

Importance of Transaction Costs for Asset Allocation in Foreign Exchange Markets*

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

January 31, 2024

Abstract

Transaction costs have a first-order effect on the performance of currency portfolios. When a fund is large, costs due to the price impact of trading quickly erode returns, leaving many popular strategies unprofitable. A mean-variance approach that accounts for transaction costs in the optimization (MVTC) efficiently tackles the problem with only relatively minor negative implications on before-cost profitability. MVTC is robust even when the price impact of trading is severe. In contrast, popular rules-of-thumb to mitigate costs are inefficient as there are unintended, adverse effects on before-cost performance that dominate the intended cost savings. Finally, we introduce an accurate extrapolation approach to expand the sample of the realized Amihud measure of [Ranaldo and Santucci de Magistris \(2022\)](#) from 12 to 26 currencies and from 2012 back in time to 1986.

JEL-Classification: F31, G11, G15.

Keywords: Transaction Costs, Price Impact, realized Amihud, Mean-Variance, Optimization, Asset Allocation, Foreign Exchange, Carry Trade.

*We are thankful to editor Nick Roussanov and one anonymous referee, as well as to Pasquale Della Corte, Jason Donaldson, Bernard Dumas, Phil Dybvig, Radha Gopalan, Nishad Kapadia, Mina Lee, Hong Liu, Dong Lou, Paolo Santucci de Magistris, Asaf Manela, Michael Melvin, George Panayotov, Ivelina Pavlova, Angelo Ranaldo, Alessio Saretto, Fabricius Somogyi, Ngoc-Khanh Tran, Adrien Verdelhan, Michael Weber, Di Woody Wu, Guofu Zhou, seminar participants at Tulane U., U. New Orleans, Washington U., and conference participants at the American Finance Association, Eastern Finance Association, Hong Kong Joint Finance Research Workshop, Multinational Finance Society, Southwestern Finance Association, Tsinghua Finance Workshop for many helpful comments and suggestions.

[†]Olin Business School, Washington University in St. Louis, IliasFilippou@wustl.edu.

[‡]The University of Hong Kong, Pokfulam Road, Hong Kong, +852 5236 2121, maurer@hku.hk.

[§]University of New Orleans, lpezzo@uno.edu.

[¶]Olin Business School, Washington University in St. Louis, and Centre for Economic Policy Research, mark.p.taylor@wustl.edu.

1 Introduction

We analyze the importance of transaction costs to investors in foreign exchange (FX) markets. The contribution of our analysis is threefold. We first document that transaction costs have first-order negative implications on the profitability of currency trading strategies despite the relatively high trading volume and liquidity in FX markets. Most importantly, while proportional costs are relatively small, the performance of many popular strategies is sensitive to the price impact of trading. For some strategies (such as the currency carry trade, momentum, and mean-variance optimized currency portfolios) costs due to the price impact quickly erode returns, rendering the strategies unprofitable.

For instance, the monthly rebalancing of a simple carry trade strategy implies a relatively large monthly turnover of 56% of its notional value. When a fund's initial assets under management (AUM) reaches USD 100 million or even 1 billion, then trade orders on average account for 3.3% or even 10.5% of the daily trading volume. This implies that sell and buy orders of such a fund significantly move bid and ask quotes, and in turn, costs due to the price impact are large. In the case of the carry trade strategy we find that the Sharpe ratio is 0.72 before accounting for costs, while for a fund with an initial AUM of USD 490 million (or bigger), the costs due to the price impact dominate the average returns and the strategy is unprofitable.

It is an important contribution to document the sensitivity of currency strategies to the price impact of trading, as the previous literature typically only shows the robustness of strategies with respect to proportional costs. In addition, our results provide a natural motivation to explain illiquidity premia in FX markets as documented by, for instance, [Mancini et al. \(2013\)](#) and [Soderlind and Somogyi \(2023\)](#). That is, investors care about FX liquidity and dislike to be exposed to liquidity risk since the price impact of trading has a first-order quantitative negative effect on the profitability of currency strategies.

Second, we provide a mean-variance-transaction-cost optimization approach (*MVTC*) that efficiently mitigates trading costs with only relatively minor negative effects on the before-cost optimality. MVTC builds on the mean-variance optimized currency portfolio of [Maurer et al. \(2023\)](#) and uses the framework of [Dybvig and Pezzo \(2020\)](#) to tackle costs.

The optimization ensures that the portfolio stays “close enough” to the optimal before-cost balance between expected return and risk, while reducing trading activity (especially in less liquid and more expensive currencies) and thus minimizing costs. Our cost-optimized portfolio performs well out-of-sample. It is robust to a price impact of trading and remains profitable even under severe conditions. Specifically, even when the initial AUM of a fund reaches USD 1 billion the after-cost Sharpe ratio is still 0.72 (which is on par with the before-cost Sharpe ratio of the simple carry trade). Moreover, its capacity is large, i.e., its after-cost average returns are positive even when the AUM of a fund exceeds USD 1 trillion.

In contrast, common rules-of-thumb to tackle transaction costs are inefficient.¹ They lead to a significant deterioration in the before-cost profitability, and this unintended adverse effect outweighs 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 *MVTC*, a fully optimized portfolio that accounts for costs in the optimization. This is an interesting contribution as the empirical finance literature as well as practitioners often employ rules-of-thumb. This is particularly important for the literature that investigates the after-cost profitability of (anomaly) portfolios. These results may be affected by the use of an (in)efficient approach to tackle trading costs. In addition, our findings provide important guidance for practitioners.

Intuitively, we can understand our findings using an analogy from basic portfolio diversification. Idiosyncratic risks can be a first-order component of an asset’s volatility. Hence, it is important for investors to tackle idiosyncratic risks and construct a diversified portfolio. Once idiosyncratic risks are diversified away, they have no (or only a negligible) effect on the performance. Similarly, we show that transaction costs (and specifically the price impact of trading) can significantly affect the profitability of currency trading strategies. Accordingly, it is important for investors to optimally tackle transaction costs in the portfolio construction. Once investors adopt efficient cost-mitigation techniques (following *MVTC*), trading costs have a negligible effect on performance.

¹ We investigate two rules-of-thumb to mitigate trading costs: (i) trade at a low frequency, and (ii) remove assets with relatively high transaction costs from the set of admissible assets. Moreover, the construction of characteristic sorted long-short strategies can also be viewed as a rule-of-thumb to tackle trading costs (and estimation errors).

Third, we contribute to the FX market literature on illiquidity by expanding the sample of the realized Amihud illiquidity measure of [Ranaldo and Santucci de Magistris \(2022\)](#) to a large cross-section of currencies against the USD (26 instead of 12) and a long time-series (starting in 1986 instead of 2012).² [Ranaldo and Santucci de Magistris \(2022\)](#) and [Lacava et al. \(2023\)](#) show that this measure is accurate and robust, and as such arguably more suitable for our analysis than alternatives in the literature.³ The daily realized Amihud measure is the ratio of the daily exchange rate growth variation and the trading volume. Olsen intraday exchange rate quotes are available for a large panel of currencies, and it is straightforward to estimate daily exchange rate growth variations. In contrast, volume data from CLS is only available for a small panel, i.e., for 12 currencies against the USD and starting in 2012. We propose novel extrapolations of the trading volume data to expand this panel. Specifically, we show that trading volume can be accurately extrapolated in the cross-section and in the time-series using the proportional bid-ask spread as a conditioning variable.⁴ The use of the bid-ask spread as a conditioning variable is consistent with the finding of [Karnaukh et al. \(2015\)](#) that the bid-ask spread is a useful proxy for liquidity in Electronic Broking Services (EBS) data.

Our paper is related to a growing literature on FX market illiquidity that uses CLS volume and Olsen intraday price data. [Fischer and Ranaldo \(2011\)](#) show that FX trading volume increases after FOMC meetings. [Ranaldo and Somogyi \(2021\)](#) estimate a VAR model similar to the stock market model of [Hasbrouck \(1991a,b\)](#) and find substantial information asymmetry across participants, time and currency pairs, and they estimate a sizeable asymmetric information risk premium. [Hasbrouck and Levich \(2021\)](#) document that more central agents trade at better conditions (centrality premium), and provide evidence that this is a compensation for providing liquidity in the market. [Hasbrouck and Levich \(2019\)](#) compute

²We make our extrapolations available on our websites.

³A possible alternative measure is the VAR approach of [Ranaldo and Somogyi \(2021\)](#). This measure has many advantages such as the possibility to differentiate the trading impact of diverse traders on market prices, i.e., the main feature exploited by the authors. However, the realized Amihud measure of [Ranaldo and Santucci de Magistris \(2022\)](#) is more precise and it is more straightforward to obtain a time-series of daily price impact measures than for VAR estimates. This makes the realized Amihud measure better suitable for our purposes.

⁴In the main text of the paper in section 2.4.1 we only show the extrapolation in the cross-section. Section A.2 in the Internet Appendix provides details about the extrapolation in the panel (cross-section and time-series) using data from CLS or the BIS Triennial Surveys.

the [Amihud \(2002\)](#) illiquidity measure and show that it is closely related to price impact estimates based on high-frequency EBS data. [Rinaldo and Santucci de Magistris \(2022\)](#) introduce the *realized* Amihud illiquidity measure – a refinement of the [Amihud \(2002\)](#) measure. They find that this novel illiquidity measure is decreasing in market depth, and increasing in bid-ask spreads, money market stress, uncertainty, risk aversion, and price inefficiencies. [Lacava et al. \(2023\)](#) further provide an in-depth theoretical discussion of the realized Amihud measure, and demonstrate its accuracy and robustness. [Cespa et al. \(2022\)](#) examine the information content of FX volume and find a stronger daily return reversal for currency pairs with abnormally low volume. [Somogyi \(2022\)](#) shows that trading against the USD is generally cheaper than other bilateral trades and argues that part of the USD dominance in terms of order volume can be attributed to the cost advantage of trading against the USD. [Roussanov and Wang \(2023\)](#) document that FX dealers trade without delay on monetary policy announcements. In contrast, funds and non-bank financials are slower and only react 3 to 5 days later. This evidence provides a channel for the delayed adjustment of exchange rates to monetary news. Finally, using CLS data on FX forwards [Bräuer and Hau \(2023\)](#) show that FX hedging demands have a significant impact on exchange rates.

Another strand of the literature uses datasets other than the volume data from CLS. [Evans and Lyons \(2002, 2008\)](#) use interdealer order flow data from Reuters Dealing 2000-1 and show that order flow is a strong determinant of exchange rates, and even more so after the release of macro news. [Froot and Ramadorai \(2005\)](#) estimate a VAR model to disentangle intrinsic-value from expected return shocks, and show that institutional-investor FX flows only correlate with expected return shocks. [Hau et al. \(2010\)](#) document that a change in the MSCI Global Equity Index has triggered large uninformed order flows in FX markets, which in turn has a significant price impact on exchange rates. Using EBS data, [Karnaukh et al. \(2015\)](#) show that FX market illiquidity correlates with funding constraints and global risk. [Mancini et al. \(2013\)](#) (using EBS data) and [Soderlind and Somogyi \(2023\)](#) (using estimates of effective bid-ask spreads) document that illiquidity in FX markets is priced in the cross-section of returns. [Breedon et al. \(2016\)](#) use data from Reuters and EBS and show that order flows have a significant effect on FX risk premia. Using disaggregated Norwegian central bank data [Evans and Rime \(2016\)](#) show that order flows predict exchange

rate fluctuations over relatively long horizons, and most of the predictability is attributed to the time-variation in the risk premium. [Menkhoff et al. \(2016\)](#) use end-user data from a top FX dealer and document that orders of demand-side investment managers forecast currency returns, while corporates and individual investors appear uninformed. [Krohn and Sushko \(2022\)](#) document a tight link between FX spot and swap market liquidity, and further show that banks' funding liquidity has a significant impact on FX market liquidity. [Hau et al. \(2021\)](#) use EU regulatory data and find that there is price discrimination across clients in forward FX markets, i.e., dealers charge significantly smaller spreads to more sophisticated clients. [Camanho et al. \(2022\)](#) use fund-level holdings data and show that funds rebalance their international allocation in response to realized returns with the aim to maintain their original risk exposure. In turn, the flows due to rebalancing have a significant impact on exchange rates. [Korsaye et al. \(2023\)](#) show that trading frictions (proxied by bid-ask spreads) are important in constructing robust global factors that correctly price international asset returns. [Bechtel et al. \(2023\)](#) propose and test a new channel that connects liquidity risk and interest rates in short-term funding markets. Finally, [Krohn et al. \(2023\)](#) show that the USD tends to appreciate in the run-up to FX fixes and depreciates afterwards, leading to a W-shaped intraday return pattern.

While this literature makes an effort to understand liquidity in FX markets,⁵ we are agnostic about the determinants of liquidity. We use the realized Amihud measure of [Rinaldo and Santucci de Magistris \(2022\)](#) for its simplicity and robustness to quantify the price impact of trading, and then analyze its implications on the profitability of popular FX trading strategies.

Finally, our paper is related to the literature on portfolio optimization in the presence of transaction costs. For brevity, we refer to [Dybvig and Pezzo \(2020\)](#) for a review of this literature.

The rest of the paper is structured as follows. Section 2 provides details about the data and the investment opportunity set in FX markets, including the specification of transaction costs. In section 3, we briefly describe seven popular trading strategies from the literature,

⁵That is, quantifying the impact of order flows on exchange rates, measuring illiquidity, understanding the heterogeneity of market participants, risk sharing, and the role of dealers, etc.

and introduce our novel mean-variance-transaction-cost optimized strategy *MVTC*. Section 4 provides the main results, and documents that the price impact has first-order effects on the profitability of most trading strategies. Section 5 shows that rules-of-thumb to mitigate costs are not effective, and *MVTC* is the dominant strategy in tackling the problem of transaction costs. In section 6, we outline the robustness checks detailed in the Internet Appendix A. Section 7 concludes. Details about the theory of *MVTC* and the data are in Internet Appendices B and C.

2 Investment Opportunity Set in FX Markets

2.1 Data

We take the view of a US investor. The short-term bond denominated in the U.S. dollar (USD) is the risk-free asset. For the universe of risky assets, we use currency returns for N_t currencies against the USD. Depending on data availability at time t , the number of currencies N_t changes through time.

Following the literature, we include the currencies of 26 countries in our analysis. Of these, 13 belong to the set of countries classified as developed in [Lustig et al. \(2011\)](#), while the remaining 13 are classified as emerging. The 13 developed countries and currency codes are: Australia (AUD), Canada (CAD), the Euro Area (EUR), France (FRF), Germany (DEM), Italy (ITL), Japan (JPY), Netherlands (NLG), New Zealand (NZD), Norway (NOK), Sweden (SEK), Switzerland (CHF), United Kingdom(GBP). The 13 emerging countries and currency codes are: Brazil (BRL), Czech Republic (CZK), Greece (GRD), Hungary (HUF), Ireland (IEP), Mexico (MXN), Poland (PLN), Portugal (PTE), Singapore (SGD), South Africa (ZAR), South Korea (KRW), Spain (ESP), and Taiwan (TWD). We provide details about data sources and filters in Internet Appendix C.

We obtain daily bid, ask and mid quotes for spot and forward exchange rates with 1- and 3-month maturities from Barclays Bank International and Reuters (B&R) via Datastream. In the following, we focus on forwards with 1-month maturity but all computations are analogous for 3-month contracts. We denote by $X_{i,t}^{BR}$ and $F_{i,t}^{BR}$ the spot and forward exchange

rates from B&R as USD per unit of currency i on day t . We add superscripts b and a to indicate bid and ask quotes. We also collect hourly spot exchange rates from Olsen, and denote by $X_{i,t,\tau}^O$ the USD per unit of currency i at hour τ on day t . Olsen provides spot quotes against the USD from January 1986 to September 2021, while the B&R data exchange rates are available against the USD starting from October 1983 for both spot and forwards. Our analysis is conducted over the sample covered by both sources, i.e., starting in January 1986 and ending in September 2021.

Spot quotes from Olsen are more comprehensive in terms of market coverage, and these data are usually regarded as more reliable and of superior quality compared to the B&R data. Comparing daily spot quotes of the two datasets, we find that mid-prices are the same

$$X_{i,t} = X_{i,t}^{BR} \approx X_{i,t}^O,$$

but proportional bid-ask spreads are generally narrower in the Olsen data for the set of developed currencies. This is consistent with the literature suggesting that bid-ask spreads in the B&R data are too wide (Lyons, 2001). However, the same is not necessarily true for the emerging currencies where Olsen spreads are often wider. Acknowledging this feature of the data we reconstruct the time series of spot bid and ask quotes as follows. For currency i against the USD we compute a daily estimate for the proportional bid-ask spread of Olsen by averaging the hourly spreads within day t , and denote it as $PBA_{i,t}^{S,O} = \frac{1}{H_t} \sum_{\tau} \ln \left(\frac{X_{i,t,\tau}^{a,O}}{X_{i,t,\tau}^{b,O}} \right)$, where H_t is the number of trading hours on day t . For the B&R data we use the daily bid and ask spot quotes to compute the spread $PBA_{i,t}^{S,BR} = \ln \left(\frac{X_{i,t}^{a,BR}}{X_{i,t}^{b,BR}} \right)$. We then define the daily ask and bid spot prices

$$X_{i,t}^a = X_{i,t} \exp \{0.5 PBA_{i,t}^S\}, \quad X_{i,t}^b = X_{i,t} \exp \{-0.5 PBA_{i,t}^S\},$$

where $PBA_{i,t}^S = \min(PBA_{i,t}^{S,O}, PBA_{i,t}^{S,BR})$ is the minimum between the daily proportional bid-ask spreads of Olsen and B&R spot quotes.

Forward quotes are only available from B&R, and therefore we can only compute the daily $PBA_{i,t}^{F,BR} = \ln \left(\frac{F_{i,t}^{a,BR}}{F_{i,t}^{b,BR}} \right)$ from daily bid and ask forward quotes. Accordingly, we cannot

use the same procedure as for spots to address the concern that B&R spreads for forwards are too wide. Instead, we construct an adjusted bid-ask spread for forwards by re-scaling $PBA_{i,t}^{F,BR}$ by the average ratio of our adjusted spot spreads $PBA_{i,t}^S$ and the B&R spot spreads $PBA_{i,t}^{S,BR}$,

$$PBA_{i,t}^F = PBA_{i,t}^{F,BR} \frac{1}{T} \sum_t \frac{PBA_{i,t}^S}{PBA_{i,t}^{S,BR}}.$$

The assumption is that the issue of too wide B&R spreads is roughly the same across spots and forwards. We then define the daily mid, ask and bid forward prices:

$$F_{i,t} = F_{i,t}^{BR}, \quad F_{i,t}^a = F_{i,t} \exp \{0.5 PBA_{i,t}^F\}, \quad F_{i,t}^b = F_{i,t} \exp \{-0.5 PBA_{i,t}^F\}.$$

In addition to the price data we use quantity data to estimate the price impact of large trade orders on the bid and ask quotes (details on the price impact measure are in section 2.4). We obtain spot order flows from the CLS Group (CLS) via Quandl.com. Flows are defined as buy or sell orders of foreign currency against the USD, and they are measured in USD. The order flow data span the period of September 2012 until September 2021, and include the following currencies: AUD, CAD, EUR, JPY, MXN, NZD, NOK, SGD, ZAR, SEK, CHF, GBP.⁶ CLS provides data across four categories: corporates, funds, non-bank financials, and total buy-side. We follow [Rinaldo and Somogyi \(2021\)](#) to identify banks acting as price takers by subtracting signed flows of corporates, funds, and non-bank financials from the total buy-side. Finally, on day t we aggregate the orders for currency i across all participants (i.e., corporates, funds, non-bank financials, and banks acting as price takers) to obtain the total spot trading volume $v_{i,t}^S$.

2.2 Monthly Returns and Trading Positions

Following the literature, we define the 1-month realized currency return as the return of an uncovered long position in the 1-month forward exchange rate contract of currency i against

⁶Following [Rinaldo and Somogyi \(2021\)](#) we exclude HUF as its coverage starts only on November 7, 2015. We further exclude KRW due to the limited number of trades per price taker category.

the USD (denominated in USD),

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

where $fd_{i,t} = \ln \left(\frac{X_{i,t}}{F_{i,t}} \right)$ is the forward discount which is known at time t , and $\Delta x_{i,t+1} = \ln \left(\frac{X_{i,t+1}}{X_{i,t}} \right)$ is the spot exchange rate growth which is realized at time $t+1$. Note that $r_{i,t+1}$ is an excess return over the risk-free rate in USD. We use the last mid-price quotes of the month to compute monthly returns before accounting for trading costs.

For trading strategy j , currency i and time t we denote by $\theta_{i,t}^{0,j}$ the initial position before trading takes place, and by $\theta_{i,t}^j$ the portfolio weight after trading. We define θ as a fraction of a fund's assets under management (AUM) measured in USD, or more generally, a fraction of invested wealth. We denote the AUM or wealth at time t by W_t . Then, the trading activity in currency i is given by

$$\theta_{i,t}^j - \theta_{i,t}^{0,j} = \Delta_{i,t}^{P+} + \Delta_{i,t}^{P-} - \Delta_{i,t}^{S-} - \Delta_{i,t}^{S+}. \quad (1)$$

which keeps track of:

- opening new long forward positions $\Delta_{i,t}^{P+} = \max \{ \theta_{i,t}^j - \max(\theta_{i,t}^{0,j}, 0), 0 \}$,
- closing existing short forward positions $\Delta_{i,t}^{P-} = \max \{ \min(\theta_{i,t}^j, 0) - \theta_{i,t}^{0,j}, 0 \}$,
- opening new short forward positions $\Delta_{i,t}^{S-} = \max \{ \min(\theta_{i,t}^{0,j}, 0) - \theta_{i,t}^j, 0 \}$,
- closing existing long forward positions $\Delta_{i,t}^{S+} = \max \{ \theta_{i,t}^{0,j} - \max(\theta_{i,t}^j, 0), 0 \}$.

Similar to θ , Δ are fractions of W_t . To obtain the USD amount traded we multiply Δ by W_t .

We consider two types of transaction costs: (i) proportional costs PC_t in equation (5), and (ii) price impact of trading PI in equation (6). Both are measured in terms of a fraction of the invested wealth. To obtain trading costs in USD we need to multiply these quantities by the fund's current AUM, W_t . Proportional costs are characterized by the bid-ask spread, and capture the costs that investors pay for small trading orders. The price impact of trading

accounts for the widening in the bid-ask spread in response to large orders. When an investor places a large buy (sell) order and the market is not deep enough to fully absorb it at the current ask (bid) quote, then this order puts upward (downward) pressure on the ask (bid) price, and the marginal cost to execute the order increases.

We denote by $\tilde{W}_t = W_t(1 - PC_t - PI_t)$ the wealth at time t after trading costs have been paid. Then, the wealth process is given by

$$\begin{aligned} W_t &= \sum_{i=1}^{N_t} (1 + r_{US,t-1} + r_{i,t}) \theta_{i,t-1}^j \tilde{W}_{t-1} + (1 + r_{US,t-1}) \left(1 - \sum_{i=1}^{N_t} \theta_{i,t-1}^j \right) \tilde{W}_{t-1} \\ &= \left(1 + r_{US,t-1} + \sum_{i=1}^{N_t} \theta_{i,t-1}^j r_{i,t} \right) (1 - PC_{t-1} - PI_{t-1}) W_{t-1}, \end{aligned}$$

where $r_{US,t-1}$ denotes the risk-free rate in USD, and we use the fact that $1 - \sum_{i=1}^{N_t} \theta_{i,t-1}^j$ is invested in the risk-free asset in the USD. Therefore, we define the realized excess return before costs r_t^j and the equivalent after cost return $r_t^{j,c}$ as

$$r_t^j = \sum_{i=1}^{N_t} r_{i,t} \theta_{i,t-1}^j, \quad (2)$$

$$r_t^{j,c} = r_t^j - (1 + r_{US,t-1} + r_t^j) (PC_{t-1} + PI_{t-1}). \quad (3)$$

From post-trading date $t - 1$ to before-trading date t , the USD amount invested in currency i changes according to

$$\theta_{i,t}^{0,j} W_t = (1 + r_{US,t-1} + r_{i,t}) \theta_{i,t-1}^j \tilde{W}_{t-1}.$$

Accordingly, the initial position in currency i at time t is given by

$$\theta_{i,t}^{0,j} = \frac{1 + r_{US,t-1} + r_{i,t}}{1 + r_{US,t-1} + r_t^j} \theta_{i,t-1}^j.$$

Our setup only considers trades θ and Δ against the USD. This ignores the possibility that it may be cheaper to trade directly between two non-USD currencies rather than twice against the USD. We address this limitation in robustness tests in section 6 and Internet Appendix A.2, and confirm that our main results are unaffected.

2.3 Proportional Costs

For the proportional costs, we follow the literature (Menkhoff et al., 2012b; Della Corte et al., 2016; Maurer et al., 2023), and account for bid and ask spreads. 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.

Our measure of the per dollar proportional transaction costs to open new long positions ($\mathbf{C}_{i,t}^{\mathbf{P}^+}$), close existing short positions ($\mathbf{C}_{i,t}^{\mathbf{P}^-}$), open new short positions ($\mathbf{C}_{i,t}^{\mathbf{S}^-}$) and close existing long positions ($\mathbf{C}_{i,t}^{\mathbf{S}^+}$) are

$$\mathbf{C}_{i,t}^{\mathbf{P}^+} \equiv \ln\left(\frac{F_{i,t}^a}{F_{i,t}^b}\right), \quad \mathbf{C}_{i,t}^{\mathbf{P}^-} \equiv \ln\left(\frac{X_{i,t}^a}{X_{i,t}^b}\right), \quad \mathbf{C}_{i,t}^{\mathbf{S}^-} \equiv \ln\left(\frac{F_{i,t}}{F_{i,t}^b}\right), \quad \mathbf{C}_{i,t}^{\mathbf{S}^+} \equiv \ln\left(\frac{X_{i,t}}{X_{i,t}^b}\right). \quad (4)$$

In the data, bid-ask spreads of forward contracts are wider than spreads of spot exchange rates. Our measure of proportional costs accounts for these differences. The proportional costs to trade the N_t currencies at time t reduce the portfolio return by

$$PC_t = \sum_{i=1}^{N_t} \sum_{k \in \mathcal{K}} \underbrace{\mathbf{C}_{i,t}^k}_{\text{cost per dollar traded}} \overbrace{\Delta_{i,t}^k}^{\text{quantity traded}}, \quad \text{where } \mathcal{K} = \{P^+, P^-, S^-, S^+\}. \quad (5)$$

Recall that the proportional costs PC_t are measured in terms of a fraction of the invested wealth. To obtain trading costs in USD we need to multiply PC_t by the AUM W_t of the fund.

Figure 1 plots the time-series of the cross-currency median of the average proportional costs for the spot market, $\text{med}(0.5C_{i,t}^{\mathbf{S}^+} + 0.5C_{i,t}^{\mathbf{P}^-})$ (solid line), and the forward market, $\text{med}(0.5C_{i,t}^{\mathbf{P}^+} + 0.5C_{i,t}^{\mathbf{S}^-})$ (dashed line). Proportional costs are roughly between 0.5 and 7 basis points (bps), are slightly higher and more volatile for the 1-month forwards, and decrease over time after the year 2000. The highest spike (around 6.5 bps) in the forward time-series occurs during the 2008 financial crisis, while the most illiquid period for the spot market is between the 1997 Asian crisis and the Dot-com crisis in the early 2000s. Figures A2 and A3 in the Internet Appendix show the country specific time-series, confirming the fact that

costs are generally lower in spot markets.

2.4 Price Impact of Trading

In our analysis we rebalance trading strategies once on the last day of the month and we implicitly assume that the execution of a trade is spread out over the course of the day. Therefore, the price impact captures the average pressure during the day.

We use the price impact functions $\Pi_{i,t}^k(\Delta_{i,t}^k)$ for $\mathcal{K} = \{P+, P-, S-, S+\}$ to quantify changes in spot and forward ask and bid quotes induced by trading currency i on day t . These changes in quotes increase costs and are additive to the proportional costs. We separately construct price impact functions for ask and bid quotes in spot and forward markets, each currency i , and point in time t , as liquidity starkly varies across markets, currencies and time.

Closing existing short positions ($\Delta_{i,t}^{P-}$) requires a buy order in the spot market. In response to buy order $\Delta_{i,t}^{P-}$ the log-ask quote in the spot market for currency i at time t increases by $\Pi_{i,t}^{P-}(\Delta_{i,t}^{P-})$. Recall that costs are defined as the difference between the effective ask price (at which trades are executed) and the mid-quote. Accordingly, an increase in ask quotes implies higher trading costs. Similarly, closing existing long positions requires a sell order in spot markets, and sell order $\Delta_{i,t}^{S+}$ decreases the associated spot market log-bid quote by $\Pi_{i,t}^{S+}(\Delta_{i,t}^{S+})$. Finally, opening new long ($\Delta_{i,t}^{P+}$) or short positions ($\Delta_{i,t}^{S-}$) requires trading in forward contracts, and forward log-ask quotes increase by $\Pi_{i,t}^{P+}(\Delta_{i,t}^{P+})$ and log-bid quotes decrease by $\Pi_{i,t}^{S-}(\Delta_{i,t}^{S-})$.

Equation (6) formalizes the price impact cost PI_t in terms of a fraction of the invested wealth to trade the N_t currencies at time t ,

$$PI_t = \sum_{i=1}^{N_t} \sum_{k \in \mathcal{K}} \underbrace{\Pi_{i,t}^k(\Delta_{i,t}^k)}_{\substack{\text{price impact per} \\ \text{dollar traded}}} \overbrace{\Delta_{i,t}^k}^{\substack{\text{quantity} \\ \text{traded}}} \quad \text{where } \mathcal{K} = \{P+, P-, S-, S+\}. \quad (6)$$

Equation (13) in section 2.4.2 provides more details and a linear functional form of $\Pi_{i,t}^k(\Delta_{i,t}^k)$.

2.4.1 Amihud Price Impact Measure

Our estimation of the price impact is based on the measure of Amihud (2002). The Amihud measure captures the percentage change in the price per dollar traded in the spot market. It can be interpreted as a proxy for the price impact parameter λ in Kyle (1985). In the context of FX markets, Rinaldo and Santucci de Magistris (2022) propose a daily realized Amihud measure for spot exchange rates using intraday price data from Olsen and trading volume data from CLS. Lacava et al. (2023) further provide an in-depth theoretical discussion of the realized Amihud measure, and demonstrate its accuracy and robustness. We follow these authors and construct the realized Amihud measure of the spot exchange rate of currency i on day t as,

$$A_{i,t}^S = \frac{RPV_{i,t}}{v_{i,t}^S}, \quad (7)$$

where $RPV_{i,t} = \sum_{\tau=1}^{H_t} |\Delta x_{i,t,\tau}|$ is the sum of the H_t hourly absolute spot exchange rate growths (using mid-price quotes) within day t , and $v_{i,t}^S$ is the daily USD trading volume for spot transactions of currency i against the USD. Note that our computation of the realized Amihud measure slightly differs from that of Rinaldo and Santucci de Magistris (2022). While we use hourly data from Olsen, they use minute-level quotes. However, we show in Figure 2, and explain in the following paragraphs, that our estimates are similar.

We have a good coverage of hourly price data from Olsen, and it is straightforward to estimate $RPV_{i,t}$. In contrast, as explained in section 2.1, we only have CLS volume data for spot transactions for a subset of 12 out of 26 currencies and starting in September 2012 instead of January 1986. Accordingly, we need to extrapolate $v_{i,t}^S$ backward in time, and across currencies. In the following, we denote the set of the 12 currencies with CLS coverage by \mathcal{I} , and we define the complement set \mathcal{I}^c containing the remaining 14 currencies without CLS volume data.⁷

For the extrapolation in the time dimension we simply use the time-series average of the

⁷Specifically, $\mathcal{I} = \{AUD, EUR, GBP, NZD, CAD, CHF, JPY, NOK, SEK, MXN, SGD, ZAR\}$, and $\mathcal{I}^c = \{BRL, CZK, DEM, ESP, FRF, GRD, HUF, IEP, ITL, KRW, NLG, PLN, PTE, TWD\}$.

available volume data. That is, we set

$$v_{i,t}^S = \bar{v}_i^S \quad \forall t, \quad (8)$$

where \bar{v}_i^S is the average of the daily trading volume of currency i from September 2012 to September 2021. This is conservative, as trading volume is generally increasing over time. For instance, the BIS Triennial Survey suggests that the trading volume of EUR/USD spot transactions for the survey years of 2013, 2016 and 2019 was more than twice the trading volume in the years 2001, 2004, and 2007. Other currencies have experienced a similar and often larger increase in trading volume over time. Thus, prior to 2012 our extrapolated Amihud measure $A_{i,t}^S$ in (7) is likely lower when we use $v_{i,t}^S = \bar{v}_i^S$, suggesting that we assume more liquidity and a lower price impact of trading than the actual market conditions. Accordingly, in the early years of our sample we likely underestimate the importance of the price impact, and in reality transaction costs likely had a more severe impact on the profitability of trading strategies than what we report. In section 6 and Internet Appendix A.2, we consider alternative extrapolations of $v_{i,t}^S$ that allow for time-variation. As expected, the time-series of the extrapolated volume data is increasing over time. In turn, this implies that the price impact in our baseline analysis is smaller and more conservative than in the robustness tests. However, we find that the choice of method for extrapolating the volume data does not have a material effect on our final results on the after-cost profitability of our strategies.

Setting $v_{i,t}^S = \bar{v}_i^S$ for all t means that the Amihud measure $A_{i,t}^S$ in (7) is only dynamic due to the time-variation in $RPV_{i,t}$. To get a sense of whether this restriction (and the use of hourly rather than minute-level Olsen data for exchange rate quotes) is problematic, we compare our $A_{i,t}^S$ imposing a time-invariant $v_{i,t}^S = \bar{v}_i^S$ to the realized Amihud measure of spot exchange rates of [Ranaldo and Santucci de Magistris \(2022\)](#).⁸ Figure 2 reports the correlation for daily estimates from September 2012 to September 2021 for each currency $i \in \mathcal{I}$. The correlations range between 0.6 and 0.8, with the exception of the JPY, which has the lowest correlation of 0.49. The generally high correlations indicate that most of the time-series variation in the realized Amihud measure stems from $RPV_{i,t}$, and it is unproblematic

⁸We are thankful to Angelo Ranaldo and Paolo Santucci de Magistris for sharing their realized Amihud estimates for the 12 currencies $i \in \mathcal{I}$ for the sample period from September 2012 to September 2021.

to impose a time-invariant $v_{i,t}^S = \bar{v}_i^S$ for the purpose of our extrapolations. Moreover, the high correlations further confirm that it is not a major concern that we use hourly data instead of minute-level data as employed by [Rinaldo and Santucci de Magistris \(2022\)](#).

Recall that we have CLS volume data and we can compute \bar{v}_i^S only for the 12 currencies $i \in \mathcal{I}$. To obtain \bar{v}_i^S (and construct $v_{i,t}^S = \bar{v}_i^S$ and ultimately $A_{i,t}^S$ in (7)) for the 14 currencies $i \in \mathcal{I}^c$, we extrapolate in the cross-section. [Karnaikh et al. \(2015\)](#) show that bid-ask spreads are a good proxy for liquidity. Motivated by this insight, we use the proportional bid-ask spread as a conditioning variable for the cross-sectional extrapolation of \bar{v}_i^S . Specifically, for currency i we denote by \overline{PBA}_i^S the time-series average from September 2012 to September 2021 of the daily proportional spot bid-ask spread $PBA_{i,t}^S$ (see definition in section 2.1). Then, we regress the natural logarithm of \bar{v}_i^S on a constant and the natural logarithm of \overline{PBA}_i^S for the 12 currencies $i \in \mathcal{I}$,

$$\ln(\bar{v}_i^S) = a + b \ln(\overline{PBA}_i^S) + \varepsilon_i, \quad (9)$$

where a and b are the regression coefficients estimated via OLS, and ε_i is the residual. Figure 3 plots $\ln(\bar{v}_i^S)$ against $\ln(\overline{PBA}_i^S)$ (indicated by asterisks), and visualizes the striking linear relation. The regression fit $R^2 = 0.8$ is astounding. The negative slope coefficient $b = -1.47$ is sensible and has a clear interpretation: lower levels of liquidity proxied by higher spreads are associated with lower traded volume. This provides confidence for the accuracy of a cross-sectional extrapolation based on bid-ask spreads.

Given the linear regression (9) we construct extrapolated average volumes for the 14 currencies $i \in \mathcal{I}^c$,

$$\bar{v}_i^S = \exp\{a + b \ln(\overline{PBA}_i^S)\}. \quad (10)$$

One complication is the availability of bid and ask quote data to compute \overline{PBA}_i^S . We only have daily bid and ask quotes between September 2012 and September 2021 for 6 of the 14 currencies $i \in \mathcal{I}^c$, namely BRL, CZK, HUF, KRW, PLN, and TWD. For these 6 currencies, we can readily use (10) to obtain extrapolated values \bar{v}_i^S , and thus construct $v_{i,t}^S = \bar{v}_i^S$ and

$A_{i,t}^S$ in (7). We plot the extrapolated \bar{v}_i^S in Figure 3 (indicated by circles). A nice feature of the extrapolation is that $\ln(PBA_i^S)$ for $i \in \mathcal{I}$ spans a relatively large range from -9.74 to -7.39. BRL, KRW and TWD lie well within that range, and CSK, HUF and PLN are located just outside with HUF taking the most extreme value of -6.92. Therefore, the extrapolations do not require us to move far out on the estimated line.

The remaining 8 currencies (namely DEM, ESP, FRF, GRD, IEP, ITL, NLG, PTE; which we denote as the subset \mathcal{I}_{EU}^c) have joined the EUR, and it is only possible to compute average proportional bid-ask spreads based on data prior to January 1999. We denote this by $\overline{PBA}_{i,\text{pre-EUR}}^S$. As bid-ask spreads are generally decreasing over time, $\overline{PBA}_{i,\text{pre-EUR}}^S$ for $i \in \mathcal{I}_{EU}^c$ computed from data prior to 1999 are not comparable to \overline{PBA}_i^S for $i \in \mathcal{I}$ constructed from data from September 2012 to September 2021. That is, if we simply plug in the pre-1999 $\overline{PBA}_{i,\text{pre-EUR}}^S$ in (10) then we obtain comparatively low extrapolated values of \bar{v}_i^S for $i \in \mathcal{I}_{EU}^c$. Accordingly, we re-scale the extrapolated volumes such that they add up to the average volume of the EUR, $\sum_{i \in \mathcal{I}_{EU}^c} \bar{v}_i^S = \bar{v}_{EUR}^S$. The idea is that the currencies that join the EUR should collectively have a trading volume that is comparable to that of the EUR. Note that this implicitly assumes the same time-invariant extrapolation as in (8). The adjusted extrapolation formula for currencies $i \in \mathcal{I}_{EU}^c$ is,⁹

$$\bar{v}_i^S = \frac{\bar{v}_{EUR}^S \exp\left\{a + b \ln\left(\overline{PBA}_{i,\text{pre-EUR}}^S\right)\right\}}{\sum_{j \in \mathcal{I}_{EU}^c} \exp\left\{a + b \ln\left(\overline{PBA}_{j,\text{pre-EUR}}^S\right)\right\}}. \quad (11)$$

Finally, we construct the Amihud measure for forwards $A_{i,t}^F$ for each of the 26 currencies by re-scaling $A_{i,t}^S$ by $PBA_i^{F/S} = \frac{1}{T} \sum_t \frac{PBA_{i,t}^F}{PBA_{i,t}^S}$, which is the time-series average over the entire sample of the ratio between the proportional forward and the spot bid-ask spreads,

$$A_{i,t}^F = PBA_i^{F/S} A_{i,t}^S. \quad (12)$$

⁹To the extent that GDP is a good proxy for the traded volume, we compared our re-scaling factor $\frac{\bar{v}_{EUR}^S}{\sum_{j \in \mathcal{I}_{EU}^c} \exp\left\{a + b \ln\left(\overline{PBA}_{j,\text{pre-EUR}}^S\right)\right\}}$ to the average GDP of the EUR countries in our sample relative to the GDP of the Euro area over the period 1995 to 2021 using the data available from the OECD, https://stats.oecd.org/Index.aspx?DatasetCode=SNA_TABLE1#. This comparison produces results in line with our assumptions. First, relative GDP is very stable across time. Second, the correlation between our re-scaling factor and the average relative GDP is 0.7.

$PBA_i^{F/S}$ is greater than 1 for every currency, reflecting the well-known fact that the forward market is less liquid than the spot market. In turn, this consistently reflects a higher price impact per dollar traded in the forward market, i.e. $A_{i,t}^F > A_{i,t}^S \forall \{i, t\}$. The tight link between liquidity in spot and forward markets in (12) is closely related to the empirical findings of Krohn and Sushko (2022), albeit their data is on swaps rather than forwards.

Note that it might be problematic to use forward price and volume data to construct $A_{i,t}^F$ analogously to $A_{i,t}^S$ in (7). The reason is that forward quotes are tightly linked to spot quotes due to no-arbitrage, and much of the variation in forward quotes is indirectly driven by fluctuations in spot prices (and thus spot order flow) rather than by order flows in forward markets. Indeed, hourly forward spreads (i.e., forward minus spot quotes) from Olsen have little variation, implying that spot and forward prices closely move together and have approximately the same variation. Bräuer and Hau (2023) further confirm in their analysis that the correlation is 0.99 between dollar forward and spot. To the extent that order flows in spot markets predominately drive spot and therefore also forward prices, a forward Amihud measure analogous to (7) would be grossly misspecified. As a simple example, consider a forward market that has no trading on a particular day. However, transactions in the corresponding spot market move spot prices, and market makers adjust forward quotes to preclude arbitrage. The Amihud measure for forwards according to (7) would be infinite on that day. In principle, this problem is also present in the construction of the Amihud measure for spot markets. In common with one strand of the literature, we therefore assume that order flows in spot markets are the main drivers while flows in corresponding forward markets are of secondary importance.¹⁰ Accordingly, our construction of $A_{i,t}^F$ is arguably more sensible as it is based on the more precisely measured $A_{i,t}^S$, and uses a simple liquidity adjustment based on bid-ask spreads.

To visualize the price impact estimates, we plot the following time-series: $A_{i,t}^S \bar{z}_i^{funds}$ (for spot transactions) and $A_{i,t}^F \bar{z}_i^{funds}$ (for forward transactions). $\bar{z}_i^{funds} = \frac{1}{T} \sum_t z_{i,t}^{funds}$ is the time-series average of $z_{i,t}^{funds} = |buy_{i,t}^{funds} - sell_{i,t}^{funds}|$, the absolute value of the daily order flow imbalance of funds. $buy_{i,t}^{funds}$ and $sell_{i,t}^{funds}$ represent the USD value of the sum of all CLS

¹⁰For instance, Ranaldo and Santucci de Magistris (2022) implicitly assume the same. However, we acknowledge that for instance Bräuer and Hau (2023) provide evidence that FX hedging demands via forwards can have a significant impact on exchange rates. This introduces some errors in the spot Amihud measure.

registered spot buy respectively sell orders of funds for currency i on day t .^{11,12} The idea is to get a sense of the magnitude of the price impact (i.e., the increase in the ask or the decrease in the bid quote) in response to the average buy and sell imbalance of funds within a day. Figure 4 plots the cross-currency median price impact (reported in bps) for the spot market (solid line) and the forward market (dashed line). As expected, the price impact is higher for the forward market (hovering around 7-8 bps) than the spot market (hovering around 4-5 bps). The price impact in spot and forward markets peaks around the 2008 financial crisis. We further notice a higher volatility of the price impact in the forward market. Finally, there is no downtrend in the price impact as we use a time-invariant volume to extrapolate the volume data in the time dimension. This feature, if anything, makes it more difficult to detect the effects of the price impact on the profitability of our trading strategies. The same insights are confirmed at the country-level as reported in Figures A4 and A5 in the Internet Appendix.

2.4.2 Linear Price Impact Function

Next, we construct a linear price impact function of the trade order size and derive a mapping to the Amihud measure. The linear price impact function is:

$$\Pi_{i,t}^k(\Delta_{i,t}^k) \equiv \pi_{i,t}^k \Delta_{i,t}^k, \quad \forall k \in \{P+, P-, S-, S+\}, \quad (13)$$

with $\pi_{i,t}^k$ specified in (14).

A linear functional form is popular in the literature (Kyle, 1985; Leland, 1985; Back, 1993; Kyle et al., 2018). Consider for now the buy order $\Delta_{i,t}^{P-}$. The mapping between per dollar price impact measure $A_{i,t}^S$ and the parameter $\pi_{i,t}^{P-}$ is derived as follows. By definition of the Amihud measure $A_{i,t}^S$ buy order $\Delta_{i,t}^{P-}$ of a fund with a current AUM of W_t pushes up

¹¹For currencies $i \in \mathcal{I}^c$ (for which we do not have CLS data) we extrapolate \bar{z}_i^{funds} using an analogous approach as for the extrapolation of \bar{v}_i^S . Similar to the volume data we show in Figure A1 in the Appendix that we can reliably extrapolate \bar{z}_i^{funds} .

¹²We only use the order imbalance \bar{z}_i^{funds} for spot transactions. We do not generate an estimate for the forward order flows. This helps us to keep the price impact comparable across the two markets.

the current log-ask price $\ln(X_{i,t}^a)$ to

$$\ln(\tilde{X}_{i,t}^a) = \ln(X_{i,t}^a) + A_{i,t}^S \Delta_{i,t}^{P-} W_t.$$

It follows that the cost (as a fraction of the AUM) that the fund has to pay over and above the mid quote $X_{i,t}$ is

$$\ln\left(\frac{\tilde{X}_{i,t}^a}{X_{i,t}}\right) \Delta_{i,t}^{P-} = \underbrace{\ln\left(\frac{X_{i,t}^a}{X_{i,t}}\right) \Delta_{i,t}^{P-}}_{=C_{i,t}^{P-}} + \underbrace{\ln\left(\frac{\tilde{X}_{i,t}^a}{X_{i,t}^a}\right) \Delta_{i,t}^{P-}}_{=\Pi_{i,t}^{P-}(\Delta_{i,t}^{P-})}.$$

The first term is the proportional costs as described in equation (5), while the second term identifies the costs due to the price impact of trading related to equation (6). Dividing the second term by $\Delta_{i,t}^{P-}$, and comparing it to the right-hand side of equation (13), we obtain

$$\pi_{i,t}^{P-} = A_{i,t}^S W_t.$$

Using a similar argument, we obtain price impact parameters $\pi_{i,t}^{S+}$ to sell currency i in the spot market to close existing long positions, and $\pi_{i,t}^{P+}$ and $\pi_{i,t}^{S-}$ to open new long, respectively short positions in the forward market,

$$\pi_{i,t}^{S+} = \pi_{i,t}^{P-} = A_{i,t}^S W_t, \quad \text{and} \quad \pi_{i,t}^{P+} = \pi_{i,t}^{S-} = A_{i,t}^F W_t, \quad (14)$$

where $A_{i,t}^S$ and $A_{i,t}^F$ are given in (7) and (12) and the discussion in section 2.4.1.

3 Trading Strategies

We construct seven currency trading strategies that are well-known for delivering high returns (section 3.1). Six of these seven strategies rank currencies based on characteristics and construct (mostly equally weighted) long-short portfolios. The seventh strategy is a mean-variance optimized currency portfolio.

As a short preview, our main results suggest that trading costs have first-order impli-

cations on the profitability of currency trading strategies. In certain cases, the impact of transaction costs completely erodes all gains, rendering these strategies unprofitable. To address this problem, in section 3.2 we construct a strategy that aims to achieve the optimal balance between rebalancing the portfolio to earn a high before-cost Sharpe ratio while reducing trading activity to lower transaction costs. This approach is able to efficiently mitigate costs, and has an outstanding after-cost performance.

Following the literature, all our strategies are rebalanced monthly. Our strategies use the information available at the end of month t to construct portfolios that we then hold until the end of the subsequent month $t + 1$. Accordingly, the returns are out-of-sample, and the trading strategies do not suffer from a look-ahead bias.

3.1 Strategies from the Literature

DOL: The dollar strategy (*DOL*) borrows in USD and equally invests in all other currencies (Lustig et al., 2011), $\theta_{i,t}^{DOL} = \frac{1}{N_t}$.

DDOL: The dollar carry (*DDOL*) takes a long position in *DOL* if the median exchange rate forward discount is positive, and a short position otherwise (Lustig et al., 2014), $\theta_{i,t}^{DDOL} = \text{sign}(\text{median}(\{fd_{j,t}\}_{j=1}^{N_t})) \frac{1}{N_t}$.

HML: The carry (*HML*) sorts currencies according to the forward discount into quintiles and borrows in the bottom quintile and invests in the top quintile (Lustig and Verdelhan, 2007). Define Q_t^j as the set of forward discount sorted currencies in quintile j . $1_{i \in Q_t^j}$ is an indicator function equal to 1 if currency i is in set Q_t^j , and 0 otherwise. $Q_t^j = \sum_{i=1}^{N_t} 1_{i \in Q_t^j}$ counts the number of currencies in quintile j at time t . Then, $\theta_{i,t}^{HML} = \frac{1_{i \in Q_t^5}}{Q_t^5} - \frac{1_{i \in Q_t^1}}{Q_t^1}$.

RB: As an alternative to *HML* we also consider the ranked-based (*RB*) version of the carry trade (Asness et al., 2013). In this strategy, the rescaled rank is used as the portfolio weight, $\theta_{i,t}^{RB} = \frac{\vartheta_{i,t}}{0.5 \sum_{i=1}^{N_t} |\vartheta_{i,t}|}$ with $\vartheta_{i,t} = \text{rank}(fd_{i,t}) - \sum_{j=1}^{N_t} \frac{\text{rank}(fd_{j,t})}{N_t}$. As for *HML*, the notional value of the *RB* is equal to 2 by construction. The rank-based weights are not sensitive to the magnitude of forward discounts. As such, *RB* avoids the issue of extreme portfolio weights due to outliers, a desirable property that is shared with *HML*. In contrast to *HML*, *RB* may benefit more from diversification as all currencies are used in the portfolio construction,

while *HML* only uses currencies in the top and bottom quintiles.

MOM: The momentum strategy *MOM* sorts currencies according to their past 12-month performance into quintiles and borrows in the bottom and invests in the top quintile (Burnside et al., 2011; Menkhoff et al., 2012b). Formally, the portfolio weights $\theta_{i,t}^{MOM}$ are defined analogously to $\theta_{i,t}^{HML}$.

VAL: The value strategy (*VAL*) sorts currencies according to the purchasing-power-parity-adjusted (PPP) or real 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) (Bilson, 1984; Menkhoff et al., 2017). We estimate the real exchange rate of currency i against the USD by multiplying the PPP (which is the ratio of prices in currency i and the USD of a representative consumption bundle) by the nominal exchange rate $X_{i,t}$. Formally, the portfolio weights $\theta_{i,t}^{VAL}$ are defined analogously to $\theta_{i,t}^{HML}$.

MV: The mean-variance optimized portfolio (*MV*) solves,

$$\max_{\{\theta_t \in \mathbb{R}^{N_t}\}} \left\{ \theta_t' \mu_t^e - \frac{\lambda}{2} \theta_t' \mathbf{V}_t \theta_t \right\}$$

where $\mu_t^e = E_t[r_{t+1}]$ is the $N_t \times 1$ vector of conditional expected currency returns, and \mathbf{V}_t is the $N_t \times N_t$ covariance matrix. Maurer et al. (2023) show that this mean-variance optimization performs well out-of-sample, and outperforms other types of mean-variance optimized portfolios that impose risk limits or other constraints.

We follow the literature and use the current forward discount $fd_{i,t}$ as a proxy for $\mu_{i,t}^e$ (Baz et al., 2001; Della Corte et al., 2009; Daniel et al., 2017; Ackermann et al., 2016; Maurer et al., 2023, 2022). This is motivated by the random walk conjecture, i.e., the well-known empirical finding that exchange rate changes are difficult to predict over a short horizon, $E_t[\Delta x_{i,t+1}] \approx 0$ (Meese and Rogoff, 1983). Moreover, as demonstrated for instance by Maurer et al. (2023), although there is evidence of predictability in $\Delta x_{i,t+1}$, it does not seem to hold any economic value. They employ a set of 18 predictors and conduct elastic net estimations to forecast $\Delta x_{i,t+1}$. Their results reveal an average out-of-sample R^2 of 0.36% for the monthly predictions. However, when applied to optimized portfolios, they document

that the portfolio based on the predicted $\Delta x_{i,t+1}$ performs worse out-of-sample than the portfolio based on the random walk assumption.

With regard to the conditional covariance matrix \mathbf{V}_t , [Maurer et al. \(2023\)](#) show that the out-of-sample performance of *MV* substantially improves if we eliminate principal components that explain less than 1% of the common variation of the currency returns. This is motivated by the results of [Lustig et al. \(2011\)](#), who document a strong factor structure in FX markets. We first use daily returns over the previous 6 months to obtain the rolling window sample covariance matrix $\hat{\mathbf{V}}_t$ at the end of month t . Then, we define the robust version of the covariance matrix as $\tilde{\mathbf{V}}_t = \mathbf{W}_t \mathbf{\Lambda}_t \mathbf{W}_t'$, and its inverse as $\tilde{\mathbf{V}}_t^{-1} = \mathbf{W}_t \mathbf{\Lambda}_t^{-1} \mathbf{W}_t'$, where \mathbf{W}_t is the matrix of eigenvectors and $\mathbf{\Lambda}_t$ is the diagonal matrix of eigenvalues of $\hat{\mathbf{V}}_t$ after removing eigenvectors and eigenvalues which explain less than 1% of the common variation of the returns. Accordingly, the portfolio weights of *MV* are $\theta_t^{\text{MV}} = \frac{1}{\lambda} \tilde{\mathbf{V}}_t^{-1} f d_t$.

Finally, we follow [Maurer et al. \(2023\)](#) and set $\lambda = 25$ to roughly match the unconditional volatility of 8% of the other currency trading strategies. Section 6 and A11 in the Internet Appendix analyze the effects of using other values for λ .

3.2 Mean-Variance-Trading-Cost Optimization (*MVTC*)

We use the framework of [Dybvig and Pezzo \(2020\)](#) to optimize over costs within a myopic mean-variance setup. We denote this strategy as the mean-variance-trading-cost optimized portfolio (*MVTC*).¹³ We ignore potentially interesting implications of a dynamic model to keep the solution tractable. Therefore, our strategy in general leads to a sub-optimal outcome. However, *MVTC* is still resilient even to a severe price impact of trading and outperforms all of our strategies by a large margin. Accordingly, we are not concerned about its theoretical sub-optimality.

The optimization problem is:

¹³[Olivares-Nadal and DeMiguel \(2016\)](#), [DeMiguel et al. \(2020\)](#) and [Hautsch and Voigt \(2019\)](#) provide alternative solution methods to transaction cost problems.

Problem 1

$$\begin{aligned} & \max_{\{\Delta_{i,t}^k \geq 0\}_{i=\{1,\dots,N_t\}}^{k \in \mathcal{K}}} \left\{ \theta_t^{\text{MVTC}'} \mu_t^e - \frac{\lambda}{2} \theta_t^{\text{MVTC}'} \mathbf{V}_t \theta_t^{\text{MVTC}} - PC_t - PI_t \right\} \\ & \text{s.t. } \theta_{i,t}^{\text{MVTC}} - \theta_{i,t}^{0,\text{MVTC}} = \Delta_{i,t}^{\text{P}+} + \Delta_{i,t}^{\text{P}-} - \Delta_{i,t}^{\text{S}-} - \Delta_{i,t}^{\text{S}+}, \end{aligned}$$

where PC_t and PI_t are defined in (5) and (6).

3.2.1 Economic Insights

The economics of Problem 1 are well established in the literature. For convenience, we briefly review some features here and refer to Internet Appendix B and Dybvig and Pezzo (2020) for additional details. The mean-variance portfolio MV , which does not account for transaction costs in the optimization, is a special case of Problem 1. If we set $PC_t = 0$, $PI_t = 0$, we obtain strategy MV , which maximizes the before-cost Sharpe ratio. The portfolio of MV is independent of the initial position $\theta_t^{0,\text{MV}}$, and it is always optimal to trade all the way to θ_t^{MV} . In contrast, if there are transaction costs $PC_t > 0$ or $PI_t > 0$, then θ_t^{MVTC} crucially depends on the origin $\theta_t^{0,\text{MVTC}}$. Intuitively, there is a trade-off between paying transaction costs and utility gains to move towards θ_t^{MV} .

Consider first the case of only proportional costs ($PC_t > 0$, $PI_t = 0$). If the initial allocation $\theta_t^{0,\text{MVTC}}$ is close enough to θ_t^{MV} , it is optimal not to trade at all as the linear marginal cost required to move towards θ_t^{MV} is higher than the quadratic marginal utility. Thus, there is a no-trading region around θ_t^{MV} . If the initial allocation $\theta_t^{0,\text{MVTC}}$ is far enough from θ_t^{MV} (i.e., outside of the no trading region), then it is optimal to move towards θ_t^{MV} but only until θ_t^{MVTC} . This is because the marginal utility of moving towards θ_t^{MV} is diminishing, and at the boundary of the no-trading region, where θ_t^{MVTC} is located, the marginal utility is equal to the relevant entries of the marginal cost vectors \mathbf{C}_t^k , $k \in \mathcal{K}$. Appendix B.2 illustrates such no-trading regions for the case of two risky assets.

A price impact ($PI_t > 0$) does not affect the shape or size of the no-trading region, but the optimal trades (which are initiated outside of the no-trading region) are different

and we never end up at the border of the no-trading region.¹⁴ Furthermore, [Dybvig and Pezzo \(2020\)](#) show that the directions of the trades are neither guaranteed to be towards the no-trading region nor that it is optimal to trade less than if there is no price impact. The starting point $\theta_t^{0,MVTC}$ and the other model parameters determine the size and the direction of the optimal trades. Therefore, it is ultimately an empirical question how a price impact affects trading.

The theoretical setup provides us with three predictions characterizing *MVTC*. First, we expect *MV* to outperform *MVTC* if the performance is measured in returns before transaction costs. This is because, by definition, *MV* is the optimal portfolio when evaluated before transaction costs (in a single-period model). Second, we expect transaction costs to be larger for *MV* than for *MVTC*. Third, we expect *MVTC* to outperform *MV* after transaction costs. Finally, the size of these differences between *MV* and *MVTC* are expected to depend on the magnitude of the transaction costs.

3.2.2 Implementation

Problem 1 does not have a closed-form solution. However, it is a well-behaved quadratic program for which accurate and fast converging numerical solutions are readily available. We solve Problem 1 for θ_t^{MVTC} with three adjustments.

First, we set all proportional costs in (4) equal, $\mathbf{C}_{i,t}^k = \mathbf{C}_{i,t} = 0.25 \sum_{h \in \mathcal{K}} \mathbf{C}_{i,t}^h, \forall k \in \mathcal{K}$. Second, we equalize the price impact parameters in (14), $\pi_{i,t}^k = \pi_{i,t} = 0.25 \sum_{h \in \mathcal{K}} \pi_{i,t}^h, \forall k \in \mathcal{K}$, and thus, the price impact functions (13) are identical across order types, $\Pi_{i,t}^k(\Delta_{i,t}^k) = \Pi_{i,t}(\Delta_{i,t}^k), \forall k \in \mathcal{K}$. These adjustments simplify our optimization problem as we can now combine $\Delta_{i,t}^{\mathbf{P}+}$ and $\Delta_{i,t}^{\mathbf{P}-}$ to $\Delta_{i,t}^{\mathbf{P}} = \Delta_{i,t}^{\mathbf{P}+} + \Delta_{i,t}^{\mathbf{P}-}$, and similarly $\Delta_{i,t}^{\mathbf{S}} = \Delta_{i,t}^{\mathbf{S}-} + \Delta_{i,t}^{\mathbf{S}+}$. Then, equation (1) becomes $\theta_{i,t}^{MVTC} - \theta_{i,t}^{0,MVTC} = \Delta_{i,t}^{\mathbf{P}} - \Delta_{i,t}^{\mathbf{S}}$. Furthermore, the performance reduction due to proportional costs (5) simplifies to $PC_t^{(adj)} = \sum_{i=1}^{N_t} \mathbf{C}_{i,t} (\Delta_{i,t}^{\mathbf{P}} + \Delta_{i,t}^{\mathbf{S}})$, and the cost due to the price impact (6) reduces to $PI_t^{(adj)} = \sum_{i=1}^{N_t} \pi_{i,t} (\Delta_{i,t}^{\mathbf{P}} + \Delta_{i,t}^{\mathbf{S}})$.

In principle, it is possible to implement Problem 1 in its original formulation, but our empirical results are qualitatively unaffected. Accordingly, it is efficient to simplify the cost

¹⁴This is shown in Theorem 2 of [Dybvig and Pezzo \(2020\)](#) for the case of a linear price impact.

structure. Note that we use the approximate proportional costs $PC_t^{(adj)}$ and price impact $PI_t^{(adj)}$ in the optimization to construct θ_t^{MVTC} , but we use the actual cost function PC_t in (5) and PI_t in (6) (with the linear specification detailed in (13) and (14) in section 2.4.2) when we compute after-cost returns in our out-of-sample performance evaluation.

Third, the solution approach of Problem 1 requires the covariance matrix to be full rank. Unfortunately, $\tilde{\mathbf{V}}_t$ is not full rank. Therefore, we re-write Problem 1 to explicitly incorporate θ_t^{MV} as the target portfolio,

Problem 2 (Strategy MVTC)

$$\begin{aligned} & \max_{\{\Delta_{i,t}^P \geq 0, \Delta_{i,t}^S \geq 0\}_{i=\{1, \dots, N_t\}}} \left\{ \begin{array}{l} \frac{1}{2} \theta_t^{MV'} f d_t - \frac{\lambda}{2} (\theta_t^{MVTC} - \theta_t^{MV})' \hat{\mathbf{V}}_t (\theta_t^{MVTC} - \theta_t^{MV}) \\ - PC_t^{(adj)} - PI_t^{(adj)} \end{array} \right\} \\ & s.t. \quad \theta_{i,t}^{MVTC} - \theta_{i,t}^{0,MVTC} = \Delta_{i,t}^P - \Delta_{i,t}^S \quad and \quad \theta_t^{MV} = \frac{1}{\lambda} \tilde{\mathbf{V}}_t^{-1} f d_t. \end{aligned}$$

$\hat{\mathbf{V}}_t$ is the sample covariance matrix estimated from daily currency returns over the past 6 months (which is full rank), and $\tilde{\mathbf{V}}_t^{-1}$ is the robust inverse of the covariance matrix $\hat{\mathbf{V}}_t$ after removing principal components that explain less than 1% of the common variation of currency returns (see section 3.1 for details). In Internet Appendix B we provide details of the numerical algorithm used to solve Problem 2. Intuitively, θ_t^{MVTC} identifies the optimal trade-off between (i) minimizing the distance to the maximum (before-cost) Sharpe ratio portfolio θ_t^{MV} , and (ii) minimizing trading costs when rebalancing from $\theta_t^{0,MVTC}$ to θ_t^{MVTC} .

Formally, if we set $\tilde{\mathbf{V}}_t^{-1} = \hat{\mathbf{V}}_t^{-1}$, then Problem 2 is equivalent to Problem 1. This is because the objective function in Problem 2 in this case is an algebraic re-arrangement of that in Problem 1. Intuitively, in the formulation of Problem 2 the maximum (before-cost) Sharpe ratio portfolio θ_t^{MV} is surrounded by the no-trading region from Problem 1. In more practical terms, when our portfolio construction approach requires the inverse of the covariance matrix then we use the robust version $\tilde{\mathbf{V}}_t^{-1}$. Moreover, when our approach needs an estimate of the covariance matrix, and full rank is required, then we resort to $\hat{\mathbf{V}}_t$. With this transformation we retain the feasibility of the approach of Dybvig and Pezzo (2020) while exploiting the superior performance of the robust mean-variance optimized portfolio

of [Maurer et al. \(2022, 2023\)](#).

4 Performance of Trading Strategies

Section 4.1 summarizes the before cost performance of the eight strategies, and confirms the findings in the literature that they are profitable and achieve high Sharpe ratios. Section 4.2 quantifies the trading costs depending on the AUM of a fund, and section 4.3 analyzes the after-cost performance. The main results are in Table 1. Panel A considers the case when there are only proportional costs but no price impact of trading. This applies to funds with a small AUM and trading volume. Panel B and C show the performance for funds with initial AUM of USD 100 million and USD 1 billion in February 1986. These funds place sizable orders and the price impact of trading is non-trivial.

4.1 Performance Before Costs

We start our discussion with a comparison of the before-cost performance of *MV*, *HML*, *RB*, *DOL*, *DDOL*, *VAL*, and *MOM*. The first three rows of panel A in Table 1 report the annualized before-cost Sharpe ratio, average return, and volatility. Consistent with the literature, we find that most strategies have an attractive before-cost performance. *MV* dominates all other strategies with a before-cost Sharpe ratio of 1.10. The interest-rate-sorted *RB* and *HML* deliver substantially lower but still high before-cost Sharpe ratios of 0.74 and 0.72. The *VAL*, *DDOL*, and *MOM* have relatively modest Sharpe ratios of 0.46, 0.35, and 0.30. Finally, the *DOL* is essentially unprofitable, with a Sharpe ratio of only 0.12. The stark before-cost outperformance of *MV* resonates with previous results in the literature ([Baz et al., 2001](#); [Della Corte et al., 2009](#); [Daniel et al., 2017](#); [Ackermann et al., 2016](#); [Maurer et al., 2023, 2022](#); [Chernov et al., 2023](#)).

The Sharpe ratios of *DDOL* (0.35) and *MOM* (0.30) are lower than the values in the original papers in the literature. [Lustig et al. \(2014\)](#) report a Sharpe ratio of 0.56 for *DDOL* (and 0.66 if it is constructed from currencies of only developed countries), and [Menkhoff](#)

et al. (2012b) find a Sharpe ratio of 0.61 for *MOM*.¹⁵ The main difference is the sample period. The former study uses data from 1983 to 2010, the sample of the latter is from 1976 to 2010, while our sample starts in 1986 and ends in 2021. We find that *DDOL* and *MOM* performed well until the global financial crisis (GFC), but since 2008 their average returns were low or even negative in the case of *MOM*.

In robustness checks reported in Section 6 and in Internet Appendix A.1, we analyze different sample periods. As a brief preview, the first plot in Figure 5 illustrates that the cumulative return of *MOM* is comparable to *HML* until 2008. After that, *HML* continues to earn relatively high average returns while *MOM* consistently makes losses. In addition, Table A3 in the Internet Appendix reproduces Table 1 for the sample period ending in December 2007.¹⁶ The first row in panel A shows that the before-cost Sharpe ratios of *DDOL* and *MOM* are 0.56 and 0.66, respectively, which are close to the values reported by Lustig et al. (2014) and Menkhoff et al. (2012b).

By construction, the annual volatilities are similar across all strategies, and range between 6.9% and 8.7%. The characteristic-sorted long-short strategies maintain a constant notional value of either one (*DOL*, *DDOL*) or two (*HML*, *RB*, *VAL*, *MOM*) resulting in comparable volatilities. *RB* has the lowest volatility of 6.88%. As pointed out in Section 3.1 *RB* invests in all currencies, resulting in better diversification. This explains the lower volatility compared to *HML* or other characteristic-sorted long-short strategies that only invest in the top and bottom quintiles. In the construction of *MV*, we set $\lambda = 25$ to target a volatility that is comparable to the characteristic-sorted long-short strategies. Specifically, *MV* has a volatility of 8.1% (see Appendix A.3.2 for a sensitivity analysis of *MV* to different values of λ).

So far, our discussion refers to the values reported in panel A of Table 1. These results apply to sufficiently small funds that face no price impact of trading. However, the results generalize to any fund size and transaction costs. The reason is that *MV*, *HML*, *RB*, *DOL*,

¹⁵They further find that *MOM* based on a shorter formation period is more profitable. For instance, if they use a 1-month formation period *MOM* achieves a Sharpe ratio of 0.95. In our sample we find that a 12-month formation period yields a more profitable *MOM* than a 1-month formation period.

¹⁶2008 appears to identify a structural break for the profitability of many FX trading strategies. After 2008 many strategies perform substantially worse than in the sample before the GFC.

DDOL, *VAL*, and *MOM* do not account for transaction costs in their construction. As such, the portfolio holdings and the before-cost performance of these strategies are unaffected by the AUM or trading costs. Indeed, the values in the top three rows across panels A (small fund without price impact), B (initial fund size of USD 100 million) and C (initial fund size of USD 1 billion) are identical for the seven strategies. This illustrates that there are no implications of the fund size or transaction costs on the before-cost performance.

Next, we analyze the before-cost performance of *MVTC*, which is novel and has not been analyzed in the literature. As *MVTC* accounts for trading costs in the optimization, theory predicts that it decreases the trading order size to reduce transaction costs. In terms of before-cost performance, this reduction in trading activity implies a sub-optimal rebalancing policy and asset allocation. Therefore, we expect that *MVTC* underperforms *MV* when evaluated in before-cost returns, as the latter targets the highest before-cost Sharpe ratio irrespective of trading costs. We further conjecture that an increase in the AUM and costs due to the price impact of trading have negative implications on the before-cost performance of *MVTC*.

We confirm these predictions in the data. The before-cost Sharpe ratio of *MVTC* is 1.06 in panel A of Table 1 when there is no price impact of trading. That is, *MVTC* underperforms *MV*. However, the difference between the Sharpe ratios is small and insignificant, 1.06 for *MVTC* and 1.10 for *MV*. Indeed, if there is no price impact of trading then *MV* and *MVTC* are highly correlated, and they are very similar in terms of all performance measures listed in panel A of Table 1. It is interesting that *MV* and *MVTC* are almost indistinguishable when there is no price impact. This means that proportional costs do not have a noticeable negative effect on the before-cost performance of *MVTC*.

In contrast, the negative implications of the price impact of trading on the before-cost performance of *MVTC* are sizable. The before-cost Sharpe ratio of *MVTC* decreases to 0.93 and 0.81 when the initial fund size is USD 100 million in panel B and USD 1 billion in panel C, respectively. In other words, the sub-optimal allocation of *MVTC* in terms of before-cost performance becomes noticeable when a fund reaches a sufficiently large size and trading costs due to price impact are considerable.

The annual volatility of *MVTC* is decreasing in the initial AUM. It is 7.9% when there is no price impact of trading (panel A), decreases to 7.4% when the initial AUM is USD 100 million (panel B), and further drops to 6.6% when the initial fund size is USD 1 billion (panel C). The decrease in volatility can be explained by the reduction in the trading activity of *MVTC* in response to higher transaction costs. A lower turnover means more stability in the portfolio weights. This is similar to shrinkage methods that are designed to tackle estimation errors and reduce out-of-sample volatility (Olivares-Nadal and DeMiguel, 2016). However, this shrinkage effect is secondary, and the impact on volatility is relatively small. We observe a stronger effect of the fund size or trading costs on the before-cost average return of *MVTC*. When there is no price impact of trading the average annual return is 8.4% (panel A). It decreases to 6.9% and 5.4% when the initial AUM are USD 100 million and USD 1 billion, respectively. The stark decline in the average return dominates the relatively small effect on volatility, resulting in a before-cost Sharpe ratio that is decreasing in the fund size.

Our results are opposite to the key insights of Yoshimoto (1996), Olivares-Nadal and DeMiguel (2016), Hautsch and Voigt (2019) and Pezzo et al. (2023), who analyze the implications of transaction costs on portfolio choice in stock markets. These authors find that accounting for transaction costs in the construction of stock portfolios primarily introduces stability (similar to a shrinkage estimator), and therefore mitigates estimation error problems. The actual mitigation of transaction costs appears secondary in their analyses. It is particularly important to note that the negative effect of transaction costs on the before-cost Sharpe ratio of *MVTC* is the opposite. That is, if an estimation error mitigation effect was present in our analysis (i.e., the dominant effect in the stock market literature), then the before-cost Sharpe ratio of *MV* would be smaller than the one of *MVTC*, and likely the before-cost Sharpe ratio of *MVTC* would be increasing in transaction costs. Therefore, the insights of the stock market literature do not carry over to the present analysis.

The reason why the insights of the stock literature are not applicable is that estimation errors are not a (first-order) concern in FX markets, while it is a serious issue in stock markets. First of all, the literature has shown that interest rate differentials are a good proxy for conditional expected currency returns, and optimized portfolios based on this proxy perform well out-of-sample (Meese and Rogoff, 1983; Baz et al., 2001; Della Corte

et al., 2009; Daniel et al., 2017; Ackermann et al., 2016; Maurer et al., 2023, 2022; Chernov et al., 2023). In contrast, it is difficult to find good estimates of conditional expected returns in stock markets. Second, FX strategies typically use a relatively small set of currencies (less than 30) compared to typical stock portfolios (several thousands). In addition, Lustig et al. (2011) show that there is a strong factor structure in FX markets while there is a large “factor zoo” in stock markets.¹⁷ Accordingly, it is feasible to obtain precise estimates of the covariance matrix of returns in FX markets, while this is a formidable problem in stock markets.

Finally, we benchmark our currency strategies to the US stock market.¹⁸ The buy-and-hold value-weighted stock market portfolio has a before-cost Sharpe ratio of 0.58 during our sample period (February 1986 to September 2021). The annual average excess return and volatility are 9.0% and 15.5%, respectively. Accordingly, *MV*, *MVTC*, *RB* and *HML* outperform the US stock market in terms of before-cost Sharpe ratio. In contrast, *DOL*, *DDOL*, *VAL*, and *MOM* deliver an inferior performance.

The inferior performance does not necessarily mean that these strategies are undesirable. First, recall that before the GFC, *DDOL* and *MOM* performed well with before-cost Sharpe ratios of 0.60 and 0.68, respectively. They only perform poorly after 2008.¹⁹ As a comparison, during the sample period from 1986 to December 2007, the US stock market earned a Sharpe ratio of only 0.50. Accordingly, pre-GFC *DDOL* and *MOM* outperform the US stock market in terms of before-cost Sharpe ratio. Second, if the correlation between the strategies is sufficiently low, then a combination of the strategies may be more efficient than investing in a single strategy. In other words, even a strategy with a relatively low Sharpe ratio may improve the performance of the portfolio. This is particularly relevant when we consider currency portfolios and the US stock market, as the correlations between them are low.²⁰

¹⁷Subsequent to Lustig et al. (2011) the FX literature has introduced many additional factors beyond the *DOL* and *HML*. However, recently Sarno et al. (2023) and Liu et al. (2023) show that these additional factors in the literature are subsumed by *DOL* and *HML*.

¹⁸Stock market excess returns are downloaded from Kenneth French’s website, https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

¹⁹The performance of the *VAL* appears unaffected by the GFC, and the before-cost Sharpe ratio is 0.46 in the subsample before and after 2008.

²⁰However, among currency portfolios Maurer et al. (2023, 2022) show that *MV* outperforms combinations of other currency strategies, and other strategies do not earn an abnormal return when *MV* is used as a pricing factor. As such *MV* on its own appears to be efficient (in terms of before cost returns) within the

To sum up, the before-cost performance of *MVTC* is virtually indistinguishable from *MV* when there are only proportional costs. However, when a fund has a large AUM and there is a price impact of trading, then the deviation is economically large of *MVTC* from the optimal balance of before-cost expected return and risk, resulting in a noticeably lower before-cost Sharpe ratio. Furthermore, independent of the fund size, *MV* and *MVTC* significantly outperform popular characteristic-sorted currency trading strategies in terms of before-cost returns. Finally, within our full sample, only *MV*, *MVTC* and the carry trade strategies *HML* and *RB* deliver a higher before-cost Sharpe ratio than the value-weighted US stock market portfolio. *DOL*, *DDOL*, *VAL*, and *MOM* underperform the US stock market.

4.2 Transaction Costs

Table 1 reports the average annualized trading costs, the monthly turnover, and the average notional value in the fourth to sixth rows in panels A, B, C. The costs are measured by $PC_t + PI_t$ and are measured as a percentage of the AUM of a fund. The monthly turnover of strategy S is defined as $\sum_{i=1}^{N_t} |\theta_{i,t}^S - \theta_{i,t}^{0,S}|$, and quantifies the fraction of the AUM that is traded every month. The notional value is $\sum_{i=1}^{N_t} |\theta_{i,t}^S|$, and identifies the combined notional value of all long and short forward positions as a fraction of the AUM. The ratio of the turnover to the notional value (TO-NV) determines the fraction of the portfolio's notional value that is traded every month.

MV has an average monthly turnover of 2.36, an average notional value of 3.15, and a TO-NV of 0.75. These quantities are independent of the fund size. The average annual costs incurred by *MV* are 1.1% when a fund is small and only pays proportional costs. The costs increase to 6.0% and even 12.0% when the initial AUM of a fund are USD 100 million or USD 1 billion. These costs are considerable given that the average before cost return of *MV* is 8.9%. While small funds (no price impact) are still profitable after costs, for larger funds the before-cost performance is quickly eroded by trading costs.

By construction, the portfolio holdings of a characteristic-sorted, long-short strategy are

space of currency returns.

in general less sensitive to the time-series variation in forward discounts (conditional expected returns) and covariances than *MV* with its fine-tuned weights. Thus, it is reasonable to expect that the turnover and transaction costs are smaller for equally weighted strategies compared to *MV*. To illustrate this, we provide an intuitive argument. Consider a mean-variance optimized portfolio and an equally weighted portfolio which sorts assets according to expected returns and buys the top quintiles and sells the bottom quintiles. A time-series variation in the covariance matrix will only affect the portfolio weights of the optimized portfolio whereas 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, whereas 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.

Indeed, we document this insight in the data for all strategies. Among the characteristic-sorted strategies, *HML* has the highest average monthly turnover of 1.12 with a TO-NV of 0.56. These values are substantially lower than in the case of *MV*. *RB* and *MOM* have a slightly lower turnover of 0.94 each, and a TO-NV of 0.48. The turnover of *DDOL* is roughly half (0.41), but given that its notional value is also only half of the other strategies, then the TO-NV remains comparable (0.41). The turnover and TO-NV are substantially lower for *VAL* (0.24 and 0.12) and *DOL* (0.03 and 0.03). This confirms our conjecture that the characteristic-sorted portfolios trade less aggressively than *MV*.

The average costs paid by the characteristic-sorted strategies are also considerably lower than *MV*. Nevertheless, the costs are far from negligible for most strategies. When the fund size is large (initial AUM of USD 1 billion) *HML*, *MOM*, and *RB* incur average costs of 7.8%, 5.6%, and 5.3%, respectively. These costs surpass the average before-cost returns, implying that they are making losses after accounting for costs. Even for a fund with initial AUM of USD 100 million, costs erode roughly half of the before-cost returns of these strategies. Thus, investors should be concerned about the price impact of trading.

DOL, *DDOL*, and *VAL* are more resilient to the price impact of trading and the costs are considerably lower, amounting to only 0.04%, 1.9%, and 2.4%, respectively, even in the

case of a large fund with an initial AUM of USD 1 billion.

Next, recall that *MVTC* is designed to optimally reduce costs compared to *MV*, which trades irrespective of costs. Accordingly, we expect that *MVTC* trades less and incurs lower trading costs than *MV*. We further expect that *MVTC* reduces its trading activity (measured in terms of the fraction of the portfolio's notional value) as the AUM becomes larger and the price impact more severe. It is, however, ambiguous how costs of *MVTC* are affected by the fund size. Holding trading activity constant, an increase in the AUM leads to higher costs due to the larger price impact. On the other side, if the reduction in trading activity is sufficiently large, then costs (as a percentage of the AUM) may still decrease as we increase the fund size, although the price impact becomes more severe.

We confirm our predictions in the data. When there is no price impact, *MVTC* has an average monthly turnover of 1.5, an average notional value of 3, implying a TO-NV of 0.5. This is two-thirds of *MV* and comparable to *HML*, *RB*, and *MOM*. The trading activity drastically decreases with an average monthly turnover of 0.7 or 0.3, an average notional value of 2.7 or 2.2, and a TO-NV of only 0.26 or 0.14 when the initial AUM of a fund is USD 100 million or USD 1 billion. This means that *MVTC*'s trading aggressiveness is comparable to *VAL* when the fund size is large and the price impact of trading is severe.

The average costs incurred by *MVTC* are 0.6% per year when there is no price impact, and 0.59% or 0.58% when the initial fund size is USD 100 million or USD 1 billion. With the exception of the *DOL* (which incurs almost zero trading costs) these costs are low compared to any of the other strategies, especially considering a large fund that is subject to a severe price impact. It is interesting that *MVTC* reduces its trading activity to the extent that the AUM appears orthogonal to the average costs incurred by the fund. This means that the increase in the per-USD trading costs due to the more severe price impact is perfectly offset by the reduction in the trading aggressiveness in response to the higher per-USD costs. These costs are comparably small considering that the average before-cost return of *MVTC* ranges between 7% and 10% (depending on the initial AUM). This is in stark contrast to *MV*, for which trading costs are a major concern.

Figure 6 provides an illustration of these findings and plots the cumulative costs incurred

by *MVTC*, *MV*, *HML*, and *MOM*.²¹ It is apparent that the costs of *MV* are always higher than the costs of any other strategy. *HML* and *MOM* incur reasonably low costs when there is no price impact of trading, but costs are sizable for large funds which face a non-trivial price impact. In contrast, *MVTC* manages to keep its costs low at all times, even if the price impact of trading is severe.

Table 1 further reports, on rows seven and eight in panels B and C, the average amount traded by the strategy (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market (row seven), or in absolute terms as millions of USD (row eight). These metrics provide a reference to understand the trading activity of each strategy. In particular, the trade size relative to the daily trading volume in the market allows a cross-strategy comparison.

Most notably, *MV*, *HML*, *RB*, and *MOM* engage in large trades in the order of 1.5-3.3% or even 5.9-10.5% of the average daily trading volume in the market when the AUM of a fund is USD 100 million or 1 billion. This resonates with the finding that the price impact of trading is severe for these strategies. In comparison for the fund size of USD 100 million or 1 billion these percentages of the daily volume in the market are 0.5% or 3.8% for *DDOL* and *VAL*, 0.6% or 1.4% for *MVTC*, and only 0.02% or 0.2% for *DOL*. This illustrates again the resilience of *MVTC* and *DOL* to the problem of transaction costs even when the size of a fund is large.

It is furthermore interesting that the average trade size of *MVTC* is roughly half that of *MV* when measured in USD (USD 33 million versus 62 million for AUM of USD 100 million; USD 90 million versus 152 million for AUM of USD 1 billion), but it is less than a quarter when expressed as a percentage of the daily trading volume in the market (0.6% vs 3.0% for AUM of USD 100 million; 1.4% vs 5.9% for AUM of USD 1 billion). Similar patterns are also observed when we compare *MVTC* to all other strategies. The reason is that *MVTC* (compared to *MV* or other strategies) does not only reduce the overall trade size, it also specifically reduces the order size of expensive currencies. These expensive currencies are illiquid and the daily trading volume in the market is typically low. In contrast, *MVTC* still

²¹We do not plot the other strategies to keep the figure legible.

trades aggressively (and comparable to *MV* or other strategies) in cheap currencies, which are liquid and have a high daily volume. Therefore, the average order size as a percentage of the daily trading volume should be much lower for *MVTC* than *MV* or other strategies, while the USD difference in average order size may be relatively small.

As a final comparison, we consider the trading costs of the US stock market. For that purpose, we use the S&P 500 SPDR ETF (SPY) as a reference.²² We have data for the SPY starting in April 1993. The correlation of monthly excess returns between the US value-weighted stock market portfolio and the SPY is 98% for the sample period from April 1993 to September 2021. Moreover, after scaling the SPY to match the volatility of the US stock market, we find that the US stock market earns on average 0.11% more per year than the SPY. This is close to the SPY’s advertised expense ratio of 0.09%. A buy-and-hold investment in the SPY does not incur additional costs beyond the expense ratio.²³ Accordingly, we use 0.09% as the relevant costs for the US stock market, independent of the AUM. This is a small fraction of the costs incurred by our FX trading strategies.

To sum up, we document that the price impact of FX trading is a major concern for large funds. For most strategies, trading costs erode a large fraction of the average before-cost returns. Specifically, when the initial AUM of a fund is USD 1 billion *MV*, *HML*, *RB*, and *MOM* incur costs that surpass the average before-cost returns, resulting in strategies that, on average, make losses after accounting for costs. Moreover, we observe that *MVTC* significantly reduces its trading activity in response to transaction costs. It incurs substantially lower costs than *MV* and other popular FX trading strategies. This effect intensifies when we consider a higher initial AUM of a fund. As a result *MVTC* efficiently tackles the transaction cost problem, resulting in costs that are a small fraction of its average returns.

4.3 Performance After Costs

As documented in the previous section, when the size of a fund is large, the price impact and trading costs are considerable and in some cases larger than the average before-cost returns.

²²We download SPY data from Yahoo finance, <https://finance.yahoo.com/>.

²³The bid-ask spread and the price impact only have to be paid twice; once at the beginning and once at the end of the investment. Considering an investment period of 35 years these costs are negligible.

Therefore, the price impact of trading has a first-order impact on the profitability of most trading strategies. In the following, we analyze the after-cost performance of our strategies in more detail. The last seven rows in panel A, B, and C in Table 1 report the after-cost Sharpe ratio, the difference between the after-cost Sharpe ratio of a strategy and that of *MVTC* (denoted by ΔSR), the after-cost annualized average return, annualized volatility, monthly skewness, maximum drawdown (MDD) measuring the largest negative return from peak to trough during our sample period, and the total return of the strategy over the sample period.

4.3.1 Mean-Variance Optimized Portfolio *MV*

We start with *MV*, which is the most profitable strategy in terms of before-cost returns. When there are only proportional costs, the after-cost Sharpe ratio of *MV* is 0.97, which is a 12% drop from the before-cost Sharpe ratio of 1.10. This suggests that trading costs are not a major concern for small funds that invest in *MV*.

In contrast, when a fund has an initial AUM of USD 100 million, then the after-cost Sharpe ratio of *MV* drops to 0.33. This is a 70% reduction from the before-cost Sharpe ratio. Interestingly, while *MV* outperforms the US stock market in terms of before-cost returns (1.10 vs 0.58) or after accounting for only proportional costs (0.97 vs 0.57), it performs considerably worse than the stock market (0.33 vs 0.57) when the initial fund size is USD 100 million.²⁴ Even worse, when the initial AUM is USD 1 billion then *MV* turns into a losing strategy with a negative after-cost Sharpe ratio of -0.31 (or an average excess return of -3.25% per year). Accordingly, when a fund is large, the costs due to the price impact of trading have a first-order impact on the profitability of *MV*.

The reduction in the Sharpe ratio of the *MV* is almost entirely explained by the direct implications of costs on the average return. The annual average return decreases from 8.9% before costs to 7.8% after deducting proportional costs, and 2.8% or even -3.25% after accounting for the price impact when a fund has an initial AUM of USD 100 million or 1

²⁴The after-cost Sharpe ratio of 0.57 of the US stock market is almost identical to its before-cost Sharpe ratio of 0.58. The reason is that the US stock market has a low expense ratio of 0.09%. In addition, we consider a buy-and-hold investment in the US stock market, so that there is no price impact of trading and the after-cost Sharpe ratio is independent of the initial AUM.

billion, respectively.

There is also a small effect of the price impact on volatility. The before-cost volatility of MV is 8.1%, and it slightly decreases to 8.0% after accounting for proportional costs. Interestingly, it increases to 8.3% or 10.3% when the initial AUM is USD 100 million or 1 billion, respectively. This increase in volatility negatively affects the after-cost Sharpe ratio but, as mentioned above, this effect is small compared to the direct effect of costs on the average return.

Similar to the effect on volatility, we observe a negative effect of the price impact on crash risks. However, while the effect on volatility is relatively modest, the crash risk implications are quantitatively large and concerning. The price impact leads to a reduction in the monthly skewness from -0.83 and the MDD from -26% (when there are only proportional costs) to -1.23 and -38% when the initial AUM is USD 100 million, or to -2.40 and -157% when the initial AUM is USD 1 billion.

Finally, we analyze the total increase in the AUM from February 1986 to September 2021 for a fund that invests in MV . A fund that started out with an AUM of USD 100 million achieves a total return of 521%. In comparison, a fund with an initial AUM of USD 1 billion has a total return of -27%. Note that these calculations account for the risk-free rate in the USD while the average returns reported above are excess returns. If these funds did not have to pay any transaction costs, then their total return would be 5,350%. Moreover, if they had to pay only proportional costs but there was no price impact of trading, then the total return would be 3,395%. Clearly, the cumulative costs over the 35-year horizon add up significantly, and the price impact of trading has a first-order effect.

Overall, trading costs are a major concern for large funds that face a non-trivial price impact of trading and intend to invest in MV . Transaction costs erode a large fraction of or even the entire expected return, and increase the crash risk at the same time.

4.3.2 Characteristic Sorted Long-Short Strategies

Next, we turn to the six characteristic-based long-short strategies HML , RB , DOL , $DDOL$, VAL , and MOM . In the previous section we established that, by construction, these strate-

gies trade less aggressively and incur lower costs than *MV*. The proportional costs borne by these strategies are small, and after-cost Sharpe ratios are close to the respective before-cost Sharpe ratios. As such, a small fund (which does not face any price impact of trading) does not need to be concerned about trading costs. This is the standard conclusion from the literature, which only assesses the effect of proportional costs on the profitability of trading strategies.

In contrast, the after-cost Sharpe ratios substantially decrease when the initial AUM of the fund is large. The effect is the largest for *HML*, *RB* and *MOM*. In the case of *HML*, the before-cost Sharpe ratio is 0.72, while the after-cost Sharpe ratio decreases to 0.67 when there is no price impact, and to 0.39 or even -0.20 when the initial AUM is USD 100 million or USD 1 billion, respectively. Similarly, for *RB* and *MOM* the corresponding values are 0.74 and 0.30 before cost, 0.68 and 0.27 after cost without price impact, 0.52 and 0.12 when the initial AUM is USD 100 million, and -0.03 and -0.34 when the initial fund size is USD 1 billion. Also similarly to *MV*, all three strategies underperform the US stock market (which has an after-cost Sharpe ratio of 0.57) when the fund size is USD 100 million (albeit *RB* performs only slightly worse than the stock market), and all three strategies are unprofitable (with negative average returns) after accounting for costs when the fund has an initial AUM of USD 1 billion.

The implications of the price impact of trading are less severe but still concerning for the *DDOL* and *VAL*. These strategies have before-cost Sharpe ratios of 0.35 for *DDOL* and 0.46 for *VAL*. Proportional costs reduce the Sharpe ratios of *DDOL* and *VAL* to 0.32 and 0.44, while the price impact of a fund with an initial AUM of USD 100 million or USD 1 billion further reduces these ratios to 0.29 or 0.10 for *DDOL* and 0.40 or 0.11 for *VAL*. While the after-cost Sharpe ratios remain positive even when the initial AUM reaches USD 1 billion, the performance is not attractive to an investor and these strategies underperform the US stock market by a large margin.

The costs incurred by *DOL* are negligible irrespective of the fund size, and the after-cost Sharpe ratio of *DOL* is virtually identical to the before-cost Sharpe ratio. As such, the *DOL* is very resilient to trading costs. However, we acknowledge that the *DOL* is not an

attractive strategy to begin with.

It is clear that the intuitive advantage of characteristic-based long-short strategies to mitigate costs (as the investment policy is less sensitive to state variables than MV) does not work well. The reason is twofold. First, the before-cost performance of these strategies is much worse than the performance of MV , and we would need a lot of cost savings to make up for this. Second, although the characteristic-based long-short strategies are less sensitive to the price impact of trading than MV , they are not resilient enough. That is, when the size of a fund is large, they still suffer greatly from the price impact of trading and incur large transaction costs.

Similarly to MV , for the characteristic-based long-short strategies we also observe an increase in crash risk due to the price impact, albeit this effect appears only relevant when the fund size is sufficiently large. For instance, the skewness and MDD of HML are -0.77 and -32%, respectively, when there is no price impact, and these values decline to -0.83 and -33% when the initial AUM is USD 100 million, and to -0.91 and -115% when the initial AUM is USD 1 billion. The crash risk for a fund with an initial AUM of USD 100 million is almost identical to a small fund with no price impact. However, the crash risk rapidly increases as we increase the initial fund size from USD 100 million to USD 1 billion. The effect is similar for the other five characteristic-based long-short strategies.

To sum up, for large funds, transaction costs have first-order implications on the profitability of characteristic-based long-short strategies.

4.3.3 Mean-Variance-Transaction-Cost Optimized Portfolio $MVTC$

Finally, we investigate the after-cost performance of $MVTC$, which is designed to optimally tackle costs. In the previous sections, we showed that, compared to the high before-cost returns, average annual costs are relatively low even when the fund size is large. Accordingly, the Sharpe ratios are comparable before and after costs. When there are only proportional costs, the after-cost Sharpe ratio of $MVTC$ is 0.99, which is a 7% drop from the before-cost Sharpe ratio of 1.06. For a fund with an initial AUM of USD 100 million or USD 1 billion, the after-cost Sharpe ratios are 0.86 or 0.72, respectively, which is an 8% or 11% decrease

from the before-cost Sharpe ratios of 0.93 or 0.81. This suggests that trading costs are not a major concern for investors in *MVTC*, independent of the fund size.

Consistent with the theory, we also observe that *MVTC* always outperforms *MV* after accounting for costs, despite the fact that *MV* has a higher before-cost Sharpe ratio. In the case of only proportional costs, the difference between the after-cost Sharpe ratios of *MV* and *MVTC*, i.e., $\Delta SR = SR^{MV} - SR^{MVCT} = -0.02$, is small and statistically insignificant.²⁵ However, when the initial AUM is USD 100 million or USD 1 billion, these differences (i.e., ΔSR) are -0.52 or -1.04, respectively, and statistically significant at the 1% level. The differences in Sharpe ratios between all other strategies and the *MVTC* are always (i.e., independent of the fund size) economically large and statistically significant, as indicated by ΔSR in Table 1. Furthermore, in comparison to the US stock market, *MVTC* delivers a much more attractive Sharpe ratio.

Figure 5 further illustrates the after-cost outperformance of *MVTC*. We plot the cumulative after-cost excess returns of *MVTC*, *MV*, *HML*, and *MOM*.²⁶ *MVTC* performs well independent of the fund size. We generally observe a weakening of the FX strategies after October 2008.²⁷ However, *MVTC* is still profitable as illustrated by the upward trajectory of the solid black lines in figure 5. Moreover, table A4 shows that post-GFC the after cost Sharpe ratio of *MVTC* is still 0.68 (no price impact), 0.56 (initial AUM of USD 100 million), and 0.31 (initial AUM of USD 1 billion).

The after-cost cumulative returns of *MV* are almost indistinguishable from *MVTC* when there is no price impact of trading. In contrast, when the fund size is large and there is a non-trivial price impact of trading the outperformance of *MVTC* over *MV* is obvious. When the initial AUM is USD 100 million, *MV* performs relatively well for the first 10 years but is essentially unprofitable after 1997. Even worse, when the initial AUM is USD 1 billion, *MV* makes large losses, especially in the 1990s, reports some small gains from 2000 to 2006, and is a losing strategy ever since.

²⁵To test for significance, we employ the test proposed by Ledoit and Wolf (2008), which uses block bootstrapping and is robust to heteroskedasticity and cross- and auto-correlation. We choose a block size of 5 observations for the block bootstrapping.

²⁶We do not plot the other strategies to keep the figure legible.

²⁷We provide a more detailed analysis of pre- and post-2008 performance in robustness tests in Section 6 and Internet Appendix A.1.

MVTC also outperform *HML* and *MOM* by a large margin. *HML* performs well throughout the sample when there is no price impact of trading. It performs relatively well until 2008 when the initial AUM is USD 100 million, but after that the performance is weak. Moreover, when the initial AUM is 1 billion, *HML* is a losing strategy almost throughout the entire sample. *MOM* performs similarly to *HML* until 2008. After 2008, even its before-cost average return turns negative, and it is unprofitable before and after costs.

Compared to the other strategies, we do not observe any worrisome effects of the fund size and trading costs on the risk of *MVTC*. The before-cost and after-cost volatilities are almost identical. The skewness and MDD are unaffected by the fund size. When there is no price impact, the monthly skewness of *MVTC* is -0.77 and the MDD is -28%. When the initial fund size is USD 100 million or 1 billion, the skewness coefficients are -1.01 or -0.92, respectively, and the MDD are -25% or -29%. Moreover, this insight is also confirmed in Figure 5. Essentially, there is only a trend adjustment in the cumulative returns of *MVTC* when we compare different fund sizes.

Finally, a fund that started out with an AUM of USD 100 million in February 1986 yields a total return of 2,076%. In comparison, a fund with an initial AUM of USD 1 billion has a total return of 1,193%. These total returns are a lot more attractive than the total returns of *MV*, i.e., 521% or -27%.

Overall, *MVTC* efficiently reduces its trading activity to save costs while maintaining a strong performance. First, trading costs are much less problematic for *MVTC* than all other strategies. Second, *MVTC* is resilient to transaction costs even when the price impact of trading is severe.

Finally we note that *MVTC* is the solution in a single-period model. It does not take into account potentially interesting dynamics in a multi-period model. It is not obvious at the outset that *MVTC* performs well out-of-sample when implemented in the data with many trading dates. It is an important empirical finding that (i) *MVTC* performs significantly better than all other strategies, (ii) achieves a high after-cost Sharpe ratio, and (iii) is resilient even when the size of a fund is large and the price impact of trading is severe.

4.3.4 Capacities of Strategies: Effect of the Fund Size

Next, we vary the size of a fund to further investigate the resilience of *MVTC* to the price impact of trading. Table 2 reports the maximum initial AUM (top panel) and the maximum terminal AUM (bottom panel) that a fund can have such that it delivers an after-cost Sharpe ratio that is larger than the threshold \overline{SR} . We consider threshold values \overline{SR} ranging between 0 and 0.9. Confirming the findings, in Table 1, when the initial AUM of a fund is less than USD 125 million or 1.25 billion, *MVTC* achieves an after-cost Sharpe ratio of at least 0.85 or at least 0.7, respectively.

Recall that the US stock market has an after-cost Sharpe ratio of 0.57. Table 2 shows that a fund investing in *MVTC* performs better than the US stock market as long as the initial AUM does not surpass USD 3 billion (or equivalently an AUM of USD 25.18 billions in September 2021). Moreover, *MVTC* only turns unprofitable (i.e., it earns a negative after-cost average return) once the initial AUM is larger than USD 1000 billions. These numbers demonstrate the *MVTC* has a large capacity and is resilient to even a very large fund size and a severe price impact of trading.

In contrast, all other strategies perform much worse. For instance, *MV* only outperforms the US stock market when the initial AUM is smaller than USD 40 million. Moreover, *MV* is unprofitable as soon as the initial AUM surpasses USD 300 million. Similarly, *HML* and *RB* underperform the US stock market or even turn completely unprofitable once the initial AUM is larger than USD 30 million or USD 70 million, respectively, or USD 490 million or USD 910 million, respectively. As such, these strategies have relatively small capacity, especially if we take the US stock market as the benchmark.

Figure 7 graphically illustrates the insights of Table 2. We plot the after-cost Sharpe ratio against the initial AUM. All strategies but *MVTC* quickly turn unprofitable. *MVTC* clearly stands out as the only currency trading strategy that is resilient to a severe price impact of trading.

We conclude that transaction costs, and specifically the price impact of trading, are a first-order concern for currency investors. Costs quickly erode the expected returns of popular strategies in the literature, and these strategies are left unprofitable. Thus, addressing

trading costs in the construction of the portfolio is essential to improve the overall performance. *MVTC* efficiently reduces trading activity to mitigate costs. Accordingly, it is resilient to even a severe price impact of trading.

5 Rules-of-Thumb to Address Transaction Costs

In this section, we investigate whether intuitive rules-of-thumb are helpful in mitigating transaction costs. Specifically, we consider two rules-of-thumb: (i) trading at a lower frequency, or (ii) removing high-cost currencies from our asset universe. Such rules reduce costs to some extent, but they are inefficient as they have significant negative implications on the performance of our strategies. Thus, our findings advise against the use of intuitive rules-of-thumb, as deviations from the optimal (before-cost) balance between expected return and risk have adverse consequences for the performance. If we invest in *MVTC*, which accounts for costs in portfolio optimization, the out-of-sample performance is significantly better. Moreover, *MVTC* is far more resilient than the other analyzed strategies when we consider large funds that face a severe price impact.

5.1 Trading at Lower Frequencies

It is intuitive that frequent trading leads to more turnover and higher transaction costs. Therefore, we investigate how the performance is affected if we reduce the trading frequency from monthly to quarterly in the hope of mitigating costs. Table 3 reports the quarterly performances of our eight strategies. The first two columns provide the before-cost and after-cost Sharpe ratios. We further report the difference between the Sharpe ratios of the strategies and *MVTC* in the third column denoted by ΔSR . The last two columns report the quarterly turnover and average annual transaction costs.

We start our discussion with *MV* and the six characteristic-based long-short strategies (*HML*, *RB*, *DOL*, *DDOL*, *VAL*, *MOM*). As expected, trading at the quarterly frequency significantly reduces costs. In general, for the seven strategies, costs are roughly cut in half, or even more when price impact is considered. This reflects well on our simple and intuitive

cost mitigation rule.

On the other hand, there is also a negative effect of quarterly trading. Generally, we find lower before-cost Sharpe ratios when we implement our strategies at the quarterly instead of the monthly frequency. This is due to both economic and statistical reasons. From an economic point of view, if a trading signal has only short-term predictive power then a strategy based on this signal is more (less) likely to be successful when we trade at a high (low) frequency. Consistent with this explanation, we observe lower average returns when we implement our strategies at the quarterly frequency. From a statistical point of view, the monthly variance is less than a third of the quarterly variance if monthly returns are positively auto-correlated (Lo, 2002). Consistent with this explanation, we observe that the annualized volatility of the quarterly strategies is higher compared to monthly rebalancing.

When there is no price impact, the negative effect on the before-cost performance outweighs the benefits of mitigating costs. That is, the after-cost Sharpe ratios of all strategies (except for *VAL*) are higher when we implement them at the monthly instead of at the quarterly frequency. The reason is that proportional costs are relatively small and, as such, are not a first-order issue. Accordingly, a small fund is better off investing at the monthly frequency.

The same is not always true for larger funds with a non-trivial price impact. In this case, costs are a major concern and cost reduction measures are relevant. For instance, when the initial AUM is USD 100 million, *MV* achieves an after-cost Sharpe ratio of 0.55 when implemented at the quarterly frequency, while it is only 0.33 when we trade monthly. Another example is *RB*. When the initial AUM is USD 1 billion then the monthly *RB* is unprofitable, with an after-cost Sharpe ratio of -0.03, while the quarterly version of *RB* is profitable, with a Sharpe ratio of 0.29. These examples illustrate that trading at a lower frequency can be helpful in mitigating costs and increasing the resilience of strategies to the price impact of trading. However, this is not generally true for all strategies. For instance, *MOM* performs consistently worse when implemented at the quarterly frequency.

The after-cost performance of the quarterly versions of *MV* and the six characteristic-based long-short strategies are relatively poor when the initial AUM is USD 1 billion. The

best-performing and most resilient among the seven strategies is *RB*, with an after-cost Sharpe ratio of 0.29. This is only half of the Sharpe ratio achieved by investing in the US stock market. Moreover, as in our baseline analysis we find that *MV* and *MOM* are unprofitable (with negative average returns) when implemented at the quarterly frequency. Therefore, our previous insight remains unchanged. The price impact of trading is a major concern for sizable funds, as costs quickly erode returns and, as a result, popular strategies in the literature perform poorly after accounting for costs.

Next, we analyze the effect of the trading frequency on *MVTC*. We observe cost savings when we trade quarterly instead of monthly. This effect is, however, quantitatively small, especially relative to the high average returns earned by *MVTC*. The reason is that *MVTC* efficiently reduces its trading activity and costs whether we implement it at a monthly or a quarterly frequency. As such, there is no benefit from overlaying the strategy with an intuitive rule-of-thumb in an attempt to manage costs. The (before-cost) performance somewhat suffers as we trade less frequently, and the after-cost Sharpe ratio of the quarterly *MVTC* is slightly lower than its monthly counterpart.

Finally, a comparison between strategies clearly suggests that it is best to invest in the monthly *MVTC*. The after-cost Sharpe ratio of *MVTC* always significantly dominates all other strategies (with the exception of *MV* when there is no price impact; in this case, *MV* and *MVTC* are almost indistinguishable). Moreover, *MVTC*, and in particular the monthly version, is the only strategy that (i) is resilient to even a large fund size with a severe price impact, and (ii) consistently outperforms the US stock market.

To sum up, we find that trading at a lower frequency decreases the turnover and transaction costs, but changing the trading frequency also has a negative first-order effect on performance. Overall, it is not advisable to trade at the quarterly frequency. Instead, investors are better off trading at the monthly frequency and accounting for costs in the construction of the strategy, as suggested by *MVTC*.

5.2 Removing Currencies with High Transaction Costs

Our second rule-of-thumb builds on the intuition that transaction costs decrease if we restrict trading to currencies with low costs. We focus on the monthly trading frequency and either (i) restrict trading to the 13 developed currencies (as defined in Section 2.1), or (ii) sequentially eliminate the most expensive currencies from the set of admissible assets in every month t .

Table 4 shows the performance and the trading costs of our strategies when we invest in only the 13 developed currencies. The first two columns report the before-cost and after-cost Sharpe ratios. The third column reports the difference between the after-cost Sharpe ratios of all strategies and the *MVTC*. The last two columns report the monthly turnover and annual costs.

We find that the costs and after-cost Sharpe ratios are generally comparable if we use the 13 developed or the full set of 26 currencies. The exceptions are *VAL*, which improves, and *MOM*, which performs substantially worse when we restrict the set to the 13 developed currencies. While developed currencies are cheaper to trade than emerging currencies, reducing the asset universe by half (from 26 to 13 currencies) implies that the portfolio weights are twice the size. In turn, the increase in the portfolio concentration leads to larger trade orders or a more severe price impact, albeit in only developed and arguably more liquid currencies. It turns out that it does not matter whether we implement our strategies using only developed or the larger set of developed and emerging currencies; the after-cost performances are comparable.

To sum up, limiting the portfolio to a set of developed currencies is not an efficient rule to mitigate the issue of transaction costs. Popular strategies in the literature are still sensitive to the price impact of trading, and costs quickly erode the average returns, leaving popular strategies unprofitable. *MVTC* (implemented either in the set of 13 developed or all 26 currencies) always earns the highest after-cost Sharpe ratio, and is the only strategy that is resilient when the size of a fund is large and the price impact is severe.

An alternative rule-of-thumb is to limit trading to only the cheapest currencies in every month t . Table 5 reports the Sharpe ratios before and after costs for *MV*, *MVTC* and

HML in columns 1 to 3 and 4 to 6.²⁸ The first row reports the performance of the strategies constructed from the full set of 26 currencies. The second row reports the results when we reduce our set of admissible currencies by one currency, i.e., in month t we drop the currency with the highest median transaction cost over the past 6 months.²⁹ Each subsequent row reduces the set of admissible currencies by an additional currency, i.e., in month t row i removes the $i - 1$ currencies with the highest median costs over the past 6 months. Note that the data availability changes through time and we do not have data for all 26 currencies in every month. The maximum number of currencies we can drop is 10, in order to have at least two currencies with which to construct a portfolio every month.

For *MV* and *HML*, the differential between the before-cost and after-cost Sharpe ratios narrows as we remove expensive currencies (difference between columns 1 and 4 respectively 3 and 6). This indicates that costs are decreasing, which confirms that the rule-of-thumb mitigates costs as expected. For *MVTC*, however, we do not observe a similar monotonic decrease in the differences between before-cost and after-cost Sharpe ratios. This implies that *MVTC* efficiently tackles costs and does not benefit from overlaying it with the intuitive cost-reducing rule-of-thumb.

More importantly, we document that the before-cost performance of all three strategies rapidly worsens as we remove currencies. *MV*, *MVTC* and *HML* achieve before-cost Sharpe ratios of 1.10, 1.06 and 0.72, respectively, in our baseline analysis when we use all 26 currencies. These values almost monotonically decrease as we remove currencies. In the extreme case when we remove the 10 most expensive currencies, the Sharpe ratios are only 0.29, 0.29 and 0.41.³⁰ The reduction in the before-cost performance is of first-order importance, and outweighs the cost savings. For *MVTC*, the after-cost Sharpe ratio is consistently decreasing in the number of removed currencies. Thus, an investor should consider the full set of currencies when constructing *MVTC*. The same pattern holds for *MV* and *HML* when there is either no price impact or the initial AUM is USD 100 million. When the fund size is large

²⁸We limit our attention to *MV*, *MVTC*, and *HML* for brevity. The results for the other strategies are similar.

²⁹Changing the rolling window length does not affect our results.

³⁰When we remove 10 currencies, for *MVTC* the before-cost Sharpe ratio is 0.29 when there is no price impact, 0.28 when the initial AUM is USD 100 million, and 0.30 when the initial AUM is USD 1 billion.

(with an initial AUM of USD 1 billion) the after-cost Sharpe ratios of *MV* and *HML* are not decreasing in the number of dropped currencies, but they are consistently negative and, as such, these strategies are unprofitable in any case. Again, there is no value for investors to follow a rule-of-thumb to remove high cost currencies from the asset universe.

We conclude that the rule-of-thumb to remove high cost currencies is inefficient. The exclusion of expensive currencies reduces costs, but it also leads to a decrease in before-cost performance. The latter effect generally dominates. The most sensible approach (which delivers the best out-of-sample performance) is the cost-optimized portfolio *MVTC* implemented in the unrestricted set of 26 currencies.

6 Robustness and Extensions

We perform several robustness checks and extensions, which we report in detail in the Internet Appendix [A](#). It is reassuring that our main results continue to hold in all robustness analyses. In the following we provide a brief overview.

A longer sample period generally provides more accurate estimates of expected returns. Thus, we use the results in Section [4](#) as our baseline. In Appendix [A.1](#), we further show that our results are robust and continue to hold in the following sub-samples: (i) pre- and post-the 2008 global financial crisis (GFC), (ii) post-introduction of the Euro (January 1999), and (iii) in and out of NBER-defined US recessions. These robustness checks are important to ensure that our results are generally true, and not driven by a specific sub-sample under special economic conditions. First, the GFC defines a regime shift, as interest rates across many currencies were held close to zero for over a decade afterwards and, in particular, *MOM* and *DDOL* were very profitable before accounting for costs pre-GFC but performed poorly afterwards. Noticeably, even in the pre-GFC sample after accounting for costs due to the price impact of a large fund, *MOM* and *DDOL* (unlike *MVTC*) perform substantially worse than the US stock market. Therefore, our results of the baseline analysis are not driven by the poor post-GFC (before-cost) performance of these strategies. Second, we show that our results are unaffected by the introduction of the Euro. This is important as the Euro has

replaced many currencies and as such has crucially altered the investment opportunity set. Finally, we document that our results hold in and out of NBER-defined recession periods. Most strikingly, *MVTC* performs extremely well and is resilient even during NBER-defined US recessions.

In Appendix [A.2](#) we show that our results are robust and continue to hold for alternative estimations of the price impact functions. The choice of method used to estimate cost functions has little effect on the after-cost profitability of our strategies and, therefore, our results are not driven by the particular estimation of the costs. First, we use the time-series average of our Amihud measure to construct a time-invariant price impact function. This mitigates the concern that the results might be driven by certain episodes of high illiquidity. Second, in the baseline analysis we keep the trading volume fixed through time and all time-series variation in the Amihud measure stems from the variation in absolute returns. This is to avoid potentially large errors when we attempt to extrapolate trading volume data in the time-series from 2012 back to 1986. In robustness checks, we do a more elaborate extrapolation back in time and allow for time-varying volume. As expected, the time-series of the extrapolated volume data is decreasing over time. In turn, this implies that the price impact in our baseline analysis is smaller and more conservative than in this robustness test. Third, we use trading volume data from the Triennial FX Survey of the BIS instead of the CLS data in our baseline analysis. An advantage of the BIS data is that it goes back until 1992, and therefore allows for a more accurate estimation of the Amihud measure before 2012. On the other hand, we need to interpolate between dates as the BIS data is only available for average trading volume in the month of April and only every three years. Finally, we cut all costs in half to conservatively address the concerns that our setup cannot handle indirect trades via the currency triangle. The idea is that it might be cheaper (and, in the best-case scenario, reduce costs by up to a half) to trade directly between two non-USD currencies rather than twice against the USD.³¹ Halving the costs implies that the capacity of every strategy approximately doubles. While this provides a boost to the strategies, the capacities of *MV* and the six characteristic-based long-short strategies are still small, and they all perform poorly when the initial AUM of a fund reaches USD 1 billion.

³¹Details to motivate this robustness test are in the Internet Appendix [A.2](#).

A potential concern about *MVTC* is that it instantly adjusts its portfolio holdings when there is a major shock to the investment opportunity set, while in reality this may not be possible. Therefore, in Appendix [A.3.1](#), we implement an event study and illustrate that the crash risk exposure of *MVTC* does not substantially increase even if it is impossible to unwind its risk exposure in turbulent periods.

Finally, in Appendix [A.3.2](#) we show that reasonable choices of the λ parameter for the *MVTC* strategy do not alter our results.

7 Conclusion

We show that, despite the high trading volume in FX markets, transaction costs have a first-order impact on the performance of currency portfolios. Specifically, while proportional costs are a minor concern, the price impact of trading is a first-order problem for popular FX trading strategies in the literature. Costs due to the price impact quickly erode average returns, and cause these strategies to underperform the US stock market or, even worse, turn them unprofitable. This is an important contribution, as the FX market literature has hitherto considered the profitability of currency strategies in the absence of price impact.

We investigate whether the issue of transaction costs can be mitigated if we (i) trade at a low frequency, or (ii) restrict trading to only liquid and low cost currencies. We confirm that these intuitive rules-of-thumb reduce transaction costs, but we also find that they are inefficient, as deviations from the optimal before-cost balance between expected return and risk have adverse consequences on performance.

In contrast, if we optimally account for costs in the construction of our portfolio (*MVTC*), we find that the after-cost performance is attractive, and *MVTC* consistently outperforms the US stock market. *MVTC* is resilient to even a large fund size and a severe price impact of trading. Accordingly, transaction costs are a first-order problem, and it is vital for investors to optimally tackle them in the portfolio construction in order to obtain a robustly profitable strategy.

Our paper further contributes to the literature on liquidity in FX markets. Most im-

portantly, we propose novel extrapolations of the realized Amihud measure introduced by [Ranaldo and Santucci de Magistris \(2022\)](#). This allows us to obtain estimates of the price impact for 26 instead of 12 currencies (against the USD) and for a sample going back to 1986 instead of 2012. In particular, we show that trading volume can be accurately extrapolated in the cross-section and in the time-series using the bid-ask spread as a conditioning variable.³²

References

- Ackermann, Fabian, Walt Pohl, and Karl Schmedders, 2016, Optimal and Naive Diversification in Currency Markets, *Management Science* .
- Amihud, Yakov, 2002, Illiquidity and stock returns: cross-section and time-series effects, *Journal of financial markets* 5, 31–56.
- Asness, Clifford S., Tobias J. Moskowitz, and Lasse H. Pedersen, 2013, Value and Momentum Everywhere, *Journal of Finance* 68, 929–985.
- Back, K., 1993, Asymmetric information and options, *Review of Financial Studies* 6.
- Baz, Jamil, Frances Breedon, Vasant Naik, and Joel Peress, 2001, Optimal Portfolios of Foreign Currencies, *The Journal of Portfolio Management* 28, 102–111.
- Bechtel, Alexander, Angelo Ranaldo, and Jan Wrampelmeyer, 2023, Liquidity risk and funding cost, *Review of Finance* 27, 399–422.
- Bilson, John F. O., 1984, Purchasing Power Parity as a Trading Strategy, *Journal of Finance* 39, 715–724.
- Bräuer, Leonie, and Harald Hau, 2023, Can time varying currency risk hedging explain exchange rates?, Working paper.
- Breedon, Francis, Dagfinn Rime, and Paolo Vitale, 2016, Carry trades, order flow, and the forward bias puzzle, *Journal of Money, Credit and Banking* 48, 1113–1134.
- Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 2011, Carry trade and momentum in currency markets, *Annual Review of Financial Economics* 3, 511–535.
- Camanho, Nelson, Harald Hau, and Helene Rey, 2022, Global rebalancing and the exchange rate, *Review of Financial Studies* 35, 5228–5274.
- Cespa, Giovanni, Antonio Gargano, Steven J Riddiough, and Lucio Sarno, 2022, Foreign exchange volume, *The Review of Financial Studies* 35, 2386–2427.

³²The results of the time-series extrapolation are in Section [A.2](#) in the Internet Appendix.

- Chernov, Mikhail, Magnus Dahlquist, and Lars Lochstoer, 2023, Pricing currency risks, *Journal of Finance* 78, 693–730.
- Chinn, Menzie D., and Hiro Ito, 2006, What Matters for Financial Development? Capital Controls, Institutions, and Interactions, *Journal of Development Economics* 81, 163–192.
- Corwin, Shane A., and Paul Schultz, 2012, A Simple Way to Estimate Bid-Ask Spreads from Daily High and Low Prices, *Journal of Finance* 67, 719–759.
- Daniel, Kent, Robert J. Hodrick, and Zhongjin Lu, 2017, The Carry Trade: Risks and Drawdowns, *Critical Finance Review* 6, 211–262.
- 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, Alberto Martin-Utrera, Francisco J. Nogales, and Raman Uppal, 2020, A Transaction-Cost Perspective on the Multitude of Firm Characteristics, *Review of Financial Studies* 33, 2180–2222.
- Dybvig, Phillip H., and Luca Pezzo, 2020, Mean-Variance Portfolio Rebalancing with Transaction Costs, Working paper, Washington University.
- Evans, Martin DD, and Richard K Lyons, 2002, Order flow and exchange rate dynamics, *Journal of political economy* 110, 170–180.
- Evans, Martin DD, and Richard K Lyons, 2008, How is macro news transmitted to exchange rates?, *Journal of Financial Economics* 88, 26–50.
- Evans, Martin D.D., and Dagfinn Rime, 2016, Order flow information and spot rate dynamics, *Journal of International Money and Finance* 69, 45–68.
- Fischer, Andreas M., and Angelo Ranaldo, 2011, Does fomc news increase global fx trading?, *Journal of Banking Finance* 35, 2965–2973.
- Froot, Kenneth A., and Tarun Ramadorai, 2005, Currency Returns, Intrinsic Value, and Institutional-Investor Flows, *Journal of Finance* 60, 1535–1566.
- Hasbrouck, Joel, 1991a, Measuring the Information Content of Stock Trades, *The Journal of Finance* 46, 179–207.
- Hasbrouck, Joel, 1991b, The Summary Informativeness of Stock Trades: An Econometric Analysis, *The Review of Financial Studies* 4, 571–595.
- Hasbrouck, Joel, and Richard M Levich, 2019, Fx liquidity and market metrics: New results using cls bank settlement data, Working paper.

- Hasbrouck, Joel, and Richard M. Levich, 2021, Network structure and pricing in the fx market, *Journal of Financial Economics* 141, 705–729.
- Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, 2021, Discriminatory pricing of over-the-counter derivatives, *Management Science* 67, 6629–7289.
- Hau, Harald, Massimo Massa, and Joel Peress, 2010, Do demand curves for currencies slope down? evidence from the msci global index change, *Review of Financial Studies* 23, 1681–1717.
- Hautsch, Nikolaus, and Stefan Voigt, 2019, Large-scale portfolio allocation under transaction costs and model uncertainty, *Journal of Econometrics* 212, 221–240.
- Karnaukh, Nina, Angelo Ranaldo, and Paul Soederlind, 2015, Understanding FX Liquidity, *The Review of Financial Studies* 28, 3073–3108.
- Korsaye, Sofonias Alemu, Fabio Trojani, and Andrea Vedolin, 2023, The global factor structure of exchange rates, *Journal of Financial Economics* 148, 21–46.
- Krohn, Ingomar, Philippe Mueller, and Paul Whelan, 2023, Foreign exchange fixings and returns around the clock, *Journal of Finance* forthcoming.
- Krohn, Ingomar, and Vladyslav Sushko, 2022, Fx spot and swap market liquidity spillovers, *Journal of International Money and Finance* 120, 1–18.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53.
- Kyle, Albert S., Anna A. Obizhaeva, and Yajun Wang, 2018, Smooth trading with overconfidence and market power, *Review of Economic Studies* 85.
- Lacava, Demetrio, Angelo Ranaldo, and Paolo Santucci de Magistris, 2023, Realized Illiquidity, Working paper.
- Ledoit, Olivier, and Michael Wolf, 2008, Robust Performance Hypothesis Testing with the Sharpe Ratio, *Journal of Empirical Finance* 15, 850–859.
- Leland, Hayne E., 1985, Option pricing and replication with transactions costs, *The Journal of Finance* 40, 1283–1301.
- Liu, Hong, 2004, Optimal Consumption and Investment with Transaction Costs and Multiple Risky Assets, *Journal of Finance* 59, 289–338.
- Liu, Sining, Thomas Maurer, Andrea Vedolin, and Yaoyuan Zhang, 2023, Dollar and Carry Redux, Working paper.
- Lo, Andrew W., 2002, The Statistics of Sharpe Ratios, *Financial Analysts Journal* 58, 34–52.
- Lustig, Hanno, Nick Roussanov, and Adrien Verdelhan, 2011, Common Risk Factors in Currency Returns, *Review of Financial Studies* 24, 3731–3777.

- Lustig, Hanno, Nick Roussanov, and Adrien Verdelhan, 2014, Countercyclical Currency Risk Premia, *Journal of Financial Economics* 111, 527–553.
- Lustig, Hano, and Adrien Verdelhan, 2007, The Cross-section of Foreign Currency Risk Premia and Consumption Growth Risk, *American Economic Review* 97, 89–117.
- Lyons, R. K.I, 2001, The Microstructure Approach to Exchange Rates, in *The Microstructure Approach to Exchange Rates* (Cambridge, MA: MIT Press).
- Mancini, Lorian, Angelo Ranaldo, and Jan Wrampelmeyer, 2013, Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums, *Journal of Finance* 68, 1805–1841.
- Maurer, Thomas A., Thuy-Duong To, and Ngoc-Khanh Tran, 2022, Pricing Implications of Covariances and Spreads in Currency Markets, *Review of Asset Pricing Studies* 12, 336–388.
- Maurer, Thomas A., Thuy-Duong To, and Ngoc-Khanh Tran, 2023, Market Timing and Predictability in FX Markets, *Review of Finance* 27, 223–246.
- 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, 2012a, Carry Trades and Global Foreign Exchange Volatility, *Journal of Finance* 67, 681–718.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012b, Currency Momentum Strategies, *Journal of Financial Economics* 106, 620–684.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2016, Information flows in foreign exchange markets: Dissecting customer currency trades, *Journal of Finance* 71, 601–634.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2017, Currency Value, *Review of Financial Studies* 30, 416–441.
- Olivares-Nadal, Alba V., and Victor DeMiguel, 2016, A Robust Perspective on Transaction Costs in Portfolio Optimization, *Operation Research* 66, 733–739.
- Pezzo, Luca, Lei Wang, and Duygu Zirek, 2023, Mean-Variance Market Timing the U.S. Stock Market, Technical report.
- Ranaldo, Angelo, and Paolo Santucci de Magistris, 2022, Liquidity in the global currency market, *Journal of Financial Economics* 146, 859–883.
- Ranaldo, Angelo, and Fabricius Somogyi, 2021, Asymmetric information risk in fx markets, *Journal of Financial Economics* 140, 391–411.
- Roussanov, Nikolai, and Xiaoliang Wang, 2023, Following the fed: Limits of arbitrage and the dollar, Working paper.

Sarno, Lucio, Federico Nucera, and Gabriele Zinna, 2023, Currency Risk Premia Redux, *Review of Financial Studies* forthcoming.

Soderlind, Paul, and Fabricius Somogyi, 2023, FX Liquidity Risk and Carry Trade Premia, *Management Science* forthcoming.

Somogyi, Fabricius, 2022, Dollar Dominance in FX Trading, Working paper.

Yoshimoto, Atsushi, 1996, The mean-variance approach to portfolio optimization subject to transaction costs, *Journal of the Operations Research Society of Japan* 31, 99–117.

Table 1: Performance of FX Strategies

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL A: NO PRICE IMPACT								
Before cost SR	1.10	1.06	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	8.40	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	7.89	8.43	6.88	7.79	7.69	6.97	8.68
Mean Costs (%)	1.10	0.60	0.36	0.40	0.02	0.20	0.11	0.30
Turnover	2.36	1.51	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	3.02	2.00	2.00	1.00	0.99	2.00	2.00
SR	0.97	0.99	0.67	0.68	0.12	0.32	0.44	0.27
ΔSR	-0.02	-	-0.32*	-0.31*	-0.87***	-0.67**	-0.55**	-0.72***
Mean (%)	7.77	7.79	5.68	4.70	0.95	2.49	3.08	2.32
Vol (%)	8.04	7.86	8.43	6.88	7.79	7.70	6.96	8.68
Skew	-0.83	-0.77	-0.77	-0.68	-0.44	-0.33	-0.18	-0.33
MDD (%)	-26.43	-28.17	-32.27	-31.43	-39.09	-17.69	-22.40	-50.31
Total return	33.95	34.41	16.05	11.73	2.45	4.89	6.37	4.33
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	1.10	0.93	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	6.93	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	7.43	8.43	6.88	7.79	7.69	6.97	8.68

Continued on next page

Table 1 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	6.04	0.59	2.68	1.47	0.02	0.42	0.44	1.54
Turnover	2.36	0.70	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	2.67	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	2.97	0.63	3.26	1.53	0.02	0.50	0.53	1.74
Avg USD Trade Size (millions)	62.13	33.43	62.36	31.26	0.41	9.42	9.69	35.00
SR	0.33	0.86	0.39	0.52	0.12	0.29	0.40	0.12
ΔSR	-0.52***	-	-0.46***	-0.33*	-0.73***	-0.56**	-0.46**	-0.73***
Mean (%)	2.77	6.33	3.35	3.63	0.94	2.28	2.75	1.08
Vol (%)	8.32	7.40	8.50	6.92	7.79	7.72	6.94	8.70
Skew	-1.23	-1.01	-0.83	-0.70	-0.44	-0.35	-0.20	-0.31
MDD (%)	-38.32	-25.26	-33.30	-31.91	-39.09	-18.66	-22.57	-74.81
Total return	5.21	20.76	6.61	7.79	2.45	4.47	5.58	2.48

PANEL C: $AUM_0 = 1$ billion

Before cost SR	1.10	0.81	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	5.35	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	6.58	8.43	6.88	7.79	7.69	6.97	8.68
Mean Costs (%)	11.95	0.58	7.82	5.31	0.04	1.87	2.40	5.63
Turnover	2.36	0.30	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	2.22	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	5.87	1.41	10.48	7.15	0.20	3.81	3.76	7.63
Avg USD Trade Size (millions)	151.72	89.65	224.94	153.49	4.09	73.45	70.31	170.01

Continued on next page

Table 1 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
SR	-0.31	0.72	-0.20	-0.03	0.12	0.10	0.11	-0.34
ΔSR	-1.04***	-	-0.92***	-0.75***	-0.60**	-0.62***	-0.61***	-1.06***
Mean (%)	-3.25	4.76	-1.82	-0.24	0.92	0.83	0.76	-3.03
Vol (%)	10.31	6.60	8.94	7.20	7.79	7.96	7.12	8.91
Skew	-2.40	-0.92	-0.91	-0.82	-0.44	-0.49	-0.39	-0.25
MDD (%)	-156.51	-29.45	-114.59	-62.63	-39.14	-46.76	-36.46	-136.40
Total return	-0.27	11.93	0.26	1.31	2.42	2.31	2.30	-0.16

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from February 1986 to September 2021. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Table 2: Capacities of the Strategies: Effect of the Fund Size

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Panel A: Asset under Management in February 1986 (USD billions)								
0	0.30	1,000.00	0.49	0.91	45.00	2.00	1.75	0.23
0.05	0.25	1,000.00	0.41	0.77	25.00	1.50	1.50	0.17
0.1	0.21	130.00	0.35	0.66	7.00	1.25	1.25	0.12
0.15	0.18	60.00	0.29	0.56	0.00	0.74	0.83	0.08
0.2	0.16	35.00	0.24	0.47	0.00	0.49	0.64	0.04
0.25	0.13	25.00	0.19	0.39	0.00	0.27	0.48	0.01
0.3	0.11	20.00	0.16	0.32	0.00	0.08	0.34	0.00
0.35	0.10	15.00	0.13	0.26	0.00	0.00	0.21	0.00
0.4	0.08	9.50	0.10	0.21	0.00	0.00	0.09	0.00
0.45	0.07	6.75	0.07	0.16	0.00	0.00	0.00	0.00
0.5	0.06	5.00	0.05	0.12	0.00	0.00	0.00	0.00
0.55	0.05	3.50	0.04	0.08	0.00	0.00	0.00	0.00
0.57	0.04	3.00	0.03	0.07	0.00	0.00	0.00	0.00
0.6	0.04	2.50	0.02	0.05	0.00	0.00	0.00	0.00
0.65	0.03	1.75	0.01	0.02	0.00	0.00	0.00	0.00
0.7	0.02	1.25	0.00	0.00	0.00	0.00	0.00	0.00
0.75	0.02	0.76	0.00	0.00	0.00	0.00	0.00	0.00
0.8	0.01	0.39	0.00	0.00	0.00	0.00	0.00	0.00
0.85	0.01	0.12	0.00	0.00	0.00	0.00	0.00	0.00
0.9	0.01	0.03	0.00	0.00	0.00	0.00	0.00	0.00
Continued on next page								

Table 2 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Panel A: Asset under Management in September 2021 (USD billions)								
0	0.70	2,885.71	1.17	2.26	111.17	4.58	4.05	0.55
0.05	0.69	2,885.71	1.14	2.19	70.78	4.06	3.86	0.48
0.1	0.68	389.07	1.11	2.11	22.76	3.72	3.62	0.39
0.15	0.67	195.62	1.07	2.02	0.00	2.75	3.02	0.30
0.2	0.66	126.80	1.02	1.92	0.00	2.08	2.65	0.18
0.25	0.65	98.53	0.97	1.80	0.00	1.32	2.24	0.06
0.3	0.63	84.06	0.91	1.68	0.00	0.45	1.77	0.00
0.35	0.62	69.18	0.84	1.53	0.00	0.00	1.23	0.00
0.4	0.60	51.40	0.75	1.37	0.00	0.00	0.61	0.00
0.45	0.57	41.42	0.66	1.19	0.00	0.00	0.00	0.00
0.5	0.55	34.37	0.55	0.99	0.00	0.00	0.00	0.00
0.55	0.52	27.63	0.42	0.76	0.00	0.00	0.00	0.00
0.57	0.5	25.18	0.36	0.66	0.00	0.00	0.00	0.00
0.6	0.48	22.57	0.27	0.50	0.00	0.00	0.00	0.00
0.65	0.44	18.21	0.15	0.21	0.00	0.00	0.00	0.00
0.7	0.40	14.86	0.00	0.00	0.00	0.00	0.00	0.00
0.75	0.34	10.82	0.00	0.00	0.00	0.00	0.00	0.00
0.8	0.28	6.71	0.00	0.00	0.00	0.00	0.00	0.00
0.85	0.24	2.51	0.00	0.00	0.00	0.00	0.00	0.00
0.9	0.24	0.77	0.00	0.00	0.00	0.00	0.00	0.00

Notes: The table reports the maximum initial (panel A) and terminal (panel B) AUM (USD billions) of a fund such that its after cost Sharpe ratio is larger or equal than $\overline{SR} \in [0, 0.9]$. The strategies are described in Section 3. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Table 3: Trading at the Quarterly Frequency

Strategies	Before TC	After TC		Turnover	Costs (%)
	SR	SR	Δ SR		
PANEL A: NO PRICE IMPACT					
MV	0.90	0.85	-0.02	2.25	0.43
MVTC	0.91	0.87	-	1.79	0.32
HML	0.51	0.50	-0.37**	0.86	0.11
RB	0.53	0.51	-0.36***	0.71	0.12
DOL	0.11	0.11	-0.76***	0.06	0.02
DDOL	0.22	0.22	-0.65***	0.22	0.05
VAL	0.53	0.52	-0.35*	0.44	0.08
MOM	0.06	0.04	-0.83***	1.60	0.20
PANEL B: $AUM_0 = 100$ millions					
MV	0.90	0.55	-0.23***	2.25	3.00
MVTC	0.84	0.78	-	1.14	0.48
HML	0.51	0.43	-0.35**	0.86	0.77
RB	0.53	0.48	-0.30**	0.71	0.34
DOL	0.11	0.11	-0.67***	0.06	0.02
DDOL	0.22	0.21	-0.57***	0.22	0.10
VAL	0.53	0.48	-0.30*	0.44	0.33
MOM	0.06	-0.04	-0.81***	1.60	0.79
PANEL C: $AUM_0 = 1$ billion					
MV	0.90	-0.06	-0.72***	2.25	8.00
MVTC	0.75	0.67	-	0.61	0.57
HML	0.51	0.11	-0.55***	0.86	3.78
RB	0.53	0.29	-0.38***	0.71	1.82
DOL	0.11	0.10	-0.56**	0.06	0.06
DDOL	0.22	0.15	-0.51**	0.22	0.54
VAL	0.53	0.26	-0.41**	0.44	1.96
MOM	0.06	-0.36	-1.03***	1.60	3.64

The table summarizes the performance of our FX market strategies when they are implemented at a quarterly frequency. The strategies are described in Section 3. Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Column 2 and 3 report the annualized before (Before TC SR) and after cost (After TC SR) Sharpe ratios. Δ SR in column 4 is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of Δ SR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Turnover in column 5 reports the average monthly turnover. Costs (%) in column 6 reports the annual transaction costs as a percentage of the AUM. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Table 4: Trading only Developed Currencies

Strategies	Before TC	After TC		Turnover	Costs (%)
	SR	SR	Δ SR		
PANEL A: NO PRICE IMPACT					
MV	1.09	1.00	-0.04	1.97	0.61
MVTC	1.10	1.04	-	1.41	0.40
HML	0.70	0.66	-0.37**	1.57	0.35
RB	0.66	0.61	-0.42***	1.35	0.36
DOL	0.08	0.08	-0.95***	0.02	0.01
DDOL	0.39	0.38	-0.66***	0.38	0.11
VAL	0.62	0.60	-0.43*	0.35	0.10
MOM	0.17	0.15	-0.89***	0.95	0.21
PANEL B: $AUM_0 = 100$ millions					
MV	1.09	0.37	-0.53***	1.97	4.46
MVTC	0.98	0.91	-	0.59	0.38
HML	0.70	0.27	-0.64***	1.57	3.89
RB	0.66	0.36	-0.54***	1.35	2.22
DOL	0.08	0.08	-0.82***	0.02	0.01
DDOL	0.39	0.35	-0.55**	0.38	0.30
VAL	0.62	0.56	-0.35	0.35	0.46
MOM	0.17	0.06	-0.84***	0.95	1.00
PANEL C: $AUM_0 = 1$ billion					
MV	1.09	-0.39	-1.14***	1.97	10.14
MVTC	0.85	0.75	-	0.24	0.36
HML	0.70	-0.32	-1.07***	1.57	9.61
RB	0.66	-0.25	-1.00***	1.35	7.01
DOL	0.08	0.08	-0.67**	0.02	0.02
DDOL	0.39	0.18	-0.57**	0.38	1.63
VAL	0.62	0.29	-0.46*	0.35	2.57
MOM	0.17	-0.28	-1.03***	0.95	4.32

The table summarizes the performance of our FX market strategies when they are constructed from only developed currencies. The strategies are described in Section 3. Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Column 2 and 3 report the annualized before (Before TC SR) and after cost (After TC SR) Sharpe ratios. Δ SR in column 4 is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of Δ SR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Turnover in column 5 reports the average monthly turnover. Costs (%) in column 6 reports the annual transaction costs as a percentage of the AUM. The sample includes 13 developed currencies against the USD for the period from February 1986 to September 2021.

Table 5: **Removing High Cost Currencies**

Drop # Top TC Currencies	Before TC SR			After TC SR		
	MV	MVTC	HML	MV	MVTC	HML
PANEL A: NO PRICE IMPACT						
keep all	1.10	1.06	0.72	0.97	0.99	0.67
drop 1	1.15	1.10	0.76	1.03	1.02	0.72
drop 2	0.99	0.95	0.71	0.88	0.88	0.67
drop 3	0.96	0.98	0.65	0.86	0.91	0.61
drop 4	0.71	0.71	0.59	0.61	0.64	0.55
drop 5	0.97	0.92	0.56	0.88	0.85	0.52
drop 6	0.69	0.59	0.56	0.61	0.54	0.53
drop 7	0.32	0.31	0.56	0.27	0.27	0.52
drop 8	0.26	0.25	0.45	0.22	0.22	0.42
drop 9	0.11	0.11	0.41	0.07	0.08	0.39
drop 10	0.29	0.29	0.41	0.26	0.26	0.38
PANEL B: $AUM_0 = 100$ millions						
keep all	1.10	0.93	0.72	0.33	0.86	0.39
drop 1	1.15	0.95	0.76	0.32	0.87	0.42
drop 2	0.99	0.84	0.71	0.27	0.75	0.40
drop 3	0.96	0.84	0.65	0.18	0.74	0.32
drop 4	0.71	0.64	0.59	0.13	0.54	0.30
drop 5	0.97	0.62	0.56	0.24	0.53	0.28
drop 6	0.69	0.43	0.56	0.17	0.34	0.30
drop 7	0.32	0.42	0.56	0.07	0.33	0.31
drop 8	0.26	0.40	0.45	0.05	0.32	0.23
drop 9	0.11	0.16	0.41	-0.03	0.08	0.23
drop 10	0.29	0.28	0.41	0.09	0.20	0.26

Continued on next page

Table 5 – continued from previous page

Drop # Top TC Currencies	Before TC SR			After TC SR		
	MV	MVTC	HML	MV	MVTC	HML
PANEL C: $AUM_0 = 1$ billion						
keep all	1.10	0.81	0.72	-0.31	0.72	-0.20
drop 1	1.15	0.79	0.76	-0.33	0.68	-0.18
drop 2	0.99	0.74	0.71	-0.37	0.63	-0.19
drop 3	0.96	0.65	0.65	-0.33	0.53	-0.25
drop 4	0.71	0.48	0.59	-0.38	0.36	-0.27
drop 5	0.97	0.50	0.56	-0.44	0.35	-0.27
drop 6	0.69	0.38	0.56	-0.39	0.24	-0.25
drop 7	0.32	0.52	0.56	-0.15	0.43	-0.23
drop 8	0.26	0.50	0.45	-0.12	0.40	-0.27
drop 9	0.11	0.26	0.41	-0.06	0.16	-0.23
drop 10	0.29	0.30	0.41	0.00	0.21	-0.15

Notes: The table summarizes the performance of the three FX market strategies *MV*, *MVTC*, and *HML*. The strategies are described in Section 3. Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Annualized before (Before TC SR) and after cost (After TC SR) Sharpe ratios are reported in columns 2 to 4 respectively 5 to 7. The top row includes all available currencies. In subsequent rows in every month t we remove one-by-one the currency with the highest median transaction cost over the previous 6 months. The sample (before removing expensive currencies) includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Proportional Costs




Figure 1: The figure plots the cross-currency median of the average proportional costs (section 2.3) for the spot market, $med(0.5C_{i,t}^{S+} + 0.5C_{i,t}^{P-})$ (solid line), and the forward market, $med(0.5C_{i,t}^{P+} + 0.5C_{i,t}^{S-})$ (dashed line). The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Correlation between Alternative Realized Amihud Measures




Figure 2: The figure shows for the spot exchange rates of each currency $i \in \mathcal{I}$ (AUD, EUR, GBP, NZD, CAD, CHF, JPY, NOK, SEK, MXN, SGD, ZAR) the correlations between our daily $A_{i,t}^S$ in (7) (imposing time-invariant volume $v_{i,t}^S = \bar{v}_i^S$ and using hourly instead of minute level Olsen data for exchange rate quotes) and the daily realized Amihud measure of [Rinaldo and Santucci de Magistris \(2022\)](#) for the sample period from September 2012 to September 2021.

Cross-sectional Extrapolation of Trading Volume Data




Figure 3: The asterisks in the figure plot the natural logarithm of the average volume $\ln(\bar{v}_i^S)$ against the natural logarithm of the average proportional bid-ask spread $\ln(\overline{PBA}_i^S)$ for the 12 currencies $i \in \mathcal{I}$ (AUD, EUR, GBP, NZD, CAD, CHF, JPY, NOK, SEK, MXN, SGD, ZAR). The average is taken over the daily sample between September 3rd 2012 through September 24th 2021. The linear line is the fitted regression $\ln(\bar{v}_i^S) = 9.76 - 1.47 \ln(\overline{PBA}_i^S)$, and the circles on the line represent the fitted values $\ln(\bar{v}_i^S)$ corresponding to $\ln(\overline{PBA}_i^S)$ for the 6 currencies BRL, CZK, HUF, KRW, PLN, and TWD (for which CLS trading volume data is unavailable).

Price Impact


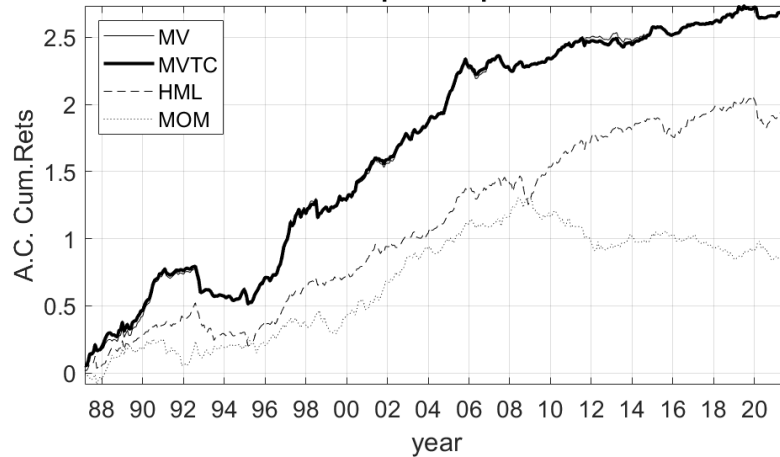


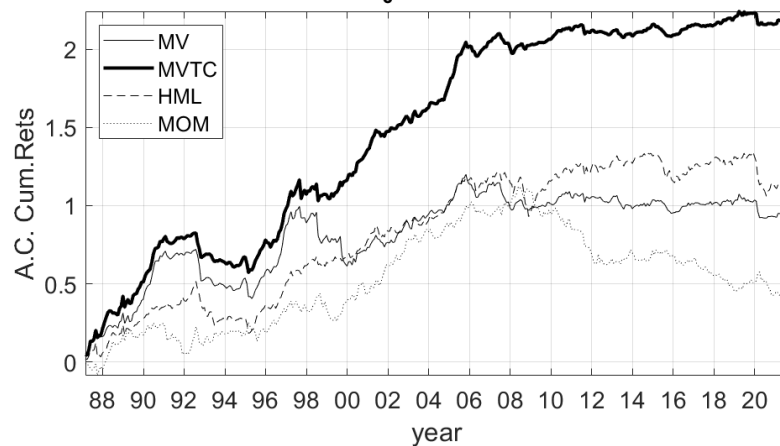
Figure 4: The figure plots the cross-currency median Amihud price impact described in Section 2.4 for the spot (solid line) and the forward market (dashed line). The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Cumulative After-Cost Excess Returns

No price impact



$AUM_0 = 0.1$ billions



$AUM_0 = 1$ billions




Figure 5: The figures plot the after-cost cumulative excess returns of *MV* (solid line), *MVTC* (bold solid line), *HML* (dashed line), and *MOM* (dotted line). The strategies are described in Section 3. The top graph assumes no price impact, while the middle and bottom graphs assume an initial AUM of USD 100 millions respectively 1 billion. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Cumulative Trading Costs

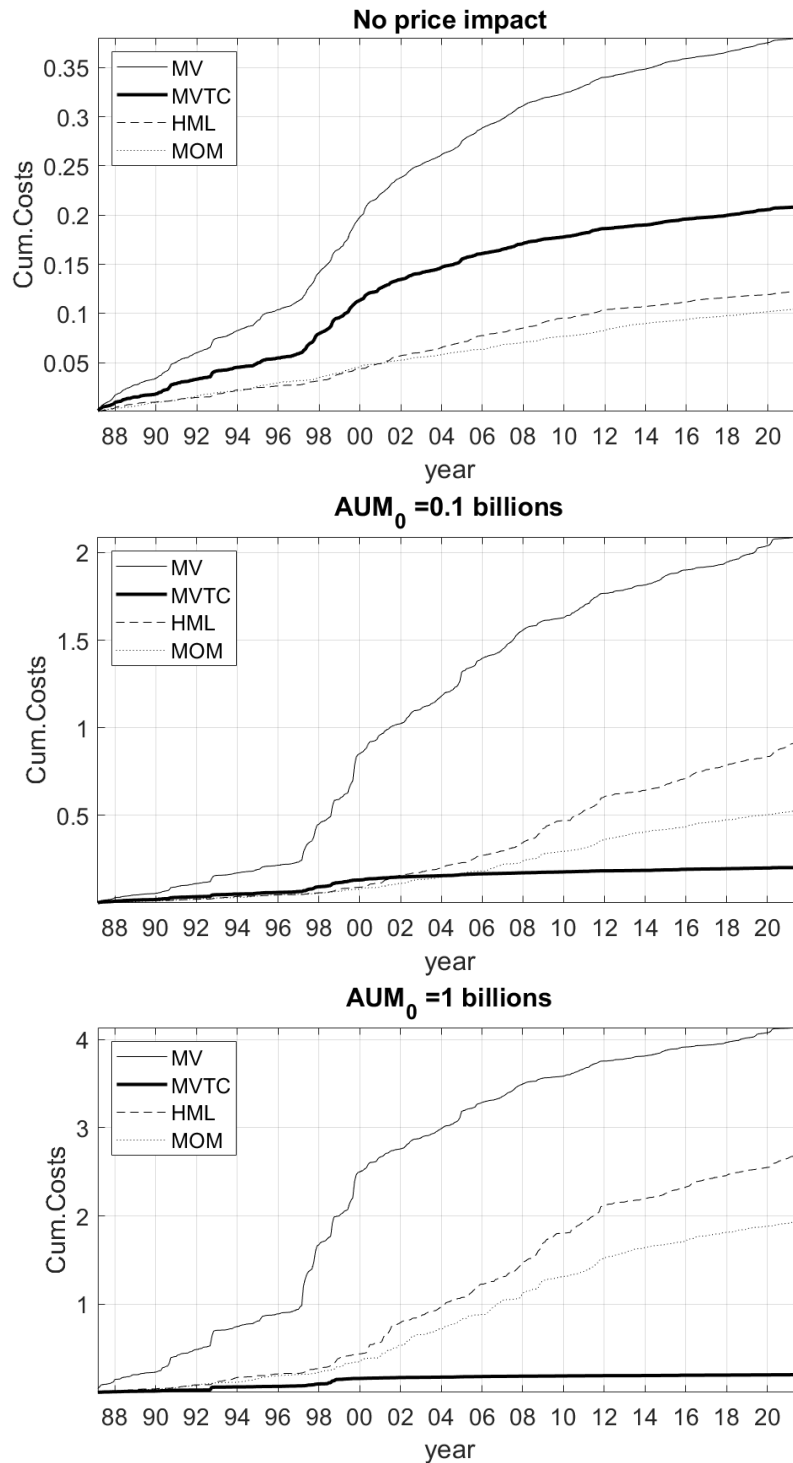


Figure 6: The figures plot the cumulative costs of *MV* (solid line), *MVTC* (bold solid line), *HML* (dashed line), and *MOM* (dotted line). The strategies are described in Section 3. The top graph assumes no price impact, while the middle and bottom graphs assume an initial AUM of USD 100 millions respectively 1 billion. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Capacities of the Strategies: Effect of the Fund Size




Figure 7: The figure plots the annualized after-cost Sharpe ratio as a function of the initial AUM of a fund for our FX trading strategies. The strategies are described in Section 3. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Internet Appendix

A Robustness

A.1 Sub-samples

A longer sample period generally provides more accurate estimates of expected returns. Thus, we use the results in section 4 as our baseline. We further show that our results are robust and continue to hold in the following sub-samples: (i) pre- and post- the 2008 Global Financial Crisis (GFC), (ii) post-introduction of Euro (January 1999), and (iii) in and out of NBER recessions. These robustness checks are important to ensure that our results are generally true, and not driven by a specific sub-sample under special economic conditions.

A.1.1 Pre- versus Post-Global Financial Crisis

As pointed out in section 4.1 the GFC marks a regime shift. In particular *MOM* was profitable in terms of before cost returns right up to the GFC, and turns into a losing strategy afterwards. Similarly, we observe that the before cost average returns of *DDOL* are significantly lower (but still positive) after the GFC. In this section we check whether the post-GFC sample is driving our baseline result that *MV* and the six characteristic based long-short strategies perform poorly after accounting for costs. In other words, we check whether our results are robust in the sub-sample from February 1986 to December 2007. A particular focus is on *DDOL* and *MOM*.

Table A3 reports the results. All strategies but *VAL* deliver noticeably higher before cost Sharpe ratios before the GFC than in the full sample. In particular *DDOL* and *MOM* have a before cost Sharpe ratios of 0.56 and 0.66, which compares to 0.35 and 0.30 in the full sample. As a comparison the Sharpe ratio of the US stock market is 0.50 during the sample of 1986 to 2007. Similar to the baseline analysis proportional costs are relatively small in the pre-GFC sample, while costs due to the price impact are a first order concern when a fund is large. When the initial AUM is USD 100 millions most strategies perform

relatively well after costs. In contrast, when the initial AUM is USD 1 billion *MV*, *HML* and *MOM* are essentially unprofitable. *RB*, *DOL*, *DDOL* and *VAL* are more resilient with after cost Sharpe ratios between 0.27 and 0.37. However, despite that they are profitable they all underperform the US stock market, which has a Sharpe ratio of 0.50. *MVTC* is by far the most resilient strategy outperforming all other strategies (including the US stock market). When the initial AUM is USD 1 billion the after cost Sharpe ratio of *MVTC* is 0.96, and it is above 1.13 when the initial fund size is smaller than USD 100 millions.

For completeness we also report the performance of the strategies in the sub-sample from January 2008 to September 2021 (table A4). All strategies perform worse after the GFC. *DOL*, *DDOL* and *MOM* are essentially unprofitable even before costs. *VAL* has a before cost Sharpe ratio of 0.41 and the after cost Sharpe ratio decreases to 0.24 when the initial fund size is USD 1 billion. *MV*, *MVTC*, *HML*, *RB* perform comparably well with after cost Sharpe ratios slightly above 0.60 when there is no price impact. These four strategies are relatively resilient with after cost Sharpe ratios between 0.43 and 0.54 when the initial AUM is USD 100 millions. When the initial fund size is USD 1 billion *MV* and *HML* turn unprofitable (negative Sharpe ratio), *RB* is essentially also unprofitable with a Sharpe ratio of 0.07, while *MVTC* has a relatively attractive after cost Sharpe ratio of 0.35. Accordingly, our main results continue to hold post-GFC. As a final note it is important to understand that the sub-sample from January 2008 to September 2021 is short (only 165 monthly observations) and expected returns (or Sharpe ratios) are imprecisely measured. To this extend we have to be cautious about drawing strong conclusions from this sub-sample analysis.

A.1.2 Introduction of the Euro

The introduction of the Euro in January 1999, and thus, the replacement of several European currencies had a non-trivial effect on the investment opportunity set in FX markets. We show the robustness of our results in table A2. Our strategies achieve very similar (before and after costs) Sharpe ratios whether we use the full sample of the sub-sample starting in February

1999.³³

A.1.3 NBER Recessions

Finally, our results are robust during and outside of NBER recession periods as documented in Table A5. *HML*, *DOL*, *DDOL* and *MOM* deliver low before cost Sharpe ratios during recessions. Interestingly, *MV*, *MVTC*, *RB* and *VAL* perform well before accounting for costs and appear recession-proof. However, when the fund size is large and the price impact is severe, all strategies except for *MVTC* turn unprofitable. It is striking that *MVTC* performs well and is resilient to even a severe price impact even during recession periods. This resonates and further strengthens our insight that investors should opt for a rigorous portfolio and cost optimization approach rather than characteristic sorted long-short strategies and intuitive rules-of-thumb.

Finally, the results outside of recession periods are almost identical to the results of the full sample. This is not surprising since there are only very few recession months in our sample.

A.2 Alternative Estimations of Transaction Costs

Our findings remain robust when employing different approaches to estimate the price impact functions. This robustness reinforces the notion that the performances of our strategies and the impact of transaction costs is largely independent of the specific cost estimation method employed, indicating that our conclusions are not driven by a specific estimation of costs.

A.2.1 Time-Invariant Price Impact

In this section we use the time-series average of our Amihud measure $\bar{A}_i^k = \frac{1}{T} \sum_t A_{i,t}^k$, for $k \in \{S, F\}$ to construct a time-invariant price impact function. Since the volume $v_{i,t}$ is time-invariant in section 2.4.1 the time-series average over the Amihud measure is equivalent to replacing the numerator in (7) by the time-series average of the exchange rate growth

³³The only noticeable difference is the after cost Sharpe ratio of *MV* which is 0.70 post-introduction of the Euro versus 0.33 in the full sample when the initial AUM is USD 100 millions.

variation $\overline{RPV}_i = \frac{1}{T} \sum_t RPV_{i,t}$. This mitigates the concern that the results might be driven by certain episodes of high illiquidity or volatility. Table A6 reports the results, which are very similar to the results in table 1. Accordingly, our conclusions are the same whether we use the price impact estimates as described in section 2.4, or the time-invariant estimates of this robustness analysis.

A.2.2 Time-Series Extrapolation of Volume Data

In the baseline analysis in section 2.4.1 we keep the trading volume fixed through time $v_{i,t} = \bar{v}_i$, and all time-series variation in the Amihud measure stems from the variation in $RPV_{i,t}$. This is to avoid potentially large errors when we attempt to extrapolate trading volume data in the time-series from 2012 back to 1986. In this section we perform such an extrapolation back in time.

We denote the average of the daily trading volume and proportional bid-ask spread within a quarter q by $\tilde{v}_{i,q}^S = \frac{1}{Q_q} \sum_{t \in \mathcal{Q}_q} v_{i,t}^S$ and $P\tilde{B}A_{i,q}^S = \frac{1}{Q_q} \sum_{t \in \mathcal{Q}_q} PBA_{i,t}^S$, where \mathcal{Q}_q is the set of days in quarter q , $1_{\{t \in \mathcal{Q}_q\}}$ is an indicator function equal to 1 if day t is in quarter q and zero otherwise, and $Q_q = \sum_t 1_{\{t \in \mathcal{Q}_q\}}$ counts the number of days in quarter q . Then, for all currency-quarter (i, q) for which we have observations $\tilde{v}_{i,q}^S$ and $\widehat{PBA}_{i,q}^S$ we run the following panel regression:

$$\ln(\tilde{v}_{i,q}^S) = \tilde{a} + \tilde{b} \ln(\widehat{PBA}_{i,q}^S) + \tilde{\varepsilon}_{i,q}. \quad (15)$$

Figure A7 plots $\ln(\tilde{v}_{i,q}^S)$ against $\ln(\widehat{PBA}_{i,q}^S)$ (indicated by asterisks). The regression fit $R^2 = 65\%$, and the slope coefficient $\tilde{b} = -1.18$ are similar to the cross-sectional regression (9).³⁴

Then, for all currency-quarter (i, q) for which we do not observe $\tilde{v}_{i,q}^S$ but have an estimate

³⁴Adding time fixed effects does not noticeably change the regression. The regression fit increases to $R^2 = 72\%$, the time fixed effects are close to zero, and the slope coefficient $\tilde{b} = -1.32$ is only slightly different from the regression without fixed effects.

of $\widehat{PBA}_{i,q}^S$, we use the following extrapolation:

$$\tilde{v}_{i,q}^S = \exp \left\{ \tilde{a} + \tilde{b} \ln \left(\widehat{PBA}_{i,q}^S \right) \right\}.$$

Note that we can do this also for the currencies $i \in \mathcal{I}_{EU}^c$, and we do not have to do the re-scaling in (11). Next, we use the quarterly volume as a proxy for the daily volume on each day within the quarter, $v_{i,t}^S = \tilde{v}_{i,q}^S, \forall t \in \mathcal{Q}_q$. Finally, we plug the constructed panel of $v_{i,t}^S$ (together with the $RPV_{i,t}$ computed using Olsen data) into (7) to get the realized Amihud $A_{i,t}^S$. We then use the approach in (12) to obtain $A_{i,t}^F$ for forwards.

As expected the time-series of the extrapolated volume data is increasing over time. In turn, this implies that the realized Amihud measure and therefore the price impact in our baseline analysis in section 2.4 is smaller and more conservative than in this robustness section. Accordingly, the results in table A7 indicate that costs due to the price impact of trading have a slightly more detrimental effect on the profitability of the strategies. As before *MVTC* is still robust to even a large fund size and a severe price impact. However, the difference in the results (table 1 vs A7) is relatively small, implying that the extrapolation method for the volume data does not affect our conclusions.

A.2.3 BIS Volume Data

In this section we use volume data from the BIS Triennial Survey instead of the CLS data. An advantage of the BIS data is that it goes back until 1992, and thus, may be useful to obtain more accurate extrapolations before 2012. On the flip side, we need to interpolate between dates as the BIS data is only available for average trading volume in the month of April and only every three years. We denote by $\hat{v}_{i,y}^S$ the average daily volume in spot markets in April of year y , where we only have data for every third year.³⁵ Moreover, we denote by $\widehat{PBA}_{i,y}^S$ the average proportional bid-ask spread for spot transactions in year y , $\widehat{PBA}_{i,y}^S = \frac{1}{Y_y} \sum_{t \in \mathcal{Y}_y} PBA_{i,t}^S$, where \mathcal{Y}_y is the set of days in year y , $1_{\{t \in \mathcal{Y}_y\}}$ is an indicator function equal to 1 if day t is in year y and zero otherwise, and $Y_y = \sum_t 1_{\{t \in \mathcal{Y}_y\}}$ counts the number of days in year y . Then, for all currency-year (i, y) for which we have observations

³⁵The survey years are 1992, 1995, 1998, 2001, 2004, 2007, 2010, 2013, 2016, 2019, 2022.

$\hat{v}_{i,y}^S$ and $\widehat{PBA}_{i,y}^S$ we implement an analogous panel regression to (15),

$$\ln(\hat{v}_{i,y}^S) = \hat{a} + \hat{b} \ln(\widehat{PBA}_{i,y}^S) + \varepsilon_{i,y}^{BIS}. \quad (16)$$

Figure A8 plot $\ln(\hat{v}_{i,y}^S)$ against $\ln(\widehat{PBA}_{i,y}^S)$ (indicated by asterisks). The regression fit $R^2 = 51\%$, and the slope coefficient $\hat{b} = -1.6$ are similar to the cross-sectional regression (15).³⁶

Then, for all currency-year (i, y) (where y includes only BIS Triennial survey years) for which we do not observe $\hat{v}_{i,y}^S$ but have an estimate of $\widehat{PBA}_{i,y}^S$, we use the following extrapolation:

$$\hat{v}_{i,y}^S = \exp\left\{\hat{a} + \hat{b} \ln(\widehat{PBA}_{i,y}^S)\right\}.$$

Next, we linearly interpolate between BIS Triennial survey years to obtain daily volume data $v_{i,t}^S$. Finally, we plug the constructed panel of $v_{i,t}^S$ (together with the $RPV_{i,t}$ computed using Olsen data) into (7) to get the realized Amihud measure $A_{i,t}^S$. We then use the approach in (12) to obtain $A_{i,t}^F$ for forwards.

The results in table A8 are almost identical to the results in table A7. Accordingly, it does not make a difference whether we construct the realized Amihud measure from extrapolated CLS data or from BIS data, and our results are robust.

A.2.4 Cutting Costs in Half

In the main text we only consider trades θ and Δ against the USD. This ignores the possibility that it may be cheaper to trade directly between two non-USD currencies rather than twice against the USD.

As an illustration suppose that we want to be long AUD and short JPY. In our analysis we open (i) a long position in an AUD forward against the USD, and (ii) a short position in a JPY forward against the USD.³⁷ Then, we add up the costs of the two trades. However, the

³⁶As with the panel regression with the CLS data, adding time fixed effects as no material effect.

³⁷Alternatively, if there were any open AUD shorts, then we close a short position in an AUD against

USD positions in the two trades are offsetting, and we can directly enter a long AUD forward against JPY. The alternative trade only incurs transaction costs for one trade instead of two trades.

This does however not imply that our analysis necessarily overestimates trading costs. First, it is an empirical question whether it is cheaper to directly trade AUD against JPY, or indirectly AUD against USD and USD against JPY. [Rinaldo and Santucci de Magistris \(2022\)](#) and [Somogyi \(2022\)](#) document in the data that trading against the USD is generally cheaper than other bilateral trades. [Somogyi \(2022\)](#) further argues that part of the USD dominance in terms of order volume can be contributed to the cost advantage of trading against USD. Second, our illustrative example can be overturned if we consider multiple trading periods. Following the literature, in our analysis we assume there are no costs to roll a contract over from month to month. Suppose at time $t - 1$ we are long AUD and short JPY, and at time t we want to remain long AUD but close our short JPY position. Effectively this means that we need to convert the short JPY position into a short USD position. If we have entered long AUD against USD and short JPY against USD forwards at $t - 1$, then at t we roll over the AUD against USD contract and close the existing short forward in JPY against USD. In contrast, if we have implemented the direct AUD against JPY at $t - 1$, then at t we either roll the contract over and open a new long position in a JPY forward against USD, or we close the previous position and open a new long forward in AUD against USD. Clearly at this point the latter trading strategy also requires two trades in forwards (one at $t - 1$, and another at t). Thus, it loses its efficiency compared to trading AUD against USD and JPY against USD at $t - 1$. Moreover, in the data it is generally the case that trading costs of AUD against USD plus USD against JPY are less than costs of AUD against JPY plus costs of JPY against USD (or AUD against USD). Accordingly, if we consider the total costs over a longer time horizon (rather than just at one point in time) and given that every position will be unwound eventually, the costs as computed in our analysis may not be higher than if we allow to trade all bilateral currency pairs.

From the point of view of implementing our analysis, using the currency triangle to

USD contract. Similarly, we close a long position in a JPY against USD contract if we were previously long. Whether we open new or close existing positions does not matter for the purpose of this illustration.

reduce costs is a non-trivial problem. First, we have N_t currencies. Considering all possible combinations leads to $\frac{N_t(N_t-1)}{2}$ bilateral trades. In contrast in our analysis we only work with N_t possible trades against the USD. Second, as illustrated above the problem becomes difficult when we have an inter-temporal setting. To sum up, our assumption to trade only against the USD is not innocuous but it clearly simplifies our analysis, and if we consider multiple time periods it may be close to the efficient policy.

Nevertheless, the robustness analysis in this section addresses the limitation of our baseline setting. Following the logic of our illustrative example in the best case scenario bilateral trades can reduce the number of trades by 50% compared to our baseline analysis where only trading against the USD is possible. If all trades had the same costs, this would imply that total trading costs can be reduced at most by 50%. Recall that the literature suggests that trades against the USD are generally cheaper than trades against other currencies. Accordingly, if we consider a strategy that trades only against the USD but is subject to artificial costs that are only 50% of the true costs, then we obtain a lower bound to the total costs of the optimal cost minimizing strategy that trades in all bilateral currency pairs and is subject to the true transaction costs. Table A9 reports the results. Cutting costs in half implies that the capacity of every strategy approximately doubles. While this provides a boost to the strategies, the capacities of *MV* and the six characteristic based long-short strategies are still small, and they all perform poorly when the initial AUM of a fund reaches USD 1 billion. To this extent our conclusions do not differ from our baseline analysis in table 1.

A.3 Sensitivity Analyses

A.3.1 Event Study: Performance of *MVTC* in Turbulent Times

Ceteris paribus in periods of high volatility a mean-variance optimized strategy unwinds its risk exposure. This is because the optimal weights are inversely related to risk. Nonetheless, unwinding positions in turbulent times might be costly or even impossible in practice. Accordingly, a potential concern is that the performance of *MVTC* may be too optimistic as reported in table 1. Although we have shown that *MVTC* is very resilient to the price impact, we now impose a stricter constraint. That is, we do not allow any rebalancing dur-

ing turbulent times. This allows us to check whether *MVTC* experiences big losses when unwinding or deleveraging is impossible during bad times.

We identify turbulent or bad times as episodes of high FX market volatility or illiquidity. We use the FX market volatility measure of [Menkhoff et al. \(2012a\)](#), and aggregate [Corwin and Schultz \(2012\)](#) effective exchange rate spreads to estimate illiquidity as suggested by [Karnaukh et al. \(2015\)](#). As illustrated in Figure [A9](#) by the blue and red dots, we select the five highest (non-consecutive) peaks in volatility and illiquidity over a rolling window of 13 months. Moreover, we add four months which represent important and well-known crises that involve the FX markets (highlighted by the vertical black lines): September 1992 which contains the Black Wednesday, July 1997 which signifies the onset of the Asian Financial Crisis with the collapse of the Thai baht, September 2008 which marks the onset of the credit crunch with the collapse of Lehman Brothers, and March 2020 the recognition that COVID is a global pandemic and the start of lock-downs across the world. As it is visually apparent, all these events form 9 distinct points in time which represent turbulent months in our analysis.

Table [A10](#) compares the Maximum Draw Downs (MDD) of *MVTC* and *BH*, which is a buy and hold strategy that keeps the position of *MVTC* fixed starting 1 month before and ending 12 months after a turbulent month. Column 2 through 4 of Table [A10](#) report whether the month is turbulent because of a peak in the volatility, illiquidity or because it is a well-known crisis. Columns 5 to 10 report the MDD for *MVTC* and *BH* when there is no price impact (columns 5 and 6), and when the initial AUM is USD 100 millions (columns 7 and 8) and USD 1 billion (columns 9 and 10). The last row shows the average MDD for each strategy across the different event studies.

On average *MVTC* experiences a maximum loss for *BH* between 1.6 to 1.8 times the loss of *MVTC*. The worst drawdowns are reported following the crisis in September 1992. The MDD of *MVTC* is roughly -10% (slightly differing across the different AUMs) while it is between -28% and -21% for *BH*. Accordingly, crash risk increases if rebalancing is impossible for an extended period of time. However, it is still manageable and is relatively mild compared to high crash risk levels of other strategies. We conclude that *MVTC* continues

to perform well with manageable crash risk exposure even if it is impossible to rebalance within up to 12 months after a spike in volatility or illiquidity, or after the onset of a crisis.

A.3.2 Risk Aversion Coefficient λ

Finally, in Table A11 we show that the Sharpe ratio of *MVTC* is increasing in the coefficient of risk aversion λ . This effect is stronger with a large fund size and a more severe price impact. The intuition is that when the risk aversion is large, then the investor does not take much risk and the portfolio holdings are close to zero. As such rebalancing the portfolio does not require large trade orders and little costs are incurred, implying that the after cost Sharpe ratio is close to the before cost Sharpe ratio of *MV*. We believe $\lambda = 25$ is a reasonable value as it implies volatilities of *MV* and *MVTC* that are comparable to the characteristic based long-short strategies. If we decrease λ to 10 or even 5, *MVTC* delivers a lower Sharpe ratio than if $\lambda = 25$, but it is still resilient to a large fund size and a severe price impact and outperforms the US stock market. Decreasing λ even further would imply a volatility and a crash risk (the chance of losing everything is high) that probably no investor would be willing to bear.

B Theory

B.1 Algorithm

We can rewrite Problem 2 as

$$\min_{\mathbf{x}_t \geq 0} \mathbf{q}_t' \mathbf{x}_t + \frac{1}{2} \mathbf{x}_t' \mathbf{H}_t \mathbf{x}_t$$

where $\mathbf{q}_t' \equiv \mathbf{C}_t' \bar{\mathbf{I}}_{2,t}' + \lambda(\theta_t^0 - \theta_t^{\text{MV}})' \hat{\mathbf{V}}_t \bar{\mathbf{I}}_t$, with $\theta_t^{\text{MV}} = \frac{1}{\lambda} \tilde{\mathbf{V}}_t^{-1} \mathbf{f} \mathbf{d}_t$, $\mathbf{H}_t = \bar{\mathbf{I}}_t' (\lambda \hat{\mathbf{V}}_t + \mathbf{\Pi}_t) \bar{\mathbf{I}}_t$, $\bar{\mathbf{I}}_t \equiv [I_{N_t}, -I_{N_t}]$. $\mathbf{\Pi}_t$ is an $N_t \times N_t$ diagonal matrix with $2\pi_{i,t}$ (defined in section 3.2.2) as the i th element on the diagonal. \mathbf{C}_t is an $N_t \times 1$ column vector with $\mathbf{C}_{i,t}$ (defined in section 3.2.2) as the i th element.

$\bar{\mathbf{I}}_{2,t} \equiv [I_{N_t}, I_{N_t}]$, and I_{N_t} is the $N_t \times N_t$ identity matrix for the N_t assets available at time t . The program returns $\mathbf{x}_t \equiv [\Delta_t^{\text{P}'}, \Delta_t^{\text{S}'}]$ and the unique optimal portfolio θ_t^{MVTC} is obtained

by

$$\theta_t^{MVTC} = \theta_t^0 + \bar{\mathbf{I}}_t \mathbf{x}_t.$$

Notice that solving this program at every t produces the strategy *MVTC* (i) with no-price impact if $\mathbf{\Pi}_t = 0$, (ii) with linear price impact if $\mathbf{\Pi}_t$ is positive definite. In any case this is a well-behaved convex program, and we solve it using the Matlab Optimization ToolBox.

B.2 Illustration of No-Trading Region

Figure A10 illustrates the trade-off of *MVTC* between optimal rebalancing and cost minimization in a setting with two risky assets. 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 in the center represents the optimal portfolio θ_t^{MV} if there were no transaction costs. The blue parallelogram surrounding θ_t^{MV} defines the no trading region when the two assets are positively correlated. If the initial allocation θ_t^0 is inside the no trading region (i.e., within the blue parallelogram), then there is no trade and $\theta_t^{MVTC} = \theta_t^0$, as the marginal cost to trade towards θ_t^{MV} exceeds the marginal utility. If the initial portfolio θ_t^0 lies outside of the no trading region, then the investor wants to move towards θ_t^{MV} but stops trading once she reaches the boundary of the no trading region. The arrows indicate the direction of trade and the arrow heads show how far to trade. The purple, brown or red colors of the arrows indicate that only asset 1, 2 or both assets are traded.

If we construct the no trading region around *MV* ignoring the correlation between the assets, the region reduces to the yellow square in Figure A10. In a continuous time setup where assets are independent, the optimal dynamic strategy of Liu (2004) shares many qualitative properties with this approximate solution. It can be shown that ignoring correlations in the construction of the region is sub-optimal both theoretically and empirically.

In Figure A10 if the two assets are positively correlated, then the blue parallelogram no trading region of *MVTC* is larger along the -45° line than the yellow square no trading region if we were to set correlations equal to zero. This is because the two assets are substitutes if they are positively correlated, while they are not substitutable if they are

uncorrelated. 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 leaves 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 0). Since an initial position θ_t^0 close to the -45° line requires the investor to buy one and sell the other asset, the marginal utility from trading towards θ_t^{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 *MVTC* is relatively smaller along the -45° line.

C Data Sources and Filters

In Table [A12](#) we list the Datastream mnemonics for spot and forward exchange rate quotes against the USD. We collect daily data from Barclays and Reuters from January 1986 to December 2021. For the Australian dollar, Japanese yen, New Zealand dollar, Norwegian krone, Swedish krona, and Swiss franc, we replace the Barclays forward series with Reuters forward rates (in parenthesis) in January 1997 as the former are discontinued in the recent period. To obtain bid and ask exchange rates, the suffixes (EB) and (EO) are added to the corresponding mnemonics. We also collect hourly bid and ask quotes for spot exchange rates from Olsen from February 1986 to September 2021.

Following [Menkhoff et al. \(2012b\)](#) and [Della Corte et al. \(2016\)](#) we do not include currencies that have a negative score on the capital account openness index of [Chinn and Ito \(2006\)](#) as those currencies are viewed as subject to major trading frictions coming from capital controls. In addition, we do not consider the Danish krone due to its low trading volume and low liquidity proxied by the Amihud measure. Details about the Amihud measure are in section [2.4](#). This is in line with [Ranaldo and Santucci de Magistris \(2022\)](#) who show that the Danish krone is very illiquid and has the largest price impact among all currencies in their sample. We further exclude the Icelandic krona, which has an average bid-ask spread of 45 bps (roughly 10 times the median spread across all currencies), and the Belgian franc due

to limited data availability. Finally, following Maurer et al. (2023) we exclude a currency at time t if it is pegged to another currency, has proportional costs bigger than 0.5%, more than 20% of its daily exchange rate changes are missing over the past 6 months, or if the absolute value of the annualized forward discount $|fd_{i,t}|$ is larger than 20%. Forward discounts of more than 20% 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 extraordinarily 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.

Data for the US stock market and the USD risk-free rate are from the Kenneth French data library, https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. Data for the S&P 500 SPDR ETF (SPY) are from Yahoo finance, <https://finance.yahoo.com/>.

D Tables and Figures

D.1 Analyses of Sub-samples

Table A1: Performance of FX Strategies: Sample Before the Introduction of the Euro (January 1999)

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL A: NO PRICE IMPACT								
Before cost SR	1.13	1.10	0.71	0.70	0.20	0.44	0.37	0.36
Before cost Mean (%)	11.86	11.20	6.03	5.08	1.43	3.20	2.90	3.37
Before cost Vol (%)	10.48	10.21	8.53	7.22	7.28	7.24	7.92	9.47
Mean Costs (%)	1.39	0.81	0.32	0.33	0.03	0.12	0.14	0.33
Turnover	2.75	1.68	0.77	0.62	0.04	0.22	0.27	0.86
Notional	4.34	4.15	2.00	2.00	1.00	1.00	2.00	2.00
SR	1.00	1.02	0.67	0.66	0.19	0.43	0.35	0.32
ΔSR	-0.02	-	-0.35	-0.36	-0.83*	-0.59	-0.67	-0.70
Mean (%)	10.45	10.39	5.70	4.76	1.41	3.08	2.76	3.04
Vol (%)	10.48	10.20	8.53	7.23	7.28	7.24	7.92	9.46
Skew	-1.23	-1.16	-1.06	-0.74	-0.37	-0.58	-0.14	-0.56
MDD (%)	-26.43	-28.17	-32.27	-31.43	-21.40	-14.85	-22.40	-19.90
Total Return	4.89	4.89	2.41	2.16	1.15	1.58	1.54	1.67
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	1.13	1.00	0.71	0.70	0.20	0.44	0.37	0.36
Before cost Mean (%)	11.86	9.96	6.03	5.08	1.43	3.20	2.90	3.37
Before cost Vol (%)	10.48	9.93	8.53	7.22	7.28	7.24	7.92	9.47

Continued on next page

Table A1 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	4.97	0.98	0.67	0.47	0.03	0.13	0.23	0.57
Turnover	2.75	1.08	0.77	0.62	0.04	0.22	0.27	0.86
Notional	4.34	3.96	2.00	2.00	1.00	1.00	2.00	2.00
Avg Relative Trade Size (%)	1.32	0.41	0.55	0.23	0.01	0.07	0.24	0.41
Avg USD Trade Size (millions)	46.12	20.02	18.56	7.86	0.40	2.42	5.73	17.51
SR	0.64	0.91	0.63	0.64	0.19	0.42	0.34	0.30
ΔSR	-0.26	-	-0.28	-0.27	-0.71	-0.48	-0.57	-0.61
Mean (%)	6.80	8.97	5.36	4.61	1.40	3.07	2.67	2.80
Vol (%)	10.58	9.90	8.54	7.23	7.28	7.25	7.91	9.47
Skew	-1.40	-1.32	-1.05	-0.74	-0.37	-0.58	-0.14	-0.56
MDD (%)	-31.51	-25.26	-33.30	-31.91	-21.41	-14.88	-22.57	-20.24
Total return	2.82	4.09	2.27	2.11	1.15	1.58	1.51	1.60

PANEL C: $AUM_0 = 1$ billion

Before cost SR	1.13	0.79	0.71	0.70	0.20	0.44	0.37	0.36
Before cost Mean (%)	11.86	7.23	6.03	5.08	1.43	3.20	2.90	3.37
Before cost Vol (%)	10.48	9.12	8.53	7.22	7.28	7.24	7.92	9.47
Mean Costs (%)	16.79	1.25	3.26	1.63	0.06	0.29	1.02	2.44
Turnover	2.75	0.51	0.77	0.62	0.04	0.22	0.27	0.86
Notional	4.34	3.51	2.00	2.00	1.00	1.00	2.00	2.00
Avg Relative Trade Size (%)	6.83	1.42	4.76	2.14	0.13	0.67	2.32	3.67
Avg USD Trade Size (millions)	252.37	79.44	164.01	74.47	3.96	24.05	55.73	159.75

Continued on next page

Table A1 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
SR	-0.38	0.65	0.31	0.47	0.19	0.40	0.24	0.10
ΔSR	-1.03***	-	-0.33	-0.18	-0.46	-0.25	-0.41	-0.55
Mean (%)	-5.26	5.97	2.75	3.44	1.37	2.91	1.87	0.92
Vol (%)	13.84	9.21	8.75	7.31	7.29	7.26	7.86	9.53
Skew	-1.76	-1.15	-0.99	-0.72	-0.37	-0.57	-0.14	-0.57
MDD (%)	-108.17	-29.45	-41.50	-36.11	-21.53	-15.14	-24.01	-23.59
Total return	-0.13	2.58	1.41	1.71	1.14	1.53	1.28	1.08

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from February 1986 to December 1998. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to December 1998.

Table A2: Performance of FX Strategies: Sample After the Introduction of the Euro (January 1999)

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL A: NO PRICE IMPACT								
Before cost SR	1.14	1.10	0.71	0.75	0.12	0.28	0.52	0.23
Before cost Mean (%)	7.22	6.81	6.04	5.09	0.95	2.29	3.40	1.94
Before cost Vol (%)	6.31	6.20	8.54	6.81	8.18	8.07	6.49	8.26
Mean Costs (%)	0.84	0.44	0.37	0.42	0.02	0.21	0.10	0.28
Turnover	2.04	1.34	1.32	1.11	0.03	0.48	0.24	0.98
Notional	2.38	2.27	2.00	2.00	1.00	0.99	2.00	2.00
SR	1.02	1.03	0.66	0.69	0.11	0.26	0.51	0.20
ΔSR	-0.01	-	-0.37*	-0.34	-0.92***	-0.77***	-0.52**	-0.83***
Mean (%)	6.37	6.36	5.67	4.67	0.93	2.08	3.29	1.66
Vol (%)	6.26	6.17	8.54	6.80	8.18	8.07	6.49	8.26
Skew	0.02	-0.02	-0.60	-0.65	-0.47	-0.21	-0.21	-0.18
MDD (%)	-12.59	-11.90	-24.45	-17.73	-39.09	-17.69	-17.75	-50.31
Total return	4.19	4.20	3.33	2.62	0.60	1.05	1.74	0.84
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	1.14	0.97	0.71	0.75	0.12	0.28	0.52	0.23
Before cost Mean (%)	7.22	5.80	6.04	5.09	0.95	2.29	3.40	1.94
Before cost Vol (%)	6.31	5.99	8.54	6.81	8.18	8.07	6.49	8.26

Continued on next page

Table A2 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	2.84	0.42	1.60	0.95	0.02	0.33	0.30	1.01
Turnover	2.04	0.82	1.32	1.11	0.03	0.48	0.24	0.98
Notional	2.38	2.08	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	1.59	0.32	1.72	0.75	0.01	0.25	0.28	1.01
Avg USD Trade Size (millions)	29.32	13.08	30.62	14.57	0.21	4.60	4.81	18.34
SR	0.70	0.90	0.52	0.61	0.11	0.24	0.48	0.11
ΔSR	-0.20**	-	-0.38*	-0.29	-0.78**	-0.65**	-0.42	-0.79**
Mean (%)	4.35	5.37	4.44	4.14	0.93	1.96	3.10	0.93
Vol (%)	6.21	5.98	8.56	6.81	8.18	8.08	6.48	8.26
Skew	-0.12	-0.49	-0.63	-0.66	-0.47	-0.21	-0.23	-0.17
MDD (%)	-13.88	-15.25	-26.22	-18.49	-39.09	-18.04	-18.19	-60.66
Total return	2.37	3.21	2.33	2.23	0.60	1.00	1.62	0.58

PANEL C: $AUM_0 = 1$ billion

Before cost SR	1.14	0.88	0.71	0.75	0.12	0.28	0.52	0.23
Before cost Mean (%)	7.22	4.55	6.04	5.09	0.95	2.29	3.40	1.94
Before cost Vol (%)	6.31	5.18	8.54	6.81	8.18	8.07	6.49	8.26
Mean Costs (%)	9.41	0.31	6.82	4.09	0.03	1.26	1.75	4.90
Turnover	2.04	0.37	1.32	1.11	0.03	0.48	0.24	0.98
Notional	2.38	1.72	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	6.16	0.85	8.76	5.18	0.12	2.28	2.37	6.25
Avg USD Trade Size (millions)	114.98	50.88	159.19	101.53	2.14	41.38	40.31	114.09

Continued on next page

Table A2 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
SR	-0.36	0.82	-0.09	0.14	0.11	0.13	0.25	-0.36
ΔSR	-1.18***	-	-0.91***	-0.68***	-0.71**	-0.69**	-0.57**	-1.18***
Mean (%)	-2.31	4.24	-0.81	0.98	0.91	1.04	1.63	-2.98
Vol (%)	6.41	5.17	8.80	6.95	8.18	8.19	6.59	8.35
Skew	-0.59	-0.07	-0.78	-0.76	-0.47	-0.28	-0.46	-0.13
MDD (%)	-56.10	-10.68	-53.08	-27.29	-39.12	-26.47	-21.38	-107.24
Total return	-0.20	2.34	0.07	0.64	0.59	0.64	0.91	-0.32

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from February 1999 to September 2021. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1999 to September 2021.

Table A3: Performance of FX Strategies: Sample Ending in December 2007

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL A: NO PRICE IMPACT								
Before cost SR	1.33	1.29	0.93	0.94	0.35	0.56	0.47	0.66
Before cost Mean (%)	12.47	11.80	7.34	6.05	2.42	3.80	3.29	5.92
Before cost Vol (%)	9.37	9.14	7.93	6.47	6.85	6.80	7.05	9.02
Mean Costs (%)	1.49	0.82	0.41	0.44	0.03	0.20	0.13	0.34
Turnover	2.83	1.74	1.00	0.84	0.04	0.35	0.24	0.83
Notional	4.02	3.83	2.00	2.00	1.00	1.00	2.00	2.00
SR	1.17	1.20	0.87	0.87	0.35	0.53	0.45	0.62
ΔSR	-0.03	-	-0.33*	-0.34*	-0.85**	-0.67**	-0.75**	-0.58*
Mean (%)	10.96	10.97	6.93	5.61	2.40	3.60	3.16	5.58
Vol (%)	9.35	9.12	7.92	6.47	6.86	6.81	7.04	9.02
Skew	-1.02	-0.96	-0.94	-0.73	-0.19	-0.50	-0.10	-0.41
MDD (%)	-26.43	-28.17	-32.27	-31.43	-30.76	-14.85	-22.40	-19.90
Total return	20.86	20.97	8.73	6.61	2.95	4.02	3.56	6.30
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	1.33	1.22	0.93	0.94	0.35	0.56	0.47	0.66
Before cost Mean (%)	12.47	10.63	7.34	6.05	2.42	3.80	3.29	5.92
Before cost Vol (%)	9.37	8.70	7.93	6.47	6.85	6.80	7.05	9.02

Continued on next page

Table A3 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	7.44	0.83	1.63	0.92	0.03	0.31	0.27	1.13
Turnover	2.83	0.92	1.00	0.84	0.04	0.35	0.24	0.83
Notional	4.02	3.46	2.00	2.00	1.00	1.00	2.00	2.00
Avg Relative Trade Size (%)	1.78	0.38	0.72	0.51	0.01	0.20	0.11	0.40
Avg USD Trade Size (millions)	64.14	29.63	40.35	18.55	0.43	6.48	6.69	26.11
SR	0.50	1.13	0.72	0.79	0.35	0.51	0.43	0.53
ΔSR	-0.62***	-	-0.41**	-0.34	-0.78**	-0.62**	-0.70**	-0.60*
Mean (%)	4.92	9.79	5.69	5.12	2.39	3.49	3.02	4.79
Vol (%)	9.76	8.68	7.91	6.46	6.86	6.81	7.01	9.02
Skew	-1.31	-1.15	-0.94	-0.72	-0.19	-0.50	-0.11	-0.38
MDD (%)	-38.32	-25.26	-33.30	-31.91	-30.78	-14.88	-22.57	-20.24
Total return	5.25	16.51	6.58	5.90	2.95	3.90	3.43	5.23

PANEL C: $AUM_0 = 1$ billion

Before cost SR	1.33	1.08	0.93	0.94	0.35	0.56	0.47	0.66
Before cost Mean (%)	12.47	8.51	7.34	6.05	2.42	3.80	3.29	5.92
Before cost Vol (%)	9.37	7.90	7.93	6.47	6.85	6.80	7.05	9.02
Mean Costs (%)	16.73	0.89	6.95	3.88	0.06	1.21	1.38	5.26
Turnover	2.83	0.42	1.00	0.84	0.04	0.35	0.24	0.83
Notional	4.02	2.88	2.00	2.00	1.00	1.00	2.00	2.00
Avg Relative Trade Size (%)	4.36	1.00	3.72	3.65	0.12	1.81	0.98	2.53
Avg USD Trade Size (millions)	204.01	93.08	239.76	138.19	4.31	59.65	61.11	185.61

Continued on next page

Table A3 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
SR	-0.36	0.96	0.04	0.33	0.34	0.37	0.27	0.07
ΔSR	-1.32***	-	-0.92***	-0.63***	-0.61*	-0.58*	-0.68**	-0.89***
Mean (%)	-4.54	7.61	0.32	2.14	2.36	2.58	1.89	0.62
Vol (%)	12.54	7.95	8.34	6.54	6.86	6.88	6.92	9.30
Skew	-2.10	-1.07	-0.87	-0.62	-0.19	-0.48	-0.09	-0.30
MDD (%)	-147.27	-29.45	-41.50	-36.11	-31.01	-15.71	-24.01	-23.89
Total return	-0.18	10.30	1.51	2.73	2.92	3.06	2.51	1.67

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from February 1986 to December 2007. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to December 2007.

Table A4: Performance of FX Strategies: Sample Starting in January 2008

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL A: NO PRICE IMPACT								
Before cost SR	0.71	0.68	0.67	0.66	0.02	0.14	0.41	-0.47
Before cost Mean (%)	3.60	3.43	5.65	4.65	0.17	1.13	2.80	-3.61
Before cost Vol (%)	5.08	5.07	8.45	7.10	8.35	8.19	6.76	7.67
Mean Costs (%)	0.48	0.26	0.26	0.32	0.01	0.18	0.07	0.24
Turnover	1.65	1.14	1.31	1.10	0.02	0.48	0.22	1.11
Notional	1.86	1.81	2.00	2.00	1.00	0.98	2.00	2.00
SR	0.61	0.62	0.64	0.61	0.02	0.12	0.40	-0.50
ΔSR	-0.01	-	0.01	-0.01	-0.61*	-0.51	-0.22	-1.13***
Mean (%)	3.11	3.17	5.39	4.33	0.16	0.95	2.73	-3.85
Vol (%)	5.07	5.07	8.44	7.09	8.35	8.20	6.75	7.67
Skew	-0.45	-0.61	-0.47	-0.62	-0.20	-0.38	-0.36	-0.49
MDD (%)	-8.48	-9.40	-24.45	-17.73	-39.09	-13.57	-17.75	-49.29
Total return	0.54	0.56	1.01	0.79	0.06	0.17	0.49	-0.37
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	0.71	0.63	0.67	0.66	0.02	0.14	0.41	-0.47
Before cost Mean (%)	3.60	3.16	5.65	4.65	0.17	1.13	2.80	-3.61
Before cost Vol (%)	5.08	5.00	8.45	7.10	8.35	8.19	6.76	7.67

Continued on next page

Table A4 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	1.12	0.28	1.00	0.66	0.02	0.24	0.18	0.52
Turnover	1.65	0.85	1.31	1.10	0.02	0.48	0.22	1.11
Notional	1.86	1.65	2.00	2.00	1.00	0.98	2.00	2.00
Avg Relative Trade Size (%)	0.62	0.17	1.10	0.47	0.01	0.15	0.17	0.41
Avg USD Trade Size (millions)	11.60	6.02	19.66	9.15	0.15	2.78	2.96	7.76
SR	0.49	0.57	0.55	0.56	0.02	0.11	0.39	-0.54
ΔSR	-0.09	-	-0.02	-0.01	-0.56	-0.47	-0.19	-1.11**
Mean (%)	2.47	2.87	4.64	3.98	0.16	0.89	2.62	-4.13
Vol (%)	5.08	5.01	8.44	7.08	8.35	8.20	6.75	7.68
Skew	-0.53	-1.51	-0.48	-0.62	-0.20	-0.38	-0.38	-0.48
MDD (%)	-9.53	-10.00	-25.41	-18.14	-39.09	-13.57	-18.00	-52.78
Total return	0.42	0.50	0.83	0.72	0.06	0.16	0.47	-0.39

PANEL C: $AUM_0 = 1$ billion

Before cost SR	0.71	0.41	0.67	0.66	0.02	0.14	0.41	-0.47
Before cost Mean (%)	3.60	1.81	5.65	4.65	0.17	1.13	2.80	-3.61
Before cost Vol (%)	5.08	4.46	8.45	7.10	8.35	8.19	6.76	7.67
Mean Costs (%)	5.32	0.24	5.70	3.21	0.04	0.83	1.16	2.68
Turnover	1.65	0.45	1.31	1.10	0.02	0.48	0.22	1.11
Notional	1.86	1.37	2.00	2.00	1.00	0.98	2.00	2.00
Avg Relative Trade Size (%)	4.61	0.53	7.75	3.96	0.08	1.47	1.61	3.56
Avg USD Trade Size (millions)	85.95	28.57	141.16	76.76	1.47	26.64	27.36	67.30

Continued on next page

Table A4 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
SR	-0.33	0.35	-0.01	0.20	0.02	0.04	0.24	-0.81
ΔSR	-0.68***	-	-0.36	-0.15	-0.34	-0.32	-0.11	-1.16***
Mean (%)	-1.74	1.58	-0.09	1.42	0.13	0.30	1.63	-6.27
Vol (%)	5.28	4.47	8.51	7.07	8.36	8.24	6.78	7.77
Skew	-0.96	-1.86	-0.53	-0.61	-0.20	-0.39	-0.50	-0.46
MDD (%)	-32.84	-9.11	-35.93	-20.77	-39.11	-13.58	-20.10	-79.67
Total return	-0.17	0.28	0.01	0.24	0.05	0.08	0.29	-0.54

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from January 2008 to September 2021. The sample includes 13 developed and 13 emerging currencies against the USD for the period from January 2008 to September 2021.

Table A5: NBER Recessions

Strategies	NBER Recessions			non-NBER Recessions		
	Before TC	After TC		Before TC	After TC	
	SR	SR	Δ SR	SR	SR	Δ SR
PANEL A: NO PRICE IMPACT						
MV	0.82	0.62	0.04	1.10	0.98	-0.03
MVTC	0.70	0.58	-	1.07	1.00	-
HML	0.28	0.23	-0.35	0.77	0.73	-0.28
RB	0.68	0.61	0.03	0.74	0.68	-0.32*
DOL	-0.18	-0.18	-0.76	0.17	0.17	-0.83***
DDOL	0.03	-0.01	-0.59	0.40	0.37	-0.63***
VAL	1.16	1.13	0.55	0.37	0.36	-0.64***
MOM	0.11	0.08	-0.50	0.32	0.29	-0.71***
PANEL B: $AUM_0 = 100$ millions						
MV	0.82	-0.06	-0.74	1.10	0.36	-0.49***
MVTC	0.78	0.68	-	0.93	0.85	-
HML	0.28	-0.17	-0.85	0.77	0.47	-0.38**
RB	0.68	0.36	-0.32	0.74	0.54	-0.31**
DOL	-0.18	-0.18	-0.86	0.17	0.17	-0.68***
DDOL	0.03	-0.07	-0.75	0.40	0.35	-0.50**
VAL	1.16	0.98	0.30	0.37	0.33	-0.52***
MOM	0.11	-0.10	-0.78	0.32	0.15	-0.70***
PANEL C: $AUM_0 = 1$ billion						
MV	0.82	-0.76	-1.20**	1.10	-0.28	-1.01***
MVTC	0.52	0.45	-	0.82	0.73	-
HML	0.28	-0.93	-1.38**	0.77	-0.09	-0.82***
RB	0.68	-0.47	-0.93	0.74	0.03	-0.70***
DOL	-0.18	-0.19	-0.64	0.17	0.16	-0.56**
DDOL	0.03	-0.46	-0.91	0.40	0.19	-0.53**
VAL	1.16	-0.16	-0.61	0.37	0.14	-0.59**
MOM	0.11	-0.82	-1.27**	0.32	-0.28	-1.01***

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 the strategy on the current row and *MVTC* (After TC Δ SR). Standard errors of Δ SR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Columns 2 to 4 report results for the subsample of months during NBER recession, while columns 5 to 7 use the subsample of non-recession months. The strategies are described in Section 3. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

D.2 Alternative Estimates of the Price Impact

Table A6: Performance of FX Strategies: Price Impact with Time-Invariant Amihud Measure

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	1.11	0.88	0.71	0.74	0.12	0.35	0.47	0.30
Before cost Mean (%)	8.90	6.47	5.98	5.05	0.94	2.69	3.28	2.58
Before cost Vol (%)	8.02	7.32	8.40	6.87	7.79	7.70	6.93	8.70
Mean Costs (%)	5.82	0.59	2.35	1.30	0.02	0.36	0.34	1.37
Turnover	2.35	0.71	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.19	2.71	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	3.16	0.61	3.44	1.55	0.02	0.49	0.53	1.74
Avg USD Trade Size (millions)	65.98	31.00	65.52	31.67	0.38	9.41	9.95	35.40
SR	0.37	0.80	0.43	0.54	0.12	0.30	0.42	0.14
ΔSR	-0.43***	-	-0.37**	-0.26	-0.69***	-0.50**	-0.38*	-0.66***
Mean (%)	3.00	5.87	3.62	3.74	0.92	2.33	2.93	1.21
Vol (%)	8.01	7.30	8.43	6.88	7.79	7.71	6.91	8.73
Skew	-1.03	-0.93	-0.80	-0.68	-0.43	-0.33	-0.18	-0.31
MDD (%)	-31.79	-26.43	-33.18	-31.87	-39.09	-17.98	-22.55	-70.52
Total return	5.79	17.60	7.38	8.16	2.42	4.56	6.02	2.64
PANEL C: $AUM_0 = 1$ billion								
Before cost SR	1.11	0.76	0.71	0.74	0.12	0.35	0.47	0.30
Before cost Mean (%)	8.90	4.57	5.98	5.05	0.94	2.69	3.28	2.58
Before cost Vol (%)	8.02	6.00	8.40	6.87	7.79	7.70	6.93	8.70

Continued on next page

Table A6 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	11.77	0.49	7.24	4.85	0.03	1.54	1.84	5.20
Turnover	2.35	0.29	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.19	2.14	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	5.94	1.25	11.73	7.77	0.17	3.98	4.10	7.99
Avg USD Trade Size (millions)	146.47	81.41	248.30	166.08	3.81	77.00	77.75	177.79
SR	-0.34	0.68	-0.15	0.02	0.12	0.15	0.20	-0.30
ΔSR	-1.02***	-	-0.83***	-0.65***	-0.56**	-0.53**	-0.47**	-0.97***
Mean (%)	-3.07	4.08	-1.31	0.18	0.90	1.15	1.42	-2.63
Vol (%)	9.06	6.02	8.64	7.01	7.79	7.81	6.93	8.91
Skew	-1.51	-0.67	-0.78	-0.65	-0.43	-0.36	-0.22	-0.25
MDD (%)	-127.39	-29.64	-94.01	-52.06	-39.14	-37.32	-23.87	-126.93
Total return	-0.19	9.38	0.52	1.67	2.40	2.72	3.16	-0.04

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. Panel A is omitted as it is identical to panel A in table 1. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. The price impact is based on the time-invariant Amihud measure as described in section A.2.1. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from February 1986 to September 2021. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Table A7: Performance of FX Strategies: Time-Series Extrapolation of Volume Data

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	1.10	0.93	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	6.52	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	6.98	8.43	6.88	7.79	7.69	6.97	8.68
Mean Costs (%)	7.09	0.55	3.25	1.76	0.02	0.54	0.65	1.99
Turnover	2.36	0.45	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	2.51	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	9.39	4.91	11.44	6.44	0.09	2.09	2.11	6.97
Avg USD Trade Size (millions)	41.09	21.46	50.04	28.16	0.41	9.15	9.24	30.51
SR	0.20	0.86	0.33	0.48	0.12	0.28	0.37	0.07
ΔSR	-0.66***	-	-0.53***	-0.38**	-0.74***	-0.58**	-0.49**	-0.79***
Mean (%)	1.69	5.97	2.78	3.33	0.94	2.16	2.54	0.63
Vol (%)	8.34	6.95	8.53	6.91	7.79	7.74	6.93	8.69
Skew	-1.38	-0.81	-0.85	-0.70	-0.44	-0.36	-0.21	-0.32
MDD (%)	-50.72	-26.68	-35.83	-33.26	-39.09	-20.38	-22.83	-74.46
Total Return	3.28	18.43	5.25	6.93	2.44	4.24	5.12	1.98
PANEL C: $AUM_0 = 1$ billion								
Before cost SR	1.10	0.72	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	4.17	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	5.80	8.43	6.88	7.79	7.69	6.97	8.68

Continued on next page

Table A7 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	13.14	0.53	8.71	6.03	0.07	2.55	3.33	6.61
Turnover	2.36	0.16	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	2.12	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	18.09	9.95	31.27	24.48	0.93	14.13	12.70	25.49
Avg USD Trade Size (millions)	79.13	43.53	136.84	107.11	4.07	61.82	55.57	111.53
SR	-0.42	0.62	-0.30	-0.13	0.12	0.02	-0.02	-0.45
ΔSR	-1.04***	-	-0.92***	-0.76***	-0.51**	-0.60**	-0.65***	-1.08***
Mean (%)	-4.45	3.64	-2.72	-0.96	0.90	0.15	-0.17	-4.02
Vol (%)	10.58	5.85	9.00	7.21	7.79	8.09	7.24	8.85
Skew	-2.61	-0.70	-0.88	-0.75	-0.44	-0.56	-0.42	-0.33
MDD (%)	-182.23	-32.41	-127.31	-79.13	-39.14	-55.60	-50.15	-148.65
Total Return	-0.52	7.96	-0.08	0.80	2.39	1.61	1.38	-0.41

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. Panel A is omitted as it is identical to panel A in table 1. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. The price impact is based on the Amihud measure constructed from the extrapolated panel (cross-section and time-series) of CLS volume data as described in section A.2.2. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from February 1986 to September 2021. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Table A8: Performance of FX Strategies: Price Impact based on BIS Triennial Survey Volume Data

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	1.10	0.96	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	7.10	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	7.37	8.43	6.88	7.79	7.69	6.97	8.68
Mean Costs (%)	7.33	0.62	2.59	1.69	0.02	0.45	0.50	1.49
Turnover	2.36	0.86	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	2.72	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	1.50	1.10	2.73	1.31	0.02	0.49	0.51	1.63
Avg USD Trade Size (millions)	33.68	48.82	52.44	26.76	0.41	9.23	9.35	32.86
SR	0.15	0.88	0.41	0.49	0.12	0.29	0.39	0.13
ΔSR	-0.73***	-	-0.47***	-0.38**	-0.75***	-0.59**	-0.49**	-0.75***
Mean (%)	1.39	6.47	3.44	3.40	0.94	2.24	2.68	1.13
Vol (%)	9.28	7.40	8.48	6.90	7.79	7.73	6.94	8.68
Skew	-2.09	-1.12	-0.80	-0.70	-0.44	-0.35	-0.18	-0.33
MDD (%)	-79.82	-27.54	-34.81	-33.14	-39.09	-19.10	-23.19	-62.74
Total return	2.75	21.82	6.86	7.14	2.44	4.40	5.43	2.53
PANEL C: $AUM_0 = 1$ billion								
Before cost SR	1.10	0.76	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	5.29	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	6.95	8.43	6.88	7.79	7.69	6.97	8.68

Continued on next page

Table A8 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	12.72	0.66	7.63	5.84	0.07	2.08	2.73	5.50
Turnover	2.36	0.46	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	2.41	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	2.33	2.66	6.32	4.11	0.20	3.30	2.96	5.78
Avg USD Trade Size (millions)	64.20	159.65	139.71	91.36	4.06	64.07	56.32	130.77
SR	-0.29	0.66	-0.18	-0.11	0.12	0.08	0.06	-0.32
ΔSR	-0.95***	-	-0.83***	-0.76***	-0.54**	-0.58**	-0.60**	-0.98***
Mean (%)	-4.18	4.63	-1.65	-0.78	0.90	0.61	0.43	-2.91
Vol (%)	14.51	7.04	9.39	7.37	7.80	7.99	7.28	9.12
Skew	-7.43	-1.37	-0.97	-0.66	-0.44	-0.50	-0.39	-0.53
MDD (%)	-186.02	-39.39	-131.67	-100.54	-39.10	-39.09	-31.00	-114.37
Total return	-0.61	11.26	0.32	0.90	2.39	2.06	1.94	-0.14

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. Panel A is omitted as it is identical to panel A in table 1. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. The price impact is based on the Amihud measure constructed from the extrapolated panel (cross-section and time-series) of BIS Triennial Survey volume data as described in section A.2.3. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from February 1986 to September 2021. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Table A9: Performance of FX Strategies: Cutting Costs in Half

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
PANEL A: NO PRICE IMPACT								
Before cost SR	1.10	1.07	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	8.61	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	8.08	8.43	6.88	7.79	7.69	6.97	8.68
Mean Costs (%)	0.55	0.39	0.18	0.20	0.01	0.10	0.06	0.15
Turnover	2.36	1.82	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	3.07	2.00	2.00	1.00	0.99	2.00	2.00
SR	1.03	1.02	0.70	0.71	0.12	0.34	0.45	0.29
ΔSR	0.01	-	-0.32*	-0.31*	-0.90***	-0.68***	-0.57**	-0.73***
Mean (%)	8.33	8.22	5.86	4.90	0.95	2.59	3.14	2.47
Vol (%)	8.06	8.06	8.43	6.88	7.79	7.70	6.96	8.68
Skew	-0.79	-0.86	-0.77	-0.68	-0.44	-0.33	-0.17	-0.33
MDD (%)	-24.73	-29.08	-31.80	-30.87	-39.07	-17.53	-22.26	-48.68
Total return	41.28	39.83	17.13	12.63	2.46	5.10	6.52	4.62
PANEL B: $AUM_0 = 100$ millions								
Before cost SR	1.10	0.96	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	7.33	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	7.61	8.43	6.88	7.79	7.69	6.97	8.68

Continued on next page

Table A9 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
Mean Costs (%)	4.27	0.52	1.62	0.81	0.01	0.21	0.23	0.86
Turnover	2.36	0.99	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	2.79	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	4.73	1.13	4.08	1.76	0.02	0.52	0.55	1.99
Avg USD Trade Size (millions)	95.31	52.29	76.77	35.55	0.41	9.80	10.07	39.65
SR	0.56	0.90	0.52	0.62	0.12	0.32	0.43	0.20
ΔSR	-0.34***	-	-0.37**	-0.27*	-0.77***	-0.57**	-0.47**	-0.69***
Mean (%)	4.56	6.80	4.41	4.29	0.95	2.48	2.96	1.77
Vol (%)	8.15	7.59	8.46	6.90	7.79	7.71	6.95	8.69
Skew	-1.01	-0.95	-0.80	-0.69	-0.44	-0.34	-0.19	-0.32
MDD (%)	-27.47	-24.91	-32.32	-31.12	-39.07	-18.03	-22.34	-63.29
Total return	10.57	24.41	10.00	10.03	2.46	4.87	6.08	3.40

PANEL C: $AUM_0 = 1$ billion

Before cost SR	1.10	0.87	0.72	0.74	0.12	0.35	0.46	0.30
Before cost Mean (%)	8.88	6.07	6.04	5.10	0.96	2.69	3.19	2.63
Before cost Vol (%)	8.07	7.00	8.43	6.88	7.79	7.69	6.97	8.68
Mean Costs (%)	9.95	0.53	5.98	3.71	0.02	1.07	1.44	3.99
Turnover	2.36	0.45	1.12	0.94	0.03	0.41	0.24	0.94
Notional	3.15	2.41	2.00	2.00	1.00	0.99	2.00	2.00
Avg Relative Trade Size (%)	10.28	2.66	16.18	10.10	0.20	4.43	4.46	10.85
Avg USD Trade Size (millions)	246.37	158.76	330.82	211.40	4.11	84.70	82.69	230.18

Continued on next page

Table A9 – continued from previous page

	<i>MV</i>	<i>MVTC</i>	<i>HML</i>	<i>RB</i>	<i>DOL</i>	<i>DDOL</i>	<i>VAL</i>	<i>MOM</i>
SR	-0.13	0.79	0.00	0.20	0.12	0.21	0.25	-0.16
ΔSR	-0.92***	-	-0.79***	-0.60***	-0.67***	-0.59**	-0.54***	-0.95***
Mean (%)	-1.21	5.54	0.03	1.38	0.94	1.62	1.74	-1.38
Vol (%)	9.58	6.98	8.75	7.07	7.79	7.82	6.98	8.80
Skew	-2.03	-1.07	-0.90	-0.77	-0.44	-0.41	-0.27	-0.27
MDD (%)	-102.74	-26.45	-73.74	-39.56	-39.10	-30.73	-23.08	-114.66
Total return	0.51	15.73	1.40	3.04	2.44	3.36	3.64	0.49

Notes: The table reports summary statistics of monthly excess returns of our strategies in the FX market. The strategies are described in Section 3. All trading costs are cut in half as a conservative resolution of the concern that table 1 does not allow for triangular trades (see section A.2.4 for details). Panel A summarizes the performance when there are only proportional costs. Panel B and C report results when the initial AUM of a fund is USD 100 millions respectively 1 billion and costs account for the price impact of trading. Before cost SR, Mean (%) and Vol (%) measure the annualized Sharpe ratio, and annualized average and volatility of the excess returns (reported in percentage points). Mean Costs (%) measures the average annualized trading costs as a percentage of the AUM. Turnover and Notional report the monthly turnover and the notional value as a fraction of the AUM. Avg Relative Trade Size (%) and Avg USD Trade Size (millions) measures the average amount traded per month (average across currency-month observations with non-trivial trade) as a percentage of the average daily trading volume in the market respectively in absolute terms as millions of USD. SR is the annualized after cost Sharpe ratio. ΔSR is the difference in the after-cost Sharpe ratio of a strategy and that of *MVTC*. Standard errors of ΔSR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Mean (%) and Vol (%) are the annualized average and volatility of the after cost excess returns (reported in percentage points). Skew and Kurt are the monthly skewness and kurtosis of the after cost excess returns, MDD (%) is the Maximum Draw Down (measured in percentage points). Total return is the return (including the risk-free rate) from February 1986 to September 2021. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

D.3 Sensitivity Analyses of *MVTC*

Table A10: Event Study of Crash Risks in Turbulent Times

Date	Nature of the Event			No PI		$AUM_0 = 100$ mills		$AUM_0 = 1$ bill	
	Liquidity	Volatility	Crises	MVTC	BH	MVTC	BH	MVTC	BH
Sep-92	1	1	1	-12.25	-24.88	-8.27	-21.14	-11.23	-27.73
Jul-97	0	0	1	-13.09	-20.35	-13.52	-19.76	-14.68	-14.03
Sep-08	0	0	1	-1.02	-4.8	-0.46	-1.54	-0.87	-1.42
Oct-08	1	0	0	-1.02	-1.97	-0.46	-1	-0.87	-1.15
Oct-09	1	0	0	-0.63	-0.72	-0.66	-0.93	-0.54	-2.06
May-10	0	1	0	-2.36	-3.39	-2.08	-2.08	-1.28	-1.27
Oct-10	1	0	0	-5.68	-8.92	-6.18	-4.34	-4.34	-2.83
Sep-11	0	1	0	-2.59	-3.77	-2.25	-4.49	-2.42	-3.12
Mar-20	1	1	1	-0.45	-2.06	-1.52	-2.16	-2.31	-3.16
Mean				-4.34	-7.87	-3.93	-6.38	-4.28	-6.31

Notes: The table reports Maximum Draw Downs (MDD (%); measured in percentage points) for the time period starting 1 month before and ending 12 months after the turbulent month under investigation on each row. Turbulent months are defined as months with high volatility, illiquidity, or a known crisis as described in Figure A9. Each row corresponds to a separate turbulent month. The last row reports the average MDD across the different events (Mean). Column 1 identifies the date of the turbulent month. Column 2 to 4 indicates with 1 whether the turbulent month corresponds to a spike in volatility, illiquidity, or a crisis. Columns 5, 7 and 9 report the MDD from *MVTC* and Columns 6, 8 and 10 the MDD of *BH*. *MVTC* is the mean-variance-transaction-cost optimized portfolio described in section 3. *BH* is a buy and hold strategy that invests in the position of *MVTC* at the beginning of the event sample and does not rebalance the initial position until the end of the event sample. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Table A11: Sensitivity to the Risk Aversion Coefficient λ

Risk Aversion Coefficient λ	Before TC SR		After TC SR			<i>MVTC</i> After TC	
	<i>MV</i>	<i>MVTC</i>	<i>MV</i>	<i>MVTC</i>	Δ SR	Vol (%)	MDD (%)
PANEL A: NO PRICE IMPACT							
5	1.10	1.06	0.96	0.98	-0.02	39.07	-150.64
10	1.10	1.06	0.97	0.99	-0.02	19.58	-72.18
25	1.10	1.06	0.97	0.99	-0.02	7.86	-28.17
50	1.10	1.06	0.97	0.99	-0.02	3.94	-14.00
100	1.10	1.06	0.97	0.99	-0.02	1.97	-6.97
200	1.10	1.06	0.97	0.99	-0.02	0.99	-3.48
PANEL B: $AUM_0 = 100$ millions							
5	1.10	0.68	0.19	0.57	-0.38**	28.35	-169.71
10	1.10	0.83	0.09	0.75	-0.66***	16.10	-69.85
25	1.10	0.93	0.33	0.86	-0.52***	7.40	-25.26
50	1.10	0.99	0.69	0.91	-0.22***	3.83	-12.43
100	1.10	1.05	0.86	0.97	-0.10	1.93	-6.56
200	1.10	1.06	0.92	0.99	-0.06	0.97	-3.41
PANEL C: $AUM_0 = 1$ billion							
5	1.10	0.56	0.21	0.40	-0.19	23.64	-137.83
10	1.10	0.65	-0.17	0.51	-0.69***	13.99	-85.95
25	1.10	0.81	-0.31	0.72	-1.04***	6.60	-29.45
50	1.10	0.88	-0.23	0.80	-1.02***	3.61	-13.13
100	1.10	0.92	0.15	0.84	-0.69***	1.89	-6.41
200	1.10	0.97	0.55	0.89	-0.34***	0.96	-3.09

Notes: The table reports the performance of *MV* and *MVTC* for different values $\lambda \in \{5, 10, 25, 50, 100, 200\}$. *MV* and *MVTC* are the mean-variance respectively mean-variance-transaction-cost optimized portfolios. The strategies are described in section 3. Column 2 and 3 provide Sharpe ratios before transaction costs (Before TC SR), and columns 4 and 5 Sharpe ratios after costs (After TC SR). Column 6 shows the difference in the after cost Sharpe ratios between *MV* and *MVTC*. Standard errors of Δ SR are estimated using block bootstrapping with a block size of 5 observations to account for heteroskedasticity, cross- and auto-correlation (Ledoit and Wolf, 2008). ***, **, * indicate a statistical significance at the 1%, 5%, 10% level. Column 7 and 8 report (in percentage points) the after cost volatility (Vol (%)) and maximum draw-down (MDD (%)) for *MVTC*. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Table A12: **Datastream Mnemonics for Exchange Rate Quotes against USD**

Currency	Spot Rate	Forward Rate
Australian dollar	BBAUDSP	BBAUD1F (USAUD1F)
Brazilian real	BRACRU\$	USBRL1F
British pound	BBGBPSP	BBGBP1F
Canadian dollar	BBCADSP	BBCAD1F
Czech koruna	CZECHC\$	USCZK1F
Euro	BBEURSP	BBEUR1F
French franc	BBFRFSP	BBFRF1F
German mark	BBDEMSP	BBDEM1F
Greek Drachma	GREDDRA\$	USGRD1F
Hungarian forint	HUNFOR\$	USHUF1F
Irish punt	BBIEPSP	BBIEP1F
Italian lira	BBITLSP	BBITL1F
Japanese yen	BBJPYSP	BBJPY1F (USJPY1F)
Mexican peso	MEXPES\$	USMXN1F
Netherland guilder	BBNLGSP	BBNLG1F
New Zealand dollar	BBNZDSP	BBNZD1F (USNZD1F)
Norwegian krone	BBNOKSP	BBNOK1F (USSEK1F)
Polish zloty	POLZLO\$	USPLN1F
Portuguese escudo	PORTES\$	USPTE1F
Singapore dollar	BBSGDSP	BBSGD1F
South Africa rand	BBZARSP	BBZAR1F
South Korean won	KORSWO\$	USKRW1F
Spanish peseta	SPANPE\$	USESP1F
Swedish krona	BBSEKSP	BBSEK1F (USSEK1F)
Swiss franc	BBCHFSP	BBCHF1F (USCHF1F)
Taiwan new dollar	TAIWDO\$	USTWD1F

Notes: The table provides the Datastream mnemonics (data from Barclays and Reuters) for spot and forward exchange rate quotes against the USD. To obtain bid and ask exchange rates, the suffixes (EB) and (EO) are added to the corresponding mnemonics.

Net-Order Flow Extrapolation

$$\ln(\bar{z}_i^{funds}) = 7.31^{***} - 1.35^{***} \times \ln(\overline{PBA}_i^S), R^2 = 0.79$$




Figure A1: The asterisks in the figure plot the natural logarithm of the daily average of the absolute value of the net-order flow (daily buy minus sell orders) of funds $\ln(\bar{z}_i^{funds})$ against the natural logarithm of the average proportional bid-ask spread $\ln(\overline{PBA}_i^S)$ for the 12 currencies $i \in \mathcal{I}$ (AUD, EUR, GBP, NZD, CAD, CHF, JPY, NOK, SEK, MXN, SGD, ZAR). The linear line is the fitted regression $\ln(\bar{z}_i^{funds}) = 9.76 - 1.47 \ln(\overline{PBA}_i^S)$, and the circles on the line represent the fitted values $\ln(\bar{z}_i^{funds})$ corresponding to $\ln(\overline{PBA}_i^S)$ for the 6 currencies BRL, CZK, HUF, KRW, PLN, and TWD (for which CLS order flow data is unavailable). The average is taken over the daily sample between September 3rd 2012 through September 24th 2021.

Country Specific Proportional Costs: Developed Countries

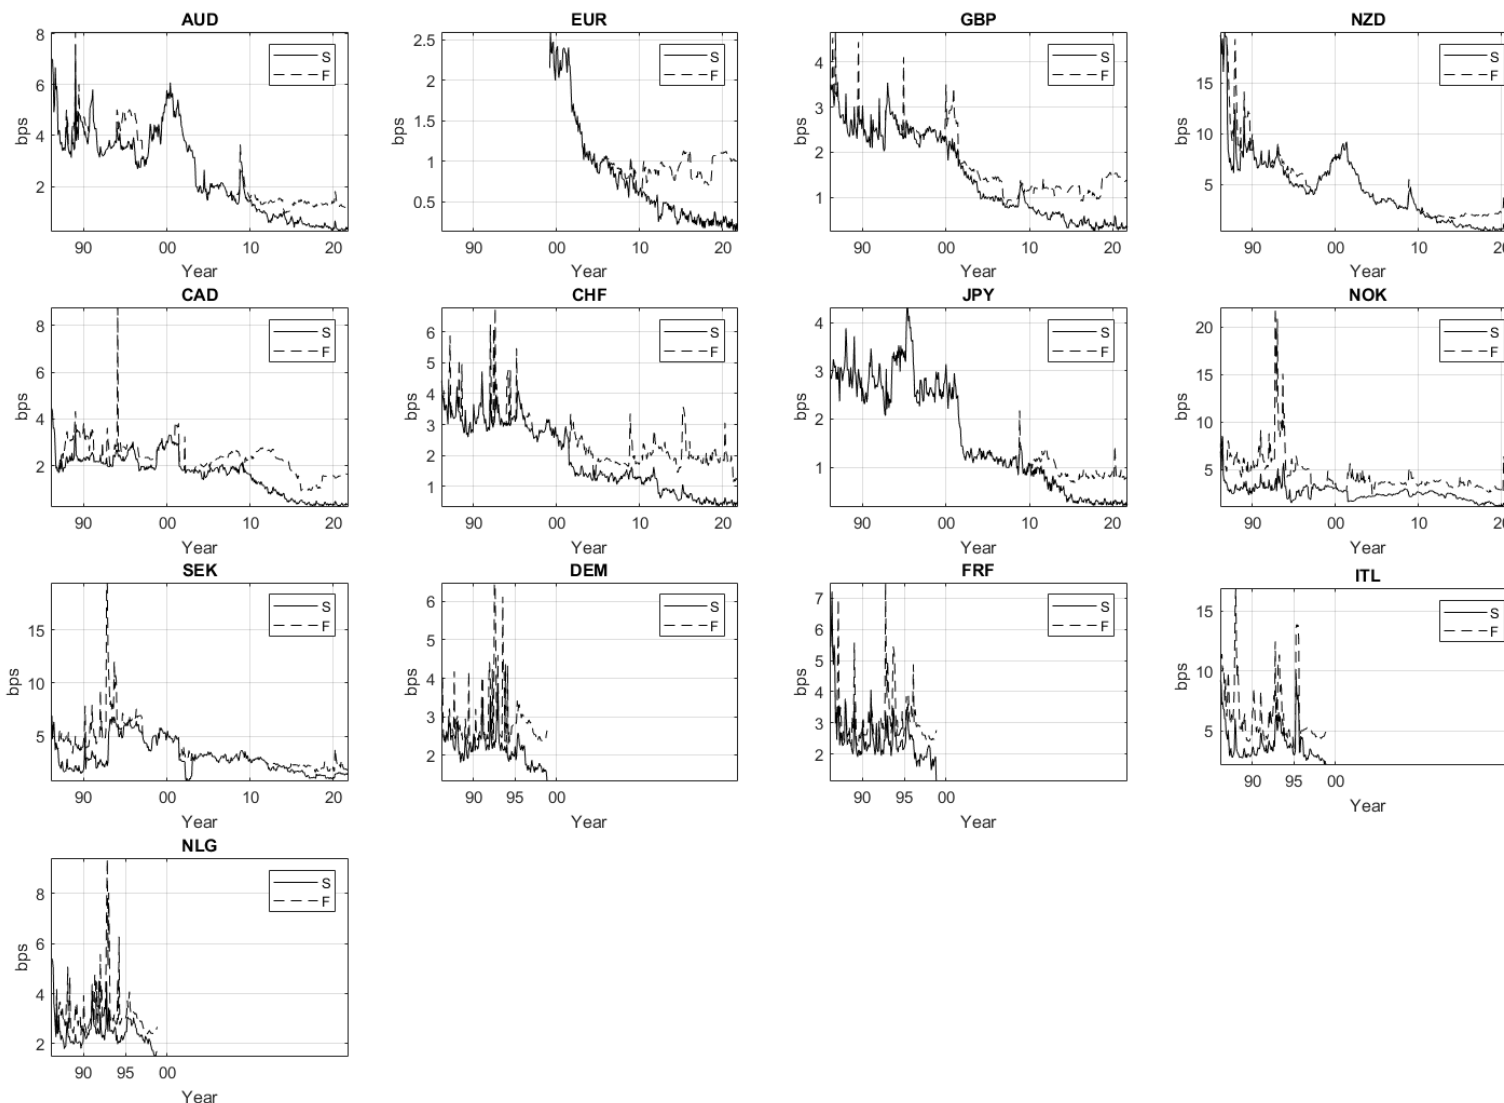


Figure A2: The figures plot the monthly average proportional costs for the spot market $0.5C_{i,t}^{S+} + 0.5C_{i,t}^{P-}$ (sold line), and for the forward market, $0.5C_{i,t}^{P+} + 0.5C_{i,t}^{S-}$ (dashed line). The sample are the 13 developed currencies from February 1986 to September 2021.

Country Specific Proportional Costs: Emerging Countries

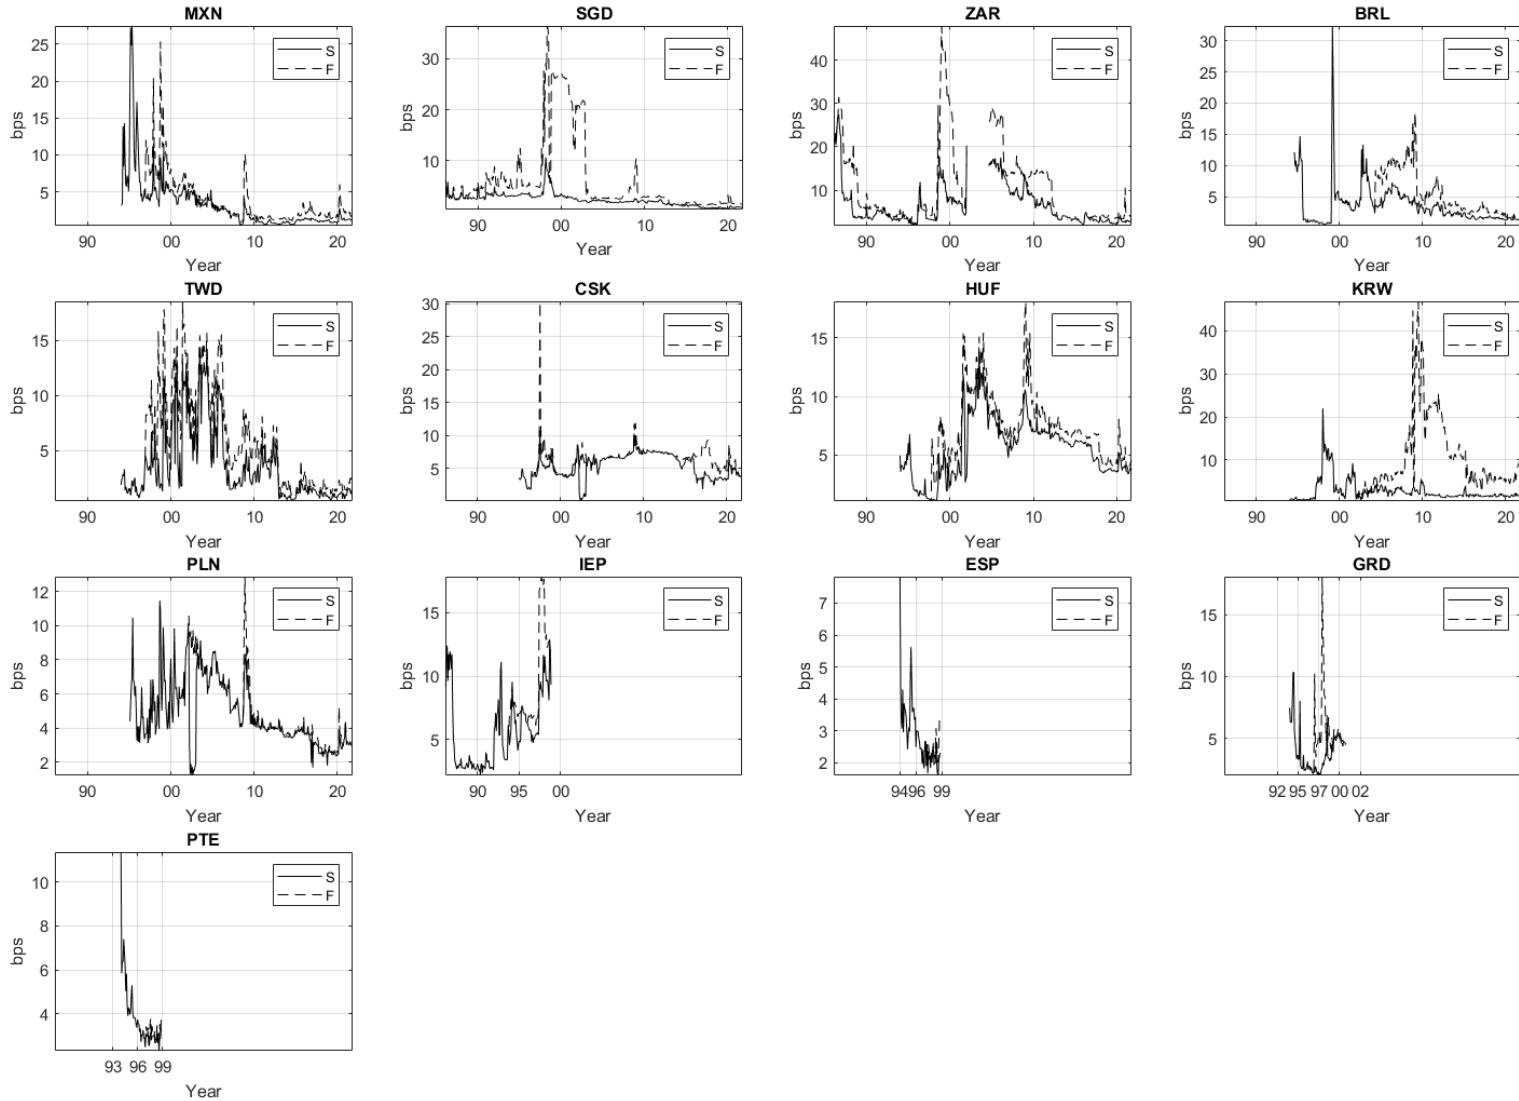


Figure A3: The figures plot the monthly average proportional costs for the spot market $0.5C_{i,t}^{S+} + 0.5C_{i,t}^{P-}$ (sold line), and for the forward market, $0.5C_{i,t}^{P+} + 0.5C_{i,t}^{S-}$ (dashed line). The sample are the 13 emerging currencies from February 1986 to September 2021.

Country Specific Price Impact: Developed Countries

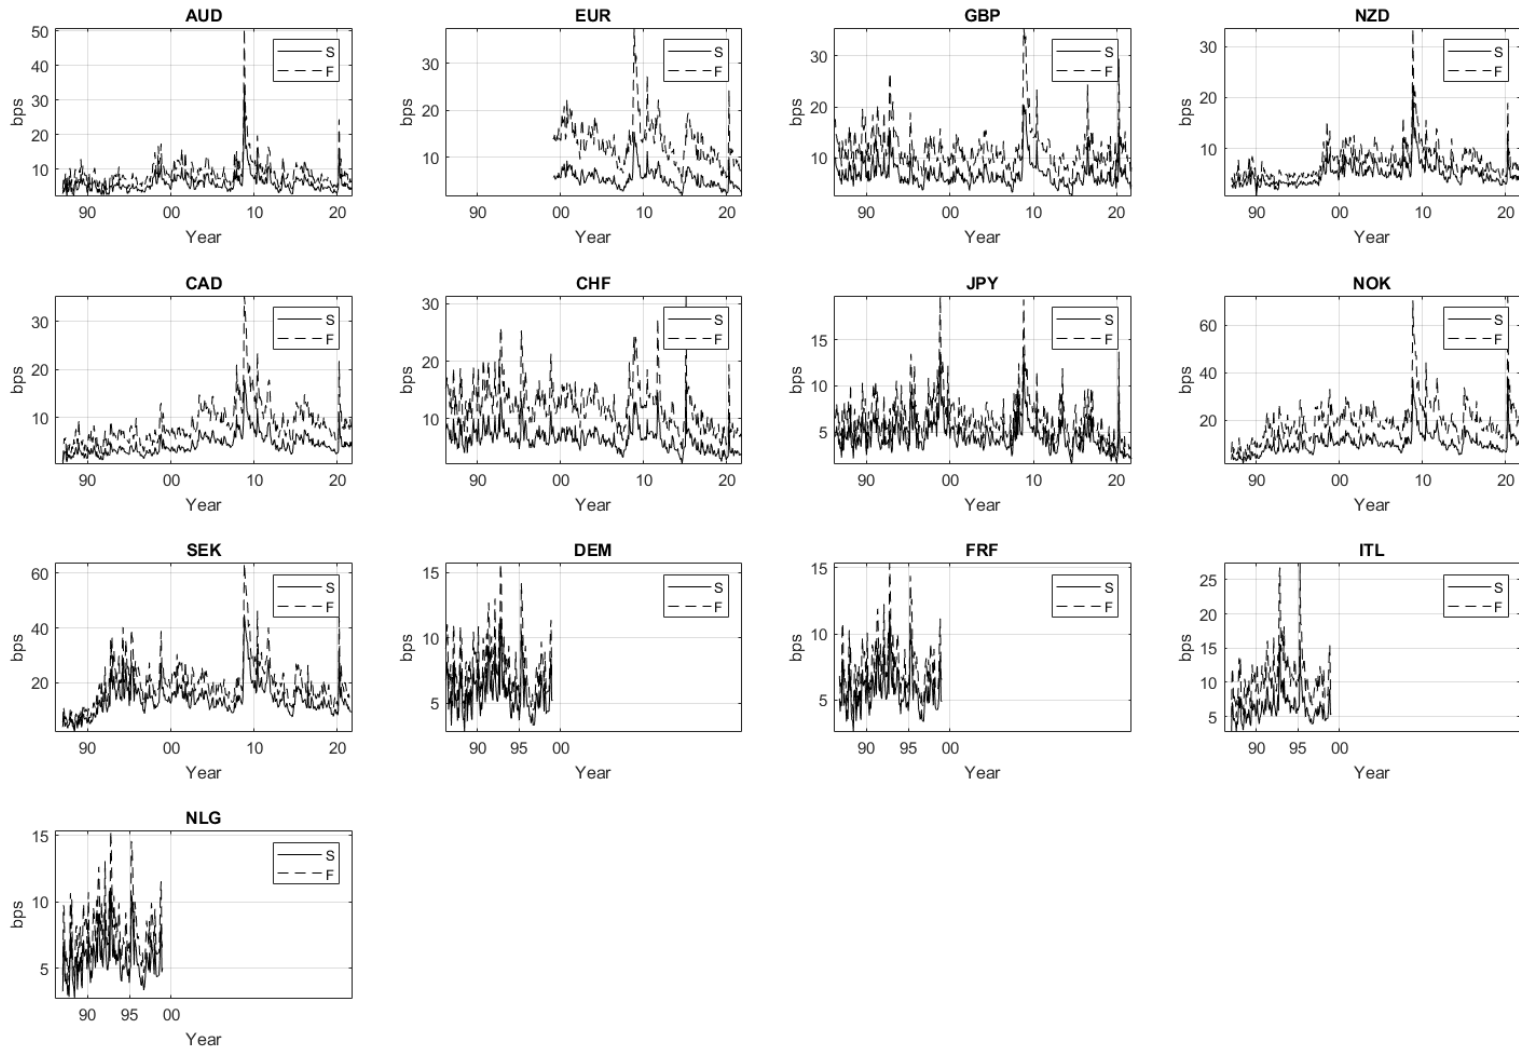


Figure A4: The figures plot the monthly price impact for the average fund in the spot market (solid line), and for the forward market (dashed line). The sample are the 13 developed currencies from February 1986 to September 2021.

Country Specific Price Impact: Emerging Countries

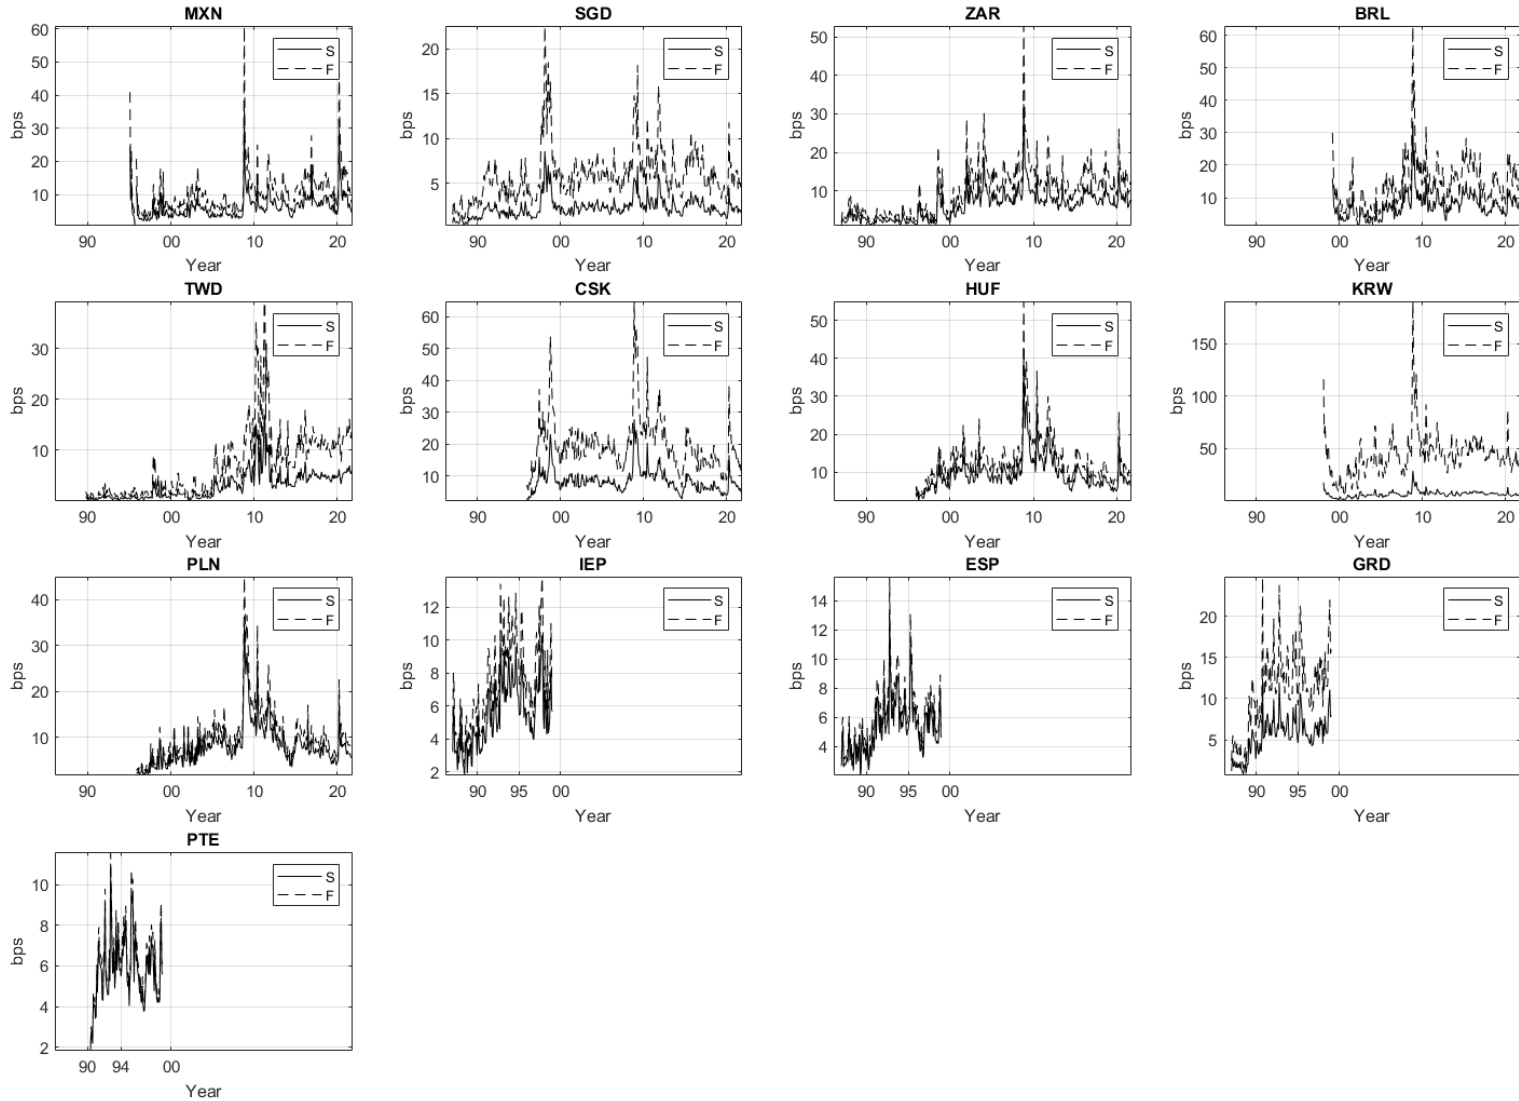


Figure A5: The figures plot the monthly price impact for the average fund in the spot market (solid line), and for the forward market (dashed line). The sample are the 13 emerging currencies from February 1986 to September 2021.

Cumulative Before-Cost Excess Returns




Figure A6: The figures plot the before-cost cumulative excess returns of *MV* (solid line), *MVTC* (bold solid line), *HML* (dashed line), and *MOM* (dotted line). For the construction of *MVTC* we set the price impact to zero. The top graph assumes no price impact, while the middle and bottom graphs assume an initial AUM of USD 100 millions respectively 1 billion. The strategies are described in Section 3. The sample includes 13 developed and 13 emerging currencies against the USD for the period from February 1986 to September 2021.

Trading Volume Panel Extrapolation using CLS Data




Figure A7: The asterisks in the figure plot the natural logarithm of the quarterly volume $\ln(\tilde{v}_{i,q}^S)$ against the natural logarithm of the quarterly proportional bid-ask spread $\ln(\widetilde{PBA}_{i,q}^S)$ for the 12 currencies $i \in \mathcal{I}$ (AUD, EUR, GBP, NZD, CAD, CHF, JPY, NOK, SEK, MXN, SGD, ZAR) from September 2012 to September 2021. The linear line is the fitted regression $\ln(\tilde{v}_{i,q}^S) = 9.76 - 1.47 \ln(\widetilde{PBA}_{i,q}^S)$, and the circles on the line represent the fitted values $\ln(\tilde{v}_{i,q}^S)$ corresponding to $\ln(\widetilde{PBA}_{i,q}^S)$ for all currency-quarter observations for which we observe $\widetilde{PBA}_{i,q}^S$ but not $\tilde{v}_{i,q}^S$. Quarterly observations are constructed as daily averages within the quarter.

Trading Volume Panel Extrapolation using BIS Spot Data




Figure A8: The asterisks in the figure plot the natural logarithm of the annual volume $\ln(\hat{v}_{i,y}^S)$ against the natural logarithm of the annual proportional bid-ask spread $\ln(\widehat{PBA}_{i,y}^S)$ for currency-year observations for which we have BIS Triennial volume survey data. The linear line is the fitted regression $\ln(\hat{v}_{i,y}^S) = 9.76 - 1.47 \ln(\widehat{PBA}_{i,y}^S)$, and the circles on the line represent the fitted values $\ln(\hat{v}_{i,y}^S)$ corresponding to $\ln(\widehat{PBA}_{i,y}^S)$ for all currency-year observations for which we observe $\widehat{PBA}_{i,y}^S$ but not $\hat{v}_{i,y}^S$. Annual observations of the bid-ask spread are constructed as daily averages within the year.

Turbulent Months




Figure A9: The plot shows the time series of the FX market volatility index following [Menkhoff et al. \(2012a\)](#) (dashed-blue), and the FX market illiquidity index following [Karnaukh et al. \(2015\)](#) (orange bold) over the sample period from February 1986 to September 2021. The blue and red dots represent the five highest non-consecutive peaks over a period of 13 months in these time series. Additionally, the vertical black lines report the dates of four known crises involving FX markets: the 1992 Black Wednesday (Sep-1992), the 1997 Asian Financial Crisis (Jul-1997), the 2008 Financial Crisis (Sep-2008), and the COVID crisis (Mar-2020)

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

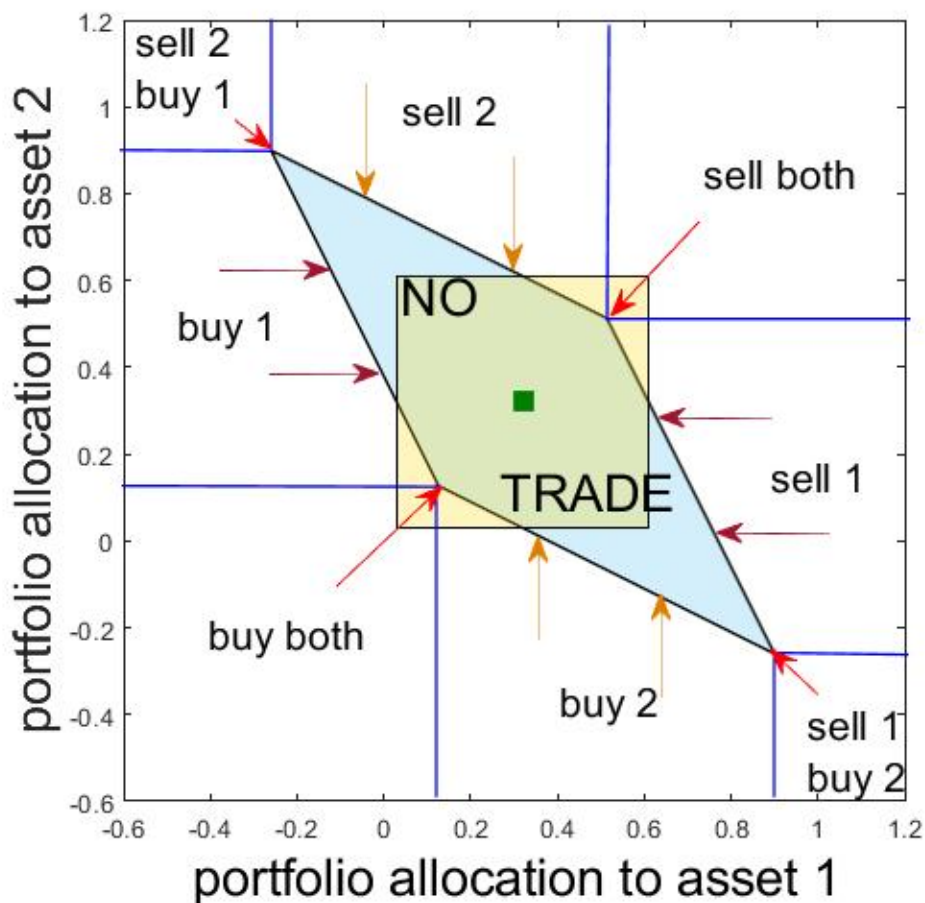


Figure A10: The investment opportunity set consists of two positively correlated risky assets. The horizontal axis measures the weight a portfolio places on asset 1 and the vertical axis the weight on asset 2. The green square is the optimal portfolio θ_t^{MV} if there are no transaction costs. The blue parallelogram illustrates the no trading region of MV_{TC} , which optimizes over transaction costs. The yellow square determines the no trading region of a strategy that 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 it is optimal to trade. In particular, when there are only proportional costs, the investor trades towards θ_t^{MV} until she reaches the boundary of the no trading region as indicated by the arrows. If we additionally model costs from a price impact, it is never optimal to trade all the way to the boundary of the no trading region. Purple, brown or red colors of the arrows indicate that only asset 1, 2 or both assets are traded.