
| Czech koruna | CZECHC$ | USCZK1F | FCU/USD |
| Danish krone | BBDKKSP | BBDKK1F | FCU/USD |
| Euro | BBEURSP | BBEUR1F | FCU/USD |
| French franc | BBFRFSP | BBFRF1F | FCU/USD |
| German mark | BBDEMSP | BBDEM1F | FCU/USD |
| Greek Drachma | GREDRA$ | USGRD1F | FCU/USD |
| Hungarian forint | HUNFOR$ | USHUF1F | FCU/USD |
| Icelandic krona | ICEKRO$ | USISK1F | FCU/USD |
| Irish punt | BBIEPSP | BBIEP1F | USD/FCU |
| Italian lira | BBITLSP | BBITL1F | FCU/USD |
| Japanese yen | BBJPYSP | BBJPY1F | FCU/USD |
| Mexican peso | MEXPES$ | USMXN1F | FCU/USD |
| Netherland guilder | BBNLGSP | BBNLG1F | FCU/USD |
| New Zealand dollar | BBNZDSP | BBNZD1F | FCU/USD |
| Norwegian krone | BBNOKSP | BBNOK1F | FCU/USD |
| Polish zloty | POLZLO$ | USPLN1F | FCU/USD |
| Portuguese escudo | PORTES$ | USPTE1F | FCU/USD |
| Singapore dollar | BBSGDSP | BBSGD1F | FCU/USD |
| South Africa rand | BBZARSP | BBZAR1F | FCU/USD |
| South Korean won | KORSWO$ | USKRW1F | FCU/USD |
| Spanish peseta | SPANPE$ | USESP1F | FCU/USD |
| Swedish krona | BBSEKSP | BBSEK1F | FCU/USD |
| Swiss france | BBCHFSP | BBCHF1F | FCU/USD |
| Taiwan new dollar | TAIWDO$ | USTWD1F | FCU/USD |

Table 31: Datastream mnemonics for currency quotes against the British pound

| Currency | Spot rate | Forward rate | Quote convention |
|---|---|---|---|
| Belgian franc | BELGLUX | BELXF1F | FCU/GBP |
| Canadian dollar | CNDOLLR | CNDOL1F | FCU/GBP |
| Danish krone | DANISHK | DANIS1F | FCU/GBP |
| French franc | FRENFRA | FRENF1F | FCU/GBP |
| German mark | DMARKER | DMARK1F | FCU/GBP |
| Irish punt | IPUNTER | IPUNT1F | FCU/GBP |
| Italian lira | ITALIRE | ITALY1F | FCU/GBP |
| Japanese yen | JAPAYEN | JAPYN1F | FCU/GBP |
| Netherlands guilder | GUILDER | GUILD1F | FCU/GBP |
| Norwegian krone | NORKRON | NORKN1F | FCU/GBP |
| Portuguese escudo | PORTESC | PORTS1F | FCU/GBP |
| Spanish peseta | SPANPES | SPANP1F | FCU/GBP |
| Swedish krona | SWEKRON | SWEDK1F | FCU/GBP |
| Swiss franc | SWISSFR | SWISF1F | FCU/GBP |
| U.S. dollar | USDOLLR | USDOL1F | FCU/GBP |

## Heuristic Adjustments to Approximate a Dynamic $MV_{TC}$

**29 Currencies:**

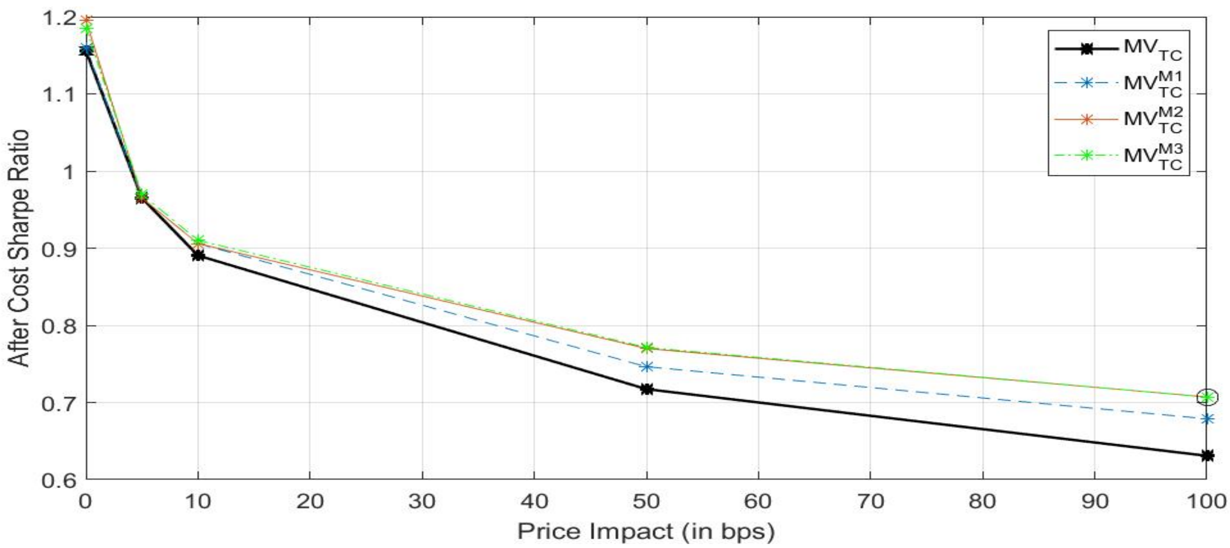


**15 Currencies:**

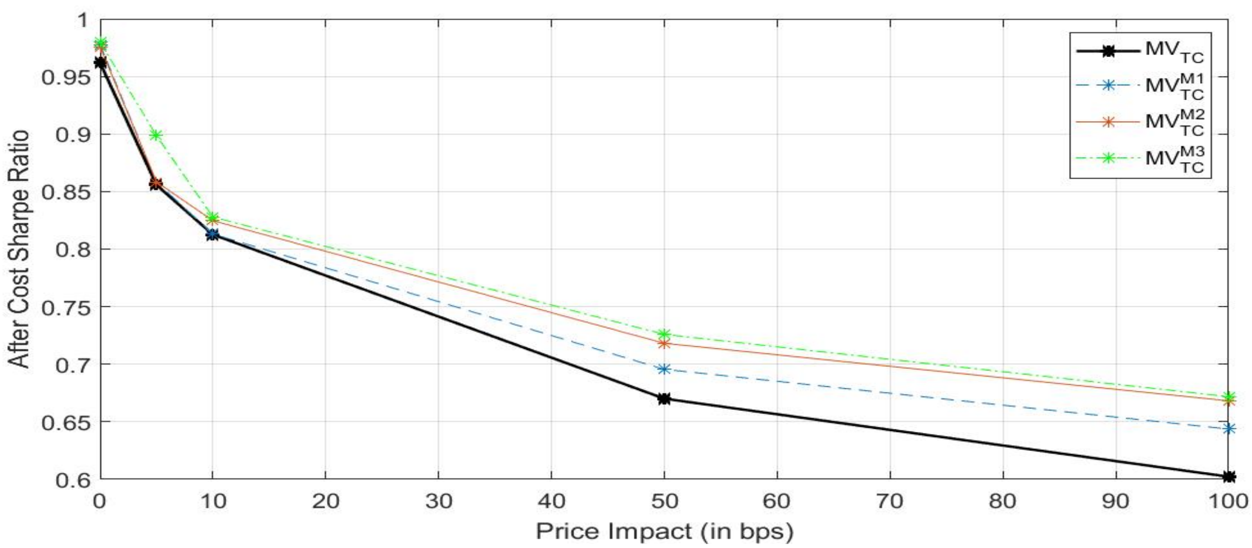


Figure 12: The plot shows the annualized after cost Sharpe ratios of four trading strategies as a function of the price impact parameter $\pi \in [0, 100]$, which is measured in basis points. The price impact is linear and constant across currencies $i$ and time $t$, i.e., $\pi = \pi^z_{i,t}$ $\forall i, t$ and $z \in \{P+, P-, S+, S-\}$, where $\pi^z_{i,t}$ are the diagonal elements of $\mathbf{\Pi_t}^{z,L}$. The results for $MV_{TC}$ are indicated by the black solid line. The other lines show the results for three different parameterizations of $MV^M_{TC}$, which optimize over the adjusted costs described in equations (2) and (3). The three cases are $c_1 \in \mathbb{R}, c_2 = a_1 = a_2 = 0$ (blue dashed line), $(c_1, c_2) \in \mathbb{R}^2, a_1 = a_2 = 0$ (orange solid line), and $(c_1, c_2, a_1, a_2) \in \mathbb{R}^4$ (green dashed-dotted line). If a data-point is circled its Sharpe ratio is statistically different from that of $MV_{TC}$ at the 10% level according to the (Ledoit and Wolf, 2008) Sharpe ratio test performed using block bootstrapping with block sizes of 5 observations to account for heteroskedasticity, cross- and auto-correlation. The data are monthly returns for our full set of 29 currencies from January 1976 to February 2016.

## Mean-Variance Problem with Directional Transaction Costs: The Case of 2 Risky Assets

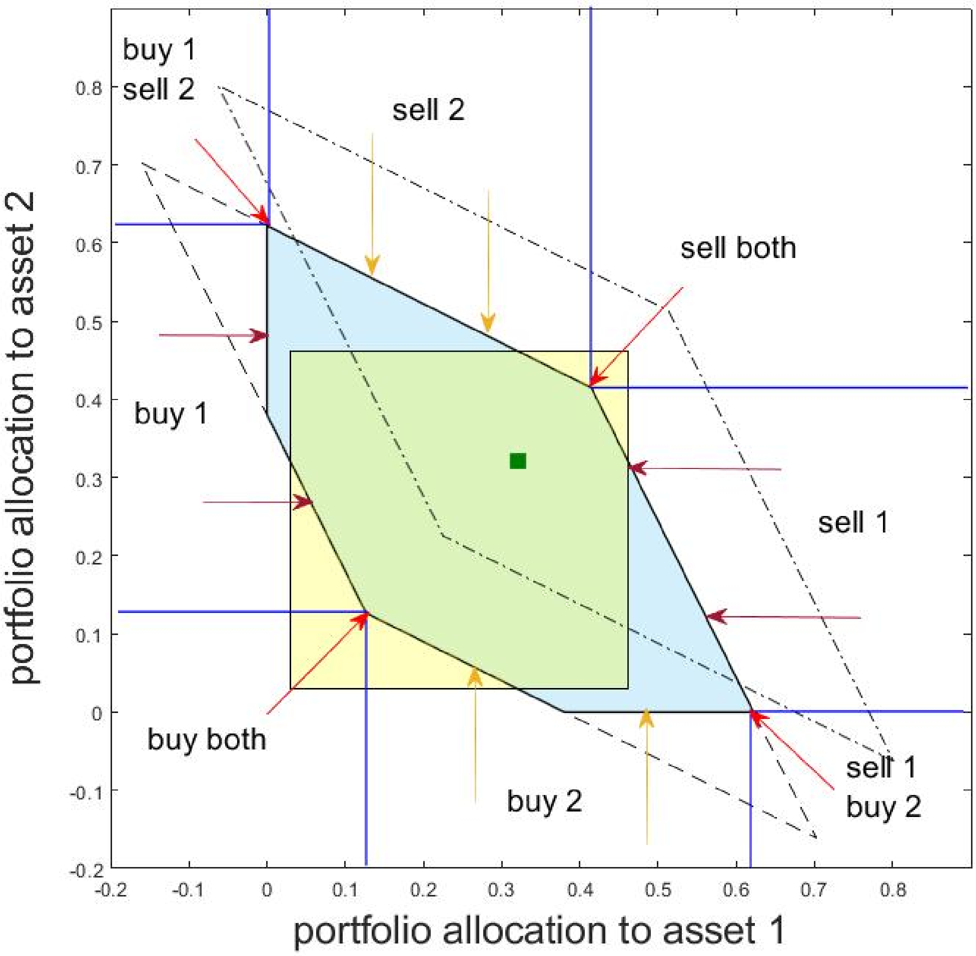


Figure 13: No trading regions for setting of two risky assets with $\mu_{\mathbf{t}}^{\mathbf{e}} = 2.4\%$, $\rho = 0.5$, $\sigma_{\mathbf{t}} = 10\%$ (the diagonal elements of $\mathbf{V_t}$), $\mathbf{C_t^{P+}} = 1.45\%$, $\mathbf{C_t^{S+}} = 0.71\%$, $\mathbf{C_t^{P-}} = 0.71\%$ , $\mathbf{C_t^{S-}} = 1.45\%$, and $\lambda = 5$. The no trading region for $MV_{TC}$ is the blue polygon, the one for $MV_{TC\backslash Corr}$ is the yellow square, and the standard mean-variance optimum $\theta_{\mathbf{t}}^{\mathbf{MV}}$ (represented by the green square) lies inside both. If there were only two costs, $\mathbf{C_t^{P-}}$ and $\mathbf{C_t^{S-}}$, the no trading region for $MV_{TC}$ would be the dash-dotted parallelogram. In contrast, if the two costs were $\mathbf{C_t^{S+}}$ and $\mathbf{C_t^{P+}}$, then the no trading region for $MV_{TC}$ would be the dashed parallelogram. The arrows indicate the optimal trades $\mathbf{\Delta_t^{P+}}$, $\mathbf{\Delta_t^{S+}}$, $\mathbf{\Delta_t^{P-}}$, $\mathbf{\Delta_t^{S-}}$ according to the $MV_{TC}$ strategy from any initial position $\theta_{\mathbf{t}}^{\mathbf{0}}$ outside the blue no trading region. Purple, brown or orange colors of the arrows indicate that only asset 1, 2 or both assets are traded.