CESifo WORKING PAPERS

8143
2020

March 2020

Retaining Alpha: The Effect of Trade Size and Rebalancing Frequency on FX Strategy Returns

Michael Melvin, Wenqiang Pan, Petra Wikstrom

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 8143*

Retaining Alpha: The Effect of Trade Size and Rebalancing Frequency on FX Strategy Returns

Abstract

The literature on currency investing that incorporates transaction costs uses costs relevant for small trade sizes. Using the entire order book of the major electronic brokerages for FX, we compute sweep-to-fill costs for trades of different sizes and illustrate the reduction in post-cost returns as trade size increases. Researchers should consider trade size *and* frequency to create realistic forecasts of post-tcost returns to gauge the capacity of a strategy. We show how incorporating tcosts in the construction of a portfolio improves performance for both high and low frequency strategies and retains a larger portion of the alpha.

JEL-Codes: G150, F310.

Keywords: transaction costs, FX microstructure, exchange rates, portfolio construction.

Michael Melvin
Rady School of Management
U.C. San Diego
Otterson Hall, 9500 Gilman Dr.
USA – La Jolla, CA 92093-0553
mmelvin@ucsd.edu

Wenqiang Pan
Columbia Business School
Columbia University
USA – New York, NY 10027
wenqiang.pan1994@gmail.com

Petra Wikstrom
Streamline Analytix LLC
USA – New York, NY 10011
petra.maria.w@gmail.com

A revised version of this paper has been accepted for publication at the Journal of Financial Markets.

1. **Introduction**

Preserving alpha occurs when portfolio managers and traders consider accurate trading costs when constructing and executing strategies. A recent *Financial Times* article states "Buy-side traders currently spend the bulk of their time focusing on alpha preservation, often called 'best execution'" (Merrin, 2019). Beyond traders' seeking best execution, portfolio managers should incorporate implementation costs and constraints when turning alpha insights into tradable strategies. Realization of the alpha from a trading insight comes when the portfolio manager generates trades that are feasible to execute in a cost-effective manner. Best practice for both investors and academics is to incorporate accurate transaction costs into simulations or backtests of currency investment strategies. Not all scholars and not all investors will always do so, and this failure to consider trade costs can have large implications for the expected versus realized performance of net-of-trade-cost returns. An investment team may say "we are too small to worry about transaction costs" thinking that only the largest firms, trading very large amounts of currency, need to worry about managing transaction costs (tcosts).[1] Scholars may think that one can either treat currency tcosts as too small to worry about or a constant that can

[1] Even at large firms, sophisticated investors may not appreciate how important trading costs are to a successful FX investment strategy. One of the authors worked at a large asset management firm and managed active currency strategies. More than once, in conversations about microstructure issues and transaction costs, equity investors remarked how lucky he was to be in currencies where one did not have to worry about trading costs.

simply be subtracted from the gross return. We show examples of how much tcosts matter for the success of a currency strategy.

Modeling tcosts all starts with the data. Even careful scholars who seek to incorporate tcosts into their work will typically use bid-ask spread data sampled from the top of the order book. It is common to refer to bid-ask spreads measured by the best bid and ask. However, such measurement will typically only apply to trades of small size, in the order of magnitude of $1 million to $5 million, depending on the currency pair. Larger trades will incur higher costs and, consequently, lower returns, other things equal. In this paper, we analyze the effects of trade size on tcosts and the consequent influence of tcosts on returns to a systematic speculative strategy. We show how post-cost returns fall as trade size rises. We also show how the incorporation of tcosts into the portfolio construction methodology can help preserve the alpha of trading signals. This turns out to be important for both high-turnover momentum strategies and low-turnover value strategies.

A popular portfolio construction technique used in many studies is to rank currencies according to some return forecast and then equal weight them in a portfolio. One problem with this approach is that it ignores the tcosts associated with different currencies and may overweight high-cost currencies and lower post-cost returns. We show how a portfolio optimization that uses the same return forecast but incorporates a tcost penalty can enhance returns relative to the equal-weight portfolio. It is interesting to note that such gains are even relevant for quite small portfolios. Investors can preserve alpha by incorporating tcosts into pre-trade optimizations in order to better size currency positions and related trade size as a function of trade costs.

In addition to trade size, we also show how sensitive returns are to portfolio rebalance frequency. This is illustrated by backtesting momentum strategies with lookback windows ranging from a very fast 1-day lookback, then increasing one day at a time and backtesting each different lookback window out to one year of return history. The results clearly show how a very fast strategy of trading on a few days returns is much too costly for even small portfolios. Post-cost returns are shown to be increasing as the momentum lookback period is lengthened out to about six months. We also examine a low-turnover value strategy to see if pre-trade cost consideration is important for a slow strategy. Even though the value strategy has much less trading, alpha is preserved by incorporating tcosts in a pre-trade optimization. Finally, we also compare a daily rebalance with a monthly rebalance strategy to examine the importance of pre-trade tcost consideration in preserving alpha if only trading once a month. We find that for all but the very smallest portfolios, investors can retain more alpha from their trading signals if tcosts are considered pre-trade. A major lesson is that investors would benefit from considering the likely trade size *and* frequency in order to create realistic forecasts of post-tcost returns to gauge the capacity of a strategy. Such accounting for the costs of trading will tend to result in lower turnover and/or a tilt into more liquid, lower-cost currencies.

One might wonder if trade size is so important any more, given the availability of algorithmic strategies offered by banks and other vendors. Such strategies automate trading so that large trades are broken into smaller sizes and traded over time rather than all at once to reduce trade costs. Of course, the cost reduction through an algorithm must be considered relative to the speed of alpha decay and associated urgency. Our study is based upon aggressing trades that sweep down the order book to fill the desired trade size. Alternatively, one could have a risk-transfer trade with a single broker, which may have better or worse execution, depending on the ability of the broker to internalize the trade. Finally, one could, for example, execute passively with an algorithmic strategy that posts limit orders in order to earn the spread. Importantly, a recent study of buy-side foreign exchange traders by Greenwich Associates (2017) finds that on the largest trading volume FX desks, trading more than $50 billion a year, 23% use algorithmic trading (algos) and smaller volume desks use algos less. For institutional fund managers, the largest funds, with annual FX volumes exceeding $50 billion, had 34% using algos. Again, smaller funds use algos less frequently.[2] This should be placed in context with the trend of increased electronification and growth potential of automated execution in certain segments of the market, while some products and markets are still rather manual. This paper includes a subset of the FX spot market with G10 and deliverable EM currencies that are highly electronic to semi-electronic in trading. It focuses on estimating the value of incorporating tcosts in the construction of a systematic trading strategy through aggressing all orders of sizes ranging between $1 million and $25 million

The paper proceeds in the following manner. In Section 2, we offer a brief review of the data sets used by earlier authors. In Section 3, we discuss our data and methodology. In

[2] For more on FX algorithmic trading, see King and Rime (2011). A recent press release by Greenwich Associates asserts that algo use rose 25 percent in 2018 compared to the earlier Greenwich survey year.

Section 4, we present and discuss the results. Finally, we provide a summary and conclusions in Section 5.

2. Measuring transaction costs in currency markets

While we believe that we are the first to explicitly study the effect of trade size on returns to currency investing, many studies have incorporated trading costs when computing returns to currency investing. These studies typically use quoted bid-ask spreads from various sources that represent the best bid and ask at the top of the order book. Our brief literature review provides an illustration of popular approaches.

2.1. WM/R spreads

A data source that is frequently used for academic studies is the Datastream 4 pm London prices. The popularity of this source is consistent with the use of the London 4 pm prices as the standard benchmark for the daily mark-to-market of global investment portfolios. The underlying data are the World Market Reuters (WM/R) WM/R fixing prices, which are averaged over a short window of time around 4 pm using trades or quotes on the major electronic crossing networks (ECNs) for currencies. The exchange rates in the fixing window are median bids and asks reflecting trades at or near the top of the order book. Below we examine the WM/R spreads relative to our order book data to assess their use as a transaction cost proxy, but for now we review papers that have used this source. Menkhoff et al. (2012) conduct an exhaustive study of currency momentum trading. They sample bid-ask spreads from the usual Datastream source used for academic studies as the quoted spreads on the last day of each month. They conclude that accounting for bid-ask spreads lowers the profitability of momentum strategies significantly, since momentum portfolios skew towards minor currencies with high transaction costs, which account for roughly 50% of momentum returns. Cenedese et al. (2014) condition a carry trade strategy on a measure of volatility and use the WM/R spread data to measure transaction costs. They find that their volatility-conditioned trading strategy yields positive post-cost returns. Lustig et al. (2011) identify a common factor that accounts for most of the cross-sectional variation in excess returns between high and low interest rate currencies. They use Datastream WM/R data, stating that the data provides conservative estimates of transaction costs, as actual costs should be smaller [they cite Lyons (2001) for this claim]. In Subsection 4.2.1, we examine whether WM/R spreads overestimate the cost of trading, and we show that this is generally true for small trades, like $1 million, but not for larger trades. Many other studies have used WM/R spreads to proxy for transaction costs (e.g., Darvas, 2009; Burnside et al., 2011; Banti et al., 2012; Della Corte et al., 2016; Maurer et al., 2019).

2.2. Electronic brokerage spreads

Apart from Datastream daily data, some authors have used quotes or trades directly from electronic brokerages like EBS or Reuters. For example, Mancini et al. (2013) provide a detailed examination of intraday pricing and liquidity in FX, along with returns to a carry trade strategy. They use the best bid and ask on the EBS electronic brokerage system for FX, which is the most popular interbank trading platform for the euro and yen (against the U.S. dollar). Their measure of the top of the order book is fine for small trades, but larger trades will reach further down the order queue to execute the full amount. Of course, the more liquidity one takes at a point in time, the wider the spreads will be. Many other studies use EBS spreads, including Breedon and Ranaldo (2013), Chaboud et al. (2014), Yamada and Ito (2017), and Ito and Yamada (2018). Breedon and Vitale (2010) use spreads from both EBS and Reuters electronic brokerages and find that the correlation between order flow and exchange rates is largely due to portfolio balance effects. Akram et al. (2008) and Rime et al. (2010) also use Reuters data.

2.3. Other transaction cost approaches

In addition to the data sources listed above, some authors have used alternative sources or methods. For instance, Ranaldo and Somogyi (2018) use spreads from the OANDA retail platform. Gargano et al. (2019) use spreads from OANDA and Dukascopy, another retail platform. Gilmore and Hayashi (2011) use proprietary spread data recorded by AIG. Della Corte et al. (2008) do not employ any proxies for transaction costs, but instead calculate the break-even cost for different strategies. They estimate that break-even costs fall largely in the range of about 100 bps to 500 bps, a high enough range so that their strategies would be profitable even under high transaction costs. Finally, Dahlquist and Hasseltoft (2019) study momentum in economic output and inflation variables as a currency trading strategy and incorporate costs by assuming constant spreads. At a constant spread of 10 bps, their Sharpe ratio falls modestly for developed market currencies. At an assumed constant spread of 30 bps for EM currencies, the Sharpe ratio falls only slightly.

The studies cited above are representative of a literature that has used a variety of approaches to incorporating trading costs. None of them have any consideration of trade size and associated higher costs. In general, such analyses will be fine for trading strategies that execute small trades. However, for larger trades, such methodologies may seriously underestimate the cost of trading and, therefore, overestimate the post-cost returns. We demonstrate how much trade size can impact returns for a generic currency trading strategy, as well as the value of embedding costs in the trade signal construction.

3. Data and methods

This section begins with a description of the data set utilized and then turns to an examination of the costs of trading each currency. We then specify two different portfolio construction approaches that we use to illustrate the effects of transaction costs on realized returns to investors.

3.1 Currencies and transaction costs

We select 18 currencies: 9 developed market currencies (EUR, JPY, GBP, CAD, AUD, NZD, SEK, NOK, and CHF) and 9 emerging market currencies (CNH, SGD, MXN, CZK, TRY, PLN, HUF, ILS, and ZAR).[3] The selection was made based upon the availability of data on the major electronic brokerages. For this study, we use two kinds of data for these 18 currencies. First, we need a daily price reference and use the WM/R daily fixing spot rate of each currency at 4 pm London. Second, we require data on bid-ask spreads for

[3] ISO codes for the currencies we use are: euro, EUR; Japanese yen, JPY; British pound, GBP; Canadian dollar, CAD; Australian dollar, AUD; New Zealand dollar, NZD; Swedish kronor, SEK; Norwegian kroner, NOK; Swiss franc, CHF; Chinese renminbi, CNH; Singapore dollar, SGD; Mexican peso, MXN; Czech koruna, CZK; Turkish lira, TRY; Polish zloty, PLN; Hungarian forint, HUF; Israeli shekel, ILS; and South African rand, ZAR.

each currency to measure transaction costs. These data come from EBS or Reuters, whichever venue is more liquid for a particular currency, and are daily averages over liquid trading hours. Our data include tradable prices so that we can calculate "sweep to fill" costs for trades of different sizes in the order book. Normally, when trade size increases, the transaction costs would also increase as the order must reach further down the order book to fill the trade. For our analysis, we compute transaction costs for trades of $1 million, $10 million, and $25 million, created as the top of the order book for $1 million, and sweep-to-fill costs for $10 million and 25 million.[4] The sample period is January 2, 2014 to March 22, 2018. Our goal is not to provide a study for a particular time of day, but to illustrate the general impact of trade size and tcosts on a systematic strategy.

3.1.1. FX trade size

Since we calculate sweep-to-fill costs of $1, $10, and $25 million, it is useful to put that size in context. However, since the FX market is relatively opaque compared to other asset classes, information on size of trades is not easily found. With regard to actual trade size, Hasbrouck and Levich (2019) analyze CLS settlement order size for April 2010, 2013, and 2016 and report mean spot settlement size in 2013 across all currencies of $1.1 million with a max trade size of $3.0 billion. Of course, the more liquid currencies will have larger mean size than the less liquid currencies. To pursue this further, we analyzed CLS data available to us on daily trade volume and number of trades for investment funds. From these data, we can infer the average trade size by day and currency. The

[4] There were a few cases where EM currencies did not have $25 million of depth in the order book. In these cases, we extrapolate out to $25 million using the available prices. The following currencies costs were against EUR rather than USD: SEK, NOK, CZK, PLN, and HUF.

results differ across currency pairs. For instance, for EURUSD, the mean trade size is inferred to be EUR4,422,533 with a max daily average of EUR53,012,885. Again, these are daily averages, so to have a day when the average trade size was greater than EUR53 million, there were some very large trades. There are some days where extremely large flows go through for each currency, relative to normal liquidity. Average daily trade size for other currencies are as follows: GBPUSD, GBP2,049,605; USDJPY, USD3,455,456; USDCAD, USD4,783,720; USDCHF, USD2,430,335; USDDKK, USD2,477,472; USDHKD, USD1,116,649; USDILS, USD1,111,143; USDMXN, USD1,927,775; USDNOK, USD1,590,312; USDSEK, USD2,182,414; USDSGD, USD1,207,076; and USDZAR, USD1,647,871. We see that average daily trade size across currencies is in the range of around USD1-4 million. Again, this is the average across all days and will certainly reflect some much larger trades and many smaller trades. The maximum average daily trade size is indicative of how some days have very large trades occurring. The maximum daily average ranges from about USD16 million for USDDKK to USD88 million for USDCAD.

We chose trade sizes with an upper bound of $25 million, which is about the largest size we can use for calculating sweep to fill costs for all the currencies we study. This may not be a particularly large size for a given fund. However, it is the relative impact of trades of different sizes that we want to capture.

3.2 Momentum strategy

In order to illustrate the impact of different trade sizes on the realized post-cost returns, we employ a generic investment strategy popular in currency markets, momentum. Momentum, or trend, strategies are particularly well-suited to analyzing tcost impacts since they are relatively high turnover strategies compared to the much slower carry or value strategies. The latter strategies tend to trade much less often than a typical momentum strategy. Of course, there are many different momentum constructions, including several benchmark momentum approaches provided by banks and index providers and performance tends to be episodic.[5] The empirical studies on currency investing in recent years have frequently created portfolios by sorting currencies into baskets using some methodology, like momentum, and then equally weighting the currencies in each basket. This equal weighting without regard to tcosts can give equal weight to currencies with very different tcosts and deliver less than optimal returns as a result. After illustrating performance using this common portfolio construction methodology, we turn to optimization using a tcost penalty and show how results change between the two approaches.

First, we use the equal-weighting portfolio sort methodology. For instance, a baseline strategy could be to use the simple approach of calculating 12-month returns and then, ranking the currencies once a month, going long the three currencies with the highest returns over the last 12 months and short the bottom three currencies with the lowest

[5] A discussion of generic currency investment strategies is provided in Pojarliev and Levich (2012). An example of a benchmark generic approach is in Sarevelos et al. (2018). Many researchers have looked at currency momentum portfolios (e.g., Okunev and White, 2003; Pojarliev and Levich, 2010; Burnside et al., 2011; Melvin and Shand, 2011; Asness et al., 2013).

returns over the past 12 months. In order to demonstrate the sensitivity of results to higher-turnover strategies, we construct past returns starting with a 1-day window, incrementing by one day at a time, up to 259 days. So we have 259 different momentum constructions ranging from the very fast past day's return to the much slower past year's return. We then rebalance the portfolio daily for each of the 259 momentum constructions to demonstrate how performance varies over the sample period depending on which momentum measure is used. Momentum returns are calculated as:

$$\forall i \in \{1,...,259\}, Mom_{t,i} = \left(P_{t-1} / P_{t-1-i}\right) - 1, \tag{1}$$

where P is the exchange rate mid-price.

3.3 Optimizing subject to transaction costs

In the initial benchmark approach, a simple equal weighting is used for the top three and bottom three momentum currencies in the portfolio. Next, we use portfolio optimization methods, incorporating transaction costs in the utility function to find the utility maximizing optimal weights for the currencies. A standard utility function for active investors is (see Melvin et al., 2013):

$$U = h'\alpha - \lambda(h'Vh) - \theta TC(\Delta h)\,. \tag{2}$$

In the above utility function, h denotes optimal holdings of assets and α is the expected return of the currencies. We use historical returns as calculated in equation (1) for our estimate of expected return. V is the covariance matrix of currency returns. We use the historical covariance for V, where the sample window used to estimate covariance is the same as the momentum strategy's construction history length for α. We constrain volatility to be less than 10% per annum. TC denotes transaction costs, which are a function of trade size Δh, and calculated from the sweep to fill costs from the order book. λ is a coefficient of risk aversion. We set λ equal to 2.5, which is a moderate degree of risk aversion. θ is the coefficient on transaction costs and is sometimes called the transaction cost amortization factor; the higher the θ, the greater the cost intimidation of trade size. We specify our baseline model with a theta of 1. However, we also use different values of θ to see the sensitivity of our results to optimization settings. Particularly, setting θ to 0 means that we do not consider transaction costs when optimizing, so we can compare results between optimizations with and without transaction costs.

For both the simple benchmark approach and the optimization approach, each period we calculate and rank the historical returns of currencies. For the benchmark approach, we choose the top three for the long portfolio and the bottom three for the short portfolio. However, optimization will tend to hold only one currency in a long position and only one currency in a short position if we do not set any constraint on the portfolio weights of currencies. As a result, we set the maximum weight of each currency as 1/3. Under this portfolio weight constraint, we hold at least three currencies in long or short positions. A key objective is to find the optimal weight h with respect to different θ and different bid-ask spreads using different trade volumes.

4. Results and evaluation

4.1. Determinants of transaction costs

Different momentum strategy settings have different returns and transaction costs. The total transaction costs are determined by two factors: the number of trades during our investment horizon and the transaction costs of each trade. First, the rebalancing frequency and construction history length influence the number of trades. A short rebalancing frequency and short construction history length means changing the composition of the portfolio more often, leading to a larger number of trades. Second, for the transaction costs of each trade, different currencies have different costs, and costs will rise with trade size. Simply put, market factors that affect spreads across currencies and time are volume and volatility. Greater volatility and smaller volume are associated with wider spreads.

4.2. Spreads, volatility, and volume

Transaction costs are measured by half spreads in basis points for trades of $1 million, $10 million, and $25 million from tradable prices on the major FX electronic brokerages by "sweep to fill" aggregation down the order book.[6] It is well known that spreads are a function of volatility as the market maker faces greater risk of an adverse price movement in more volatile times. Spreads are also a function of volume traded. Currencies with

[6] The major limit order books in FX are EBS and Reuters Matching. As mentioned above, other researchers have used the top of the order book from these sources for tcost proxies. The FX market is fragmented with alternative venues for trading. As a result, there is no market-wide single source with which to measure depth. While the market share of EBS and Reuters have fallen slowly over time as competing venues are introduced, they are generally seen as key sources of information for price discovery.

higher volume are generally more liquid and have smaller spreads than less-traded currencies. We summarize the relationship between volatility, volume, and transaction costs of the currencies in our sample in Tables 1 and 2. They present the relationship between the transaction costs of trades of $1 million, $10 million, and $25 million, along with data on average daily trading volume and volatility. Transaction cost (bps) is calculated by TC=(1/2)*(Ask Price – Bid Price)/Mid Price. The data for spreads and volatility are averaged over all days in the sample. The data on volume are from the BIS Triennial Survey, and represent a one-month sample of global trading.[7]

Tables 1 and 2 show that over the sample period, some currencies with relatively low volatility against the USD have relatively large spreads due to the low volume traded. This is better seen in Figure 1, where we plot the bid-ask spread for a $25 million trade against the average daily volume traded for the emerging market currencies. Volatility is depicted by the size of the bubble for each currency. There is generally a downward sloping relationship between spread size and volume traded, so that the larger the volume traded, the smaller the spread. Some EM currencies have relatively low volatility but relatively large spreads due to the low volume traded.

The row labeled 25/1 in Tables 1 and 2 shows the ratio of the spread for a $25 million trade to a $1 million trade. Since computing the sweep-to-fill costs of larger trades requires data that many researchers do not have and is very computationally intensive, the ratios in the tables may serve as a guide to researchers who have top-of-order-book spreads, but want to consider how larger portfolios, as traded by many investors, would

[7] See the Bank for International Settlements (2016).

be affected by tcosts. Table 1 indicates that the ratio of tcosts for large to small trades ranges from a low of 2.7 for EUR to 4.4 for NOK. Table 2 shows ratios ranging from 3.5 to 4.6. In general, a researcher could use a factor of 3 for G10 currencies and 4 for EM currencies to adjust the top-of-book spreads to account for larger trade size.

4.2.1. WM/R spreads

As discussed in Section 1, many researchers and practitioners have used the WM/R fixing spreads as a proxy for currency transaction costs. The WMR row in Tables 1 and 2 gives sample average half-spreads for the WM/R daily fixing taken at 4 pm London time each day. The WM/R London 4 pm price is popular as a benchmark price for marking-to-market global equity indexes like those of MSCI and many other professionally managed portfolios. The daily fixing price is also popular as a daily benchmark price for currency investing. The results in Tables 1 and 2 allow us to assess whether the WM/R spreads are representative of actual transaction costs. Referring to early work by Lyons (2001), where WM/R spreads were asserted to be wider than actual trading costs, some authors have reduced the WM/R spreads by 50% to better reflect the actual costs faced by market participants.[8] The results in Tables 1 and 2 allow us to assess this further using our data on actual costs of trades of different sizes.

[8] Gargano et al. compare WM/R spreads to spreads on FX aggregator sites Olsen Financial Technologies and Dukascopy and find that WM/R spreads are significantly larger than spreads on the other sites. It is not clear what size trade is associated with the spreads from the two aggregator sites, but presumably it is the top of their order books. Menkhoff et al. (2012) decrease the WM/R spreads by 50%, which they believe will bring them more in line with their expectation of actual costs of trading.

Table 1 shows WM/R spreads range from 1.15 bps for EUR to 3.54 bps for NOK. WM/R spreads are compared to costs of trading $1 million from our electronic brokerage data in the row labeled WMR/1 . WM/R spreads are larger than the actual costs in all cases. The results are supportive of the conclusions drawn by Menkhoff et al. (2012) and Gargano et al. (2019) that WM/R spreads overstate the actual trading costs. However, the row labeled WMR/25 in Table 1 shows, that for a larger trade size like $25 million, WM/R spreads understate the trading costs. If one is modeling trading strategies for small trades, it is reasonable to reduce WM/R spreads. However, if one constructs a model for trades of reasonably larger size, the WM/R spreads understate the trading costs. The WMR/25 row of Table 1 shows that the largest WM/R spread underestimate occurs for the highest cost G10 currencies.

Table 2 displays similar information for EM currencies. While WM/R spreads tend to overstate the trade costs for small trades for some EM currencies, costs are understated for MXN, TRY, PLN, ZAR, and ILS. The data suggest that it may not be correct to apply a WM/R spread reduction factor across all EM currencies. The WMR/25 row of Table 2 shows that for larger trades, the WM/R spreads understate the actual cost of trading by a considerable amount. Considering that WM/R spreads are constructed from sampled trades on the major electronic brokerages (actual executed buys and sells to which a "normal" spread is added by WM/R), this is consistent with most trades being at or near the top of the order book, so that the WM/R fixing spreads will be an underestimate of trade costs for larger sizes.

The WM/R spread comparisons drawn from Tables 1 and 2 suggest that researchers using the WM/R spread as a trade cost proxy should exercise caution before applying a fixed reduction across all currencies. If a model is intended to capture more than top of the order book trades, then the WM/R spreads will understate the costs for all currencies studied and should be increased rather than reduced.

4.3 Speed of momentum strategy and returns

Next, we explore the effect of the speed of momentum strategies on the costs and returns for the simple benchmark strategy with equally-weighted currencies. For this analysis, we use a daily rebalance frequency for momentum constructions as described in Subsection 3.2, ranging from a lookback period of one day, increasing by one day at a time, up to one year of past return history. The backtest simulation starts at 1/3/2014 and ends at 3/22/2018. The portfolio is equal-weighted, so the weights of the three long and short currencies are the same. To capture the effect of trade size, we create portfolios of \$3 million, \$30 million, and \$75 million and constrain individual currency trade size to less than \$1 million, \$10 million, or \$25 million, respectively, for each size portfolio. Summary results are presented in Figure 2. To be clear, the exercise is to first construct a 1-day return history as the measure of momentum returns used in ranking the top- and bottom-3 currencies to include in the portfolio each day. Then that portfolio is backtested over the sample period to determine the net-of-trade-costs return. The first point plotted in Figure 2 represents the return to the 1-day momentum strategy over the sample period. Then the exercise is repeated by adding an additional day to the momentum construction history or lookback period. The second point plotted in Figure 2 represents the return to the 2-day momentum strategy. We iterate over the momentum portfolio simulation repeatedly, by adding one additional day to the lookback period and determining the net return. In Figure 2, we plot the returns for all such portfolios from 1 day to 259 days (about one year).

From Figure 2, we find that net returns to the equal-weight benchmark momentum strategy are generally negative for the sample period, regardless of trade size. This was not a good sample period for simple generic momentum. Figure 2 also shows that all portfolios tend to have increasing returns with longer horizon momentum constructions up to around 100 days, regardless of trade size and corresponding transaction costs. The fastest constructions, like 1 or 2 days, trade much too often and performance is poor. So slowing down the turnover by lengthening the momentum construction window enhances performance until around 100 days. Beyond this, performance is fairly flat until 128 days after which net return begins to fall. For longer histories, there is no systematic gain. Thus, using past performance, we select 128 days as the construction history length for the remaining analysis. One could just as easily pick a few days less or more and the results would not change significantly. To dive deeper into the net return effects, we also analyze the relationship between transaction costs and different construction histories. Similar to Figure 2 for return, Figure 3 shows the total transaction costs from 1/3/2014 to 3/22/2018 with respect to different construction histories.

Figure 3 illustrates the impact of construction history for all three levels of trade size. Larger trades incur higher costs. Since returns for daily rebalances will have slowly evolving change over longer horizons, the longer the construction history, the lower the transaction costs. While it is obvious that short construction periods will have more volatile measured momentum returns and lead to a larger number of trades, what is interesting in Figure 3 is how fast the total costs fall as model construction horizons increase from the initial one day out to around 20 days, after which there are smaller decreases in costs associated with further slowing.

4.4 Transaction costs and net return with portfolio optimization

In the previous subsection, we show that simple generic momentum with equal weighting on each currency was not a successful investment strategy during our sample period. Now we turn to optimizing portfolios, where momentum is still used for the return forecast, but we solve for the utility maximizing portfolio weights on individual currencies using the utility function presented in Subsection 3.3:

$$U = h'\alpha - \lambda(h'Vh) - \theta TC(\Delta h) \tag{3}$$

$$\alpha_{t,i} \rightarrow Mom_{t,i} = \left(P_{t-1} / P_{t-1-128}\right)_i \tag{3a}$$

$$V \rightarrow Cov_{t,i,j}(\alpha_{t,i}, \alpha_{t,j}) \tag{3b}$$

$$TC \rightarrow \left\{ \begin{array}{l} Spread_t^* = Spread_{t-1}, \quad t = 1 \\ Spread_t^* = \gamma Spread_{t-1} + (1-\gamma) Spread_{t-1}^*, \quad t > 1 \end{array} \right\} \tag{3c}$$

The forecast of returns for each currency i, $\alpha_{t,i}$, is created from the 128-day momentum construction identified in the last subsection. The covariance matrix V is also based on the rolling past 128 days of returns. To avoid “peaking ahead,” tcost forecasts denoted *Spread**, are generated from lagged data using an exponentially-weighted moving average with a half-life of 10 days (determining the value of γ). As discussed earlier, this utility function includes a theta term, which is the coefficient on transaction costs. A theta of zero, results in a zero cost utility function so that the utility maximizing portfolio holdings will be determined without trade cost consideration. Alternatively, the higher theta, the greater the penalty applied to costs in the optimization. We vary theta in increments of 0.1 to examine the sensitivity of results to alternative cost intimidation in the utility function. Figure 4 shows the total dollar costs relative to portfolio size over alternative thetas. Note that at the far left is the cost of trading the equal-weighted generic benchmark momentum portfolio, labeled B. An optimized portfolio without a transaction cost penalty (theta of 0) has lower costs compared to the benchmark as the optimization tends to pick currencies with lower volatility, which tend to have lower costs. Both the benchmark portfolio construction and the zero-cost portfolio optimization do not penalize for costs, but the net return always accounts for the realized costs from trading. As theta increases from zero to impose a cost penalty ex ante, costs initially drop sharply for theta of 0.1. Beyond 0.1, increasing theta results in small reductions in costs for smaller trade sizes, out to around a theta of 1, after which costs are almost constant as theta increases.

We explored the effect of alternative theta on net return from our optimized portfolio (where net return equals gross return minus transaction costs). We simulated portfolios for theta from 0 to 3 and find that results differ across portfolio sizes. The smallest portfolio of \$3 million reaches its highest annualized net return of 3.2% for a theta of 0.8. The \$30 million portfolio reaches its highest annualized net return of 2.6% at a theta of 1. The \$75 million portfolio achieves highest annualized net return of 2.7% for a theta of 0.7. We use three small portfolio sizes to demonstrate that a fund does not have to be large for managers to worry about tcosts. While costs in Figure 4 monotonically decrease with higher theta, net return is not similarly increasing with higher theta beyond a small positive theta that differs for different sized portfolios. This reflects the fact that, at times, the high cost currencies are also high return currencies, so that greater cost intimidation in the optimization will penalize high cost currencies and lower their position size, reducing net return.

Table 3 shows information ratios for the benchmark portfolio and for optimized portfolios using a theta of zero (no tcost penalty) and the return-maximizing theta values described in the prior paragraph for the three different portfolios with maximum trade sizes of \$1 million, \$10 million, and \$25 million. Alternatively, we could simply impose a theta of 1 for all and the results would be qualitatively the same as reported below. Note that a portfolio with individual currency trade sizes of \$1 million is profitable for all three strategies. Portfolios with larger trade sizes earn negative returns for the benchmark and zero-theta portfolios. This is a caution to using the top of the order book as a sort of generic measure of transaction costs. The top of the book is only useful for small trades, so that what seems to work for small trade sizes may result in negative outcomes for larger portfolios. Having realistic cost estimates for the relevant portfolio size is important if backtests are to be credible indicators of strategy performance. The optimized portfolio with return-maximizing thetas raises the information ratio substantially relative to the benchmark and zero-theta portfolios. This cuts losses for the larger portfolios and enhances the smallest portfolio performance to an information ratio of 0.39 from 0.14 for the benchmark and 0.10 for the zero-theta optimization. The results in Table 3 reinforce the importance of including transaction costs in the optimization.[9]

4.5 Impact of transaction costs on portfolio composition and trade frequency

Table 4 shows the volume of trades over the full sample period, trade costs at different maximum trade size and different thetas, and annual portfolio turnover for the benchmark portfolio and optimized portfolios. The benchmark and zero-cost optimization portfolio are quite similar. However, when we consider transaction costs in the optimization with a theta of 1, the volume of trades and total transaction costs decrease substantially, regardless of whether the maximum trade size is $1 million, $10 million, or $25 million. Total costs will be high with theta of zero due to the much greater number of trades. Perhaps a more meaningful look at trade costs, in addition to total costs of trading, is the average transaction cost per dollar traded (total tcosts/total trade volume). The results in Table 4 indicate that the average cost per dollar traded drops substantially if a cost penalty is included in the optimization. Finally, the annual portfolio turnover is about 47 times the portfolio size each year if we ignore the trading costs. With an optimized

[9] Novy-Marx and Velikov (2016) show how rules-based cost mitigation strategies, like trading only low-cost stocks, can boost performance for equity portfolios. Maurer et al. (2019) study currency markets and show that an optimization with a transaction cost penalty outperforms rules-based cost mitigation strategies for currencies. They use the WM/R spreads from Datastream as their tcost proxy.

portfolio that accounts for transaction costs, turnover is in the more reasonable range of 6 times per year for this momentum strategy.

In Table 5, we compare the percentage of time a currency appears in the portfolio when moving from a theta of zero to theta of one for a portfolio with a maximum trade size of $1 million. First currencies are ranked in order of average costs of trading $1 million. Then, the No column shows the percentage of times each currency appears in the portfolio over the simulation period studied above for the optimized momentum strategy with no tcost penalty. The Yes column shows the percentage of time each currency is in the portfolio when the tcost penalty is used in the optimization (with theta=1). The Change column gives the change in the percentage of time a currency appears in the portfolio when we move from a zero-cost portfolio to one with a cost penalty. For instance, with no tcost penalty, the ILS is in the portfolio 34% of the days. With a tcost penalty in the optimization, the ILS appears in the portfolio only 10% of the time, a fall of 24 percentage points. The figure shows that moving from a zero-cost optimization to a cost penalty tends to lower the frequency with which one holds the most expensive currencies and raises the portfolio holding frequency for the relatively low-cost currencies, like CNH, EUR, JPY, and others. Ignoring trading costs results in portfolio managers holding high-cost currencies more often.

4.6. Robustness analysis of alternative portfolios

4.6.1. Tcost effects with lower portfolio rebalance frequency

The results reported so far are for a daily portfolio rebalance of the momentum strategy. It is useful to examine the importance of accurate transaction costs if, instead of daily, an investor rebalances monthly. We use the same momentum strategy as above, with a 128-day look-back period for expected return and covariance. Now we rebalance the portfolio once a month. It is fairly common for researchers to rebalance monthly strategies at month-end. Since we have accurate data on daily transaction costs, we examine the average spreads for each currency at month end, a week prior to month-end (5 trading days prior), and mid-month (11 trading days before the end of the month) to determine if month-end trading would incur higher costs relative to other times in the month. In addition, to further explore seasonalities in the data, we examine Fridays and end-of-quarter. The results are reported in Table 6. Our primary interest is in comparing month-end with other times in the month. For some currencies, month-end trading incurs higher than average costs, but on average across all currencies, month-end trading costs are close to overall average costs. Trading earlier in the month generally lowers costs relative to month-end. For instance, EM trades with a maximum 1 million size see month-end trading with a value of 1.01, or 1% above average costs. Trading mid-month, 11 days prior to end-of-month, has a value of 0.97, or 97% of average costs. Our monthly strategy is rebalanced at mid-month to realize the cost savings indicated by the seasonalities in the spreads. We also find that Fridays tend to have higher costs than other days and quarter-end tends to have the highest costs of all.

Portfolio diagnostics from our monthly simulations are reported in Table 7. Interestingly, the optimized portfolios outperform the benchmark portfolios, even for the smallest trade size. Since the optimization will downweight more volatile currencies, turnover and tcosts fall considerably relative to the benchmark. The question of particular interest here is whether tcosts still matter if rebalancing only once a month. For the smallest portfolio, optimized results are the same with and without the tcost penalty. For the middle portfolio, imposing a tcost penalty enhances performance a little. However, the largest portfolio experiences a substantial gain in performance when tcosts are accounted for in portfolio construction.

Comparing the monthly rebalance with the earlier daily rebalance, we infer the following:

- The monthly rebalance has a zero tcost optimized portfolio substantially outperforming the benchmark strategy of ranking currencies and trading the top and bottom.
- There are small differences between the zero tcost optimization and the benchmark strategy when portfolios are rebalanced on a daily basis.
- With a daily rebalance, tcosts are important for even the smallest portfolio.
- A monthly rebalance shows that tcosts are not so important for the smallest portfolio, but very important for the largest portfolio.

A pretrade optimization that incorporates a tcost penalty lowers turnover and realized tcosts by a small amount for smaller portfolios, but as portfolio size rises, the tcost penalty reduces turnover and tcosts substantially relative to a portfolio constructed with no pretrade cost consideration. For instance, comparing the annual turnover of the optimizations with no cost penalty with those incorporating the cost penalty, turnover is the following: a) $1 million, 5.71 vs. 5.40; b) $10 million, 5.72 vs. 5.04; and c) $25 million, 7.17 vs. 4.62.

Intuitively, less frequent trading should lower trading costs, but even for portfolios where $25 million is the largest trade, there is a gain from portfolio construction that incorporates a cost penalty. For small portfolios where $1 million is the largest trade, pretrade consideration of tcosts is not so important if trading only once a month.

4.6.2. Tcost effects with slower value strategy

We have focused on a momentum currency strategy, as momentum is a relatively fast, high-turnover strategy where transaction costs should play an important role. What if instead, one pursues a slower, low-turnover strategy like value? To explore the tcost effects of a generic value strategy, we create portfolios where the expected return is the percentage deviation of the spot rate from the IMF Purchasing Power Parities (PPPs) as given in the IMF World Economic Outlook (WEO). The covariance matrix is based on a rolling 128-day window and the portfolio is rebalanced daily, as in the earlier momentum results summarized in Table 3. Table 8 provides the portfolio performance of our value strategy.

The obvious takeaway from Table 8 is that tcosts still matter for a value strategy. By incorporating the tcost penalty in the optimization with $\theta=1$, the return and information ratio of the strategy both increase relative to the no-tcost simulation. The turnover for the momentum portfolio with no tcost penalty was about 47 and with tcost in an optimization fell to about 5. The turnover for the value strategy with no tcost penalty is about 10 and with tcost in an optimization fell to about 1. Even for such a low-turnover strategy, pretrade consideration of tcosts enhances performance.

5. Conclusion

It is common for studies to use bid-ask spreads at the top of the order book as a proxy for the costs of trading currencies. However, such a measure is only representative of the trade costs for small trades of around $1 million. Larger trades will incur wider spreads as larger trades exhaust the limit orders at the top of the order book and are filled at increasingly higher cost. We have data on the entire order book for the major FX electronic brokerages. These data allow us to compute sweep-to-fill costs to determine an estimate of the costs of different trade sizes with which we demonstrate the effect of trade size on trade costs, portfolio returns, and composition of asset holdings to currency investors.

A simple generic momentum strategy with no consideration of costs generally yields negative returns over the sample period studied. By including a penalty for costs in a portfolio optimization, positive returns are realized for larger portfolios. We find that as we move from the zero cost portfolio to optimizations that include greater penalties for costs, there is a limited range over which performance improves. Beyond a certain point, further trade intimidation reduces portfolio returns. One lesson is that the zero cost portfolio trades too much. However, if one intimidates trading too much, higher cost currencies will tend to disappear from the portfolio and their return potential is lost. Transaction costs and portfolio returns are also sensitive to different strategy settings with regard to frequency and size of trade rebalances: the faster the trading strategy (the shorter the momentum construction lookback history), the higher the transaction costs. For the sample considered, a momentum model based on approximately the past 6 months of returns yielded the best results.

To explore whether transaction costs matter for lower-frequency portfolio rebalancing, we conduct simulations for the momentum strategy using a monthly trade frequency. For the smallest portfolio with the smallest trade size, consideration of transaction costs is not so important. However, for larger portfolios with larger trade sizes, tcost consideration is still important, even with a monthly rebalancing. Momentum tends to be a fast, high-turnover strategy, so we simulate a slower, lower turnover value strategy to analyze whether transaction costs matter in this case. We find that accurate tcost consideration matters for the value strategy, even though the turnover is much lower than for a momentum strategy.

While many believe that since the currency market is highly liquid, transaction costs are a minor concern and one can trade small or large amounts without significant consequences.

We demonstrate that trade size matters a lot. If one measures costs using the top of the order book, this will underestimate the costs of larger trades and can lead to an investor pursuing a strategy that yields negative returns. We construct portfolios with associated trade sizes of around $1 million, $10 million, and $25 million. Our momentum construction is profitable for a small portfolio. However, the same portfolio construction approach with larger portfolios generates negative returns. This evidence of "large trades, small alpha" underscores the importance of having accurate cost measurements for backtests so that portfolio managers can understand the capacity of their strategy in terms of portfolio size and associated trades.

References

Akram, Q., Rime, D., and Sarno, L. (2008). Arbitrage in the foreign exchange market: turning on the microscope. Journal of International Economics 76: 237-253.

Asness, C., Moskowitz, T., and Pederson, L. (2013). Value and momentum everywhere. The Journal of Finance 68: 929-985.

Bank for International Settlements. (2016). Triennial Central Bank Survey. BIS Basle.

Banti, C., Phylaktis, K. and Sarno, L. (2012). Global liquidity risk in the foreign exchange market. Journal of International Money and Finance 31: 267-291.

Breedon, F. and Ranaldo, A. (2013). Intraday patterns in FX returns and order flow. Journal of Money, Credit and Banking 45: 953-965.

Breedon, F. and Vitale, P. (2010). An empirical study of portfolio-balance and information effects of order flow on exchange rates. Journal of International Money and Finance 29: 504-524.

Burnside, C., Eichenbaum, M., and Rebelo, S. (2011). Carry trade and momentum in currency markets. Annual Review of Financial Economics, Annual Reviews 3: 511-535.

Burnside, C., Eichenbaum, M., Kleshchelski, I., and Rebelo, S. (2011). Do peso problems explain the returns to the carry trade? Review of Financial Studies 24: 853–891.

Cenedese, G., Sarno, L., and Tsiakas, I. (2014). Foreign exchange risk and the predictability of carry trade returns. Journal of Banking and Finance 42: 302-313.

Chaboud, A., Chiquoine, B., Hjalmarsson, E., and Vega, C. (2014). Rise of the machines: algorithmic trading in the foreign exchange market. Journal of Finance 69: 2045-2084.

Darvas, Z. (2009). Leveraged carry trade portfolios. Journal of Banking and Finance 33: 944-957.

Della Corte, P., Ramadorai, T., and Sarno, L. (2016). Volatility risk premia and exchange rate predictability. Journal of Financial Economics 120: 21-40.

Della Corte, P., Sarno, L. and Tsiakas, I. (2009). An economic evaluation of empirical exchange rate models. The Review of Financial Studies 22: 3491-3530.

Gargano, A., Riddiough, S., and Sarno, L. (2019). Foreign exchange volume. Working Paper, Cass Business School.

Gilmore, S. and F. Hayashi. (2011). Emerging market currency excess returns. American Economic Journal: Macroeconomics 3: 85-111.

Greenwich Associates (2017), Long-Term Investors Embrace FX Algos.

Hasbrouck, J. and Levich, R. (2019). FX liquidity and market metrics: new results using CLS bank settlement data. Working Paper, NYU.

Ito, T., and Yamada, M. (2018). Did the reform fix the London fix problem? Journal of International Money and Finance 80: 75-95.

King, M., and Rime, D. (2011). Algorithmic trading and FX market liquidity. CFA Magazine 22: 15-17.

Lustig, H., Roussanov, N., and Verdelhan, A. (2011). Common risk factors in currency markets. The Review of Financial Studies 24: 3731-3777.

Lyons, R. (2001). The microstructure approach to exchange rates. MIT Press.

Mancini, L., Ranaldo, A. and Wrampelmeyer, J. (2013). Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums. The Journal of Finance 68: 1805–1841.

Maurer, T., Pezzo, L., and Taylor, M. (2019). Importance of transaction costs for asset allocations in FX markets. Working Paper, Washington University.

Melvin, M., and Shand, D. (2011). Active currency investing and performance benchmarks. Journal of Portfolio Management 37: 46-59.

Melvin, M., Prins, J. and Shand, D. (2013), Forecasting Exchange Rates: An Investors Perspective. Handbook of Economic Forecasting (G. Elliott and A. Timmermann, eds.), Elsevier.

Menkhoff, L., Sarno, L., Schmeling, M. & Schrimpf, A. (2012). Currency momentum strategies. Journal of Financial Economics 106: 660-684.

____________. (2017). Currency value. Review of Financial Studies 30: 416-441.

Merrin, S. (2019). Let slip the traders in alpha. Financial Times FTfm, June 3, 2019: 9.

Novy-Marx, R., and Velikov, M. (2016). A taxonomy of anomalies and their trading costs. Review of Financial Studies 29: 104-147.

Okunev, J., and White, D. (2003). Do momentum-based strategies still work in foreign currency markets? Journal of Financial and Quantitative Analysis 38: 425-447.

Pojarliev, M., and Levich, R. (2010). Trades of the living dead: style differences, style persistence and performance of currency fund managers. Journal of International Money and Finance 29: 1752-1775.

Pojarliev, M., Levich, R. (2012). A New Look at Currency Investing. CFA Institute.

Pukthuanthong-Le, K., Levich, R., and Thomas III, L. (2007). Do foreign exchange markets still trend? Journal of Portfolio Management 34: 114-118.

Ranaldo, A., and Somogyi, F. (2018). Heterogeneous information content of global FX trading. Working Paper University of St. Gallen.

Rime, D., Sarno, L., and Sojli, E. (2010). Exchange rate forecasting, order flow and macroeconomic information. Journal of International Economics 80: 72-88.

Saravelos, G., Gopal, S., Grover, R., Natividade, C., Harvey, Ol, Anand, V., Winkler, R., and Kalani, G. (2018). A Guide to FX as an Asset Class. Deutsche Bank.

Serban, A. (2010). Combining mean reversion and momentum trading strategies in foreign exchange markets. Journal of Banking and Finance 34: 2720-2727.

Yamada, M. and Ito, T. (2017). The forex fixing reform and its impact on cost and risk of forex trading banks. Finance Research Letters 21: 157-162.

Table 1: Annualized volatility and transaction costs of g10 currencies

Transaction costs are measured by half spreads in basis points for trades of $1 million, $10 million, and $25 million from spot prices on the EBS or Reuters FX electronic brokerages by "sweep to fill" aggregation down the order book. Values for volatility and spreads are sample averages. Volume data are billions of USD from the BIS Triennial Survey, and include spot and forward dated transactions against the USD. Spreads are in basis points.

Currency	EUR	JPY	GBP	CAD	CHF	AUD	NZD	SEK	NOK
Volatility	8.60%	10.20%	9.40%	7.90%	10.50%	10.00%	11.10%	6.30%	8.40%
Volume	$1,172b	$901b	$470b	$218b	$180b	$262b	$78b	$66b	$48b
1mn	0.47	0.53	0.9	0.96	1.07	1.09	1.62	2.14	2.81
10mn	0.8	0.98	1.49	1.76	2.2	1.99	3.41	4.32	6
25mn	1.29	1.68	2.51	3.1	4.24	3.52	6.41	8.53	12.31
25/1	2.7	3.2	2.8	3.2	4.0	3.2	4.0	4.0	4.4
WMR	1.15	1.46	1.55	1.40	2.88	2.15	3.01	2.78	3.54
WMR/1	2.45	2.75	1.72	1.46	2.69	1.97	1.86	1.30	1.26
WMR/25	0.89	0.87	0.62	0.45	0.68	0.61	0.47	0.33	0.29

Table 2: Annualized volatility and transaction costs of emerging market currencies

Transaction costs are measured by half spreads in basis points for trades of $1, 10, and 25 million from spot prices on the EBS or Reuters FX electronic brokerages by "sweep to fill" aggregation down the order book. Values for volatility and spreads are sample averages. Volume data are billions of USD from the BIS Triennial Survey, and include spot and forward dated transactions against the USD. Volume data for HUF and CZK are against EUR. Volume data for CNH and CZK are from the respective central bank sources, as their volumes are not reported separately in the BIS Triennial Survey. Spreads are in basis points.

Currency	CNH	SGD	MXN	CZK	TRY	PLN	HUF	ZAR	ILS
Volatility	3.6%	5.4%	11.6%	3.6%	11.4%	5.8%	6.2%	15.7%	6.6%
Volume	$67b	$81b	$90b	$18b	$64b	$19b	$5b	$40b	$7b
1mn	0.43	0.99	2.30	2.50	3.14	3.80	4.29	4.92	6.02
10mn	0.95	2.22	5.20	5.19	6.90	7.32	7.90	10.38	12.85
25mn	1.94	4.58	10.53	10.26	14.22	13.96	14.80	21.55	26.66
25/1	4.5	4.6	4.6	4.1	4.5	3.7	3.5	4.4	4.4
WMR	1.17	2.39	1.17	6.48	1.81	3.57	5.88	3.99	2.52
WMR/1	2.71	2.42	0.51	2.59	0.58	0.94	1.37	0.81	0.42
WMR/25	0.60	0.52	0.11	0.63	0.13	0.26	0.40	0.19	0.09

Table 3: Information ratio using different theta and trade size (annualized results)

Strategy	Max Trade Size	Net Return	Volatility	Information Ratio
Benchmark	$1 million	1.3%	9.0%	0.14
	$10 million	-0.9%	9.0%	-0.10
	$25 million	-4.5%	9.0%	-0.50
Θ=0.00	$1 million	0.9%	8.4%	0.10
	$10 million	-0.4%	8.4%	-0.04
	$25 million	-2.7%	8.4%	-0.32
Θ=0.8	$1 million	3.2%	8.4%	0.39
Θ=1	$10 million	2.6%	8.2%	0.32
Θ=0.7	$25 million	2.7%	8.2%	0.32

Table 4: Trade costs and turnover for different theta and maximum trade size

	Total trade volume (million $)			Total transaction costs (million $)			Average transaction costs per dollar traded (bps)			Annual turnover		
Benchmark	1mn	10mn	25mn	1mn	10mn	25mn	1mn	10mn	25mn	1mn	10mn	25mn
	594	5709	13173	0.28	5.30	22.41	2.33	4.64	8.51	47.10	47.43	47.45
Theta=0	1mn	10mn	25mn	1mn	10mn	25mn	1mn	10mn	25mn	1mn	10mn	25mn
	586	5714	13562	0.27	4.19	16.67	2.29	3.67	6.15	46.75	46.86	46.84
Theta=1	1mn	10mn	25mn	1mn	10mn	25mn	1mn	10mn	25mn	1mn	10mn	25mn
	83	526	1910	0.03	0.22	1.11	1.54	2.05	2.91	6.56	4.10	6.07

Table 5: Percentage of time a currency appears in the portfolio with and without a tcost penalty

Currency	Spread (bps)	Tcost Penalty		Change
		No	Yes	
CNH	0.43	40%	49%	*9%pts*
EUR	0.47	15%	39%	*25%pts*
JPY	0.53	47%	56%	*9%pts*
GBP	0.90	60%	62%	*2%pts*
CAD	0.96	34%	40%	*6%pts*
SGD	0.99	14%	20%	*6%pts*
CHF	1.07	23%	30%	*7%pts*
AUD	1.09	27%	37%	*10%pts*
NZD	1.62	64%	64%	*0%pts*
SEK	2.14	38%	35%	*-3%pts*
MXN	2.30	48%	59%	*11%pts*
CZK	2.50	36%	31%	*-5%pts*
NOK	2.81	36%	42%	*6%pts*
TRY	3.14	64%	59%	*-6%pts*
PLN	3.80	25%	21%	*-4%pts*
HUF	4.29	19%	2%	*-17%pts*
ZAR	4.92	71%	57%	*-14%pts*
ILS	6.02	34%	10%	*-24%pts*

Table 6: Seasonality in transaction costs

The table reports ratios of average spreads for each seasonal feature relative to overall sample average spreads.

ECN Spread Seasonalities Relative to Sample Averages

$1 million max trade size

	EUR	JPY	GBP	CAD	CHF	AUD	NZD	SEK	NOK	Average
Friday	1.02	1.02	1.03	1.03	1.05	1.02	1.04	1.03	1.03	1.03
Month end	1.00	1.03	1.00	1.00	0.97	0.99	0.99	1.04	1.02	1.00
Quarter end	1.01	1.02	1.05	1.04	0.93	0.98	0.99	1.29	1.14	1.05
5 trading days prior to month end	0.99	0.99	0.99	0.98	0.97	0.98	0.98	0.95	0.98	0.98
11 trading days prior to month end	1.00	1.00	1.00	0.98	1.01	0.99	1.00	0.95	0.98	0.99

	CNH	SGD	MXN	CZK	TRY	PLN	HUF	ZAR	ILS	Average
Friday	1.04	1.03	1.05	0.99	1.01	1.01	0.99	1.04	1.05	1.02
Month end	0.98	1.00	0.99	1.03	1.10	0.99	0.93	1.05	1.01	1.01
Quarter end	1.04	1.01	1.01	1.08	1.09	0.91	1.03	1.14	1.14	1.05
5 trading days prior to month end	0.98	0.99	0.98	0.97	0.97	0.98	1.00	0.94	1.04	0.98
11 trading days prior to month end	0.99	1.00	1.00	0.94	0.98	0.99	0.96	0.94	0.94	0.97

$25 million max trade size

	EUR	JPY	GBP	CAD	CHF	AUD	NZD	SEK	NOK	Average
Friday	1.04	1.04	1.04	1.06	1.05	1.05	1.05	1.02	1.01	1.04
Month end	0.99	1.04	0.98	0.99	0.98	0.97	0.99	1.04	1.02	1.00
Quarter end	1.02	1.03	1.04	1.03	0.95	0.97	0.99	1.22	1.07	1.04
5 trading days prior to month end	0.98	0.97	0.98	0.96	0.96	0.95	0.96	0.95	1.01	0.97
11 trading days prior to month end	0.98	0.97	1.00	0.98	0.99	0.95	0.98	0.95	1.00	0.98

	CNH	SGD	MXN	CZK	TRY	PLN	HUF	ZAR	ILS	Average

Friday	1.03	1.05	1.03	1.03	1.02	1.01	1.02	1.03	1.07	1.03
Month end	0.97	1.01	1.00	1.03	1.02	1.00	0.96	1.02	0.98	1.00
Quarter end	1.00	1.07	1.00	1.04	1.06	1.03	0.99	1.07	1.03	1.03
5 trading days prior to month end	0.98	1.01	1.01	0.97	1.01	1.04	0.96	0.97	1.05	1.00
11 trading days prior to month end	0.99	1.01	1.00	1.01	1.02	1.00	0.98	0.96	1.02	1.00

Table 7: Information ratio using different theta and trade size: monthly rebalance (annualized results)

Strategy	Max Trade Size	Net Return	Volatility	Information Ratio
Benchmark	$1 million	1.3%	9.2%	0.14
	$10 million	1.0%	9.3%	0.10
	$25 million	0.2%	9.3%	0.02
Θ=0.00	$1 million	2.8%	8.3%	0.34
	$10 million	2.6%	8.3%	0.32
	$25 million	0.3%	8.3%	0.03
Θ=1.00	$1 million	2.8%	8.3%	0.34
	$10 million	3.1%	8.3%	0.37
	$25 million	3.4%	8.4%	0.41

Table 8: Information ratio for value strategy using different theta and trade size (annualized results)

Strategy	**Max Trade Size**	**Net Return**	**Volatility**	**Information Ratio**
Benchmark	$1 million	0.4%	8.9%	0.04
	$10 million	-0.3%	8.9%	-0.04
	$25 million	-1.7%	8.9%	-0.19
Θ=0.00	$1 million	-0.6%	8.6%	-0.07
	$10 million	-0.9%	8.6%	-0.10
	$25 million	-1.5%	8.6%	-0.17
Θ=1.00	$1 million	1.4%	8.4%	0.16
	$10 million	-0.3%	8.3%	-0.03
	$25 million	-0.4%	8.0%	-0.05

Figure 1: Spread, volume, and volatility for emerging market currencies
Bid-ask spreads are plotted against average daily volume. The size of the bubble represents volatility.

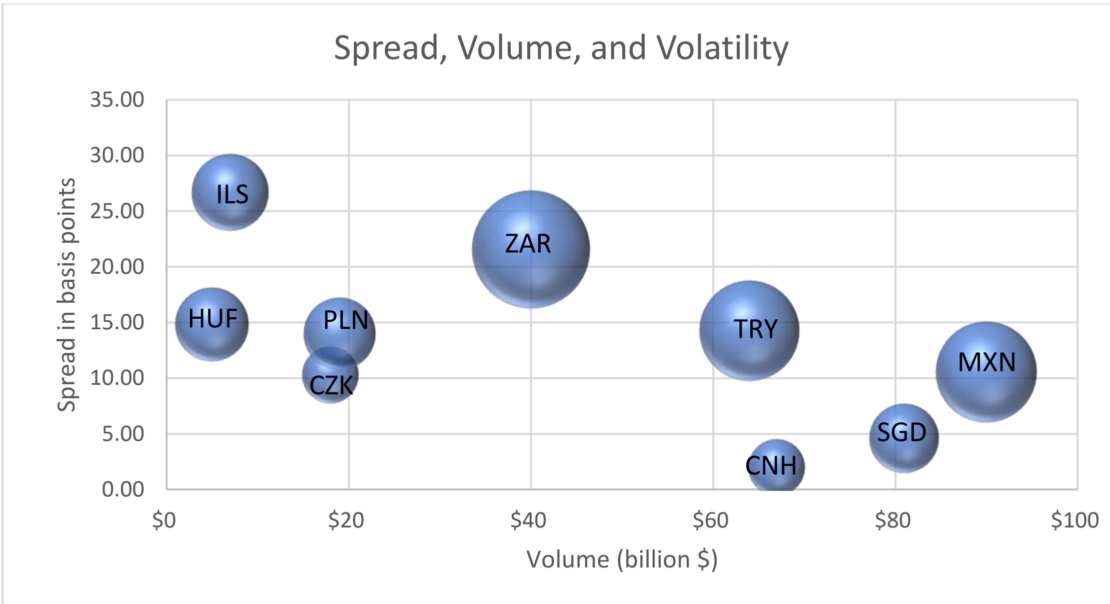

Figure 2: Total return with respect to different construction history length
Backtest portfolio simulations for momentum strategies ranging from historical returns of one day to one year are plotted for portfolios with trade sizes of $1 million, $10 million, and $25 million maximum size.

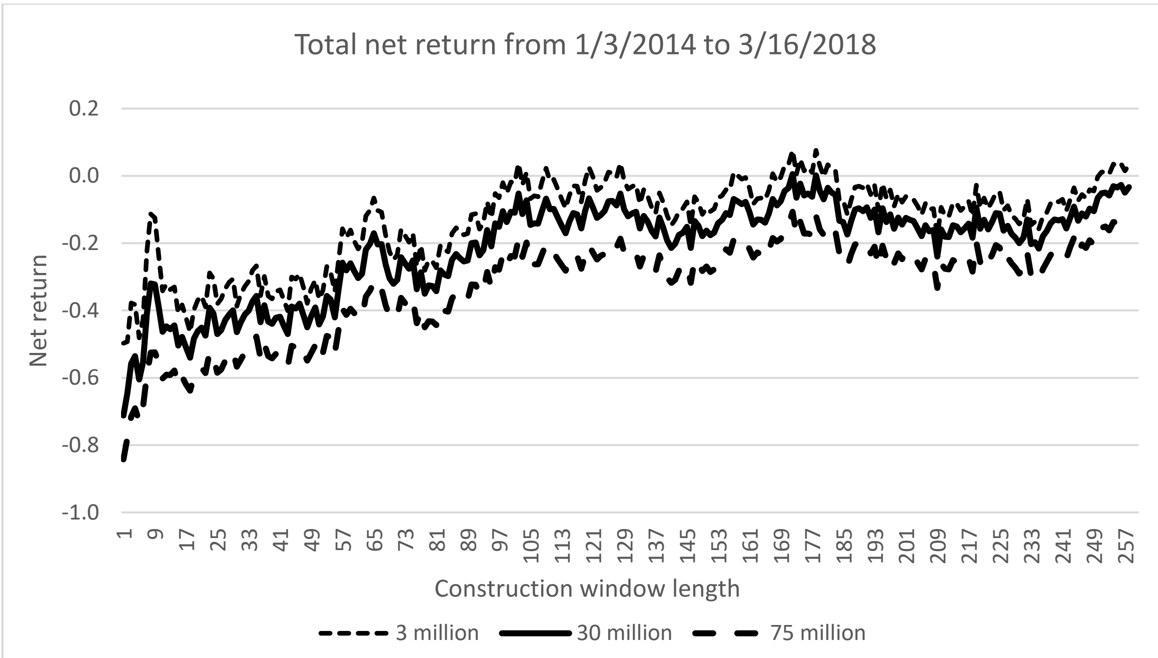

Figure 3: Total transaction costs with respect to different construction histories
Backtest portfolio simulations are constructed for momentum strategies ranging from historical returns of one day to one year with maximum trade sizes of $1 million, $10 million, and $25 million. Total costs, measured as (total costs/total portfolio size) in percent over the sample period, are plotted for each strategy.

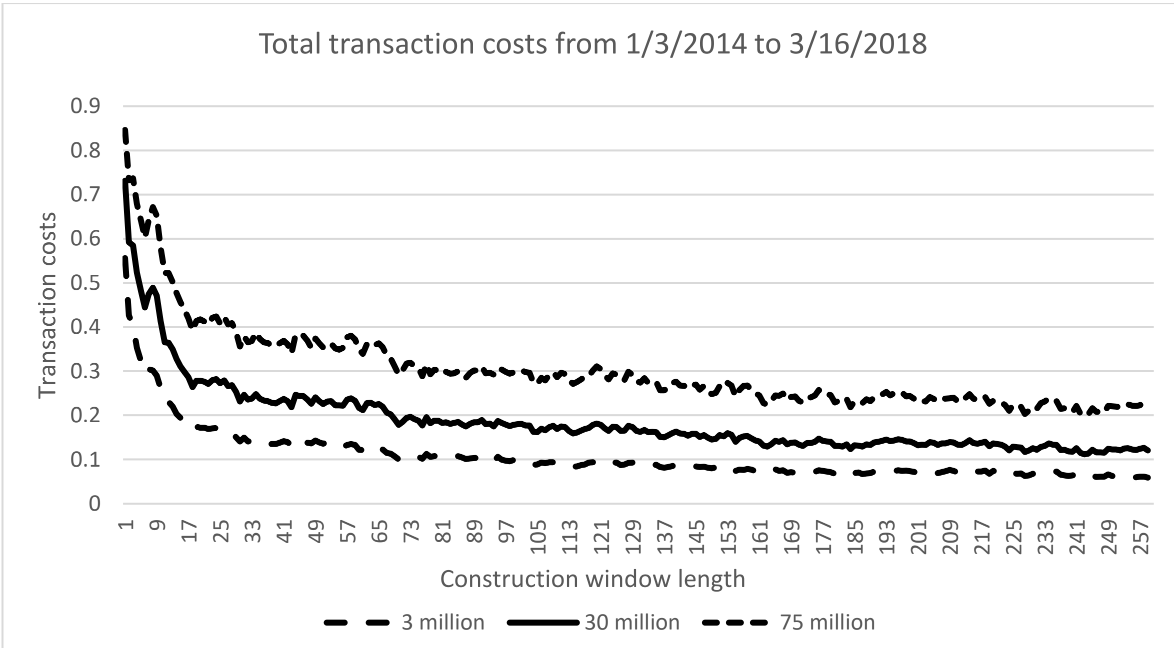

Figure 4: Transaction costs using different theta and portfolio size