The Value of Volume in Foreign Exchange*

Antonio Gargano† University of Melbourne

Steven J. Riddiough‡ University of Melbourne

Lucio Sarno§ Cass Business School and CEPR

August 16, 2017

We investigate the information content of foreign exchange (FX) volume using a rich dataset from the leading FX settlement platform. Consistent with theoretical predictions on the economic content of volume, we find that FX volume contains predictive information for both the time-series and cross-section of currency returns, which generates substantial economic value. A contrarian investment strategy that conditions on past daily volume generates an annualized return of over 19% and a Sharpe ratio of 1.82. We show the returns remain high after accounting for bid-ask spreads and are unrelated to other common currency strategies and risk factors.

Keywords: foreign exchange volume, currency returns, information asymmetry, transaction costs.
JEL Classification: F31, G12, G14, G15.

*We would like to thank Ric Colacito for helpful comments. Special thanks to Johan Bryssinck, Clayton Feick, Rob Franolic, Ray McTaggart, Chris Murray and Yue Yang for answering many data related questions. We thank the Accounting & Finance Association of Australia and New Zealand for generous financial support. All errors are our own.

†University of Melbourne, Faculty of Business and Economics. Email: antonio.gargano@unimelb.edu.au

‡University of Melbourne, Faculty of Business and Economics. Email: steven.riddiough@unimelb.edu.au.

§Cass Business School, City, University of London and Centre for Economic Policy Research (CEPR). Email: lucio.sarno@city.ac.uk.

1 Introduction

Foreign exchange (FX) market volume is enormous but little understood due to sparsity of data. In 2016, an average of $5 trillion was traded daily – over 10 times the size of daily global equity market volume. But trading volume in the FX market is opaque due to its decentralized multi-dealer structure, with no central exchange recording volume. Researchers are thus hampered in their efforts to understand the properties and information content of FX volume, which is lamentable given that volume has been shown to embed powerful predictive information about the time-series and cross-section of returns in other financial markets.[1]

In this paper, we make progress towards understanding FX volume and its link to currency returns by introducing a novel and comprehensive dataset to the literature. The data is from CLS Bank in New York, the largest FX settlement platform in the world, which has an unparalleled view of daily FX market transactions.[2] The data provides a rich source of high-frequency information on volume at the hourly level, across 31 currency pairs over a 5-year period from November 2011 to December 2016, and is split across three FX market instruments: spot, swaps and outright forwards.

We address two primary research questions. First, does FX volume contain information that is statistically and economically relevant for understanding future currency returns? Theory postulates a link between volume and returns because volume proxies for unobservable market conditions, such as the relative amount of liquidity and privately informed trading, that impact subsequent price changes. The FX market provides a unique testbed for this mechanism because the market inherently generates information asymmetries due to opaque flows of information. While in a centralized market – such as

[1] In the *cross section*, Datar, Naik, and Radcliffe (1998) find that stock returns are decreasing in turnover, while Chordia, Subrahmanyam, and Anshuman (2001) document a similar negative cross-sectional relationship using the variability of dollar trading volume. In contrast, Gervais, Kaniel, and Mingelgrin (2001) find a positive cross-sectional relationship between abnormal volume and future stock returns. In the *time series*, volume has been shown to contain information regarding continuation and reversal patterns caused by the strategic interaction between informed and liquidity traders (e.g., Admati and Pfleiderer, 1988). Cooper (1999) finds that stocks with decreasing volume experience larger reversals, while Conrad, Hameed, and Niden (1994) find this is true for high-transaction securities. Volume tends to increase with disagreement (e.g., Harris and Raviv, 1993), and therefore also predicts future return volatility (Gallant, Rossi, and Tauchen, 1992) and skewness (Chen, Hong, and Stein, 2001). Finally, Chordia and Swaminathan (2000) find that, consistent with models featuring sequential arrival of information (e.g., Jennings, Starks, and Fellingham, 1981), returns on high-volume portfolios lead the returns on low-volume portfolios, since low-volume stocks are less responsive to new information.

[2] Hasbrouck and Levich (2017) also make use of CLS data. Unlike this study, the authors obtain individual trades over a single month in order to compare CLS data with the Triennial Surveys of the Bank for International Settlements (BIS). Reassuringly, the authors find the CLS data provides a close match with the BIS Triennial Surveys, which we also confirm later in this paper.

the equity market – trading volume is easily observable and thus measurable, in an over-the-counter market – such as the FX market – volume and its information content are only observable to market dealers. Also, the FX market is populated by vastly different participants, ranging from hedge funds and currency asset managers that speculate on the movements in exchange rates to multinational corporations that trade to hedge FX exposure. This means that total trading volume is the aggregate of trades which reflect very different information sets, skills and objectives.

Our second question is whether the information content of volume varies across traded instruments. This question is particular to the FX market, where the composition of market participants differs substantially *across* instruments traded on the same underlying currency pairs. Hedge funds and some Principal Trading Firms (PTFs), for example, usually enter speculative trades either in the spot market due to the standardization and liquidity of the instrument, or the forward market to gain enhanced leverage, while multinational corporations and institutional investors are more likely to use FX swaps to hedge their currency risk exposure (e.g., Moore, Schrimpf, and Sushko, 2016).

The answers to these questions are relevant to a broad audience: to global investors seeking novel sources of returns and diversification; to academics searching for new insights into nominal exchange rate determination and the economics of volume; and to regulators and market designers wishing to better understand the motives for FX trading when designing optimal transparency regimes.[3]

In the theoretical literature, volume provides predictive information for future returns because it reflects information about the amount of "liquidity" and "informed" trading in the market. In a model of symmetric information in which only liquidity needs generate trading, Campbell, Grossman, and Wang (1993) show that return reversals are more likely after high-volume periods to compensate risk-averse dealers for providing liquidity.[4] In models with information asymmetry, however, private information also generates trading (Wang, 1994; Llorente et al., 2002). Trades based on private information are mimicked by uninformed investors, resulting in smaller expected return reversals (and possibly even return continuations) following high volume. Thus in markets characterized by high

[3] A central objective for regulators is to choose the optimal amount of information available to market participants. Making FX volume fully transparent is not *a priori* optimal if it reduces overall liquidity provision. A drop in liquidity could occur, for example, if revealing transactions exposes dealers' positions and disincentivizes a build-up of inventory (see, e.g., Lyons, 1996).

[4] In the model, liquidity demand is generated due to changes in risk aversion among a sub-population of investors that alters their optimal holding of risky assets.

levels of information asymmetry, return reversals should be less pronounced when volume is high.

The decentralized multi-dealer structure of the FX market inherently generates both liquidity and informed trading, while the disparate reasons for FX trading mean that the relationship between volume and currency returns is unclear and necessitates empirical exploration. Unlike other financial markets, a large fraction – approximately 40% according to the BIS (2016) – of all volume in the FX market is between dealers. Of the remaining volume generated by trades between dealers and customers, the breakdown of customers in the FX market is unusual relative to other markets. In addition to institutional investors and hedge funds, the market is driven by a large share of trading by central banks, governments, multinational corporations and retail investors.

Within the dealer-to-dealer segment of volume, part of the trading is driven by liquidity demands. This trading results from inventory management considerations, essentially the desire to minimize net FX exposure (see Lyons (1995)). Lyons (1997) refers to this as "hot potato" trading in which risk is passed around the interbank system, although recent evidence suggests the degree of "hot potato" trading is decreasing over time as liquidity becomes more concentrated in just a few large banks (Moore, Schrimpf, and Sushko (2016)).[5] In the customer segment of the market, individual trades will often be made without an explicit desire to profit from future exchange rate moves. Central banks, for example, trade FX when managing their foreign reserve account and when intervening in the FX market in an attempt to impact the evolution of exchange rates, while multinational corporations transact in FX as part of the day-to-day business of importing and exporting goods and services.

Informed trading can arise because, unlike centralized markets, customer-dealer FX trades are not publicly disclosed.[6] A dealer therefore has a partial but potentially informative view of FX trades (Lyons (2001); Evans and Lyons (2002)). These trades may be privately informative to dealers because they originate from customers with information about future macro fundamentals – such as central banks (Peiers (1997)) – or because in aggregation the orders help predict the direction of future macro fundamentals, such as multinational corporation orders signaling a shift in a country's trade

[5]The desire to finish the day "flat" is part of a risk management strategy that prevents dealers from maintaining risky positions overnight when trading continues in other time zones. Using the BIS Triennial Survey, Moore, Schrimpf, and Sushko (2016) report that across dealer banks, the warehousing of risk is common at the largest FX dealers. This pooling of liquidity across only a few locations reduces the need to quickly reduce positions. Furthermore, the authors find that liquidity provision is increasingly being provided by non-bank electronic market makers.

[6]On the Nasdaq stock market, for example, trades in U.S. equities are reported through Nasdaq's 'trade reporting facility' usually within 10 seconds of trade execution.

Exhibit A:
Market Participants by FX Instrument

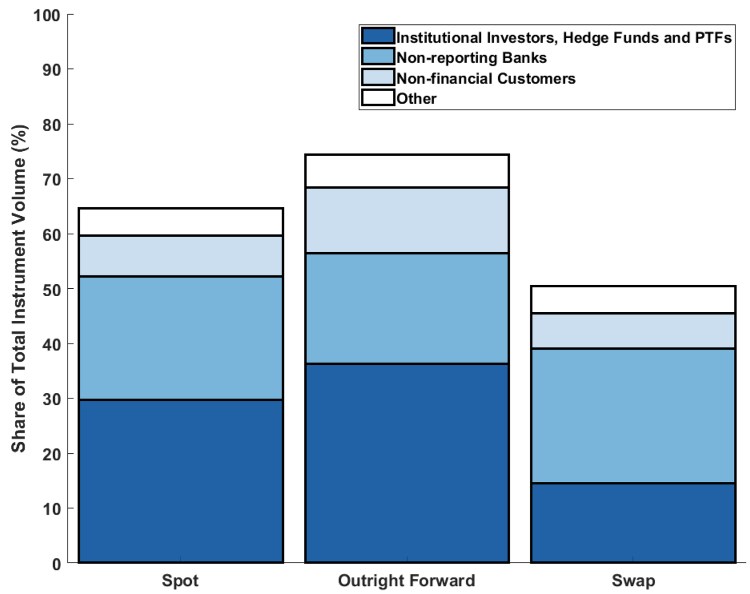

balance (Lyons (1997), Rime, Sarno, and Sojli (2010)). Indeed consistent with the view that banks gather information from customer orders, empirical evidence finds that bank trades embed predictive information about exchange rates (Osler and Vandrovych (2009)). Some investor groups may also have better information processing that generates more accurate forecasts of future fundamentals. Menkhoff et al. (2016), for example, show that the order flow of demand-side investors (such as hedge funds and "real money") contain predictive information about future FX rates.[7]

Since the level of informed trading in the FX market is driven by (i) the amount of customer trading (which dealers convert into private signals) and (ii) the amount of customer trading based on 'private signals' that results from better information processing, the level of informed trading is likely to vary across volume generated in different FX market instruments. We make this conjecture because across FX instruments, the level and composition of customer trading differs substantially. In Exhibit A, we present details about the relative share of customer trades in the spot, outright forward

[7] A number of studies have found proximate evidence of information asymmetries in the FX market. Ito, Lyons, and Melvin (1998) identify private information using changes in intraday FX market volatility before and after a lunchtime trading ban in Tokyo. Lyons (1995) and Payne (2003) find dealers adjust spreads in response to privately informed order flow, which can account for up to 60% of average bid-ask spreads. Survey evidence also indicates that dealers with access to more customer flow are considered to be better informed (Cheung and Wong, 2000).

and swap segments of the market. The highest level of all customer-based trading is in the spot and outright forward markets, indicating that more dealer originated private information is likely to accrue from this volume. Moreover, the incremental customer-orientated volume in those instruments is principally originated by hedge funds and PTFs, which is known to provide economically valuable information about future exchange rates (Menkhoff et al. (2016)).

In our empirical analysis, we begin by assessing the nature of the relationship between currency returns and aggregate FX volume (the sum of spot, swap and outright forward volume). Employing fixed-effects panel regressions, we find that FX volume is a key determinant of future currency returns once interacted with the current return. The interaction coefficient between returns and volume is positive and highly statistically significant, consistent with the prediction that FX volume contains high levels of informed trading (Wang (1994); Llorente et al. (2002)). Moreover, the first-order auto-correlation coefficient is negative and statistically significant, indicating a reversal effect when volume is low. We repeat the analysis separately on spot, outright forward and swap data, and find the effect is strongest for spot and weakest for swap FX volume. In fact, the interaction term for swap volume is nine times lower than for spot volume and not statistically different from zero.

We assess the economic significance of these results using a portfolio approach. Portfolios are constructed by performing a conditional double sort in which, each day, currency pairs are sorted by their daily return and then by their respective volume over the previous 24 hours. From the panel regression results, the returns on low volume currency pairs are expected to reverse, while the effect should be dampened for high volume returns. The resulting strategy, which we term "Reversal Low" (*RevL*), has a long (short) position in currency pairs with low (high) prior returns, from the perspective of the base currency, *and* low FX volume.

The *RevL* strategy generates impressive investment performance: its annualized return is over 19%, while the Sharpe ratio is 1.82. Consistent with the panel regressions, the performance of the strategy is strongest when conditioned on spot and outright forward volume. We assess whether these returns are robust to the inclusion of transaction costs. Because our proposed *RevL* strategy requires daily rebalancing, we pay more attention to the measurement of transaction costs than is typically done in the FX literature, which is largely dominated by monthly rebalancing strategies such as carry. Obtaining data from a variety of sources, including a large retail platform, the single-bank trading platform of a large global bank, and the filtered quotes from Olsen Financial Technologies, we show that bid-ask spreads typically used in the FX literature are overly conservative and likely to substantially under-estimate the returns available to investors. After incorporating these costs, we find the (*RevL*) strategy return remains sizable (13.1% per annum) and highly statistically significant. The returns to the strategy conditioned on swap volume, however, are no longer statistically different from zero.

In additional analysis, we find the strategy is profitable when rebalanced at most points during the trading day, except for the lowest liquidity hours when only the Australian FX market is open and spreads are widest. The economic information in volume is also shown to be short-lived – the strategy needs to be implemented within a few hours of the signal being realized to generate positive returns after incorporating transaction costs. Furthermore, the strategy is found to be uncorrelated with other strategies such as currency carry and momentum, and therefore offers investors a potential new source of diversification.

We run a battery of robustness tests on the strategy and find the returns are unaffected when: forming either eight or 12 (rather than nine) conditionally sorted portfolios; measuring volume using various detrending procedures; and standardizing volume over different horizons. Moreover, we find the returns are not driven by isolated currency pairs and remain statistically significant when p-values are calculated using a bootstrap procedure.

Related literature. We contribute to three strands of the literature. First, we build on recent studies casting new light on trading and liquidity in the FX market. Mancini, Ranaldo, and Wrampelmeyer (2013) and Hasbrouck and Levich (2017) use novel data to investigate the behavior and characteristics of FX market *liquidity*, while Menkhoff et al. (2016) use proprietary data on FX market *order flow* to show that customer orders contain predictive information about future currency returns. We differ from these papers by studying FX *volume*. FX volume is admittedly a close relative (conceptually) to both liquidity and order flow, yet the precise relationship is ambiguous. Volume in the FX market is likely to rise when liquidity is high but also rises during periods of low liquidity, as witnessed during the global financial crisis (Melvin and Taylor (2009)), while order flow is a record of 'signed' transactions and can therefore equal zero even when aggregate volume is high.[8]

The few prior studies analyzing FX volume address different research questions and use less comprehensive data. For example, Grammatikos and Saunders (1986) study the relationship between volume and volatility using daily FX futures volume for a smaller cross-section of five currency pairs, while Fisher and Ranaldo (2011) and Levich (2012) focus on the impact of Federal Open Market Committee announcements and changes in financial regulation on FX volume. In these studies, Levich (2012) employs data on interbank outright forwards from the periodic survey conducted by six "FX Committees", whereas Fisher and Ranaldo (2011) use an earlier and coarser version of the CLS data, aggregated across five currency pairs at daily frequency.

Second, we contribute to the growing literature studying the cross-sectional predictability of currency returns and the strategies that exploit this predictability, including: carry (Lustig, Roussanov, and Verdelhan, 2011; Menkhoff et al., 2012a), value (Asness, Moskowitz, and Pedersen, 2013; Menkhoff et al., 2017) and momentum (Menkhoff et al., 2012b; Asness, Moskowitz, and Pedersen, 2013).[9] The strategy we propose has three desirable characteristics: (i) the investment performance is very attractive relative to other common strategies during the sample period we examine; (ii) it has theoretical underpinnings and thus has economic justification for continuing to offer a source of desirable returns; and (iii) it is unrelated to other common strategies and thus offers a novel source of diversification. Furthermore, we contribute to this literature more generally by providing guidance on the use of bid-ask spreads. Researchers frequently use average quoted spreads reported by WM/Reuters, which are thought to be much larger than the effective spreads paid in markets (Gilmore and Hayashi (2011)). To compensate, authors have typically employed ad-hoc rules, such as assuming a 50% WM/Reuters spread (Goyal and Saretto (2009); Menkhoff et al. (2016)). We show that even this rule is overly punitive and that the scaling factor should be closer to 25% of the quoted WM/R spread.

[8] The relationship between liquidity and volume is known to be ambiguous from the equity literature. Chordia, Roll, and Subrahmanyam (2001) find a strong positive correlation between the 'dollar depth' measure of liquidity and volume, whereas other studies find trading activity increases in periods of high volatility (see e.g., Karpoff (1987)), which often correspond to periods of low market liquidity and increased trading spreads. Indeed, Chordia, Roll, and Subrahmanyam (2001) find a positive correlation between measures of effective spreads and volume. Since the seminal work of Lyons (1995), the FX literature has focused on understanding the information contained in order flow – the net number of buyer-initiated versus seller-initiated orders – and its role in determining price behavior (Evans, 2002; Payne, 2003).

[9] Carry, value and momentum are three strategies often used as the basis for exchange traded funds. Other recently proposed strategies include mixing carry and value signals (Jordá and Taylor, 2012), enhancing carry trade profits (Bekaert and Panayotov, 2016) and conditioning on signals from volatility risk premia extracted from currency options (Della Corte, Ramadorai, and Sarno, 2016).

Third, we contribute to the literature on volume in financial markets. Llorente et al. (2002) find time-series evidence that volume embeds information about the level of "speculative" trading, while Conrad, Hameed, and Niden (1994) find cross-sectional evidence in favor of the information asymmetry theory of Wang (1994). In contrast, Campbell, Grossman, and Wang (1993) and Cooper (1999) find time-series and cross-sectional evidence in support of the liquidity demand theory of volume. These studies focus on equity market spot volume. Our study explores a different market in which information asymmetries have a clearer motivation, and across different market instruments in which investor types are known to differ. We thus provide evidence on the properties of volume in the largest over-the-counter market and find support for the information asymmetry hypothesis, which is highlighted by investigating the cross-section of FX market instruments to show the effect is weakest when the traded instrument has less "speculative" trading activity.

The rest of the paper is organized as follows. Section 2 describes the data. Section 3 provides summary statistics. Section 4 presents the statistical evidence on the relationship between returns and volume. Section 5 explores the question of economic value in a portfolio setting. Section 6 investigates the source of *RevL* returns. Section 7 analyzes the diversification benefits from adding the strategy to other common currency strategies or within existing currency hedge funds. Section 8 provides robustness tests. Section 9 concludes. An Internet Appendix is available at the end of the document.

2 Data

2.1 Foreign Exchange Market Volume

We obtain data on FX volume from CLS Bank (CLS), which is available from Quandl – a financial, economic and alternative data provider. CLS is a financial institution that operates the world's largest FX cash settlement system. Following an FX transaction, Settlement Members of CLS submit the details of the order for authentication and matching by CLS. Volume is recorded once instructions are received from both counterparties to the trade. In the data, CLS records the transaction as occurring at the timestamp of the first instruction being received.

According to CLS, the vast majority of trade instructions are received within two minutes of trade execution. In recent work, Hasbrouck and Levich (2017) find that CLS receives the majority of its instructions less than ten seconds after trade execution. On the settlement date (at time t+2 for most spot transactions), CLS simultaneously settles both sides of the FX transaction, thus mitigating FX settlement risk.[10] There are currently 66 Settlement Members, mainly comprised of the world's largest banks, although settlement services are also extended to 'third party' clients of the Settlement Members, including other banks, funds, non-bank financial institutions and multinational corporations. In 2015, there were over 20,000 'third parties' using CLS's settlement services, providing CLS with a unique daily view of FX trading activity.

The dataset has several characteristics that make it uniquely suitable for investigating the information contained in FX volume and its statistical and economic value. First, volume is recorded at hourly intervals from 9pm Sunday to 9pm Friday (London time), which matches the full FX trading week from the open in Sydney on Monday morning to the close in New York on Friday. The hourly data enables us to aggregate volume across 24-hour periods to match with daily currency returns and assess the impact liquidity has on returns throughout the trading day. The relative high-frequency of the data is also important for assessing the *statistical* value of volume in a time series setting. The current gold standard for learning about aggregate FX volume is from the survey of central banks conducted by the BIS.[11] The BIS data is unsuitable for our purposes, however, because of its low frequency: the survey provides a snapshot of FX market volume, covering one month of trading (usually in April) at three year intervals. At the other extreme, researchers have used ultra-high-frequency data from inter-dealer trading platforms. This data is generally obtained for short samples, making it better suited to answering microstructure-related questions, e.g., pertaining to the information content of order flow.[12]

Second, for the full sample period at our disposal from 1 November 2011 to 31 December 2016,

[10]See Galati (2002) and Lindley (2008) for details of the CLS settlement process and the systemic impact on settlement risk.

[11]The survey has therefore been used to explore longer-term trends in the FX market; see e.g., Galati and Heath (2008), King and Schrimpf (2010), Rime and Schrimpf (2013), Moore, Schrimpf, and Sushko (2016).

[12]Lyons (1995), for example, observes direct interdealer transactions of a single USDDEM dealer over five trading days. For the same currency pair, and over a similar trading interval, Payne (2003) observes the activity of multiple dealers on the Reuters D2000-2 electronic brokerage system. Berger et al. (2008) use one minute data from the EBS brokerage platform to study EURUSD and USDJPY, while Evans (2002) uses four months of direct interdealer data from the Reuters D2000-1 platform. Data on high-frequency trading activity has also been used to test various microstructure models of dealer behavior, as, e.g., in the study by Bjønnes and Rime (2005). Broader studies have recently been conducted by Evans and Lyons (2013) and Menkhoff et al. (2016) using customer-dealer transactions at two of the largest global banks to assess the predictive information contained in order flow.

we employ data for 17 major currencies and 31 currency pairs.[13] This large cross-section is important for evaluating the *economic* value of volume by allowing us to construct more granular portfolio strategies, in which currency pairs are grouped based on past returns and volume, to assess the investment performance that can be generated using the predictive information in volume.

Third, the volume data is split across spot, outright forward and swap transactions, which form the three primary traded instruments. This cross-sectional information is important and unique to the FX market because it is known that the composition of market participants differs considerably across instruments, thus providing an opportunity to perform sharper tests of the economic theory linking volume with future returns.

We acknowledge two main limitations of the data. First, the sample is relatively short, containing around five years of data. Nonetheless, it is not uncommon to use shorter data samples when casting new light on the FX market (Mancini, Ranaldo, and Wrampelmeyer, 2013) or when studying the cross-sectional importance of volume (Conrad, Hameed, and Niden, 1994). Furthermore, the fact the data is higher-frequency provides greater statistical power when estimating time-series coefficients, and offers a large sample of portfolio returns for the daily rebalanced strategy. We also allay concerns regarding the representativeness of our results using a bootstrap procedure. The approach allows us to estimate a robust p-value on the *RevL* strategy's main performance measures (further details are provided in Section 8.3). In addition, we measure profitability using the out-of-sample performance measure of Ingersoll et al. (2007), which is less affected by outliers than the commonly reported Sharpe Ratio.

Second, while the CLS data provides the most comprehensive time series of FX volume, it does not contain complete market coverage. According to the 2016 BIS Triennial Survey, \$5 trillion dollars are traded daily across the FX market. In contrast, CLS settles around \$1.5 trillion dollars or 30% of total FX volume. The data does not offer complete market coverage for a variety of reasons: certain currencies with relatively high volume, such as the Chinese yuan and Russian ruble, are not settled

[13]The full dataset contains data for 18 major currencies and 33 currency pairs. The Hungarian forint (HUF) enters the dataset later, on 17 November 2015, and therefore we decide to maintain a balanced panel that excludes the USDHUF and EURHUF currency pairs. The remaining 31 currency pairs are: EURUSD, USDJPY, GBPUSD, AUDUSD, USDCAD, USDCHF, NZDUSD, USDDKK, USDNOK, USDSEK, USDKRW, USDSGD, USDHKD, USDILS, USDZAR, USDMXN, EURJPY, EURGBP, EURCAD, EURCHF, EURDKK, EURNOK, EURSEK, EURAUD, GBPCAD, GBPCHF, GBPAUD, GBPJPY, AUDJPY, AUDNZD, and CADJPY.

by CLS; the data does not cover settlement of FX options and non-deliverable forwards; and many smaller regional banks opt-out from becoming settlement members due to insufficient FX turnover.

It should be noted, however, that the CLS coverage is also understated when compared with the BIS survey. A large fraction of the volume reported to BIS is comparatively uninformative internal ("related party") trades that take place across desks within the same bank, while prime brokered "give-up" trades are double-counted in the BIS figures.[14] After accounting for these effects, total FX volume is nearer to \$3 trillion per day and thus CLS settles around 50% of all FX market volume. In the Internet Appendix we further mitigate concerns about the representativeness of the sample by providing evidence that an almost perfect relationship exists between the share of currency-pair volume in the BIS Triennial Surveys and the CLS data.[15]

2.2 Exchange Rate Returns and Currency Excess Returns

We supplement the FX volume data with daily WM/Reuters spot and one-week forward exchange rates obtained from Thompson Reuters, available from Datastream. The WM/Reuters FX rates are recorded at 4pm in London each trading day. The choice of FX data is important to ensure that both the FX volume and exchange rate return are measured over precisely the same 24-hour period. We calculate the exchange rate return as the log difference in the exchange rate over a trading day

$$\Delta s_{t+1} = s_{t+1} - s_t, \tag{1}$$

where lowercase letters refer to logs. Because we use both U.S. dollar and cross-rates in our analysis, when we refer to 'high' and 'low' returns we do so from the perspective of the base currency. We also calculate excess returns that take into account the interest rate differential between the two currencies. The FX excess return is given by

$$rx_{t+1} = \Delta s_{t+1} + (i_t^* - i_t), \tag{2}$$

[14] In the 2016 BIS report, of the \$5.07 trn in daily turnover, \$0.94 trn was generated via "related party trades", while \$0.89 trn of volume was "prime brokered".

[15] Around 20% of the trades settled by CLS are customer initiated trades (as opposed to inter-dealer trades), which is slightly less than the amount of customer-dealer trading recorded in the BIS Triennial Surveys and is likely caused by prime brokered trades being indistinguishable from other inter-dealer trades when reported to CLS.

where i_t and i_t^* are the overnight interest rates in the quote and base currencies. In our core analysis, we scale the excess returns by the spot exchange return of the quote currency against the U.S. dollar, in order to take the perspective of an American investor (we consider alternative investors' perspectives in the Internet Appendix). In practice we do not observe overnight money market rates for all 17 currencies in our sample. We therefore extract information about interest rate differentials from forward rates using the covered interest rate parity condition that $s_t - f_{k,t} \approx i_t^* - i_t$, where $f_{k,t}$ is the k-period forward rate observed at time t.[16]

3 Summary Statistics

In this section, we present summary statistics for our data on FX volume. In Table 1, we report summary statistics disaggregated across individual currency pairs. The first four columns report the sample mean, median, standard deviation and first order autocorrelation for aggregate volume. Unsurprisingly, the currency pairs with the most volume in our dataset are EURUSD, USDJPY and GBPUSD, for which CLS settles on average \$500bn, \$244bn, and \$186bn, respectively each day. The ranking of currency pairs by volume is largely in line with prior expectations from the BIS Triennial Surveys.[17] We find the distribution of daily volume is largely symmetric, as evidenced by the median and mean values being similar for each currency pair. The results also reveal a positive relationship between the level and variability of trading volume, which suggests that a normalization of volume is required to make meaningful comparisons across currency pairs in our subsequent analysis. Turning to the fourth column, we observe that most of the series display a mild first order autocorrelation; however, there is non-trivial cross-sectional variation ranging from EURJPY, which displays comparatively persistent volume (0.73 autocorrelation coefficient) to EURCAD that displays relatively fast

[16]We use one-week forwards for our main analysis but find qualitatively identical results using one-month forwards. Moreover, due to large deviations from CIP since the global financial crisis (e.g., Baba and Packer (2009), Du, Tepper, and Verdelhan (2017), Rime, Schrimpf, and Syrstad (2017)) we also use the available euro-currency deposit rates for 13 currencies to estimate interest rate differentials. We find the results remain virtually unchanged. All results are available upon request.

[17]In the Internet Appendix we confirm the reliability of the CLS data by comparing the average daily volume recorded by CLS with the 2013 and 2016 results from the BIS Triennial Central Bank Survey. The BIS conducts its market wide assessment of FX market volume in April every three years. It therefore provides the most comprehensive snapshot of trading activity at periodic intervals. We compare our data using the equivalent average trading volumes in April of 2013 and 2016 (the two survey windows covered within our dataset). The results show a clear agreement between the two datasets and an almost perfect cross-sectional correlation in the share of market volume across currency pairs.

mean reversion (0.22 autocorrelation coefficient).

The next three columns report pairwise correlations across spot, forward and swap volume.[18] In most cases, spot volume is correlated more with forward volume than swap volume – providing an early indication that swap market activity may embed different information.

The last three columns report correlations between dollar volume and the total number of trades – the two measures of trading activity in our dataset. In the spot market, the two measures are usually highly correlated (the average correlation is 84%), which perhaps reflects the standardization of trading size in FX spot (a 'standard' trade is for $10mn for major currencies). A noticeably lower average correlation is observed across the forward (45%) and swap (59%) markets. This finding is perhaps indicative of the large and variable FX transactions that multinational corporates make when hedging FX exposure in forward and swap markets. More generally, it reflects the bespoke nature of forwards and swaps, in which banks accommodate different trade amounts in the over-the-counter market.[19]

In the top panel of Figure 1, we report the average aggregate volume for each hour of the trading day (based on the time in London). In the early hours of the day, when only the Asian markets are open, volume levels are comparatively low. FX volume picks up noticeably when European and London markets open at 6am and 7am, respectively. After a fall in trading around lunchtime in Europe, volume rises again when New York traders enter the market around 1pm in London. The lowest volume is recorded between 10pm and 11pm when only the Australian market is open.

In the remainder of Figure 1, we report the same figure for spot, outright forward and swap volume. Consistent with the empirical evidence in the equity market (see Gerety and Mulherin, 1992; Jain and Joh, 1988) we find that volume concentrates in the early and later parts of the trading day in London and New York.[20] These are known to be higher volatility periods in which more economic

[18]We compute correlations across the one-day (log) growth in volume for each instrument.

[19]The majority of FX swap transactions are for next-day settlement ("tom-next") and therefore used to fund overnight interest rate positions (see, e.g., King, Osler, and Rime (2012)). In the remainder of the paper we focus on dollar volume as our measure of FX market activity. We make this decision for two reasons. First, theories on the economic content of volume focus specifically on dollar turnover, not the total number of trades. Second, traders frequently apply order splitting strategies to achieve better execution, making the number of trades a poor reflection of overall volume.

[20]However, this pattern is not consistent across all currency pairs. For example, trading activity in the Asia-Pacific pairs (e.g. AUDUSD, NZDUSD, AUDJPY, USDHKD) show three peaks that include the opening hours in Sydney and Tokyo. Most of the trading activity in European currency-pairs is concentrated between 7am and 4pm in London, while trading activity in USDMXN is highest during New York trading hours. These figures are available upon request.

news are publicly released (see Foster and Viswanathan, 1993; Berry and Howe, 1994). Interestingly, these patterns are not uniform across instruments. Activity in forward trading peaks at 4pm when markets close in London, while trading activity in FX swaps peaks around 8am when the London market opens. Admati and Pfleiderer (1988) propose that intraday patterns are due to the interaction between informed and uninformed investors. These patterns thus provide additional tentative evidence that the economic information in FX volume might differ across FX instruments.

4 The Predictive Power of Volume: Statistical Significance

In this section, we explore the statistical relationship between returns and volume. To do so, we run a series of panel regressions to assess if volume, once interacted with the current excess return, contains predictive information about future currency excess returns. We also run bilateral time-series regressions to explore if the panel regression results are consistent across our universe of 31 currency pairs.

4.1 Panel Regressions

We first test the relationship between volume and returns by estimating fixed-effects panel regressions in which currency returns are regressed on lagged currency returns and the interaction between lagged returns and volume. When working with volume data we want to ensure stationarity and to measure volume in relation to the market's capacity for absorbing volume in a particular currency pair. We therefore follow the literature (see e.g., Campbell, Grossman, and Wang, 1993; Llorente et al., 2002) and define our measure of aggregate volume for currency pair i as the (log) deviation from its recent trend, defined as

$$v_{i,t} = log(volume_t) - log\left(\frac{\sum_{s=1}^{N} volume_{t-s}}{N}\right) \tag{3}$$

where N is the number of days over which we estimate the moving average in daily volume.[21] The

[21] In our core analysis we use a one-month (21 day) window. In further analysis we provide evidence that the results are qualitatively unaffected when the window is expanded to 63, 126 and 252 days.

baseline model is thus

$$rx_{i,t+1} = \alpha_i + \tau_t + \beta_1 rx_{i,t} + \beta_2\Big(rx_{i,t} * v_{i,t}\Big) + \beta_3 v_{i,t} + \boldsymbol{\gamma}'\boldsymbol{x}_{i,t} + \boldsymbol{\delta}'\Big(\boldsymbol{x}_{i,t} \odot rx_{i,t}\Big) + \epsilon_{i,t+1}, \tag{4}$$

where $rx_{i,t}$ is the log currency excess return for currency pair i at time t, $\boldsymbol{x}_{i,t}$ is a vector of controls relative to pair i, α_i and τ_t denote currency-pair and time fixed effects, $\odot$ is the element-wise product operator used to generate interaction terms with controls. We control for daily FX market volatility and liquidity. Volatility is measured by fitting a GARCH(1,1) to each excess return series, while liquidity is estimated using the daily bid-ask spread for each currency pair.

Predictions vary as to the sign and significance of the coefficient β_2. In a model with informed trading, Llorente et al. (2002), show that β_2 is positive and increasing in the level of information asymmetry (informed trading) in the market. The intuition is that high volume indicates the presence of informed trades, which are mimicked by less informed investors in the following period – leading to a return continuation effect following high volume days. In contrast, Campbell, Grossman, and Wang (1993) propose a model with only liquidity-driven trading. Trading in the model is driven by exogenous shifts in risk aversion in a subset of the population. If risk aversion increases for one group, the marginal investor becomes more risk averse and only willing to hold a security at a lower price. Other investors accommodate this shift in risk aversion by requiring a higher rate of return. It follows that high volume days lead to strong return reversals, and thus β_2 is predicted to be negative. In both models there is no direct role for volume to influence future returns – volume only serves to provide information about the state of the world beyond market prices.[22]

We report coefficient estimates in Table 2. The first column presents results for the specification that includes only the lagged values of returns and volume. As predicted by theory, volume alone is not informative about future currency excess returns. Moreover, unlike the equity market, we find no evidence of return autocorrelation in the FX market. In the second specification we test if volume becomes informative once interacted with the current excess return. Supporting the theory of

[22] The first-order autocorrelation term (β_1) has a less pronounced role in the models. When volume is zero, Llorente et al. (2002) find returns should reverse although they acknowledge that other models would also generate the same prediction. Wang (1994), for example, shows that returns are likely to reverse in the presence of significant 'noise trading'. In Campbell, Grossman, and Wang (1993), the autocorrelation term could be either positive or negative depending on the magnitude of various parameters in the model.

Llorente et al. (2002), we find that β_1 is negative and statistically significant at the 10% level, while the coefficient on the interaction term β_2 is positive (0.20) and highly statistically significant. The interaction term is also economically significant. A one standard deviation shift in volume below its mean results in a reversal effect that is around three times stronger. The result indicates the presence of speculative FX trading activity leading to predictable currency excess returns. The inclusion of controls (column 3) has no qualitative effect on the interaction coefficient but results in a stronger overall marginal effect from past returns when volume is zero. We find the results are robust when limiting the analysis to just U.S. dollar- (column 4) or euro-base pairs (column 5).

The last three columns of Table 2 present results for the full baseline specification for volume across the three FX instruments. In each case, the coefficient on the lagged return is negative and statistically significant, indicative of privately informed 'speculative trading' playing a role in each instrument. Nonetheless, the coefficient estimates on the interaction term display a clear declining pattern across instruments: 0.19 for spot volume (significant at the 1 percent level); 0.07 (significant at the the 5 percent level) for outright forward volume; down to 0.02 for swap volume, which is not statistically significant and suggests the role of information asymmetries is less dominant in the FX swap market.[23]

Summing up, the results in this section suggest that the interaction between returns and volume predicts future currency excess returns in the time series. This feature is consistent across currency pairs and is stronger in the spot market than in the forward market, while it is not statistically significance in the swap market. We turn to the economic significance of this result in the next section.

[23] In additional analysis, we split the sample in half and find the interaction term remains positive and highly statistically significant in both the first and second parts of the sample. We also explore if the results are consistent across all 31 currency pairs by estimating individual bilateral models, in which currency excess returns are regressed on their lagged excess returns and the interaction between lagged excess returns and volume. In Figure A.1 of the Internet Appendix, we report ordinary least squares coefficient estimates for the interaction term across currency pairs. The results are strikingly consistent. The interaction term is positive for 26 of the 31 currency pairs when conditioning on total volume. When conditioning on spot, outright forward and swap volume, we find that 28, 26 and 20 interaction terms are positive.

5 The Predictive Power of Volume: Economic Significance

The previously documented statistical relationship between FX volume and currency returns does not necessarily imply that FX volume provides economic value. To quantify economic value, we employ a portfolio approach that allows us to assess the gains arising from a strategy that exploits the statistical relationship. The portfolio approach has been used recently in various currency market studies exploring novel currency trading signals (see e.g., Della Corte, Ramadorai, and Sarno, 2016; Menkhoff et al., 2016).

Time-series to cross-section. We learn from the time-series analysis that the estimated coefficient on the interaction term (β_2), between past volume and excess returns, is positive. This finding implies that if excess returns in the past period were positive, then excess returns in the next period are also expected to be positive *conditional* on our measure of volume being positive (i.e., we expect a continuation). In contrast, if our volume measure is negative, we expect a negative return (i.e., a reversal effect). Our measure of volume has a mean value close to zero (-0.04), while its standard deviation is around 0.4. The full marginal effect of past returns on future returns, however, goes beyond the interaction term and is equal to $\beta_1 + \beta_2 v_t$, and hence the effect when volume is zero (β_1) also plays a key role.

In our main specification, β_1 is approximately -0.07 and β_2 is around 0.2. These coefficients imply that the "break-even" point for volume (when past returns generate a zero expected future return) is given by $v_t^* = -\beta_1/\beta_2 \approx -0.35$. It thus follows that volume needs to be around one standard deviation above its mean value before continuations take effect. Up to that level, returns are expected to reverse, with the effect being stronger the lower the level of volume.

These findings from the time-series analysis have clear implications for building a portfolio strategy. A reversal (or contrarian) strategy should generate a positive return. But the effect should be stronger when conditioning the strategy on *low* volume currency pairs. An investor should therefore hold currencies with low prior volume; forming long positions in recently *depreciated* currencies and a short position in recently *appreciated* currencies. We refer to this portfolio as the 'Reversal-Low' (*RevL*) strategy. If volume is economically important, then the excess returns to this strategy should not only be significant, but should also be higher on average than those generated by the alternative reversal

Exhibit B:
3x3 Conditional Double Sort

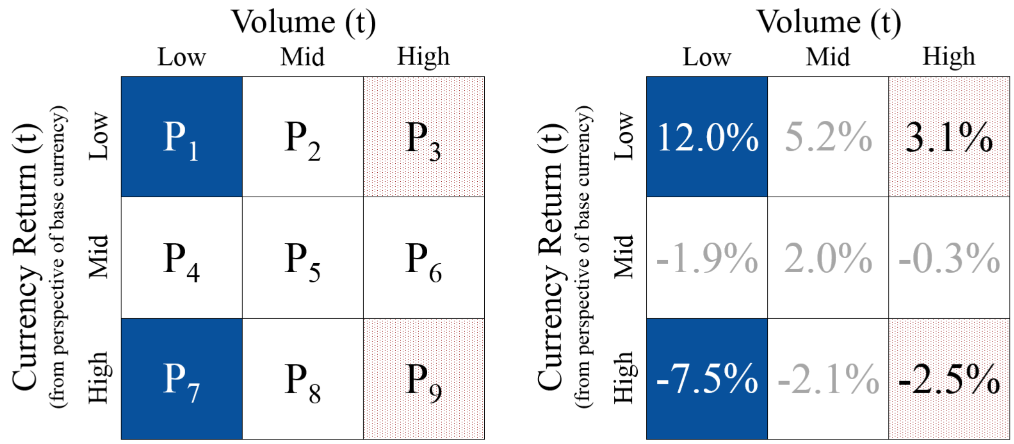

strategy constructed using high volume currency pairs (*RevH*).

The *RevL* strategy. In Exhibit B, we provide a graphical depiction of our conditional double sorting procedure. In the 3x3 case, currency pairs are initially allocated at time t into three groups (from low to high) conditional on their returns between times t-1 and t (in which the return is measured from the perspective of the base currency). Within these groupings, the currency pairs are then allocated into a further three sub-groups (from low to high) conditional on volume between times t-1 and t, and thus we form nine groups, i.e., portfolios, in total. We rebalance these portfolios daily.

In Exhibit B, we also report the annualized mean returns for each portfolio. Beginning from the first row, currencies which previously depreciated the most are likely to appreciate in the following period. The magnitude of the effect is, however, strongly decreasing as volume shifts from low to high. P_1 generates a large annualized mean return of 12%, while P_3 generates a milder return, estimated at 3.1% per annum. Similarly, in the final row, currencies which appreciated over the preceding 24 hours are likely to depreciate subsequently. But again, the effect is amplified if volume was previously low. In this case, P_7 generates a negative annualized return that is 5% lower than that of P_9.

The returns indicate that forming strategies that condition on FX volume could generate substantial economic value. Specifically, the *RevL* strategy should provide strong investment performance and outperform the analogous *RevH* strategy. The *RevL* strategy is formed by taking a long position in P_1 (low return, low volume) and short position in P_7 (high return, low volume). In contrast, the

RevH strategy is formed as the difference between P_3 and P_9.

In Table 3, we report results on the investment performance of the reversal strategies. Starting from the 3x3 conditional double sort with total volume, the *RevL* strategy is found to deliver strong investment performance, consistent with the results presented in Exhibit B. The out-of-sample annualized average return is 19.5%, which is statistically different from zero at the one percent level. In contrast, the return of the *RevH* strategy is 5.6% and not statistically distinguishable from zero.

The outperformance of the *RevL* strategy is also clear when computing the Sharpe ratios of the *RevL* and *RevH* portfolios, which are 1.82 and 0.46 respectively, with the difference being highly statistically significant.[24] While the Sharpe ratio is commonly used to assess investment performance, it exhibits certain drawbacks. For example, the statistic does not take into account the effects of non-normality (Jondeau and Rockinger (2012)), which may be particularly important in a small-sample setting. We therefore also report the *theta* (Θ) performance measure proposed by Ingersoll et al. (2007), which re-estimates the sample mean by putting less weight on outlier returns. We show that for both strategies, Θ is only slightly lower than the average return, indicating that extreme outliers and non-normalities are not driving the strategies' returns. The finding that the returns to *RevL* and *RevH* are not driven by a few outliers, is confirmed in the cumulated returns reported in Figure 2 (upper plots). The absence of significant jumps and the consistent positive slope illustrate that the strong performance of the *RevL* strategy is observed across the entire sample period.[25]

The results remain qualitatively similar when we instead form eight portfolios (2x4 conditional double sort in which currency pairs are sorted into four volume baskets after initially being sorted into two return baskets) or 12 portfolios (3x4 conditional double sort). Interestingly, in each case the maximum drawdown (MDD) is always substantially smaller for the *RevL* strategy compared to the *RevH* strategy.

Finally in Table 3, we report equivalent results for volume in each individual FX instrument. We find equally strong results when replacing total volume with either spot or outright forward volume. In each case the returns are high, while for spot (outright forward) volume, the Sharpe ratio on *RevL*

[24]We test whether two Sharpe ratios are statistically different using the procedure proposed by Ledoit and Wolf (2008). We thank Michael Wolf for kindly making the code available on his website at www.econ.uzh.ch/en/people/faculty/wolf.

[25]In Section 8.3, we further allay concerns surrounding the length of the sample by basing inference on a bootstrap procedure for each performance measure reported in Table 3.

ranges from 1.41 to 1.78 (1.44 to 1.70). The results for swap volume are slightly weaker. The reversal effect is not as pronounced and, in fact the *RevH* strategy even displays a negative return in the 3x4 sort. Nonetheless, the returns are high and the Sharpe ratios on the *RevL* strategy are consistently over one and always statistically greater than the Sharpe ratios on the equivalent *RevH* strategy.

In Internet Appendix Table A.2, we report summary statistics for the nine portfolios formed using the 3x3 conditional double sort. Three key results stand out. First, most portfolio returns are not statistically different from zero and thus sorting on past returns is not guaranteed to generate positive returns. In fact, apart from P_1 and P_7 – the two portfolios that comprise the *RevL* strategy – only P_2 (low return, mid volume) generates a statistically significant return. Second, the turnover of the strategies is extremely high. On average, each portfolio exhibits over 80% turnover each day. It follows that no single currency pair dominates either the P_1 or P_7 portfolios. Third, the high turnover and currency compositions make the *RevL* strategy clearly different from other common currency strategies such as carry, value and momentum.[26]

6 Understanding the *RevL* Returns

In this section, we investigate potential reasons for why the gross returns to the *RevL* strategy are so high. We consider three primary explanations: (i) transaction costs make the strategy unprofitable; (ii) the returns reflect compensation for exposure to risk; and, (iii) risk-oriented capital cannot exploit the trading strategy because the information signal is short lived.

6.1 Transaction Costs

Bid-ask spreads. The majority of trading strategies proposed in the FX literature have typically been rebalanced monthly and thus the issue of transaction costs has been largely innocuous (see, e.g., Lustig, Roussanov, and Verdelhan (2011), Menkhoff et al. (2012b) and Menkhoff et al. (2017)). The

[26]It is useful to note that, because we consider cross-rates in our core analysis, the currency weights in our portfolios are not necessarily equal. In the literature it is common to employ equally weighted "dollar neutral" strategies that strip out any U.S. dollar base effect (see, e.g., Lustig, Roussanov and Verdelhan, 2011). In Internet Appendix Table A.8, we consider this case by presenting the equivalent results for the sub-sample of 16 U.S. dollar currency pairs, in which we ensure we are neutral to the U.S. dollar throughout the sample. We find the results remain qualitatively similar, with high average returns and Sharpe ratios displayed across all FX instruments.

RevL strategy, however, requires daily rebalancing and thus transaction costs may have a material impact on the returns available to investors.

Incorporating transaction costs in the FX market is not trivial. The decentralized over-the-counter nature of the FX market means that traded prices are not publicly available. Instead researchers have typically relied on using an average of dealer quoted spreads recorded at 4pm London time as provided by WM/Reuters (WM/R). It is known that these quoted spreads are much higher than the effective spreads actually paid in the FX market, and thus much of the literature has employed an arbitrary scaling of 50% of the quoted bid-ask spread to proxy for the effective spread available (e.g., Goyal and Saretto (2009), Menkhoff et al. (2012b) and Menkhoff et al. (2017)). Even this number has been viewed as conservative. Gilmore and Hayashi (2011) find transaction costs due to bid-ask spreads are likely to be much lower than 50% of the quoted spread.

We overcome the arbitrary 50% scaling of WM/R bid-ask spreads by collecting data on FX spreads from various dealer platforms. Specifically we employ data from three sources: (i) inter-dealer quotes provided by Olsen Financial Technologies (Olsen), the leading provider of interbank FX quotes across a range of platforms; (ii) quoted spreads from the retail aggregator platform of Dukascopy Bank (Dukascopy), a Swiss based FX broker that services active traders, hedge funds and banks; and, (iii) quoted spreads charged on a single-bank platform by a large global bank, the identity of which we keep anonymous for confidentiality purposes. The collection of this additional data serves two primary purposes. First, it allows us to report with a higher degree of confidence the returns an investor could have achieved by investing in the *RevL* strategy. Second, it allows us to compare the WM/R spreads with alternative data sources to assess whether a scaling of the WM/R spreads is appropriate and, if so, to provide quantitative evidence on the necessary value for the scaling coefficient.

The data from Dukascopy Bank and Olsen Financial Technologies is at hourly frequency across the full sample period. The data from the single-bank platform is available as the average quote across London trading hours (9am to 5pm) for a portion of our sample in 2015. We therefore begin our analysis by comparing the bid-ask spreads reported at 4pm by WM/R with those reported by Dukascopy and Olsen and then assess the reliability of these data through comparison with the quotes from the single-bank platform.

Figure 3 displays the daily median bid-ask spread (as a percentage of the mid price) across currency pairs at 4pm for WM/R, Dukascopy and Olsen. The time-series pattern in spreads is similar across each series – falling during the second half of 2013 before rising and falling again at the start of 2015. The level of spreads is, however, markedly different across series. While Dukascopy and Olsen spreads largely overlap, WM/R spreads are, on average, substantially higher. From 2014 to the end of 2016, the median spread was around 0.05% according to WM/R, but around 0.01% according to Dukascopy and Olsen. Indeed from the figure it can be seen that even using a 50% WM/R scaling, the spread remains around twice the actual market level.

In Table A.4 of the Internet Appendix, we provide a currency-by-currency breakdown of the average bid-ask spread at 4pm for WM/R, Dukascopy and Olsen. In addition we report the ratio of the average spread for each pair. On average Dukascopy and Olsen spreads are around 25% the level of WM/R spreads, although the ratio varies from a low of 0.08 (EURUSD) to 1.02 (USDMXN). Furthermore, the overlap in Dukascopy and Olsen spreads seen in Figure 3 is again observable in Table A.4 – the average ratio of spreads is very close to one (0.94). We conclude that the use of a WM/R 50% scaling rule is likely to substantially understate the returns an investor can expect by following a currency investment strategy.

In Table A.5 of the Internet Appendix, we provide a similar comparison between the Dukascopy and Olsen spreads with the quotes from the single-bank platform. Reassuringly, the quotes line up closely. Across the currency pairs that are tradable on the single-bank platform, the median quotes are only marginally wider (average ratios of 0.95 and 0.94 relative to Olsen and Dukascopy). We view these results as further confirmation of the need to use alternatives to WM/R spreads when incorporating bid-ask spreads in empirical research. We conclude that, at least for strategies covering recent years, researchers should consider halving the rule-of-thumb again – a 25% scaling rule appears to more accurately reflect the bid-ask faced by both retail and institutional investors. In the following analysis we compare the returns available to an investor when incorporating Dukascopy, Olsen and WM/R 25% bid-ask spreads.[27]

[27]We do, however, acknowledge a caveat associated with incorporating transaction costs: the inability to fully incorporate *all* trading costs. These additional costs include price 'slippage' and spreads widening when transaction sizes are particularly large. Both of these effects would lower our reported returns and should be incorporated by market participants when entering higher frequency strategies. Nonetheless, our returns should still provide a fair reflection of the returns available to many market participants that do not cause major price movements through trading nor trade

Results. The results from incorporating bid-ask spreads are reported in Table 4. As expected, the transaction costs substantially reduce the profitability of the trading strategies. When forming portfolios based on total volume, we observe a reduction in the annualized out-of-sample mean returns of between 6% and 7.5% depending on the choice of spreads. Nevertheless, the returns to the *RevL* strategy remain positive and highly statistically significant – at the one-percent level using Dukascopy spreads and at the five-percent level using Olsen and WM/R 25%. Moreover, the Sharpe ratios are consistently over 1.0. In contrast, the average returns to the *RevH* strategy turn negative indicating there are no economic gains from conditioning on high volume. Finally, the *theta* values confirm that non-linearities and outliers do not significantly affect our results.

Compared to the results reported in Table 3, the difference between FX instruments is more evident. Across the three sorts, the average returns to the *RevL* strategy are usually significant at the five percent level when conditioning on spot and outright forward volume. The *RevL* strategy does not, however, generate a statistically significant return when conditioning on swap volume alone. The results are consistent with the earlier findings from the panel regressions: swap volume contains less predictive information for currency excess returns than spot and forward volume.

6.1.1 Implementation of the *RevL* strategy at different hours of the trading day

The previous section indicates that the 4pm bid-ask spreads do not wipe out the returns to the *RevL* strategy. We explore whether this finding reflects the fact that bid-ask spreads are low during this part of the trading day, when both London and New York traders are at their desks and thus whether the returns disappear at other points of the day. To do so, we re-run the previous exercise by forming portfolios at each hour of the trading day and then calculate gross returns as well as net returns using Olsen bid-ask spreads.[28]

We report the results in Figure 4. The gross returns can be seen to be high throughout the entire day, such that the return is always above 10% and ranges to over 20% on five occasions. The effect of transaction costs, however, is substantial and more apparent with the introduction of low-volume periods when bid-ask spreads are at their widest. Most strikingly, the net returns during Sydney

in especially large volume.

[28] We also perform the analysis using Dukascopy spreads and achieve almost identical results, which are available upon request.

opening hours are negative, which is particularly surprising given that those hours also correspond with the highest gross returns. To see why, we also plot the median bid-ask spread across each hour (white boxes, right-hand-side axis). It can be seen that bid-ask spreads during Sydney open are between 200% and 300% higher than those observed during peak volume periods. Interestingly, the highest net returns are observed at 4pm London time despite it not being the point of highest gross returns for the strategy. We thus conclude that while bid-ask spreads cannot generally account for the large gross returns, they do place constraints on when the strategy can be successfully implemented.

6.2 Risk

A second possible explanation for the high returns to the *RevL* strategy is that the returns reflect compensation for exposure to risk. In the model of Campbell, Grossman, and Wang (1993), for example, a spike in risk aversion among a sub-population results in a fall in the price of risky securities and a rise in expected returns.

We explore the risk compensation hypothesis by running a series of OLS regressions, in which we first regress gross *RevL* returns and then net *RevL* returns on the daily returns associated with recently proposed currency risk factors and strategies. Specifically, we consider: dollar and slope risk (DOL and CAR) proposed by Lustig, Roussanov, and Verdelhan (2011); global FX volatility risk (VOL) as constructed in Menkhoff et al. (2012a); 12-month currency momentum (MOM) (see Menkhoff et al. (2012b) and Asness, Moskowitz, and Pedersen (2013) for details); and illiquidity risk (ILL) using the Corwin-Schultz measure (see Mancini, Ranaldo, and Wrampelmeyer (2013) and Karnaukh, Ranaldo, and Söderlind (2015) for further details). In each model we include the DOL ('market') factor with one of the other factors. The use of a two-factor linear stochastic discount factor, in which DOL is the first factor is standard in much of the recent literature on currency market risk (Lustig, Roussanov, and Verdelhan (2011); Mancini, Ranaldo, and Wrampelmeyer (2013); Della Corte, Riddiough, and Sarno (2016)).

We report results in Table 5. In columns 1-4 we present results when regressing gross *RevL* returns on the various factors and strategies. In columns 5-8 we do the same for net returns. The most immediate observation from columns 1-4 is that none of the factors explain much of the daily variation in *RevL* returns. In each case the constant (α) remains high and statistically significant, indicative of large model 'mispricing'. The goodness of fit test supports this conclusion: the adj-R^2 statistic is always close to zero. The smallest pricing error is observed for the model that includes the illiquidity factor. Interestingly, the sign on the coefficient is negative – indicating that the strategy is negatively related to illiquidity risk. This implies that when the illiquidity factor of Mancini, Ranaldo, and Wrampelmeyer (2013) is low (i.e., when there is *less* liquidity available in the FX market), the returns to the *RevL* strategy improve. This finding may be because higher expected returns are required by liquidity providers during times of low liquidity, which generates even larger subsequent return reversals.

The inclusion of bid-ask spreads in columns 5-8 (only for the *RevL* strategy) reduces the mispricing generated by the models. Once again, however, the mispricing persists in the case of the models including CAR, VOL and MOM. There is evidence, however, that the inclusion of both transaction costs and liquidity may reduce the mispricing of the strategy to a level at which α is not statistically different from zero. As noted above, however, the sign of the coefficient implies the strategy offers a possible *hedge* against illiquidity risk, rather than reflecting compensation for exposure to it. We therefore provide tentative evidence that the size of the strategy's return is also partly conditional on the level of liquidity on the day the strategy is implemented. The alpha available on the strategy still remains economically high, however, at around 7% and thus the lack of statistical significance is partly due to the combination of a short sample period and comparatively high volatility of the *RevL* strategy, which prevents us from drawing firmer conclusions.

6.3 Tradeability

We have so far learned that transactions costs and commonly accepted sources of risk cannot fully account for the large returns generated by the *RevL* strategy. These findings could imply that the returns simply reflect compensation for some yet unidentified risk factor – perhaps relating to the risk-bearing capacity of market dealers – or alternatively, the returns could reflect a market inefficiency – providing a possible 'free lunch' to investors that are able to enter into the strategy. A possible reason for this 'inefficiency' is that comparatively little risk-orientated capital *can* currently exploit the strategy. The volume data was unavailable to investors for much of our sample period and, even when it did become available during 2016, it was released by CLS with a one-day lag. If the information content of the volume signal is short lived, then even a single day's delay could render the economic content of the signal worthless.

In this section, we examine the speed by which the signal associated with volume loses value. To do so, we calculate the signal at 9am each day and, unlike in previous sections in which we assume the portfolio is instantly formed, we allow for delay in portfolio formation in incremental one-hour time intervals. We display the results of the exercise in Figure 5. The height of each bar in the figure represents the total (gross) return associated with a *RevL* strategy that uses the 9am conditioning information but whose formation time is delayed by between one and 12 hours (+1h implies formation at 10am, +2h reflects formation at 11am and so forth).

The clear downward sloping pattern implies that the economic usefulness of the signal is short lived. After seven hours the expected return is only half its original size. This value falls to just one-quarter following a 12 hour delay. The speed of decay in the signal's economic value becomes more pronounced when incorporating bid-ask spreads. After an eight hour formation delay, an investor is no longer expected to profit from the strategy and, in fact, would likely generated post-transaction-cost losses if the signal was received any later. These results imply that the returns are currently unavailable for most investors. But if the volume data becomes available at a more timely interval, it will be interesting to assess the continuing performance of the strategy and to reflect again on the underlying drivers of the return process. Alternatively, timely volume information could be captured by a top FX dealer as long as their market share is large enough to be able to infer market volume from the order flows that are executed by the dealer.

7 Diversification with *RevL*

The high returns to the *RevL* strategy, which we documented in the previous sections, raise the prospect that the strategy can be employed by global investors to improve the investment performance of a currency portfolio. In this section, we explore this possibility in two ways. First, we analyze the performance of a portfolio that combines standard currency market strategies with the *RevL* strategy, and assess the economic and statistical improvement in the portfolio's performance. Second, we employ data on the returns generated by currency hedge funds to assess the extent to which the *RevL* strategy explains those returns and by doing so provide evidence on the extent to which the strategy (or a related strategy) is adopted in current practice.

Common currency strategies. We construct U.S. dollar-base portfolios against the other 16 major currencies in our study. The portfolios are constructed by sorting currencies into one of five portfolios by either forward premia (Carry), past one-month exchange rate returns (momentum) or deviations from the real exchange rate (value).[29] We form individual strategies from these five portfolios by constructing an equally weighted 'high-minus-low' strategy that goes long in Portfolio 5 and short in Portfolio 1. The strategies are rebalanced monthly.

In the case of forward premia, we go long currencies with the highest interest rates and short currencies with the lowest interest rates. We refer to the strategy as *CAR*. Momentum is based on going long currencies that appreciated the most over the previous year and short those which depreciated the most over the previous year. We refer to this strategy as *MOM*. Finally, value is formed by going long currencies that are deemed to be the most undervalued and short those deemed to be the most overvalued. We refer to the strategy at *VAL*. We also construct a 'market' portfolio that allocates an equally weighted position in all currencies against the U.S. dollar. We refer to the strategy as *DOL*.

It is widely recognized that since the global financial crisis, common FX strategies including carry, momentum and value have markedly underperformed, leading to substantial underperformance by currency hedge funds (Wong and Arnold (2017)). Table 6 presents evidence consistent with this observation. In Panel A, we present the average pre-transaction cost returns of the four trading strategies between December 2011 and December 2016. Carry was the best performer, generating an annualized return of 3.6% and Sharpe ratio of 0.50. Momentum and value, however, have been notably

[29] Carry, momentum and value have been studied extensively in the currency market literature; see, e.g., Lustig, Roussanov, and Verdelhan (2011); Asness, Moskowitz, and Pedersen (2013); Menkhoff et al. (2012a) and Menkhoff et al. (2017). We build value portfolios following the procedure proposed in Asness, Moskowitz, and Pedersen (2013). To do so, we collected Consumer Price Index (CPI) data from the International Monetary Fund's *International Financial Statistics* database. Value is measured as the difference between the five-year log change in the CPI differential and the five-year log exchange rate return (where the return is relative to the average exchange rate observed over the period 4.5 to 5.5 years earlier).

less successful. In fact, the returns to the momentum portfolio were negative over the period. The U.S. dollar has appreciated strongly since 2014, which accounts for the underperformance of the *DOL* strategy. In the third row of Panel A, we present the monthly return correlations between the four strategies and *RevL* strategy, for which we construct monthly returns simply by summing the daily returns in each month. In each case the correlation is low, indicating that substantial diversification benefits could be enjoyed from including the *RevL* strategy in the portfolio, which may substantially help in improving the investment performance of currency funds.

In Panel B, we present the investment performance of three currency portfolios that combine the currency strategies in Panel A: an equally weighted portfolio (EW), a global minimum variance portfolio (MV), and the 'tangency' portfolio (TG). The portfolios are rebalanced monthly. The equally weighted portfolio simply assigns 25% weight to each portfolio every month. The global minimum volatility portfolio is the portfolio with the lowest return volatility, representing the solution to the following optimization problem: min $w'\Sigma w$ subject to the constraint that the weights sum to unity $w'\iota = 1$, where w is the $N \times 1$ vector of portfolio weights on the risky assets, ι is a $N \times 1$ vector of ones, and Σ is the $N \times N$ covariance matrix of the asset returns. The weights of the global minimum volatility portfolio are given by $w = \frac{\Sigma^{-1}\iota}{\iota'\Sigma^{-1}\iota}$. The tangency portfolio maximizes the Sharpe ratio, the weights for which are given by $w = \frac{\mu}{\mu'\Sigma^{-1}\mu}\Sigma^{-1}\mu$ where μ is the $N \times 1$ vector of expected strategy returns. We compute the optimal weights across our entire sample. Apart from the equally weighted strategy, the returns are therefore measured in-sample.

The Sharpe ratios are found to fluctuate between -0.03 for the naive equally weighted strategy to over 0.55 for the tangency portfolio. These values are dwarfed, however, when we include the *RevL* strategy within the set of possible currency strategies. The strategy (even though it is included based on the *after* t-cost returns) substantially increases the investment performance. For the equally weighted portfolio, the Sharpe ratio is 0.58, while in the case of the optimized portfolio, the *RevL* strategy is found to increase the Sharpe ratio by around two to three times. Overall, the results point towards the *RevL* strategy providing a novel and highly desirable source of diversification for currency investors.[30]

[30]One may wonder whether the fact that the *RevL* returns originate from a daily rebalancing setting while the other strategies only rebalance monthly over-estimates the relative performance of *RevL* or alters the correlations across strategies. This is not the case: we also calculate correlations and portfolios returns when all currency strategies are

Currency hedge funds. In the preceding analysis, we implicitly assumed that currency investors do not currently follow the *RevL* strategy and thus, when combined with other commonly implemented strategies, would benefit from the properties of its returns. We now turn to the returns generated by currency hedge funds to assess if this conclusion remains valid in the context of the actual behavior of currency funds. To do so, we collect monthly currency hedge fund returns from BarclayHedge, a data provider that specializes in hedge funds.[31]

Within the database, funds self-categorize their currency trading strategies as either "fundamental" or "technical". A fund that relies on technical indicators may have a greater propensity to follow recent trends and place less emphasis on value strategies. Since the returns, *ex-ante*, are potentially quite different across the two groups, in the following analysis we explore the returns for three groups: all currency funds; fundamental funds; and technical funds. In particular, we are interested in assessing if the returns to these currency funds are related to the *RevL* strategy. If so, the addition of the strategy may provide less overall benefit to the portfolio than implied by the previous analysis. To test whether this is the case, we run the following panel regression:

$$r_{i,t} = \alpha_i + \tau_t + \beta RevL_t + \boldsymbol{\gamma}'\boldsymbol{x}_t + \epsilon_{i,t+1}, \tag{5}$$

where α_i and τ_t denote fund and time fixed effects, $r_{i,t}$ is the return of the currency hedge fund i in month t, $RevL_t$ is the return of the *RevL* strategy, and the vector $\boldsymbol{x}_t$ contains the following currency factors and strategies: the currency trend-following factor (TFRX) of Fung and Hsieh (2001); the dollar and slope factors (DOL and CAR) proposed by Lustig, Roussanov, and Verdelhan (2011); momentum and value (MOM and VAL) as calculated by Asness, Moskowitz, and Pedersen (2013)); the FX illiquidity factor (ILL) of Mancini, Ranaldo, and Wrampelmeyer (2013), which we construct using bid-ask spreads (ba); and the Corwin-Schultz (cs) measure (see Karnaukh, Ranaldo, and Söderlind (2015) for further details).

We report results in Table 7. The results in column 1 are based on all currency hedge funds reporting to BarclayHedge at December 2016, which are then further split between currency funds

rebalanced daily and find no qualitative impact on the results. These results are available upon request.

[31] The currency hedge fund returns are self-reported to BarclayHedge and have been used in other recent studies (see Nucera and Valente (2013) and the references herein).

that trade according to technical indicators (column 2) and currency funds that trade according to fundamental signals (column 3). We first note that across 'All Funds' the *RevL* strategy does not appear to be related to currency hedge fund returns – signaling a potentially large performance gain if the strategy were adopted within the hedge fund community. The technical fund returns are, on average, negatively related to the strategy. This finding is not altogether surprising. The *RevL* strategy is contrarian by nature – buying (selling) recently depreciated (appreciated) currencies. In contrast, 'chartists' or 'technical' traders are thought to follow trends or momentum in returns. This presumption is further confirmed by the positive relationship between technical fund returns and the trend-following factor.

Introducing a contrarian trading strategy within the broader portfolio could therefore generate substantial gains within technically-orientated funds. We also note that, while fundamental fund returns are not related to the *RevL* strategy, they also appear to follow trend-following strategies in addition to the usual carry and dollar-directional trading. Surprisingly, fundamental firm returns are not systematically driven by the currency value strategy, suggesting that fundamentally orientated firms only selectively trade towards the 'anchor' exchange rate implied by purchasing power parity or possibly use more sophisticated measures of currency valuation. Finally, we find the illiquidity factor is also largely unrelated with the variation in currency hedge fund returns, suggesting these funds do not systematically seek exposure to liquidity risk.

8 Robustness tests

In this section we present a series of tests to assess the robustness of the *RevL* returns. First, we measure volume using alternative trends as well as by standardizing the volume series. Second we explore the performance of the strategy using various sub-samples of currency pairs. Finally, we extend our examination of the *RevL* returns by constructing bootstrapped p-values.

8.1 *RevL* using alternative detrending and standardization procedures

In Equation (3) we define our measure of aggregate volume as the (log) deviation from its previous one month (21 trading days) trend. We first investigate if this choice of window over which to measure the trend has a material impact on the returns to the *RevL* strategy. In Panel A of Table A.6 in the Internet Appendix, we compare the returns to the *RevL* strategy when using the 21-day window with the returns to the strategy when estimating the trend over 63, 126 and 252 trading days. We find that in each case the returns remain comparably high (between 17.3% and 18.6% per annum) and the Sharpe ratios are consistently over 1.6. Moreover, the performance of the *RevL* strategy is always significantly stronger (in terms of both returns and Sharpe ratios) than the *RevH* strategy.

We next consider whether our sorting procedure is influenced by the cross-sectional variation in volume. If the variation in volume is high across a few currency pairs, it increases the likelihood of those currencies driving the returns to the *RevL* and *RevH* strategies. We control for the cross-sectional variability in volume by standardizing the expression in Equation (3) by the series' standard deviation, calculated over the same period as the trend. The results, reported in Panel B of Table A.6, are largely comparable with Panel A. The results become slightly stronger in the case of the 21 day and 63 day standardizations and moderately weaker in the 126 day and 252 day cases. Across all specifications, however, the investment performance remains strong. The Sharpe ratios are always above 1.5 and the difference between *RevL* and *RevH* is pronounced across each measure of volume.

8.2 *RevL* across subsamples of currency pairs

A potential concern arising from our earlier analysis is that the economic profitability of the *RevL* strategy is driven by just a few currency pairs. We investigate this possibility by running the analysis on subsets of the 31 currency pairs in our sample. We report results in Table A.7 of the Internet Appendix. We consider four cases: (i) only USD base pairs (16 currency pairs); (ii) only EUR and GBP base pairs (13 currency pairs); (iii) all pairs excluding emerging market, fixed and pegged pairs (24 currency pairs); (iv) and the G10 plus most liquid EUR crosses (14 currency pairs).

In each case the investment performance of the *RevL* strategy remains impressive. The returns range from around 15% to almost 19% and are statistically significant at the 1% level in each case. The Sharpe ratios also remain high and, except for the EUR and GBP base currency subsample, are statistically higher than those observed for the *RevH* strategy. We thus conclude that our results are robust to the choice of currency pairs in the sample.

8.3 *RevL* returns and bootstrapped p-values

While the *RevL* strategy generates substantial returns in our sample, a possible concern is that the sample itself is relatively short and, given the returns are positively skewed, may lead to an overstatement of the statistical significance of the strategy's returns. We address this concern by adopting a bootstrap procedure, similar to Goyal and Welch (2008) and Mark (1995), to generate p-values for our various measures of investment performance. Specifically, we generate data under the null hypothesis that neither past volume nor past returns have any predictive ability, and that instead the data are generated according to the following system:

$$rx_{i,t+1} = \alpha_i + \varepsilon_{i,t+1} \tag{6}$$

$$v_{i,j,t+1} = \mu_{i,j} + \rho_{i,j} \times v_{i,j,t} + \eta_{i,j,t+1} \tag{7}$$

where $rx_{i,t}$ is the log excess return for currency pair i at time t, and $v_{i,j,t}$ is the dollar trading volume for currency pair i at time t for instrument j, where j={total, spot, forward, swap}. We begin by estimating the system via OLS equation-by-equation for each currency pair, and obtain a vector of residuals $\{\boldsymbol{\hat{E}_t} = (\boldsymbol{\hat{\epsilon}}_t,\ \boldsymbol{\hat{\eta}}_{j=total,t},\ \boldsymbol{\hat{\eta}}_{j=spot,t},\ \boldsymbol{\hat{\eta}}_{j=forward,t},\ \boldsymbol{\hat{\eta}}_{j=swap,t})'\}_{t=1}^{T}$ where $\boldsymbol{\hat{\epsilon}_t} = (\hat{\epsilon}_{1,t},\ \hat{\epsilon}_{2,t}, ...,\ \hat{\epsilon}_{31,t})$, and $\boldsymbol{\hat{\eta}_{j,t}} = (\hat{\eta}_{1,j,t},\ \hat{\eta}_{2,j,t}, ...,\ \hat{\eta}_{31,j,t})$. In order to generate a series of disturbances for our bootstrapped sample, we randomly draw with replacement $T + 100$ times from the residuals $\boldsymbol{\hat{E}_t}$, yielding a bootstrapped series of residuals $\{\boldsymbol{\hat{E}_t^b}\}_{t=1}^{T+100}$.[32] We draw from the residuals in tandem to preserve the contemporaneous correlation between the disturbances in the original sample across instruments and currency pairs.

Using $\{\boldsymbol{\hat{E}_t^b}\}_{t=1}^{T+100}$ and setting the initial observations in Equation (7) equal to the unconditional means of the respective series, we build a bootstrapped sample of $T + 100$ observations, $\{\boldsymbol{rx_t^b}, \boldsymbol{v^b_{j=total,t}},\ \boldsymbol{v^b_{j=spot,t}},\ \boldsymbol{v^b_{j=forward,t}},\ \boldsymbol{v^b_{j=swap,t}}\}_{t=1}^{T+100}$ where $\boldsymbol{rx_t^b} = (rx^b_{1,t},\ rx^b_{2,t}, ...,\ rx^b_{31,t})$, and $\boldsymbol{v^b_{j,t}} = (v^b_{1,j,t},\ v^b_{2,j,t}, ...,\ v^b_{31,j,t})$ are the vectors containing the bootstrapped values across the 31 currency pairs in our sample. For each bootstrap sample we perform the portfolio exercise described in Section 5. We repeat this process 5,000 times, providing an empirical distribution of the performance

[32] We drop the first 100 observations to randomize the initial observations and thus consider a bootstrapped sample of T observations that matches the original sample length.

measures used in Table 3 – the average annualized return, Sharpe Ratio, the performance fee measure of Ingersoll et al. (2007), and maximum drawdown – under the null hypothesis of no predictability. For each performance measure, the bootstrapped p-value is calculated using the proportion of the bootstrap values greater than the ones computed using the original sample.

Figure A.3 reports the result for the 3x3 *RevL* strategy based on total volume.[33] The vertical dashed line in each sub-figure denotes the value we record in Table 3. In each case, the realized investment performance of the *RevL* strategy is found to be substantially different from the expected value under the null of no predictability. In only 0.02% of simulations is a higher average return, Sharpe ratio or Θ observed. Furthermore, in almost 99% of the simulations, the maximum drawdown of the strategy was higher than the one we documented in our sample. Overall, we view these findings as consistent with our earlier conclusion that the returns are highly statistically significant and suggestive of a dynamic relationship between FX volume and currency excess returns.

9 Conclusions

Using a novel dataset from the leading FX settlement platform, we find that FX volume contains information that is both statistically and economically powerful for understanding price behavior in the currency market. Consistent with theories that view volume as a proxy for otherwise unobservable market characteristics, we show that volume – once interacted with current excess returns – can help predict future currency excess returns in both the time-series and cross-section. When volume is low, prices tend to subsequently reverse. However, this reversal effect is reduced when volume is high, consistent with high volume indicating the presence of more informed trading. Exploiting this pattern in a portfolio setting generates substantial investment gains, including an annualized return of over 19% and a Sharpe ratio of 1.82. The investment performance remains robust to the incorporation of daily transaction costs, for which we provide new evidence on accurate measurement of bid-ask spreads from a retail platform, a reputable provider of interdealer quotes, and quotes from the single platform of a global bank. Moreover, the returns are unrelated to other commonly traded currency factors and strategies and thus offer substantial diversification benefits, which we highlight using hedge

[33] Results for the 2x4 and 3x4 sorts, as well as for the individual FX instruments, are available upon request.

fund return data.

The findings contribute to the recent literature casting new light on the comparatively opaque FX market and documenting new sources of currency return predictability. Moreover, the richness of the volume data allows us to show that markets with more 'speculative' activity – including the FX spot and forward markets – generate sharper and more statistically robust results than the FX swap market, which is characterized by a larger proportion of inter-dealer volume and hedging activity driven by institutional investors and multinational corporations.

Much work remains to be done to understand volume in FX markets and, more generally, in over-the-counter markets. Future theoretical contributions could seek to understand the sources of information contained in volume across different instruments and to provide new insights into the optimal level of volume transparency. Empirical work may seek to explore the determinants of volume and investigate its interaction with other publicly available information, to provide a more nuanced perspective on the relationship between volume and returns.

References

Admati, A. R., and P. Pfleiderer. 1988. A theory of intraday patterns: Volume and price variability. *Review Financial Studies* 1:3–40.

Asness, C. S., T. J. Moskowitz, and L. H. Pedersen. 2013. Value and momentum everywhere. *Journal of Finance* 68:929–85.

Baba, N., and F. Packer. 2009. Interpreting deviations from covered interest parity during the financial market turmoil of 2007–08. *Journal of Banking & Finance* 33:1953–62.

Bekaert, G., and G. Panayotov. 2016. Good carry, bad carry. *Columbia Business School Research Paper* No. 15-53.

Berger, D. W., A. P. Chaboud, S. V. Chernenko, E. Howorka, and J. H. Wright. 2008. Order flow and exchange rate dynamics in electronic brokergae system data. *Journal of International Economics* 75:93–109.

Berry, T. D., and K. M. Howe. 1994. Public information arrival. *Journal of Finance* 49:1331–46.

Bjønnes, G. H., and D. Rime. 2005. Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics* 75:571–605.

Campbell, J. Y., S. J. Grossman, and J. Wang. 1993. Trading volume and serial correlation in stock returns. *Quarterly Journal of Economics* 108:905–39.

Chen, J., H. Hong, and J. C. Stein. 2001. Forecasting crashes: Trading volume, past returns and conditional skewness in stock prices. *Journal of Financial Economics* 61:345–81.

Cheung, Y.-W., and C. Y.-P. Wong. 2000. A survey of market practictioners's views on exchange rate dynamics. *Journal of International Economics* 401–19.

Chordia, T., R. Roll, and A. Subrahmanyam. 2001. Market liquidity and trading activity. *Journal of Finance* 56:501–30.

Chordia, T., A. Subrahmanyam, and R. V. Anshuman. 2001. Trading activity and expected stock returns. *Journal of Financial Economics* 59:3–32.

Chordia, T., and B. Swaminathan. 2000. Trading volume and cross-autocorrelations in stock returns. *Journal of Finance* 55:913–35.

Conrad, J. S., A. Hameed, and C. Niden. 1994. Volume and autocovariances in short-horizon individual security returns. *Journal of Finance* 49:1305–29.

Cooper, M. 1999. Filter rules based on price and volume in individual security overreaction. *Review of Financial Studies* 12:901–35.

Datar, V. T., N. Y. Naik, and R. Radcliffe. 1998. Liquidity and stock returns: An alternative test. *Journal of Financial Markets* 1:203–2019.

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

Della Corte, P., J. S. Riddiough, and L. Sarno. 2016. Currency premia and global imbalances. *Review of Financial Studies* 29:2161–93.

Du, W., A. Tepper, and A. Verdelhan. 2017. Deviations from covered interest rate parity. *Journal of Finance (forthcoming)* .

Evans, M. D. D., and K. R. Lyons. 2002. Order flow and exchange rate dynamics. *Journal of Political Economy* 110:170–80.

Evans, M. D. D. 2002. FX trading and exchange rate dynamics. *Journal of Finance* 57:2405–47.

Fisher, A. M., and A. Ranaldo. 2011. Does FOMC news increase global FX trading? *Journal of Banking & Finance* 35:2965–73.

Foster, D. F., and S. Viswanathan. 1993. Variations in trading volume, return volatility, and trading costs: Evidence on recent price formation models. *Journal of Finance* 48:187–211.

Fung, W., and A. D. Hsieh. 2001. The risk in hedge fund strategies: Theory and evidence from trend followers. *Review of Financial Studies* 14:313–41.

Galati, G., and A. Heath. 2008. What drives the growth in FX activity? Interpreting the 2007 triennial survey. *BIS Quarterly Review* 63–72.

Gallant, R. A., P. E. Rossi, and G. Tauchen. 1992. Stock prices and volume. *Review of Financial Studies* 5:199–242.

Gerety, M. S., and H. J. Mulherin. 1992. Trading halts and market activity: An analysis of volume at the open and the close. *Journal of Finance* 47:1765–84.

Gervais, S., R. Kaniel, and D. H. Mingelgrin. 2001. The high-volume return premium. *Journal of Finance* 56:877 – 919.

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

Goyal, A., and A. Saretto. 2009. The cross-section of option returns and volatility. *Journal of Financial Economics* 94:310–26.

Goyal, A., and I. Welch. 2008. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 21:1455–508.

Grammatikos, T., and A. Saunders. 1986. Futures price variability: A test of maturity and volume effects. *Journal of Business* 59:319–30.

Harris, M., and A. Raviv. 1993. Differences of opinion make a horse race. *Review of Financial Studies* 6:473–506.

Hasbrouck, J., and R. M. Levich. 2017. FX market metrics: New findings based on CLS bank settlement data. *NBER Working Paper* No. 23206.

Ingersoll, J., M. Spiegel, W. Goetzmann, and I. Welch. 2007. Portfolio performance manipulation and manipulation-proof performance measures. *Review Financial Studies* 20:1503–46.

Ito, T., K. R. Lyons, and T. M. Melvin. 1998. Is there private information in the FX market? The Tokyo experiment. *Journal of Finance* 52:1111–30.

Jain, P. C., and G.-H. Joh. 1988. The dependence between hourly prices and trading volume. *Journal of Financial and Quantitative Analysis* 23:269– 283.

Jennings, R. H., L. T. Starks, and J. C. Fellingham. 1981. An equilibrium model of asset trading with sequential information arrival. *Journal of Finance* 36:143–61.

Jondeau, E., and M. Rockinger. 2012. On the importance of time variability in higher moments for asset allocation. *Journal of Financial Econometrics* 10:84–123.

Jordá, O., and M. A. Taylor. 2012. The carry trade and fundamentals: Nothing to fear but feer itself. *Journal of International Economics* 88:74–90.

Karnaukh, N., A. Ranaldo, and P. Söderlind. 2015. Understanding FX liquidity. *Review of Financial Studies* 28:3073–108.

Karpoff, J. M. 1987. The relation between price changes and trading volume: A survey. *Journal of Financial and Quantitative Analysis* 22:109–26.

King, M. R., L. C. Osler, and D. Rime. 2012. Foreign exchange market structure, players, and evolution. In J. James, I. W. Marsh, and L. Sarno, eds., *Handbook of Exchange Rates*. Wiley & Sons, Hoboken, New Jersey.

King, M. R., and A. Schrimpf. 2010. The $4 trillion question: what explains FX growth since the 2007 survey? *BIS Quarterly Review* 27–42.

Ledoit, O., and M. Wolf. 2008. Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance* 15:850–9.

Levich, R. M. 2012. FX counterparty risk and trading activity in currency forward and futures markets. *Review of Financial Economics* 21:102–10.

Llorente, G., R. Michaely, G. Saar, and W. Jiang. 2002. Dynamic volume-return relation of individual stocks. *Review of Financial Studies* 15:1005–47.

Lustig, H., N. Roussanov, and A. Verdelhan. 2011. Common risk factors in currency markets. *Review of Financial Studies* 24:3731–77.

Lyons, K. R. 1997. A simultaneous trade model of the foreign exchange hot potato. *Journal of International Economics* 42:275–98.

———. 1995. Test of microstructural hypothesis in the foreign exchange market. *Journal of Financial Economics* 39:321–51.

———. 2001. The microstructure approach to exchange rate. *MIT Press* .

Mancini, L., A. Ranaldo, and J. Wrampelmeyer. 2013. Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums. *Journal of Finance* 68:1805–41.

Mark, N. C. 1995. Exchange rates and fundamentals: Evidence on long-horizon predictability. *American Economic Review* 85:201–2018.

Melvin, T. M., and M. Taylor. 2009. The crisis in the foreign exchange market. *Journal of International Money and Finance* 28:1317–30.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf. 2012a. Carry trades and global foreign exchange volatility. *Journal of Finance* 67:681–718.

———. 2012b. Currency momentum strategies. *Journal of Financial Economics* 106:660 – 684.

———. 2017. Currency value. *Review of Financial Studies* 30:416–41.

———. 2016. Information flows in foreign exchange markets: Dissecting customer currency trades. *Journal of Finance* 71:601 – 634.

Moore, M., A. Schrimpf, and V. Sushko. 2016. Downsized FX markets: Causes and implications. *BIS Quarterly Review* 35–52.

Nucera, F., and G. Valente. 2013. Carry trades and the performance of currency hedge funds. *Journal of International Money and Finance* 33:407–25.

Osler, L. C., and V. Vandrovych. 2009. Hedge funds and the origins of private information in currency markets. *Brandeis International Business School Working Paper* .

Payne, R. 2003. Informed trade in spot foreign exchange markets: An empirical investigation. *Journal of International Economics* 61:307–29.

Peiers, B. 1997. Informed traders, intervenation, and price leadership: A deeper view of the microstructure of the foreign exchange market. *Journal of Finance* 52:1589–614.

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

Rime, D., and A. Schrimpf. 2013. The anatomy of the global FX market through the lens of the 2013 triennial survey. *BIS Quarterly Review* 27–43.

Rime, D., A. Schrimpf, and O. Syrstad. 2017. Segmented money markets and covered interest parity arbitrage. *BIS Working Paper* No. 651.

Wang, J. 1994. A model of competitive stock trading volume. *Journal of Political Economy* 102:127–68.

Wong, A., and A. Arnold. 2017. Classic models are failing FX hedge funds desperate for return. *Bloomberg On-line Edition (17th-March)* .

Table 1: A First Look at the Foreign Exchange Volume Data

	Summary Statistics				Correlations (%)					
					Instruments			N. Trades and Dollar Volume		
	Mean	Median	StDev	AC(1)	Spot and Forward	Swap and Forward	Spot and Swap	Spot	Forward	Swap
G10										
EURUSD	499.95	497.70	83.75	0.41	48.53	26.62	7.83	92.74	53.22	71.30
USDJPY	244.00	241.43	62.02	0.49	51.92	33.78	39.55	97.08	52.13	78.05
GBPUSD	186.25	183.77	31.44	0.41	44.98	18.45	-3.61	89.48	45.17	67.96
AUDUSD	123.98	123.76	27.45	0.58	50.70	27.10	27.90	96.81	46.36	73.14
USDCHF	86.81	86.57	16.96	0.38	42.52	21.59	9.35	81.85	31.79	82.67
USDCAD	80.92	80.25	14.92	0.39	57.61	40.66	52.61	84.26	41.11	54.62
NZDUSD	32.61	31.63	8.17	0.59	45.96	17.77	14.54	95.76	74.69	41.76
USDSEK	27.17	26.73	6.57	0.43	32.83	17.23	14.94	57.99	33.23	63.33
USDNOK	22.48	22.58	4.88	0.25	34.15	16.13	19.25	48.63	37.03	65.50
USD Pairs										
USDMXN	28.27	28.11	6.96	0.46	60.11	37.06	47.23	96.49	55.28	68.76
USDSGD	26.46	26.37	5.94	0.45	60.26	30.96	32.57	91.02	55.17	81.64
USDHKD	25.59	25.65	6.26	0.40	43.49	30.63	40.58	83.17	40.38	90.79
USDZAR	20.94	20.91	4.52	0.43	57.67	28.96	26.68	90.50	57.49	82.79
USDKRW	17.69	17.63	3.16	0.48	66.59	66.71	98.24	84.66	60.50	85.95
USDDKK	14.50	14.42	3.69	0.31	28.84	20.82	16.97	39.71	23.92	72.81
USDILS	4.31	4.24	1.41	0.34	43.56	30.89	38.57	81.83	49.65	71.56
EUR Pairs										
EURGBP	35.21	34.55	7.79	0.38	41.20	15.97	11.79	94.16	37.42	52.94
EURJPY	23.74	21.70	10.05	0.73	39.49	23.76	21.22	96.83	36.49	51.73
EURCHF	22.56	21.46	7.85	0.60	47.16	28.74	23.02	93.61	45.85	70.65
EURSEK	9.87	9.48	2.89	0.34	47.86	15.34	28.29	93.23	41.52	42.58
EURNOK	8.11	7.70	2.65	0.40	48.66	21.56	20.04	93.89	42.33	46.33
EURAUD	5.87	5.57	2.05	0.30	39.89	10.73	7.53	91.85	47.49	40.04
EURDKK	5.40	5.01	2.49	0.44	29.71	19.58	37.17	76.28	26.50	60.49
EURCAD	3.42	3.23	1.11	0.22	33.16	9.54	7.74	85.82	36.13	41.08
GBP Pairs										
GBPJPY	6.05	5.57	2.61	0.61	34.21	16.12	11.76	90.36	45.89	48.93
GBPCHF	2.13	1.99	0.89	0.26	40.26	10.78	10.38	71.89	40.63	42.78
GBPAUD	1.97	1.83	0.76	0.39	22.17	15.01	3.62	81.02	48.61	51.11
GBPCAD	1.38	1.28	0.65	0.25	24.52	9.33	5.90	69.17	30.79	40.04
Other Pairs										
AUDJPY	5.45	5.10	1.94	0.57	33.61	11.60	13.91	93.12	46.43	45.32
AUDNZD	2.98	2.71	1.26	0.29	40.36	26.00	13.37	92.58	68.54	23.64
CADJPY	0.84	0.77	0.38	0.35	22.28	7.42	6.66	76.53	43.37	37.62
Average					42.05	22.51	22.30	84.29	45.10	58.97

This table presents summary statistics of daily trading activity for each currency pair in our sample. The first three columns display the sample mean ($Mean$), median ($Median$) and standard deviation ($StDev$) of volume, in USD billions. The fourth column reports the first-order autocorrelation coefficient ($AC(1)$). The fifth, sixth and seventh columns report the pairwise correlation (in percent) between spot and outright forward volume; swap and outright forward volume; and spot and swap volume, respectively. The last three columns display the correlation between the number of trades and dollar volume for each instrument: Spot, Outright Forward and Swap. The final row displays the average correlation across all currency pairs.

Table 2: The Value of Volume: Statistical Significance

	All Pairs			Only US	Only EUR	All Pairs		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Total Volume					*Spot*	*Forward*	*Swap*
$Return_t$	-0.003	-0.027*	-0.069**	-0.060**	-0.128***	-0.070**	-0.053*	-0.059*
	(0.89)	(0.09)	(0.05)	(0.03)	(0.00)	(0.02)	(0.09)	(0.08)
$Volume_t$	-0.002	-0.002	-0.002	-0.002	-0.012	0.003	0.000	-0.011
	(0.92)	(0.91)	(0.93)	(0.91)	(0.82)	(0.88)	(0.98)	(0.54)
$Return_t * Volume_t$		0.196***	0.194***	0.187***	0.293***	0.185***	0.073**	0.016
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.02)	(0.53)
Time FE	YES	YES	YES	YES	YES	YES	YES	YES
Pair FE	YES	YES	YES	YES	YES	YES	YES	YES
Controls	NO	NO	YES	YES	YES	YES	YES	YES
R-Square	0.02	0.02	0.02	0.03	0.10	0.02	0.02	0.02
N	38,139	38,139	38,077	19,628	11,070	38,078	38,078	38,078

This table presents coefficient estimates and associated p-values (reported in parentheses) for the following panel regression:

$$rx_{i,t+1} = \alpha_i + \tau_t + \beta_1 rx_{i,t} + \beta_2\Big(r_{i,t} * v_{i,t}\Big) + \beta_3 v_{i,t} + \boldsymbol{\gamma}'\boldsymbol{x}_{i,t} + \boldsymbol{\delta}'\Big(\boldsymbol{x}_{i,t} \odot rx_{i,t}\Big) + \epsilon_{i,t+1},$$

where α_i and τ_t denote currency-pair and time fixed effects, $rx_{i,t}$ is the log currency excess return for currency pair i at time t,$\boldsymbol{x}_{i,t}$ is a vector of controls relative to pair i, $\odot$ is the element-wise product operator used to generate interaction terms with controls, and ϵ_{t+1} is the model error term. Following the literature, for each currency pair i we measure volume at time t as the (log) deviation from its recent trend, define as $v_{i,t} = log(Volume_t) - log\left(\frac{\sum_{s=1}^{21} Volume_{t-s}}{21}\right)$. The values reported in columns (1) to (5) are based on total volume (computed as the sum of spot, outright forward and swap volume) while values reported in columns (6) to (8) are based on spot, outright forward and swap volume respectively. The values reported in columns (1) to (3) and (6) to (8) are based on all 31 currency pairs in our sample. The values reported in columns (4) and (5) are calculated for samples including only USD- and EUR-base pairs, respectively. p-values are based on double-clusters standard errors (see Petersen, 2009). Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 3: The Value of Volume: Economic Significance

	3 × 3 sort			**2 × 4 sort**			**3 × 4 sort**		
				Total Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	19.51***	5.61	13.90***	16.46***	3.78	12.68***	20.22***	5.55	14.67**
SR	1.82	0.46	1.36***	1.89	0.37	1.53***	1.84	0.36	1.47***
Θ (%)	17.79	3.45		15.33	2.21		18.40	2.03	
MDD	9.00	19.04		5.73	18.98		8.66	19.72	
				Spot Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	19.22***	3.62	15.60***	11.95***	2.97	8.99**	17.41***	3.36	14.06**
SR	1.78	0.29	1.48***	1.41	0.28	1.13**	1.60	0.21	1.38***
Θ (%)	17.47	1.38		10.87	1.24		15.64	-0.27	
MDD	6.80	18.74		5.65	20.62		8.34	26.91	
				Forward Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	18.70***	4.66	14.05***	13.49***	6.17	7.31*	16.35***	1.13	15.22***
SR	1.70	0.42	1.29***	1.60	0.67	0.93*	1.44	0.08	1.35***
Θ (%)	16.91	2.77		12.42	4.89		14.43	-1.67	
MDD									
				Swap Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	14.59***	4.06	10.53**	13.01***	2.46	10.55**	12.71**	-0.35	13.05**
SR	1.32	0.37	0.95**	1.41	0.27	1.14**	1.11	-0.03	1.14**
Θ (%)	12.76	2.27		11.75	1.23		10.76	-3.19	
MDD	16.50	15.53		7.18	17.32		18.68	30.82	

This table presents the out-of-sample economic performance of the Reversal-Low (*RevL*) and Reversal-High (*RevH*) strategies before transaction costs. Each day we perform a conditional double sort by first sorting currency pairs by the previous day's return and then by the previous day's volume. The *RevL* strategy takes positions in currency pairs with low prior volume; long currencies which previously depreciated and short currencies which previously appreciated. The *RevH* strategy is the analogous strategy that takes positions in currency pairs with high prior volume. Results are reported separately for total volume as well as spot, forward and swap instruments. The left-most column is based on a 3x3 conditional double sort, the middle column is based on a 2x4 conditional double sort, while the right-most column is based on a 3x4 conditional double sort. We report the annualized average return (*mean*), annualized Sharpe ratio (SR), the Θ performance measure of Ingersoll et al. (2007) and the maximum drawdown (MDD). The first (second) value in the Δ-column denotes the difference between the annualized average return (Sharpe Ratio) of the *RevL* and *RevH* strategies. We test whether the individual annualized average returns (and their difference) are statistically different from zero with Newey-West (1987) adjusted t-statistics. We test whether the two Sharpe ratios are statistically different from each other using the procedure proposed by Ledoit and Wolf (2008). Values marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 4: Understanding *RevL*: Can Transaction Costs Explain the Returns?

	Dukascopy			**Olsen**			**WM/R 25%**		
				Total Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	13.07***	-1.30	14.36***	11.90**	-2.05	14.11***	11.59**	-0.72	12.31**
SR	1.14	-0.10	1.24***	1.04	-0.16	1.20***	1.01	-0.06	1.07**
Θ (%)	11.10	-3.70		9.94	-4.41		9.64	-3.10	
MDD	14.39	34.12		17.27	37.99		16.18	32.49	
				Spot Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	11.08**	-1.59	12.68**	9.09*	-2.24	11.55**	10.22**	-1.55	11.77**
SR	0.98	-0.12	1.10**	0.80	-0.18	0.97**	0.90	-0.12	1.02**
Θ (%)	9.16	-4.03		7.16	-4.64		8.30	-3.98	
MDD	14.41	32.61		18.82	36.82		15.38	32.83	
				Forward Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	11.14**	-1.67	12.81***	10.04**	-2.69	13.06***	10.40**	-1.40	11.79***
SR	0.97	-0.14	1.11***	0.88	-0.22	1.10***	0.90	-0.12	1.02**
Θ (%)	9.17	-3.87		8.08	-4.88		8.42	-3.59	
MDD	11.45	34.89		14.53	39.81		12.63	33.78	
				Swap Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	5.52	0.99	4.53	4.45	-0.74	5.41	4.26	0.54	3.73
SR	0.47	0.09	0.39	0.38	-0.06	0.45	0.37	0.05	0.32
Θ (%)	3.49	-1.03		2.41	-2.69		2.23	-1.49	
MDD	28.05	22.37		33.37	29.97		31.32	22.36	

This table presents the out-of-sample economic performance of the Reversal-Low (*RevL*) and Reversal-High (*RevH*) strategies after transaction costs. Each day we perform a conditional double sort by first sorting currency pairs by the previous day's return and then by the previous day's volume. The *RevL* strategy takes positions in currency pairs with low prior volume; long currencies which previously depreciated and short currencies which previously appreciated. The *RevH* strategy is the analogous strategy that takes positions in currency pairs with high prior volume. Results are reported separately for total volume as well as spot, forward and swap instruments. The left-most column is based on a bid-ask spreads from Dukascopy, the middle column is based on bid-ask spreads from Olsen, while the right-most column is based on bid-ask spread from WM/R. All results are based on a 3x3 conditional double sort. We report the annualized average return (*mean*), annualized Sharpe ratio (*SR*), the Θ performance measure of Ingersoll et al. (2007) and the maximum drawdown (*MDD*). The first (second) value in the Δ-column denotes the difference between the annualized average return (Sharpe Ratio) of the *RevL* and *RevH* strategies. We test whether the individual annualized average returns (and their difference) are statistically different from zero with Newey-West (1987) adjusted t-statistics. We test whether the two Sharpe ratios are statistically different from each other using the procedure proposed by Ledoit and Wolf (2008). Values marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 5: Understanding *RevL*: Can Currency Factors Explain the Returns

	Excluding Bid-Ask Spreads				Including Bid-Ask Spreads			
	1	2	3	4	5	6	7	8
α	19.27***	19.01***	19.26***	11.54**	12.88**	12.80***	12.89**	6.98
DOL	-0.09	-0.03	-0.07	-0.06	-0.09	-0.01	-0.05	-0.03
CAR	0.10**				0.14***			
VOL		4.38**				2.66		
MOM			-0.11**				-0.12**	
ILL				-0.11**				-0.09**
Adj R^2	0.01	0.01	0.01	0.02	0.01	0.00	0.01	0.01

The table presents the results from ordinary least-square time-series regressions, in which the daily returns to the *RevL* strategy are regressed on currency factors. In the first four columns, the returns to the *RevL* strategy are gross of costs, whereas in columns five to eight, the returns are net of bid-ask spreads. In each regression, we include a constant, the 'dollar' (DOL) risk factor and a second currency factor from the list of carry (CAR); currency volatility (VOL); 3-month currency momentum (MOM); and foreign exchange illiquidity (ILL). In each case we report the time-series coefficients on the variables and constant (α). We calculate standard errors according to Newey-West (1987). We report regression adjusted R^2 statistics in the final row. Values marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 6: The Value of Volume: Diversification and Investment Performance.

Panel A: Performance and correlation with $RevL$

	CAR	DOL	MOM	VAL
μ (%)	3.58	-3.48	-1.74	1.14
SR	0.50	-0.54	-0.30	0.23
ρ (Total)	-0.16	0.11	0.14	-0.03

Panel B: Portfolio Results

	Without $RevL$			With $RevL$		
	EW	MV	TG	EW	MV	TG
$mean$ (%)	-0.12	0.53	2.60	2.15	1.46	5.96
SR	-0.03	0.17	0.55	0.58	0.48	1.19
ω_{RevL} (%)				20.00	7.22	38.05

This table presents summary statistics on the average returns and Sharpe ratio of currency strategies and their correlation ($corr$) with the Reversal Low ($RevL$) strategy (Panel A). The strategies we consider are carry (CAR), dollar (DOL), momentum (MOM) and value (VAL). In Panel B, we report the investment performance of broad currency portfolios that include the aforementioned strategies excluding and including the $RevL$ strategy. The portfolios we form place different weights on the strategies and include an equal weighting scheme (EW) plus an estimation of optimal weights when minimizing the portfolio variance (MV) and maximizing the Sharpe ratio at the tangency point (TG) on the efficient frontier. In the final row, we report the average weight allocated to the $RevL$ strategy (ω_{RevL}) within each portfolio strategy.

Table 7: The Value of Volume: Currency Hedge Funds

	(1) All Funds	(2) Technical Funds	(3) Fundamental Funds
$RevL$	-0.053	-0.107*	0.011
	(0.14)	(0.08)	(0.53)
$TFRX$	0.036***	0.046**	0.021***
	(0.00)	(0.01)	(0.00)
CAR	0.070	0.046	0.111***
	(0.15)	(0.60)	(0.00)
DOL	-0.217**	0.106	-0.613***
	(0.03)	(0.47)	(0.00)
MOM	0.042	0.171	-0.109
	(0.70)	(0.31)	(0.15)
VAL	0.012	0.002	0.004
	(0.92)	(0.99)	(0.96)
$ILL(ba)$	-0.014	0.014	-0.049
	(0.80)	(0.89)	(0.18)
$ILL(cs)$	0.017	0.029	0.008
	(0.52)	(0.46)	(0.65)
R^2	0.09	0.09	0.15
N	3,087	1,686	1,401

This table presents coefficient estimates and the relative p-values (in round brackets) from the following panel regression:

$$r_{i,t} = \alpha_i + \tau_t + \beta RevL_t + \boldsymbol{\gamma}' \boldsymbol{x}_t + \epsilon_{i,t+1}$$

where α_i and τ_t denote fund and time fixed effects, $r_{i,t}$ is the return of the Currency Hedge Fund i in month t, $RevL_t$ is the return of the $RevL$ strategy, and the vector $\boldsymbol{x}_t$ contains the following currency factors and strategies: currency trend factor ($TFRX$), carry factor (CAR), dollar factor (DOL), momentum (MOM), currency value (VAL) and the illiquidity factor, calculated using the bid-ask spread ($ILL(ba)$) and Corwin-Schultz technique ($ILL(ca)$) (see Karnaukh, Ranaldo and Sodelind (2015) for details). The results in column 1 are based on all currency hedge funds reporting to BarclayHedge at December 2016, which are then further split between currency funds that trade according to technical indicators (column 2) and currency funds that trade according to fundamental signals (column 3). p-values are based on double-cluster standard errors (see Petersen, 2009). Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Figure 1: A Day in the Life of Foreign Exchange Volume

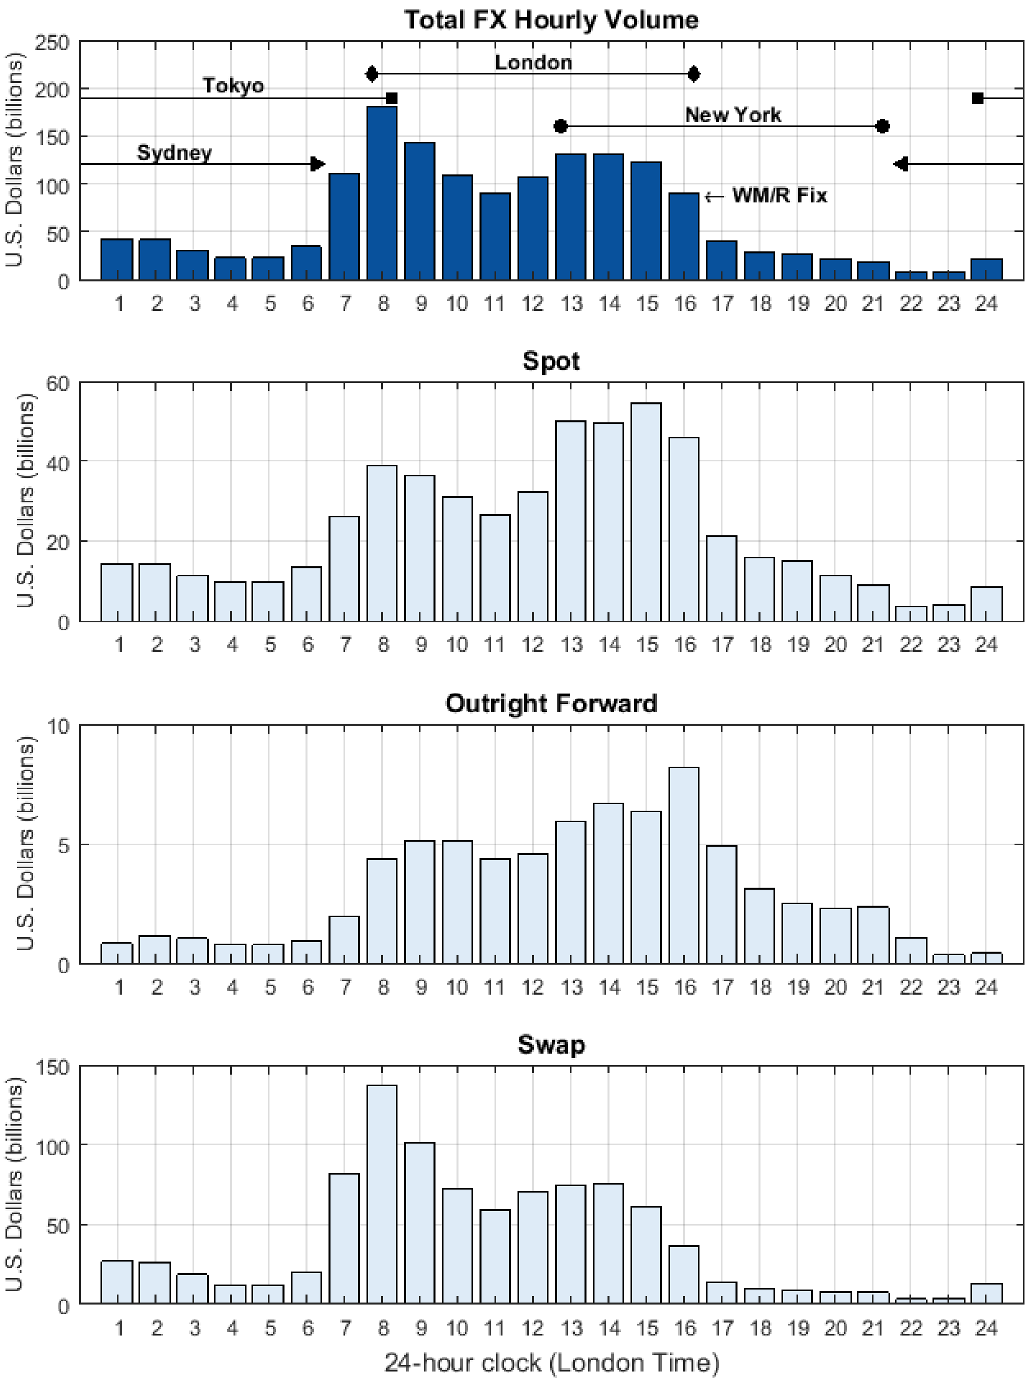

This figure displays the average hourly volume (in U.S. dollars) across the entire day (London time). The average is computed across all trading days in our sample, from November 2011 to December 2016. The first plot reflects total volume, computed as the sum of spot, outright forward and swap volume; the second, third and fourth plots reflect spot, outright forward and swap volume respectively. In all plots volume is aggregated across 31 currency pairs, the full list of pairs is reported in the caption to Figure 2. The number underneath each bar denotes the closing time, e.g. the bar denoted 16 refers to volume between 3pm and 4pm (London time). The solid horizontal lines in the first plot indicate market trading hours in London (from 7am to 4pm), New York (from noon to 9pm), Sydney (from 9 p.m to 6am) and Tokyo (from 11pm to 8am).

Figure 2: Cumulative Returns to the *RevL* Strategy Across Alternative Conditional Double Sorts

This figure displays out-of-sample daily cumulative returns. The upper plots are based on total volume, computed as the sum of spot, outright forward and swap market volume, and include returns for the Reversal-Low (*RevL*) (blue solid line) and Reversal-High (*RevH*) (red dotted line) strategies. Each day we perform a conditional double sort by first sorting currency pairs by the previous day's return and then by the previous day's volume. The *RevL* strategy takes positions in currency pairs with low prior volume; long currencies which previously depreciated and short currencies which previously appreciated. The *RevH* strategy is the analogous strategy that takes positions in currency pairs with high prior volume. The lower plots displays disaggregated results for the *RevL* strategy across each FX instrument: spot (blue solid line), outright forward (blue dashed line) and swap (blue dotted line). The plots in the first, second and third columns are based on 3x3, 2x4 and 3x4 conditional double sorts.

Figure 3: Comparing Bid-Ask Spreads: How Appropriate is the 50% WM/R Scaling Factor?

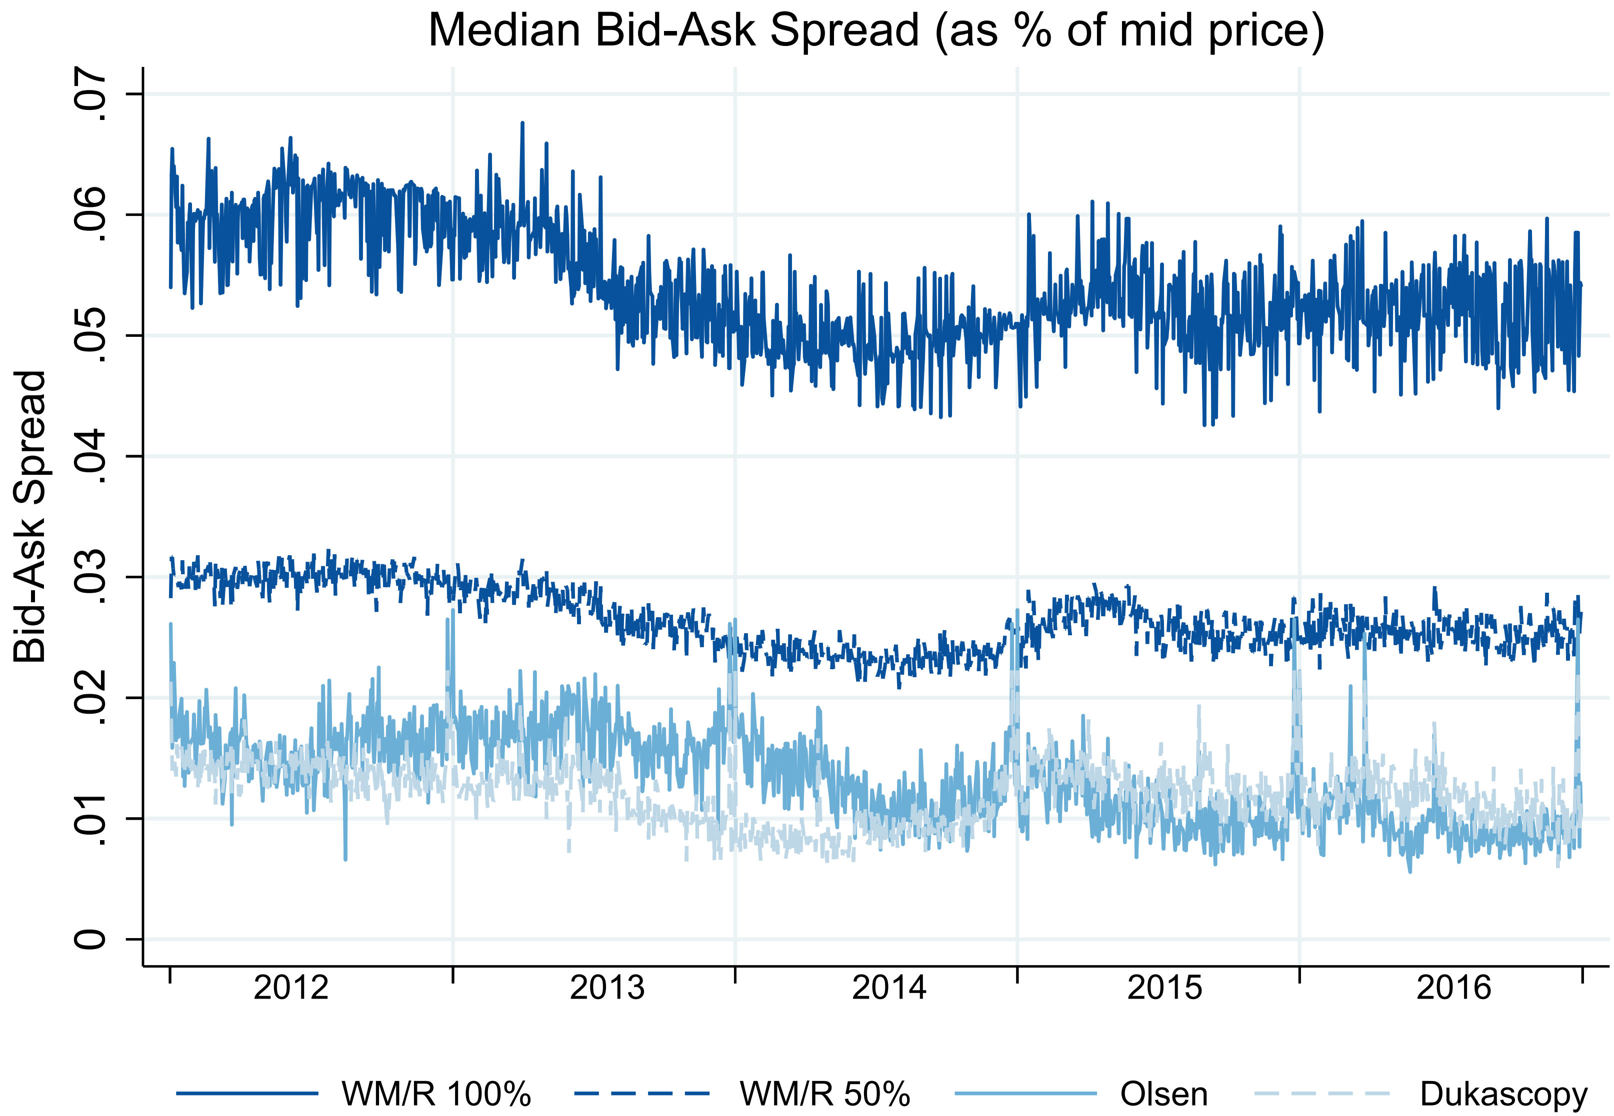

This figure presents the time-series of median bid-ask spreads (as a percentage of the mid price) across all currency pairs in our sample using three different data sources: WM/Reuters (WM/R), Olsen Financial Technologies (Olsen) and Dukascopy Bank (Dukascopy). All series are based on prices recorded at 4pm (London Time).

Figure 4: Is the *RevL* Strategy Profitable Independent of the Formation Hour?

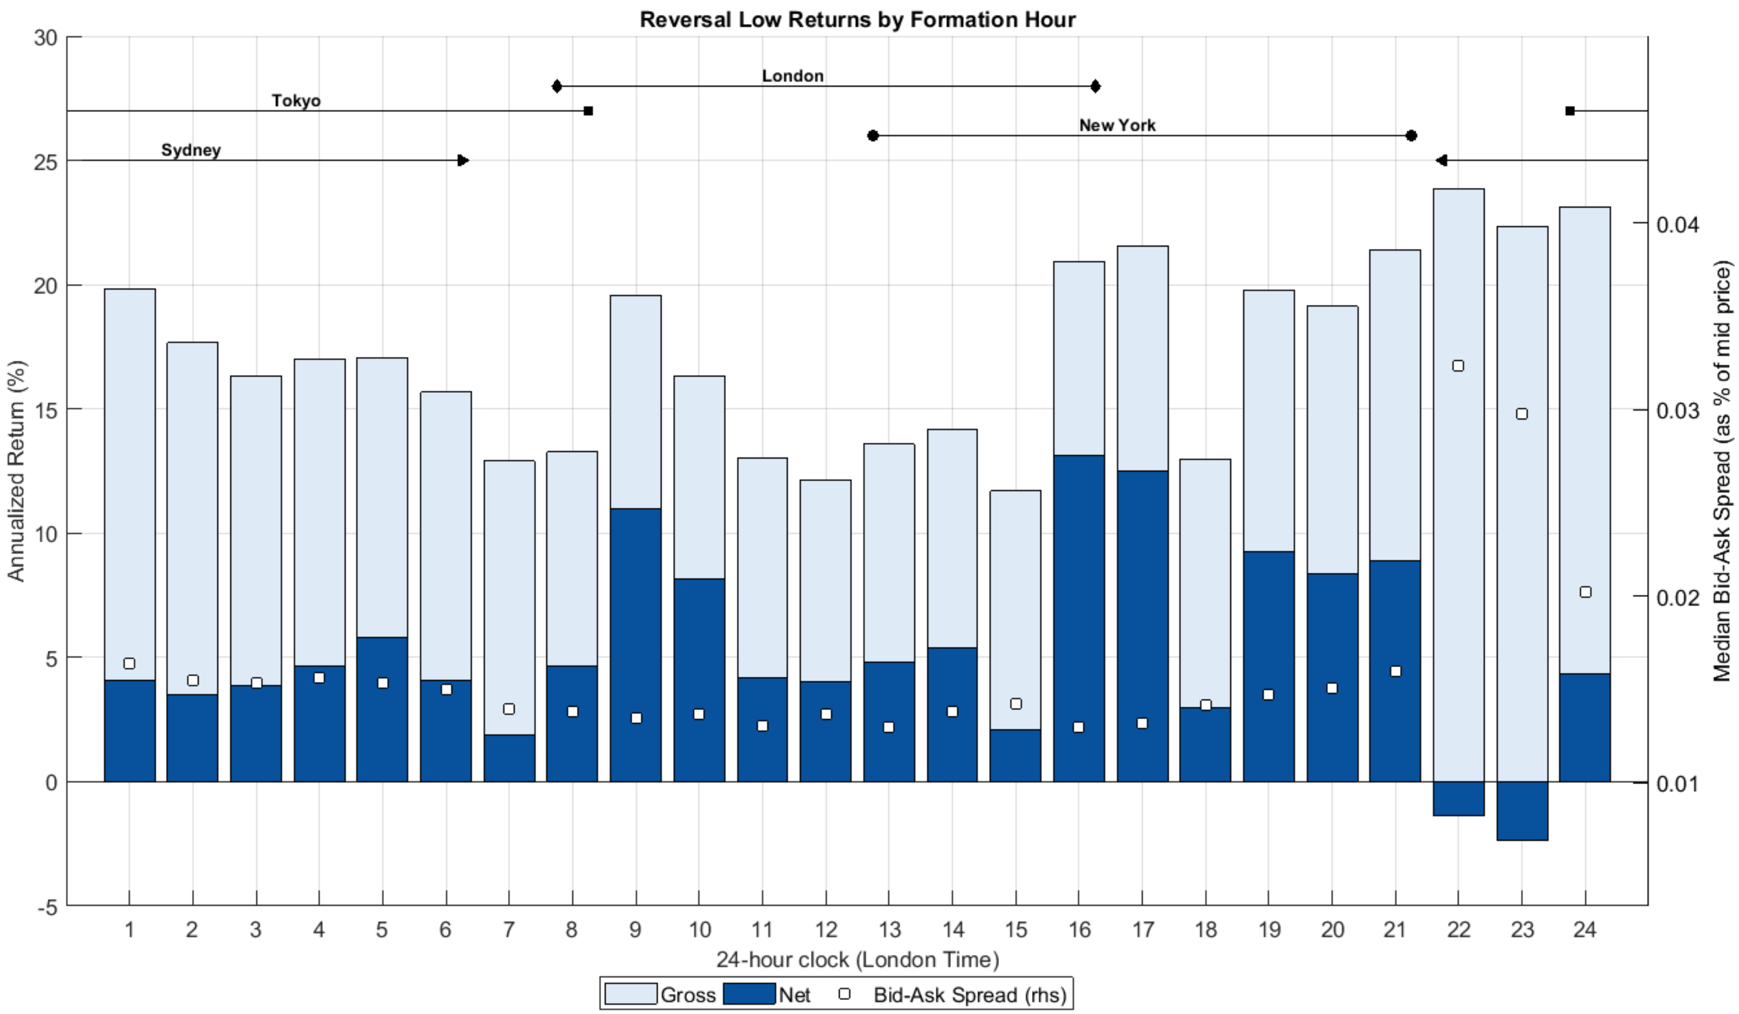

This figure presents the gross and net returns to the *RevL* strategy when implemented at different times of the trading day (London time). The height of each bar reflects the average daily gross return to the strategy, while the RED bars reflect the net return to the strategy after incorporating bid-ask spreads (left-hand-size axis). The white squares reflect the median bid-ask spread (as a percentage of the mid price) at each hour of the day (right-hand-side axis).

Figure 5: How Long does the Economic Information in the *RevL* Trading Signal Survive?

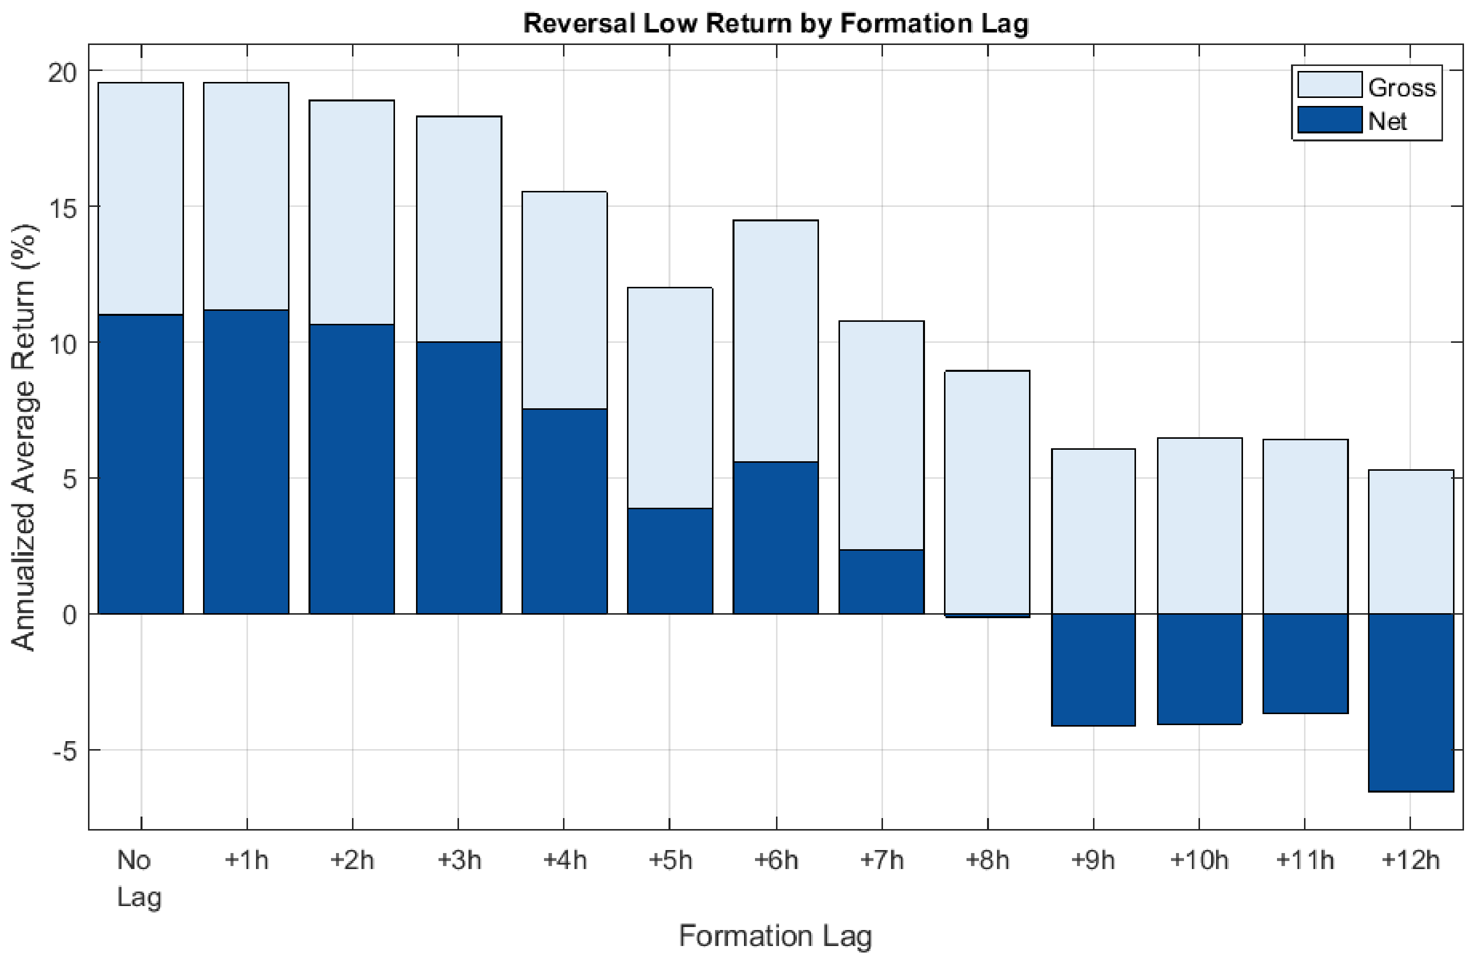

This figure presents the gross and net returns to the *RevL* strategy when implemented with a formation lag. The main trading signal is generated at 9am (London time) each day. The strategy is then implemented using the signal but with a lag of between one hour (+1h) and 12 hours (+12h). The RED bars represent the returns to the strategy after incorporating bid-ask spreads.

Internet Appendix to "The Value of Volume in Foreign Exchange"

Table A.1: Comparison Between CLS and BIS Volume Data

	April 2013			April 2016		
Pair	CLS (Billions)	CLS (%)	BIS(%)	CLS (Billions)	CLS (%)	BIS(%)
EURUSD	490.45	29.89	28.57	455.84	30.61	28.38
USDJPY	290.17	17.69	21.67	285.94	19.20	21.82
GBPUSD	188.69	11.50	10.46	165.08	11.08	11.38
AUDUSD	151.62	9.24	8.05	100.11	6.72	6.35
USDCHF	87.50	5.33	4.07	79.56	5.34	4.36
USDCAD	76.16	4.64	4.42	81.53	5.47	5.28
EURJPY	40.85	2.49	3.27	16.11	1.08	1.91
USDMXN	34.64	2.11	2.83	25.68	1.72	2.18
EURGBP	34.59	2.11	2.26	30.45	2.04	2.42
NZDUSD	34.02	2.07	1.81	31.28	2.10	1.89
USDSEK	27.00	1.65	1.22	32.36	2.17	1.60
USDSGD	26.20	1.60	1.44	30.00	2.01	1.96
USDNOK	25.37	1.55	1.08	19.78	1.33	1.16
EURCHF	24.14	1.47	1.57	15.43	1.04	1.07
USDZAR	23.19	1.41	1.13	17.59	1.18	0.97
USDHKD	22.46	1.37	1.53	26.19	1.76	1.86
USDKRW	18.18	1.11	1.33	15.79	1.06	1.89
EURSEK	9.29	0.57	0.62	10.76	0.72	0.87
AUDJPY	8.56	0.52	1.02	5.67	0.38	0.75
EURNOK	8.30	0.51	0.44	9.51	0.64	0.68
EURAUD	5.99	0.37	0.46	4.58	0.31	0.39
EURDKK	4.44	0.27	0.29	3.87	0.26	0.31
EURCAD	3.19	0.19	0.33	3.59	0.24	0.34
CADJPY	0.79	0.05	0.13	1.12	0.08	0.17

This table presents summary statistics for the CLS volume data in April of 2013 and 2016 for the currency pairs in the sample. For each year, we report the average daily volume settled by CLS for each currency pair (column 1), the volume as a percentage of the total volume (column 2) and the equivalent percentage share reported by the Bank for International Settlements (BIS) in their 2013 and 2016 Triennial Surveys of Central Banks.

Table A.2: Descriptive Statistics and Composition of Double-Sorted Return-Volume Portfolios

PANEL A: Descriptive Statistics

	Low Returns			*Mid Returns*			*High Returns*		
	Low Vol	*Mid Vol*	*High Vol*	*Low Vol*	*Mid Vol*	*High Vol*	*Low Vol*	*Mid Vol*	*High Vol*
	P1	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9**
mean (%)	11.97***	5.22*	3.09	-1.87	2.00	-0.25	-7.54**	-2.14	-2.52
Sharpe	1.71	0.77	0.44	-0.32	0.39	-0.06	-1.20	-0.32	-0.33
std (%)	7.00	6.76	7.02	5.81	5.19	4.44	6.26	6.76	7.67
skew	0.44	0.31	-0.49	-3.14	-0.94	-0.47	-0.73	-0.89	-3.19
kurt	6.7	6.7	15.9	49.2	20.8	7.5	9.4	11.7	52.4
ac(1)	-0.02	-0.03	-0.05	0.01	-0.02	-0.01	-0.01	-0.05	-0.04
mdd (%)	5.7	10.4	11.2	20.2	7.5	10.5	41.8	14.3	26.0
turnover (%)	87.9	88.7	81.5	82.5	86.7	77.7	84.7	83.7	84.7

PANEL B: Main Currencies Entering each Portfolio (and percentage of sample in the portfolio)

	P1	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9**
Highest	CADJPY	GBPUSD	USDNOK	EURDKK	USDHKD	USDHKD	AUDJPY	AUDUSD	CADJPY
	(14%)	(17%)	(16%)	(30%)	(25%)	(44%)	(19%)	(23%)	(15%)
2nd Highest	EURAUD	EURUSD	CADJPY	EURCHF	EURCHF	EURDKK	CADJPY	USDJPY	USDILS
	(14%)	(17%)	(16%)	(18%)	(20%)	(39%)	(18%)	(21%)	(14%)
3rd Highest	GBPAUD	AUDUSD	NZDUSD	USDHKD	USDSGD	EURCHF	GBPJPY	EURUSD	AUDNZD
	(14%)	(17%)	(16%)	(15%)	(16%)	(23%)	(17%)	(20%)	(14%)

This table presents descriptive statistics for currency portfolios sorted by past returns and volume (Panel A) and the main currencies entering each portfolio (Panel B). Portfolios are rebalanced daily. We report summary statistics for the annualized mean, with the superscripts *, **, *** represent significance of the portfolio return at the 10%, 5% and 1% confidence levels using Newey-West (1987) corrected standard errors. In addition we report the Sharpe ratio (*Sharpe*), standard deviation (*std*), maximum drawdown (*mdd*), skewness (*skew*), kurtosis (*kurt*), first-order autocorrelation coefficient (*ac(1)*). We also report the maximum drawdown (*mdd*) and average turnover (*t/o*) for each portfolio. In Panel B we report the main currencies entering the portfolios. The percentage of days in the sample in which a currency enters a portfolio is reported in parentheses.

Table A.3:

The Impact on *RevL* Returns from Scaling WM/R Bid-Ask Spreads

	WM/R 20%			WM/R 30%			WM/R 50%		
				Total Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	13.34***	0.57	12.78**	10.31**	-2.25	12.55**	5.09	-7.50	12.59**
SR	1.17	0.04	1.12**	0.90	-0.18	1.08**	0.44	-0.59	1.03**
Θ (%)	11.39	-1.82		8.35	-4.63		3.10	-9.89	
MDD	14.24	29.57		17.59	35.94		24.21	50.25	
				Spot Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	11.84**	-0.28	12.11**	8.87*	-2.66	11.53**	3.45	-7.80	11.24**
SR	1.04	-0.02	1.07**	0.78	-0.21	0.99**	0.30	-0.61	0.91**
Θ (%)	9.92	-2.71		6.95	-5.10		1.51	-10.24	
MDD	13.97	30.14		16.85	35.20		26.40	48.13	
				Forward Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	11.99**	0.22	11.77***	9.21*	-2.85	12.06***	3.48	-8.60*	12.07***
SR	1.04	0.02	1.02***	0.80	-0.23	1.04**	0.30	-0.71	1.01**
Θ (%)	10.02	-1.99		7.24	-5.05		1.50	-10.81	
MDD	11.46	30.93		14.05	36.76		20.75	51.70	
				Swap Volume					
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	5.60	2.30	3.30	2.52	-0.94	3.46	-3.07	-6.48	3.41
SR	0.48	0.20	0.28	0.22	-0.08	0.30	-0.26	-0.56	0.30
Θ (%)	3.57	0.28		0.49	-2.95		-5.13	-8.50	
MDD	28.34	19.33		34.03	25.60		45.08	40.75	

This table presents the out-of-sample economic performance of the Reversal-Low (*RevL*) and Reversal-High (*RevH*) strategies after transaction costs. Each day we perform a conditional double sort by first sorting currency pairs by the previous day's return and then by the previous day's volume. The *RevL* strategy takes positions in currency pairs with low prior volume; long currencies which previously depreciated and short currencies which previously appreciated. The *RevH* strategy is the analogous strategy that takes positions in currency pairs with high prior volume. Results are reported separately for total volume as well as spot, forward and swap instruments. The left-most column is based on a bid-ask spreads from Dukascopy, the middle column is based on bid-ask spreads from Olsen, while the right-most column is based on bid-ask spread from WM/R. All results are based on a 3x3 conditional double sort. We report the annualized average return (*mean*), annualized Sharpe ratio (*SR*), the Θ performance measure of Ingersoll et al. (2007) and the maximum drawdown (*MDD*). The first (second) value in the Δ-column denotes the difference between the annualized average return (Sharpe Ratio) of the *RevL* and *RevH* strategies. We test whether the individual annualized average returns (and their difference) are statistically different from zero with Newey-West (1987) adjusted t-statistics. We test whether the two Sharpe ratios are statistically different from each other using the procedure proposed by Ledoit and Wolf (2008). Values marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table A.4: A Comparison of Bid-Ask Spreads Part 1

	Bid-Ask			Ratios		
	WM/R	Olsen	Dukascopy	$\frac{Dukascopy}{WM/R}$	$\frac{Olsen}{WM/R}$	$\frac{Dukas}{Olsen}$
AUDJPY	7.93	1.34	1.24	0.16	0.17	0.93
AUDNZD	10.39	2.19	2.21	0.21	0.21	1.01
AUDUSD	4.73	1.04	1.04	0.22	0.22	1.00
CADJPY	5.98	1.64	1.25	0.21	0.27	0.76
EURAUD	7.07	1.11	1.15	0.16	0.16	1.04
EURCAD	5.50	1.30	1.27	0.23	0.24	0.98
EURCHF	3.33	0.96	0.97	0.29	0.29	1.02
EURDKK	0.54	0.92	0.35	0.65	1.73	0.38
EURGBP	5.43	1.05	1.01	0.19	0.19	0.97
EURJPY	5.37	0.87	0.48	0.09	0.16	0.56
EURNOK	5.04	2.73	2.05	0.41	0.54	0.75
EURSEK	3.41	2.23	1.80	0.53	0.65	0.81
EURUSD	2.38	0.48	0.19	0.08	0.20	0.39
GBPAUD	7.55	1.44	1.53	0.20	0.19	1.06
GBPCAD	6.07	1.52	1.79	0.29	0.25	1.18
GBPCHF	8.90	1.56	1.50	0.17	0.18	0.96
GBPJPY	6.15	1.08	0.88	0.14	0.17	0.81
GBPUSD	3.04	0.65	0.51	0.17	0.21	0.79
NZDUSD	5.92	1.64	1.54	0.26	0.28	0.94
USDCAD	3.04	0.77	0.77	0.25	0.25	1.01
USDCHF	6.07	1.01	1.09	0.18	0.17	1.08
USDDKK	2.94	0.93	0.66	0.23	0.32	0.72
USDHKD	0.64	0.59	0.32	0.50	0.92	0.54
USDJPY	3.00	0.58	0.31	0.10	0.19	0.53
USDMXN	1.88	1.52	1.91	1.02	0.81	1.25
USDNOK	7.38	2.61	2.28	0.31	0.35	0.88
USDSEK	5.86	2.04	1.91	0.33	0.35	0.94
USDSGD	4.94	2.37	1.75	0.35	0.48	0.74
USDZAR	9.41	3.80	3.67	0.39	0.40	0.97
Average				0.23	0.25	0.94

This table compares the bid-ask spread across three data sources: WM/Reuters (WM/R), Olsen Financial Technologies (Olsen) and Dukascopy Bank (Dukascopy). The first three columns report the time-series median bid-ask spread (as percentage of mid-price, i.e. $\frac{p_{bid}-p_{ask}}{0.5*(p_{bid}+p_{ask})}$) scaled by 10,000. These values are based on prices recorded at 4pm (London Time). The fourth column displays the ratio between the median bid-ask spread from Dukascopy and the one from WM/R, the fifth column displays the ratio between the median bid-ask spread from Olsen and the one from WM/R, the sixth column displays the ratio between the median bid-ask spread from Dukascopy and the one from Olsen.

Table A.5: A Comparison of Bid-Ask Spreads Part 2

	\$ Bid-Ask			Ratios	
	Bank	Olsen	Dukascopy	$\frac{Olsen}{Bank}$	$\frac{Dukascopy}{Bank}$
AUDJPY	106.02	107.78	132.22	1.02	1.25
AUDNZD	2.61	1.95	2.47	0.75	0.95
AUDUSD	0.67	0.70	1.09	1.05	1.63
EURCHF	2.21	1.41	1.82	0.64	0.83
EURDKK	5.63	9.77	4.02	1.73	0.71
EURGBP	0.69	0.76	0.99	1.10	1.43
EURJPY	104.45	98.89	80.04	0.95	0.77
EURNOK	32.26	27.01	30.19	0.84	0.94
EURSEK	30.92	22.66	25.33	0.73	0.82
EURUSD	0.55	0.52	0.32	0.95	0.58
GBPJPY	188.21	158.89	182.21	0.84	0.97
GBPUSD	0.99	0.88	0.92	0.89	0.94
NZDUSD	1.29	1.01	1.23	0.78	0.96
USDCAD	1.22	0.94	1.12	0.78	0.92
USDCHF	1.66	1.16	1.52	0.69	0.91
USDDKK	5.20	6.44	5.35	1.24	1.03
USDHKD	1.64	4.90	3.57	2.98	2.17
USDJPY	56.63	52.22	37.77	0.92	0.67
USDMXN	26.69	37.22	39.86	1.39	1.49
USDNOK	42.25	37.04	36.17	0.88	0.86
USDSEK	35.14	27.79	26.76	0.79	0.76
USDSGD	2.40	2.64	2.98	1.10	1.24
USDZAR	36.19	42.44	54.03	1.17	1.49
Average				0.95	0.94

This table compares the bid-ask spread across three data sources: an anonymous dealer (Bank), Olsen Financial Technologies (Olsen) and Dukascopy Bank (Dukascopy). The first three columns report the time-series median of the daily average bid-ask spread (i.e. $p_{bid} - p_{ask}$) – between 9am and 5pm London time– scaled by 10,000. The fourth column displays the ratio between the median bid-ask spread from Olsen and the one from "Bank", the fifth column displays the ratio between the median bid-ask spread from Dukascopy and the one from "Bank". The sample period goes from January 2015 to July 2015.

Table A.6: Alternative Measures of Volume

Panel A: Without Standardization

	21 Days			**63 Days**			**126 Days**			**252 Days**		
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	19.51***	5.61	13.90***	17.34***	4.74	12.60**	17.59***	3.59	14.01***	18.58***	2.68	15.90***
SR	1.82	0.46	1.36***	1.63	0.39	1.24***	1.71	0.28	1.42***	1.77	0.21	1.56***
Θ (%)	17.79	3.45		15.65	2.53		16.00	1.23		16.93	0.27	
MDD	9.00	19.04		9.51	20.25		8.00	23.64		8.18	23.14	

Panel B: With Standardization

	21 Days			**63 Days**			**126 Days**			**252 Days**		
	RevL	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ	*RevL*	*RevH*	Δ
mean (%)	20.29***	4.92	15.37***	18.25***	5.59	12.66**	16.45***	0.51	15.94***	16.82***	-0.08	16.90***
SR	1.85	0.40	1.45***	1.69	0.46	1.24***	1.56	0.04	1.52***	1.59	-0.01	1.60***
Θ (%)	18.49	2.73		16.52	3.39		14.79	-1.91		15.14	-2.52	
MDD	7.99	15.45		9.61	19.81		8.62	28.14		10.19	25.22	

This table presents the out-of-sample economic performance of the Reversal-Low (*RevL*) and Reversal-High (*RevH*) strategies. Each day we perform a conditional double sort by first sorting currency pairs by the previous day's return and then by the previous day's volume. The *RevL* strategy takes positions in currency pairs with low prior volume; long currencies which previously depreciated and short currencies which previously appreciated. The *RevH* strategy is the analogous strategy that takes positions in currency pairs with high prior volume. Results are based on total volume and on a 3x3 conditional double sort. In Panel A, we change our measure of volume reported in Equation (3) by changing the number of days over which we estimated the moving average in daily volume between 21 days and 252 days. In Panel B, we maintain the different moving average windows but also divide the measure of volume reported in Equation (3) by the standard deviation of the measure over the same sample window. We report the annualized average return (*mean*), annualized Sharpe ratio (*SR*), the Θ performance measure of Ingersoll et al. (2007) and the maximum drawdown (*MDD*). The first (second) value in the Δ-column denotes the difference between the annualized average return (Sharpe Ratio) of the *RevL* and *RevH* strategies. We test whether the individual annualized average returns (and their difference) are statistically different from zero with Newey-West (1987) adjusted t-statistics. We test whether the two Sharpe ratios are statistically different from each other using the procedure proposed by Ledoit and Wolf (2008). Values marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table A.7: Currency Pair Subsamples

	USD Base Pairs (16)			EUR and GBP Base Pairs (13)		
	RevL	**RevH**	**Δ**	**RevL**	**RevH**	**Δ**
mean (%)	15.39***	2.43	12.96**	14.83***	8.33	6.50
SR	1.24	0.17	1.07**	1.35	0.63	0.73
Θ (%)	13.08	-0.66		13.03	5.76	
MDD	13.84	22.27		14.98	18.89	

	Ex EM and Fixed (24)			G10 plus Major Crosses (14)		
	RevL	**RevH**	**Δ**	**RevL**	**RevH**	**Δ**
mean (%)	17.99***	6.41	11.59**	18.63***	8.59	10.04
SR	1.53	0.46	1.07**	1.34	0.57	0.78*
Θ (%)	15.93	3.50		15.77	5.15	
MDD	16.84	25.18		9.85	28.89	

This table presents the out-of-sample economic performance of the Reversal-Low (*RevL*) and Reversal-High (*RevH*) strategies. Each day we perform a conditional double sort by first sorting currency pairs by the previous day's return and then by the previous day's volume. The *RevL* strategy takes positions in currency pairs with low prior volume; long currencies which previously depreciated and short currencies which previously appreciated. The *RevH* strategy is the analogous strategy that takes positions in currency pairs with high prior volume. Results are based on total volume and on a 3x3 conditional double sort. In each sub-table we change the sample of currency pairs. In the top-left sub-table we include all currency pairs that include the U.S. dollar (USD), in the top-right sub-table we include all currency pairs that include the euro (EUR) or British pound (GBP); in the bottom-left sub-table we exclude emerging market and currency pairs that are either fixed or pegged to another currency. Finally, in the bottom-right sub-table we include the G10 currency pairs plus major EUR crosses (EURJPY, EURGBP, EURCHF, EURNOK and EURSEK). We report the number of currency pairs in each sample in parentheses. We report the annualized average return (*mean*), annualized Sharpe ratio (*SR*), the Θ performance measure of Ingersoll et al. (2007) and the maximum drawdown (MDD). The first (second) value in the Δ-column denotes the difference between the annualized average return (Sharpe Ratio) of the *RevL* and *RevH* strategies. We test whether the individual annualized average returns (and their difference) are statistically different from zero with Newey-West (1987) adjusted t-statistics. We test whether the two Sharpe ratios are statistically different from each other using the procedure proposed by Ledoit and Wolf (2008). Values marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table A.8: The Value of Volume: Economic Significance – Dollar Neutral

	Total			Spot		
	RevL	**RevH**	Δ	**RevL**	**RevH**	Δ
mean (%)	15.32***	-0.71	16.02***	17.03***	-1.38	18.41***
SR	1.78	-0.06	1.85***	1.98	-0.12	2.11***
Θ (%)	14.21	-2.50		15.92	-3.20	
MDD	6.94	25.20		7.95	19.18	
	Forward			**Swap**		
mean (%)	14.58***	9.77**	4.81	13.76***	1.65	12.11**
SR	1.65	0.93	0.72	1.51	0.16	1.35**
Θ (%)	13.41	8.13		12.52	-0.02	
MDD	4.97	12.96		8.69	20.21	

This table presents the out-of-sample economic performance of the Reversal-Low (*RevL*) and Reversal-High (*RevH*) strategies before transaction costs for 16 U.S. dollar currency pairs (the U.S. dollar is the quote currency in each case). Each day we perform a 3x3 conditional double sort by first sorting currency pairs by the previous day's return and then by the previous day's volume. The *RevL* strategy takes positions in currency pairs with low prior volume; long currencies which previously depreciated and short currencies which previously appreciated. The *RevH* strategy is the analogous strategy that takes positions in currency pairs with high prior volume. Results are reported separately for total volume as well as spot, forward and swap instruments. We report the annualized average return (*mean*), annualized Sharpe ratio (SR), the Θ performance measure of Ingersoll et al. (2007) and the maximum drawdown (MDD). The first (second) value in the Δ-column denotes the difference between the annualized average return (Sharpe Ratio) of the *RevL* and *RevH* strategies. We test whether the individual annualized average returns (and their difference) are statistically different from zero with Newey-West (1987) adjusted t-statistics. We test whether the two Sharpe ratios are statistically different from each other using the procedure proposed by Ledoit and Wolf (2008). Values marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table A.9: Alternative Investor Perspectives

	EUR			JPY		
	RevL	**RevH**	Δ	**RevL**	**RevH**	Δ
mean (%)	19.54***	5.82	13.72***	19.31***	5.53	13.79***
SR	1.82	0.48	1.34***	1.80	0.46	1.34***
Θ (%)	17.83	3.67		17.60	3.40	
MDD	9.16	19.23		9.38	19.52	
	GBP			**CHF**		
mean (%)	19.56***	5.90	13.66***	19.41***	5.69	13.72***
SR	1.82	0.49	1.34***	1.81	0.47	1.34***
Θ (%)	17.84	3.74		17.69	3.54	
MDD	9.15	19.26		9.10	19.22	

This table presents the out-of-sample economic performance of the Reversal-Low (*RevL*) and Reversal-High (*RevH*) strategies. We take the perspective of an investor in one of four different locations: Euro-area, Japan, the United Kingdom and Switzerland. To do so, we scale returns by the spot exchange rate return between quote currencies and the currency of the home investor. Each day we perform a conditional double sort by first sorting currency pairs by the previous day's return and then by the previous day's volume. The *RevL* strategy takes positions in currency pairs with low prior volume; long currencies which previously depreciated and short currencies which previously appreciated. The *RevH* strategy is the analogous strategy that takes positions in currency pairs with high prior volume. Results are based on total volume and on a 3x3 conditional double sort. We report the annualized average return (*mean*), annualized Sharpe ratio (SR), the Θ performance measure of Ingersoll et al. (2007) and the maximum drawdown (MDD). The first (second) value in the Δ-column denotes the difference between the annualized average return (Sharpe Ratio) of the *RevL* and *RevH* strategies. We test whether the individual annualized average returns (and their difference) are statistically different from zero with Newey-West (1987) adjusted t-statistics. We test whether the two Sharpe ratios are statistically different from each other using the procedure proposed by Ledoit and Wolf (2008). Values marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Figure A.1 Bilateral Regression Interaction Coefficients Across FX Instruments and Currency Pairs

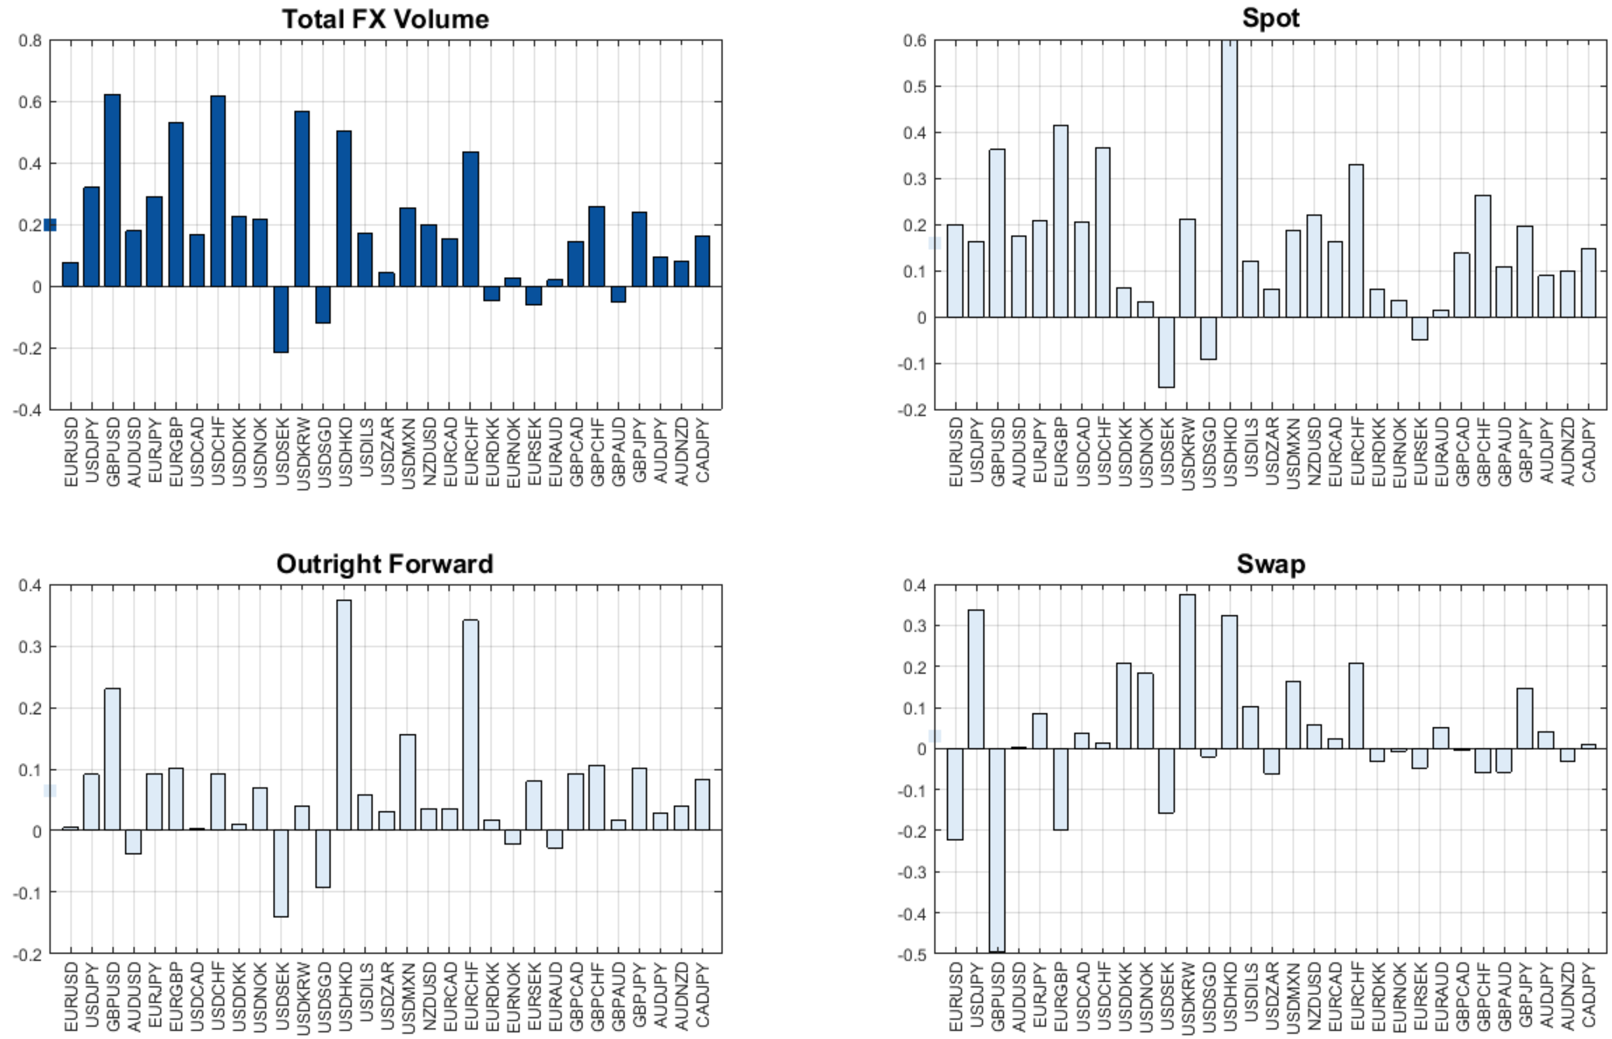

The figure presents the ordinary least square coefficient estimates of the interaction term (β_2) from the following regression:

$$rx_{t+1} = \alpha_i + \beta_1 rx_t + \beta_2 \Big(rx_t * v_t \Big) + \beta_3 v_t + \epsilon_{t+1}, \tag{A-1}$$

where rx is the bilateral currency excess return and v is is volume as defined in Equation (3). The upper left figure presents results for total volume, the remaining figures present results for spot, forward and swap volume.

Figure A.2: Bid-Ask Spreads Across Currency Pairs

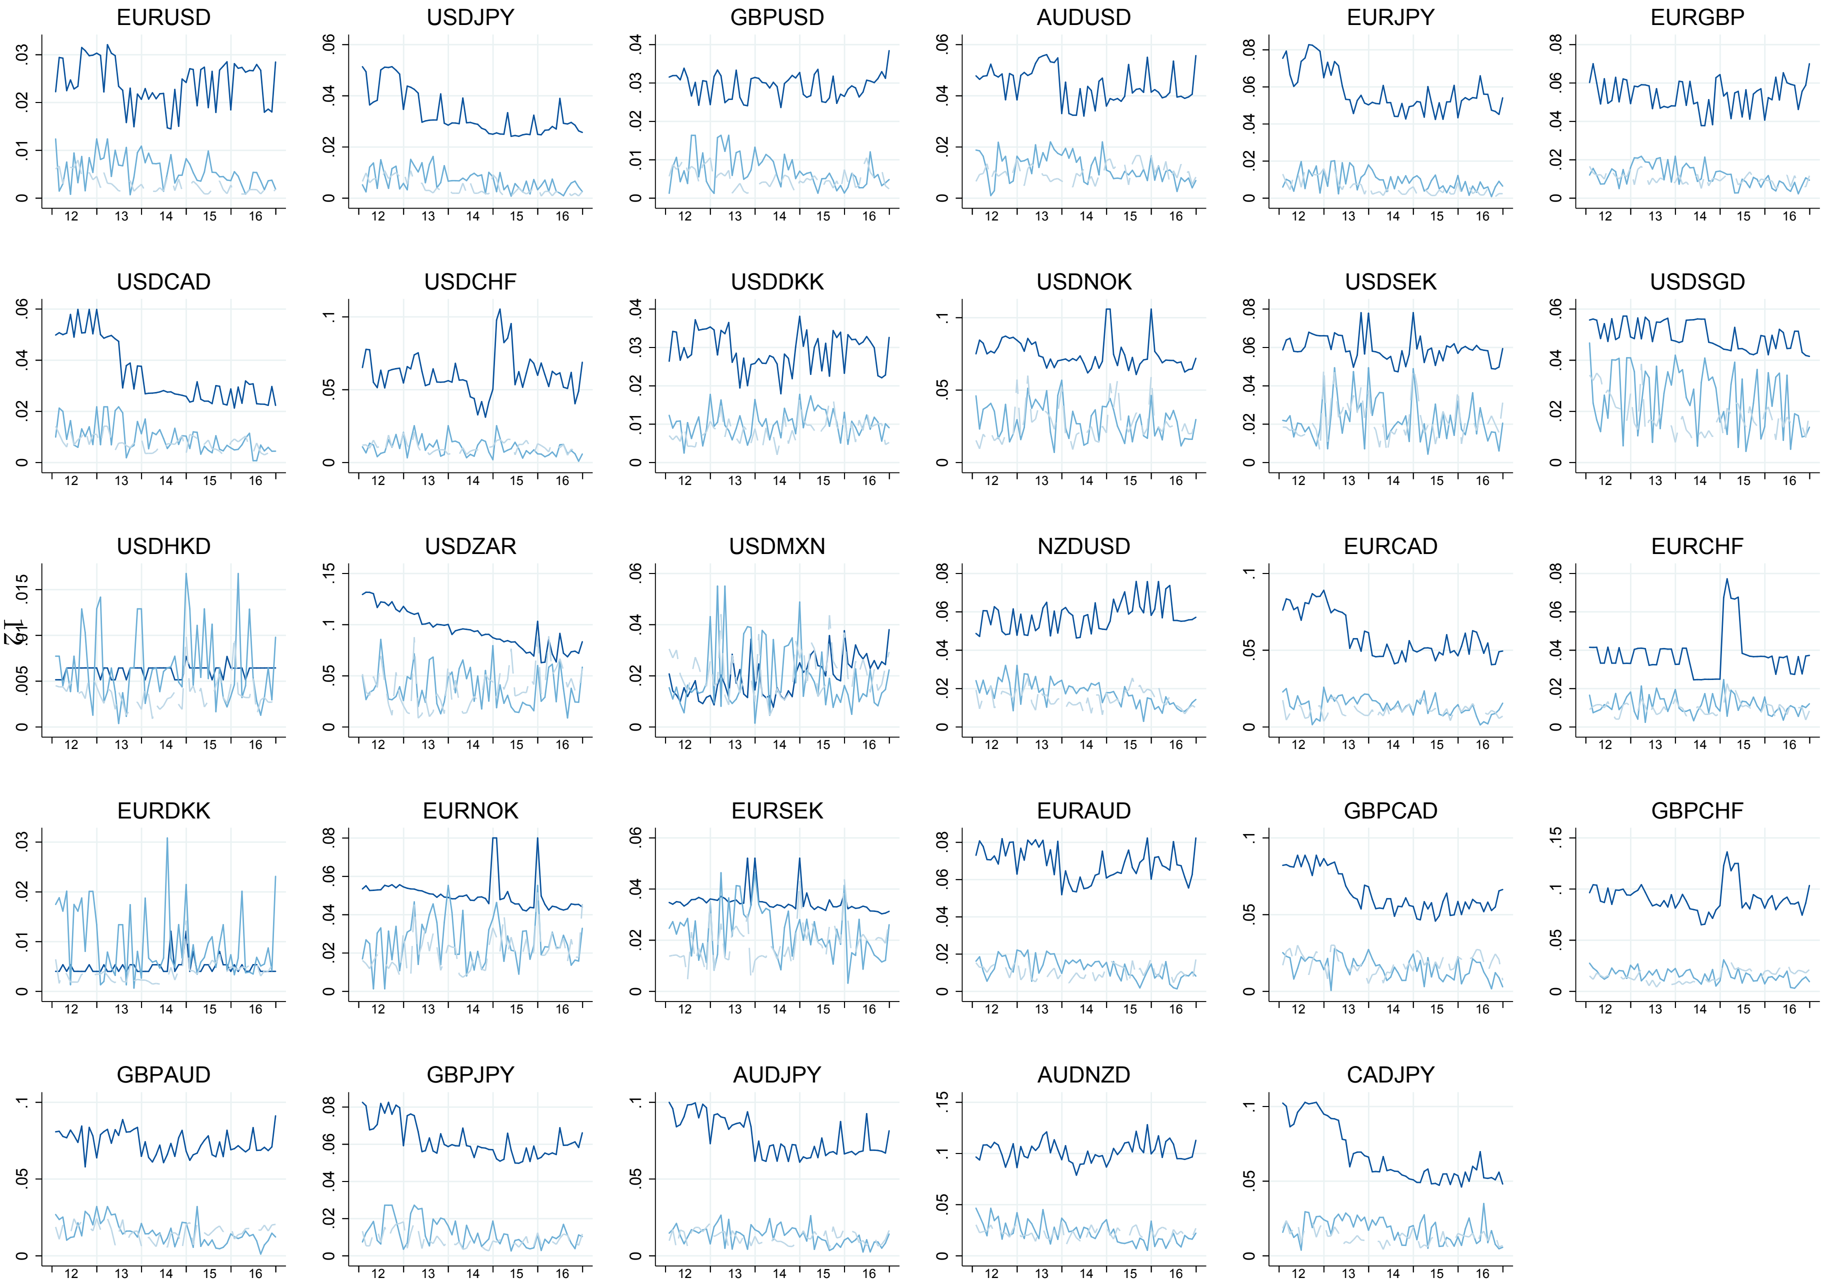

The figure presents the time-series of bid-ask spreads (as a percentage of the mid price) across currency pairs using different data sources including WM/Reuters (WM/R), Olsen Financial Technologies (Olsen) and Dukascopy Bank (Dukascopy).

Figure A.3: Bootstrapped Economic Performance Under the Null of No Cross-Sectional Predictability

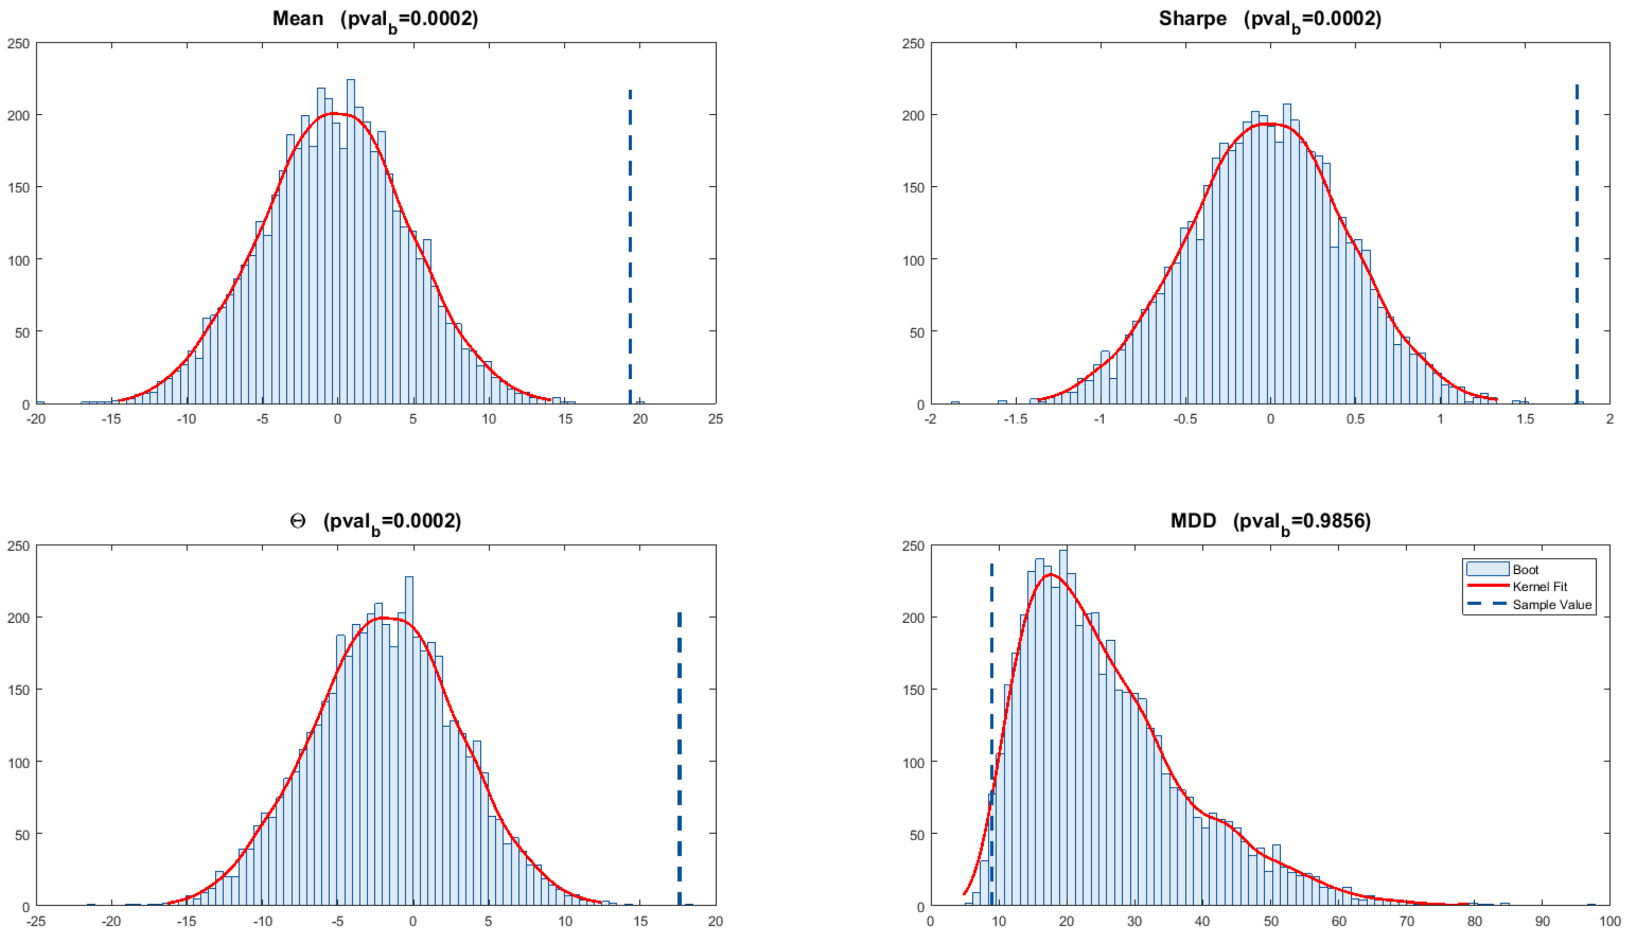

This figure compares the economic performance of the *RevL* strategy computed from 5,000 bootstrapped samples generated under the null of no predictability (light blue histogram), with the one computed from the original sample (vertical dashed line). All results are based on total volume and a 3x3 conditional double-sort. The upper-left plot displays the annualized average return, the upper-right plot displays the annualized Sharpe Ratio, the lower-left plot displays the θ performance measure of Ingersoll et al., and the lower-right plot displays the Maximum Drawdown, the bootstrapped p-values ($pval_b$) are reported in parenthesis in the titles.