Exchange Rates and Liquidity Risk

Evans, Martin

Georgetown University Department of Economics

31 August 2020

Exchange Rates and Liquidity Risk

Martin D.D. Evans*

August 31, 2020

First Draft: Comments Welcome

Abstract

I use Forex trading data to study how risks associated with the lack of liquidity contribute to the dynamics of 17 spot exchange rates through their time-varying contributions to risk premia. I find that liquidity risk matters. All the foreign exchange risk premia compensate investors for exposure to liquidity risk; and, for many currencies, exposure to liquidity risk appears to be more important than exposure to the traditional carry and momentum risk factors. I also find that variations in the price of liquidity risk make economically important contributions to the behavior of individual foreign currency returns: they account for approximately 34 percent, on average, of the variability in currency returns compared to the contribution of approximately 8 percent from the prices of carry and momentum risk.

Keywords: Foreign Currency Trading, Liquidity, Returns, Risk Premia, and Risk Factors.

JEL Codes: F3; F4; G1.

*Georgetown University, Department of Economics, Washington DC 20057. Tel: (202) 687-1570 email: evansm1@georgetown.edu. This research was undertaken independently without compensation.

Introduction

Foreign currencies are widely viewed as trading in a highly liquid market that is characterized by large trading volumes and low transaction costs. However, in reality, actual trading conditions varying considerably across currency pairs, trading venues, and instruments. For example, wholesale spot trading among major currencies such as the US Dollar, the Euro, British Pound, and Japanese Yen takes place under very different conditions than trading in minor currencies like the Russian Ruble or Turkish Lira. These differences are understood by professional traders and guide their trading decisions, but they do not figure prominently in recent academic research on currency returns. In particular, a sizable literature has developed studying the returns on portfolios comprising a large number of currencies that largely overlooks the cross-currency differences in trading conditions. This paper explores whether variations in liquidity across currencies and time affect the behavior of currency returns. In particular, I use high-frequency trading data from the foreign exchange (FX) market to examine whether risks associated the different liquidity measures contribute to the dynamics of 17 spot exchange rates through their time-varying contributions to foreign currency risk premia.

My analysis proceeds in three steps. In the first, I construct three different measures of market liquidity from an electronic trading platform, a limit order book, used by professional FX traders. The measures are based on the spread between the best bid and ask prices available on the limit order book, the depth of limit orders, and the volatility of prices, all measured at the 30-second frequency. These high-frequency measures are aggregated to monthly series. In the second step, I construct portfolios of currency returns sorted on the different liquidity measures as well as traditional carry and momentum risk factors. I then construct a no-arbitrage pricing model with a stochastic discount factor that accurately accounts for the behavior of these factor-sorted portfolio returns and delivers a beta representation for the risk premia on individual foreign currencies. In the final step, I use the beta representation to estimate the exposure of individual currencies to liquidity and other risk factors and quantify how changes in the factor risk prices contribute to the variability of risk premia and individual spot rates.

The main finding to emerge from this analysis is that liquidity risk matters. More precisely, I show that the risk premia on all the 17 foreign currencies compensate investors for exposure to liquidity risk, measured by the betas on one or more of the three liquidity risk factors. This finding applies to the risk premia on major currency pairs (e.g., EURUSD, JPYUSD, and GBPUSD) that are widely considered to trade in highly liquid markets. Furthermore, for many currencies, exposure to liquidity risk appears to be more important than exposure to the traditional carry and momentum risk factors that have been the focus of earlier research. I also find that variations in the price of liquidity risk make economically important contributions to the behavior of individual foreign currency returns. These variations account for approximately 34 percent, on average, of the variability in currency returns compared to the contribution of approximately 8 percent from variations in the prices of carry and momentum risk.

My findings arise because it is possible to construct profitable trading strategies that involve borrowing in some currencies and lending in others based on the liquidity of spot trading. These liquidity-based strategies are analogous to those that characterize the carry trade, expect they use measures of spot trading liquidity rather than interest rates to choose the borrowing and lending currencies. I argue that these strategies are profitable because the liquidity of spot trading embeds information about the future behavior of spot rates that is not incorporated into forward prices. In particular, when the fear of a foreign currency crash rises among spot traders, there will be changes in the structure of limit orders that reflect the perceived increased risk of supplying liquidity. These changes are only partially reflected in the prices of forward contracts because they are determined by adding forward/swap points to the best limit prices. Importantly, forward points are determined by money market conditions not liquidity conditions in spot currency trading. Consequently, an increase in the risk of a currency crash that only lowered depth in the spot market would have no effect on forward prices. So, if a reduction in depth is a precursor of a future fall in the price of foreign currency, a strategy of selling foreign currency forward funded by future spot purchases will make a profit on average. Conversely, if crash risk falls below the norm and depth increases, a strategy of buying foreign currency forward and selling in the spot market will make a profit on average.[1] The data shows that a depth-based portfolio strategy combining both these elements produces a positive average return over the 10 year sample period.

This is not the first paper to suggest that liquidity risk affects the behavior of currency returns. Brunnermeier, Nagel, and Pedersen (2008) use the CBOE VICX and LIBOR spreads to examine the link between reductions in liquidity funding and losses on carry trades. In a similar vein, Menkhoff, Sarno, Schmeling, and Schrimpf (2012a) use an aggregate measure of foreign currency bid-ask spreads, the TED spread, and an equity-based liquidity measure (from Pástor and Stambaugh, 2003) to proxy for global liquidity risk in the foreign currency trading. Banti, Phylaktis, and Sarno (2012) study a global liquidity measure from the estimated price-impact of order flow on individual currency returns. The returns to the carry trade are also linked to a global liquidity measure constructed from "noise" in the US Treasury Market by Hu, Pan, and Wang (2013). A key feature distinguishing my analysis from these papers is that I focus on differences in liquidity across currency pairs rather than an aggregate economy-wide measure of liquidity. This approach has two important benefits. First, it directly ties trading conditions for particular currency pairs to the behavior of returns. Second, it allows me to identify how different liquidity measures contribute to systematic

[1] Under covered interest parity, these strategies are equivalent to (i) borrowing foreign currency to lend domestically, and (ii) borrowing domestically to lend in foreign currency, respectively.

risk through the construction of liquidity-sorted portfolios.[2]

Research in market microstructure has long recognized that liquidity is a multifaceted concept that cannot be represented by a single variable. Alternative measures considered in the literature include the price-impact of order flow as in Kyle (1985), Evans and Lyons (2002), and Banti, Phylaktis, and Sarno (2012); return reversal in Campbell, Grossman, and Wang (1993), Pástor and Stambaugh (2003), and Mancini, Ranaldo, and Wrampelmeyer (2013); and price dispersion in Chordia, Subrahmanyam, and Anshuman (2001). I use three measures constructed with high-frequency data from an electronic limit order book used by professional foreign currency traders. The first measure is the spread between the best bid and ask prices on the order book.[3] The second is the depth of the limit orders at the best bid and ask prices. I also use intraday price volatility computed as the standard deviation of price changes over 30-second intervals.[4]

Earlier research on currency returns (cited below) has used bid and ask prices to more accurately represent the monthly returns investors would receive on actual currency positions. These bid and ask prices are typically based on benchmark prices computed at 4:00 pm (London time), known as the WMR Fix. In contrast, I compute daily and monthly averages of the spreads between bid and ask prices sampled every 30 seconds. These aggregate spread measures are likely to be more informative about trading conditions than the benchmark spreads at the end of each month.[5] As a complementary liquidity measure, I also compute daily and monthly averages of (top-of-book) depth sampled every 30 seconds. Variations in bid-ask spreads can be a poor indicator of changing liquidity when there is very little depth at the top of the limit order book because the execution prices for all but the smallest trades will differ from the best bid or ask prices. One important feature of my data is that it comes from a trading platform that aggregates limit prices from multiple sources, so both the spread and depth liquidity measures are representative of trading conditions for a currency pair across a variety of trading venues. I used intraday volatility to capture another aspect of liquidity, the price impact of order flow. When the flow of incoming market orders are matched with outstanding limit orders beyond the top of the order book and those limit orders are not immediately replaced, order flow induces price volatility. Of course, limit prices can change for other reasons (such as

[2]Of course, changes in aggregate liquidity could have differential impacts on the trading conditions for individual currency pairs, so my analysis also accommodates variations in aggregate liquidity. Banti, Phylaktis, and Sarno (2012) also examine portfolio returns constructed according to the exposure to their measure of global liquidity. The effects of changing aggregate liquidity on stock returns have been studied by Pástor and Stambaugh (2003), Acharya and Pedersen (2005), Korajczyk and Sadka (2008) and Hasbrouck (2009), among others.

[3]Earlier research on bid-ask spreads in the FX market include Bessembinder (1994), Bollerslev and Melvin (1994), and Hsieh and Kleidon (1996).

[4]Consistent with the link between liquidity and volatility suggested by Copeland and Galai (1983), Bollerslev and Melvin (1994) found a significant positive relationship between the bid-ask spread and exchange rate volatility in the interbank market trading of Deutsche mark-US dollar. Melvin and Taylor (2009) study FX market liquidity during the Global Financial Crisis.

[5]See Evans (2018) and Evans, O'Neill, Rime, and Saakvitne (2018) for comparisons of trading at the WMR Fix with other time periods.

the arrival of news), so intraday volatility is an imperfect measure of price-impact. It nevertheless provides another source of information about trading conditions in particular currency pairs.[6]

My analysis extends research exploring the sources of foreign currency risk through the analysis of factor-sorted portfolios. This literature originates with the work of Lustig and Verdelhan (2007) and Lustig, Roussanov, and Verdelhan (2011) who studied the properties of portfolios sorted on interest rates that emulated a version of the carry trade.[7] Subsequent research by Menkhoff et al. (2012b) and Della Corte, Ramadorai, and Sarno (2016) extended this line of research to momentum-sorted portfolios. I include portfolios sorted on both interest rates and momentum in my analysis to reduce the chance that the liquidity measures I focus on are not proxying for other omitted risk factors. However, my results show that carry and momentum are less important risk factors driving currency returns than were found in earlier work. Indeed, consistent with the findings in Burnside (2019), interest-sorted portfolios emulating the carry trade produce much less impressive returns in my sample period that includes years following the 2007-2008 financial crisis, than was found in pre-crisis data. Thus, the importance of liquidity risk in my data does not provide an explanation for the returns to the carry trade along the lines suggested by Burnside, Eichenbaum, and Rebelo (2011), since it appears that these returns have diminished.

One other feature of my analysis deserves to be mentioned. In contrast to much of the literature, I do not focus on the role of liquidity or other risk factors in accounting for the cross-currency pattern of unconditional expected returns. Because the 2007-2008 financial crisis covers a sizable portion of my data sample, it is unclear that average returns from this sample are particularly reliable estimates of unconditional expected returns. This makes testing the cross-sectional implication of a no-arbitrage pricing model a challenge. My focus, instead, is on the implications of the model for the time-series behavior of currency risk premia and returns. In so doing I attempt to exploit the changes in foreign currency trading conditions associated with the 2007-2008 financial crisis.

The remainder of the paper is organized as follows. Section 1 describes the data, the construction of the liquidity measures, and the characteristics of the factor-sorted portfolios. Section 2 develops the no-arbitrage model that identifies how the various risk factors determine individual foreign currency risk premia. Section 3 presents the empirical results. Section 4 provides concluding comments.

[6] My use of volatility is foreshadowed by prior research. In particular, Della Corte, Ramadorai, and Sarno (2016) consider portfolios sorted on volatility, while Menkhoff et al. (2012a) show that a global volatility proxy contains important information which can be used to price returns of carry trade portfolios.

[7] Other papers studying the properties of interest-rate sorted portfolios of excess currency returns include: De Santis and Fornari (2008), Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2009), Verdelhan (2010), Burnside, Eichenbaum, and Rebelo (2011), Christiansen, Ranaldo, and Söderlind (2011), Gilmore and Hayashi (2011), Hassan et al. (2012), Menkhoff et al. (2012a), Menkhoff, Sarno, Schmeling, and Schrimpf (2012b), Mueller, Stathopoulos, and Vedolin (2017), Gavazzoni, Sambalaibat, and Telmer (2013), Hu, Pan, and Wang (2013) Jurek (2014), Lettau, Maggiori, and Weber (2014), Daniel, Hodrick, and Lu (2014), and Dobrynskaya (2014).

1 Data

My analysis uses intraday trading data on 17 currency pairs spanning the period from January 1st, 2006 until December 31st, 2015. The trading data comes from the Hotspot FX trading platform, which is an electronic brokerage widely used by professional Forex traders, owned by Cboe Global Markets. The trading platform aggregates quotes from a large number of banks and other financial institutions into a limit order book that provides uses of the platform with the best (tradable) bid and offer prices at which they can sell and buy the base currency. The raw data contains four series for each currency pair sampled at the 30-second frequency: the best bid price, the best offer price, the depth of bids at the best bid price, and the depth of offers at the best offer price. I use these intraday data to construct different measures of liquidity and price dynamics at a monthly frequency for use in my analysis.[8] I also make use of daily forward FX prices from Bloomberg.

My analysis makes use of four monthly measures constructed from the intraday trading data: average depth, average spread, price trend, and price volatility. To describe the construction of these measures, let P_i^{B} and P_i^{O} denote the best bid and offer prices at instant i, with corresponding depths D_i^{B} and D_i^{O}. For each day n in the sample, I compute the average daily depth $D_n = \frac{1}{2880}\sum_{i=1}^{2880}(D_i^{\text{B}} + D_i^{\text{O}})$ over the 2880 30-second intervals in the 24 hours ending at 16:00 hours (London time) on day n. Similarly, I compute the average daily spread, expressed in basis points, as $sprd_n = \frac{10000}{2880}\sum_{i=1}^{2880}(\ln(P_i^{\text{O}}) - \ln(P_i^{\text{B}}))$. I also construct two measures of the intraday dynamics in mid-point prices $P_i = \frac{1}{2}P_i^{\text{B}} + \frac{1}{2}P_i^{\text{O}}$. The first measure is the daily trend, expressed in basis points, as $\mu_n = \frac{1}{2880}\sum_{i=1}^{2880}\Delta p_i$, where $\Delta p_i = 10000(\ln(P_i) - \ln(P_{i-1}))$. The second measure is the daily standard deviation in price changes, $\sigma_n = (\frac{1}{2880}\sum_{i=1}^{2880}(\Delta p_i - \mu_n)^2)^{1/2}$. The monthly series for the depth, the spread, the price trend, and price volatility are computed by averaging each daily measure between the last trading days of each month. Thus, the depth measure for month t, D_t is computed as the average of the D_n's between the day after the last trading day on month $t-1$, and the last trading day of month t. The monthly series for the spread $sprd_t$, price trend μ_t, and price volatility σ_t, are computed analogously from $sprd_n$, μ_n, and σ_n, respectively.

Spot rates, forward rates, and forex returns are computed from end-of-month prices. In particular, the month t spot rate S_t is identified by the mid-point price P_i at 16:00 hours on the last trading day of month t. Similarly, I construct one month forward prices F_t from the mid-point of the bid and offer forward prices reported by Bloomberg on the last trading day of month t. I follow the standard academic practice of defining spot and forward rates in terms of the U.S. dollar price of foreign currency.[9] Under this definition,

[8] Many empirical studies of FX trading use data from either the EBS or Reuter's electronic trading systems. One advantage of the Hotspot data is that it provides representative information on prices and liquidity across a wide number of currency pairs. In contrast, trading on the EBS and Reuters systems is concentrated in different currency pairs, so one would have to aggregate information from both systems to obtain representative information on liquidity.

[9] By market convention, some currencies are quoted on Hotspot in terms of the foreign currency price of U.S. dollars. For

a rise (fall) in S_t represents a depreciation (appreciation) in the U.S. dollar and the log excess return on a foreign currency position between the end of month t and $t+1$ is $er_{t+1} = \ln S_{t+1} - \ln F_t$. I also make use of the contemporaneous difference between the log spot rate and the log forward rate, $fd_t = \ln S_t - \ln F_t$, which I term the forward discount. Under covered interest parity (CIP), fd_t is the difference between the logs of the foreign and U.S. short-term interest rates, so I refer to foreign currencies selling at a forward discount (i.e. $fd_t > 0$) as high-interest currencies vis-a-vis the U.S. dollar.[10]

Table 1 presents summary statistics for the monthly variables used in the analysis. For readability, the table reports statistics for three groups of the currency pairs:[11] Group A comprises the so-called major currencies. The mean excess returns and forward discounts for these currency pairs were generally smaller (in absolute value) than in Groups B and C. Group A currencies also appear more "liquid" insofar as their mean depth is (generally) larger and mean spreads are smaller. Notice, however, that the standard deviation of depth is much larger for the EURUSD and JPYUSD than all the other currency pairs, while the standard deviation for the spreads are much smaller. This suggests that time-series variations in "liquidity" for Group A currencies are characterized more by varying depth than changing spreads, whereas for Group C currencies, variations in "liquidity" appear more in the form of varying spreads (particular for the RUBUSD and PLNUSD). The right-hand columns of the table show how monthly price trends and intraday volatility vary across the currency pairs. For Group A currency pairs, mean trends are generally small and the standard deviations are all close to 3.4 percent. Across the Group B and C currency pairs, there is greater heterogeneity in both trends and their variability. At one extreme, the HKDUSD depreciation rate has a mean and standard deviation of 0.020 and 0.145, while at the other extreme the RUBUSD rate has a mean and standard deviation of -1.251 and 9.055. The mean value for intraday volatility σ_t are close to one basis point for most currencies; with the exceptions of the HKDUSD, CNHUSD, and RUBUSD. Month-by-month variations in average intraday volatility are identified by the standard deviation in σ_t. Again the RUBUSD stands out as a currency pair exhibiting large month-by-month changes in intraday volatility.

these currency pairs, monthly spot and forward rates S_t and F_t are computed from the reciprocals of mid-point spot and forward prices, and the daily price trends are computed as $\mu_n = -\frac{1}{2880}\sum_{i=1}^{2880}\Delta p_i$.

[10]Note that this is purely for pedagogical convenience. I am not assuming that CIP holds continuously over the sample period; see, e.g., Du, Tepper, and Verdelhan (2018) for evidence to the contrary.

[11]Individual pairs are identified by their ISO standard abbreviations: AUD=Australian Dollar,CAD=Canadian Dollar, CHF=Swiss Franc, CNH=Chinese Yuan, EUR=Euro, GBP=British Pound, HKD=Hong Kong Dollar, JPY=Japanese Yen, MXN=Mexican Peso, NOK=Norwegian Kroner, NZD=New Zealand Dollar, PLN=Polish Zloty, RUB=Russian Rouble, SEK=Swedish Krona, SGD=Singapore Dollar, TRY=Turkish Lira, USD=United States Dollar, and ZAR=South African Rand.

Table 1: Summary Statistics

	N	Excess Returns		Forward Discount		Depth		Spread		Trend		Volatility	
Currency Pair		mean	std	mean	std	mean	std	mean	std	mean	std	mean	std
EURUSD	120	-0.727	38.382	0.406	1.898	6.933	5.497	1.145	1.524	-0.067	3.713	0.984	0.386
JPYUSD	120	-0.260	32.997	-0.014	0.019	6.841	4.877	1.586	2.107	0.024	3.645	1.082	0.438
CHFUSD	120	1.311	40.328	-1.151	1.267	5.314	2.854	1.982	1.361	0.298	3.756	1.097	0.442
GBPUSD	120	-2.985	31.924	-1.079	2.840	4.240	2.242	1.574	1.455	0.089	3.185	0.938	0.436
										0.086	3.575		
AUDUSD	120	-2.863	50.655	-2.450	1.532	5.468	3.019	1.983	1.761			1.397	0.699
CADUSD	120	-1.681	36.100	0.281	0.725	4.927	2.501	2.590	2.886	0.002	3.009	1.025	0.423
HKDUSD	612	0.060	1.188	-0.198	0.263	4.882	3.301	0.501	0.441	0.020	0.145	0.042	0.027
NOKUSD	120	-1.766	42.882	1.071	1.444	4.160	1.815	8.607	5.196	-0.111	3.673	1.392	0.558
NZDUSD	120	-1.686	53.855	-1.663	0.662	4.379	2.191	4.142	3.251	-0.155	5.830	1.558	0.685
SEKUSD	120	-0.929	42.998	0.130	1.465	4.342	2.124	7.538	3.787	0.260	4.468	1.348	0.548
SGDUSD	88	1.533	21.518	-0.272	1.153	3.810	1.605	3.527	4.464	0.236	2.488	0.594	0.198
CNHUSD	24	1.635	11.052	3.099	1.564	3.613	1.311	1.870	1.634	-0.005	0.961	0.207	0.129
MXNUSD	116	-0.685	39.344	3.675	1.851	2.784	1.153	11.000	17.679	-0.405	2.863	0.964	0.522
RUBUSD	44	-17.707	83.841	9.132	4.863	3.353	2.759	75.297	168.845	-1.251	9.055	2.350	5.007
TRYUSD	86	-4.698	46.182	8.571	2.969	2.410	0.933	52.764	99.021	-0.539	3.611	1.352	0.585
PLNUSD	112	-1.021	69.911	2.128	1.805	1.854	1.334	216.024	630.594	0.096	3.819	1.102	0.833
ZARUSD	80	-8.983	54.689	0.396	0.237	2.489	0.997	21.025	24.150	-0.732	5.779	1.573	0.488

Notes: Summary statistics for monthly variables over the 10 year sample period: Jan 1st 2006 through Dec 31st 2015, subject to data availability. Statistics are reported for each of the 17 currency pairs listed in the left-hand column. Individual pairs are identified by their ISO standard abbreviations. The number of monthly observations for each currency pair is shown in the column headed N. The table reports the sample mean and standard deviation (std.) for log excess returns $er_t^i \times 1200$, forward discount $fd_t \times 1200$, depths D_t (millions of USD), spreads $sprd_t$ (basis points), trends $\mu_t \times 2880 \times 365/1200$ (percent per month), and volatility σ_t (basis points).

1.1 Risk Factors

I identify the risk factors driving excess returns on individual foreign currencies by building monthly portfolios of currencies sorted by forward discounts, price trends, intraday volatility, depth, and spreads. The portfolio for factor x_t is constructed as follows. At the end of month t, I allocate all the currency pairs to three portfolios based on the rank of the factor x_t for each currency pair. For the case where the factor is the forward discount, the four currency pairs with the highest implied foreign interest rate are assigned to the high portfolio, and the four pairs with lowest foreign interest rate are assigned to the low portfolio. The remaining currency pairs are assigned to the medium portfolio. In the case of the two liquidity measures, I assign the four currency pairs with the greatest liquidity, measured either by the largest depth or narrowest spread, to the high portfolio, and the four pairs with the least liquidity (smallest depth or widest spread) to the low portfolio. With respect to the price trend and volatility factors, the high portfolio comprises currency pairs with highest price trend or least volatility, whereas the low portfolio comprises currency pairs with the lowest price trend or highest volatility.

Portfolio returns are computed from the average of the excess returns on the individual currencies assigned to each portfolio at the end of month $t+1$. For example, the log excess return on the high portfolio are $er^{\text{H}}_{t+1} = \ln\left(\frac{1}{4}\sum_{i=1}^{4}\exp(er^{i}_{t+1})\right)$ where er^{i}_{t+1} are the log excess returns for currency pair i assigned on the high portfolio at the end of month t. Log excess returns on the low portfolio, er^{L}_{t+1}, are identified in an analogous manner. Portfolios are rebalanced at the end of each month. I will use the difference between log returns on the high and low portfolios $er^{\text{X}}_{t+1} = er^{\text{H}}_{t+1} - e^{\text{L}}_{t+1}$ for each of the five factors X in the model developed below.[12]

Table 2 reports summary statistics for the factor portfolios. Panel A reports the mean log excess return for the high and low portfolios associated with each of the five factors, while Panel B reports statistics for the difference in log excess returns $er^{\text{X}}_{t+1} = er^{\text{H}}_{t+1} - e^{\text{L}}_{t+1}$. All excess returns are measured in annual percentage points. As the table shows, there are sizable differences between the mean returns on the high and low portfolios for several of the factors. In particular, the mean return differences for the depth and price trend factors are approximately 3.5 and 3.9 percent (with p-values for the null of a zero mean of 0.08 and 0.02), respectively. To interpret these results, notice that er^{X}_{t+1} is approximately equal to the log excess return from a strategy that borrows in the four foreign currencies assigned to the low portfolio and lends in the four foreign currencies assigned to the high portfolio. So, according to the statistics in Table 2, a strategy of

[12] Since the Hotspot data only contains information on 17 USD currency pairs, I only construct three portfolios for each factor. By comparison, Lustig et al. (2011) use 26 currencies (after the introduction of the euro) to construct six portfolios, and then use the difference between the returns on portfolios one and six as factor returns. These factor returns contain approximately the same number of individual currency pairs as I use here when constructing $er^{\text{X}}_{t+1} = er^{\text{H}}_{t+1} - e^{\text{L}}_{t+1}$. Section 3.4 discusses the robustness of my results to different methods for constructing the factor portfolios.

Table 2: FACTOR PORTFOLIO RETURNS

Portfolio Return and Composition		Factors: fd_t	D_t	$sprd_t^{-1}$	μ_t	σ_t^{-1}
A:						
high: er^{H}_{t+1}	mean	-0.626	-0.125	-1.330	1.426	-1.235
	std	(3.546)	(2.796)	(2.195)	(3.098)	(2.482)
	turnover	60.0%	75.0%	50.0%	95.8%	70.0%
	duration	4.707	5.286	9.005	1.488	3.702
low: er^{L}_{t+1}	mean	-2.211	-3.666	-2.247	-2.477	-2.451
	std	(3.340)	(3.745)	(3.761)	(3.307)	(3.610)
	turnover	32.5%	78.3%	48.3%	99.2%	71.7%
	duration	6.538	3.003	5.445	1.766	3.110
B:						
Difference						
$er^{\text{X}}_{t+1} = er^{\text{H}}_{t+1} - e^{\text{L}}_{t+1}$	mean	1.585	3.541	0.917	3.904	1.216
	std	(2.190)	(2.394)	(2.823)	(1.864)	(2.193)
	pval	0.213	0.076	0.401	0.018	0.294
C:						
Correlation Matrix						
	fd_t		0.000	0.000	0.009	0.654
	D_t	-0.634		0.000	0.001	0.299
	$sprd_t^{-1}$	-0.526	0.713		0.032	0.000
	μ_t	-0.240	0.302	0.197		0.185
	σ_t^{-1}	0.042	0.096	0.371	-0.122	

Notes: Panel A reports the mean log excess return (measured in annual percent) on the high and low portfolios sorted by each of the factors shown in the heading of each column: the forward discount fd_t, depth D_t, the reciprocal of the spread $sprd_t^{-1}$, the price trend μ_t, and the reciprocal of volatility σ_t^{-1}. Standard errors for the mean log excess return are reported in parenthesis. The turnover statistics report the fraction of the months in the sample period for which there is a change in at least one currency in the high/low portfolios. The duration statistics report the mean duration in months of a currency pair in a portfolio, averaged across all 17 currency pairs. Panel B reports the mean and standard error for the difference between the log excess return on the high and low portfolios, er^{HL}_{t+1}, together with the p-value for the null that the difference is equal to zero. Panel C reports the correlation matrix for er^{HL}_{t+1}, correlations below the leading diagonal and p-values for the null of a zero correlation above the diagonal.

borrowing in the four foreign currencies with the least depth and lending in the four with the most depth, on average produces a return of 3.5 percent. Similarly, a strategy of borrowing in the four foreign currencies that were depreciating most against the US Dollar in the prior month and lending in the four currencies that were appreciating most against the US Dollar in the prior month, generate returns of 3.9 percent on average. This is a type of momentum strategy similar to those studied by Burnside (2011), Lustig, Roussanov, and Verdelhan (2011), and Menkhoff et al. (2012b).

Strategies based on the other three factors also produced positive returns on average, but they are smaller and are not statistically significant. In particular, a carry-trade strategy of borrowing in the four foreign currencies with the lowest interest rates and lending in the four with the highest rates, on average produces a return of 1.59 percent. This estimate is smaller than the average returns on interest-rate sorted portfolios reported in earlier research (see, e.g., Lustig and Verdelhan, 2007, Lustig, Roussanov, and Verdelhan, 2011), but it is consistent with more recent results reported by Burnside (2019). He finds that returns on similarly constructed portfolios declined after the 2007-2008 recession, which covers much of the sample period studied here. Strategies based on the spread and intraday volatility also generate smaller mean returns. When the spread is used to select low and high liquidity currencies for borrowing and lending, the mean return is 0.9 percent. The mean return from lending to low volatility currencies with funds borrowed from high volatility currencies gives a mean return of 1.2 percent.

Table 2 also provides information on the changing composition of the factor portfolios. The turnover statistics report the fraction of the sample period where there is a change in at least one of the currency pairs assigned to either the high of low portfolios. These statistics range from 33 to 99 percent. The duration statistics report the mean number of months an individual currency pair remains in either the high or low portfolio, averaged across the 17 currency pairs in the sample. It is clear from both of these statistics that the composition of the factor portfolios varies considerably from month-to-month.

Panel C of Table 2 shows the correlations between the different factors returns. The largest correlations appear between the liquidity factors, depth and spread, and between the forward discount and liquidity factors. The correlation between the returns on the depth and spread portfolios is 0.71, which is highly statistically significant. It seems that both depth and spread provide information on the liquidity of a currency pair, but neither measure is completely informative. The correlations between the returns on the liquidity measures and the forward discount are large, negative, and statistically significant. This means that the carry trade strategy is to some extent borrowing in high liquidity currencies where foreign interest rates are low and lending in low liquidity currencies where foreign interest rates are high. Since interest rates and liquidity are both endogenous, these findings are insufficient to argue whether the returns on carry-trade or liquidity strategies have a more straightforward structural interpretation, but I will return to this question below. Finally, Panel C also shows that the correlations between the other factor returns are relatively small.

Table 3: Factor Return ARMA Models

	AR Coefficients			MA Coefficients					Tests	
Factor	ϕ_1	ϕ_2	ϕ_3	θ_1	θ_2	θ_3	σ^2	R^2	Sig	ARCH
Interest Diff.	-0.481** (0.216)	-0.506*** (0.141)		0.375** (0.164)	0.766*** (0.113)		489.208*** (46.505)	0.116	64.925 (0.000)	0.056 (0.812)
Depth	1.225*** (0.337)	-0.231 (0.374)	-0.225** (0.111)	-1.242*** (0.354)	0.297 (0.340)		558.820*** (55.786)	0.163	631.339 (0.000)	0.047 (0.828)
Spread	0.619* (0.382)	0.153 (0.361)		-0.613* (0.360)	-0.052 (0.368)	-0.239*** (0.086)	855.712*** (95.023)	0.059	102.275 (0.000)	0.414 (0.520)
Trend	-0.369* (0.209)	0.162 (0.158)		0.549*** (0.206)	-0.023 (0.220)	0.313*** (0.107)	349.496*** (36.379)	0.136	139.000 (0.000)	0.565 (0.452)
Volatility	0.773*** (0.155)			-0.893*** (0.171)	0.246*** (0.084)		521.019*** (68.121)	0.075	34.784 (0.000)	0.650 (0.420)

Notes: Table reports estimates of ARMA models for factor returns

$$(1-\phi_1 L-\phi_2 L^2-\phi_3 L^3)(er^{\text{HL}}_{t+1}-\overline{er}^{\text{HL}}_{t+1})=(1+\theta_1 L+\theta_2 L^2+\theta_3 L^3)e_{t+1}$$

where $er^{\text{HL}}_{t+1}=er^{\text{H}}_{t+1}-e^{\text{L}}_{t+1}$ is the difference between the log excess return on the high and low portfolios for a particular factor. $\overline{er}^{\text{HL}}_{t+1}$ is the mean value for er^{HL}_{t+1} reported in Table 1. Estimates computed by maximum likelihood under the assumption that $e_{t+1}\sim i.i.d.N(0,\sigma^2)$ from 119 monthly observations; asymptotic standard errors are reported in parenthesis under parameter estimates. Statistical significance at the 10, 5, and 1 percent levels indicated by *, **, and ***, respectively. The column headed R^2 reports the variance ratio $\mathrm{V}(\hat{er}^{\text{HL}}_{t+1})/\mathrm{V}(er^{\text{HL}}_{t+1})$, where $\hat{er}^{\text{X}}_{t+1}$is the one-month-ahead forecast for er^{X}_{t+1} computed from the ARMA estimates. The column headed "Sig" reports Wald statistics for the null that that all the coefficients in the ARMA model equal zero. The column headed "ARCH" reports the LM statistics for first-order ARCH in the ARMA residuals. P-values for both tests are shown in parenthesis.

1.2 Return Predictability

The asset-pricing implications of the factor returns depend on both their unconditional averages reported in Table 2 and their time-series predictability. To quantify the degree of predictability, I estimated ARMA models for each of the five factor returns. Starting from an ARMA(3,3) specification, I followed a "testing-down" approach where statistically insignificant coefficients on the highest AR and MA terms are dropped until a parsimonious specification is found that adequately represents the autocorrelations in the factor returns. The chosen ARMA model estimates for each of the five factor returns are reported in Table 3. As the table shows, many of the estimated AR and MA coefficients are highly statistically significant. Indeed, Wald tests for the joint significance of all the AR and MA coefficients in each model, reported under the column headed “Sig” , produce P-values of less than 0.001 for all five models. To quantify the degree of predictability in the factor returns implied by the model estimates, I compare the variance of the one-month-ahead ARMA forecasts for the factor returns $\hat{er}^{\text{HL}}_{t+1}$ against the variance of the actual factor returns er^{HL}_{t+1}, with the R^2 statistic $\mathrm{V}(\hat{er}^{\text{HL}}_{t+1})/\mathrm{V}(er^{\text{HL}}_{t+1})$. The ARMA models for depth-sorted returns and trend-sorted returns have moderate forecasting power with R^2 statistics of 0.17 and 0.14, respectively. The forecasting power of the ARMA models for the other factor returns is more limited. The models for the interest differential, spread, and volatility factors produce R^2 statistics of 0.12, 0.06, and 0.07, respectively. Finally, the right-hand column of the table reports the results of LM tests for first-order ARCH in the ARMA model residuals. These tests reveal no statistically significant evidence of conditional heteroskedasticity, a finding I will make use of in my analysis below.

1.3 Why is Spot Liquidity a Risk Factor?

The empirical results in Table 2 and 3 suggest that different measures of liquidity in spot currency trading can be used to construct trading strategies that produce positive returns. Although these strategies are analogous to carry and momentum strategies that have been studied in earlier research, it seems surprising that choosing to borrow or lend in particular currencies based on the liquidity of spot trading could be an attractive proposition.

My proposed explanation is based on the microstructure of foreign currency trading. Currency dealers at major banks are typically split into two groups (or “desks”). One group specializes in spot trading; that is, executing orders from the bank’s customers and trading on behalf of the bank. These dealers continually monitor and participate on multiple trading venues that comprise the wholesale spot market and so will be aware of changing liquidity conditions. The second group of dealers focus on foreign currency forward and swap contracts. These dealers run an entirely separate trading book from the spot dealers and typically are housed in the bank’s money market department. This separation between the dealer groups is reflected in the way forward prices (swaps) are quoted; namely in terms of points that are added or subtracted from spot prices to obtain outright forward rates. Market practice is to quote points based on the prevailing rates for borrowing and lending in money markets because these rates determine the terms for swap contracts that trade in far greater volume than outright forward contracts.[13]

The key implication of these observations is that variations in the supply of liquidity for spot trading, which are captured by changes in the structure of limit orders on the wholesale trading venues, will not

[13] According to the BIS survey (Bank of International Settlements, 2016), trading in outright forwards and swaps accounted for 14 and 47 per cent of daily trading volume, respectively.

be fully reflected in the prices of forward contracts. In particular, changes in the depth of limit orders or the structure of bid and ask prices below the top of the order book, will not move forward prices if there is no change in the best bid and ask prices at the top of the book. This means that changes in the supply of liquidity that reflect revisions in traders' views about the future behavior of spot prices need not affect forward prices even if traders' views are correct on average. In sum, therefore, changes in the structure of limit orders may contain more information about the future behavior of spot prices than is reflected in the variations of forward prices.

The results in Table 2 suggest that depth represents a particularly important source of the extra information about future spot prices. One likely explanation for this is that depth falls when traders perceive that the risk of a foreign currency crash (characterized by an abrupt and substantial fall in the price of foreign currency) rises. So if a reduction in depth is a precursor of a large fall in the price of foreign currency on average, the excess return on a portfolio that comprises currencies with low depth should produce negative average returns, consistent with the results in Table 2. This same logic applies to the other liquidity measures based on the spreads and intraday volatility, but it appears that these measures contain less extra information about future foreign currency prices, so portfolios sorted on these measures have smaller average (absolute) returns.

2 Factor Returns and Currency Risk Premia

I now develop a no-arbitrage model for the risk premia on individual foreign currencies that uses the properties of the factor returns discussed above. Earlier research on currency portfolio returns examined the implications of no-arbitrage models for the cross-section of (unconditional) expected returns (see, e.g., Lustig and Verdelhan, 2007 and Lustig, Roussanov, and Verdelhan, 2011), but here I focus on the time-series implications. In particular, the model developed here provides a simple framework for identifying how different risk factors contribute to time-varying foreign currency risk premia. I will then use this framework in Section 3 to empirically examine the behavior of the risk premia in each of the 17 foreign currencies.

The starting point for the model is the no arbitrage condition $1 = \mathrm{E}_t[\exp(\kappa_{t+1} + r^j_{t+1})]$, where κ_{t+1} is the log Stochastic Discount Factor (SDF) and r^j_{t+1} is the log return on asset j denominated in US Dollars. $\mathrm{E}_t[.]$ denotes expectations conditioned on information available at the end of month t. I assume that this condition applies to all foreign currency returns and the return on U.S. T-bills, r^{TB}_t. My focus is on excess currency returns, so it is useful to rewrite the no-arbitrage condition as

$$1 = \mathrm{E}_t[\exp(m_{t+1} + er^j_{t+1})], \tag{1}$$

where $m_{t+1} = \kappa_{t+1} + r^{\text{TB}}_t$ is the adjusted SDF and $er^j_{t+1} = r^j_{t+1} - r^{\text{TB}}_t$ is the log excess return on asset j. In the case of currency j, the log excess return is $s^j_{t+1} - f^j_t$. Under the assumption that the conditional distribution of log excess returns and the adjusted SDF is jointly normal, (1) implies that

$$\mathrm{E}_t er^j_{t+1} + \tfrac{1}{2}\mathrm{V}_t(er^j_{t+1}) = -\mathrm{CV}_t(m_{t+1}, er^j_{t+1}), \tag{2}$$

where $\mathrm{V}_t(.)$ and $\mathrm{CV}_t(.,.)$ denote the conditional variance and covariance, respectively. The left-hand-side of this expression identifies the foreign exchange risk premium for currency j: the expected log excess return on currency j plus one half the conditional variance to account for the fact that we are working with logs rather than levels of returns. Equation (2) shows that the risk premium is determined by the conditional covariance between the SDF m_{t+1} and the excess return.

I now derive a beta representation for the right-hand-side of (2) that can be computed from the moments of the factor returns. Let $\mathbf{er}_{t+1}$ denote the 17×1 vector of log excess currency returns. Following Campbell and Viceira (2002), we can well-approximate the log excess returns on the high and low factor portfolio as

$$\begin{aligned} er^{\text{H}}_{t+1} &\simeq w^{\text{H}}_t(\text{X})'\mathbf{er}_{t+1} + \tfrac{1}{2}w^{\text{H}}_t(\text{X})'\left(\mathrm{diag}[\Omega_t] - \Omega_t w^{\text{H}}_t(\text{X})\right) \qquad \text{and} \\ er^{\text{L}}_{t+1} &\simeq w^{\text{L}}_t(\text{X})'\mathbf{er}_{t+1} + \tfrac{1}{2}w^{\text{L}}_t(\text{X})'\left(\mathrm{diag}[\Omega_t] - \Omega_t w^{\text{L}}_t(\text{X})\right), \end{aligned} \tag{3}$$

where $\Omega_t = \mathrm{V}_t(\mathbf{er}_{t+1})$. Here $w^{\text{H}}_t(\text{X})$ and $w^{\text{L}}_t(\text{X})$ are the vectors of portfolio weights that assign currencies to the high and low portfolios based on their rank for factor X in month t.[14] These approximations make er^{H}_{t+1} and er^{L}_{t+1} linearly dependent on the vector individual log excess returns, $\mathbf{er}_{t+1}$. As a result, if the individual log excess returns in $\mathbf{er}_{t+1}$ satisfy (2), then $\mathrm{E}_t er^{\text{H}}_{t+1} + \frac{1}{2}\mathrm{V}_t(er^{\text{H}}_{t+1}) \simeq -\mathrm{CV}_t(m_{t+1}, er^{\text{H}}_{t+1})$ and $\mathrm{E}_t er^{\text{L}}_{t+1} + \frac{1}{2}\mathrm{V}_t(er^{\text{L}}_{t+1}) \simeq$ -$\mathrm{CV}_t(m_{t+1}$,$\mathrm{er}^{\text{L}}_{t+1})$. Taking the difference between these two approximations produces

$$\mathrm{E}_t er^{\text{X}}_{t+1} + \tfrac{1}{2}\mathrm{V}_t(er^{\text{H}}_{t+1}) - \tfrac{1}{2}\mathrm{V}_t(er^{l}_{t+1}) = -\mathrm{CV}_t(m_{t+1}, er^{\text{X}}_{t+1}), \tag{4}$$

where $er^{\text{X}}_{t+1} = er^{\text{H}}_{t+1} - e^{\text{L}}_{t+1}$.

[14] For example, in the case with the high portfolio comprises currencies 1-5, $w^{\text{H}}_t(x)' = [\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, ...,]$.

To derive the beta representation, I propose a specification for the SDF m_{t+1} that satisfies (4) for all five factor returns. Let $\mathbf{er}^{\text{X}}_{t+1} = \mathbf{er}^{\text{H}}_{t+1} - \mathbf{er}^{\text{L}}_{t+1}$ denote the 5×1 vector of high minus low excess returns for the five factors. The proposed SDF is given by

$$m_{t+1} = \psi_t - \lambda_t'(\mathbf{er}^{\text{X}}_{t+1} - E_t\mathbf{er}^{\text{X}}_{t+1}), \tag{5}$$

where ψ_t is a time-varying scalar, and λ_t is a time-varying 5×1 vector, which needs to be determined. Since the no-arbitrage condition applies to the log return on U.S. T-bills, $1 = \text{E}_t[\exp(m_{t+1})]$, and so $\text{E}_t m_{t+1} + \frac{1}{2}\text{V}_t(m_{t+1}) = 0$. Consequently,

$$\psi_t + \frac{1}{2}\lambda_t'\Omega^{\text{X}}_t\lambda_t = 0, \tag{6}$$

where $\Omega^{\text{X}}_t = \text{V}_t(\mathbf{er}^{\text{X}}_{t+1})$. Furthermore, we can stack the 5 equations in (4) to give

$$\text{E}_t\mathbf{er}^{\text{X}}_{t+1} + \tfrac{1}{2}\text{diag}[\Omega^{\text{H}}_t] - \tfrac{1}{2}\text{diag}[\Omega^{\text{L}}_t] = \Omega^{\text{X}}_t\lambda_t, \tag{7}$$

where $\Omega^{\text{H}}_t = \text{V}_t(\mathbf{er}^{\text{H}}_{t+1})$ and $\Omega^{\text{L}}_t = \text{V}_t(\mathbf{er}^{\text{L}}_{t+1})$. Equations (6) and (7) pin down ψ_t and λ_t, so we can rewrite (5) as

$$m_{t+1} = -\frac{1}{2}\mu_t' \left(\Omega^{\text{X}}_t\right)^{-1} \mu_t - \mu_t' \left(\Omega^{\text{X}}_t\right)^{-1} (\mathbf{er}^{\text{X}}_{t+1} - E_t\mathbf{er}^{\text{X}}_{t+1}), \tag{8}$$

where

$$\mu_t = \text{E}_t\mathbf{er}^{\text{X}}_{t+1} + \tfrac{1}{2}\text{diag}[\Omega^{\text{H}}_t] - \tfrac{1}{2}\text{diag}[\Omega^{\text{L}}_t].$$

Finally, substituting for m_{t+1} from (8) in (2) gives the beta representation for the currency risk premium:

$$\begin{aligned} \text{E}_t er^j_{t+1} + \tfrac{1}{2}\text{V}_t(er^j_{t+1}) &= \mu_t' \left(\Omega^{\text{X}}_t\right)^{-1} \text{CV}_t(\mathbf{er}^{\text{X}}_{t+1}, er^j_{t+1}) \\ &= \mu_t'\beta^j_t \end{aligned} \tag{9}$$

This equation shows that the risk premium for currency j is equal to the beta-weighted sum of the factor risk prices that comprise the vector μ_t. The betas for currency j are identified by the slope coefficients from the (conditional) projection of er^j_{t+1} on five factor return differences in $\mathbf{er}^{\text{X}}_{t+1}$.

The beta-representation in (9) implies that cross-currency differences in risk premia at a point in time are due to differences in their betas. For example, the difference between the risk premia on foreign currencies i and j at month t, is equal to $\mu_t'(\beta^i_t - \beta^j_t)$. In this model there are five betas for each currency, so differences between risk premia depend on the weighted sum of the differences in each of the individual betas. The beta representation also allows for time-variation in the risk premia for individual currencies; either via changes in the currency's beta or through movements in the price of risk. In the empirical analysis that follows, I find that estimates of the betas are quite stable over the sample period, so variations in the price of risk are the main drivers of currency risk premia through time.

Several features of the model deserve comment: First, the beta representation in (9) applies to log excess returns rather simple excess currency returns (i.e., $S_{t+1} - F_t$), which is the norm in the literature. This feature facilitates the integration of the beta representation in the analysis of individual currencies below. Second, the beta representation in (9) is conditional, in the sense that the betas and factor risk prices are derived from the conditional first and second moments of log excess returns rather than the unconditional moments. Of course, the no-arbitrage condition $1 = \mathrm{E}_t[\exp(\kappa_{t+1} + r^j_{t+1})]$ also implies the existence of an unconditional beta representation, but that is better suited to the analysis of average risk premia across currencies.

The third feature of the model concerns the specification for the log SDF in (8). Earlier research on currency portfolio returns considered no-arbitrage models in which SDFs for different counties follow particular stochastic processes driven by multiple shocks. For example, Verdelhan (2018) proposes a model where the log SDFs are driven by country-specific and global shocks. Such models have implications for the behavior of factor-sorted portfolios of currency returns which can be used to interpret average returns on actual factor-sorted portfolios.[15] Other studies use specifications for the SDF that include the risk factors directly. For example, Menkhoff et al. (2012a) include their FX volatility risk factor in their specification for the SDF. In contrast, here I specify the log SDF as a function of the returns on the five factor portfolios, but the actual form of the function (i.e., the values for ψ_t and λ_t) is pinned down by the requirement that each of the factor portfolio returns satisfies the no-arbitrage condition in (4). This approach is analogous to the common practice of using the factor portfolio returns to estimate the parameters of a linear specification for the SDF (see, Lustig and Verdelhan, 2007, Lustig, Roussanov, and Verdelhan 2011, Menkhoff et al., 2012a, and others). The key difference is that I find analytic expressions for ψ_t and λ_t that ensure the factor portfolios exactly meet the no-arbitrage conditions. Thus, the model is not designed to test whether the factor portfolios (or the individual currency returns) satisfy the no-arbitrage conditions, but rather to examine the sources of time-series variation in the risk premia for individual currencies *within a no-arbitrage framework*.

Finally, it is worth emphasizing that the specification for the log SDF incorporates information on the five risk factors, but only through their impact on the portfolio returns. This is a particularly useful feature of the model when considering the possible role of liquidity as a risk factor. As I noted in the Introduction,

[15] To derive these implications, the models typically assume that markets are complete. In contrast, my approach does not make an assumption about the degree of international risk sharing, see Evans (2017) for a discussion.

liquidity is a multi-faceted concept that cannot be fully represented by a single variable. Here I consider two measures of liquidity, depth and the bid-ask spread; as well as intraday volatility, which may also be related to liquidity. My specification for the log SDF includes all three measures without an assumption about how accurately they each represent “true liquidity”.

3 Empirical Results

I use the beta representation in equation (9) to estimate the dynamics of the currency risk premia. In Section 1.1 I showed that there was no statistically significant evidence of heteroskedasticity in the residuals of the ARMA models for the factor returns, er^{HL}_{t+1}. My empirical implementation, therefore, proceeds under the assumption that the factor returns are conditionally homoskedastic. This makes the currency betas constant, so they can be estimated from simple regressions. I also test for stability in the beta estimates to check that the homoskedasticity assumption is reasonable. The complete model for the currency risk premia can then be obtained by combining the betas with forecasts of the factor returns from the ARMA models.

3.1 Currency Betas

According the equation (9), the vector of betas for currency j are given by

$$\beta^j_t = \mathrm{V}_t(\mathbf{er}^{\text{x}}_{t+1})^{-1}\mathrm{CV}_t(\mathbf{er}^{\text{x}}_{t+1}, er^j_{t+1}),$$

where $\mathbf{er}^{\text{x}}_{t+1}$ is the vector of log portfolio returns for the five factors. If $\mathbf{er}^{\text{x}}_{t+1}$ is conditionally homoskedastic, as was indicated by the results in Section 1.1, the betas for currency j are constant, and can be estimated as the slope coefficients from a regression of the excess return for currency j on a constant and the vector $\mathbf{er}^{\text{x}}_{t+1}$:

$$er^j_{t+1} = \beta_o + (\mathbf{er}^{\text{x}}_{t+1})'\beta^j + u^j_{t+1}. \tag{10}$$

In principle, the excess return for currency j can appear in the high and low portfolios used to construct the elements of $\mathbf{er}^{\text{x}}_{t+1}$, so er^j_{t+1} could be present on both sides of the regression equation. This could give rise to spurious estimates of the betas if it occurs frequently throughout the sample. In practice, this does not appear to be a problem because the composition of the high and low portfolios changes so much from month to month. Indeed, if I construct the $\mathbf{er}^{\text{x}}_{t+1}$ vector from high and low portfolio returns that exclude currency j, I obtain very similar results. These results are discussed in Section 3.4.

Table 4 reports the beta estimates obtained by estimating regression (10) for each of the 17 currency

Table 4: Currency Betas

		Factors					
		fd_t (i)	D_t (ii)	$sprd_t^{-1}$ (iii)	μ_t (iv)	σ_t^{-1} (v)	R^2/SE
Group A:	EURUSD	0.383*	-0.131	0.461**	0.433**	0.632***	0.376
		(0.205)	(0.187)	(0.203)	(0.211)	(0.198)	30.204
	JPYUSD	-0.037	-0.427***	-0.194	0.017	0.099	0.204
		(0.162)	(0.174)	(0.201)	(0.163)	(0.197)	29.313
	CHFUSD	0.905***	-0.160	0.567**	0.307	0.251	0.315
		(0.189)	(0.185)	(0.235)	(0.205)	(0.237)	33.232
	GBPUSD	0.080	0.402**	-0.096	0.079	0.517***	0.211
		(0.191)	(0.188)	(0.174)	(0.150)	(0.128)	28.236
Group B:	AUDUSD	0.774***	0.061	0.960***	0.293	0.838***	0.610
		(0.169)	(0.169)	(0.196)	(0.183)	(0.188)	31.499
	CADUSD	-0.241	-0.067	0.404***	0.094	0.696***	0.493
		(0.150)	(0.153)	(0.134)	(0.129)	(0.105)	25.592
	HKDUSD	-0.003	-0.010	0.009	0.012*	-0.004	0.067
		(0.006)	(0.008)	(0.007)	(0.007)	(0.006)	1.138
	NOKUSD	-0.173	-0.129	0.473***	0.203	0.896***	0.488
		(0.158)	(0.154)	(0.192)	(0.219)	(0.204)	30.557
	NZDUSD	1.147***	0.193	0.958***	0.305	0.897***	0.632
		(0.217)	(0.187)	(0.171)	(0.220)	(0.164)	32.549
	SEKUSD	0.185	-0.192	0.615***	0.355	0.764***	0.442
		(0.174)	(0.176)	(0.184)	(0.219)	(0.187)	31.983
	SGDUSD	0.339**	0.082	0.325**	-0.014	0.333***	0.463
		(0.141)	(0.100)	(0.159)	(0.113)	(0.125)	15.676

		Factors					
		fd_t (i)	D_t (ii)	$sprd_t^{-1}$ (iii)	μ_t (iv)	σ_t^{-1} (v)	R^2/SE
Group C:	CNHUSD	-0.395***	-0.053	-0.262***	-0.076	-0.008	0.363
		(0.135)	(0.180)	(0.097)	(0.115)	(0.084)	8.625
	MXNUSD	0.100	0.110	0.796***	0.307**	0.325*	0.605
		(0.137)	(0.126)	(0.166)	(0.155)	(0.173)	24.607
	RUBUSD	-2.623***	-1.506***	1.252***	-0.652*	1.114***	0.774
		(0.507)	(0.392)	(0.346)	(0.389)	(0.405)	39.403
	TRYUSD	0.025	0.531***	0.561***	0.282	0.480**	0.670
		(0.183)	(0.170)	(0.186)	(0.195)	(0.201)	26.385
	PLNUSD	0.355	0.065	1.609***	0.401	-0.163	0.471
		(0.326)	(0.314)	(0.498)	(0.337)	(0.495)	50.600
	ZARUSD	0.820***	0.674**	0.660***	0.233	0.844***	0.589
		(0.275)	(0.307)	(0.253)	(0.257)	(0.230)	34.823
	Average Betas						
	Group A	0.333	-0.079	0.185	0.209	0.375	0.276
	Group B	0.290	-0.009	0.535	0.178	0.631	0.456
	Group C	-0.286	-0.030	0.770	0.083	0.432	0.579
	All	0.097	-0.033	0.535	0.152	0.501	0.457

Notes: The table reports estimates of currency betas for each of the five factors. The betas for currency j are estimated from the regression:

$$er^j_{t+1} = \beta_0 + (\mathbf{er}^{\mathrm{x}}_{t+1})'\beta^j + u^j_{t+1}$$

where $\mathbf{er}^{\mathrm{x}}_{t+1}$ is the 5×1 vector of log excess returns differences between high and low factor portfolios for each of the five factors x: the forward discount $fd_t = s_t - f_t$, depth D_t, the reciprocal of the spread $sprd_t^{-1}$, the price trend μ_t, and the reciprocal of volatility σ_t^{-1}. Robust standard errors are shown in parenthesis below the beta estimates. The R^2 statistic and regression standard error (SE) are reported in the right-hand column of each block. Results are displayed for the same groupings of currencies as in Table 1. The lower right-hand block of the table reports averages of the beta estimates for each each group of currencies and across all currencies.

pairs. Again, for readability, the table shows results for the same currency groups as Table 1. Overall, two features of the estimates stand out. First, the betas for each of the five factors are large and statistically significant for many of the currencies. Thus individual currency returns reflect, in part, exposure to multiple sources of risk. Second, there are significant cross-currency differences in the size and sign of the betas for a particular factor. Individual currency returns reflect very different exposures to particular sources of risk. Together, these features give rise to a very heterogenous pattern of beta estimates across factors and currencies.

It is instructive to consider how the estimated betas for particular factors differ across currencies. Consider first the betas for the forward discount shown in column (i). According to the estimates, the carry-trade risk captured by this factor makes a significant contribution to the returns on seven of the currencies. The estimates are positive for the CHFUSD, AUDUSD, NZDUSD, SGDUSD, and ZARUSD, and negative for both the CNHUSD and RUBUSD. The former group of currencies is traditionally associated with the carry trade, so it is unsurprising that their betas are positive. By contrast, the negative beta estimates for the CNHUSD and particularly the RUBUSD indicate that both currencies produce low excess returns when their interest rates are high, which is contrary the logic of the carry trade. It is also worth noting that the carry-trade betas are small and insignificant for a majority of the currencies. This implies that cross-currency differences in exposure to carry-trade risk do not account for much of the behavior of individual currency returns. These findings are consistent with the results in Burnside (2019) after the 2007-2008 regression.

The estimated betas for the depth and spread factors, reported in columns (ii) and (iii), indicate how exposure to "liquidity" risk differs across currencies. Recall that liquidity risk is measured by the return on borrowing in the least liquid currencies (smallest depth or widest spread) and lending in the most liquid. Although the factor returns for depth and spreads are positively correlated (see Table 2), many more currencies appear to have significant exposure to the spread measure of liquidity risk than the depth measure. The spread betas are positive and statistically significant in 13 currency pairs, whereas the depth betas are significant and positive for only three currency pairs. It is also interesting to note that the estimated liquidity betas are quite different among the major currencies in Group A. In particular both the EURUSD and CHFUSD have positive spread betas and insignificant depth betas. In contrast, the JYPUSD and GBPUSD have insignificant spread betas but their depth betas have significant different signs; positive for GBPUSD and negative for JPYUSD. The estimated liquidity betas display a more consistent pattern among the Group B currencies. None of the depth estimates are significant, while all but one of the spread betas are positive and significant. The spread betas are also significant for all the currencies in Group C, but in the CNHUSD case, the estimate is negative. Four of these currency pairs also have significant depth betas, but as in Group

A, the estimates have different signs.

Columns (iv) and (v) of Table 4 report the betas for trend and volatility factor. The trend beta measures the risk of following a momentum strategy that involves borrowing in currencies that are appreciating most against the USD and lending in those that are depreciating most. It appears that this is only an important of risk for the EURUSD and MXNUSD, since these are the only currencies with significant beta estimates. Recall that the volatility factor is based on average intraday volatility. The estimated volatility betas are positive and statistically significant in 11 of the currency pairs. Because spreads are strongly correlated with intraday volatility, the spread and volatile betas are generally similar across the currencies.

Table 5: Beta Stability

		Factors					
		fd_t	D_t	$sprd_t^{-1}$	μ_t	σ_t^{-1}	Joint
A:	EURUSD	0.263	0.260	0.744	0.539	0.876	0.317
	JPYUSD	0.109	0.483	0.479	0.001	0.330	0.004
	CHFUSD	0.488	0.683	0.211	0.922	0.748	0.641
	GBPUSD	0.721	0.102	0.076	0.441	0.534	0.336
B:	AUDUSD	0.311	0.190	0.822	0.213	0.302	0.570
	CADUSD	0.897	0.125	0.270	0.048	0.982	0.323
	HKDUSD	0.114	0.020	0.632	0.434	0.693	0.338
	NOKUSD	0.808	0.395	0.580	0.598	0.584	0.786
	NZDUSD	0.349	0.142	0.565	0.067	0.417	0.211
	SEKUSD	0.740	0.213	0.997	0.366	0.395	0.715
	SGDUSD	0.341	0.444	0.366	0.919	0.987	0.713
C:	CNHUSD	0.459	0.881	0.177	0.859	0.649	0.141
	MXNUSD	0.266	0.481	0.224	0.445	0.565	0.542
	RUBUSD	0.050	0.994	0.068	0.346	0.005	0.000
	TRYUSD	0.001	0.453	0.660	0.539	0.409	0.021
	PLNUSD	0.285	0.286	0.465	0.818	0.730	0.306
	ZARUSD	0.637	0.018	0.236	0.646	0.690	0.148

Notes: The table reports p-values for Wald tests of the null that the betas estimated in regression (10) are equal in the first half and second half of the sample period. The p-values for the Wald tests of the betas for individual factors are show under the columns headed by each factor. The left-had column reports the p-value for the null that there is sub-sample stability in all the betas jointly. The factors shown in the headings are: the forward discount fd_t, depth D_t, the reciprocal of the spread $sprd_t^{-1}$, the price trend μ_t, and the reciprocal of volatility σ_t^{-1}.

The estimates in Table 4 assume that there is little time-series variation in the betas. To check on the validity of this assumption, Table 5 reports the results of stability test on the betas estimated from regression

(4). In particular, the table reports p-values for the null hypothesis that estimates of the betas from the first and second half of the sample are equal. These stability tests are applied to the betas on the individual factors and across the betas for all the factors jointly. As the table shows, there is generally very little evidence of instability, either in the individual betas or jointly across the betas. There are, however, three exceptions to this general pattern. In the case of the RUBUSD and TRYUSD some of the p-values for the individual betas are very small, and we would reject the joint stability test at the 5 percent significance level. However, since the available data for both currency pairs is somewhat limited (see Table 1), it is unclear whether these asymptotic inferences are entirely reliable. The other exception is the JYPUSD, where the Wald test strongly rejects stability in the beta on the price-trend factor. Since there is little evidence of instability in this beta from many of the other currency pairs, or in the other factor betas for the JYPUSD, there is no obvious explanation for this result. I will nevertheless allow for the possibility that this beta is poorly estimated in the analysis below.

3.2 Currency Risk Premia

I now combined the beta estimates from Table 4 with forecasts of the factor returns implied by the ARMA models in Table 2 to study the properties of the risk premia on individual currencies. For this purpose, I rewrite the beta representation in (9) as

$$rp_t^j \equiv \mathrm{E}_t er_{t+1}^j = \alpha^j + (\mathrm{E}_t \mathbf{er}_{t+1}^{\mathrm{X}})' \beta^j \tag{11}$$

where $\alpha^j = (\frac{1}{2}\mathrm{diag}[\Omega_t^{\mathrm{H}}] - \frac{1}{2}\mathrm{diag}[\Omega_t^{\mathrm{L}}])\beta^j - \frac{1}{2}\mathrm{V}_t(er_{t+1}^j)$. I interpret the results of the stability tests in Table 5 to mean that the conditional variances and covariances of log excess returns exhibit little heteroskedasticity, so it is reasonable to treat the α^j term as a constant. The other terms on the right-hand-side of (11) are computed from the estimated betas for currency j from Table 4 $\hat{\beta}^j = [\hat{\beta}_{\mathrm{X}}^j]$ for $\mathrm{X} = \{1, ., 5\}$, and the one-month ahead ARMA forecasts for the factor returns $\hat{\mathrm{E}}_t \mathbf{er}_{t+1}^{\mathrm{X}} = [\mathrm{E}_t er_{t+1}^{\mathrm{X}}]$. Henceforth, I refer to the expected log excess returns, $\mathrm{E}_t er_{t+1}^j$, as the currency risk premia for simplicity.

According to (11) time-series variations in the currency risk premia originate from change in the price of risk for each of the five risk factors. The relative importance of changing risk prices for individual currency risk premia can be measured by their variance contributions. In particular, (11) implies that $\mathrm{V}(rp_t^j) = \sum_{\mathrm{X}=1}^{5} \beta_{\mathrm{X}}^j \mathrm{CV}(\mathrm{E}_t \mathrm{er}_{t+1}^{\mathrm{X}}, \mathrm{rp}_t^j)$, so the variance contribution of the risk price for factor X is given by $\beta_{\mathrm{X}}^j \mathrm{CV}(\mathrm{E}_t er_{t+1}^{\mathrm{X}}, rp_t^j)/\mathrm{V}(rp_t^j)$. The estimated variance contributions of each of the five risk-prices to the individual currency risk premia are shown in columns (i) to (v) of Table 6. The right-hand columns of the

Table 6: Risk Price Variance Contributions

		Risk Prices					Variance Ratios	
		fd_t (i)	D_t (ii)	$sprd_t^{-1}$ (iii)	μ_t (iv)	σ_t^{-1} (v)	$\mathcal{R}_1$ (vi)	$\mathcal{R}_2$ (vii)
A:	EURUSD	0.224	-0.045	0.155	0.273	0.392	0.028	0.075
	JPYUSD	-0.001	0.808	0.197	-0.001	-0.004	0.014	0.071
	CHFUSD	0.711	-0.030	0.143	0.105	0.071	0.040	0.128
	GBPUSD	0.023	0.594	-0.048	0.031	0.400	0.044	0.210
B:	AUDUSD	0.281	0.029	0.362	0.058	0.270	0.051	0.084
	CADUSD	0.080	-0.050	0.303	0.026	0.642	0.034	0.069
	HKDUSD	0.027	0.463	-0.031	0.492	0.049	0.010	0.143
	NOKUSD	0.012	-0.061	0.241	0.075	0.733	0.027	0.055
	NZDUSD	0.398	0.074	0.261	0.047	0.220	0.068	0.107
	SEKUSD	0.066	-0.074	0.296	0.186	0.526	0.019	0.043
	SGDUSD	0.325	0.107	0.308	-0.002	0.263	0.024	0.051
C:	CNHUSD	0.621	0.071	0.267	0.038	0.003	0.052	0.143
	MXNUSD	0.012	0.112	0.659	0.109	0.108	0.042	0.069
	RUBUSD	0.659	0.249	0.017	0.069	0.006	0.022	0.029
	TRYUSD	0.001	0.497	0.309	0.064	0.129	0.056	0.084
	PLNUSD	0.033	0.037	0.880	0.055	-0.005	0.020	0.043
	ZARUSD	0.199	0.373	0.203	0.032	0.194	0.116	0.197
	Average	0.216	0.186	0.266	0.097	0.235		

Notes: The table reports the variance contribution of each of the risk price to the individual currency risk premia, computed by $\hat{\beta}_{\mathrm{x}}^{j}\mathrm{CV}(\hat{\mathrm{E}}_t er_{t+1}^{\mathrm{x}}, \hat{\mathrm{E}}_t er_{t+1}^{j})/\mathrm{V}(\hat{\mathrm{E}}_t er_{t+1}^{j})$, where $\hat{\beta}_{\mathrm{x}}^{j}$ is the estimated beta on factor x for currency j shown in Table 4, and $\hat{\mathrm{E}}_t er_{t+1}^{\mathrm{x}}$ is the one-month-ahead ARMA forecast for the factor x return computed from the ARMA models shown in Table 3. The estimated risk premia $\hat{\mathrm{E}}_t er_{t+1}^{j}$ are computed from the right-hand-side of (11) with $\alpha^j = 0$. $CV(.,.)$ and $V(.)$ denote the sample covariance and variance. The variance ratios are $\mathcal{R}_1 = \mathrm{V}(\hat{\mathrm{E}}_t er_{t+1}^{j})/\mathrm{V}(er_{t+1}^{j})$ and $\mathcal{R}_2 = \mathrm{V}(\hat{\mathrm{E}}_t er_{t+1}^{j})/\mathrm{V}(\mathrm{E}[er_{t+1}^{j}|\mathbf{er}_{t+1}^{\mathrm{x}}])$, where $\mathrm{E}[er_{t+1}^{j}|\mathbf{er}_{t+1}^{\mathrm{x}}]$ is the fitted value from regression (10) in Table 4.

table also report two variance ratios: $\mathcal{R}_1 = \mathrm{V}(rp_t^j)/\mathrm{V}(er_{t+1}^j)$ and $\mathcal{R}_2 = \mathrm{V}(rp_t^j)/\mathrm{V}(\mathrm{E}[er_{t+1}^j|\mathbf{er}_{t+1}^{\mathrm{x}}])$, where $\mathrm{E}[er_{t+1}^j|\mathbf{er}_{t+1}^{\mathrm{x}}]$ denotes the projection of the excess currency return on the five factor returns (i.e., the fitted value from regression (10) in Table 4). The $\mathcal{R}_1$ measures the variance contribution of the risk premium to realized log excess returns, while $\mathcal{R}_2$ measures the premium's contribution to the variance of excess returns that are perfectly correlated with the five factor returns.

The estimates in Table 6 show that variations in the price of risk for all five factors make significant contributions to the variability of individual currency risk premia, but the importance of their contributions differs considerably across currency pairs. In particular, variations in the price of carry-risk are the most important source of variation in the risk premia for just four pairs: the CHFUSD, NZDUSD, SGFUSD, CNHUSD, and RUBUSD. For the majority of risk premia, changes in the price of carry-risk have very little effect. The average variance contribution across all 17 currencies is 22 percent.

Variations in the price of liquidity risk make more substantial variance contributions. For example, changes in the price of liquidity-risk measured by depth account for approximately 81 and 59 percent of the variability in the risk premia for the JPYUSD and GBPUSD, respectively. This is a surprising result in light of the fact that spot trading in both currencies is thought to be highly liquid. Among the other currencies, changes in the price of liquidity-risk contribute significantly to variability of the RUBUSD, TRYUSD, and ZAR risk premia. Changes in the price of liquidity-risk measured by the spread make sizable contributions to the volatility of the risk premia for a different set of currency pairs; notably the AUDUSD, MXNUSD, and PLNUSD. Changes in this risk price make the largest variance contribution on average across the 17 currencies, accounting for approximately 27 percent of the volatility in the risk premia. Changes in both liquidity risk prices on average account for approximately 45 percent of premia volatility, which is more than twice the contribution of carry-risk prices.

Columns (iv) and (v) show the contributions of the momentum and intraday volatility risk prices. Changes in the price of momentum risk make small contributions to the volatility of most risk premia, with the exceptions of the EURUSD and HKDUSD, where the contributions are 27 and 49 percent, respectively. The average contribution across all currencies is 10 percent, which is the lowest contribution of the five risk prices. The variance contribution of intraday volatility is large on average at 24 percent, and is particularly significant in the case of the CADUSD, NOKUSD, and SEKUSD. As was noted earlier, because high intraday volatility may reflect the large price-impact of trades when depth is low, variations in the price of volatility risk may represent a third facet of changing liquidity risk pricing. Under this interpretation, the total variance contribution of all forms of liquidity risk is 69 percent on average across all 17 currencies.

Finally, the estimated variance ratios in columns (vi) and (vii) compare the variability of the risk premia with realized excess returns. The estimates of $\mathcal{R}_1$ show that variations in the risk premia are quite small compared to the variability in realized returns. The average value for $\mathcal{R}_1$ across all currencies is 3.9 percent. These estimates are in line with the R^2 statistics obtained from regressions of log excess returns on forward discounts. The estimates of $\mathcal{R}_2$ are larger, averaging 9.4 percent across all currencies. The R^2 statistics from the ARMA models in Table 3 showed changes in the price of risk account for between 8 and 16 percent of the variations in factor returns, so the estimates for $\mathcal{R}_2$ imply that the individual risk premia are approximately as variable as the five risk prices.

3.3 News about Risk

The results in Table 6 show that while the changing price of liquidity risk is the most important driver of the variations in currency risk premia, it does not produce much predictability in realized excess returns. I now consider another channel through which excess returns are affected by changes in the price of liquidity risk, namely through the information they convey about future risk premia.

My analysis is based on an identity that links excess currency returns to changing expectations about the paths for future forward discounts and risk premia. The identity is derived from the definition of the risk premium $rp_t = \mathrm{E}_t s_{t+1} - s_t + fd_t$. Re-writing this identify as a difference equation in s_t, solving forwards H periods, and applying the Law of Iterated Expectations produces

$$s_t = -\mathrm{E}_t \sum_{i=0}^{H-1} rp_{t+i} + \mathrm{E}_t \sum_{i=0}^{H-1} fd_{t+i} + \mathrm{E}_t s_{t+H},$$

so the error in forecasting the spot rate one-month ahead is

$$s_{t+1} - \mathrm{E}_t s_{t+1} = -(\mathrm{E}_{t+1} - \mathrm{E}_t) \sum_{i=1}^{H} rp_{t+i} + (\mathrm{E}_{t+1} - \mathrm{E}_t) \sum_{i=1}^{H} fd_{t+i} + (\mathrm{E}_{t+1} - \mathrm{E}_t) s_{t+1+H}. \tag{12}$$

By definition, the realized excess return equals the risk premium and forecast error: $er_{t+1} = rp_t + s_{t+1} - \mathrm{E}_t s_{t+1}$. Combining this identity with (12) and taking the limit as $H \to \infty$ gives the following equation for the log excess return

$$er_{t+1} = rp_t - (\mathrm{E}_{t+1} - \mathrm{E}_t) \sum_{i=1}^{\infty} rp_{t+i} + (\mathrm{E}_{t+1} - \mathrm{E}_t) \sum_{i=1}^{\infty} fd_{t+i} + \zeta_{t+1}, \tag{13}$$

where $\zeta_{t+1} = lim_{H \to \infty} (\mathrm{E}_{t+1} - \mathrm{E}_t) s_{t+1+H}$.

Equation (13) identifies all the proximate factors that can drive log excess returns. The first term on the right identifies the expected log excess return; i.e., the risk premium. The remaining terms identify the factors that contribute to the error in forecasting the spot rate one month. These factors are: (i) news about future risk premia identified by the second term, (ii) news about expected interest differentials implicit in the forward discount in the third term, and (iii) revisions in expectations concerning the spot rate in the distant future.

My focus is on the term identifying news about future risk premia. According to (9) the risk premium for currency j has the beta representation: $rp_t^j = \mu_t' \beta^j = \sum_{\mathrm{x}=1}^{5} \beta_{\mathrm{x}}^j \mu_t^{\mathrm{x}}$, where μ_t^{x} is the price of risk for factor x and β_{x}^j is the factor beta for currency j. Substituting this expression for the risk premium allows us to write news concerning future risk premia for currency j as

$$(\mathrm{E}_{t+1}-\mathrm{E}_t)\sum_{i=1}^{\infty} rp_{t+i}^j = \sum_{\text{x}=1}^{5}\left(\beta_{\text{x}}^j \sum_{i=1}^{\infty}(\mathrm{E}_{t+1}-\mathrm{E}_t)\mu_{t+i}^{\text{x}}\right). \tag{14}$$

Thus, news about the future path of the risk premium for currency j is equal to the beta-weighted average of news concerning the future price of risk for the five factors, $(\mathrm{E}_{t+1}-\mathrm{E}_t)\mu_{t+i}^{\text{x}}$. Under the homoskedasticity assumption employed above, news about a risk-price $(\mathrm{E}_{t+1}-\mathrm{E}_t)\mu_{t+i}^{\text{x}}$ is equal to the revision in the forecasts for factor returns $(\mathrm{E}_{t+1}-\mathrm{E}_t)er_{t+1+i}^{\text{x}}$. Thus, to investigate how changing risk prices affect currency returns through the news they convey about future risk premia, I estimate regressions of the form

$$\begin{aligned} er_{t+1}^j \quad = \quad & \hat{rp}_t^j + b_1\sum_{i=1}^{\infty}(\hat{\mathrm{E}}_{t+1}-\hat{\mathrm{E}}_t)er_{t+1+i}^{\text{x}_1} + b_2\sum_{i=1}^{\infty}(\hat{\mathrm{E}}_{t+1}-\hat{\mathrm{E}}_t)er_{t+1+i}^{\text{x}_2} + b_3\sum_{i=1}^{\infty}(\hat{\mathrm{E}}_{t+1}-\hat{\mathrm{E}}_t)er_{t+1+i}^{\text{x}_3} \\ & + b_4\sum_{i=1}^{\infty}(\hat{\mathrm{E}}_{t+1}-\hat{\mathrm{E}}_t)er_{t+1+i}^{\text{x}_4} + b_5\sum_{i=1}^{\infty}(\hat{\mathrm{E}}_{t+1}-\hat{\mathrm{E}}_t)er_{t+1+i}^{\text{x}_5} + \eta_{t+1}. \end{aligned} \tag{15}$$

The first term on the right-hand-side is the estimated risk premium studied above. The next five terms identify news about the future risk premia originating from revisions in the forecasts of future risk-prices. The forecast revisions $(\hat{\mathrm{E}}_{t+1}-\hat{\mathrm{E}}_t)er_{t+1+i}^{\text{x}}$ are computed from the ARMA model estimates in Table 3. I truncate the infinite horizon sums after 60 months because the estimated forecast revisions are negligible beyond this horizon. The regression coefficients b_i quantify the degree to which the ARMA-based estimates of news about risk prices affect (unexpected) excess returns. If news about these risk prices is uncorrelated with news about future interest differentials and long-horizon spot rates, which are both represented by the error term η_{t+1}, the regression coefficients should equal the currency betas; i.e., $b_{\text{x}} = \beta_{\text{x}}^j$ for factors $\text{x} = \{1,..,5\}$ identified in the beta representation for the risk premium in (9). I estimate the b_i's freely, and test whether the estimates satisfy this restriction.

The results from estimating regression (15) for each of the 17 currencies are shown in Table 7. The left-hand panel of the table reports the variance contributions of the risk-prices, while the right-hand reports the results from tests of the coefficient restrictions. I split the variance contributions into three groups of the risk-prices. Column (i) reports the variance contribution of the carry-risk and momentum risk price, column (ii) the contributions of the two liquidity risk prices, and column (iii) the contribution of the volatility risk-price.[16] The variance contribution of all five risk prices are shown in column (iv). The average estimated

[16]The variance contributions are computed in an analogous fashion to those in Table 6. For example, the variance contribution of risk price x_1 is computed as $CV(\hat{b}_1\sum_{i=1}^{60}(\hat{\mathrm{E}}_{t+1}-\hat{\mathrm{E}}_t)er_{t+1+i}^{\text{x}_1}, er_{t+1}^j)/V(er_{t+1}^j)$ where $\hat{b}_1$ are the OLS estimates of b_1 from

Table 7: News Regressions

		I: News R^2				II: Beta Tests			
		C & M (i)	Liquidity (ii)	Volatility (iii)	All (iv)	C & M (v)	Liquidity (vi)	Volatility (vii)	All (viii)
Group A:	EURUSD	0.028	0.080	0.249	0.357	5.205	0.029	29.100	47.718
		(0.021)	(0.018)	(0.040)	(0.039)	(0.074)	(0.986)	(0.000)	(0.000)
	JPYUSD	0.017	0.206	0.002	0.225	0.428	6.458	0.004	14.970
		(0.006)	(0.041)	(0.004)	(0.045)	(0.807)	(0.040)	(0.952)	(0.010)
	CHFUSD	0.133	0.077	0.055	0.265	13.410	0.094	3.408	22.364
		(0.056)	(0.027)	(0.023)	(0.057)	(0.001)	(0.954)	(0.065)	(0.000)
	GBPUSD	0.004	0.022	0.158	0.184	0.745	0.429	40.210	48.759
		(0.007)	(0.020)	(0.039)	(0.044)	(0.689)	(0.807)	(0.000)	(0.000)
Group B:	AUDUSD	0.024	0.222	0.278	0.524	10.583	0.135	42.525	94.244
		(0.033)	(0.040)	(0.042)	(0.042)	(0.005)	(0.935)	(0.000)	(0.000)
	CADUSD	0.077	0.116	0.267	0.460	5.966	0.053	92.481	114.936
		(0.029)	(0.019)	(0.042)	(0.046)	(0.051)	(0.974)	(0.000)	(0.000)
	HKDUSD	0.030	0.032	-0.002	0.060	12.104	0.181	0.362	14.118
		(0.019)	(0.036)	(0.006)	(0.043)	(0.002)	(0.913)	(0.547)	(0.015)
	NOKUSD	0.040	0.136	0.324	0.499	3.411	0.299	49.050	96.743
		(0.019)	(0.024)	(0.043)	(0.040)	(0.182)	(0.861)	(0.000)	(0.000)
	NZDUSD	0.067	0.232	0.255	0.553	23.990	0.620	55.954	116.135
		(0.051)	(0.048)	(0.035)	(0.046)	(0.000)	(0.733)	(0.000)	(0.000)
	SEKUSD	0.018	0.130	0.267	0.415	3.137	0.462	36.834	69.595
		(0.013)	(0.021)	(0.041)	(0.045)	(0.208)	(0.794)	(0.000)	(0.000)
	SGDUSD	0.045	0.189	0.243	0.477	8.512	0.721	25.832	59.560
		(0.065)	(0.057)	(0.052)	(0.056)	(0.014)	(0.697)	(0.000)	(0.000)
Group C:	CNHUSD	0.385	-0.026	0.007	0.365	11.935	0.083	0.835	35.785
		(0.167)	(0.088)	(0.026)	(0.103)	(0.003)	(0.960)	(0.361)	(0.000)
	MXNUSD	0.061	0.301	0.113	0.475	6.502	1.223	14.565	23.723
		(0.021)	(0.031)	(0.024)	(0.040)	(0.039)	(0.542)	(0.000)	(0.000)
	RUBUSD	0.400	0.160	0.211	0.771	41.589	1.910	17.368	86.231
		(0.061)	(0.063)	(0.031)	(0.044)	(0.000)	(0.385)	(0.000)	(0.000)
	TRYUSD	0.054	0.438	0.174	0.665	9.354	5.336	29.092	57.555
		(0.020)	(0.050)	(0.040)	(0.046)	(0.009)	(0.069)	(0.000)	(0.000)
	PLNUSD	0.007	0.341	0.000	0.348	1.074	0.897	0.029	3.881
		(0.008)	(0.040)	(0.000)	(0.041)	(0.584)	(0.638)	(0.864)	(0.567)
	ZARUSD	-0.028	0.353	0.203	0.528	8.709	3.035	37.548	68.443
		(0.046)	(0.082)	(0.034)	(0.059)	(0.013)	(0.219)	(0.000)	(0.000)
	Average	0.080	0.177	0.165	0.422				

Notes: The left-hand panel of the table reports variance contributions of different risk prices based on the estimate of regression (15). The variance contribution of risk price x_1 is $CV(\hat{b}_1 \sum_{i=1}^{60}(\hat{\text{E}}_{t+1} - \hat{\text{E}}_t)er^{\text{x}_1}_{t+1+i}, er^j_{t+1})/V(er^j_{t+1})$, where $\hat{b}_1$ are the OLS estimates of b_1 from regression (15), and $(\hat{\text{E}}_{t+1} - \hat{\text{E}}_t)er^{\text{x}_1}_{t+1+i}$ are the ARMA-based forecast revisions. The contributions are computed as the estimated slope coefficient from a regression of $\hat{b}_1 \sum_{i=1}^{60}(\hat{\text{E}}_{t+1} - \hat{\text{E}}_t)$ on er^j_{t+1}. Robust standard errors are shown in parenthesis below the slope estimates. Contributions from multiple risk prices (i.e., for x_1 and x_2) are computed from regressions of $\hat{b}_1 \sum_{i=1}^{60}(\hat{\text{E}}_{t+1} - \hat{\text{E}}_t)er^{\text{x}_1}_{t+1+i} + \hat{b}_2 \sum_{i=1}^{60}(\hat{\text{E}}_{t+1} - \hat{\text{E}}_t)er^{\text{x}_2}_{t+1+i}$ on er^j_{t+1}. Column (i) reports contributions of the prices of carry and momentum risk, column (ii) reports the contributions of the two liquidity risk prices, and column (iii) reports the contribution of the volatility risk-price. The right-hand panel of the table reports Wald tests for the null that $b_{\text{x}} = \beta^j_{\text{x}}$ for factors $\text{x} = \{1, .., 5\}$, with P-value in parenthesis. Columns (v) - (vii) report tests on pairs of coefficients that correspond to the risk prices in columns (i) - (iii).

contribution is 42 percent, but there is a significant range of estimates from 6 percent for the HKDUSD to 77 percent for the RUBUSD. With the exception of the HKDUSD, these estimated contributions are statistically significant. This finding confirms the idea that for most currencies excess returns respond to news about future risk prices. Although there are some variations across currencies, on average news about liquidity risk makes the largest variance contribution to excess returns, estimated at 18 percent. The average contribution for volatility risk is 16 percent with carry and momentum risk accounting for 8 percent. There are also noteworthy differences in the variance contributions across the three Groups of currencies. In particular, the contribution of liquidity risk rises from 10 percent in Group A, to 15 percent in Group B to 26 percent in Group C. In contrast, there are no significant differences in the contributions of carry and moment risk prices between Groups A and B, or between the contributions of volatility risk-prices between Groups A and C.

The tests for the coefficient restrictions reported in the right-hand-panel of the table reveal that there is an inconsistency between how revisions in expected future risk prices affect currency returns, and the betas that measure the contemporaneous impact of factor returns. Column (viii) shows that tests of the null hypothesis that $b_{\text{x}} = \beta^{j}_{\text{x}}$ for factors $\text{x} = \{1, .., 5\}$ are strongly rejected across 16 of the 17 currency pairs. However, in many cases, these test results reflect large differences between the coefficient on the volatility risk price estimated in regression (15) and the volatility beta estimated in regression (10). As column (vii) shows, the equality of these estimates is strongly rejected in all but two currencies. Tests of the restrictions on the carry and momentum coefficients generate more mixed results; the restrictions are rejected at the 5 percent level in 10 of the 17 currency pairs. In contrast, tests of the restrictions on the liquidity risk coefficients are insignificant at the 5 percent level except for the JPYUSD.

The results in Table 7 show that news about future risk prices, identified by the ARMA models, make economically meaningful contributions to the month-by-month movements in individual currency returns. Furthermore, for most currency pairs, news about the future prices of liquidity and volatility risk appear to be more important drivers of excess returns than the prices of the carry and momentum risk factors, which have been the focus of earlier research. These findings show that FX prices respond to news, but it is not the news about the fundaments driving interest rates that featured in textbook models. Here news about the future price of carry risk embodies information about future interest rates insofar as they determine the composition of carry portfolio returns, but this does not appear to be an economically important driver of most FX prices. According to Table 7, news about future liquidity and intraday volatility are far more

regression (15), and $(\hat{\text{E}}_{t+1} - \hat{\text{E}}_t)er^{\text{x}_1}_{t+1+i}$ are the ARMA-based forecast revisions. Variance contributions from two risk prices x_1 and x_2 are computed as $CV(\hat{b}_1 \sum_{i=1}^{60}(\hat{\text{E}}_{t+1} - \hat{\text{E}}_t)er^{\text{x}_1}_{t+1+i}, er^{j}_{t+1})/V(er^{j}_{t+1}) + CV(\hat{b}_2 \sum_{i=1}^{60}(\hat{\text{E}}_{t+1} - \hat{\text{E}}_t)er^{\text{x}_2}_{t+1+i}, er^{j}_{t+1})/V(er^{j}_{t+1})$.

important.

3.4 Further Results

I extended the analysis in several directions to check on the robustness of the results presented above. First, I considered whether the behavior of the factor portfolio returns varied significantly with the exclusion of particular currency pairs. Second, I examined whether the risk exposures measured by the currency betas apply equally to investors with long or short foreign currency positions. Third, I compared the in-sample ARMA forecasts for the risk prices with out-of-sample forecasts.

3.4.1 Factor Portfolio Composition

The currency beta estimates in Table 4 are computed from regression of individual currency returns er^j_{t+1} on the factor portfolio returns er^{x}_{t+1} that are constructed from all 17 of the currency pairs. Consequently, in some times periods the excess return on currency j, er^j_{t+1}, can appear on both sides of the regression. Although the factor portfolios have a high turnover of individual currencies, and the mean duration of individual currencies in the portfolios is short, it is still possible that the incidence of er^j_{t+1} on both sides of the regression equation is high enough to produce spurious coefficient estimates. To investigate this possibility, I re-estimated the regressions for each currency pair j, using alternate versions of the five factor portfolios that excluded pair j. The results from this exercise are very similar to the results in Table 4 (see Appendix Table A.1 for details). Differences between the beta estimates are typically smaller than the standard errors, and the statistically significant betas are generally similar. One notable exception to this pattern appears in the RUBUSD regression, where the beta for the volatility factor is smaller and no longer statistically significant. However, as I noted earlier, these estimates are based on comparatively few monthly observations, so it is unclear whether this difference is due to a spurious correlation or small-sample instability in the beta estimates.

3.4.2 Risk Exposures

All the empirical results presented above are based on currency returns compute from the mid-points of future bid and ask spot FX prices. As such, these returns do not accurately capture the price an investor would received when closing either a short or long position in the foreign currency. Unfortunately, it is impossible to compute returns that capture the transaction prices that are relevant for closing all positions because execution prices depend on the size of the trade (and in many instances the identity of the investor).[17] Nevertheless, we can use the bid and ask prices to investigate whether investors with different foreign currency positions would face markedly different amounts of risk than is implied by the betas in Table 4.

Let $er_{t+1}(\mathrm{L})$ and denote the log excess return on a long foreign currency position. According to the beta representation in equation (A.1), the expect return on this position satisfies $\mathrm{E}_t er_{t+1}(\mathrm{L}) + \frac{1}{2}\mathrm{V}_t(er_{t+1}(\mathrm{L})) = \mu_t' \left(\Omega_t^{\mathrm{x}}\right)^{-1} \mathrm{CV}_t(\mathbf{er}_{t+1}^{\mathrm{x}}, er_{t+1}(\mathrm{L}))$. Notice that the covariance term only depends on the covariance between the factor returns and the transaction price $t+1$, which in this case is a bid price P_{t+1}^b. So relevant measure of exposure is given by $\mathrm{CV}_t(\mathbf{er}_{t+1}^{\mathrm{x}}, lnP_{t+1}^b)$. For an short position in foreign currency, the relevant exposure is $\mathrm{CV}_t(\mathbf{er}_{t+1}^{\mathrm{x}}, lnP_{t+1}^a)$, where P_{t+1}^a is an ask price. Consequently, the difference in exposure between investors holding long and short foreign currency positions depends on $\mathrm{CV}_t(\mathbf{er}_{t+1}^{\mathrm{x}}, lnP_{t+1}^b - lnP_{t+1}^a) = -\mathrm{CV}_t(\mathbf{er}_{t+1}^{\mathrm{x}}, sprd_{t+1})$. To investigate the size of these covariances, I estimate regressions of the bid-ask spread for currency j on a constant and the vector of factor returns $\mathbf{er}_{t+1}^{\mathrm{x}}$:

$$sprd_{t+1}^j = \beta_o + (\mathbf{er}_{t+1}^{\mathrm{x}})' \beta_{sprd}^j + u_{t+1}^j. \tag{16}$$

If all investors face the same exposure to each risk factor X irrespective of whether they hold long or short foreign currency positions, the spread betas in the vector β_{sprd}^j should be insignificantly different from zero.

The results from estimating these regressions are reported in Appendix Table A.2.[18] Overall, the spread beta estimates are very close to zero and statistically insignificant for most currencies. One notable exception are the NOKUSD betas for the forward discount and depth betas. The estimates imply that investors with long NOK positions have approximately 20 percent greater exposure to carry risk than those with short NOK positions, while the opposite is true for exposure to liquidity risk identified by the depth factor. That said, estimated betas for the forward discount and depth reported in Table 4 are only equal to -0.17 and -0.13, so in absolute terms the differences in exposure identified by the spread betas are not particularly significant from an economic perspective. In sum, therefore, the estimates of regression (16) show that the estimated betas in Table 4 are quite representative of the risk exposures faced by investors with reasonably long or short foreign currency positions.

[17] Spot FX trading on the wholesale trading platforms run by EBS and Reuters allow for pre-trade anonymity, so a trader's identity does not effect the execution price, but on the trading platforms operated by large banks (e.g. Barclays' BARX, or Deutsche Bank's Autobahn), which dominate trading between banks and their customers, algorithms quote prices based on the identity of the customer.

[18] The estimates use the best bid and ask prices on the Hotspot trading platform, so strictly speaking the results only apply for long and short positions that are smaller than the depth at the top of the limit order book (see Table 1 for summary statistics).

3.4.3 Risk Price Dynamics

The results in Tables 6 and 7 use forecasts for the factor portfolio returns computed from the ARMA models estimated over the entire sample period. In view of the long tradition in exchange-rate research of using out-of-sample forecasts, it seems appropriate to consider whether out-of-sample forecasts from the ARMA models would produce significantly different results. To this end, I compared out-of-sample and in-sample ARMA forecasts for each of the factor portfolio returns from January 2010 onwards. In particular, the out-of-sample forecasts for portfolio return $er^{\mathrm{x}}_{\tau+i}$ for $i > 0$ are computed from an ARMA model estimated on data from $t = 1$ to τ, whereas the in-sample forecasts combined data from $t = 1$ to τ with the ARMA parameters estimated over the entire sample. The correlations between the two one-month ahead forecasts are high across the five factor portfolios, ranging from 0.72 for the forward discount to 0.92 for the depth factor. At longer horizons, the correlations are even larger. While these findings indicate that there are differences between the in-sample and out-of-sample forecasts, the differences do not greatly affect the relative contributions of the factor risk prices to the variability of individual currency returns. Using the out-of-sample forecasts, news concerning the prices of liquidity and intraday volatility risk contribute more to the variability of individual currency returns than news concerning the risk prices for carry and momentum, which is consistent with the results in Table 7.

4 Conclusion

This paper has studied how the risks associated with the lack of liquidity in spot FX trading contribute to the dynamics of currency returns. The main finding to emerge is that liquidity risk matters. More precisely, I have shown that the risk premia on all the 17 foreign currencies studied compensate investors for exposure to liquidity risk, measured by the betas on one or more of the three liquidity risk factors. This finding applies to the risk premia on major currency pairs that are widely considered to trade in highly liquid markets. Furthermore, for many currency pairs, exposure to liquidity risk appears to be more important than exposure to the traditional carry and momentum risk factors that have been the focus of earlier research. I also found that variations in the price of liquidity risk make economically important contributions to the behavior of individual foreign currency returns. These variations account for approximately 34 percent, on average, of the variability in currency returns compared to the contribution of approximately 8 percent from variations in the prices of carry and momentum risk.

References

Acharya, Viral V and Lasse Heje Pedersen. 2005. "Asset pricing with liquidity risk." *Journal of Financial Economics* 77 (2):375–410. 3

Bank of International Settlements. 2016. "Triennial Central Bank Survey of foreign exchange and OTC derivatives markets in 2016." Tech. rep., Bank of International Settlements. 12

Banti, Chiara, Kate Phylaktis, and Lucio Sarno. 2012. "Global liquidity risk in the foreign exchange market." *Journal of International Money and Finance* 31 (2):267 – 291. 2, 3

Bessembinder, Hendrik. 1994. "Bid-Ask Spreads in the Interbank Foreign Exchange Markets." *Journal of Financial Economics* 35 (3):317–348. 3

Bollerslev, Tim and Michael Melvin. 1994. "Bid-Ask Spreads and Volatility in the Foreign Exchange Market: An Empirical Analysis." *Journal of International Economics* 36:355–372. 3

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen. 2008. "Carry Trades and Currency Crashes." *NBER Macroeconomics Annual* 23 (1):pp. 313–348. 2

Burnside, Craig. 2011. "The cross section of foreign currency risk premia and consumption growth risk: Comment." *American Economic Review* 101 (7):3456–76. 10

———. 2019. "Exchange Rates, Interest Parity, and the Carry Trade." In *Oxford Research Encyclopedia of Economics and Finance*. 4, 10, 19

Burnside, Craig, Martin S. Eichenbaum, and Sergio Rebelo. 2011. "Carry Trade and Momentum in Currency Markets." *Annual Review of Financial Economics* 3:511–535. 4

Campbell, John Y, Sanford J Grossman, and Jiang Wang. 1993. "Trading volume and serial correlation in stock returns." *The Quarterly Journal of Economics* 108 (4):905–939. 3

Campbell, John Y. and Luis M. Viceira. 2002. *Strategic Asset Allocation*. Oxford University Press. 14

Chordia, Tarun, Avanidhar Subrahmanyam, and V Ravi Anshuman. 2001. "Trading activity and expected stock returns." *Journal of financial Economics* 59 (1):3–32. 3

Christiansen, Charlotte, Angelo Ranaldo, and Paul Söderlind. 2011. "The time-varying systematic risk of carry trade strategies." *Journal of Financial and Quantitative Analysis* 46 (04):1107–1125. 4

Copeland, Thomas E. and Dan Galai. 1983. "Information effects on the bid-ask spread." *Journal of Finance* 38 (5):1457–1469. 3

Daniel, Kent, Robert J Hodrick, and Zhongjin Lu. 2014. "The carry trade: Risks and drawdowns." Tech. rep., National Bureau of Economic Research. 4

De Santis, Roberto A and Fabio Fornari. 2008. "Does Business Cycle Risk Account for Systematic Returns from Currency Positioning? The International Perspective." *The International Perspective (June 3, 2008)* . 4

Della Corte, Pasquale, Tarun Ramadorai, and Lucio Sarno. 2016. "Volatility risk premia and exchange rate predictability." *Journal of Financial Economics* 120 (1):21–40. 4

Dobrynskaya, Victoria. 2014. "Downside market risk of carry trades." *Review of Finance* 18 (5):1885–1913. 4

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan. 2018. "Deviations from covered interest rate parity." *The Journal of Finance* 73 (3):915–957. 6

Evans, Martin D. D. and Richard K. Lyons. 2002. "Time-Varying Liquidity in Foreign Exchange." *Journal of Monetary Economics* 49 (5):1025–1051. 3

Evans, Martin D.D. 2017. "External balances, trade and financial conditions." *Journal of International Economics* 107:165 – 184. 16

Evans, Martin DD. 2018. "Forex trading and the WMR fix." *Journal of Banking & Finance* 87:233–247. 3

Evans, Martin DD, Peter O'Neill, Dagfinn Rime, and Jo Saakvitne. 2018. "Fixing the Fix? Assessing the Effectiveness of the 4pm Fix Benchmark." *FCA Occasional Paper* . 3

Farhi, Emmanuel, Samuel Paul Fraiberger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan. 2009. "Crash Risk in Currency Markets." Working Paper 15062, National Bureau of Economic Research. 4

Gavazzoni, Federico, Batchimeg Sambalaibat, and Chris Telmer. 2013. "Currency risk and pricing kernel volatility." . 4

Gilmore, Stephen and Fumio Hayashi. 2011. "Emerging market currency excess returns." *American Economic Journal: Macroeconomics* 3 (4):85–111. 4

Hasbrouck, Joel. 2009. "Trading costs and returns for US equities: Estimating effective costs from daily data." *The Journal of Finance* 64 (3):1445–1477. 3

Hassan, Tarek et al. 2012. "Carry trade: Beyond the forward premium puzzle." In *2012 Meeting Papers*, 124. Society for Economic Dynamics. 4

Hsieh, David A. and Allan W. Kleidon. 1996. "Bid-Ask Spreads in Foreign Exchange Markets: Implications for Models of Assymetric Information." In *The Microstructure of Foreign Exchange Markets*, edited by Jeffrey A. Frankel, Giampaolo Galli, and Alberto Giovannini. Chicago, 41–65. 3

Hu, Grace Xing, Jun Pan, and Jiang Wang. 2013. "Noise as information for illiquidity." *The Journal of Finance* 68 (6):2341–2382. 2, 4

Jurek, Jakub W. 2014. "Crash-neutral currency carry trades." *Journal of Financial Economics* 113 (3):325–347. 4

Korajczyk, Robert A and Ronnie Sadka. 2008. "Pricing the commonality across alternative measures of liquidity." *Journal of Financial Economics* 87 (1):45–72. 3

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." 53 (6):1315–1335. 3

Lettau, Martin, Matteo Maggiori, and Michael Weber. 2014. "Conditional risk premia in currency markets and other asset classes." *Journal of Financial Economics* 114 (2):197 – 225. 4

Lustig, H., H. Roussanov, and A. Verdelhan. 2011. "Common Risk Factors in Currency Markets." *Review of financial studies* 24 (11):3731–3777. 4, 10, 13, 16

Lustig, Hanno and Adrien Verdelhan. 2007. "The cross section of foreign currency risk premia and consumption growth risk." *The American economic review* 97 (1):89–117. 4, 10, 13, 16

Mancini, Loriano, Angelo Ranaldo, and Jan Wrampelmeyer. 2013. "Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums." *The Journal of Finance* 68 (5):1805–1841. 3

Melvin, Michael and Mark P Taylor. 2009. "The crisis in the foreign exchange market." *Journal of International Money and finance* 28 (8):1317–1330. 3

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. 2012a. "Carry trades and global foreign exchange volatility." *The Journal of Finance* 67 (2):681–718. 2, 4, 16

———. 2012b. "Currency momentum strategies." *Journal of Financial Economics* 106 (3):660 – 684. 4, 10

Mueller, Philippe, Andreas Stathopoulos, and Andrea Vedolin. 2017. "International correlation risk." *Journal of Financial Economics* 126 (2):270–299. 4

Pástor, L'uboš and Robert F Stambaugh. 2003. "Liquidity risk and expected stock returns." *Journal of Political economy* 111 (3):642–685. 2

Pástor, Ľuboš and Robert F. Stambaugh. 2003. "Liquidity Risk and Expected Stock Returns." 111 (3):642–685. 3

Verdelhan, Adrien. 2010. "A Habit-Based Explanation of the Exchange Rate Risk Premium." *The Journal of Finance* 65 (1):123–146. 4

———. 2018. "The share of systematic variation in bilateral exchange rates." *The Journal of Finance* 73 (1):375–418. 16

Appendix

Additional empirical results are shown below.

Table A.1: Currency Betas from Factors with Excluded Currencies

		Factors					
		$\hat{r}_t - r_t$ (i)	D_t (ii)	$1/\nabla_t$ (iii)	μ_t (iv)	$1/\sigma_t$ (v)	R^2/SE
Group A:	EURUSD	0.388***	-0.140	0.567***	0.312*	0.544***	0.476
		(0.150)	(0.155)	(0.142)	(0.171)	(0.155)	27.672
	JPYUSD	0.279**	-0.275*	0.074	-0.077	-0.063	0.115
		(0.137)	(0.160)	(0.156)	(0.155)	(0.128)	30.905
	CHFUSD	0.585***	-0.222	0.457*	0.245	0.317	0.230
		(0.207)	(0.214)	(0.245)	(0.199)	(0.229)	35.231
	GBPUSD	-0.048	0.020	0.233	0.125	0.447***	0.251
		(0.141)	(0.172)	(0.160)	(0.128)	(0.110)	27.504
Group B:	AUDUSD	0.655***	0.070	0.977***	0.297*	0.687***	0.559
		(0.178)	(0.175)	(0.172)	(0.165)	(0.168)	33.499
	CADUSD	-0.195	-0.023	0.382***	0.123	0.657***	0.488
		(0.154)	(0.145)	(0.131)	(0.120)	(0.108)	25.713
	HKDUSD	-0.004	-0.014*	0.009	0.013**	-0.006	0.090
		(0.005)	(0.008)	(0.007)	(0.007)	(0.005)	1.124
	NOKUSD	-0.136	-0.075	0.411***	0.102	0.782***	0.387
		(0.183)	(0.173)	(0.152)	(0.223)	(0.191)	33.440
	NZDUSD	0.681**	0.002	0.924***	0.355	0.727***	0.376
		(0.281)	(0.267)	(0.199)	(0.256)	(0.176)	42.369
	SEKUSD	0.208	-0.026	0.415**	0.345*	0.735***	0.354
		(0.186)	(0.219)	(0.216)	(0.193)	(0.195)	34.422
	SGDUSD	0.296***	0.046	0.344***	-0.030	0.338***	0.484
		(0.122)	(0.091)	(0.138)	(0.107)	(0.109)	15.361

		Factors					
		$\hat{r}_t - r_t$ (i)	D_t (ii)	$1/\nabla_t$ (iii)	μ_t (iv)	$1/\sigma_t$ (v)	R^2/SE
Group C:	CNHUSD	-0.270**	0.106	-0.352***	-0.140	0.074	0.319
		(0.129)	(0.173)	(0.125)	(0.133)	(0.084)	8.919
	MXNUSD	0.102	0.147	0.552***	0.340	0.505***	0.553
		(0.137)	(0.123)	(0.146)	(0.161)	(0.142)	26.192
	RUBUSD	-2.092**	-2.116***	2.145***	-1.406**	0.146	0.403
		(0.941)	(0.847)	(0.786)	(0.654)	(0.591)	64.003
	TRYUSD	0.054	0.268	0.640***	0.388	0.398*	0.540
		(0.195)	(0.193)	(0.221)	(0.250)	(0.212)	31.154
	PLNUSD	-0.221	-0.672*	0.586**	0.784***	1.378***	0.406
		(0.335)	(0.367)	(0.279)	(0.269)	(0.332)	53.617
	ZARUSD	0.964***	0.417	0.814***	0.399*	0.538**	0.446
		(0.327)	(0.275)	(0.262)	(0.233)	(0.225)	40.435

Average Betas

	(i)	(ii)	(iii)	(iv)	(v)	R^2/SE
Group A	0.301	-0.155	0.333	0.151	0.311	0.268
Group B	0.215	-0.003	0.494	0.172	0.560	0.391
Group C	-0.244	-0.308	0.731	0.061	0.507	0.445
All	0.073	-0.146	0.540	0.128	0.483	0.381

Notes: The table reports estimates of currency betas for each of the five factors. The betas for currency j are estimated from the regression:

$$er^j_{t+1} = \beta_0 + (\mathbf{er}^{\mathrm{x}}_{t+1})'\beta^j + u^j_{t+1}$$

where $\mathbf{er}^{\mathrm{x}}_{t+1}$ is the 5×1 vector of log excess returns differences between high and low factor portfolios for each of the five factors x: the the implied interest differential $\hat{r}_t - r_t$, depth D_t, the reciprocal of the spread $1/\nabla_t$, the price trend μ_t, and the reciprocal of volatility $1/\sigma_t$. Unlike Table 4, the returns on the factor portfolios are constructed without currency j, so er^j_{t+1} only appears on the left-hand-side of the regression. Robust standard errors are shown in parenthesis below the beta estimates. The R^2 statistic and regression standard error (SE) are reported in the right-hand column of each block. The lower right-hand block of the table reports averages of the beta estimates for each each group of currencies and across all currencies.

Table A.2: Currency Spread Betas

		Factors					
		fd_t (i)	D_t (ii)	$sprd_t^{-1}$ (iii)	μ_t (iv)	σ_t^{-1} (v)	R^2/SE
Group A:	EURUSD	0.000	0.001	-0.001	0.001	0.000	0.012
		(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	0.237
	JPYUSD	-0.001	0.002*	-0.002	0.001	0.000	0.004
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	0.849
	CHFUSD	-0.002	0.002	-0.002	0.001	-0.001	0.125
		(0.002)	(0.001)	(0.002)	(0.001)	(0.001)	0.208
	GBPUSD	-0.001	0.001	-0.001	0.001	0.000	0.079
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	0.138
Group B:	AUDUSD	0.000	0.002	-0.001	0.002	0.001	0.004
		(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	1.043
	CADUSD	-0.001	0.004*	-0.002	0.002	0.001	0.125
		(0.004)	(0.003)	(0.003)	(0.002)	(0.001)	0.300
	HKDUSD	0.000	0.000	0.000	0.001***	0.000	0.186
		(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	0.065
	NOKUSD	-0.038***	0.026***	-0.032***	0.000	0.007	0.318
		(0.015)	(0.011)	(0.012)	(0.006)	(0.007)	1.489
	NZDUSD	-0.004	0.007*	-0.005	0.003	0.000	0.061
		(0.006)	(0.004)	(0.005)	(0.003)	(0.003)	0.869
	SEKUSD	-0.014	0.008	-0.010	-0.002	0.000	0.225
		(0.009)	(0.006)	(0.007)	(0.003)	(0.004)	0.697
	SGDUSD	-0.006*	0.003	-0.002	-0.002	-0.005	0.064
		(0.003)	(0.003)	(0.003)	(0.004)	(0.006)	0.762

		Factors					
		fd_t (i)	D_t (ii)	$sprd_t^{-1}$ (iii)	μ_t (iv)	σ_t^{-1} (v)	R^2/SE
Group C:	CNHUSD	0.008	0.001	0.000	-0.007	0.006	0.138
		(0.007)	(0.003)	(0.006)	(0.004)	(0.007)	0.491
	MXNUSD	-0.012*	0.007	-0.006	0.003	-0.005	0.145
		(0.007)	(0.005)	(0.006)	(0.006)	(0.006)	0.960
	RUBUSD	-0.775*	0.398	-0.187	-0.308	0.077	0.091
		(0.431)	(0.347)	(0.167)	(0.390)	(0.221)	54.042
	TRYUSD	-0.243*	0.114	-0.160	-0.050	0.029	0.299
		(0.140)	(0.079)	(0.104)	(0.054)	(0.045)	9.522
	PLNUSD	-0.005	0.094	0.010	0.514	0.119	0.028
		(0.224)	(0.147)	(0.219)	(0.601)	(0.233)	68.753
	ZARUSD	0.024	-0.025*	0.012	-0.020*	-0.011	0.172
		(0.022)	(0.015)	(0.018)	(0.011)	(0.014)	2.547
	Average Betas						
	Group A	-0.001	0.001	-0.001	0.001	0.000	0.055
	Group B	-0.009	0.007	-0.007	0.001	0.001	0.141
	Group C	-0.167	0.098	-0.055	0.022	0.036	0.146
	All	-0.063	0.038	-0.023	0.008	0.013	0.122

Notes: The table reports estimates of the spread betas for each of the five factors. The betas for the spread on currency j are estimated from the regression:

$$sprd^j_{t+1} = \beta_o + (\mathbf{er}^{\mathrm{x}}_{t+1})'\beta^j_{sprd} + u^j_{t+1}$$

where $\mathbf{er}^{\mathrm{x}}_{t+1}$ is the 5×1 vector of log excess returns differences between high and low factor portfolios for each of the five factors. Robust standard errors are shown in parenthesis below the beta estimates. The R^2 statistic and regression standard error (SE) are reported in the right-hand column of each block. The lower right-hand block of the table reports averages of the beta estimates for each each group of currencies and across all currencies.