Volatility and the Cross-Section of Returns on FX Options

Jonathan Fullwood[a], Jessica James[b,c], Ian W. Marsh[c,*]

[a] *Bank of England, Threadneedle Street, London, EC2R 8AH, UK*

[b] *Commerzbank, 30 Gresham Street, London, EC2V 7PG, UK*

[c] *Business School (formerly Cass), City, University of London, 106 Bunhill Row, London, EC1Y 8TZ, UK*

Abstract

We study the cross-section of returns on FX options sorting currencies based on implied volatilities (IV). Long straddle positions in currencies with low (high) implied volatilities perform well (poorly). A long low IV-short high IV strategy produces large average returns after transactions costs. Total volatility matters rather than any component or transformation of volatility. The returns are distinct from those in the literature on foreign exchange returns or equity option returns and cannot be explained convincingly by standard risk factors. We argue cross-sectional differences in hedging demand combined with limits to arbitrage contribute to misspricing in FX options.

Keywords: Options returns, implied volatility, straddles, foreign exchange

JEL: G130 G150

*Corresponding author

Email addresses: jonathan.fullwood@bankofengland.co.uk (Jonathan Fullwood), jessica.james@commerzbank.com (Jessica James), i.marsh@city.ac.uk (Ian W. Marsh)

1. Introduction

A growing literature shows predictability in the cross-section of expected option returns. These papers almost exclusively focus on equity options. In this paper we examine the performance of buy and hold straddle strategies in the over-the-counter foreign exchange options market. Average daily trading volume in this market during April 2013 was estimated to be \$337bn, while non-index equity options turnover was around \$50bn per day, less than one-sixth of FX options turnover.[1] Despite the size of the market, the predictability of FX options returns has been largely neglected in the literature.

We focus on volatility mispricing and analyse the returns on at-the-money-forward straddle positions formed by simultaneously buying a call and a put on the same asset with the same maturity and strike price. Straddles pay out when the price of the underlying asset moves by a large amount in either direction. Since such straddles have very low deltas they have very little directional exposure to the underlying currencies. When we cross-sectionally sort currencies based on their implied volatilities we find that buying straddles on currencies with relatively low (high) implied volatilities has good (poor) performance. The mean return on a long straddle position in the currency with the lowest implied volatility at the initiation of the trade is 25% per month. The mean return on a long straddle on the currency with the highest implied volatility is -14% per month. The long-short strategy returns 34% per month after accounting for transactions costs, with an annualised Sharpe ratio of 1.03. More diversified long-short straddle, butterfly and delta hedged option positions reveal similar returns patterns but Sharpe ratios as high as 1.38. These returns are comparable in magnitude to those reported for long-short straddles taken in equity options by Goyal and Saretto (2009) and Vasquez (2017). Such eye-catching returns are, of course, affected by the leverage embedded in options and we discuss below alternative interpretations of option returns. One approach is the return on the notional value which is almost 40 basis points per month (or 4.7%pa) for the undiversificed long-short straddle.

In a set of robustness tests we examine the performance of butterfly positions that cap the maximum payoffs earned by a long position in a currency, reducing the influence of periods with exceptionally good performance. Butterfly positions give similar results to straddles, suggesting that our results are not driven by extreme outliers. Chernov et al. (2018) identify days on which either the spot exchange rate and/or the variance of exchange rates jump. Concerned that our results are driven by such events we exclude all months

[1] Notional values of contracts. Figures for foreign exchange options taken from Table 1 of the Bank for International Settlements Triennial Central Bank Survey, Global foreign exchange market turnover 2013, detailed statistical tables https://www.bis.org/publ/rpfxf13fxt.pdf. Figures for equity options taken from Annex 1 of the World Federation of Exchanges Derivatives Market Survey 2013 world-exchanges.org.

containing jumps and show that straddle and butterfly returns marginally increase. Jumps do not appear to be primary drivers of our findings. More generally, we show that positive returns are earned by long-short straddle and butterfly positions quite consistently throughout our sample rather than being specific to a small number of months.

We show that the level of implied volatility matters rather than some key component or transformation of volatility. The volatility risk premium, computed as the difference between historical and implied volatilities, does not predict currency straddle returns, though we confirm that it does have power to predict future exchange rate movements consistent with Della Corte et al. (2016a). We also show that while lagged and contemporaneous changes in implied volatility are related to straddle returns they do not reduce the importance of the level of implied volatility. Taken together, these findings do not support an explanation of return predictability based upon an overreaction and subsequent correction of implied volatilities. Cao and Han (2013) show that delta-hedged equity options returns decrease with the idiosyncratic volatility of the underlying assets, but total volatility has no forecasting power. In our application, we find that idiosyncratic and systematic components of volatility are equally important and that when included alongside total volatility in Fama-MacBeth regressions, idiosyncratic volatility is insignificant.

Our work contributes to a growing literature that looks at the cross-sectional relationship between asset volatilities and options returns.[2] Goyal and Saretto (2009) sort stocks on the difference between realised volatility and implied volatility and show that a large positive difference predicts a positive future return from market-hedged options positions in those stocks. Vasquez (2017) sorts stocks based upon the slope of the volatility term structure. He shows that straddles with high slopes outperform straddles with low slopes and he too fails to identify risk factors capable of explaining his results. Cao and Han (2013) form delta-hedged option positions based on idiosyncratic risk in the underlying equity. They find that high idiosyncratic volatility predicts low future hedged options returns and show that their results are consistent with market imperfections and limits to arbitrage.Cao et al. (2017) widen the search for factors that predict option returns and show that eight out of twelve stock characteristics that predict stock returns also predict delta-hedged options returns. Among these, they demonstrate that delta-hedged options gains increase with idiosyncratic stock volatility. An et al. (2014) provide evidence of bi-directional linkages between option and underlying markets. They show that stocks with large increases in implied volatilities tend to have

[2]A separate but related branch of the literature uses information from the options market to create strategies in the underlying asset. Ang et al. (2006) rank stocks according to total volatility, idiosyncratic volatility and changes in option implied volatilities and show that high ranking stocks provide poor future returns. Bali and Hovakimian (2009) sort stocks on total realised and implied volatilities. Implied volatilities are not related to future excess stock returns but realised volatilities are. The latter finding is consistent with Ang et al. (2006) but the former is not.

high future returns, and that high past returns tend to predict increases in implied volatility. Our work differs from all of these papers by considering the much more active foreign exchange options market. We re-examine the perfomance of several of these predictive variables in this new context with mixed results. Most importantly, however, the predictive power of the level of implied volatility is robust to the inclusion of all of these alternative predictive factors.

While our finding that implied volatility negatively predicts straddle returns is clearly related to the low risk anomaly, our results differ from those in the literature in crucial aspects. As noted, both idiosyncratic and systematic components of implied volatility are important in our application. The inability of idiosyncratic implied volatility to predict returns runs contrary to the arguments of Bali et al. (2017) and Liu et al. (2018) who argue that the low-risk anomaly in stocks is driven by idiosyncratic risk. Similarly, systematic risk alone is not enough to explain our results. A betting-against-(volatility)-beta factor, analogous to Frazzini and Pedersen (2014)'s betting-against-beta factor for equity markets, is priced in the cross-section of straddles but only accounts for a small proportion of returns from the long-short strategy. The betting-against-correlation factor of Asness et al. (2020) explains even less. Further, we show that unlike other examples of the low risk anomaly, the returns on the long leg of our strategy (in low implied volatility currencies) are more volatile than those from the short leg.

How might we explain our findings?

The first explanation we consider is that this is an equilibrium result and is only to be expected. Hu and Jacobs (2019) show that the expected return from holding a call (put) is a decreasing (increasing) function of volatility in both Black-Scholes and stochastic volatility models. Under certain reasonable conditions, the expected return from a long straddle position is also decreasing in volatility. While this argument carries over to foreign exchange options it can only explain a small proportion of the cross-sectional return differential we show. The returns do not appear to result from equilibrium conditions.

A second explanation is that the returns are compensation for risk exposures. In the absence of a formal model of cross-sectional risk in an options investing environment, our ability to comprehensively test for risk exposures is limited. We adopt a linear factor model approach, recognising that this is unlikely to be a well-specified characterization of the cross-section of option returns. We take this approach to test whether the returns we show are straightforwardly related to risk factors and characteristics of the options portfolios. While we find some explanatory power for both aggregate risks and higher moment characteristics a significant residual return component remains that is almost the same magnitude as the mean raw return.

While contributing, risk exposures do not explain a material proportion of returns.

Third, we conjecture - and find evidence in support of - a mispricing explanation. We argue that the combination of high hedging demand for a currency together with limits on the amount of capital available to support the writing of options can lead to implied volatility rising above fair value and hence poor subsequent straddle returns. We proxy currency-specific hedging demand using hedger positions in Chicago Mercantile Exchange futures markets, and limits to arbitrage using a standard set of proxies including the TED spread and the VIX. Using panel fixed effects regressions we show that when hedging demand for a currency is high and limits to arbitrage are more binding, the implied volatility for that currency is high and subsequent straddle returns on that currency are low. Further, it is the part of implied volatility that can be explained by excess hedging demand that contributes most to predicting poor subsequent straddle returns. Importantly, we feel there are two reasons why this explanation is unlikely to be valid for the equity market, which explains why our findings are at odds with the equity options literature. First, unlike the foreign exchange market in which both corporates and financial companies have extensive currency exposures they may wish to hedge, there is only limited demand to hedge equity positions. Second, even if hedging demand was high, no individual equity or group of equities is sufficiently large to deplete the pool of capital by enough to drive observable cross-section differences across the hundreds of stocks traded in the equity options market. In contrast, the small number of major currencies traded means that hedging demand can overwhelm the available supply when arbitrage capital is scarce, leading to the cross-sectional mispricing we detect.

Our work also contributes to the literature on cross-sectional strategies in foreign exchange markets. The foreign exchange carry trade has received most of the attention [see, among many others, Lustig and Verdelhan (2007) and Menkhoff et al. (2012a)], but a similar approach has been taken to consider the performance of currencies sorted by momentum [Menkhoff et al. (2012b) and Barroso and Santa-Clara (2015)], customer order flow [Menkhoff et al. (2016)], and other fundamental currency characteristics [Della Corte et al. (2016b) and Menkhoff et al. (2017)]. Della Corte et al. (2016a) rank currencies by their volatility risk premium – the difference between expected realised and implied volatilities – and take positions in the underlying currencies. These papers all demonstrate that excess returns on currencies can be explained by the factors used to perform the portfolio sorts. In some cases these returns are driven by the interest differentials on currencies (most notably in the case of carry), and in others the ranking predicts movements in the spot exchange rate (including, for example, the volatility risk premium). Our paper adopts a different approach and takes positions in foreign exchange options rather than the underlying spot or forward markets. The returns we show are not related to carry-type returns and are not earned by predicting the direction of spot exchange rate movements. The returns earned from long-short straddle positions based on implied volatility ranks are distinct from those detailed in the literature on foreign exchange (or equity) returns to date.

We detail our data sources and the peculiarities of the over the counter market for foreign exchange options in the next section. In Section 3 we first detail the returns earned from portfolios ranked on implied volatility levels and then show that alternative volatility-based measured used in the literature cannot explain our findings. In Section 4 we seek to understand the returns, considering leverage-based equilibrium explanations, risk-based explanations, and the role of hedging and limits to arbitrage. We conclude in Section 5.

2. Data

All data used in this paper are extracted from Bloomberg. The data commence 31 March 1999 and end 28 September 2016. We sample weekly, using New York closing prices each Wednesday. We analyse European-style options with one-month maturity on the nine most liquid currencies against the US dollar (Australian dollar, euro, British pound, Japanese yen, Swiss franc, Canadian dollar, Danish kroner, Norwegian kroner, and Swedish kronor). While the Danish kroner is closely pegged to the euro and so might not be viewed as an independent ninth currency, our results are unchanged if the kroner is excluded from the analysis. We also extract spot rates corresponding to the initiation and expiry dates of each option, one-month forward rates on the initiation date, and initiation date one-month money market interest rates.

Some comments on the dimensions of our data are in order. First, the sample spans almost 18 years, slightly shorter than the 20 years available for US equities. However, extending the sample pre-1999 produces a clear structural change in the cross-section. Several pre-euro European currencies including the German mark, French franc and Italian lira would enter the sample, while data for other currencies (such as Swedish kronor and Norwegian Kroner) are not available. Furthermore, the pre-Euro currencies were closely managed against each other meaning they are not independent assets. Second, we focus on the cross-section of G10 currencies even though options data for some emerging economies are available for at least part of the post-1999 period. Some emerging currencies with active options markets such as the Hong Kong or Taiwanese dollars are pegged to the U.S. dollar and so have very low implied volatilities. If included, such currencies would always be in the lowest volatility portfolios. Conversely, other currencies such as the Brazillian real and

Argentine peso have consistently high implied volatilites and would always lie in the high volatility portfolios. We could include currencies such as the Korean won or Mexican peso with ex-post intermediate volatility levels but this would make sample selection somewhat ad hoc. We prefer to use the full available sample of G10 currencies but recognize that this suggests we should be conservative and note that our findings could only be relevant for advanced economies whose options are traded in liquid markets. Put differently, we know that our findings are not driven by smaller countries whith less liquid options markets.

Our interest is mainly in examining the performance of options positions determined by their volatility characteristics and so we attempt to minimize the impact of movements in the underlying currencies. We do this by considering straddle positions for the majority of our analysis. Straddles are a combination of a call and a put option on a currency with the same maturity and the same strike price. In our analysis we use options with one month maturities and with deltas of 0.5 in absolute value, though these will be of opposite sign for calls and puts. These choices mean that the the strike price of both the call and put will be the one-month forward exchange rate and these options are termed at-the-money-forward (ATMF) options.

The over-the-counter foreign exchange options market follows some important pricing conventions. In particular, quotes are typically made for combinations of options with specific deltas, rather than at fixed strike prices. One such combination, and the basis for the analysis in this paper, is the at-the-money-forward delta-neutral straddle. The implied volatility of this straddle is the 'price' quoted in the market. We collect both bid and ask implied volatilities. These implied volatilities are inverted to determine the strike price and bid and ask option prices. The computed option premia are expressed as US dollar prices of options with notional values of one dollar. Option payoffs are computed in dollar terms. We construct time-series of straddle returns using the sum of the price of the call and put as the reference beginning price and the payout of the option that expires in-the-money as the terminal payoff.

Conventions in the foreign exchange market mean that for most currencies the exchange rate is quoted as foreign currency units per US dollar. A call option pays out when the quoted exchange rate rises. In the case of USD-JPY, for example, this is when the US dollar appreciates relative to the Japanese yen. However, for the Australian dollar, Euro and British pound the rate is quoted as the foreign currency price of one US dollar, and calls pay out when the US dollar depreciates relative to the foreign currency. We collect options according to whether they pay out when the currency in question appreciates against the US dollar, and term these 'long currency' positions. Long currency positions therefore group put options for the majority of currencies with call options for Australian dollar, euro and British pound. Short currency positions are defined analogously.

Finally, we note that the options examined in this paper are extremely actively traded. As detailed in Rime et al. (2013), the Bank for International Settlements estimates the average daily turnover in over-the-counter foreign exchange rate options during April 2013 to be $337bn/day, with volume dominated by the ten major currencies we examine. Percentage quoted spreads differ across currencies but average between 3-4% of premium for the most liquid pairs (EUR-USD and USD-JPY) rising to almost 7% for USD-NOK and USD-SEK and 10% for the least active USD-DKK. These magnitudes are comparable with those faced in equity options markets. This might be puzzling given the much greater levels of activity in FX options markets but note that the quotes supplied by Bloomberg are for deals with a notional value of $10m. Equity options trades are typically much smaller. Furthermore, conversations with traders in this market have suggested that depths of $20-30m at the quote has been normal for the G10 currencies we examine since the mid-1990s, and that price improvement relative to Blomberg quotes is usual in the market. However, we use the full quoted spread when calculating long-short trading strategy returns to be conservative, and consider the importance of price impact below. Goyal and Saretto (2009) and Murray (2013) note that for equity options the short leg of the straddle could require investors to satisfy margin requirements in excess of the price of the straddle. These margin requirements can be costly. For much of our sample period such adjustments did not exist in the much more lightly regulated over-the-counter FX market. Margining and collateral are now entering but depend hugely on the credit rating of the two participants to the deal and the regulatory regime under which the deal is traded. To the extent they continue to enter the FX options markets, such adjustments will negatively impact returns.

3. Analysis

3.1. Returns to straddle positions: One-way sorts on volatility

We begin our analysis by ranking currencies according to their implied volatilities each week. We allocate each currency to one of nine portfolios such that portfolio one contains the currency with the lowest implied volatility and portfolio nine contains the currency with the highest implied volatility. To be precice, currencies are sorted each Wednesday according to the implied volatility quoted in the market that day. Positions are initiated immediately and held for one-month until maturity. We do not impose any lag between sorting and portfolio formation. In the Internet Appendix we show that forming portfolios each week based instead on the *previous* Wednesday's implied volatilities has no material impact on our results. We calculate the returns to holding a long straddle position (long ATMF call, long ATMF put in the same asset) over the subsequent month until expiry. Since we are taking only long positions in options we use ask prices when computing returns for each portfolio. We report results based on ask prices when we sort by implied volatility (or other characteristics) and when we perform Fama-MacBeth cross-sectional regressions. When considering the returns to long-short strategies we use quoted bid and ask prices as appropriate to account for transactions costs.

Panel A of Table 1 reports summary statistics of the returns of the nine straddle portfolios. The first row shows that straddle returns follow a clear pattern with mean returns decreasing from 25.5% to -13.9% as implied volatilities rise. Standard deviations and high-low return ranges are naturally quite large as these are not diversified portfolios. However, a t-test shows that the returns on the two extreme portfolios are significantly different from each other (p-val$<$0.001). After accounting for transactions costs, the P1-P9 (termed low-minus-high or LMH) average monthly return spread is 34.1% with an associated t-statistic of 8.93. LMH returns are positive for 62% of the observations and we plot the cumulated returns of the LMH strategy (after accounting for transactions costs) in Figure 1. Note that this figure plots the returns of the LMH strategy implemented monthly and with a one-month holding period to avoid the effect of overlapping observations caused by weekly data and one-month holding periods. While there are some periods of exceptional performance - specifically 2008 - LMH returns are strongly positive throughout most of the sample and we demonstrate this in section 3.5.1 below. These are the key findings that motivate the rest of the analysis in this paper.

All portfolios lose essentially 100% in some months when both put and call positions in the straddle expire worthless. However, maximum returns are naturally much larger as large price movements in the underlying currency result in one of the two options paying out a multiple of the premium invested. The returns on each portfolio are positively skewed but the skew is largest for the lower volatility portfolios. Accordingly, the LMH straddle portfolio inherits a positive skew of 0.86, and records maximum losses (gains) of -311% (684%). Given the distribution of returns to straddle positions we are concerned that our positive LMH returns are driven by extreme outliers. We remove the effect of such extreme returns by examining iron butterfly positions in robustness analysis below and show that our findings also apply to iron butterfly-based returns. The non-normality, consistent with the findings of Broadie et al. (2009), means we should be cautious when interpreting the statistical significance of our results. To address this issue, we first compute a bootstrap 99% confidence interval for the mean return as follows. We draw with replacement from the original 910 weekly-sampled one-month straddle returns (after transactions costs), creating 10,000 bootstrap samples. We calculate the mean for each of these samples and use the 0.5% and 99.5% percentiles to form the confidence interval. This 99% confidence interval is [24.8%, 44.6%] which confirms that the LMH straddle return is positive and significantly different from zero. Second, we use a similar process to bootstrap the 99% confidence interval for the monthly Sharpe ratio. The interval is [0.222, 0.381], which corresponds to anualized Sharpe ratios of [0.769, 1.319]. Third, we bootstrap the t-statistic associated with the mean return. We subtract the mean monthly return from the time-series of LMH straddle returns, and draw 10,000 samples with replacement. We calculate the bootstrap t-statistic for each sample and use these to calculate critical values for the t-statistic. At the 1% level, the critical value in a two-sided test is 2.55, again confirming the significance of our key result.

Some care is needed in interpreting the mean returns to the LMH strategy. Given that we calculate returns as straddle payoff/straddle premium-1, this can be interpreted as the 'return on capital invested'. If an investor follows the strategy for one period, her unconditional expected return on capital invested is 34.1%. If she allocates an infinitessimally small proportion of her wealth to the strategy, the long-run expected return is also 34.1% per month. However, if she allocates a substantial proportion of her wealth to the strategy she risks bankruptcy before reaching the long-run. A rule-of-thumb might suggest that an amount of capital equal to twice the maximum drawdown is needed to support this strategy. In this case, the return scales to 5.5% per month, which we interpret as the 'return on total capital employed'. Finally, to account for the leverage in option positions we compute the 'return on notional' as (straddle payoff - straddle premium)/notional value of straddle. Returns using this definition are given in Panel B of Table1. The LMH return in this case is 38.9bp/month after transactions costs, or 4.67% per annum.

Scaling by notional value allows us to consider the effects of payoff and premium cost separately. By construction, the average premium cost scaled by notional value increases with implied volatility, from below 1% for the portfolio containing the currency with the lowest implied volatility to over 1.5% for the portfolio containing the currency with the highest. The payoffs earned, however, are essentially unrelated to the sort and all portfolios earn average payoffs close to 1.2% of notional amount. Neither is there any relationship between the sort and the variability of payoffs/notional value. These simple statistics suggest that the LMH straddle strategy performs well not because it predicts the cross-section of future returns but simply because it identifies large cross-sectional differences in the costs of acquiring quite similar payoffs. After demonstrating the robustness of our findings and ruling out other explanations, in Section 4.3 we outline a mispricing-based argument consistent with this simple finding.

The characteristics of the currencies sorted on implied volatility are reported in Table 2. By construction, the mean level of implied volatility increases with the sort, rising from 7.9% to 13.1%. We calculate realised volatilities as the standard deviation of realised daily spot exchange rate returns over the previous twelve months. Though sorted on implied volatilities, the second row of the Table shows that the portfolios exhibit mean levels of historical realised volatilities very similar to the level of implied volatilities. Implied volatilities are slightly higher than realised volatilities, as is typically observed in foreign exchange markets but the opposite to what has been reported by Driessen et al. (2009), among others, for equity markets. The third row of Table 2 reveals no clear relationship between implied volatility and forward premia. This is important since it suggests the LMH return spread is not related to, and certainly not driven by, the failure of uncovered interest rate parity or, equivalently, the success of the carry trade in foreign exchange markets.

The results in table 2 allow us to consider the price impact necessary to eradicate the large HML returns. The mean level of implied volatility for the currency in portfolio 1 is 7.94%. The mean percentage spread on the call (or put) premium for portfolio 1 is 5.22%, implying the implied volatility price quoted in the market would average 7.73-8.15. For the return on the marginal dollar invested in portfolio 1 to reduce to zero, the ask implied volatility would have to increase to over 10%. Muravyev and Pearson (2015) show that sophisticated traders in equity options markets can significantly reduce their transactions costs by timing their trades. The price impact needed to erode the profitability of the HML strategy in FX markets is exceptionally large for any reasonably sophisticated trader to have to pay in such an actively traded market.

The final four rows of Table 2 show that currencies with high volatility also tend to have higher risk reversal prices (and absolute prices). A positive risk reversal price for a currency means the volatility of an out of the money call on the currency is higher than the volatility of the similar put, implying the market expects the currency to appreciate. Our strategy of buying straddles on low volatility and selling straddles on high volatility currencies is related to a directional strategy which trades against market expectations. However, since we trade straddles (or, later, delta hedged calls or puts) our strategy is explicitly designed to be directionless and as we demonstrate below, returns from our strategy are not driven by directional movements in the underlying currencies.

3.1.1. Diversification effects

The LMH strategy is not well-diversified as it takes a long straddle position for the currency with the lowest ranked implied volatility and a short straddle position in the currency with the highest implied volatility. Our data set has a maximum of nine currencies in the cross section. Table 3 summarizes the returns and Sharpe ratios for progressively more diversified portfolios formed from these nine currencies. All portfolios are equally weighted and returns account for transactions costs. Mean returns decline slightly as the number of currencies in the portfolio increases but the standard deviation falls faster meaning that Sharpe ratios improve. Buying straddles on the four currencies with implied volatilities below the cross sectional median and writing straddles on the four currencies with above-median implied volatilities yields mean monthly returns of 21% and an annual Sharpe ratio of 1.31. Figure 2 plots the distributions of LMH returns for different levels of diversification. Diversification reduces but does not erase the non-normality of the returns distributions, however bootstrapping confirms the statistical significance of the mean returns from these strategies.

3.2. Returns to straddle positions: Cross-sectional regressions

We employ the Fama and MacBeth (1973) cross-sectional approach to determine the importance of implied volatility in a regression framework. We interpret our analysis as a characteristics-based approach in which returns are regressed on one or more characteristics measured at the initiation of the position for each set of weekly observations. We then calculate the time series averages of the slope coefficients. That is, we run regressions:

$$R^i_{t,t+k} = \alpha_1 + \gamma_1 IV^i_t + \epsilon^i_{t,t+k} \tag{1}$$

where $R^i_{t,t+k}$ is the return of an at-the-money straddle for currency i purchased at time t with tenor k equal to one month. IV_t is the time t implied volatility of this option position. The results of this regression are given in the first column of Table 4. Inference is very clear as there is a strong negative relationship between implied volatilities and straddle returns consistent with the one-way sort results above. In the second and third columns of the table we show that the relationship is not driven simply by the outlier currencies with the most extreme implied volatilities. In the second column, the regressions are run on the cross-section of up to seven currencies where the currencies with the highest and lowest implied volatility each week are dropped from the cross section. In the third column we drop the two currencies with the highest implied volatilities and the two currencies with the lowest implied volatilities. The estimated coefficient on implied volatility remains statistically significantly negative and in fact slightly increases in magnitude as the cross-section dimension decreases. Similarly, excluding individual currencies and estimating equation (1) based on eight currencies in the cross-section shows that our results are not driven by specific currencies. We also estimate equation (1) using panel regressions with time fixed effects. Our inferences are not affected by this change in estimation technique and results are reported in the Internet Appendix.

In the next two columns we split the sample in half in the time series dimension. Coefficient estimates from the 1999-2007 interval (column (4)) and 2008-2016 (column (5)) are both significantly negative. While the coefficient magnitude is smaller in the more recent period the difference is not statistically significant. Our finding of a negative relationship between implied volatility and straddle returns is consistent through time.

In summary, the one-way sorts of the previous section and the Fama-MacBeth regressions in this section are consistent. Currencies with high (low) implied volatility relative to other currencies are poor (good) value. It is well-known that option implied volatility mean reverts in a time-series sense. Our results are suggestive of cross-sectional mean reversion for foreign exchange options. Subsequent sections test whether option return predictability due to cross-sectional mean reversion is robust to other known drivers of options returns.

3.3. Alternative volatility measures

We now consider whether alternative volatility-based measures can explain our findings. In this section we concentrate on transformations of volatility, specifically the volatility risk premium (the difference between historical and implied volatilities), and lagged and future changes in implied volatilities. We show that neither the volatility risk premium nor lagged or future changes in volatility capture the same information as implied volatility. In the following section we decompose total implied volatility into components and see whether these improve our power to explain straddle returns. They do not, and implied volatility remains the better measure on which to sort.

Goyal and Saretto (2009) conduct an exercise very similar to ours but using returns on equity option straddles. However they note that in their sample sorting on either implied volatilities or realised volatilities does not produce economically important or statistically significant differences in returns in the cross-section. Instead, they form portfolios based on the difference between historical realised volatilities and implied volatilities. They motivate the use of this variable with an appeal to the presence of mispricing in options markets due to investor overreaction. Cao and Han (2013) report results of a similar relationship when looking at returns to hedged equity options returns.

Though the motivation is not the same, the difference between historical and implied volatilities has also been demonstrated to have power in foreign exchange markets. Della Corte et al. (2016a) examine returns to speculating in the FX forward markets based on the difference between historical and implied volatilities, a measure often referred to as the volatility risk premium.[3] They show that the time t volatility risk premium predicts future currency returns. Further, they show that returns to trading based on the volatility risk premium stem from movements in spot exchange rates (and not from gathering forward premia). The negative relationship between straddle returns and volatility levels that we observe could be due to a negative volatility risk premium combined with a positive correlation between the level of volatility and the volatility risk premium. A negative volatility risk premium has been shown in several papers, including Della Corte et al. (2016a) in FX markets. Similarly, a positive correlation between volatility levels and the volatility risk premium is observed in many option pricing models, including Heston (1993).

In the first row of Table 5 we consider the performance of straddle positions on currencies ranked according to the difference between historical realised volatility and current implied volatility. There is no clear relationship. While currencies with a high volatility risk premium outperform those with a low one, the relationship is non-monotonic and bears little resemblance to that seen based on ranking on implied (or realised) volatility. In Table 6 we use Fama-MacBeth regressions to test the power of the difference between historical realised and current implied volatilities to explain straddle returns. For convenience, the first column of this table repeats the benchmark results based on implied volatility rankings. On its own, the difference between realised and implied volatilities is unrelated to straddle returns (column (2)). When combined with implied volatilities (column (3)) the magnitude of the estimated coefficient increases but it remains far from statistical significance. This volatility difference does not seem related to straddle returns in FX options.

Conversely, and consistent with the results of Della Corte et al. (2016a), the volatility risk premium does explain future spot exchange rate returns. We run Fama-MacBeth regressions with the spot exchange rate change over the tenor of the straddle as dependent variable, and report results in Table 7. Column

[3]In their paper they focus on the difference between historical realised volatility and the synthetic volatility swap rate. This rate is a somewhat complicated model-free implied volatility measure. However, as their text and internet appendix makes clear, the results they obtain from using simple at-the-money implied volatilities in its place are virtually identical.

(1) shows that implied volatility levels do not explain spot returns but the second column shows that the volatility risk premium is significantly related to spot returns. The coefficient on the volatility risk premium is negative suggesting a negative price of volatility risk. Volatility risk premia are important for foreign exchange returns but they are directional and relate to the movements in underlying spot exchange rates. They do not explain delta-neutral straddle returns well.

The motivation for using realised minus implied volatility in Goyal and Saretto (2009) is volatility-related option mispricing. Stocks with high implied volatility could have experienced a recent increase in volatility. If investors overreact to this increase and pay too much for options on high volatility currencies then subsequent straddle returns will be low. Such an overreaction would be consistent with Stein (1989), Poteshman (2001) and Barberis and Huang (2001). In the second row of Table 5 we report the returns to straddles based on sorting on the change in implied volatility over the month leading up to the portfolio formation date. Row (3) repeats this using changes in realised volatility. Neither exhibit particularly clear patterns in returns. Consistent with this, column (4) in Table 6 shows that the change in volatility is not strongly related with straddle returns in Fama-MacBeth regressions when included individually, though we note that lagged changes in implied volatility become significant when included alongside the level of implied volatility (column (5)). The coefficient is positive, however, suggesting that recent increases in implied volatility are associated with future positive returns on the straddles, inconsistent with the overreaction story. More importantly, the coefficient on the level of implied volatility is unaffected by the inclusion of lagged changes in implied volatility. Our results do not appear to be caused by an overreaction to recent increases in volatility.

Cao and Han (2013) note that if any volatility-related option mispricing is corrected during the tenor of the option position then the observed relationship between the level of volatility and straddle returns would disappear once we control for changes in volatility over the tenor of the option. We therefore include the change in implied volatility over the one-month holding period of the straddle position in the Fama-MacBeth regressions reported in column (6) of Table 6. The coefficient on volatility changes is strongly significant and positive. The positive relation between straddle returns and changes in implied volatilities during the holding period of the position is to be expected by definition, even in the absence of mispricing. However, the negative coefficient on the level of volatility remains essentially unaltered. Again, these results are not consistent with an overreaction hypothesis and leave the basic relationship between the level of implied volatility and straddle returns unaffected.

If implied volatilities predict straddle returns but the difference between realised and implied volatilities does not, this suggests that realised volatility ought also predict straddle returns. The fourth row of Table 5 shows that this is indeed the case. We note, however, that realised volatility does not predict quite as well as implied volatility (and alternative versions of realised volatility computed over different historical windows do not improve on implied volatilities).[4] Are our results then simply driven by cross-sectional reversals in realised volatility? That is, do currencies with cross-sectionally higher than average variation subsequently exhibit lower than average lower variation? No, since Fama-MacBeth regressions of future realised volatility calculated over the one-month holding period on realised historical volatility yield a positive and highly significant coefficient (0.74, with a $t-stat$ of 17.6 and an average R^2 of 0.40). Similar regressions using implied volatility as a regressor produce a coeffcient of 0.93 with a $t-stat$ of 28.5 and average R^2 of 0.54. High (low) levels of realised and implied volatility correctly predict high (low) levels of future realised volatility, not cross-sectional reversals.

Recalling the flat profile of payoffs on notional shown in Table 1, the puzzle is that these differing levels of future realised volatility do not translate into differing payoffs. Figure 4 gives, for each portfolio sorted on implied volatility, a cross-plot of future realised volatility against moneyness at expiry. The latter term is defined as the absolute deviation of the spot price at expiry from the strike price, expressed in dollars. While there is a positive correlation in each cell, it is clear that high realised volatility does not always translate into a large moneyness at expiry. It is also clear that the slope of the relation between future realised volatility and moneyness is steeper for lower volatility portfolios – plotted in the top row – than higher volatility portfolios in the bottom row. The average ratio of moneyness to future realised volatility falls almost monotonically from 0.309 (portfolio 1) to 0.208 (portfolio 9). This pattern is not just driven by the outliers in future volatility and remains even if high realised volatility months are excluded. Fama-MacBeth regressions of moneyness on implied or historical volatility both yield slightly negative but statistically insignificant slope coefficients. High implied volatility then predicts high future realised volatility, but this does not translate into larger payoffs. There must be a greater degree of negative serial correlation between daily exchange rate changes in the higher volatility portfolios such that their higher future realized volatility translates into a similar level of moneyness at expiry to that earned by low volatility portfolios. Our finding then is not about cross-sectional reversals in realised volatility but it is, at least in part, about the lack of

[4]We put forward below an explanation for predictable returns based on high hedging demand for a currency combined with limits to arbitrage pushing up the price of options causing mispricing. It seems plausible that raised levels of realised volatility contribute to increased current hedging demand pressure. Given the small number of currencies in the cross-section, portfolio rankings based on realized or implied volatilities are likely to be very similar if realised and implied volatilities are reasonably correlated. Table 5 suggests that ralised volatility acts as a noisy proxy for implied volatility.

predictability of moneyness at expiry. Since historical volatility fails to predict straddle pay-offs, sorting based on this yields predictable returns. Sorting based on implied volatility improves performance since this is more closely related to the cost of the straddles.

Finally, for this section, we consider the information contained in the term structure of implied volatilities. Vasquez (2017) shows that high (low) volatility term structure slopes predict positive (negative) returns for equity straddle portfolios. We repeat his analysis defining the term structure to be the difference between the implied volatility of six-month and one-month ATMF options. The final row of Table 5 confirms the findings of Vasquez (2017) for the foreign exchange options market. Returns increase almost monotonically with the slope of the term structure, and the return difference between extreme currencies is over 23% per month (not accounting for transactions costs) consistent with the five percent per week reported by Vasquez (2017). Fama-MacBeth regressions (columns (7) and (8) of Table 6) show that the relationship between returns and the term structure of implied volatilities is positive and statistically significant when only the term structure is included. However, the coefficient on the term structure of implied volatility falls substantially and loses statistical significance when the *level* of implied volatility is also included. Once again, the coefficient on the level of implied volatility is unaffected by the inclusion of the additional explanatory variable and the goodness of fit statistic is markedly higher when the level of implied volatility is included in the regression.

We conclude that some alternative volatility-based measures have explanatory power in the foreign exchange options market. However, when added to the cross-sectional regressions, none of the alternatives remove the power of the level of implied volatility to predict one month straddle returns. Why the results from Goyal and Saretto (2009) and Cao and Han (2013) do not carry over from the equity options market to the foreign exchange options market is a puzzle. We leave this for future research.

3.4. Decompositions of implied volatility

In this section we consider the abilities of various decompositions of implied volatility to explain future option position returns. In particular, we first check whether our results are purely cross-sectional by considering the predictive power of average time-series levels of implied volatility. We find that while long-lasting cross-sectional differences in implied volatility levels do predict future returns, the time-varying component of implied volatilities has significant additional predictive power. Following this, we decompose implied volatility into idiosyncratic and systematic components to better understand the source of the forecasting power we have demonstrated. We find that straddle returns are best predicted by total implied volatility. Sorting based on either component of volatility does not result in a clear relationship with returns.

It is also natural to wonder whether our findings are simply another example of the low risk anomaly. To a degree they clearly are – implied volatility negatively predicts returns. However, they differ from the findings in the literature in several ways. First, as already noted, systematic and idiosyncratic components of volatility are equally important, and indeed one component does not work well in the absence of the other. It appears that total implied volatility predicts returns in our setting, rather than either simply systematic or idiosyncratic volatility. Second, we show that a betting-against-(volatility)-beta factor, while significantly priced in the cross-section, only accounts for a small fraction of LMH returns. A betting-against-correlation factor explains even less. Third, while our strategy takes a long position in low implied volatility currencies and short positions in high volatility currencies, the returns on the long leg are more volatile than those on the short leg (Table 1). Put differently, the omega of the long straddle is much higher than the omega of the short straddle, contrary to the embedded leverage hypothesis of Frazzini and Pedersen (2012).

It could be a concern that, given the relatively short sample period of around seventeen years, the predictability we have shown stems from long-lasting differences in the levels of implied volatilities across currencies. As an informal test of this, Figure 3 plots the portfolio allocation of each currency based on implied volatility rankings. It is clear that the portfolios to which each currency is allocated varies substantially through time. While the low volatility portfolio is often either the Canadian dollar or British pound, every one of the nine currencies is the currency with the highest implied volatility at some point during our sample period and is allocated to Portfolio 9. In particular contrast with the carry trade, in which the Japanese yen and Swiss franc are dominant funding currencies, the yen and franc show no particular tendency to have persistently high or low implied volatility.

To test more formally whether cross-sectional volatility differences explain our results we first take the time-series average of implied volatility currency by currency and sort based on this time-invariant measure. Note that this approach suffers from look-ahead bias. Our results are best viewed as an upper bound on the performance of an investment strategy based on average volatility levels. The first row of Table 8 details the straddle returns generated by such a sort. There is indeed a tendency for low average implied volatility currencies to outperform those with high average levels, but this relationship is non-monotonic and the LMH spread is actually negative. The relevant Fama-MacBeth regression results are given in Table 9. The first column shows that there is no relationship between straddle returns and the time-varying sensitivity of a currency's implied volatility to the cross-sectional average level of implied volatility (this sensitivity being estimated using rolling regressions). The second column suggests that average levels of implied volatility can predict straddle returns and column (3) shows that when included alongside the time-varying level of implied volatility, the average level maintains power. However, the time-varying level of implied volatility remains significant, albeit with a smaller coefficient than obtained when it is the only regressor. The goodness of fit statistic is also, again, noticeably higher when both the time-varying level of implied volatility is included alongside the average level, suggesting that the time-varying level has important additional predictive power. Our results are not solely driven by very long-lasting cross-sectional differences in volatility levels, although these do have some predictive power.

Cao and Han (2013) demonstrate that total underlying stock volatility has no power to forecast delta-hedged equity options returns. However, they find that returns decrease with the level of idiosyncratic volatility in the underlying stocks. We investigate whether a systematic/idiosyncratic decomposition of implied volatility can provide us with insight in our foreign exchange application as follows. In the spirit of An et al. (2014), for each currency we first regress the level of implied volatility on the cross-sectional average level using a rolling regression with a fifty week window, and collect the coefficient estimates. We calculate the systematic component of implied volatility for a currency each week by multiplying its estimated beta for that week with the prevailing cross-currency average level of implied volatility. The idiosyncratic component then is the difference between the level of implied volatility and the systematic component. Sorting currencies on the basis of these two components results in no clear relationship with straddle returns, as demonstrated by the final two rows of Table 8. High systematic implied volatility does appear to predict low subsequent returns (row 2), but the spread between low and high systematic volatility is much smaller than that from the total level of implied volatility and any relationship is far from monotonic. Row (3) of the table suggests that the idiosyncratic component does not forecast straddle returns at all. These findings are reinforced by the results of Fama-MacBeth regressions in Table 9. Column (4) shows that including idiosyncratic or, in unreported results, systematic volatility alongside total implied volatility results in an insignificant coefficient in the additional term and an unchanged coefficient on implied volatility. Unlike the equity market results of Cao and Han (2013) or An et al. (2014), we find that decomposing implied volatility into systematic and idiosyncratic components does not aid our understanding of straddle return predictability in foreign exchange markets.

Coval and Shumway (2001) show that zero-beta ATM equity index straddles offer consistently negative returns, even once hedged against crash risk. They interpret this apparent overpricing of calls and puts as suggesting exposure to some risk beyond market (or crash) risk. Straddles on other assets do not yield consistently negative returns (and in particular Deutsche mark straddle returns were essentially zero). However, they infer that negative returns on securities whose volatilities are positively correlated with systematic risk (proxied by VIX) suggest market risk is priced. Table 1 shows that Portfolios 6-9 with high cross-sectional implied volatilities all offer negative mean returns. The implied volatilities of these portfolios have large positive correlations of between 0.75 and 0.7625 with VIX, consistent with Coval and Shumway (2001). However, Portfolios 1-3 have large positive mean returns while their implied volatilities are also positively correlated with VIX (ranging from 0.6-0.77) which is not predicted by their hypothesis. We repeat the exercise performed above to split total volatility into systematic and idiosyncratic components but this time use the VIX rather than the cross-sectional average of implied volatilities. Results are largely the same (final row of Table 8). Portfolios with high exposure to systematic volatility underperform those with less exposure but not consistently and any effect is much smaller than found when using total implied volatilities. Vix-based betas are far from significance when included alongside the levels of implied volatility in Fama-MacBeth regressions. Exposure to systematic stochastic volatility does not appear to explain much of the returns pattern we observe.

Bali et al. (2017) and Liu et al. (2018) argue that the low-risk anomaly observed in stock returns is driven by idiosyncratic risk. The inability of idiosyncratic implied volatility to predict straddle returns suggests that their arguments do not carry over to foreign exchange straddles. Frazzini and Pedersen (2014) and Asness et al. (2020) instead claim that the systematic component is more important. If their betting-against-beta (BAB) approach is to explain the pattern of returns on straddle positions we would expect the systematic component of volatility to dominate. This is not the case in our application, but for completeness we proceed with a test of the BAB model. Using rolling regressions of each currency's implied volatility of the cross-section mean level of implied volatility as described above, we obtain time-series estimates of the beta coefficients for each currency. Following Frazzini and Pedersen (2014) we shrink these time-series estimates towards the cross-sectional mean, giving the former a weight of 0.6 and the latter a weight of 0.4. We rank currencies based on these 'shrunk' beta coefficients. The betting-against-volatility-beta (BAVB) factor is then constructed by first allocating all currencies with a shrunk beta greater (less) than the median value to the high (low) beta portfolio. These portfolios are equally weighted and are rescaled to have a beta of one at portfolio formation. The BAVB factor is the zero-beta portfolio that is long the low beta portfolio of straddles and short the high beta portfolio of straddles:

$$r_{t+1}^{BAB} = \frac{1}{\beta_{t+1}^{L}}(r_{t+1}^{L}) - \frac{1}{\beta_{t+1}^{H}}(r_{t+1}^{H}) \tag{2}$$

where β^L is the average beta of the low beta portfolio, r^L is the return on the low beta portfolio of straddles and H denotes equivalent terms for the high beta portfolio.

We first regress straddle returns from portfolios sorted by the level of implied volatility on the BAVB factor. Results are reported in Panel A of Table 10. The loadings on the BAVB factor are all statistically significantly positive and tend to decrease as the implied volatility rank increases. However, large intercept terms remain and these follow the decreasing pattern observed for raw straddle returns sorted by volatility. Second, we regress LMH returns on the BAB factor (the first column of Panel B in Table 10). Again, the BAVB factor is statistically very significant (t-stat = 4.6) but the intercept is still large at 29% (t-stat 7.6) and only slightly lower than the average LMH return of 35%. Similar findings emerge if we examine returns from more diversified LMH straddle portfolios. In unreported results we also consider a betting-against-volatility-correlation (BAVC) factor following Asness et al. (2020) but this is less able to explain our findings than BAVB. In sum, our results suggest that while a betting-against-volatility-beta and, relatedly, the systematic component of implied volatility have some power to explain returns to straddles, the larger part is unexplained.

3.5. Robustness

We perform a number of robustness tests in this subsection. We first show that our main findings regarding LMH returns are not driven by large outliers by instead considering returns on iron butterflies. Second, we demonstrate that our findings are not specific to a small part of the total sample but that they are robustly found in almost all of it. Third, we examine the role of exchange rate jumps but conclude that they are not primary drivers of LMH returns. Finally, we show that returns to delta-hedged options positions exhibit similar characteristics to returns on straddles.

3.5.1. Effect of extreme returns

Given the distribution of returns to straddle positions we could be concerned that our positive LMH returns are driven by extreme outliers. We remove the effect of such extreme returns by instead examining iron butterfly positions as follows. Starting with the long straddle position (long an ATMF call and put) we then write an out-of-the-money (OTM) call and an out-of-the-money put resulting in a (short) iron butterfly position. This has two effects. First, it reduces the cost of the strategy relative to the straddle since we now receive premia on the two options we have written. Second, and more relevant to the issue at hand, it means that payoffs beyond the strike of the out-of-the-money options are flat since the long ATMF call (put) is offset by the short OTM call (put). Extreme returns beyond the OTM strike prices are thus capped. We consider 10-delta OTM options here but results using 25-delta options are very similar and are available on request. Table 11 reports summary statistics of returns on iron butterflies ranked by implied volatility. Mean returns are very similar to those reported for straddles but the standard deviation and maximum returns are much reduced and skewness is less pronounced. The straddle-based results are clearly not driven by extreme returns.

More generally, it is important to demonstrate that our findings are not driven by specific time intervals. Our strategy is long volatility for the currency with the lowest cross-sectional implied volatility. This involves paying premiums and collecting payouts when large movements in exchange rates occur. Depending on the nature of these premia and payouts, the returns on long straddle positions could well be concentrated in a small number of months. In the other leg, where straddles are sold on currencies with the highest cross-sectional volatility, returns are less likely to be spectacular if volatility only reverts slowly.

Table 12 provides descriptive statistics of returns on LMH straddles and 25-delta butterflies broken down into the long and short legs. In this table we sample the data monthly to avoid issues with overlapping observations. There clearly are some spectacular positive returns for the straddle, driven - as expected - by the long leg. Nevertheless, both legs of the strategy and the strategy itself have positive mean and median returns. We also consider the returns of the 25-delta butterfly strategy in which the largest positive and negative returns are hedged away. The butterfly returns are reasonably Normally distributed but maintain the large positive mean and median returns. The short leg of the butterfly has a relatively small median loss, but the long leg of the butterfly and both straddle legs give more positive returns than losses even after transactions costs. Figure 5 plots rolling 12-month average returns from the straddle strategy and shows that the average return is positive in the vast majority of windows. If we consider calendar year averages, only the part-year 1999 records a mean loss to LMH straddle. We conclude that our findings are not specific to a small part of our total sample but are instead robustly found in almost all of it.

3.5.2. Exchange rate jumps

Chernov et al. (2018) discuss jumps in both exchange rates and in the variances of exchange rates. They identify specific days on which either the spot exchange rate and/or the variance of exchange rates jump for four currencies in our sample (Australian dollars, Swiss francs, British pounds and Japanese yen) and for the 21-currency dollar index of Lustig et al. (2014). Concerned once again that our results might be driven by these specific periods, we compute LMH and iron butterfly returns when the sample is restricted to exclude months in which the holding period contains one or more identified jumps. We consider jumps in spot rate, jumps in variance, jumps in the four named currencies and jumps in the dollar index separately and jointly. We also exclude the whole of 2008 since this is the best performing calendar year for the strategy and is also a year full of identified jumps. We perform calculations based on monthly-sampled data here to avoid complications from overlapping data. The full set of tabulated results are reported in the Internet Appendix.

Removing all months with identified jumps of any nature slightly increases the mean returns on both strategies. Straddle returns rise from 42.5% to 45.3% and butterfly returns from 27.3% to 32.3%. Excluding just jumps in the specific exchange rates or excluding all jumps in variance barely affects mean returns, however excluding jumps in the 21-currency index does reduce mean LMH straddle returns to 38.4% from 42.5%. Simply excluding 2008 has the largest effect on returns, reducing straddles returns by 5% per month to 37%. Nevertheless, returns on the LMH straddle and butterfly remain very large. Jumps, while naturally relevant for our strategies' returns, do not appear to be primary drivers of our findings.

Chernov et al. (2018) note that the probability of a jump in variance is increasing in the variance. To test whether this has bearing on our findings we compare portfolio returns when the levels of volatility are comparable. Specifically, we take all weeks when the implied volatility of a currency lies between 9 and 15% and compute mean straddle returns for each portfolio. For Portfolio 1 currencies, volatility between 9 and 15% corresponds approximately to 80th-95th percentile range, while for Portfolio 9 this is the 7th-80th percentile range. This subset of the data taken when volatility is unusually high for Portfolio 1 produces a mean straddle return of 17.12%. The mean straddle return on Portfolio 9 - for which this volatility range is more normal - is -16.65%. Comparing periods when both currencies have comparable levels of volatility, returns differ strongly between low ranked and high ranked portfolios.

3.5.3. Returns to delta-hedged positions

Our work has concentrated on the returns to a straddle which is, by construction, a directionless position that pays off whether the currency appreciates or depreciates as long as it changes value sufficiently to recoup the cost of the options bought. We now move to consider delta-hedged options positions in which, say, a one unit long position in a currency taken through purchasing an option is combined with a delta unit short position in the currency taken in the forward market.[5]

Panel A of Table 13 gives the returns on delta-hedged long currency portfolios sorted by implied volatility (row 1) and the key alternative volatility measures discussed above. Panel B repeats these statistics using delta-hedged short currency positions. As with the straddle portfolios, a sort on the level of implied volatility provides the strongest evidence of a relationship between returns and volatility – low volatility currencies outperform higher volatility ones. Implementing this relationship via a long-short strategy results in an after transactions costs mean return on a LMH delta-hedged long currency position of 43.5% with a monthly Sharpe ratio of 0.24, and the $LMH2$ mean return is 34.6% with a Sharpe ratio of 0.28. The performance of delta-hedged short currency positions is somewhat worse. The LMH mean return is 26.1% per month (the monthly Sharpe ratio is 0.14) while the $LMH2$ mean return is 27.5% (Sharpe ratio is 0.17).

Sorting currencies on changes in implied volatility reveals no clear pattern in delta-hedged returns (row 2). Sorting according to the volatility risk premium does suggest a large difference between the two extreme currencies, especially for long currency positions, but the intermediate currencies show no pattern to returns. The power of the volatility risk premium at least in the extremes is not surprising since, as demonstrated above, it predicts future spot returns and a static delta-hedged position is exposed to directional movements in the underlying assets. The final row of the table shows that sorting on the idiosyncratic component of implied volatility also provides a large LMH mean return (of almost 12% for a long currency position) but this is driven by the poor performance of the high idiosyncratic volatility currency and the mean returns of the other eight portfolios are not monotonic.

In summary, simple one-way sorts suggest that returns to delta-hedged positions in currencies also appear related to cross-sectional differences in implied volatility levels and, at best, only weakly related to alternative volatility measures. We test this more formally using Fama-MacBeth regressions in Tables 14

[5] As in the work of Goyal and Saretto (2009), our delta-hedging is static in that we do not continually rebalance. This reduces the impact of transactions costs but at the cost of reducing the quality of the hedge against movements in the underlying exchange rate as deltas change. Cao et al. (2017) also use a static hedge but show that their results are unchanged when they rebalance daily.

and 15 for long currency and short currency positions, respectively. In each specification, the coefficient on the level of implied volatility is negative and clearly statistically significant. The addition of alternative volatility measures leaves the coefficient on implied volatility effectively unchanged, even in the cases in which these alternative measures are statistically significant.

4. Understanding the returns

In this section we seek to understand better what drives LMH returns. We start by asking whether the relationship between volatility and straddle returns is mechanical and predicted by option pricing theory. Hu and Jacobs (2019) show that under certain conditions, the expected return to a straddle can be shown to be decreasing in volatility. However, while there could be an equilibrium relationship between volatility and expected returns, the magnitude is much smaller than we observe empirically.

Trading on the basis of volatility ranking is clearly a risky investment strategy. We therefore assess whether returns simply offer compensation for risk. While we demonstrate statistically valid relations between risk factors and straddle returns, none of those factors examined can explain a significant proportion of the returns. It does not appear that the volatility ranking-based returns are simply a compensation for risk.

We then ask whether the market simply misprices options in a predictable way. Our explanation has three key characteristics. First, we argue that high levels of hedging demand can lead the price of options (and, equivalently, the implied volatility) on those currencies to increase above fair value. Second, for hedging demand to affect prices in this way there must be frictions that limit the ability of speculators to write such options. Our expectation is that when limited speculative capital is available, high hedging demand in certain currencies can cause option prices to rise and hence returns on long straddle positions in those currencies to be poor. Third, recognising that our results are at odds with those reported for the equity options literature, it is important to note that, as discussed below, this mechanism is unlikely to apply to equity options markets.

4.1. A leverage-based explanation

Coval and Shumway (2001) study expected option returns, deriving a relation between moneyness and expected returns on calls and puts. Hu and Jacobs (2019) extend this analysis and consider the role of the (historical) volatility of the underlying asset in determining expected option returns. Their derivation makes it clear that expected returns on options are affected by leverage, which is in turn a function of maturity, moneyness and volatility. Coval and Shumway (2001) show that increasing leverage due to moneyness reduces call returns, and Hu and Jacobs (2019) show that increased leverage due to volatility has the same effect.

Refining the arguments of Hu and Jacobs (2019), both Chaudhury (2017) and Aretz et al. (2016) argue that the relationship between expected option returns and volatility is ambiguous once the effect of the expected return on the underlying asset is allowed to vary. They note that while higher idiosyncratic volatility lowers (raises) the return of calls (puts), systematic volatility, which is related to expected returns on the underlying, has an opposite-signed effect on returns.

Our results are based on straddles, positions that combine a call and a put with equal moneyness, volatility and maturity. It might be expected that the two opposite signed effects of leverage would cancel out but Hu and Jacobs (2019) and Aretz et al. (2016) demonstrate that the relation between volatility levels and the expected returns on calls and puts are not always exactly equal. Aretz et al. (2016) show that the effect of total volatility on returns is only clear when considered in conjunction with the moneyness of the options. For sufficiently low strike prices, high total volatility raises the expected return to both calls and puts, reducing expected returns for sufficiently high strikes. Nevertheless, the effect of volatility on expected straddle returns is only large when the strike price is far from being at-the-money. As our straddles are at-the-money-forward, and the forward discount over a one-month horizon is small, our results are not driven by the equilibrium relationships highlighted in these papers. We confirm this with a calibration experiment in the Internet Appendix, where we also derive the expected return for straddles in FX options using the Garman-Kohlhagen framework.[6]

4.2. Controlling for risk

In this section we attempt to establish whether the large portfolio returns are compensation for risk. We first consider aggregate risks and rely on time-series regressions of straddle and delta hedged returns on the returns to various factors. A linear factor model is unlikely to be a well-specified characterization of the cross-section of returns to options positions but we take this approach to determine whether the returns detailed in this paper are straightforwardly related to likely aggregate risks.

[6]We are grateful to Anthony Neuberger and Laura Ballotta for help in deriving expected returns.

The aggregate risks we consider are, first, an aggregate volatility proxy computed as the equally-weighted average return from the nine straddle portfolios and, second, a US dollar exposure proxy computed as the equally-weighted average one-month forward return. Ideally, we would have proxied aggregate volatility with the returns on a straddle on the market index. These exist for equities markets but unfortunately not for currency markets. Instead of the return from a straddle on a basket of currencies, we use the average of returns on single currency straddles, comparable with the common individual variance risk measure of Driessen et al. (2009) used in Cao and Han (2013). Results are reported in Table 16. The first two columns use LMH and $LMH4$ straddle returns, respectively, both of which account for transactions costs. Only the average market return to the straddle is significant, suggesting that our cross-section strategy loads significantly on time-series aggregate volatility risk. However, crucially, the constant term, which can be interpreted as the abnormal return relative to the linear pricing model, remains very significant and almost unchanged in value from the raw mean returns after transactions cost (34% for LMH and 21% for $LMH4$). While these aggregate risk factors explain a small proportion of the variation in straddle and delta-hedged returns, the bulk of returns from the long-short strategy remains unexplained.

In an attempt to broaden the set of risks considered, Table 17 presents regression results using the Fung and Hsieh (2004) lookback factors. These factors are the excess returns on lookback straddle options on bonds, currencies, commodities, short-term interest rates and stock indices. They are designed to replicate the maximum possible returns to trend-following strategies on the relevant asset class and have been widely used in studies including Bollen and Whaley (2009), Patton and Ramadorai (2013) and Della Corte et al. (2016a). Table 17 uses one month returns to the LMH implied volatility straddle strategy sampled monthly since the lookback factor data are only available monthly. The first column provides the results of regressing LMH monthly returns on these five factors ($LMH4$ results are broadly similar and so are omitted). Bond and commodity lookback factors are borderline significant for LMH returns but again a substantial unexplained average return remains. Furthermore, similar regressions using butterfly and delta hedged strategy returns as dependent variables provide quite inconsistent results (columns (2) and (3)). The final three columns use standard Fama-French three factors, augmented with the change in the credit spread of Moody's BAA bonds over the 10-year Treasury rate. Again, the results are inconsistent. Overall, while some factors are occasionally significant, all of the strategies have very large alpha relative to these factors since the constant terms remain strongly significant and essentially equal to the raw mean returns in each regression.

Finally, for this section, we use Fama-MacBeth regressions to explore whether differences in the sensitivities of returns to higher moment risks drive our results. In the first stage we regress straddle returns on the squared and cubed cross-sectional average straddle return using a rolling regression with a fifty week window. We collect the slope coefficient estimates, again filling in the initial observations with the first point estimate. We then include these slope coefficient estimates – which we interpret as sensitivities to skew and kurtosis risks – in cross-sectional regressions. Table 18 gives the results. Betas on both squared and cubed mean returns are positive and significant and have decent explanatory power. These coefficients get larger and more significant when the level of implied volatility is added to the regression (column (2)). Nevertheless, while the magnitude of the coefficient on implied volatility is reduced by around one-third compared with those reported in Tables 4 and 6 it remains large and highly statistically significant. The results suggest that while some of the power of implied volatility is due to variation in sensitivities to both skewness and kurtosis risks, implied volatility has significant explanatory power for returns over and above these effects. Again, we note that these results are somewhat different to those from the equity market. Goyal and Saretto (2009) find no evidence indicative of higher moment effects using equity options.

4.3. Hedging demand and limits to arbitrage

If exposures to risk factors can only partially explain the profitability of LMH straddle returns we are forced to consider whether the market misprices these options. A mispricing explanation would involve, first, explaining what forces might lead to mispricing and, second, explaining why arbitrage capital does not quickly return prices to their correct level. In addition, given that our results are at odds with those from the equity options literature, it would be helpful if the explanation is somewhat specific to the foreign exchange market.

The usual suspects in any mispricing explanation are retail investors. However, given the nature of the over-the-counter foreign exchange options market, retail investors are essentially absent. Instead, we consider the actions of hedgers, following Hirshleifer (1990) and, with a more recent discussion of mispricing in forward foreign exchange markets, Borio et al. (2018). We envisage a world in which high levels of demand to hedge specific currencies through the purchase of options lead the price of options (and, equivalently, the implied volatility) on those currencies to increase above fair value. As noted in Hirshleifer (1990), for hedging demand to affect prices in this way there must be frictions that limit the ability of speculators to write such options and here we follow the literature on limits to arbitrage. Our expectation is that when limited speculative capital is available, high hedging demand in certain currencies can cause option prices to rise and hence returns on long straddle positions in those currencies to be poor. Such effects will be largest when hedging demand is high and limits to arbitrage most binding. Conversely, when capital is relatively freely available, hedging demand can be accommodated without price impact, and when capital is limited it takes cross-currency asymmetries in hedging demand to create enough cross-sectional variation in implied volatilities to make the LMH straddle strategy profitable.

We proxy currency hedger demand using the Commodity Futures Trading Commissions (CFTC) Commitment of Traders reports on aggregate hedger positions in currency futures contracts traded on the Chicago Mercantile Exchange (CME). The CFTC reports classify large traders as either commercial or non-commercial. A trader is deemed to be commercial if she is "engaged in business activities hedged by the use of futures or options markets."[7] We follow Bessembinder (1992) and De Roon et al. (2000) and treat commercial traders as hedgers and non-commercial traders as liquidity providers. The Commitment of Traders reports are available for our full sample but only cover six of our nine currencies (AUD, GBP, CAD, EUR, JPY and CHF against the USD). Futures on the Scandinavian currencies were not traded on the CME for the majority of our sample and are, even now, very illiquid. The reports provide a breakdown of each Tuesday's open interest for futures, which are received from reporting firms on Wednesday and publicly released on Friday. Recall that we form our straddle portfolios based upon Wednesday's implied volatilities and so the open interest data we use marginally predates portfolio formation. We compute hedging demand in each currency i at time t as

$$h_{it} = \frac{\textit{number of contracts held by commercial traders}_{it}}{\textit{number of open contracts held by commercial and non-commercial traders}_{it}}$$

Note that this is not the more commonly-used directional measure of net hedging demand used by, for example, De Roon et al. (2000). Rather, we sum the number of long and short currency contracts and calculate the proportion of total open positions held by hedgers.

We also use a standard set of proxies for limits to arbitrage: the TED spread (a measure of funding liquidity, see Ang et al. (2011) and Nagel (2012)), the VIX index (a measure that captures changes in arbitrageurs' capital availability, see Gromb and Vayanos (2002) and Brunnermeier and Pedersen (2009)), and the St. Louis Fed financial distress indicator (which increases with illiquidity and volatility in financial markets, see Della Corte et al. (2016a)). The TED spread and VIX are measured daily and we use their values for each Wednesday when the straddle portfolios are formed. The financial distress indicator is reported

[7] Commitment of Traders Report explanatory notes available at https://www.cftc.gov/MarketReports/CommitmentsofTraders/ExplanatoryNotes/index.htm

each Friday and we use the previous Friday's value. We combine all three proxies of limits to arbitrage to strengthen their power. In the results discussed below we simply take the average of the standardised Z-scores of the proxies each period and denote this composite variable LTA_t. In the Internet Appendix we report essentially identical results using instead the first principal component of the three proxies.

We first run panel fixed effects regressions of straddle returns on hedging demand, limits to arbitrage and the product of these two terms, reporting results in the first column of Table 19. The demand proxy has a negative coefficient, indicating that subsequent returns are lower when hedging demand is high. The limits to arbitrage measure has a positive coefficient, indicating that returns from long straddle positions are higher when market conditions are less conducive to arbitrage. More importantly, the product term bears a significantly negative coefficient implying that long straddle returns are lower when hedging demand is high and limits to arbitrage are more binding, consistent with our hypothesis. Results in Section 4.2 suggest that higher moment risks partially explain straddle returns. Including higher moment terms in the panel excess demand regression only slightly changes the estimates and, in particular, the coefficient on the product term is still significantly less than zero (coefficient is -0.76 with a t-statistic of -3.2). Including time fixed effects - and therefore dropping LTA as it is common to all currencies - does not affect inference (see column (2)).

We next show that the mechanism by which excess hedging demand affects straddle returns is as conjectured. First, we run panel fixed effects regressions of implied volatility on hedging demand and the product of hedging demand and the composite limits to arbitrage variable. We also include time fixed effects to capture comment elements driving implied volatilities each period. Results are reported in column (3). High hedging demand alone is not sufficient to drive implied volatilities higher, but implied volatilities are positively related to the interaction term. When hedging demand for a currency is high *and* limits to arbitrage more binding, the implied volatility for that currency is high. Second, we compute the fitted values from this regression (excluding the fixed time effects) and denote these by IV_{fitted}. Similarly, the residuals are denoted IV_{resid}. We then regress subsequent straddle returns on that part of implied volatility that can be explained by excess hedging demand (IV_{fitted}), that part of implied volatility unrelated to excess demand (IV_{resid}) and time fixed effects. The results are reported in column (4). The coefficient on the part of implied volatility explained by excess hedging demand is strongly significantly negative and of much larger magnitude than the coefficient on the unexplained part of implied volatility. Taken together, the results of Table 19 suggest that excess hedging demand can drive implied volatility above fair value and that this explicable component of variation in implied volatilities that is related to the poor subsequent performance of straddle positions. Our explanation lacks a formal model and our evidence while supportive is preliminary. We hope that future work can shed further light on whether mispricing truly lies behind our findings.

Finally, we note that a hedging-based mispricing explanation is unlikely to hold in equity options markets, and that mechanisms put forward to explain equity options mispricing do not translate to the foreign exchange case. The different reasons for mispricing across the two markets explain why our empirical results are at odds with the existing literature on equity options returns.

Unlike the foreign exchange market in which financial and non-financial corporates have extensive currency exposures they may wish to hedge, there is only limited demand to hedge equity positions. Furthermore, given the limited cross-section of major currencies traded, high hedging demand in a few currencies can overwhelm supply when arbitrage capital is scarce, causing the cross-sectional mispricing we detect. Conversely, no individual equity - or group of equities - is sufficiently large that hedging demand could deplete the pool of arbitrage capital available by enough to cause detectable cross-sectional differences across the hundreds of stocks traded in the options market.

Cao and Han (2013) demonstrate that delta hedged equity option returns decrease with idiosyncratic volatility in the underlying stocks. Their explanation is also based on excess demand driving up options prices, but differs from ours in crucial respects. They argue that the supply of options for stocks with high idiosyncratic risk is limited since such stocks tend to be small and illiquid and hence difficult to hedge, especially for a dealer with potentially thousands of options positions on different underliers in her portfolio at any given time who may prefer to hedge with equity index products. The underlying assets in our case are a small number of currencies with exceptionally liquid spot markets, suggesting that limits on the supply of foreign exchange options are unlikely to be driven by the Cao and Han (2013) mechanisms. Instead of asset-specific characteristics limiting supply, we rely on market-wide shortages as represented by our choice of proxies for the limits to arbitrage. Cao and Han (2013) also differ by arguing that high demand for equity options comes from speculators attracted by high idiosyncratic volatility in the underlying stock, rather than our focus on corporate hedgers likely to be more concerned by total exchange rate volatility. Our measure of hedging demand is, by construction, inversely related to speculative demand and so our empirical results are also at odds with the Cao and Han (2013) mechanism.

5. Conclusion

We have considered the returns from taking long-short straddle positions in currencies sorted according to their implied volatilities. Straddles on currencies with low implied volatility significantly and consistently outperform straddles on those with high implied volatility. More sophisticated sorts based on either components of the total implied volatility or relationships between alternative measures of volatility do not perform as well and when included as additional explanatory factors do not remove the power of implied volatility to predict future straddle returns. Our findings suggest that the cross-sectional mean level of implied volatility is a very good measure of fair value for major currencies.

Returns from the appropriate long-short straddle strategy are not well-explained by standard risk factors and they are not related to other FX strategies known to produce positive returns such as the carry trade. The high average return to this strategy appears to imply that currencies with outlying levels of volatility are mispriced (or exposed to a risk that we have not considered). We present some evidence consistent with the conjecture that when the availability of speculative capital is limited, high hedging demand for a currency causes its implied volatility to rise and hence subsequent returns on a long straddle position in that currency to be poor.

6. Acknowledgements

We thank audiences at Cass Business School, Manchester Business School, Essex Business School, University of Paris Dauphine, Heriot Watt University and Swansea University, together with Kevin Aretz, Eser Arisoy, Laura Ballotta, Peter Billington, Constantin Bolz, Michael Brennan, Giulia Fantini, Thomas Flury, Divya Goel, Guenter Grimm, Aneel Keswani, Mamdouh Medhat, Nick Motson, Anthony Neuberger, Richard Payne, Ser-Huang Poon, Marco Realdon, Lucio Sarno, Maik Schmeling and Daniel Trum for helpful comments and discussions. The comments of an anonymous referee also greatly enhanced this paper. All errors are our own. The views expressed in this document are those of the authors and not those of Commerzbank or the Bank of England. The information contained in is paper may be based on trading ideas where Commerzbank may trade in such financial instruments with customers or other counterparties. Any prices provided herein are historical only, and do not represent firm quotes as to either size or price. The past performance of financial instruments is not indicative of future results.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

An, B.-J., Ang, A., Bali, T. G., Cakici, N., 2014. The joint cross section of stocks and options. The Journal of Finance 69, 2279–2337.

Ang, A., Gorovyy, S., Van Inwegen, G. B., 2011. Hedge fund leverage. Journal of Financial Economics 102, 102–126.

Ang, A., Hodrick, R. J., Xing, Y., Zhang, X., 2006. The cross-section of volatility and expected returns. The Journal of Finance 61, 259–299.

Aretz, K., Lin, M.-T., Poon, S.-H., 2016. Moneyness, total, systematic, and idiosyncratic volatility, and the cross-section of european option returns. Available at SSRN 2872168 .

Asness, C., Frazzini, A., Gormsen, N. J., Pedersen, L. H., 2020. Betting against correlation: Testing theories of the low-risk effect. Journal of Financial Economics 135, 629–652.

Bali, T. G., Brown, S. J., Murray, S., Tang, Y., 2017. A lottery-demand-based explanation of the beta anomaly. Journal of Financial and Quantitative Analysis 52, 2369–2397.

Bali, T. G., Hovakimian, A., 2009. Volatility spreads and expected stock returns. Management Science 55, 1797–1812.

Barberis, N., Huang, M., 2001. Mental accounting, loss aversion, and individual stock returns. The Journal of Finance 56, 1247–1292.

Barroso, P., Santa-Clara, P., 2015. Beyond the carry trade: Optimal currency portfolios. Journal of Financial and Quantitative Analysis 50, 1037–1056.

Bessembinder, H., 1992. Systematic risk, hedging pressure, and risk premiums in futures markets. The Review of Financial Studies 5, 637–667.

Bollen, N. P., Whaley, R. E., 2009. Hedge fund risk dynamics: Implications for performance appraisal. The Journal of Finance 64, 985–1035.

Borio, C. E., Iqbal, M., McCauley, R. N., McGuire, P., Sushko, V., 2018. The failure of covered interest parity: Fx hedging demand and costly balance sheets. BIS Working Papers .

Broadie, M., Chernov, M., Johannes, M., 2009. Understanding index option returns. Review of Financial Studies 22, 4493–4529.

Brunnermeier, M. K., Pedersen, L. H., 2009. Market liquidity and funding liquidity. Review of Financial Studies 22, 2201–2238.

Cao, J., Han, B., 2013. Cross section of option returns and idiosyncratic stock volatility. Journal of Financial Economics 108, 231–249.

Cao, J., Han, B., Tong, Q., Zhan, X., 2017. Option return predictability. Available at SSRN 2698267 .

Chaudhury, M., 2017. Volatility and expected option returns: A note. Economics Letters 152, 1–4.

Chernov, M., Graveline, J., Zviadadze, I., 2018. Crash risk in currency returns. Journal of Financial and Quantitative Analysis 53, 137–170.

Coval, J. D., Shumway, T., 2001. Expected option returns. The Journal of Finance 56, 983–1009.

De Roon, F. A., Nijman, T. E., Veld, C., 2000. Hedging pressure effects in futures markets. The Journal of Finance 55, 1437–1456.

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

Della Corte, P., Riddiough, S. J., Sarno, L., 2016b. Currency premia and global imbalances. Review of Financial Studies 29, 2161–2193.

Driessen, J., Maenhout, P. J., Vilkov, G., 2009. The price of correlation risk: Evidence from equity options. The Journal of Finance 64, 1377–1406.

Fama, E. F., MacBeth, J. D., 1973. Risk, return, and equilibrium: Empirical tests. The Journal of Political Economy pp. 607–636.

Frazzini, A., Pedersen, L. H., 2012. Embedded leverage. Tech. rep., National Bureau of Economic Research.

Frazzini, A., Pedersen, L. H., 2014. Betting against beta. Journal of Financial Economics 111, 1–25.

Fung, W., Hsieh, D. A., 2004. Hedge fund benchmarks: A risk-based approach. Financial Analysts Journal 60, 65–80.

Goyal, A., Saretto, A., 2009. Cross-section of option returns and volatility. Journal of Financial Economics 94, 310–326.

Gromb, D., Vayanos, D., 2002. Equilibrium and welfare in markets with financially constrained arbitrageurs. Journal of Financial Economics 66, 361–407.

Heston, S. L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies 6, 327–343.

Hirshleifer, D., 1990. Hedging pressure and futures price movements in a general equilibrium model. Econometrica: Journal of the Econometric Society pp. 411–428.

Hu, G., Jacobs, K., 2019. Volatility and expected option returns. Journal of Financial and Quantitative Analysis p. 1–77.

Liu, J., Stambaugh, R. F., Yuan, Y., 2018. Absolving beta of volatility's effects. Journal of Financial Economics 128, 1–15.

Lustig, H., Roussanov, N., Verdelhan, A., 2014. Countercyclical currency risk premia. Journal of Financial Economics 111, 527–553.

Lustig, H., Verdelhan, A., 2007. The cross section of foreign currency risk premia and consumption growth risk. The American Economic Review 97, 89–117.

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

Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012b. Currency momentum strategies. Journal of Financial Economics 106, 660–684.

Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2016. Information flows in foreign exchange markets: Dissecting customer currency trades. The Journal of Finance 71, 601–634.

Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2017. Currency value. Review of Financial Studies 30, 416–441.

Muravyev, D., Pearson, N. D., 2015. Options trading costs are lower than you think. The Review of Financial Studies .

Murray, S., 2013. A margin requirement based return calculation for portfolios of short option positions. Managerial Finance 39, 550–568.

Nagel, S., 2012. Evaporating liquidity. Review of Financial Studies 25, 2005–2039.

Patton, A. J., Ramadorai, T., 2013. On the high-frequency dynamics of hedge fund risk exposures. The Journal of Finance 68, 597–635.

Poteshman, A. M., 2001. Underreaction, overreaction, and increasing misreaction to information in the options market. The Journal of Finance 56, 851–876.

Rime, D., Schrimpf, A., et al., 2013. The anatomy of the global fx market through the lens of the 2013 triennial survey. BIS Quarterly Review .

Stein, J. C., 1989. Efficient capital markets, inefficient firms: A model of myopic corporate behavior. The Quarterly Journal of Economics pp. 655–669.

Vasquez, A., 2017. Equity volatility term structures and the cross section of option returns. Journal of Financial and Quantitative Analysis pp. 1–28.

Table 1
Summary Statistics of Returns on Currency Straddles Ranked by Implied Volatility

This table reports summary statistics for the one-month holding period returns to ATMF straddle positions on currencies sorted by implied volatilities. The currency with the highest implied volatility each week is allocated to Portfolio 9, the currency with the lowest volatility to Portfolio 1. Portfolios are then held for one-month until expiry. Panel A reports statistics based on returns defined as straddle payoff/straddle premium-1. Panel B reports statistics based on returns scaled by the notional value of positions.

	Portfolios sorted by IV								
	P1 Lo	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9 Hi**
Panel A:									
Mean	0.255	0.256	0.117	-0.003	0.098	-0.111	-0.081	-0.089	-0.139
Median	-0.026	0.048	-0.122	-0.177	-0.120	-0.298	-0.237	-0.264	-0.325
Std Dev.	1.161	1.052	0.909	0.808	0.943	0.728	0.723	0.709	0.704
Skew	2.142	1.751	1.453	1.289	1.514	1.416	1.202	0.985	1.292
Min	-0.997	-0.998	-0.996	-0.999	-0.998	-0.995	-0.998	-0.997	-1.000
Max	9.123	7.127	4.545	3.676	5.222	3.478	3.981	2.776	3.661
Panel B:									
Ret. on Notional	0.002	0.002	0.001	-0.000	0.001	-0.002	-0.001	-0.002	-0.002
Prem. on Not.	0.009	0.010	0.011	0.012	0.012	0.013	0.013	0.014	0.015
Payoff on Not.	0.011	0.013	0.012	0.012	0.013	0.011	0.012	0.012	0.013
Std Dev. Payoff	0.011	0.011	0.010	0.010	0.012	0.009	0.010	0.010	0.011

Table 2
Summary Statistics of Currencies Ranked by Implied Volatility

This table reports summary statistics for characteristics of currencies allocated to portfolios based upon implied volatilities. The currency with the highest implied volatility each week is allocated to Portfolio 9, the currency with the lowest volatility to Portfolio 1. The table reports the average values of implied volatility, historical realised volatility, forward discount, together with 10 and 25 delta risk reversal prices and absolute prices for each portfolio.

	Portfolios sorted by IV									
	P1 Lo	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9 Hi**	**Average**
Mean IV	0.079	0.088	0.096	0.103	0.107	0.111	0.116	0.121	0.131	0.106
Mean RV	0.082	0.088	0.095	0.101	0.105	0.107	0.114	0.118	0.124	0.104
Mean Fwd Disc	0.000	0.000	-0.000	-0.000	-0.000	-0.000	0.000	0.000	0.000	0.000
Mean RR10	0.004	0.005	0.005	0.007	0.008	0.008	0.013	0.013	0.017	0.009
Mean RR25	0.002	0.003	0.003	0.004	0.004	0.004	0.006	0.006	0.008	0.004
Mean abs(RR10)	0.010	0.012	0.012	0.014	0.017	0.017	0.020	0.021	0.024	0.016
Mean abs(RR25)	0.006	0.006	0.007	0.007	0.009	0.009	0.010	0.010	0.012	0.009

Table 3
Long-short Returns and Diversification

This table reports summary statistics for the one-month holding period returns to long-short straddle positions on currencies sorted by implied volatilities. The currency with the highest implied volatility each week is allocated to Portfolio 9, the currency with the lowest volatility to Portfolio 1. LMH denotes a portfolio that takes a long straddle position in Portfolio 1 and a short straddle position in Portfolio 9. $LMHx$ denotes a portfolio that takes a long straddle position in the x currencies with the lowest implied volatilities and a short straddle position in the x currencies with the highest implied volatilities. Positions are held for one-month until expiry. Bid and ask quotes are used to account for transactions costs.

	Mean return	**Std dev.**	**Annual Sharpe ratio**
LMH	0.341	1.151	1.02
$LMH2$	0.322	0.810	1.37
$LMH3$	0.261	0.654	1.38
$LMH4$	0.211	0.553	1.31

Table 4
Univariate Fama-MacBeth Regressions

This table reports the results of Fama-MacBeth regressions. The dependent variable in each case is the return on straddle positions in currencies. The single explanatory variable in each column is the implied volatility (IV) of the currency. Column (1) uses the full sample. In column (2) the currencies with the highest and lowest implied volatilities each week are excluded. In column (3) we exclude the two currencies with the highest volatility and the two currencies with the lowest volatility each week. Column (4) reports results using the first half of the sample while column (5) uses the second half of the data sample. Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1) **Full**	(2) **P2-P8**	(3) **P3-P7**	(4) **H1**	(5) **H2**
IV	-9.882	-11.298	-13.260	-11.523	-8.241
	(-8.437)	(-7.500)	(-3.906)	(-6.666)	(-5.255)
Observations	7,875	6,055	4,236	3,780	4,095
R-squared	0.183	0.232	0.318	0.192	0.173
Number of groups	910	910	910	455	455

Table 5
Returns on Currency Straddles Ranked by Other Volatility Measures

This table reports mean returns to straddle positions on currencies ranked by alternative volatility measures. The currency with the highest volatility measure each week is allocated to Portfolio 9, the currency with the lowest volatility to Portfolio 1. In the first row, currencies are ranked according to the difference between historical realised volatility and current implied volatility. Row 2 reports returns to the straddle portfolios on currencies ranked by the change in implied volatilities over the previous month. Row 3 reports straddle returns from rankings based on changes in historical realised volatilities over the previous month. Row 4 reports mean straddle returns from rankings based on the term structure of implied volatilities. The term structure is calculated as the difference between implied volatilities of options with six and one month to expiry. The final row reports mean straddle returns from rankings based on realised historical volatilities calculated over the previous year

	P1 Lo	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9 Hi**
Sorted by $RV - IV$	-0.031	0.031	0.025	0.072	0.143	0.049	0.040	-0.040	0.031
Sorted by dIV	0.000	-0.003	0.064	0.060	0.052	0.031	0.037	0.027	0.018
Sorted by dRV	-0.043	0.086	0.045	-0.009	0.073	0.057	0.054	0.043	-0.013
Sorted by RV	0.236	0.267	0.127	0.102	-0.014	-0.058	-0.103	-0.100	-0.133
Sorted by $TSIV$	-0.113	-0.078	-0.022	0.036	0.092	0.073	0.078	0.113	0.122

Table 6
Fama-MacBeth Regressions – Controls

This table reports the results of Fama-MacBeth regressions. The dependent variable in each case is the return on straddle positions in currencies. Explanatory variables are denoted as follows: IV is the current implied volatility, RV is the historical realised volatility, $RV - IV$ is the difference between realised and implied volatilities, dIV is the one month lagged change in implied volatility, $fdIV$ is the one month change in implied volatility measured over the life of the straddle position, and $TSIV$ is the term structure of implied volatility defined as the difference between the implied volatility of options with six months and one month to maturity. Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IV	-9.882		-10.744		-10.811	-9.612		-10.079
	(-8.437)		(-7.908)		(-8.194)	(-8.217)		(-7.555)
$RV - IV$		-0.228	-3.205					
		(-0.112)	(-1.434)					
dIV				1.611	5.195			
				(0.747)	(2.177)			
$fdIV$						21.374		
						(7.601)		
$TSIV$							17.622	6.213
							(4.821)	(1.555)
Observations	7,875	7,875	7,875	7,875	7,875	7,875	7,875	7,875
R-squared	0.183	0.160	0.334	0.177	0.343	0.390	0.166	0.338
Number of groups	910	910	910	910	910	910	910	910

Table 7
Fama-MacBeth Regressions – Spot Exchange Rate Changes

This table reports the results of Fama-MacBeth regressions. The dependent variable in each case is the one-month change in the spot exchange rate. Explanatory variables are denoted as follows: IV is the current implied volatility, and $RV - IV$ is the difference between realised and implied volatilities. Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1)	(2)	(3)
IV	0.009		-0.016
	(0.305)		(-0.483)
$RV - IV$		-0.076	-0.076
		(-1.746)	(-1.495)
Observations	7,875	7,875	7,875
R-squared	0.187	0.200	0.362
Number of groups	910	910	910

Table 8
Returns on Currency Straddles Ranked by Decompositions of Implied Volatility

This table reports mean returns to straddle positions on currencies ranked by alternative decompositions of implied volatility. The currency with the highest volatility measure each week is allocated to Portfolio 9, the currency with the lowest volatility to Portfolio 1. In the first row, currencies are sorted by the sample average level of implied volatility. In the second row, currencies are sorted by systematic volatility as defined in section 3.4. In the third row, currencies are sorted by idiosyncratic volatility defined as the difference between total and systematic volatility. In the final rwo we compute systematic volatility using exposure to the VIX.

	P1 Lo	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9 Hi**
Sorted by avg. IV	-0.092	0.607	-0.044	0.174	-0.001	-0.050	-0.198	-0.064	-0.065
Sorted by syst. IV	-0.006	0.130	0.090	0.122	0.075	-0.013	0.024	-0.028	-0.099
Sorted by idio. IV	0.007	0.012	0.046	0.061	0.146	0.026	0.070	0.002	-0.052
Sorted by exp. to VIX	0.105	0.064	0.106	0.129	-0.046	-0.016	0.028	-0.026	-0.051

Table 9
Fama-MacBeth Regressions – Implied Volatility Decompositions

This table reports the results of Fama-MacBeth regressions. The dependent variable in each case is the return on straddle positions in currencies. Explanatory variables are denoted as follows: IV is the current implied volatility, βIV is the estimated coefficient from the regression of IV on the cross-sectional average level of IV, $IVavg$ is the time series average level of implied volatility, and $IVidio$ is idiosyncratic volatility defined as the difference between total and systematic volatility as defined in section 3.4. Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1)	(2)	(3)	(4)
IV	-11.733		-7.257	-11.167
	(-8.655)		(-4.364)	(-8.703)
βIV	0.083			
	(1.264)			
$IVavg$		-11.495	-5.134	
		(-10.591)	(-3.167)	
$IVidio$				-0.567
				(-0.827)
Observations	7,875	7,875	7,875	7,875
R-squared	0.343	0.157	0.324	0.343
Number of groups	910	910	910	910

Table 10
Betting-Against-Volatility-Beta Results

Panel A of this table reports results from regressing returns to straddle positions on currencies ranked by implied volatility on a betting-against-(volatility)-beta factor constructed as described in the text. The currency with the highest implied volatility each week is allocated to Portfolio 9, the currency with the lowest volatility to Portfolio 1. Portfolios are then held for one-month until expiry. The first row reports the regression coefficient on the betting-against-beta factor. The second row reports the intercept term. Associated t-statistics are given in parentheses. The third row gives the goodness-of-fit of the regression. Panel B repeats this using LMH strategy returns using straddles with varying degrees of diversification (denoted $LMH - LMH4$) and 10-delts and 25-delta iron butterfly positions both without and with maximum diversification.

Panel A:									
	P1 Lo	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9 Hi**
BAB	0.821	0.684	0.540	0.394	0.477	0.181	0.203	0.160	0.286
	(7.05)	(8.31)	(7.44)	(6.26)	(7.03)	(3.64)	(3.61)	(3.34)	(5.46)
Intercept	0.158	0.196	0.052	-0.042	-0.022	-0.136	-0.100	-0.100	-0.174
	(4.39)	(5.79)	(1.75)	(1.57)	(0.59)	(5.53)	(4.02)	(4.01)	(7.22)
R-squared	0.202	0.172	0.148	0.099	0.142	0.028	0.034	0.021	0.067
Panel B:									
	LMH	$LMH2$	$LMH3$	$LMH4$					
BAB	0.511	0.521	0.431	0.379					
	(4.63)	(8.12)	(10.31)	(12.12)					
Intercept	0.296	0.269	0.279	0.250					
	(7.59)	(8.82)	(12.74)	(13.92)					
R-squared	0.082	0.140	0.184	0.213					

Table 11
Returns on Iron Butterflies Ranked by Implied Volatility

This table reports summary statistics for the one-month holding period returns to 10-delta iron butterfly positions on currencies sorted by implied volatilities. The currency with the highest implied volatility each week is allocated to Portfolio 9, the currency with the lowest volatility to Portfolio 1. Portfolios are then held for one-month until expiry.

	P1 Lo	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9 Hi**
Mean	0.241	0.245	0.103	0.009	0.071	-0.102	-0.083	-0.089	-0.141
Median	0.089	0.156	-0.017	-0.085	-0.019	-0.221	-0.149	-0.181	-0.243
Std Dev.	0.922	0.826	0.742	0.700	0.732	0.605	0.591	0.598	0.581
Skew	0.794	0.487	0.723	0.671	0.529	0.449	0.241	0.259	0.446
Min	-0.996	-0.998	-0.996	-0.999	-0.998	-0.994	-0.998	-0.997	-1.000
Max	3.014	2.906	3.497	3.131	2.302	2.023	1.311	2.074	2.071

Table 12
Descriptive Statistics of Straddle and Butterfly Returns, Data Sampled Monthly

This table reports summary statistics for the one-month holding period returns to LMH straddle and 25-delta iron butterfly positions on currencies sorted by implied volatilities.

	Mean	25%ile	50%ile	75%ile	Stdev	Skew	Kurt
Straddle	0.4253	-0.1729	0.2781	0.9473	1.0961	1.5633	9.6923
Long	0.3075	-0.4410	0.0556	0.7745	1.1605	2.7828	16.031
Short	0.1178	-0.3156	0.3062	0.6556	0.7183	-1.2488	4.5553
Fly25	0.2643	-0.1961	0.2146	0.8091	0.7310	-0.0145	2.7669
Long	0.2597	-0.1521	0.2561	0.6397	0.6481	0.1327	2.4085
Short	0.0046	-0.4480	-0.0577	0.4127	0.4928	0.2129	1.8233

Table 13
Returns on Delta Hedged Positions Ranked by Decompositions of Implied Volatility

This table reports mean returns to delta hedged positions on currencies ranked by alternative volatility measures. The currency with the highest volatility measure each week is allocated to Portfolio 9, the currency with the lowest volatility to Portfolio 1.

	P1 Lo	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9 Hi**
Panel A: Long curr.									
Sorted by IV	0.335	0.249	0.150	0.039	0.056	-0.066	-0.026	-0.069	-0.143
Sorted by dIV	0.100	0.064	0.104	0.119	-0.024	0.068	0.045	0.054	-0.032
Sorted by RV-IV	-0.077	0.079	0.044	0.066	0.118	0.089	0.032	0.070	0.125
Sorted by idio. IV	0.023	0.041	0.113	0.126	0.146	0.090	0.092	0.020	-0.094
Panel B: Short curr.									
Sorted by IV	0.170	0.269	0.079	-0.040	0.122	-0.146	-0.133	-0.104	-0.139
Sorted by dIV	-0.100	-0.065	0.011	0.009	0.122	0.008	0.024	-0.004	0.074
Sorted by RV-IV	0.012	-0.022	0.049	0.056	0.179	-0.005	0.036	-0.143	-0.064
Sorted by idio. IV	0.016	-0.042	-0.005	-0.014	0.148	-0.037	0.049	-0.016	-0.011

Table 14
Fama-MacBeth Regressions – Delta-Hedged Long Currency Positions

This table reports the results of Fama-MacBeth regressions. The dependent variable in each case is the return on delta-hedged long positions in currencies. Explanatory variables are denoted as follows: IV is the current implied volatility, dIV is the change in implied volatility over the month prior to portfolio formation, $RV - IV$ is the difference between realised and implied volatilities, $IVidio$ is idiosyncratic volatility defined as the difference between total and systematic volatility, and $FwdDisc$ is the forward discount of each currency relative to the US dollar (which also equals the difference between strike and spot price). Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1)	(2)	(3)	(4)	(5)
IV	-10.909	-11.387	-10.573	-11.616	-9.835
	(-5.649)	(-5.164)	(-4.855)	(-5.601)	(-4.694)
dIV		-1.318			
		(-0.332)			
$RV - IV$			0.836		
			(0.190)		
$IVidio$				-1.617	
				(-1.427)	
$FwdDisc$					24.762
					(1.104)
Observations	7,875	7,875	7,875	7,875	7,875
R-squared	0.173	0.340	0.337	0.345	0.351
Number of groups	910	910	910	910	910

Table 15
Fama-MacBeth Regressions – Delta-Hedged Short Currency Positions

This table reports the results of Fama-MacBeth regressions. The dependent variable in each case is the return on delta-hedged short positions in currencies. Explanatory variables are denoted as follows: IV is the current implied volatility, dIV is the change in implied volatility over the month prior to portfolio formation, $RV - IV$ is the difference between realised and implied volatilities, $IVidio$ is idiosyncratic volatility defined as the difference between total and systematic volatility, and $FwdDisc$ is the forward discount of each currency relative to the US dollar (which also equals the difference between strike and spot price). Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1)	(2)	(3)	(4)	(5)
IV	-8.855	-10.226	-10.899	-10.702	-9.380
	(-4.459)	(-4.680)	(-4.568)	(-4.665)	(-4.631)
dIV		11.559			
		(2.469)			
$RV - IV$			-7.182		
			(-2.309)		
$IVidio$				0.437	
				(0.338)	
$FwdDisc$					-5.118
					(-0.211)
Observations	7,875	7,875	7,875	7,875	7,875
R-squared	0.173	0.359	0.341	0.369	0.349
Number of groups	910	910	910	910	910

Table 16
Risk Adjusted Options Returns I

This table reports the results of OLS regressions of long-short options positions on explanatory variables. The dependent variables in columns (1) and (2) are single-currency and four-currency long-short straddle returns. Explanatory variables are denoted as follows: $Agg.Vol_{straddle}$ is the equaly-weighted cross-sectional average return from the straddle positions, and $Agg.Forward$ is the equally-weighted cross-sectional average return on one-month forward positions. Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Agg.Vol_{strad}$	0.508	0.260						
	(4.946)	(5.459)						
$Agg.Forward$	1.650	0.657	0.326	0.168	0.246	-0.012	11.681	-22.697
	(0.762)	(0.620)	(0.214)	(0.227)	(0.226)	(-0.022)	(2.510)	(-4.169)
$Agg.Vol_{fly10}$			0.370	0.146				
			(4.792)	(4.101)				
$Agg.Vol_{fly25}$					0.328	0.106		
					(3.407)	(2.304)		
$Agg.Vol_{DHlong}$							0.555	
							(4.516)	
$Agg.Vol_{DHshort}$								0.895
								(5.029)
Constant	0.334	0.205	0.281	0.153	0.175	0.061	0.409	0.251
	(9.412)	(11.911)	(8.822)	(9.955)	(6.661)	(4.820)	(7.203)	(4.334)
Observations	910	910	910	910	910	910	910	910
R-squared	0.062	0.068	0.024	0.017	0.011	0.005	0.061	0.120

Table 17
Risk Adjusted Options Returns II

This table reports the results of OLS regressions of long-short straddle, 10-delta iron butterfly and delta hedged positions on risk factors. Explanatory variables are denoted as follows: $LBbond$, $LBcurr$, $LBcomm$, $LBir$ and $LBstock$ denote Fung-Hsieh lookback factors calculated as the excess returns on lookback straddle options on bonds, currencies, commodities, short-term interest rates and stock indices respectively. Mkt, SMB and HML are the usual Fama-French factors and $dBaaBond$ is the change in the credit spread of Moody's Baa bonds over the 10-year Treasury yield. These regressions are at the monthly frequency. Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1) Straddle	(2) Butterfly	(3) Delta Hedge	(4) Straddle	(5) Butterfly	(6) Delta Hedge
$LBbond$	-1.025	-0.826	-1.160			
	(-1.901)	(-1.950)	(-1.604)			
$LBfx$	0.717	0.784	-0.611			
	(1.448)	(2.225)	(-0.716)			
$LBcomm$	-0.935	-0.927	-0.933			
	(-2.049)	(-2.248)	(-1.139)			
$LBir$	0.450	0.271	0.601			
	(1.472)	(1.408)	(0.839)			
$LBstock$	0.799	0.685	2.436			
	(1.574)	(1.642)	(2.547)			
Mkt				0.010	-0.007	-0.031
				(0.599)	(-0.953)	(-1.139)
SMB				-0.043	-0.012	-0.032
				(-1.784)	(-1.369)	(-0.708)
HML				0.051	0.018	0.096
				(1.685)	(1.701)	(1.702)
$dBaaBond$				0.387	0.078	0.250
				(2.364)	(1.531)	(0.929)
Constant	0.452	0.408	0.502	0.431	0.181	0.401
	(5.286)	(5.943)	(3.197)	(5.828)	(5.418)	(3.326)
Observations	211	211	211	211	211	211
R-squared	0.054	0.061	0.052	0.075	0.047	0.054

Table 18
Higher Moment Risks

This table reports the results of Fama-MacBeth regressions. The dependent variable in each case is the return on straddle positions in currencies. Explanatory variables are denoted as follows: IV is the current implied volatility, $BetaSkew$ and $BetaKurt$ denote estimated sensitivities to skew and kurtosis risks as described in the text. Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates.

	(1)	(2)
IV		-6.272
		(-6.085)
$BetaSkew$	0.251	0.329
	(3.687)	(4.885)
$BetaKurt$	0.837	0.921
	(4.612)	(5.491)
Observations	7,875	7,875
R-squared	0.404	0.548
Number of groups	910	910

Table 19
Excess Demand Regressions

This table reports the results of panel fixed effects regressions of straddle returns (columns headed 'Returns') and implied volatilities (column headed 'IV') on explanatory variables. Explanatory variables are denoted as follows: *Hedging* is the proportion of total open interest in currency futures held by hedgers, *LTA* denotes the average of the standardized values of the level of the TED spread, the level of the VIX, and the value of the St. Louis Fed financial distress indicator. Interaction terms between hedge demand and the limits to arbitrage proxy are also included. The regressions are at a weekly frequency with up to six currencies in the cross section. Newey-West standard errors are used to compute the t-statistics reported in parentheses beneath the parameter estimates. Time and currency fixed effects are included as noted in the table.

	(1) **Returns**	(2) **Returns**	(3) **IV**	(4) **Returns**
$Hedging$	-0.428	-0.494	0.019	
	(-1.957)	(-1.955)	(1.109)	
LTA	0.602			
	(3.370)			
$Hedging * LTA$	-0.916	-0.817	0.035	
	(-3.497)	(-5.120)	(2.106)	
IV_{fit}				-23.985
				(-3.526)
IV_{res}				-3.109
				(-2.359)
Constant	0.371	0.636	0.077	0.618
	(2.583)	(1.962)	(2.684)	(1.957)
Currency FE	Yes	Yes	Yes	Yes
Time FE	No	Yes	Yes	Yes
Observations	5,376	5,376	5,406	5,376
R-squared	0.010	0.378	0.778	0.411

Figure 1. Cumulative Returns to IV-rank-based LMH Straddle Strategy

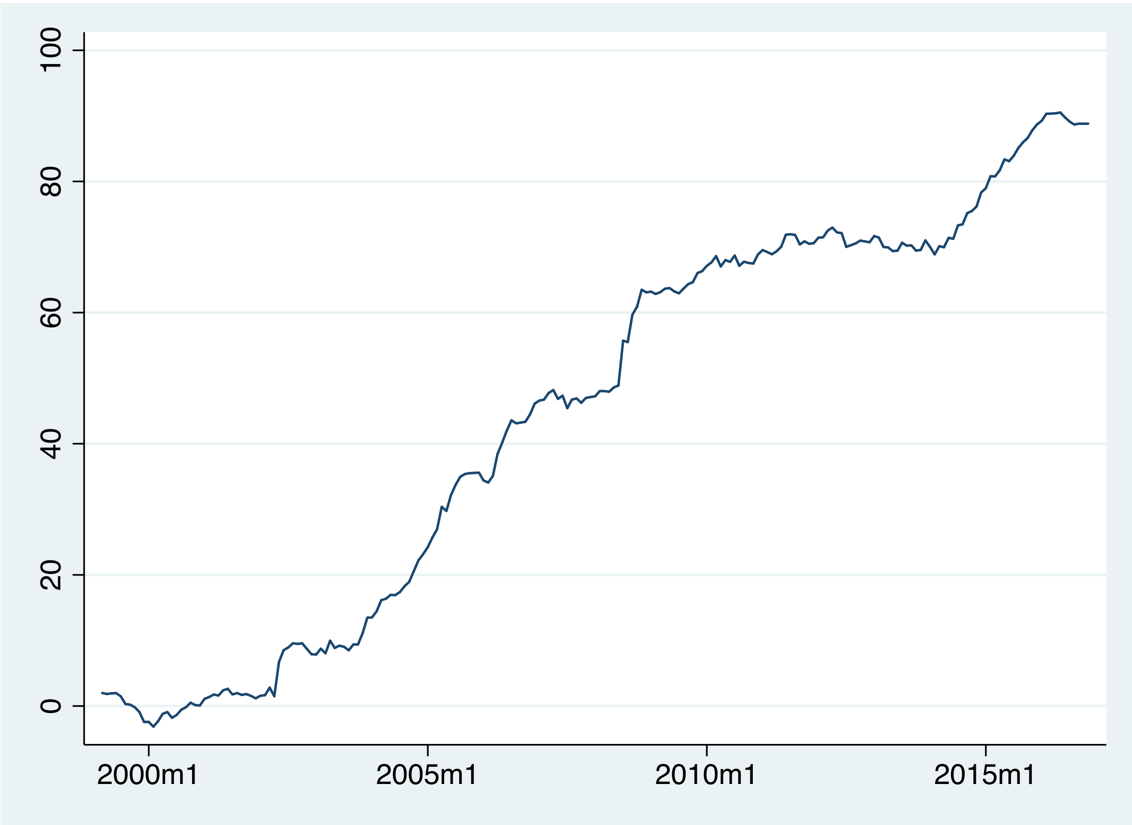

Figure 2. Distributions of Returns to IV-rank-based LMH Straddle Strategies with Different Levels of Diversification

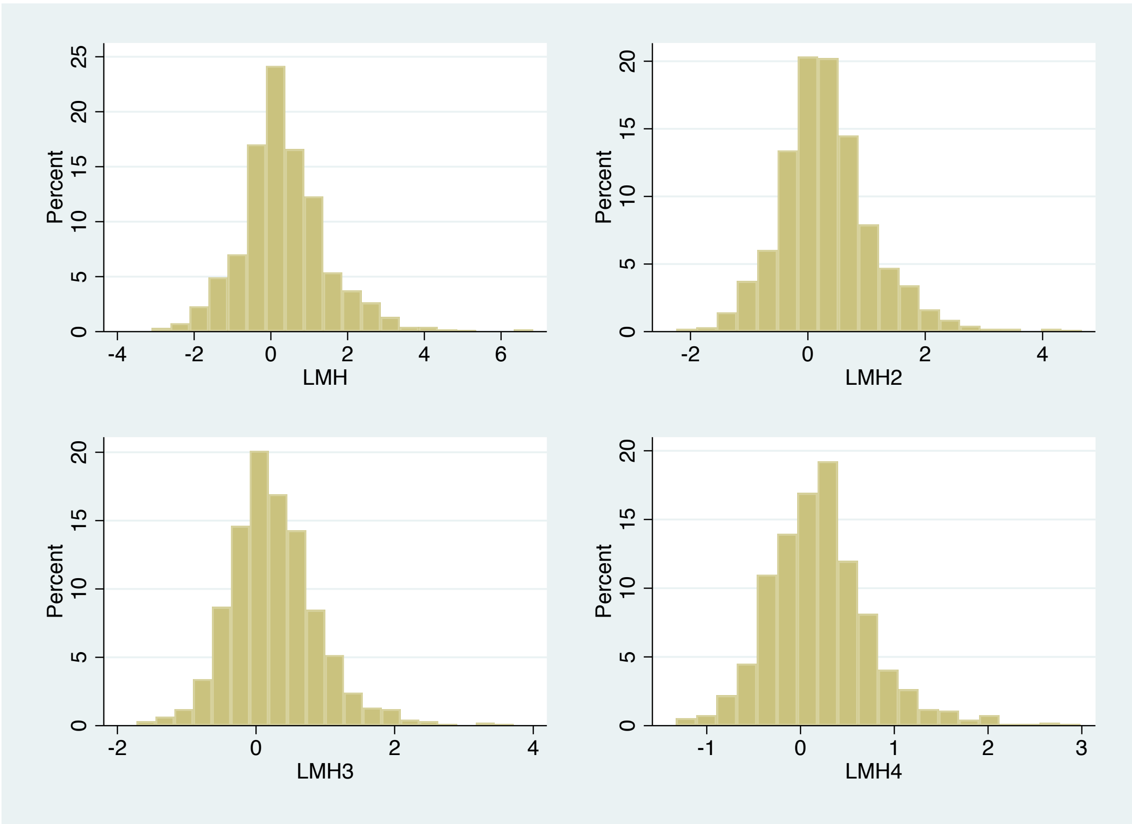

Figure 3. Portfolio Allocation by Currency Based on IV Ranking

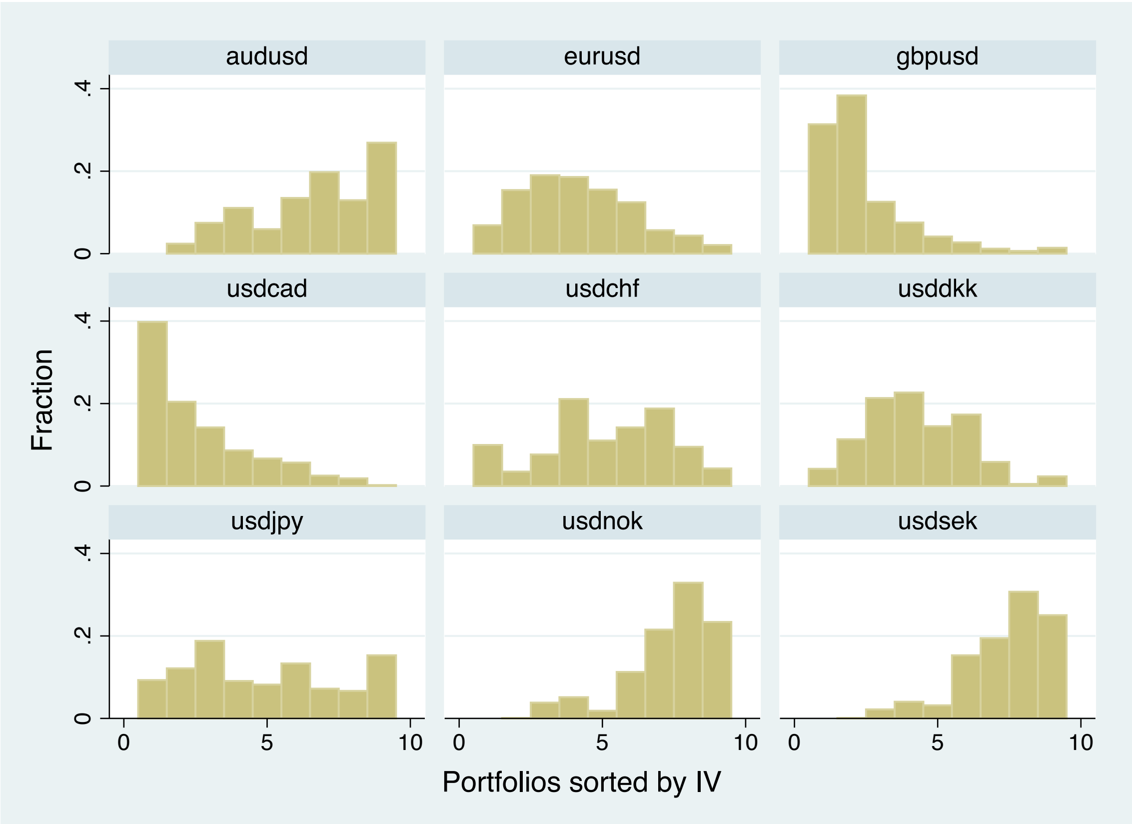

Figure 4. Moneyness at Expiry and Realised Volatility

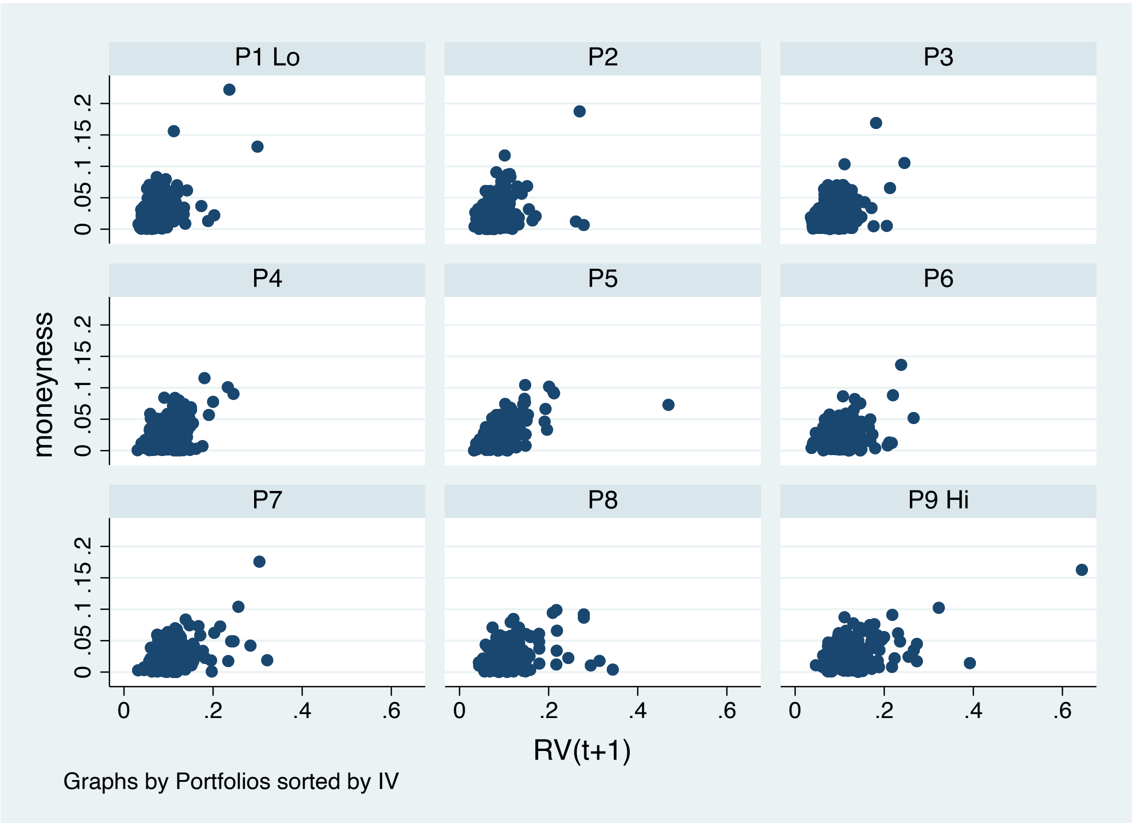

Figure 5. LMH **Straddle Returns, 12 Month Moving Averge**