Journal of Financial Economics 111 (2014) 527–553

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Countercyclical currency risk premia ☆

Hanno Lustig [a,d], Nikolai Roussanov [b,d,*], Adrien Verdelhan [c,d]

[a] Anderson School of Management, University of California at Los Angeles, Box 951477, Los Angeles, CA 90095, United States
[b] The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104, United States
[c] MIT Sloan, 100 Main Street, Cambridge, MA 02139, United States
[d] National Bureau of Economic Research, 1050 Massachusetts Ave., Cambridge, MA 02139, United States

A R T I C L E I N F O

Article history:
Received 13 June 2012
Received in revised form
28 January 2013
Accepted 25 February 2013
Available online 19 December 2013

JEL classification:
G12
G15
F31

Keywords:
Exchange rates
Forecasting
Risk

A B S T R A C T

We describe a novel currency investment strategy, the 'dollar carry trade,' which delivers large excess returns, uncorrelated with the returns on well-known carry trade strategies. Using a no-arbitrage model of exchange rates we show that these excess returns compensate U.S. investors for taking on aggregate risk by shorting the dollar in bad times, when the U.S. price of risk is high. The countercyclical variation in risk premia leads to strong return predictability: the average forward discount and U.S. industrial production growth rates forecast up to 25% of the dollar return variation at the one-year horizon. The estimated model implies that the variation in the exposure of U.S. investors to worldwide risk is the key driver of predictability.

1. Introduction

Currency carry trades, which go long in baskets of currencies with high interest rates and short in baskets of currencies with low interest rates, have been shown to deliver high Sharpe ratios. In particular, the dollar-neutral high-minus-low carry trade which only uses the ranking of foreign interest rates to build portfolios and hence ignores all information in the level of short-term U.S. interest rates, has received lots of attention. Our paper examines a different investment strategy that exploits the time-series variation in the average U.S. interest rate difference vis-à-vis the rest of the world: this strategy goes long in a basket of foreign currencies and short in the dollar whenever the average foreign short-term interest rate is above the U.S. interest rate, typically during U.S. recessions, while it shorts all foreign currencies and takes long positions in the dollar otherwise. This simple investment strategy, which we refer to as the 'dollar carry trade,' produces Sharpe ratios in excess of 0.50, higher than those on both the high-minus-low portfolio carry trades and the U.S. stock market over the same sample.

We develop a no-arbitrage asset pricing model to show how the dollar carry trade exploits the connection between U.S. short-term interest rates and the volatility of the U.S. pricing kernel. When the volatility of the U.S. pricing kernel is high, U.S. short-term interest rates tend to be low relative to the rest of the world, because of large precautionary savings and increased demand for liquidity.

☆ The authors thank Gurdip Bakshi, Pasquale Della Corte, Urban Jermann, Karen Lewis, Tarun Ramadorai, Frans De Roon, Chris Telmer, Bill Schwert (the editor), Ivan Shaliastovich, Rob Stambaugh, Amir Yaron, the anonymous referee, and seminar participants at many institutions and conferences for helpful comments. Roussanov acknowledges financial support from Wharton Global Initiative and Cynthia and Bennett Golub Endowed Faculty Scholarship.

* Corresponding author at: The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104, United States. Tel.: +1 215 746 0004.
E-mail address: nroussan@wharton.upenn.edu (N. Roussanov).

0304-405X/$ - see front matter
http://dx.doi.org/10.1016/j.jfineco.2013.12.005

As a result, U.S. investors in the dollar carry strategy are long in foreign currencies and short in the dollar when the U.S. pricing kernel is more volatile than foreign pricing kernels. This strategy is risky, because the absence of arbitrage implies that the dollar appreciates in the case of a bad shock to the U.S. pricing kernel, when its volatility is higher than abroad. U.S. investors in the dollar carry strategy thus bear the risk of a dollar *appreciation* in bad times, when they are long foreign currencies and short in the dollar. When U.S. short-term interest rates are high relative to the rest of the world, the dollar carry trade takes a short position in foreign currencies and a long position in the dollar: investors then bear the risk of a dollar *depreciation* in the case of a good innovation to the U.S. pricing kernel.

Hence, the expected excess returns on a long position in foreign currency, funded by a short position in the dollar, should be high in bad times for the U.S., but low or even negative in good times. U.S. investors collect a positive currency risk premium because they are betting against their own intertemporal marginal rates of substitution.

We document new evidence of predictability for the returns on a basket of foreign currencies funded by a short position in the dollar that is consistent with counter-cyclical variation in currency risk premia. This evidence accounts for the profitability of the dollar carry trade strategy. A version of our model estimated to match the dynamics and the cross-section of interest rates and exchange rates generates a large dollar carry trade risk premium. The parameter estimates imply that the key novel feature of our model—time-variation in U.S.-specific exposure to global risk—is the main driver of currency return predictability and the dollar carry risk premium.

The key predictor in our study is the average forward discount on foreign currency against the U.S. dollar which is the difference between the average short-term interest rate across a broad set of developed countries and the U.S. short-term interest rate. The one-month-ahead average forward discount on foreign currency against the dollar explains around 3% of the variation in the foreign currency excess returns on a basket of developed country currencies over the next month. As the horizon increases, the R^2 increases, because the average forward discount is persistent. At the 12-month horizon, the average forward discount explains up to 13% of the variation in returns over the next year. These effects are economically meaningful. As the U.S. economy enters a recession, U.S. investors who short the dollar earn a larger interest rate spread, the average forward discount on foreign currency, and they earn an additional 145 basis points per annum in currency appreciation per 100 basis point increase in the interest rate spread as well. In other words, an increase in the average forward discount of 100 basis points increases the expected excess return by 245 basis points per annum and it leads to an annualized depreciation of the dollar by 145 basis points.

Our predictability findings are not simply a restatement of those documented in the large literature on violations of the uncovered interest parity (UIP) that originated with the classic papers by Hansen and Hodrick (1980) and Fama (1984). We find that the average forward discount has forecasting power at the individual currency level above and beyond that of the currency-specific interest rate differential—both in terms of the slope coefficients and the average R^2. In fact, the average forward discount drives out the bilateral one in a panel regression for developed currencies. Consistent with our predictability results, a version of the carry trade that goes long or short individual currencies based on the sign of the individual forward discount, rather than the average forward discount, only delivers a Sharpe ratio of 0.3 (which becomes essentially zero once transaction costs are taken into account), on the same basket of currencies. The average forward discount on the U.S. dollar against a basket of developed country currencies is a strong predictor of the excess returns on a basket of foreign currencies, even when the basket consists only of emerging markets currencies. All of this evidence points to the economic mechanism behind exchange rate and currency return predictability, namely, variation in the home country-specific price of risk.

The dollar premium is driven by the U.S. business cycle, and it increases during U.S. recessions. The U.S.-specific component of macroeconomic variables such as the year-over-year rate of industrial production growth predicts future excess returns (with a negative sign) on the basket of foreign currencies, even after controlling for the average forward discount. These two predictors deliver in-sample R^2's of 23% at the one-year horizon for the average returns on the basket of developed country currencies, and 25% for the basket including all currencies. The effects are large: a 100 basis point drop in year-over-year U.S. industrial output growth raises the expected excess return, and hence increases the expected rate of dollar depreciation over the following year, by 90 basis points per annum, after controlling for the average forward discount. We also show that the volatility of U.S. consumption growth volatility forecasts dollar returns. As in the model, these macroeconomic variables do not predict the returns on the high-minus-low currency carry trade, which is consistent with the notion that the high-minus-low carry trade premium is determined by the global price of risk in financial markets.

If markets are complete, the percentage change in the spot exchange rate reflects the difference between the log of the domestic and the foreign pricing kernels. As a simple thought experiment, we can decompose the log pricing kernels, as well as the returns, into a country-specific component and a global component. In a well-diversified currency portfolio, the foreign country-specific risk averages out, and the U.S. investor holding this portfolio is compensated only for bearing U.S.-specific risk and global risk. The high-minus-low carry trade portfolio also eliminates U.S.-specific risk and, hence, the high-minus-low carry premium has to be exclusive compensation for taking on global risk. On the contrary, the dollar carry trade average returns compensate U.S. investors for taking on both U.S.-specific risk and global risk when the price of these risks is high in the U.S. Indeed, the high-minus-low carry trade returns are strongly correlated with changes in global financial market volatility, as shown by Lustig, Roussanov, and Verdelhan (2011), while the dollar carry trade is not. At the same time, the dollar carry trade returns are correlated with the average growth rate of aggregate consumption across countries, a proxy for worldwide macroeconomic risk, and the rate of the U.S.-specific component of industrial production growth.

Most of our paper focuses on the U.S. dollar, but a similar basket-level carry trade can be implemented using any base currency. We call such strategies base carry trades. These base carry trades can be implemented in other currencies, but they only 'work' for base currencies whose forward discounts are informative about the local price of risk, such as the U.S. and the U.K. In other countries, such as Japan, Switzerland, Australia, and New Zealand, the base carry trade is highly correlated with the high-minus-low currency carry trade. Our no-arbitrage model traces out a U-shaped relation between the mean of a country's average forward discount and the correlation between base carry and global carry trade returns that is confirmed in the data.

The paper proceeds as follows: Section 2 discusses the relation of our paper to the existing literature. Section 3 describes the data, the construction of currency portfolios and their main characteristics, and motivates our analysis by presenting a simple investment strategy that exploits return predictability to deliver high Sharpe ratios. Section 4 presents a no-arbitrage model of exchange rates, which belongs to the essentially affine class that is popular in the term-structure literature. Section 5 takes the model to the data. The model matches the key moments of interest rates and exchange rates in the data, reproduces the key features of the dollar carry and high-minus-low carry trade risk premia, and offers an interpretation of our predictability findings. Section 6 shows that macroeconomic variables such as the rate of industrial production growth as well as aggregate consumption growth volatility have incremental explanatory power for future currency basket returns. Section 7 concludes.

2. Related literature

Our paper relates to a large literature on exchange rate predictability that is too vast to survey here.[1] Instead, we limit our literature review to recent work that explores currency return predictability from a finance perspective.[2] In recent work on currency portfolios, Ang and Chen (2010) show that changes in interest rates and term spreads predict currency excess returns, while Chen and Tsang (2013) show that yield curve factors containing information both about bond risk premia and about future macroeconomic fundamentals have forecasting power for individual currencies as well. Adrian, Etula, and Shin (2010) show that the funding liquidity of financial intermediaries in the U.S. predicts currency excess returns on short positions in the dollar, where funding liquidity growth is interpreted as a measure of the risk appetite of these intermediaries. Hong and Yogo (2012) show that the futures open market interest has strong predictive power for returns on a portfolio of currency futures. Bakshi and Panayotov (2013) present evidence that long-short carry trade returns are predictable using a measure of foreign exchange volatility and commodity prices. Asness, Moscowitz, and Pedersen (2013) show that the predictability of foreign exchange returns using lagged exchange rates in the cross-section of currencies is systematically related to momentum and reversals in other asset classes. Our paper is the only one that explicitly links currency return predictability to U.S.-specific business cycle variation.

Our work is also closely related to the literature that documents time-varying risk premia in various asset markets. The ability of the average forward discount to forecast individual exchange rates and returns on other currency baskets echoes the ability of forward rates to forecast returns on bonds of other maturities, as documented by Stambaugh (1988) and Cochrane and Piazzesi (2005). Ludvigson and Ng (2009), Joslin, Priebsch, and Singleton (2010), and Duffee (2011) document that U.S. industrial production growth contains information about bond risk premia that is not captured by interest rates (and, therefore, forward discounts). The industrial production index is highly correlated with the output gap used by Cooper and Priestley (2009) to predict stock returns. We find similar evidence of countercyclical risk premia in currency markets.

Our model of the stochastic discount factor belongs to a class of essentially affine models common in the literature on the term structure of interest rates. Models in this class have been applied to currency markets by Frachot (1996), Backus, Foresi, and Telmer (2001), Brennan and Xia (2006), Lustig, Roussanov, and Verdelhan (2011), and Sarno, Schneider, and Wagner (2012). In our model the bulk of the stochastic discount factor variation (SDF) is common across countries, consistent with Brandt, Cochrane, and Santa-Clara (2006), Bakshi, Carr, and Wu (2008), Colacito (2008), and Colacito and Croce (2011). While ours is a no-arbitrage model, it shares some key features with equilibrium models of currency risk premia that emphasize time-varying volatility of the pricing kernel and its pro-cyclical effect on the short-term interest rate, such as Verdelhan (2010), Backus, Gavazzoni, Telmer, and Zin (2011), and Bansal and Shaliastovich (2012).[3]

While this paper focuses on Gaussian shocks and the first two moments of stochastic discount factors, Gavazzoni, Sambalaibat, and Telmer (2012) argue that future work should consider higher moments. Earlier work on higher moments includes Evans and Lewis (1995), who study the link between peso problems and the estimation of currency risk premia. More recently, Bhansali (2007) documents the empirical properties of hedged carry trade strategies. Brunnermeier, Nagel, and Pedersen (2009) show that risk reversals increase with interest rates. Jurek (2009) provides a comprehensive empirical investigation of hedged carry

[1] This literature is surveyed, for example, in Hodrick (1987) and Lewis (1995).

[2] While our paper focuses on the expected returns on currency portfolios, Campbell, Medeiros, and Viceira (2010) focus on the second moments of currency returns, because they are interested in the risk management demand for individual currencies from the vantage point of U.S. bond and equity investors.

[3] Atkeson and Kehoe (2009) argue that this effect is important for understanding the impact of monetary policy on interest rates.

trade strategies. Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011) use currency options to construct carry trade positions that are hedged at-the-money. Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2012) estimate a no-arbitrage model with disaster risk using a cross-section of currency options. Chernov, Graveline, and Zviadadze (2012) study jump risk at high frequencies.

3. Returns to timing the U.S. dollar

Currency excess returns correspond to simple investment strategies: investors pocket the interest rate difference between two countries, known at the time of their investment, but expose themselves to the risk stemming from exchange rate fluctuations over the investment horizon. The literature has mostly focused on the predictability of excess returns for individual foreign currency pairs. By shifting the focus to investments in baskets of foreign currencies, our paper shows that most of the predictable variation in currency markets is common across currencies.

In this section, we describe our primary data set and give a brief summary of currency returns at the level of currency baskets. We use the quoted prices of traded forward contracts of different maturities to study return predictability at different horizons. Hence, there is no interest rate risk in the investment strategies that we consider. Moreover, these trades can be implemented at fairly low costs.

3.1. Preliminaries

3.1.1. Currency excess returns using forward contracts

We use s to denote the log of the nominal spot exchange rate in units of foreign currency per U.S. dollar, and f for the log of the forward exchange rate, also in units of foreign currency per U.S. dollar. An increase in s means an appreciation of the U.S. dollar. The log excess return rx on buying a foreign currency in the forward market and then selling it in the spot market after one month is simply $rx_{t+1}=f_t-s_{t+1}$. This excess return can also be stated as the log forward discount on foreign currency minus the change in the spot rate: $rx_{t+1}=f_t-s_t-\Delta s_{t+1}$. In normal conditions, forward rates satisfy the covered interest rate parity condition; the forward discount on foreign currency is equal to the interest rate differential: $f_t-s_t \approx i_t^\star - i_t$, where $i^\star$ and i denote the foreign and domestic nominal risk-free rates over the maturity of the contract.[4] Hence, the log currency excess return equals the interest rate differential less the rate of depreciation: $rx_{t+1}=i_t^\star-i_t-\Delta s_{t+1}$.

3.1.2. Horizons

Forward contracts are available at different maturities. We use k-month maturity forward contracts to compute k-month horizon returns (where $k=1$, 2, 3, 6, and 12). The log excess return on the k-month contract for currency i is $rx^i_{t+k}=-\Delta s^i_{t\to t+k}+f^i_{t\to t+k}-s^i_t$, where $f^i_{t\to t+k}$ is the k-month forward exchange rate, and the k-month change in the log exchange rate is $\Delta s^i_{t\to t+k}=s^i_{t+k}-s^i_t$. For horizons above one month, our series consist of overlapping k-month returns computed at a monthly frequency.

3.1.3. Transaction costs

Profitability of currency trading strategies depends on the cost of implementing them. Since we have bid–ask quotes for spot and forward contracts, we can compute the investor's actual realized excess return net of transaction costs. The *net* log currency excess return for an investor who goes long in foreign currency is: $rx^l_{t+1}=f^b_t-s^a_{t+1}$. The investor buys the foreign currency or equivalently sells the dollar forward at the bid price (f^b) in period t, and sells the foreign currency or equivalently buys dollars at the ask price (s^a_{t+1}) in the spot market in period $t+1$. Similarly, for an investor who is long in the dollar (and thus short the foreign currency), the net log currency excess return is given by: $rx^s_{t+1}=-f^a_t+s^b_{t+1}$. For our regression-based analysis we use midpoint quotes for spot and forward exchange rates in constructing excess returns, instead of the net excess returns.

3.1.4. Data

We start from daily spot and forward exchange rates in U.S. dollars (USD). We build end-of-month series from November 1983 to June 2010. These data are collected by Barclays and Reuters and available on Datastream. Our main data set contains the following countries/currencies: Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Finland, France, Germany, Greece, Hong Kong, Hungary, India, Indonesia, Ireland, Italy, Japan, Kuwait, Malaysia, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Saudi Arabia, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, Turkey, United Arab Emirates, United Kingdom, as well as the Euro. The euro series start in January 1999. We exclude the euro area countries after this date and only keep the euro series. Some of these currencies have pegged their exchange rate partly or completely to the U.S. dollar over the course of the sample; for this reason, we exclude Hong Kong, Saudi Arabia, and United Arab Emirates. We also exclude Turkey to avoid our results being driven by near-hyperinflation episodes. Based on large failures of CIP, we deleted the following observations from our sample: South Africa from the end of July 1985 to the end of August 1985; Malaysia from the end of August 1998 to the end of June 2005; and Indonesia from the end of December 2000 to the end of May 2007.

3.1.5. Baskets of currencies

We construct three currency baskets. The first basket is composed of the currencies of developed countries: Australia, Austria, Belgium, Canada, Denmark, France, Finland, Germany, Greece, Italy, Ireland, Japan, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, and the U.K., as well as the Euro. The second basket groups all of the remaining currencies, corresponding to the emerging countries in our sample. The third basket consists of all of the currencies in our sample. All of the average log excess returns

[4] Akram, Rime, and Sarno (2008) study high frequency deviations from covered interest parity (CIP). They conclude that CIP holds at daily and lower frequencies. While this relation was violated during the extreme episodes of the financial crisis in the fall of 2008 (e.g., see Baba and Packer, 2009), including or excluding those observations does not have a major effect on our results.

and average log exchange rate changes are equally weighted within each basket.

The average log excess return on currencies in basket j over horizon k is $\overline{rx}^j_{t\to t+k} = (1/N_t^j)\sum_{i=1}^{N_t^j} rx^i_{t+k}$, where N_t^j denotes the number of currencies in basket j at time t. Similarly, the average change in the log exchange rate is $\overline{\Delta s}^j_{t\to t+k} = (1/N_t^j)\sum_{i=1}^{N_t^j} \Delta s^i_{t\to t+k}$, and the average forward discount (AFD) on foreign currency against the U.S. dollar for maturity k is $\overline{f}^j_{t\to t+k} - \overline{s}^j_t = (1/N_t^j)\sum_{i=1}^{N_t^j} f^i_{t\to t+k} - s^i_t$.

The AFDs are negatively correlated with the U.S. short-term interest rates. However, the AFD is clearly stationary, while U.S. short-term interest rates trend downward from 10% (thee-month Treasury bill rate) to essentially 0%, notably, because of the secular decline in (global) inflation over the sample. The AFDs computed on developed and emerging countries are virtually identical in the first half of the sample, but diverge dramatically during the period around the Asian financial crisis of 1997–1998, with emerging countries′ interest rates shooting up relative to both the U.S. and the developed countries' interest rates. This disparity suggests that one should expect different patterns of predictability for the different baskets.

At the 12-month horizon, the average AFD of foreign currencies against the dollar is 60 basis points per annum in the sample of developed countries. This means that, on average, 12-month interest rates are 60 basis points higher in these countries than in the U.S. In the sample of emerging countries, this average difference is 257 basis points.

The AFDs are persistent, especially at longer horizons: the AFD of developed countries using 12-month forward rates has an autocorrelation of 0.98 at monthly frequency (corresponding to an annualized autocorrelation is 0.78). At shorter horizons, the autocorrelation is much smaller: the average forward discount for developed countries based on one-month forward rates, which we will use to construct the monthly trading strategy, has an autocorrelation of 0.89 at the monthly frequency (i.e., 0.25 annualized). Therefore, the AFDs are less persistent than some of the commonly used return-forecasting variables such as the dividend yield on the U.S. stock market, which has an annualized autocorrelation of 0.96.

3.2. *The dollar carry trade*

We design a simple, implementable investment strategy that exploits the predictability of currency returns by the AFD. Investors go long all foreign currencies when the average foreign currency trades at a forward discount, i.e., when the average foreign interest rate (across all *developed* countries) is above the U.S. short-term interest rate, and short all foreign currencies otherwise. We call this investment strategy the *Dollar* carry trade strategy. We incorporate bid–ask spreads in order to account for the cost of implementing this strategy. As is clear from the top panel of Fig. 1, the dollar carry trade typically shorts the dollar during and after recessions (when the AFD is positive), and goes long the dollar during expansions (when the AFD turns negative), where recessions and expansions are determined by the National Bureau of Economic Research (NBER).

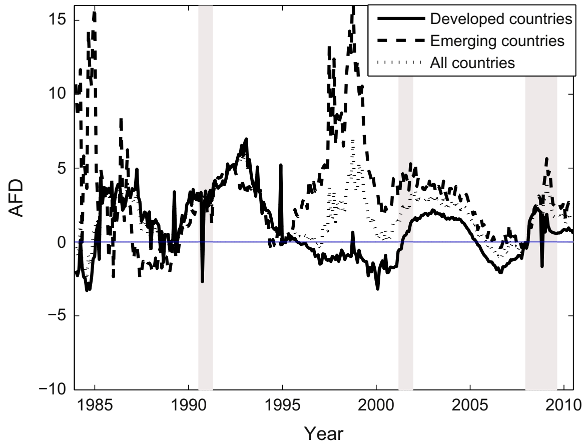

Fig. 1. Average 12-month forward discounts on three currency baskets. This figure presents the average 12-month forward discounts on three currency baskets (developing countries, developed countries, and all countries). The shaded areas are U.S. recessions according to NBER. The sample period is 11/1983–6/2010.

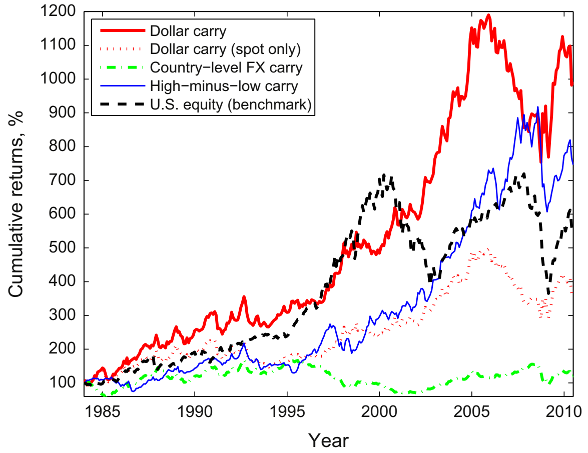

Fig. 2. Carry trade excess return indexes. This figure plots the total return index for four investment strategies, starting at $100 on November 30, 1983. The dollar carry trade goes long all one-month forward contracts in a basket of developed country currencies when the average one-month forward discount for the basket is positive, and short the same contracts otherwise. This strategy is labeled *Dollar carry*. The component of this strategy that is due to the spot exchange rate changes, i.e., excluding the interest rate differential, is *dollar carry* (*spot only*). The individual country-level carry trade is an equal-weighted average of long-short positions in individual currency one-month forward contracts that depend on the sign of the bilateral forward discounts; this strategy is labeled *Country-level FX carry*. The third strategy corresponds to dollar-neutral high-minus-low currency carry trades in one-month forward contracts (*High-minus-low carry*). The fourth strategy, *U.S. equity* (*benchmark*), is simply long the excess return on the CRSP value-weighted U.S. stock market portfolio. All strategies are levered to match the volatility of the stock market.

Fig. 2 reports the return indices on this dollar carry trade strategy compared to other currency trading strategies, as well as the aggregate equity market returns, using both the entire sample of currencies and the smaller subsample of developed countries; all of these were levered to match the volatility of U.S. stock returns.

As an alternative to the dollar carry strategy, we use a similar strategy implemented at the country (or, rather, individual currency) level. For each currency in our sample, investors go long in that currency if the corresponding forward discount is positive, and short otherwise. There is substantial heterogeneity in terms of average excess returns and Sharpe ratios at the individual currency level. We report the equal-weighted average excess return across all currencies, which is a simple measure of returns earned on a diversified portfolio comparable to investing in a broad basket. We call this the *country-level FX* carry trade strategy. We compare these strategies to another popular currency trading strategy, the high-minus-low (HML) currency carry trade, and to U.S. equity market returns. The HML carry trade strategy corresponds to currency carry trades that go long in a portfolio of high interest rate currencies and short in a portfolio of low interest currencies, with no direct exposure to the U.S. dollar. This strategy is implemented using the currency portfolios sorted by interest rate differentials in Lustig, Roussanov, and Verdelhan (2011) extended to our longer sample, with five portfolios for the subsample of developed country currencies and six portfolios including all currencies in our sample.

To compare these strategies, we scale the currency positions so that all currency returns are as volatile as equity returns—this can be thought of as constructing a levered position, except that the weights are based on in-sample estimates of volatility.[5] In our sample, the volatility of U.S. stock market excess returns is 15.5%. An investor starting with $100 in December 1983 in the dollar carry trade would have ended up with $1,467 at the end of the sample. The interest rate component (or carry) accounts for $860 and the rest ($607) is due to the fluctuations in the signed position in the U.S. dollar against the basket of foreign currencies, suggestive of predictable changes in the exchange rate. On the other hand, the HML currency carry trade delivers $356 dollars, while the country-level strategy barely breaks even.

Table 1 reports the means, standard deviations, and Sharpe ratios of the returns on these three investment strategies. We report (in brackets) standard errors on the means. The currency excess returns take into account bid–ask spreads on monthly forward and spot contracts, while equity excess returns do not take into account transaction costs. The standard errors are obtained by resampling the series using the stationary bootstrap procedure of Politis and Romano (1994) in order to capture the autocorrelation and heteroskedasticity properties of the series. The sample average of dollar carry returns, our estimate of the dollar premium, is 5.60% (4.28%) per year for the basket of developed (all) currencies. The annualized Sharpe ratios are 0.66 (0.56), respectively. The exchange rate component of the dollar carry trade strategy (i.e., the part due to the depreciation or appreciation of the dollar and not due to the interest rate differential) delivers an average return of 3.77%, which is statistically different from zero (the bootstrap standard error is 1.67%—these numbers are not reported in the table).

By comparison, the average HML carry trade returns are 3% (4.4%) for the basket of developed (all) currencies, respectively, corresponding to Sharpe ratios of about 0.3 (0.5). Interestingly, the country-level *FX* carry strategy, that has elements of both dollar and HML carry trades, does not perform nearly as well as either of these aggregate strategies with an average return of about 0.5% (which is not statistically significant) and a Sharpe ratio that is close to zero. If bid–ask spreads are not taken into account, the country-level carry strategy does exhibit positive average excess returns with a Sharpe ratio of about 0.3, but the returns on the other two carry strategies computed without transaction costs are even higher, with Sharpe ratios close to 0.9 (not reported in the table for brevity). The dollar carry trade returns also do not exhibit much negative skewness (-0.27, compared to the *HML* carry skewness of -0.98, also not reported in the table).

The right panel of Table 1 reports the mean, standard deviation, and Sharpe ratios of the returns on these three investment strategies scaled to deliver the same volatility as equity markets. The dollar carry strategy delivers an average excess return of over 10.2%, while the country-level strategies deliver an average excess return of below 1%. Using a greater number of signals contained in individual currency pairs' forward discounts does not appear to provide an advantage over a simple strategy that pits the U.S. dollar against a broad basket of currencies. The superior performance of the dollar carry trade relative to the other trading strategies is also apparent from Fig. 2. Given that the dollar carry strategy has essentially zero correlation (unconditionally) with both the HML carry trade and with the stock market, the high average returns and Sharpe ratios earned by this strategy are clear evidence that the average interest rate difference between the U.S. and a broad group of developed countries contains substantial information about risk premia in currency markets.

3.3. *Predictability in currency markets*

The dollar carry trade is highly profitable because the average forward discount on foreign currencies of developed countries against the dollar forecasts basket-level exchange rate changes and returns, even in a horse race with the individual currency pairs' forward discounts.

3.3.1. *Predictability tests*

We run the following regressions of basket-level average log excess returns on the AFD, and of average changes in spot exchange rates on the AFD:

$$\overline{rx}^{j}_{t\to t+k} = \psi_0 + \psi_{\mathbf{f}}(\overline{f}^{Dev}_{t\to t+k} - \overline{s}^{Dev}_t) + \eta_{t+k}, \tag{1}$$

$$-\overline{\Delta s}^{j}_{t\to t+k} = \zeta_0 + \zeta_{\mathbf{f}}(\overline{f}^{Dev}_{t\to t+k} - \overline{s}^{Dev}_t) + \varepsilon_{t+k}, \tag{2}$$

[5] This construction of the levered strategy is for the purpose of illustrating the risk-return trade-off in the currency strategies, and is not implementable in practice as it is based on the ex post standard deviations of exchange rate changes and stock returns. A more sophisticated construction could be based on the lagged realized or contemporaneous implied stock market and FX volatility and rebalanced dynamically.

Table 1

Currency carry trades and equity market excess returns. The table reports the mean, standard deviation, and Sharpe ratios of three carry trade investment strategies in comparison to the U.S. equity market returns. The first strategy (USD, or dollar carry trade) goes long all available one-month currency forward contracts when the average forward discount of *developed countries* is positive, and short the same contracts otherwise. The second strategy (FX, or individual currency carry trade) is similar to the first one, but implemented at the level of individual currencies. For each country, the strategy goes long that currency if the corresponding one-month forward discount is positive, and short otherwise. We report the mean excess return across all countries. The third strategy (HML, or high-minus-low carry trade) is long in a basket of the currencies with the largest one-month forward discounts, and short in a basket of currencies with the lowest one-month forward discounts, with no direct exposure to the U.S. dollar. To construct this strategy, we sort all currencies into six bins (five when we exclude emerging market countries), and we go long in the last portfolio, short in the first, as in Lustig, Roussanov, and Verdelhan (2011). The fourth (equity benchmark) strategy is long the U.S. stock market and short the U.S. risk-free rate. In the left panel, we report the raw moments. In the right panel, we scale each currency strategy such that they exhibit the same volatility as the U.S. equity market. Data are monthly, from Reuters and Barclays (available on Datastream). Equity excess returns are for the CRSP value-weighted stock market index. Excess returns are annualized (means are multiplied by 12 and standard deviations are multiplied by $\sqrt{12}$). Sharpe ratios correspond to the ratio of annualized means to annualized standard deviations. Currency excess returns take into account bid–ask spreads on monthly forward and spot contracts, while equity excess returns do not take into account transaction costs. We report standard errors for all of the quantities (in brackets) obtained by stationary bootstrap. The sample period is 11/1983–6/2010.

	Raw returns				Scaled returns			
	USD	FX	HML	Equity	USD	FX	HML	Equity
Panel A: Developed countries								
Mean	5.60	0.48	3.00	6.26	10.18	0.91	4.77	6.26
	[1.66]	[1.65]	[1.92]	[2.98]	[3.29]	[3.11]	[3.14]	[2.98]
Std. Dev.	8.53	8.24	9.73	15.49	15.49	15.49	15.49	15.49
	[0.42]	[0.39]	[0.62]	[0.92]	[0.92]	[0.92]	[0.92]	[0.92]
Sharpe Ratio	0.66	0.06	0.31	0.40	0.66	0.06	0.31	0.40
	[0.20]	[0.20]	[0.21]	[0.20]	[0.20]	[0.20]	[0.21]	[0.20]
Corr(USD,.)		0.32	−0.03	0.01				
		[0.10]	[0.09]	[0.07]				
Panel B: All countries								
Mean	4.28	0.36	4.41	6.26	8.70	0.72	7.58	6.26
	[1.48]	[1.53]	[1.80]	[3.03]	[3.23]	[3.08]	[3.26]	[3.03]
Std. Dev.	7.61	7.77	9.02	15.49	15.49	15.49	15.49	15.49
	[0.39]	[0.38]	[0.48]	[0.92]	[0.92]	[0.92]	[0.92]	[0.92]
Sharpe Ratio	0.56	0.05	0.49	0.40	0.56	0.05	0.49	0.40
	[0.20]	[0.20]	[0.21]	[0.21]	[0.20]	[0.20]	[0.21]	[0.21]
Corr(USD,.)		0.29	0.01	−0.00				
		[0.11]	[0.08]	[0.07]				

for each basket $j \in \{Dev, Emerg, All\}$. Since the log excess returns are the difference between changes in spot rates at $t+1$ and the AFD at t, for the developed countries basket, $j = Dev$, these two regressions are equivalent and $\psi_{\mathbf{f}} = \zeta_{\mathbf{f}} + 1$. The Uncovered Interest Parity (UIP) hypothesis states that expected changes in exchange rates are equal to interest rate differentials, while currency excess returns are not predictable. With our notation, UIP implies that $\zeta_{\mathbf{f}} = -1$ and $\psi_{\mathbf{f}} = 0$ for $j = Dev$ (thus, Eq. (2) is equivalent to the classic Fama, 1984 regression, up to the sign of the slope coefficient on the average, not country-specific, interest rate difference).

We report t-statistics for the slope coefficients $\psi_{\mathbf{f}}$ and $\zeta_{\mathbf{f}}$ for both asymptotic and finite-sample tests. The AFDs are strongly autocorrelated, albeit less so than individual countries′ interest rates. We use Hansen and Hodrick's (1980) methodology in order to compute asymptotic standard errors with the number of lags equal to the horizon of the forward contract plus one lag. Our results are robust to using instead Newey-West standard errors computed with the optimal number of lags following Andrews (1991) in order to correct for error correlation and conditional heteroskedasticity, as well as Newey-West standard errors using non-overlapping data.

Bekaert, Hodrick, and Marshall (1997) note that the small sample performance of these test statistics is also a source of concern. In particular, due to the persistence of the predictor variable, estimates of the slope coefficient can be biased (as pointed out by Stambaugh, 1999) as well as have wider dispersion than the asymptotic distribution. To address these problems, we compute bias-adjusted small-sample t-statistics, generated by bootstrapping 10,000 samples of returns and forward discounts from a corresponding vector auto-regression (VAR) under the null of no predictability. We resample the residuals in blocks of random lengths, following the stationary bootstrap procedure of Politis and Romano (1994).[6]

[6] Our bootstrapping procedure follows Mark (1995) and Kilian (1999) and is similar to the one recently used by Goyal and Welch (2005) on U.S. stock excess returns. It preserves the autocorrelation structure of the predictors and the cross-correlation of the predictors' and returns' shocks. The random-block resampling allows for the potential heteroskedasticity in residuals while preserving stationarity of the underlying series. Ang and Bekaert (2007) and Bakshi, Panayotov, and Skoulakis (2011) study the power and size properties of different estimation procedures in the context of persistent predictors.

Table 2

Forecasting currency excess returns and exchange rates with the average forward discount. This table reports results of forecasting regressions for average excess returns and average exchange rate changes for baskets of currencies at horizons of one, two, three, six, and 12 months. For each basket we report the R^2, and the slope coefficient $\psi_{\mathbf{f}}$ in the time-series regression of the log currency excess return on the average log forward discount of developed countries, and similarly the slope coefficient $\zeta_{\mathbf{f}}$ and the R^2 for the regressions of average exchange rate changes. The t-statistics for the slope coefficients in brackets are computed using the following methods. *HH* denotes Hansen and Hodrick (1980) standard errors computed with the number of lags equal to the length of overlap plus one lag. The *VAR*-based statistics are adjusted for the small-sample bias using the bootstrap distributions of slope coefficients under the null hypothesis of no predictability, estimated by drawing from the residuals of a VAR with the number of lags equal to the length of overlap plus one lag. Data are monthly, from Barclays and Reuters (available via Datastream). The returns do not take into account bid–ask spreads. The sample period is 11/1983–6/2010.

	Developed countries				Emerging countries				All countries			
	Excess returns		Exchange rates		Excess returns		Exchange rates		Excess returns		Exchange rates	
Horizon	$\psi_{\mathbf{f}}$	R^2	$\zeta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\zeta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\zeta_{\mathbf{f}}$	R^2
1	2.45	2.91	1.45	1.03	2.06	2.21	2.28	2.63	2.19	2.93	1.56	1.51
HH	[2.55]		[1.51]		[2.09]		[2.22]		[2.54]		[1.80]	
VAR	[2.61]		[1.53]		[2.20]		[2.47]		[2.57]		[1.93]	
2	2.49	5.00	1.49	1.86	2.09	3.96	2.34	4.70	2.25	5.08	1.64	2.75
HH	[2.52]		[1.51]		[2.14]		[2.24]		[2.50]		[1.81]	
VAR	[2.37]		[1.51]		[2.02]		[2.12]		[2.39]		[1.66]	
3	2.46	6.52	1.46	2.40	2.04	4.94	2.29	5.84	2.21	6.49	1.60	3.53
HH	[2.46]		[1.46]		[1.97]		[2.05]		[2.40]		[1.73]	
VAR	[2.19]		[1.47]		[1.94]		[2.16]		[2.39]		[1.63]	
6	2.45	10.23	1.45	3.84	2.02	6.96	2.29	8.25	2.19	9.95	1.61	5.63
HH	[2.50]		[1.48]		[1.80]		[1.87]		[2.38]		[1.73]	
VAR	[2.49]		[1.49]		[2.03]		[2.12]		[2.48]		[1.93]	
12	2.12	13.14	1.12	4.05	2.27	12.94	2.54	14.86	1.90	12.37	1.32	6.45
HH	[2.18]		[1.15]		[1.91]		[1.93]		[2.10]		[1.42]	
VAR	[2.14]		[1.20]		[2.79]		[3.19]		[2.21]		[1.51]	

The regressions in Eqs. (1) and (2) test different hypotheses. In the regression for excess returns in Eq. (1), the null states that the log expected excess currency returns are constant. In the regression for log exchange rates changes in Eq. (2), the null states that changes in the log spot rates are unpredictable, i.e., the expected excess returns are time-varying and they are equal to the interest rate differentials (i.e., forward discounts).

3.3.2. *Predictability results*

Table 2 reports the test statistics for these regressions. The left panel focuses on developed countries. There is strong evidence against UIP in the returns on the developed countries basket, at all horizons.[7] The estimated slope coefficients, $\psi_{\mathbf{f}}$, in the predictability regressions are highly statistically significant, regardless of the method used to compute the t-statistics, except for annual horizon non-overlapping returns; we have too few observations given the length of our sample. The R^2 increases from about 3% at the monthly horizon to up to 13% at the one-year horizon. This increase in the R^2 as we increase the holding period is not surprising, given the persistence of the AFD.

Moreover, given that the coefficient is substantially greater than unity, average exchange rate changes are also predictable, although statistically, we cannot reject the hypothesis that they follow a random walk. The R^2's for the exchange rate regressions are lower, ranging from just over 1% for monthly to 4% for annual horizons.

As noted in the Introduction, the impact of the AFD on expected excess returns is large. At the one-month horizon, each 100 basis point increase in the forward discount implies a 245 basis points increase in the expected return, and it increases the expected appreciation of the foreign currency basket by 145 basis points. The estimates are very similar for all maturities, except the 12-month estimate, which is 33 basis points lower. The constant in this predictability regression is 0.00 (not reported) at all maturities. This is why our naive investment rule used for implementing the dollar carry trade is actually optimal.

The central panel in Table 2 reports the results for the emerging markets basket. We use the AFD of the developed country basket to forecast the emerging markets basket returns. There is no overlap between the countries used to construct the AFD and the currencies in the portfolio of emerging market countries. There is equally strong predictability for average log excess returns and average spot rate changes for the emerging markets basket because the AFD of developed countries is not very highly correlated with the AFD of emerging countries. In fact, for the emerging countries basket, excess returns are not at all predictable using their own AFD, and the UIP condition cannot be rejected (these results are not reported here for

[7] Hansen and Hodrick (1980) and Fama (1984) conclude that predicted excess returns move more than one-for-one with interest rate differentials, implying some predictability in exchange rates, even though the statistical evidence from currency pairs is typically weak.

Table 3
Predictability using bilateral and average forward discounts: Panel regressions. This table reports results of panel regressions for average excess returns and average exchange rate changes for individual currencies at horizons of one, two, three, six, and 12 months, on both the average forward discount for developed countries and the currency-specific forward discount, as well as currency fixed effects (to allow for different drifts). For each group of countries (developed, emerging, and all), we report the slope coefficients on the average log forward discount for developed countries ($\psi_{\mathbf{f}}$) and on the individual forward discount (φ_f), and similarly the slope coefficient $\zeta_{\mathbf{f}}$ and ξ_f for the exchange rate changes. The *t*-statistics for the slope coefficients in brackets are computed using robust standard errors clustered by month and country. Data are monthly, from Barclays and Reuters (available via Datastream). The returns do not take into account bid–ask spreads. The sample period is 11/1983–6/2010.

	Developed countries				Emerging countries				All countries			
	Excess returns		Exchange rates		Excess returns		Exchange rates		Excess returns		Exchange rates	
Horizon	$\tilde{\psi}_{\mathbf{f}}$	φ_f	$\tilde{\zeta}_{\mathbf{f}}$	ξ_f	$\tilde{\psi}_{\mathbf{f}}$	φ_f	$\tilde{\zeta}_{\mathbf{f}}$	ξ_f	$\tilde{\psi}_{\mathbf{f}}$	φ_f	$\tilde{\zeta}_{\mathbf{f}}$	ξ_f
1	1.87	0.60	1.87	−0.40	1.59	1.12	1.59	0.12	1.56	0.95	1.56	−0.05
	[2.13]	[1.52]	[2.13]	[−0.99]	[1.91]	[2.30]	[1.91]	[0.25]	[2.02]	[2.47]	[2.02]	[−0.12]
2	2.10	0.51	2.10	−0.49	1.35	1.19	1.35	0.19	1.52	1.01	1.52	0.01
	[2.74]	[1.15]	[2.74]	[−1.12]	[1.81]	[2.15]	[1.81]	[0.34]	[2.22]	[2.26]	[2.22]	[0.01]
3	2.15	0.39	2.15	−0.61	1.15	1.30	1.15	0.30	1.38	1.07	1.38	0.07
	[3.02]	[0.88]	[3.02]	[−1.36]	[1.66]	[2.47]	[1.66]	[0.57]	[2.18]	[2.44]	[2.18]	[0.16]
6	2.23	0.33	2.23	−0.67	1.02	1.31	1.02	0.31	1.34	1.09	1.34	0.09
	[3.45]	[0.77]	[3.45]	[−1.53]	[1.53]	[2.77]	[1.53]	[0.66]	[2.54]	[2.75]	[2.54]	[0.23]
12	1.89	0.32	1.89	−0.68	1.00	1.56	1.00	0.56	1.10	1.22	1.10	0.22
	[4.12]	[0.99]	[4.12]	[−2.13]	[1.52]	[3.00]	[1.52]	[1.08]	[2.38]	[2.83]	[2.38]	[0.52]

brevity). This is consistent with the view that, among emerging market currencies, forward discounts mostly reflect inflation expectations rather than risk premia. It is also consistent with the findings of Bansal and Dahlquist (2000), who argue that the UIP has more predictive power for exchange rates of high-inflation countries and in particular, emerging markets (Frankel and Poonawala, 2010 report similar results). Nevertheless, as our predictability results indicate, risk premia are important for understanding exchange rate fluctuations even for high inflation currencies.

The right panel in Table 2 pertains to the sample of all countries. Not surprisingly, excess return predictability is very strong there as well.

3.4. *The average forward discount and bilateral exchange rates*

We now compare our predictability results to standard tests in the literature, which mostly focus on bilateral exchange rates. By capturing the dollar risk premium, the average forward discount is able to forecast individual currency returns, as well as their basket-level averages. In fact, it is often a better predictor than the individual forward discount specific to the given currency pair. One way to see this is via a pooled panel regression of excess returns on the AFD and the currency-specific forward discount:

$$rx^i_{t\to t+k} = \psi^i_0 + \tilde{\psi}_{\mathbf{f}}(\overline{f}^{Dev}_{t\to t+k} - \overline{s}^{Dev}_t) + \varphi_f(f^i_t - s^i_t) + \eta^i_{t+k}, \quad (3)$$

and a similar regression for spot exchange rate changes:

$$-\Delta s^i_{t\to t+k} = \zeta^i_0 + \tilde{\zeta}_{\mathbf{f}}(\overline{f}^{Dev}_{t\to t+k} - \overline{s}^{Dev}_t) + \xi_f(f^i_t - s^i_t) + \tilde{\eta}^i_{t+k}, \quad (4)$$

where ψ^i_0 and ζ^i_0 are currency fixed effects, so that only the slope coefficients are constrained to be equal across currencies.

Table 3 presents the results for the developed and emerging countries' subsamples, as well as the full sample of all currencies. The coefficients on the AFD are large, around 2 for developed countries for both excess returns and exchange rate changes (as we are controlling for individual forward discounts). The coefficients are highly statistically significant. In contrast, the coefficients on the individual forward discount are small for the developed markets sample, not statistically different from zero (and in fact, negative for spot rate changes). For emerging countries, individual forward discounts are as important as the AFD for predicting excess returns, but not for exchange rate changes.

A similar picture emerges from bivariate predictive regressions run separately for individual currencies. We do not report these results here, but they are briefly summarized as follows. Using the AFD of the developed countries basket to predict bilateral currency returns and exchange rate changes over six-month horizons yields average R^2's (across all developed country currencies) of 10% and 5%, compared to the average R^2's using the bilateral forward discount of 7.3% and 2.3%, respectively. Similarly, at 12-month horizons the average R^2's using the AFD are 15% and 9% for excess returns and spot rate changes, respectively, compared to the average R^2's of 11.6% and 4.1%, respectively, using currency-specific forward discounts. While the differences are smaller at shorter horizons and for emerging market currencies, the results are broadly consistent with the average forward discount containing information about future exchange rates above and beyond that in individual currency forward discounts.

We thus find that a single return-forecasting variable describes time-variation in currency excess returns and changes in exchange rates even better than the forward discount rates on the individual currency portfolios. This variable is the AFD of developed countries. The results are consistent across different baskets and maturities. When the AFD is high (i.e., when U.S. interest rates are lower than the average world interest rate), expected currency excess returns are high. Conditioning their investments on the level of the AFD, U.S. investors earn large currency excess returns that are not correlated with the HML currency carry trades. We turn now to a no-arbitrage model that offers an interpretation of our empirical findings.

4. A no-arbitrage model of interest rates and exchange rates

We develop a no-arbitrage model that can quantitatively account for the bulk of our empirical findings and explains the link between risk prices in the U.S. and currency return predictability.

4.1. Pricing kernel volatility and currency returns

We assume that financial markets are complete, but that some frictions in the goods markets prevent perfect risk-sharing across countries. As a result, the change in the real exchange rate Δq^i between the home country and country i is $\Delta q^i_{t+1} = m_{t+1} - m^i_{t+1}$, where m denotes the log SDF (also known as a pricing kernel or intertemporal marginal rate of substitution, IMRS) and q^i is measured in country i goods per home country good. An increase in q^i means a real appreciation of the home currency. For any variable that pertains to the home country (the U.S.), we drop the superscript.

The real expected log currency excess return equals the interest rate difference or forward discount plus the expected rate of appreciation. If pricing kernels m are log-normal, the real expected log currency excess return on a long position in a basket of foreign currency i and a short position in the dollar is equal to

$$\begin{aligned} E_t\left[rx^{basket}_{t+1}\right] &= -\frac{1}{N}\sum_i E_t\left[\Delta q^i_{t+1}\right] + \frac{1}{N}\sum_i r^i_t - r_t \\ &= \frac{1}{2}\left[Var_t(m_{t+1}) - \frac{1}{N}\sum_i Var_t\left(m^i_{t+1}\right)\right]. \end{aligned} \tag{5}$$

The dollar carry trade goes long or short depending on the magnitude of the average interest rate differential $((1/N)\sum_i r^i_t - r_t)$. In order for this strategy to earn positive average returns, investors must be long in the basket when expected returns on foreign currency investments are high, i.e., when the volatility of the U.S. pricing kernel is high (relative to foreign pricing kernels), and short in the basket when the volatility of the U.S. pricing kernel is low. Therefore, these expected returns are driven by the U.S. economic conditions. By contrast, the real expected log currency excess return on the HML currency carry trade is given by

$$E_t\left[rx^{basket}_{t+1}\right] = \frac{1}{2}\left[\frac{1}{N_L}\sum_{j\in L} Var_t\left(m^j_{t+1}\right) - \frac{1}{N_H}\sum_{j\in H} Var_t\left(m^j_{t+1}\right)\right], \tag{6}$$

where H (L) denote high (low) interest rate currencies, respectively. The expected returns on the HML currency carry trade are high when the gap between the SDF volatilities of low and high interest rate currencies increases. These expected returns are driven by global economic conditions (e.g., volatility in the world financial markets).

We use a no-arbitrage asset pricing model in the tradition of the affine term structure models to interpret the evidence on the conditional expected returns earned on the U.S. dollar basket documented above. We show that the model can replicate these empirical facts while also matching other key features of currency excess returns and interest rates. The model makes new predictions about the cross-sectional properties of average returns on currency baskets formed from the perspective of different base currencies, which we verify in the data.

4.2. Setup

We extend the no-arbitrage model developed by Lustig, Roussanov, and Verdelhan (2011) to explain high-minus-low carry trade returns. In the model, there are two types of priced risk: country-specific shocks and common shocks. Brandt, Cochrane, and Santa-Clara (2006), Bakshi, Carr, and Wu (2008), Colacito (2008), and Colacito and Croce (2011) emphasize the importance of a large common component in SDFs to make sense of the high volatility of SDFs and the relatively ‘low’ volatility of exchange rates. In addition, there is a lot of evidence that much of the stock return predictability around the world is driven by variation in the global risk price, starting with the work of Harvey (1991) and Campbell and Hamao (1992). Lustig, Roussanov, and Verdelhan (2011) show that, in order to reproduce cross-sectional evidence on currency excess returns, risk prices must load differently on this common component.

In our model, carry trade returns are driven by real variables and inflation is not priced. Hollifield and Yaron (2001) have documented that nearly all of the predictable variation in currency returns is due to real, not nominal, variables. This was confirmed by Lustig, Roussanov, and Verdelhan (2011) who found that the carry trade portfolios sorted by (nominal) forward discounts produce almost equally large spreads in real interest rates.

We consider a world with N countries and currencies. We do not specify a full economy complete with preferences and technologies; instead we posit a law of motion for the SDFs directly as being driven by both global and country-specific shocks. The risk prices associated with country-specific shocks depend only on the country-specific factors, but we allow the risk prices of world shocks to depend on world and country-specific factors. To describe these risk prices, we introduce a common state variable z^w_t shared by all countries and a country-specific state variable z^i_t. The country-specific and world state variables follow autoregressive square-root processes:

$$z^i_{t+1} = (1-\phi)\theta + \phi z^i_t - \sigma\sqrt{z^i_t}u^i_{t+1}, \tag{7}$$

$$z^w_{t+1} = (1-\phi^w)\theta^w + \phi^w z^w_t - \sigma^w\sqrt{z^w_t}u^w_{t+1}. \tag{8}$$

Intuitively, z^i_t captures variation in the risk price due the business cycle of country i, while z^w_t captures global variation in risk prices. We assume that in each country i, the logarithm of the real SDF m^i follows a three-factor conditionally Gaussian process:

$$-m^i_{t+1} = \alpha + \chi z^i_t + \sqrt{\gamma z^i_t}u^i_{t+1} + \tau z^w_t + \sqrt{\delta^i z^w_t}u^w_{t+1} + \sqrt{\kappa z^i_t}u^g_{t+1}, \tag{9}$$

where u^i_{t+1} is a country-specific SDF shock while u^w_{t+1} and u^g_{t+1} are common to all countries SDFs. All of these three innovations are Gaussian, with zero mean and unit variance, independent of one another and over time. There are two types of common shocks. The first type u^w_{t+1} is priced identically in all countries that have the same exposure δ, and all differences in exposure are *permanent*. Examples of this type of innovation would be a global financial crisis or some form of global uncertainty. This dollar-neutral innovation will be the main driving force behind the HML carry trade. The second type of common shock, u^g_{t+1}, is not, as heterogeneity with respect to this innovation is *transitory*: all countries are equally exposed to this shock *on average*, but conditional exposures vary over time and depend on country-specific economic conditions. This, for example, could be a global productivity shock that affects some economies more than others depending on each country's current consumer demand or monetary policy.[8] This innovation, in conjunction with the U.S.-specific innovation, is the main driving force behind the dollar carry trade, and is obviously not dollar-neutral. We include this last type of shock to allow the exposure of each country's intertemporal marginal rate of substitution to global shocks, and therefore the price of global risk, to vary over time with that country's economic and financial conditions (z^i_t).

To be parsimonious, we limit the heterogeneity in the SDF parameters to the different loadings δ^i on the world shock u^w_{t+1}; all the other parameters are identical for all countries. The only qualitative departure of our model from the model in Lustig, Roussanov, and Verdelhan (2011) is the separation of the world shock into two independent shocks, u^w_{t+1} and u^g_{t+1}, dictated by the correlation properties of different carry trade strategies.

4.2.1. Currency risk premia for individual currencies

In our model, the forward discount on currency i relative to USD is equal to: $r^i_t - r_t = (\chi - \frac{1}{2}(\gamma+\kappa))(z^i_t - z_t) - \frac{1}{2}(\delta^i - \delta)z^w_t$. The three same variables, z_t, z^i_t, and z^w_t, determine the time-variation in the conditional expected log currency excess returns on a long position in currency i and a short position in the home currency: $E_t[rx^i_{t+1}] = \frac{1}{2}[(\gamma+\kappa)(z_t - z^i_t) + (\delta - \delta^i)z^w_t]$.

Accounting for the variation in expected currency excess returns across different currencies requires variation in the SDFs' exposures to the common innovation. Lustig, Roussanov, and Verdelhan (2011) show that *permanent* heterogeneity in loadings, captured by the dispersion in δ's, is necessary to explain the variation in *unconditional* expected returns (why high interest rate currencies tend not to depreciate on average), whereas the *transitory* heterogeneity in loadings, captured by the κz^i_t term in Eq. (9), is necessary to match the variation in *conditional* expected returns (why currencies with currently high interest rates tend to appreciate). While much of the literature on currency risk premia focuses on the latter (conditional) expected returns (e.g., see Lewis, 1995 for a survey), the former (unconditional) average returns are also important (e.g., Campbell, Medeiros, and Viceira, 2010) and account for between a third and a half of the carry trade profits, as reported in Lustig, Roussanov, and Verdelhan (2011).

4.2.2. Currency risk premia for baskets of currencies

We turn now to the model's implications for return predictability on baskets of currencies. We use a bar superscript ($\overline{x}$) to denote the average of any variable or parameter x across all the countries in the basket. The average real expected log excess return of the basket is

$$E_t[\overline{rx}_{t+1}] = \frac{1}{2}(\gamma+\kappa)(z_t - \overline{z_t}) + \frac{1}{2}(\delta - \overline{\delta})z^w_t. \tag{10}$$

We assume that the number of currencies in the basket is large enough so that country-specific shocks average out within each portfolio. In this case, $\overline{z}$ is approximately constant (and exactly in the limit $N \to \infty$ by the law of large numbers). As a result, the real expected excess return on this basket consists of a dollar risk premium (the first term above, which depends only on z_t) and a global risk premium (the second term, which depends only on z^w_t). The real expected excess return of this basket depends only on z and z^w. These are the same variables that drive the AFD: $\overline{r_t} - r_t = (\chi - \frac{1}{2}(\gamma+\kappa))(\overline{z_t} - z_t) + \frac{1}{2}(\delta - \overline{\delta})z^w_t$. Thus, the AFD contains information about average excess returns on a basket of currencies.

4.2.3. Unconditional HML carry trades

The model has strong predictions on HML currency carry trades; Lustig, Roussanov, and Verdelhan (2011) study them in detail. Here, to keep things simple, we consider investment strategies that do not entail continuous rebalancing of the portfolios, i.e., unconditional HML carry trades.

If one were to sort currencies by their average interest rates (not their current rates) into portfolios, then, as shown by Lustig, Roussanov, and Verdelhan (2011), investors who take a carry trade position would only be exposed to common innovations, not to U.S. innovations. The return innovations on this HML investment (denoted hml^{unc}, for unconditional carry trades) are driven only by u^w shocks:

$$hml^{unc}_{t+1} - E[hml^{unc}_{t+1}] = \left(\frac{1}{N_L}\sum_{i \in L}\sqrt{\delta^i z^w_t} - \frac{1}{N_H}\sum_{i \in H}\sqrt{\delta^i z^w_t}\right)u^w_{t+1}. \tag{11}$$

We thus label the u^w shocks the HML carry trade innovations. This HML portfolio yields positive average returns if the pricing kernels of low interest rate currencies are more exposed to the global innovation than the pricing kernels of high interest rate currencies.

4.3. The dollar carry risk premium in the model

We turn now to the dollar carry strategy. In order to build intuition for dollar carry risk, it helps to a consider a special case.

[8] Gourio, Siemer, and Verdelhan (2013) propose an international real business cycle model with two common shocks: in their model, shocks to the probability of a world disaster drive the HML carry trade risk. Productivity shocks are correlated across countries and thus exhibit a common component, akin to a second type of common shock.

4.3.1. Special case: average exposure to global shocks

Consider the case of a basket consisting of a large number of developed currencies, such that the average country's SDF has the same exposure to global shocks u^w as the base country (the U.S.): $\overline{\delta} = \delta$. In this special case, the dollar, measured vis-à-vis the basket of developed currencies, does not respond to the common shocks u^w that are priced in the same way in the U.S. and, on average, in all the other countries. However, the dollar does respond to U.S.-specific shocks (u in the model) and to global shocks (u^g) that are priced differently in each country. A short position in the dollar is risky because the dollar appreciates following negative U.S.-specific shocks and negative global shocks to which the U.S. is more exposed than other countries.

In this case, the log currency risk premium on the basket only depends on the U.S.-specific factor z_t, not the global factor:

$$E_t\left[\overline{rx}_{t+1}\right] = \tfrac{1}{2}(\gamma+\kappa)(z_t - \overline{z_t}). \tag{12}$$

Hence, the currency risk premium on this basket compensates U.S. investors proportionally to their exposures to the local risk governed by γ and to their exposure to global risk governed by κ. Given the assumption of average exposure, the dollar risk premium is driven exclusively by U.S. variables (e.g., the state of the U.S. business cycle). U.S. investors expect to be compensated more for bearing this risk during recessions. We refer to the risk premium as the domestic component of the dollar carry trade risk premium or dollar premium because its variation depends only on the U.S.-specific state variable.

Similarly, given the average exposure assumption, the AFD only depends on the U.S. factor z_t:

$$\overline{r_t} - r_t = \left(\chi - \tfrac{1}{2}(\gamma+\kappa)\right)(\overline{z_t} - z_t). \tag{13}$$

Reproducing the failure of the UIP requires assuming that $\chi < \frac{1}{2}(\gamma+\kappa)$. Then during "bad times," when z_t increases, U.S. interest rates are low and the average forward discount is high.[9] Since the parameters guiding country-specific state variables are the same across countries, the mean AFD should be equal to zero. Empirically, the basket of developed countries, currencies formed from the U.S. perspective has a mean AFD of less than 1% per annum, which is not statistically different from zero. As a result, the assumption that the U.S. SDF has the same sensitivity to world shocks as the average developed country appears reasonable.

By creating a basket in which the average country shares the home country's exposure to global shocks, we have eliminated the effect of foreign idiosyncratic factors on currency risk premia *and* on interest rates. For this specific basket, the slope coefficient in a predictability regression of the average log returns in the basket on the AFD is $\psi_{\mathbf{f}} = -\frac{1}{2}(\gamma+\kappa)/\left(\chi - \frac{1}{2}(\gamma+\kappa)\right)$. Correspondingly, the slope coefficient in a regression of average real exchange rate changes on the real forward discount is $\zeta_{\mathbf{f}} = -\chi/\left(\chi - \frac{1}{2}(\gamma+\kappa)\right)$. Provided that $\chi < \frac{1}{2}(\gamma+\kappa)$ (i.e., interest rates are low in bad times and high in good times), a positive average forward discount forecasts positive future excess returns.[10]

Under the assumption that $\overline{\delta} = \delta$, the dollar carry trade strategy is long the basket when the average forward discount (and therefore the dollar premium) is positive, and short otherwise. Therefore, the dollar carry risk premium is given by

$$E_t\left[\textit{Dollar carry}_{t+1}\right] = \tfrac{1}{2}(\gamma+\kappa)(z_t - \overline{z_t})\,\mathrm{sign}(z_t - \overline{z_t}) > 0. \tag{14}$$

The dollar premium is driven entirely by the domestic state variable z_t. This state variable influences the market price of *local* risk (i.e., the compensation for the exposure to U.S.-specific shocks u_{t+1}), as well as the market price of *global* risk (i.e., the compensation for the exposure to global shocks u^g_{t+1}).

In contrast to the dollar strategy, the expected excess returns on the unconditional HML carry trade portfolio do not depend on z_t, the U.S.-specific factor, but, given our assumptions, they depend only on the global state variable z^w_t. Hence, we do not expect the AFD to predict returns on HML carry (or other currency trading strategies that are dollar-neutral) as long as $\overline{\delta} \approx \delta$. This prediction is confirmed in the data: there is no evidence of predictability of HML carry trade returns using the AFD.

4.3.2. General case

In general, the innovations to the dollar carry returns are driven by all three shocks that drive the stochastic discount factor:

$$\textit{Dollar carry}_{t+1} - E_t\left[\textit{Dollar carry}_{t+1}\right] = \begin{bmatrix} \sqrt{\gamma z_t}\, u_{t+1} \\ +\left(\sqrt{\kappa z_t} - \frac{1}{N}\sum_i \sqrt{\kappa z^i_t}\right) u^g_{t+1} \\ +\left(\sqrt{\delta z^w_t} - \dfrac{1}{N}\sum_i \sqrt{\delta^i z^w_t}\right) u^w_{t+1} \end{bmatrix} \mathrm{sign}(\overline{r_t} - r_t). \tag{15}$$

The first two parts constitute the domestic component of the dollar carry premium, driven by the domestic state variable z_t. The last part is the global component. If the U.S. has average exposure to global shocks, then the last component, driven by the global state variable, is approximately zero. However, if the U.S. investor's SDF is more exposed to the global risk than the average ($\delta > \overline{\delta}$), the AFD will tend to be higher on average, and the long position in the basket of foreign currencies is profitable more often; the converse is true when ($\delta < \overline{\delta}$).

Froot and Ramadorai (2005) show that much of the exchange rate variation is driven by transitory shocks to expected excess returns. Accordingly, in our model, the U.S.-specific shock u_t and the world shock u^w_t that drive

[9] Requiring that the state variables are always positive at least in the continuous-time approximation is not sufficient to ensure that the real interest rates are positive. Nominal interest rates are almost always positive in our simulations. See also the related discussion in Backus, Foresi, and Telmer (2001).

[10] If $\chi = 0$, the Meese and Rogoff (1983) empirical result holds in population: the log of real exchange rates follows a random walk, and the expected log excess return is equal to the real interest rate difference. On the other hand, when $\gamma+\kappa = 0$, UIP holds exactly, i.e., $\zeta_{\mathbf{f}} = -1$.

the innovations to the U.S. dollar exchange rate also drive the conditional prices of risk.

4.3.3. *Interpreting shocks*

Is there a potential economic interpretation of these shocks? Lustig, Roussanov, and Verdelhan (2011) show that the HML strategy returns are highly correlated with the changes in the volatility of the world equity market portfolio return, making it a good candidate for the u^w shock and giving the global state variable z^w the interpretation of global financial market volatility. Menkhoff, Sarno, Schmeling, and Schrimpf (2012) show that carry trade returns co-move with global exchange rate volatility. Interestingly, these variables are uncorrelated with the dollar carry returns, as is the HML portfolio itself. The dollar carry portfolio is, however, correlated with the average growth rate of aggregate consumption across countries in the Organization for Economic Co-operation and Development (OECD). The correlation is 0.19 and is statistically significant (it is also for the HML portfolio). This suggests world consumption growth as a good candidate for the u^g shock. Further, the dollar carry is correlated with the U.S.-specific component of the growth rate of U.S. industrial production (obtained as a residual from regressing the U.S. industrial production growth rate on the world average), suggesting that the innovations to the home-country state variable z capture the domestic component of the cyclical variation in the volatility of the stochastic discount factor (and therefore the price of global risk). We pursue this interpretation further in Section 6.

4.3.4. *Predictability and heterogeneity*

When the home country exposure to the global shock u^w differs from that of the average foreign country ($\delta \neq \overline{\delta}$), then the currency risk premium loads on the global factor, and so does the forward discount for that currency. Therefore, in the general case the average forward discount would have less forecasting power for excess returns and exchange rate changes because the local and global state variables may affect them differently. In the special case of average loading, $\delta = \overline{\delta}$, the presence of heterogeneity in these loadings still matters. This type of heterogeneity will invariably lower the UIP slope coefficient in a regression of exchange rate changes on the forward discount in absolute value relative to the case of a basket of currencies. The UIP slope coefficient for individual currencies using the forward discount for that currency is given by

$$\zeta_f^i = -\frac{\chi(\chi - \frac{1}{2}(\gamma+\kappa))var(z_t^i - z_t)}{(\chi - \frac{1}{2}(\gamma+\kappa))^2 var(z_t^i - z_t) + \frac{1}{4}(\delta^i - \delta)^2 var(z_t^w)}. \tag{16}$$

The UIP regression coefficient of the average exchange rate changes in the basket on the average forward discount has the same expression as the UIP coefficient for two ex ante identical countries:

$$\zeta_{\mathbf{f}} = -\tfrac{1}{2}(\gamma+\kappa)/(\chi - \tfrac{1}{2}(\gamma+\kappa)). \tag{17}$$

It follows that the basket-level slope coefficient on AFD is the largest of all individual FX slope coefficients: $\zeta_{\mathbf{f}} \geq \zeta_f^i$.

Intuitively, at the level of country-specific investments, the volatility of the forward discount is greater, relative to the case of a basket of currencies, but the covariance between interest rate differences and exchanges rate changes is not. Hence, heterogeneity in exposure to the global innovations pushes the UIP slope coefficients towards zero, relative to the benchmark case with identical exposure. Therefore, we expect to see larger slope coefficients for UIP regressions on baskets of currencies, not simply due to the diversification effect of reducing idiosyncratic noise, but because these baskets eliminate the effect of heterogeneity in exposure to global innovations that attenuates predictability. This prediction of the model is consistent with the data (subject to the sampling error). In our entire sample of developed and emerging country currencies, only two exchange rate series, the U.K. pound sterling and the Euro (moreover, on a shorter sample), exhibit slope coefficients in the UIP regressions that are slightly greater (but not statistically different) from the coefficient of 1.45 estimated for the developed country basket.

4.3.5. *Correlation between carry strategies*

To study the correlation of returns on different carry strategies, we proceed again in two steps, starting with the case of a country with average exposure to the common shock. In this case, the innovations to the dollar carry trade returns are independent of the u^w shocks, but are driven exclusively by U.S.-specific u shocks and u^g shocks:

$$\left[\sqrt{\gamma z_t}u_{t+1} + \sqrt{\kappa}\left(\sqrt{z_t} - \frac{1}{N}\sum_i \sqrt{z_t^i}\right)u_{t+1}^g\right] \text{sign}(z_t - \overline{z_t}).$$

To derive this result, we assume that the dispersion in δ^i is sufficiently small so that $(1/N)\sum_i \sqrt{\delta^i} \approx \sqrt{\overline{\delta}} = \sqrt{\delta}$. In this case, the u^w shocks simply do not affect the dollar exchange rate. Hence, if the U.S. is a country with an average δ, then the dollar carry will only depend on the second shock u_{t+1}^g and its correlation with the unconditional carry trade returns hml_{t+1}^{unc} will be zero, because the innovations there only depend on u_{t+1}^w.

If the home country's exposure δ is either well above or below the average, then the dollar carry trade returns have a positive correlation with the unconditional HML carry trade, and hence a higher correlation with the conditional HML carry trade as well. In general, the correlation between the HML currency carry trade (sorting currencies by current interest rates) and the dollar carry depends on the relative contributions of the common SDF shock u_{t+1}^g to their returns.[11] In Section 5.7, we trace out this

[11] The correlation between the HML carry trade and the dollar carry depends on the relative contributions of the common SDF shock u_{t+1}^g to their returns, as the conditional correlation between the two strategies is given by

$$Corr_t(Dollar\ carry_{t+1}, hml_{t+1}) = \frac{\kappa\left|\sqrt{z_t} - \frac{1}{N}\sum_i \sqrt{z_t^i}\right|\left(\frac{1}{N_L}\sum_{i\in L}\sqrt{z_t^i} - \frac{1}{N_H}\sum_{i\in H}\sqrt{z_t^i}\right)}{\sqrt{Var_t(Dollar\ Carry_{t+1})}\sqrt{Var_t(hml_{t+1})}}, \tag{18}$$

where $Var_t(Dollar\ carry_{t+1}) = \gamma z_t + |\sqrt{\kappa z_t} - (1/N)\sum_i \sqrt{\kappa z_t^i}|^2$, and $Var_t(hml_{t+1}) = ((1/N_L)\sum_{i\in L}\sqrt{\delta^i z_t^w} - (1/N_H)\sum_{i\in H}\sqrt{\delta^i z_t^w})^2 + ((1/N_L)\sum_{i\in L}\sqrt{\kappa z_t^i} - (1/N_H)\sum_{i\in H}\sqrt{\kappa z_t^i})^2$.

U-shaped relation between the correlation and the average forward discount, determined in the model by δ^i. But to do that, we first need to estimate the model parameters.

5. Model estimation

In this section, we estimate the model on a small sample of moments and compared its predictions to the data. To summarize, the model is defined by the following set of equations:

$$-m_{t+1}^i = \alpha + \chi z_t^i + \sqrt{\gamma z_t^i} u_{t+1}^i + \tau z_t^w + \sqrt{\delta^i z_t^w} u_{t+1}^w + \sqrt{\kappa z_t^i} u_{t+1}^g, \tag{19}$$

$$z_{t+1}^i = (1-\phi)\theta + \phi z_t^i - \sigma \sqrt{z_t^i} u_{t+1}^i, \tag{20}$$

$$z_{t+1}^w = (1-\phi^w)\theta^w + \phi^w z_t^w - \sigma^w \sqrt{z_t^w} u_{t+1}^w, \tag{21}$$

$$\pi_{t+1}^i = \pi_0 + \eta^w z_t^w + \sigma_\pi \varepsilon_{t+1}^i. \tag{22}$$

There is a single source of heterogeneity in the countries′ exposure to the global shocks. On the real side, the model thus has 11 parameters that are identical for all countries (α, χ, γ, τ, κ, ϕ, θ, σ, ϕ^w, θ^w, and σ^w). Additionally, countries differ according to their exposure (δ^i) to world shocks that are priced globally. The parameters δ^i are assumed to be distributed on the interval $[\delta_h, \delta_l]$. We are not targeting individual countries, but assume that the domestic δ is equal to the average across δ^i's, thus adding three parameters to the model (we do not need to make any other assumptions on the distribution of δ^i's for the purpose of the estimation). On the nominal side, the model has three parameters (π_0, η^w, and σ_π). The model thus has a total of 17 parameters. We use a version of the Generalized Method of Moments (GMM) estimator that targets 18 moments. We first define the moments and then describe how we obtain their empirical values.

5.1. Targets

Table 4 reports all the moments used in the estimation, along with their weights in the GMM objective function. When available, the table also reports closed-form expressions. Panel A of Table 4 lists the 12 moments used in the optimization procedure, while Panel B focuses on additional moments used to identify the remaining parameters (these moments are matched exactly). Finally, for completeness, Panel C of Table 4 reports additional moments not used in the estimation: the implied first and second moments of nominal interest rates, as well as their cross-country correlations, and the volatility of nominal changes in exchange rates.

The GMM objective function targets 12 moments to estimate the following 11 country-invariant parameters (and the home-country's δ which is assumed equal to the average of the δ^i's):

$$\{\chi, \tau, \gamma, \kappa, \phi, \theta, \sigma, \phi^w, \theta^w, \sigma^w, \delta\}.$$

These target moments are the slope coefficients for the regression of average exchange rate changes on the average forward discount ($\zeta_\mathbf{f}$), the R^2 of this regression (which is closely related to the Sharpe ratio on the dollar carry strategy) ($R^2_{\zeta_\mathbf{f}}$), the standard deviation ($Std(\overline{r_t^i} - r_t)$) and the autocorrelation of the average forward discount ($Corr(\overline{r_t^i} - r_t, \overline{r_{t+1}^i} - r_{t+1})$), the standard deviation ($Std(r)$), and autocorrelation ($Corr(r_t, r_{t+1})$) of the U.S. real short-term interest rates, the cross-country correlation of real interest rates ($Corr(r_t, r_t^i)$), the standard deviation of changes in real exchanges rates ($Std(\Delta q)$), the average excess return on the developed country currency basket ($E(rx_t^{dollar})$), the average slope of the nominal term structure of interest rates in the U.S., calculated as the average nominal ten-year yield minus the average nominal risk-free rate ($E(y^{\$,10} - r^\$)$), as well as two Feller conditions that ensure the non-negativity of the volatility processes.[12] A Feller coefficient of (at least) 20 guarantees that all of the state variables following square-root processes are positive (this is exact in the continuous-time approximation, and implies a negligible probability of crossing the zero bound in discrete time, a possibility never realized in our simulations given these parameter values, even with samples as long as 1,000,000 periods). We set both Feller coefficients to 30 (otherwise θ and θ^w are not identified).

The GMM objective function minimizes a weighted sum of the squared differences between closed-form expressions for these moments and their empirical counterparts. The minimization procedure also imposes a set of economically motivated constraints: $\chi - 0.5(\gamma+\kappa) < 0$ (procyclical real risk-free rates), $\sigma > 0$, $\theta > 0$, $\phi < 1$, as well as $\sigma^w > 0$, $\theta^w > 0$, $\phi^w < 1$ (stationary volatilities), $\gamma > 0$, and $\kappa > 0$ (positive standard deviations), and $\delta < 2 \times AFD_{max}/\theta_w$ (positive exposures to the world shock for all countries). These constraints are not binding at the optimum but they help to ensure convergence of the minimization algorithm.

Panel B of Table 4 lists the six additional moments that pin down

$$\{\alpha, \eta^W, \pi_0, \sigma^\pi, \delta_h, \delta_l\}$$

directly. The first four are country-invariant. The last two determine the domain for δ^i, the only source of heterogeneity. First, the mean U.S. real risk-free rate ($E(r)$) is uniquely determined by the parameter α, once the other real parameters (χ, γ, τ, κ, θ, and θ^w) are obtained. Second, the three parameters that determine inflation rates are determined by three simple moments. The mean inflation rate $E(\pi)$ determines π_0, once η^w and θ^w are known. The

[12] The nominal log zero-coupon yield of maturity n months in the currency of country i is given by the standard affine expression

$$y_t^{\$,i,n} = -\frac{1}{n}\left(\tilde{A}_n^i + \tilde{B}_n^i z_t^i + \tilde{C}_n^i z_t^w\right),$$

where the coefficients satisfy the recursion:

$$\tilde{A}_n^i = -\alpha - \pi_0 + \tilde{A}_{n-1}^i + \tilde{B}_{n-1}^i (1-\phi)\theta + \tilde{C}_{n-1}^i (1-\phi^w)\theta^w + \frac{1}{2}\sigma_\pi^2, \tag{23}$$

$$\tilde{B}_n^i = -\left(\chi - \frac{1}{2}(\gamma+\kappa)\right) + \tilde{B}_{n-1}^i (\phi + \sigma\sqrt{\gamma}) + \frac{1}{2}(\tilde{B}_{n-1}^i \sigma)^2, \tag{24}$$

$$\tilde{C}_n^i = -\left(\eta^w + \tau - \frac{1}{2}\delta^i\right) + \tilde{C}_{n-1}^i \left(\phi^w + \sigma^w \sqrt{\delta^i}\right) + \frac{1}{2}(\tilde{C}_{n-1}^i \sigma^w)^2. \tag{25}$$

Table 4

Estimation results. Sample moments of the data and implied population moments in the model corresponding to the minimized value of the GMM objective function. Panel A lists moments targeted in estimation, Panel B presents moments matched exactly, and Panel C presents additional moments not targeted in estimation.

Moment	Closed form expression	Data	Weight	Model
Panel A: GMM target moments for parameters $\{\chi,\tau,\gamma,\kappa,\phi,\theta,\sigma,\phi^w,\theta^w,\sigma^w,\delta\}$				
$\zeta_{\mathbf{f}}$	$-\frac{\chi}{\left(\chi-\frac{1}{2}(\gamma+\kappa)\right)}$	1.71 [0.87]	10^4	1.71 [0.87]
$R^2_{\zeta_{\mathbf{f}}}$	$\frac{2\chi^2\, var(z)}{var(\Delta q)}$	1.89% [1.97%]	10^4	0.32% [1.37%]
$Std(r)$	$\sqrt{(\chi-\frac{1}{2}(\gamma+\kappa))^2 var(z)+(\tau-\frac{1}{2}\delta^i)^2 var(z^w)}$	0.53% [0.03%]	10	0.30% [0.30%]
$Std(\overline{r_t^i}-r_t)$	$\left(\chi-\frac{1}{2}(\gamma+\kappa)\right)Std(z)$	0.61% [0.05%]	10^2	0.26% [0.26%]
$Corr(\overline{r_t^i}-r_t,\overline{r_{t+1}^i}-r_{t+1})$	ϕ	0.89 [0.04]	10^2	0.91 [0.05]
$Corr(r_t,r_{t+1})$	$\frac{\phi(\chi-\frac{1}{2}(\gamma+\kappa))^2 var(z)+(\tau-\frac{1}{2}\delta)^2\phi^w var\left(z_t^w\right)}{var(r)}$	0.96 [0.04]	10^2	0.95 [0.03]
$Corr(r_t,r_t^i)$	$\left(\tau-\frac{1}{2}\delta^i\right)^2\frac{var(z^w)}{var(r)}$	0.28 [0.05]	10^4	0.27 [0.05]
$E(y^{\$,10}-r^{\$})$		2.15% [0.12%]	10	0.85% [0.90%]
$Std(\Delta q)$	$\sqrt{2\gamma\theta+2\chi^2\, var(z)+o}$	11.01% [0.46%]	2×10^6	11.01% [0.46%]
$E(rx_t^{dollar})$	$\gamma\theta$	0.82% [1.88%]	1	0.38% [0.12%]
Feller	$2(1-\phi)\frac{\theta}{var(z)}$	30.00	1	30.00
Feller	$2(1-\phi^w)\frac{\theta^w}{var(z^w)}$	30.00	1	30.00
Panel B: Exactly matched moments for parameters $\{\alpha,\eta^w,\pi_0,\sigma^\pi,\delta_h,\delta_l\}$				
$E(r)$	$\alpha+\theta\left(\chi-\frac{1}{2}(\gamma+\kappa)\right)+\theta^w\left(\tau-\frac{1}{2}\delta^i\right)$	1.31% [0.18%]		1.31% [0.18%]
$E(r_t^l-r_t)$	$-\frac{1}{2}\left(\delta^l-\delta\right)\theta^w$	3.10% [0.28%]		3.10% [0.28%]
$E(r_t^h-r_t)$	$-\frac{1}{2}\left(\delta^h-\delta\right)\theta^w$	−2.04% [0.24%]		−2.04% [0.24%]
R^2_π	$\frac{(\eta^w)^2 var(z^w)}{var(inflation)}$	6.50% [1.56%]		6.50% [1.56%]
$E(\pi)$	$\pi_0+\eta^w\theta$	2.71% [0.19%]		2.71% [0.19%]
$Std(\pi)$	$\sqrt{(\eta^w)^2 var(z^w)+\sigma_\pi^2}$	1.32% [0.04%]		1.32% [0.04%]
Panel C: Additional implied moments				
$E(r^{\$})$	$\alpha+\theta\left(\chi-\frac{1}{2}(\gamma+\kappa)\right)+\theta^w\left(\tau+\eta^w-\frac{1}{2}\delta^i\right)-\frac{1}{2}\sigma_\pi^2$	4.24% [0.22%]		4.02% [0.30%]
$Std(r^{\$})$	$\sqrt{(\chi-\frac{1}{2}(\gamma+\kappa))^2 var(z^i)+(\tau+\eta^w-\frac{1}{2}\delta^i)^2 var(z^w)}$	0.64% [0.04%]		0.32% [0.25%]
$Corr(r_t^{\$},r_t^{\$,i})$	$\left(\tau+\eta^w-\frac{1}{2}\delta^i\right)^2\frac{Var(z^w)}{var(r)}$	0.60 [0.04]		0.33 [0.29]
$Std(\Delta s)$	$\sqrt{2\gamma\theta+2\chi^2 var(z^i)+2\sigma_\pi^2+o}$	10.91% [0.47%]		11.16% [0.46%]

common component of all inflation rates is governed by the parameter η^w; its value is derived from the average R^2 of regressions of country-specific inflation rates on world inflation and the average volatility of the country-specific inflation rates $(\eta^w = \overline{R^2} \times \overline{\sigma_\pi^2}/\sigma_{z^w}^2)$. The volatility of the country-specific inflation shocks (σ_π) is derived from the average volatility across countries, corrected for its common component. Third, the heterogeneity in the model is pinned down by two moments that capture the cross-country differences in interest rates. The home country is assumed to have the average $\delta = \overline{\delta^i}$. As a result, the time-series mean of the AFD is zero. The mean forward discount on currency i is given by $\mathrm{E}(r_t^i - r_t) = -\frac{1}{2}(\delta^i - \delta)\theta^w$,so the range of unconditional mean AFDs is used to estimate the range of δs: $\delta^i = \delta - 2\mathrm{E}(r_t^i - r_t)/\theta^w$.

5.2. Empirical implementation

The specific values we target are listed in the third column of Table 4. The table reports the annualized versions of means and standard deviations (the former scaled by 12 and the latter scaled by $\sqrt{12}$), as well as monthly autocorrelations. To compute these target moments, we use one-month forward and spot exchange rates, expressed in U.S. dollars and obtained from Barclays and Reuters (Datastream). We focus on the subset of developed countries for which long time-series are available. The sample runs from 11/1983 to 06/2010 and contains the following countries: Australia, Canada, Denmark, Germany (extended using the euro series), Japan, New Zealand, Norway, Sweden, Switzerland, and United Kingdom. We use mean interest rate differentials for the Swiss franc and the Australian dollar vis-à-vis the U.S. to determine the boundaries δ_h and δ_l, respectively.[13]

Nominal interest rate differences (with respect to the U.S. short-term interest rates) are obtained as log differences between one-month forward and spot rates (assuming that covered interest rate parity holds). Adding the one-month Fama risk-free rate available from the Center for Research in Security Prices (CRSP) leads to nominal short-term interest rates in each country. Real interest rates are defined as nominal interest rates minus expected inflation, which is approximated by the previous 12-month log change in the consumer price index (CPI). Changes in exchange rates correspond to one-month log changes in spot rates. Real exchange rates are obtained by multiplying spot rates by the ratio of home to foreign CPI indices. Basket-level series are obtained as simple averages of all series in the sample. For example, the real (nominal) average forward discount corresponds to the average real (nominal) interest rate difference across countries in the sample. Realized inflation rates correspond to the one-month log differences in CPIs.

The table also lists the weights used in the estimation. The weights are chosen on economic, not statistical grounds. For example, the exchange rate predictability produced by the AFD (characterized by the slope coefficient and R^2) is obviously less precisely estimated than the volatilities of the U.S. risk-free rate and AFD. A pure statistical objective would lead to large and similar weights for the exchange rate, U.S. interest rates, and AFD volatilities, as well as the autocorrelation of U.S. interest rates and AFD, along with lower weights on the other moments. Yet, we choose to weigh the predictability moments more than the second moments and autocorrelations because they are the main focus of our paper. We also choose to overweight the volatility of real exchange rates because matching this moment has been a challenge in international finance. By doing so, we set a high hurdle for the model in terms of matching the dollar carry risk premium. In fact, the model cannot match the one observed in the data exactly without imputing too much volatility to exchange rates. This suggests that exchange rates are puzzlingly smooth when confronted with the size of risk premia in currency markets. Brandt, Cochrane, and Santa-Clara (2006) reached the same conclusion based on equity risk premia.

We provide standard errors for both the sample moments and for the moments implied by the estimated parameter values. These standard errors are computed using a block-bootstrap procedure, whereby both the sample moments and the parameter estimates are obtained for each replication. We draw 1,000 samples of all the exchange rate and interest rate series used in the estimation from the original data. Each sample is constructed by drawing (with replacement) blocks of three observations in order to capture autocorrelation and heteroskedasticity present in some of the series. This block-bootstrap procedure delivers 1,000 values of each target moment. The estimation code is run on each set of targets. This procedure thus takes into account all the uncertainty in the data, especially the small sample size.

5.3. Model fit

We now review the target values and compare them to their model counterparts in Table 4. Overall, the model delivers a reasonable fit.

Panel A reports the moments used in the GMM estimation. The model matches the slope coefficient of the basket-level change in nominal exchange rates on the average nominal interest rate difference essentially exactly at 1.71 with the same standard error as that for the sample moment (0.87). The R^2 in this regression is more difficult to match, with a model-implied value of only 0.32% compared to 1.89% in the data. There is considerable uncertainty about this moment; the sample bootstrap standard error is 1.97 while the estimation standard error is 1.37. Thus, even though we set up our estimation in order to allow the model to replicate the main predictability result of our paper, exchange rates remain close to random walks.

The model understates the volatilities of the U.S. real risk-free rate and the AFD. Empirically, the standard deviation of the U.S. real interest rate (AFD) is 0.53% (0.61%). In the model, the standard deviation of the U.S. real interest rate (AFD) is only 0.30% (0.26%). The model,

[13] While these currencies do not produce the maximal spread in mean forward discounts over our sample (the Japanese yen and the New Zealand dollar do), given the statistical uncertainty about these estimates we opt for the more conservative targets.

however, matches the three other moments of real interest rates. The persistence of the U.S. real risk-free rate and the AFD are close to their empirical values: 0.96 vs. 0.95 for the risk-free rate and 0.89 vs. 0.91 for the AFD. The average pairwise correlation of real interest rates (with the U.S. real interest rate) is 0.28 in the data vs. 0.27 in the model. Finally, the model underestimates the slope of the nominal yield curve: the ten-year minus one-month yield spread is 0.85% in the model vs. 2.15% in the data, although the difference is not statistically significant (while the raw data moment is estimated quite precisely, with the standard error of only 12 basis points, the model-implied value is not, as it has a standard error of 90 basis points). That the model produces a lower average yield curve slope can be potentially due to the fact that unexpected inflation risk is not priced in the model (only expected inflation is priced through its dependence on z^w).

The model also matches the volatilities of real exchange rates. These values are simple averages of all dollar-based exchange rates' standard deviations in the sample. The average annual standard deviation of real exchange rate changes is 11% in the data and the model matches this value essentially exactly, given the large weight placed on this moment. The sample moment is estimated very precisely with a sample bootstrap standard error of 0.46%. The model implies a small unconditional dollar risk premium of 0.38% per annum, to be compared to a risk premium of 0.82% in the data. While the latter has substantial uncertainty around it with the raw standard error of 1.88% (and therefore not statistically different from zero), the GMM estimate is fairly precise with the standard error of only 0.12%, and thus significantly lower than observed in the data. This may be interpreted as evidence that our assumption that the U.S. has δ equal to that of the average country is only approximately valid.

Panel B reports the moments that are matched exactly. Note that the standard errors on sample moments and model-generated moments coincide in this case. The model is estimated to match the mean U.S. one-month real interest rate of 1.31%, the mean inflation rate of 2.71%, and its (cross-country) average standard deviation of 1.32%. The common inflation component across countries accounts for 6.50% of inflation dynamics in the data and in the model; these numbers correspond to the average of country-specific statistics. The mean interest rate differentials against the highest-δ and lowest-δ countries are set to match the mean interest rate differentials for the Swiss franc (-2.04%) and the Australian dollar (3.10%), respectively.

Panel C reports some additional moments that were not used in the estimation. The model slightly underestimates the mean U.S. one-month nominal interest rate (4.24% in the data vs. 4.02% in the model) because the estimation targets the average inflation rate across countries instead of the U.S. inflation rate. While the model matches closely the average correlation of real risk-free rates across countries, it potentially understates the comovement of *nominal* interest rates: the average correlation vis-á-vis the U.S. in the data is 0.6 (and very precisely estimated), whereas the average estimated moment is only 0.33, albeit with a large standard error of 0.29. Similarly, the average correlation of foreign nominal yield curve slopes with the world average, $Corr(\overline{y}^{\$,10}-\overline{r}_t^{\$}, y^{i\$,10}-r_t^i)$, is 0.62 in the data but only 0.23 if simulated out of the model using the benchmark parameter estimates (not reported in the table; empirical correlation is computed using swap rate data available from Bloomberg and Datastream). This is in line with the evidence in Sarno, Schneider, and Wagner (2012), who estimate an affine term-structure model targeting both the yield curve and exchange rate moments across countries. They find that the model can either match the term structures of the interest rates or the currency risk premia well, but not both. Finally, the standard deviation of nominal exchange rate changes, reported in Panel C, is 10.91% in the data vs. 11.16% in the model (with a standard error of 0.46%).

5.4. Parameter values

The estimated parameter values are reported in Table 5, along with their standard errors. We obtain standard errors for each of the parameters using the same bootstrap procedure as described above, by taking standard deviations of corresponding parameter estimates across the bootstrap replications. We also report the key moments of the stochastic discount factor dynamics implied by the parameter estimates at the bottom of the table (in annualized units).

The parameters α, ϕ, θ, σ, ϕ^w, as well as π_0, and σ_π, appear significantly different from zero, while χ, γ, τ, κ, θ^w, σ^w, η^w, δ, δ_h, δ_l are not. The estimation thus pins down precisely the dynamics of the country-specific state variable, as well as the persistence of the world state variable, but it does not precisely measure its mean and volatility. A clear difference between the two state variables appears though: the contribution of the world state variable to the stochastic discount factor volatility is much larger on average but less volatile conditionally than its country-specific counterpart because it is more persistent (the annualized unconditional volatility of z implied by these parameter estimates is 0.5% vs. 1.32% for z^w). This difference mirrors the distinction between the dollar carry and the high-minus-low carry risk premia: the latter is close to unpredictable and very volatile as it is driven solely by the world state variable z^w, while the former derives from the strong predictability of foreign currency returns following the country-specific state variable z^i.

The average conditional SDF volatility is estimated to be 0.59. This conditional volatility itself is highly variable, however, with annualized standard deviation of 4.21% percent. The model reproduces the volatility of exchange rates while delivering large and volatile Sharpe ratios. As a consequence, the pricing kernels are also highly correlated across countries at 0.98, on average (cf. Brandt, Cochrane, and Santa-Clara, 2006).

The fact that parameters such as χ, γ, κ, and δ's are not estimated very precisely is not very surprising, given the statistical uncertainty around the estimates of slope coefficients in predictive regressions. However, while it is hard to reject that χ is equal to zero, clearly the evidence on the average returns earned by the dollar carry and the HML carry strategies indicates that κ and δ cannot be equal to zero. In order to demonstrate this, we now use the

Table 5
Parameter estimates.

This table reports the parameter values for the estimated version of the model. The model is defined by the following set of equations:

$$-m_{t+1}^i = \alpha + \chi z_t^i + \sqrt{\gamma z_t^i} u_{t+1}^i + \tau z_t^w + \sqrt{\delta^i z_t^w} u_{t+1}^w + \sqrt{\kappa z_t^i} u_{t+1}^g,$$
$$z_{t+1}^i = (1-\phi)\theta + \phi z_t^i - \sigma \sqrt{z_t^i} u_{t+1}^i,$$
$$z_{t+1}^w = (1-\phi^w)\theta^w + \phi^w z_t^w - \sigma^w \sqrt{z_t^w} u_{t+1}^w,$$
$$\pi_{t+1}^i = \pi_0 + \eta^w z_t^w + \sigma_\pi \varepsilon_{t+1}^i.$$

The 17 parameters were obtained to match the moments in Table 4 under the assumption that all countries share the same parameter values except for δ^i, which is distributed uniformly on $[\delta_h, \delta_l]$. The home country exhibits the average δ, which is equal to 0.36. The standard errors obtained by bootstrapping are reported between brackets.

Stochastic discount factor					
α (%)	χ	τ	γ	κ	δ
0.76	0.89	0.06	0.04	2.78	0.36
[0.29]	[0.56]	[0.30]	[0.13]	[1.74]	[0.94]
State variable dynamics					
ϕ	θ (%)	σ (%)	ϕ^w	θ^w (%)	σ^w (%)
0.91	0.77	0.68	0.99	2.09	0.28
[0.04]	[0.19]	[0.25]	[0.07]	[2.30]	[0.67]
Inflation dynamics				Heterogeneity	
η^w	π_0 (%)	σ^π (%)		δ_h	δ_l
0.25	−0.31	0.37		0.22	0.49
[0.58]	[0.08]	[0.01]		[0.51]	[0.75]
Implied SDF dynamics					
$E(Std_t(m))$	$Std(Std_t(m))$ (%)	$E(Corr(m_{t+1}, m_{t+1}^i))$		$Std(z)$ (%)	$Std(z^w)$ (%)
0.59	4.21	0.98		0.50	1.32

estimated model to study the properties of currency excess returns in simulated data.

5.5. Simulated dollar and HML carry trades

We use the estimated parameters to simulate the model under the assumption that the world is populated by $N=30$ countries with the distribution of δ^i's assumed uniform (given that it is approximately symmetric around δ). We simulate the model for $T=336$ periods and use the artificial data to form the currency trading strategies analogous to those analyzed in Section 3. Table 6 reports the moments of the dollar and HML carry trade returns generated by the model. Standard deviations (in brackets, except for the exogenously fixed standard deviation of levered strategies) are obtained by simulating the model 1,000 times and taking standard deviations for each of the measures across the simulations. The unlevered average dollar carry trade return is only 1.86% in the model, and the Sharpe ratio is half of that in the data (0.24 vs. 0.56), with substantial uncertainty around both of the estimates. Levering up the return as we do in the data to match the empirically observed stock market volatility of 15.5% produces an average dollar carry excess return of 3.72% per annum. Hence, the dollar premium in the model is much smaller than that in the data. We can increase the dollar premium by increasing the volatility of country-specific risk, but only at the cost of overshooting the volatility of exchange rates. The HML carry trade, on the contrary, is roughly as profitable in the model as in the data, with a Sharpe ratio of 0.46; it is 0.31 in the data for developed countries and 0.49 in our large sample of developed and emerging countries, albeit with a higher volatility (15.30% in the model vs. around 10% in the data). The correlation between the dollar and HML currency carry trade is 0.19, which is higher than in the data but still very low. Overall, the model reproduces both of the carry trade strategies, but it has more difficulty matching their volatility (overstating it for the *HML*), while understating the dollar carry trade returns, on average.

5.6. Return predictability in simulated data

The model generates sizable returns on levered dollar carry trade strategies because the average excess returns for baskets of currencies are predictable by AFDs in the model, in particular at longer horizons.

Table 7 displays the results of the long-horizon predictability regressions for two types of samples simulated from the calibrated model. The term structure of interest rates—and therefore the long-horizon forward discounts—is computed in closed form. Panel A displays the results generated using a single long sample of length $T=33{,}600$ that is meant to closely approximate population values.

Table 6
Excess returns on carry strategies: simulated data. This table reports the mean, standard deviation and Sharpe ratios of two investment strategies using returns simulated from the model. The first strategy, the *Dollar* carry, is conditional: it goes long the basket when the average forward discount is positive, and short the same basket otherwise. The second strategy, the *HML* carry, corresponds to currency carry trades (long in a basket of high interest rate currencies, short in a basket of low interest currencies). Currency excess returns are sorted by interest rates into five portfolios. Excess returns are annualized (means are multiplied by 12 and standard deviations are multiplied by $\sqrt{12}$). Sharpe ratios correspond to the ratios of annualized means to annualized standard deviations. We report (in brackets) standard errors on the means calculated by simulating 1,000 samples of length 336 periods and taking the standard deviations of the means across simulations. Scaled returns are constructed to match the exogenously set volatility of 15.5%, which corresponds to U.S. equity return volatility in the data.

	Raw returns		Scaled returns	
	Dollar	HML	Dollar	HML
Mean	1.86 [1.34]	6.92 [2.62]	3.72 [2.72]	7.13 [2.92]
Std.dev.	7.89 [0.57]	15.30 [0.94]	15.50	15.50
Sharpe ratio	0.24 [0.18]	0.46 [0.19]	0.24 [0.18]	0.46 [0.19]

Table 7
Forecasting basket returns and exchange rates: simulated data. This table reports the slope coefficients and R^2 for the regressions of excess returns and exchange rate changes on the average forward discount implied by the calibrated model. Left panel presents the large sample regression results, generated from a single simulated sample of length $T=33{,}600$, with Hansen-Hodrick *t*-statistics in brackets. Right panel reports small sample results obtained by averaging over the $N=1{,}000$ point estimates using simulated samples of length $T=336$. The *t*-statistics in parentheses use standard deviations of the point estimates across simulations as standard errors.

	Large sample				Small sample			
	Excess returns		Exchange rates		Excess returns		Exchange rates	
Horizon	ψ_f	R^2	ζ_f	R^2	ψ_f	R^2	ζ_f	R^2
1	2.63 [11.63]	0.71	1.63 [7.20]	0.27	3.28 [1.42]	1.45	2.28 [0.99]	0.95
2	2.61 [12.61]	1.30	1.61 [7.77]	0.50	3.27 [1.41]	2.70	2.27 [0.98]	1.77
3	2.61 [13.03]	1.81	1.61 [8.04]	0.70	3.28 [1.41]	3.80	2.28 [0.98]	2.52
6	2.54 [12.94]	2.76	1.54 [7.84]	1.03	3.32 [1.42]	6.36	2.32 [0.99]	4.30
12	2.40 [11.59]	3.30	1.40 [6.76]	1.15	3.38 [1.41]	9.22	2.38 [0.99]	6.42

Panel B presents results using a large number (1,000) of small samples of length $T=336$ as in the actual data that allows us to evaluate the finite sample properties of these regressions and compare them to the empirical counterparts in Table 2.

The slope coefficients for average excess returns and average exchange rate changes are very similar in the actual and in the simulated data, as are the R^2's. In the long sample regressions the slope coefficient is about 2.4 for returns and 1.4 for exchange rate changes at the one-month horizon, declining only slightly at the 12-month horizon. The "population" R^2's appear slightly smaller than in the data whereas the small sample R^2's are of essentially the same magnitudes as those observed empirically. In the latter case, the AFD can explain 1.5% of the variation in one-month excess returns, and 9.5% of variation in 12-month returns, almost as much as in the data; similarly, it explains up to 6.7% of variation in average exchange rate changes over twelve month periods. However, comparison of the two sets of results indicate that both the slope coefficients and the R^2's of the small sample regressions may be upwardly biased compared to their population values. They are also potentially imprecisely estimated as the *t*-statistics for the small sample regressions based on the simulated distribution of estimated coefficients are in the neighborhood of 1.4 for returns and of 1.0 for exchange rate changes.

This simulation evidence indicates that even though some of the key model parameters are not estimated precisely, they are well-identified by the predictability moments that we chose to overweight in our estimation procedure. In particular, the point estimate of κ is 2.78 with a standard error of 1.74, as reported in Table 5. In our model this parameter is identified sharply by the predictability regressions of dollar excess returns with the average forward discounts, as well as the properties of returns on the dollar carry strategy. These empirical moments can be used to test the hypothesis $\kappa=0$ by simulating the model under the null (i.e., re-estimating the model under the restriction $\kappa=0$) and comparing the empirically observed values of the moments with the Monte Carlo distributions of their model analogues. Fig. 3 displays histograms of excess return predictability coefficients, R^2, dollar carry expected returns, and Sharpe ratios for 1,000 Monte Carlo simulations using both the full model and the restricted model with $\kappa=0$, along with the empirically observed values of the moments. It is apparent from these graphs that the probabilities of generating values as large as those observed in the data are negligible under the null of $\kappa=0$ (always less than 5%), while under the full model the observed moments are comfortably within the 95% confidence intervals. Therefore, even though there is substantial statistical uncertainty around the estimated parameter κ, it is clearly significantly greater than zero, statistically as well as economically.

5.7. Base carry vs. HML carry

So far we have explored the quantitative implications of the model under the assumption that the home country, on average, has exposure δ that is equal to the average exposure across all countries in the basket. Since the heterogeneity in these exposures is necessary to explain the dispersion in AFDs and unconditional currency risk premia across currencies in the data, it is interesting to explore the predictions of the model for the returns on currency baskets formed from the perspective of different base currencies. Again, we refer to strategies of going long

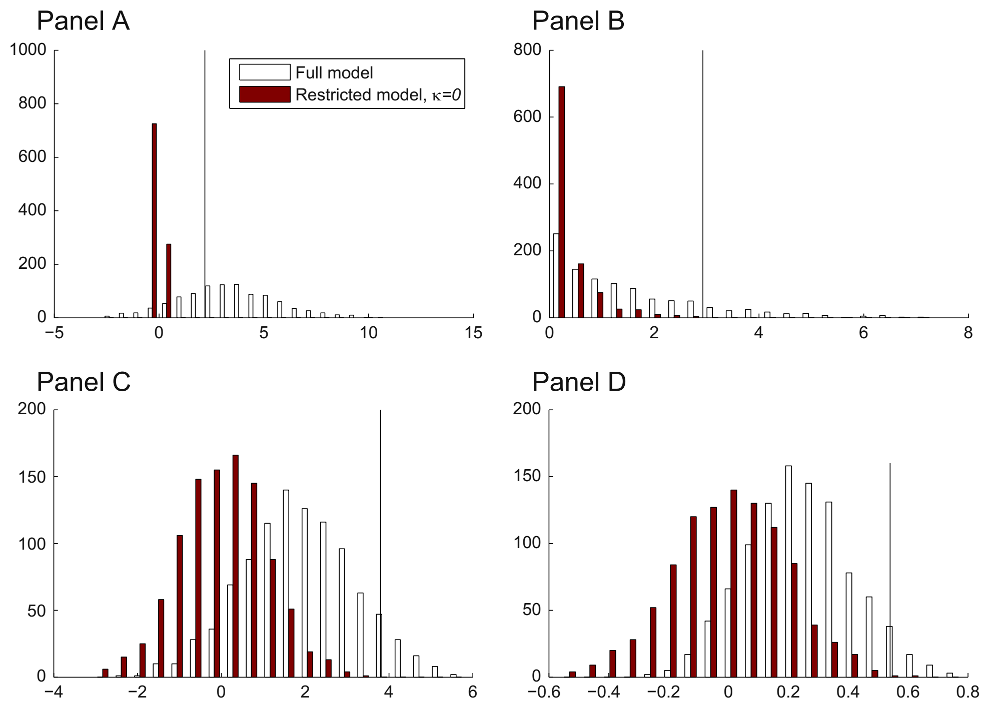

Fig. 3. Return predictability and dollar carry: Monte Carlo simulations. Histograms of estimated predictive regression coefficients $\psi_{\mathbf{f}}$ (Panel A) and the corresponding R^2 (Panel B), as well as average returns and Sharpe ratios on the dollar carry strategies (Panels C and D, respectively), simulated 1,000 times from the model using samples of length $T=336$ periods. Each moment is simulated using both the full model estimates and the estimates obtained under the restriction κ=0. The vertical lines depict the empirical estimates of the moments. Data are monthly, from Barclays and Reuters (available via Datastream). The sample period is 11/1983–6/2010. Panel A. Excess return predictability coefficient, ψ_f, Panel B. Excess return predictability −R2, Panel C. Dollar carry expected excess return and Panel D. Dollar carry Sharpe ratio.

a basket of all currencies when the AFD on foreign currency is positive and short otherwise from a perspective of a given country as base carry.

We compute the returns on the HML carry strategy from the perspective of different base currencies. The model predicts that the correlation between two strategies is U-shaped as a function of the mean of the basket's average forward discount. Countries with 'average' loadings δ will have low correlation between the two strategies as documented above, whereas countries with either high or low exposures will have higher correlations since they will tend to systematically have either lower or higher interest rates, respectively, than the average country, causing the base carry strategy to correlate with the HML strategy more often as they both load on the common shock u^g in the same direction.

Fig. 4 compares the correlations simulated from the model (solid line) to those observed in the data for the so-called G10 currencies, i.e., the ten most traded currencies: U.S. Dollar, U.K. Pound Sterling, Deutschmark/Euro, Australian, Canadian, and New Zealand Dollars, Japanese Yen, Swiss Franc, Norwegian Krone, and Swedish Krona. Consistent with the model's prediction, for base currencies that have unconditional mean of the average forward discount close to zero (e.g., U.S., U.K., Canada, Sweden), the correlation between the two carry strategies is also close to zero. At the same time, for currencies that on average exhibit high interest rates and therefore low mean AFD, these correlations are substantially higher, consistent with below-average global factor exposure δ (Australia, New Zealand). The correlations are also somewhat higher for countries with usually low interest rates, and therefore positive mean AFD, such as Japan and Switzerland, suggesting they may have above-average δs. Note that this result is not mechanical, since the HML strategy formed from the perspective of a given base currency does not include the base currency itself, whereas the base carry strategy always has the base currency in one leg of the position and an equal-weighted average of all other currencies in the other leg (otherwise the correlation would be much higher due to the overlap for the extreme currencies, such as Japan or New Zealand). This pattern of correlations therefore supports the two-common-factors structure of our proposed pricing kernel.

5.7.1. *Cross-section of predictability*

Given our calibration results, it appears that the average exposure assumption fits the data well for the U.S. vis-a-vis the group of developed countries used to form our benchmark currency basket. However, we should not expect the same results to hold for baskets formed from the perspective of any arbitrary country—only for countries that exhibit an 'average' exposure to the global shocks.

We compared the U.S. predictability results to those obtained for baskets formed from the perspectives of the other base currencies. For the U.K. and Canada—countries with a sample mean of the AFD close to zero, similar to the

U.S.—we find strong predictability of returns on long positions in foreign currencies and short positions in domestic currency, consistent with the model. However, for Japan and Switzerland, which have much higher sample means for the average forward discount, there is much weaker evidence of predictability, suggesting that these countries′ loadings on the global shocks (δ^i) are greater than the average developed country. Similarly, for Australia and New Zealand, whose AFDs on foreign currency are much lower than average as they typically have high (real) interest rates, which from the standpoint of the model implies low loading on the global shocks, there is also little evidence of predictability of excess returns or exchange rate changes.

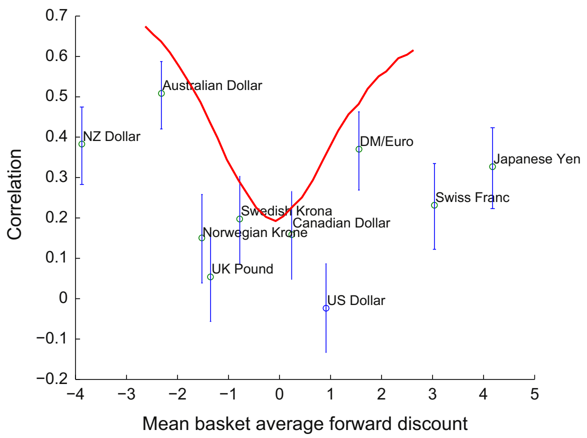

Fig. 4. Correlations Between base carry and HML carry trades. The solid line plots the predicted correlations simulated from the model using samples of length $T = 100{,}000$ periods. Each circle plots the correlation between the returns on a base carry strategy (long or short the basket of all foreign currencies depending on the sign of the average forward discount, from the perspective of a given base country) and returns on a global (*hml*) carry strategy (long the portfolio of high interest rate currencies, short the portfolio of low interest rate currencies, based on six forward-discount sorted portfolios formed from the perspective of the same base currency). The vertical error bars depict 95% confidence intervals for these correlations. Data are monthly, from Barclays and Reuters (available via Datastream). The sample period is 11/1983–6/2010.

5.8. Dollar carry and world equity returns

In Section 3 we showed that the correlation between the stock market returns and the dollar carry trade is very low. Given the fact that in our model global innovations play a major role in generating the dollar carry trade, it is a natural question whether the model produces a reasonable correlation between the dollar carry and the global stock market returns. To answer this question, we extend the model to equity markets. The aggregate stock market process in each country is a claim to a stream of dividends. We assume that the dividend growth process of country i follows:

$$\Delta d^i_{t+1} = \mu_D + \psi_w z^w_t + \sigma_D \sqrt{z^i_t} u^i_{t+1} + \sigma^g_D \sqrt{z^i_t} u^g_{t+1},$$

so that the global state variable affects the conditional mean (e.g., long-run growth is common across countries), but contemporaneous innovations are both common and idiosyncratic. The logarithmic gross real rate of return on each country's stock market index denominated in that country's currency r^{ei}_{t+1} is obtained using a standard log-linear approximation (detailed in the Supplementary On-line Appendix).

We construct an equal-weighted portfolio of all countries′ equity indexes denominated in the nominal home country units (U.S. dollars). The gross returns on this portfolio are calculated as

$$R^W_{t+1} = \frac{1}{N} \sum_{i=1}^{N} \exp\left(r^{ei}_{t+1} - \Delta q^i_{t+1} + \pi_{t+1}\right).$$

The four parameters of the dividend processes (μ_D, ψ_w, σ_D, and σ^g_D) are estimated in order to match the key moments of global dividend growth processes as well as the world equity market returns reported in Table 8. Specifically, the estimation matches the cross-country

Table 8
World equity market and the dollar carry trade.
This table compares the world equity market and dollar carry trade returns in the model and in the data. In each country, the dividend growth process is simulated as

$$\Delta d^i_{t+1} = \mu_D + \psi_w z^w_t + \sigma_D \sqrt{z^i_t} u^i_{t+1} + \sigma^g_D \sqrt{z^i_t} u^g_{t+1}. \tag{33}$$

The model parameters (μ_D, ψ_w, σ_D, and σ^g_D) are calibrated using the dividend series built by Datastream for the following countries: Australia, Canada, Denmark, Germany, Japan, New Zealand, Norway, Sweden, Switzerland, U.K., and U.S. The sample period is 11/1983–6/2010. Matching the cross-country averages of means and standard deviations of the dividend growth series, the average R^2 of the regression of each country's dividend growth on the world-average dividend growth, and the mean of the world stock market index return leads to $\mu_D = 0.01$, $\psi_w = -2.00$, $\sigma_D = 0.40$, and $\sigma^g_D = 0.09$. The mean of the monthly, nominal dividend growth rate in country i is denoted $E(\Delta d^i_t + \pi^i_t)$; its standard deviation is denoted $\sigma(\Delta d^i_t + \pi^i_t)$. Both the mean and the standard deviation are annualized (multiplied by 12 and $\sqrt{12}$). The table reports the cross-country average of those two moments, denoted by an upper bar. Each country's 12-month dividend growth is regressed on the world-average 12-month dividend growth. The world series is obtained as an equally weighted average of the country series. The table reports the cross-country average (denoted Avg. R^2) of all the corresponding R^2's. The table also reports the annualized mean (denoted $E(R^W_t)$) and standard deviation (denoted $\sigma(R^W_t)$) of the world equity return, which is constructed as an equal-weighted portfolio of country indexes with monthly rebalancing, expressed in U.S. dollars. The last column of the table reports the correlation between the monthly world equity returns and the dollar carry trade returns.

	$\overline{E(\Delta d^i_t + \pi^i_t)}$	$\overline{\sigma(\Delta d^i_t + \pi^i_t)}$	Avg. R^2	$E(R^W_t)$	$\sigma(R^W_t)$	$corr(R^W_t, Carry^\$)$
Data	10.85	12.89	7.64	12.17	14.97	0.14
Model	9.79	12.36	7.91	11.98	9.20	0.10

averages of means and standard deviations of the dividend growth series, the average R^2 of the regression of each country's dividend growth on the world-average dividend growth, and the mean of the world stock market index return. These moments are obtained by computing dividends from the dividend yield and market capitalization series built by Datastream. The sample comprises the same set of countries used before in the estimation of the model. The sample period is 11/1983–6/2010, again as in the rest of the paper. The parameters governing dividend dynamics are listed in the caption of Table 8.

We match the average mean and standard deviation of the log dividend growth processes (at monthly frequency, nominal, in local currencies) fairly closely (the mean is 9.79% in the model vs. 10.85% in the data, and the volatility is around 12% in both cases), as well as the average R^2 of the regression of each country's 12-month dividend growth on the average dividend growth rate (just under 8% in both model and data). The world equity market is obtained as an equal-weighted portfolio of these indexes (in nominal U.S. dollar terms) both in the model and in the data. Table 8 reports the mean and standard deviation for this return. We match the former very closely, at 12% per annum in both model and data, but we undershoot the global market volatility: the model implies a standard deviation of 9.20% vs. 14.97% in the data. The key implication of this calibration is that the correlation between the world equity market return (denominated in dollars) and the dollar carry return is no greater in the model than it is in the data. The correlation between world market returns and dollar carry returns is 0.1 in the model vs. 0.14 in the data. The model, thus, does not appear to overestimate the common variation in pricing kernels and is a useful framework for understanding international equity markets.

6. Countercyclical currency risk premia

Our empirical results imply that expected excess returns on currency portfolios vary over time. The no-arbitrage model in Section 4 suggests that this variation is driven by time-variation in the U.S.-specific prices of domestic and global risk. A large literature in empirical asset pricing suggests that risk premia in equity markets and bond markets increase in economic downturns. Consequently, we expect the dollar risk premium to be countercyclical with respect to the U.S.-specific component of the business cycle. Indeed, in this section we show that time-variation in conditional expected returns on U.S.-based currency baskets has a large U.S. business cycle component: expected excess returns go up in U.S. recessions and go down in U.S. expansions, which is similar to the countercyclical behavior that has been documented for bond and stock excess returns. This feature of asset markets is a key ingredient of leading dynamic asset pricing models (see Campbell and Cochrane, 1999; Bansal and Yaron, 2004 for prominent examples). The evidence at the level of currency baskets is strong enough to survive most out-of-sample tests. Consistent with the predictions of our model, this is not true for baskets formed from the perspectives of all countries. Finally, we link time-varying risk premia to time-varying aggregate consumption and inflation volatility.

6.1. Cyclical properties of the average forward discount

We use $\widehat{\mathrm{E}}_t rx_{t+1}$ to denote the forecast of the one-month-ahead excess return based on the AFD for a basket:

$$\widehat{\mathrm{E}}_t rx_{t+1} = \psi_0 + \psi_{\mathbf{f}}(\overline{f}_{t\to t+k} - \overline{s}_t). \tag{26}$$

Therefore, expected excess returns on currency baskets inherit the cyclical properties of the AFDs. To assess the cyclicality of these forward discounts, we use three standard business cycle indicators—the 12-month percentage change in U.S. industrial production index (IP), the 12-month percentage change in total U.S. non-farm payroll index (*Pay*), and the 12-month percentage change in the Help Wanted index (*Help*)—and three financial variables—the term spread (i.e., difference between the 20-year and the one-year Treasury zero-coupon yields, *Term*), the default spread (the difference between the BBB and AAA bond yields, *Def*), and the Chicago Board Options Exchange Market Volatility Index (VIX).[14] Macroeconomic variables are often revised. To check that our results are robust to real-time data, we use vintage series of the payroll and industrial production indices from the Federal Reserve Bank of Saint Louis. The results are very similar. Note that macroeconomic variables are also published with a lag. For example, the industrial production index is published around the 15th of each month, with a one-month lag (e.g., the value for May 2009 was released on June 16, 2009). In our tables, we do not take into account this publication lag of 15 days or so and assume that the index is known at the end of the month. We check our results by lagging the index an extra month. The publication lag sometimes matters for short-horizon predictability but does not change our results over longer horizons.

Table 9 reports the contemporaneous correlations of the AFDs (across horizons) with these macroeconomic and financial variables. As expected, the AFD (and, therefore, forecasted excess returns) is countercyclical: it is negatively correlated with all of the macroeconomic variables (IP, *Pay*, and *Help*). We find roughly the same business cycle variation in AFD across horizons. At every maturity we consider, the AFD appears countercyclical. Since excess returns load positively on the AFD, they are also countercyclical, i.e., high in bad times and low in good times. AFDs are positively correlated with the slope of the U.S. term structure and the default spread, again suggesting that basket excess returns are countercyclical. The AFD, however, is not correlated with the VIX index (point estimates are negative but small and not statistically significant), pointing again to the difference between dollar carry and HML carry trades, which are correlated with stock market volatility (as shown in Lustig, Roussanov, and Verdelhan, 2011).

[14] Industrial production data are from the International Monetary Fund's International Financial Statistics. The payroll index is from the Bureau of Economic Analysis. The Help Wanted Index is from the Conference Board. Zero-coupon yields are computed from the Fama-Bliss series available from CRSP. The VIX index as well as the corporate and Treasury bond yields are from Datastream.

Table 9
Contemporaneous correlations between average forward discounts and macroeconomic and financial variables. This table reports the contemporaneous correlation between average forward discounts of developed countries and different macroeconomic and financial variables: the 12-month percentage change in industrial production (*IP*), the 12-month percentage change in the total U.S. non-farm payroll (*Pay*), and the 12-month percentage change of the Help-Wanted index (*Help*), the default spread (*Def*), the slope of the yield curve (*Term*) and the CBOE volatility index (*VIX*). Standard errors in brackets are computed using stationary block-bootstrap. Data are monthly. The sample period is 11/1983–6/2010.

Horizon	*IP*	*Pay*	*Help*	*Term*	*Def*	*VIX*
1	−0.29	−0.20	−0.12	0.48	0.22	−0.09
	[0.05]	[0.06]	[0.06]	[0.04]	[0.06]	[0.05]
2	−0.31	−0.21	−0.14	0.49	0.23	−0.09
	[0.05]	[0.06]	[0.06]	[0.04]	[0.06]	[0.05]
3	−0.31	−0.21	−0.15	0.49	0.23	−0.08
	[0.05]	[0.06]	[0.06]	[0.04]	[0.06]	[0.05]
6	−0.33	−0.23	−0.20	0.49	0.20	−0.06
	[0.05]	[0.06]	[0.06]	[0.04]	[0.06]	[0.05]
12	−0.37	−0.28	−0.28	0.48	0.15	−0.05
	[0.05]	[0.06]	[0.06]	[0.05]	[0.06]	[0.05]

6.2. Macro factors and currency return predictability

So far we have focused on the predictive power of the AFD, but the countercyclical nature of excess returns suggests that macro variables themselves might help to forecast excess returns, potentially above and beyond what is captured by the AFDs. We check this conjecture by focusing on the predictive power of the IP variable, controlling for the AFD.

In the benchmark version of the model, there is only a single state variable that describes the market price of U.S. risk, and it is spanned by the average forward discount. In a model with more state variables, the average forward discount is a linear combination of these state variables, and as long as the SDF loadings on these variables (χ, γ, and κ) differ, the conditional expected returns are no longer proportional to the AFD. Adding other interest rate-related variables, such as the slope of the term structure, may or may not help identify these other factors. Evidence from the term structure of U.S. interest rates suggests that business cycle variables, such as IP growth rates, contain information about risk premia in the bond markets that is not captured by the interest rates themselves (see Ludvigson and Ng, 2009; Duffee, 2011; Joslin, Priebsch, and Singleton, 2010). We thus expect similar U.S.-specific macroeconomic variables to have forecasting power for currency excess returns, as well as spot exchange rate changes.

6.2.1. Industrial production growth

Using the same notation as in the previous sections, Table 10 reports two sets of regression results:

$$\overline{rx}_{t\to t+k}=\psi_0+\psi_{IP}\Delta\log IP_t+\psi_{\mathbf{f}}(\overline{f}_{t\to t+k}-\overline{s}_t)+\eta_{t+k}, \tag{27}$$

$$-\overline{\Delta s}_{t\to t+k}=\zeta_0+\zeta_{IP}\Delta\log IP_t+\zeta_{\mathbf{f}}(\overline{f}_{t\to t+k}-\overline{s}_t)+\varepsilon_{t+k}. \tag{28}$$

We use the developed markets′ AFD since it is the strongest predictor of returns on all baskets (developed, emerging, or all countries). All the estimated slope coefficients on industrial production are negative and, for horizons of three months and above, strongly statistically significant. The Wald tests reject the restriction that the two slope coefficients for excess returns are jointly equal to zero for all baskets at horizons of three months and above (using various methods) and, for exchange rate changes, at horizons of six and 12 months. The R^2's for average excess returns at the 12-month horizon are between 23% and 28% across different baskets, and between 15% and 34% for average exchange rate changes.

Since we are controlling for the average forward discount of the developed markets basket, the *IP* coefficient for this basket is the same for excess returns and exchange rate changes, capturing the pure effect of the countercyclical risk premium on expected depreciation of the dollar, rather than the return stemming from the interest rate differential. Thus, holding interest rates constant, a one percentage point drop in the annual change in IP raises the dollar risk premium by 50–100 basis points per annum at the monthly horizon and by as much as 90–153 basis points at the annual horizon, all coming from the expected appreciation of the foreign currencies against the dollar. Since the AFD itself is countercyclical, the total effect is even greater, implying an increase in expected returns of up to 120 basis points for annual data.

The U.S. IP appears highly correlated with similar indices in other developed countries. For example, its correlation with the average index for the G7 countries (excluding the U.S., and using 12-month changes in each index) is equal to 0.5.[15] To check that the U.S.-specific component of the U.S. IP index matters most here, we run the following predictive regressions using the residuals from the projection of these 12-month changes on the average foreign IP indices:

$$\Delta\log IP_t=\alpha+\beta\overline{\Delta\log IP_t}+IP_{res,t}, \tag{29}$$

$$\overline{rx}_{t\to t+k}=\psi_0+\psi_{IP_{res}}IP_{res,t}+\psi_{\mathbf{f}}(\overline{f}_{t\to t+k}-\overline{s}_t)+\eta_{t+k}, \tag{30}$$

where $\overline{\Delta\log IP_t}$ denotes the average of the 12-month changes in IP indices across 28 developed countries (excluding the U.S.).

The predictive power of *IP* lies mostly in the U.S.-specific component of *IP*, denoted $IP_{res,t}$, for long-horizon returns. We obtain R^2's between 16% and 25% with the IP residuals for both average excess returns and average spot exchange rate changes at the 12-month horizon. The slope coefficients are lower for the short-horizon returns, but larger for long horizons. For annual holding periods, a one percentage point decline in the U.S. IP relative to the world average implies a 145–165 basis point increase in the risk premium, even if interest rate differentials do not change.

To check that the U.S.-specific component of the AFD of developed countries matters most here, we run predictability tests using the residuals from the projection of the AFD on the average 12-month changes of foreign countries′ industrial production indices, which removes much of the

[15] The G7 countries are Canada, France, Germany, Italy, Japan, and the U.K., as well as the U.S.

Table 10

Forecasting excess returns and exchange rates with industrial production and the average forward discount. This table reports results of forecasting regressions for average excess returns and average exchange rate changes for baskets of currencies at horizons of one, two, three, six, and 12 months. For each basket we report the R^2, and the slope coefficients in the time-series regression of the log currency excess return on the 12-month change in the U.S. Industrial Production Index (ψ_{IP}) and on the average log forward discount ($\psi_{\mathbf{f}}$), and similarly the slope coefficients ζ_{IP}, $\zeta_{\mathbf{f}}$, and the R^2 for the regressions of average exchange rate changes. The *t*-statistics for the slope coefficients in brackets are computed using the following methods. *HH* denotes Hansen and Hodrick (1980) standard errors computed with the number of lags equal to the length of overlap plus one lag. The *VAR*-based statistics are adjusted for the small-sample bias using the stationary bootstrap distributions of slope coefficients under the null hypothesis of no predictability, estimated by drawing random blocks of residuals of a VAR with the number of lags equal to the length of overlap plus one lag. Data are monthly, from Barclays and Reuters (available via Datastream). We also report the Wald tests (*W*) of the hypothesis that both slope coefficients are jointly equal to zero; the percentage *p*-values in brackets are for the χ^2-distribution under the parametric cases (*HH*) and for the bootstrap distribution of the *F*-statistic under *VAR*. Data are monthly, from Barclays and Reuters (available via Datastream). The returns do not take into account bid–ask spreads. The sample period is 11/1983–6/2010.

	Developed countries								Emerging countries								All countries							
	Excess returns				Exchange rates				Excess returns				Exchange rates				Excess returns				Exchange rates			
Horizon	ψ_{IP}	$\psi_{\mathbf{f}}$	W	R^2	ζ_{IP}	$\zeta_{\mathbf{f}}$	W	R^2	ψ_{IP}	$\psi_{\mathbf{f}}$	W	R^2	ζ_{IP}	$\zeta_{\mathbf{f}}$	W	R^2	ψ_{IP}	$\psi_{\mathbf{f}}$	W	R^2	ζ_{IP}	$\zeta_{\mathbf{f}}$	W	R^2
1	−0.54	2.14	7.00	3.40	−0.54	1.14	3.16	1.53	−1.15	−0.20	3.78	2.74	−1.15	−1.20	6.73	4.93	−0.65	1.68	5.00	2.72	−0.65	0.68	2.49	1.41
HH	[−0.96]	[2.06]	[1.24]		[−0.96]	[1.10]	[29.79]		[−1.95]	[−0.27]	[28.75]		[−1.95]	[−1.57]	[7.64]		[−1.23]	[1.50]	[10.49]		[−1.23]	[0.61]	[49.35]	
VAR	[−1.02]	[2.32]	[0.00]		[−0.97]	[1.26]	[0.00]		[−2.39]	[−0.47]	[0.00]		[−2.32]	[−2.26]	[0.00]		[−1.31]	[1.66]	[0.00]		[−1.41]	[0.70]	[0.10]	
2	−0.65	2.09	10.35	6.25	−0.65	1.09	6.71	3.14	−1.17	−0.64	7.10	6.54	−1.17	−1.64	6.64	11.66	−0.74	1.64	7.53	5.24	−0.74	0.64	4.97	3.06
HH	[−1.34]	[2.02]	[0.63]		[−1.34]	[1.05]	[17.97]		[−2.38]	[−0.80]	[9.83]		[−2.38]	[−2.05]	[2.40]		[−1.65]	[1.52]	[3.43]		[−1.65]	[0.60]	[25.98]	
VAR	[−1.24]	[1.90]	[0.00]		[−1.21]	[1.02]	[0.00]		[−2.35]	[−1.14]	[0.00]		[−2.32]	[−2.78]	[0.00]		[−1.41]	[1.52]	[0.00]		[−1.60]	[0.55]	[0.10]	
3	−0.72	1.99	23.67	8.68	−0.72	0.99	19.77	4.65	−1.28	−0.54	8.01	9.81	−1.28	−1.54	7.59	15.74	−0.82	1.52	10.17	7.57	−0.82	0.52	9.45	4.77
HH	[−1.66]	[1.97]	[0.43]		[−1.66]	[0.98]	[12.21]		[−2.71]	[−0.68]	[3.00]		[−2.71]	[−1.94]	[1.33]		[−2.08]	[1.53]	[1.72]		[−2.08]	[0.53]	[13.49]	
VAR	[−1.28]	[1.69]	[0.00]		[−1.49]	[0.96]	[0.00]		[−2.64]	[−0.92]	[0.00]		[−2.67]	[−2.36]	[0.00]		[−1.76]	[1.23]	[0.00]		[−1.76]	[0.41]	[0.00]	
6	−0.87	1.84	38.02	15.58	−0.87	0.84	32.04	9.57	−1.48	−0.25	6.37	18.21	−1.48	−1.25	6.88	24.14	−0.96	1.59	11.94	15.92	−0.96	0.59	10.58	11.21
HH	[−2.60]	[2.03]	[0.00]		[−2.60]	[0.93]	[0.53]		[−3.06]	[−0.35]	[0.27]		[−3.06]	[−1.74]	[0.50]		[−3.15]	[2.06]	[0.01]		[−3.15]	[0.76]	[0.22]	
VAR	[−1.71]	[1.78]	[0.00]		[−1.85]	[0.84]	[0.00]		[−3.46]	[−0.46]	[0.00]		[−3.24]	[−1.87]	[0.00]		[−2.16]	[1.37]	[0.00]		[−2.36]	[0.51]	[0.00]	
12	−0.91	1.37	16.75	23.20	−0.91	0.37	13.05	15.16	−1.53	−0.07	7.37	28.40	−1.53	−1.07	7.35	34.51	−1.00	1.14	12.55	24.36	−1.00	0.14	10.25	18.49
HH	[−3.39]	[1.50]	[0.00]		[−3.39]	[0.41]	[0.00]		[−3.06]	[−0.08]	[0.24]		[−3.06]	[−1.24]	[0.60]		[−3.64]	[1.71]	[0.00]		[−3.64]	[0.21]	[0.01]	
VAR	[−2.15]	[1.35]	[0.00]		[−2.23]	[0.40]	[0.10]		[−5.27]	[−0.17]	[0.00]		[−5.00]	[−1.77]	[0.00]		[−2.89]	[1.18]	[0.00]		[−2.93]	[0.13]	[0.00]	

Table 11

Forecasting returns and exchange rates with realized consumption volatility and inflation volatility. This table reports results of forecasting regressions for average excess returns and average exchange rate changes for baskets of currencies at horizons of one, two, three, six, and 12 months. For each basket we report the R^2, and the slope coefficients in the time-series regression of the log currency excess return on the 36-month standard deviations of the monthly growth rates of the U.S. aggregate consumption of nondurables and services (ψ_{σ_c}) and of the corresponding inflation rate (ψ_{σ_p}), and similarly the slope coefficients ζ_{σ_c}, ζ_{σ_p}, and the R^2 for the regressions of average exchange rate changes. The *t*-statistics for the slope coefficients in brackets are computed using the following methods. *HH* denotes Hansen and Hodrick (1980) standard errors computed with the number of lags equal to the length of overlap plus one lag. The *VAR*-based statistics are adjusted for the small-sample bias using the stationary bootstrap distributions of slope coefficients under the null hypothesis of no predictability, estimated by drawing random blocks of residuals of a VAR with the number of lags equal to the length of overlap plus one lag. Data are monthly, from Barclays and Reuters (available via Datastream). The returns do not take into account bid–ask spreads. The sample period is 11/1983–6/2010.

	Developed Countries								Emerging Countries								All Countries							
	Excess Returns				Exchange Rates				Excess Returns				Exchange Rates				Excess Returns				Exchange Rates			
Horizon	ψ_{σ_c}	ψ_{σ_p}	*W*	R^2	ζ_{σ_c}	ζ_{σ_p}	*W*	R^2	ψ_{σ_c}	ψ_{σ_p}	*W*	R^2	ζ_{σ_c}	ζ_{σ_p}	*W*	R^2	ψ_{σ_c}	ψ_{σ_p}	*W*	R^2	ζ_{σ_c}	ζ_{σ_p}	*W*	R^2
1	1.94	0.91	1.94	0.66	1.26	0.75	1.00	0.31	−1.66	1.90	1.82	0.62	−1.55	2.69	2.61	0.82	1.03	1.09	0.86	0.35	0.74	1.21	0.81	0.28
HH	[1.32]	[0.42]	[52.58]		[0.88]	[0.35]	[80.04]		[−1.04]	[0.91]	[65.47]		[−0.94]	[1.26]	[53.51]		[0.78]	[0.54]	[80.91]		[0.57]	[0.60]	[85.86]	
VAR	[0.84]	[0.30]	[0.20]		[0.43]	[0.30]	[0.30]		[−0.80]	[0.63]	[0.30]		[−0.74]	[0.91]	[0.30]		[0.48]	[0.36]	[0.30]		[0.27]	[0.45]	[0.30]	
2	2.19	0.91	1.83	1.48	1.52	0.72	1.12	0.76	−1.57	2.12	2.37	1.20	−1.47	2.92	2.87	1.58	1.20	1.18	0.91	0.82	0.92	1.30	0.87	0.67
HH	[1.48]	[0.43]	[39.43]		[1.06]	[0.34]	[69.10]		[−1.02]	[1.09]	[59.50]		[−0.90]	[1.46]	[44.12]		[0.90]	[0.63]	[71.60]		[0.71]	[0.69]	[77.14]	
VAR	[1.07]	[0.36]	[0.00]		[0.72]	[0.29]	[0.30]		[−0.80]	[0.79]	[0.10]		[−0.81]	[1.13]	[0.00]		[0.73]	[0.45]	[0.10]		[0.51]	[0.51]	[0.30]	
3	2.38	1.28	2.72	2.65	1.71	1.08	1.89	1.52	−1.49	2.58	3.63	1.93	−1.40	3.38	4.15	2.56	1.36	1.57	1.60	1.63	1.08	1.68	1.65	1.42
HH	[1.56]	[0.61]	[24.99]		[1.17]	[0.51]	[52.34]		[−0.93]	[1.38]	[47.94]		[−0.83]	[1.77]	[29.20]		[0.98]	[0.85]	[55.14]		[0.81]	[0.92]	[59.31]	
VAR	[1.13]	[0.47]	[0.00]		[0.81]	[0.37]	[0.10]		[−0.83]	[1.02]	[0.00]		[−0.75]	[1.36]	[0.00]		[0.75]	[0.65]	[0.00]		[0.52]	[0.66]	[0.10]	
6	2.95	1.36	4.57	6.73	2.33	1.12	4.46	4.54	−1.08	2.77	7.58	2.59	−0.96	3.59	6.90	3.68	1.86	1.62	2.43	4.23	1.63	1.75	2.77	3.95
HH	[1.96]	[0.62]	[6.48]		[1.61]	[0.51]	[19.50]		[−0.68]	[1.42]	[50.78]		[−0.58]	[1.80]	[30.16]		[1.34]	[0.84]	[34.89]		[1.22]	[0.91]	[32.85]	
VAR	[1.53]	[0.58]	[0.00]		[1.18]	[0.49]	[0.00]		[−0.66]	[1.17]	[0.10]		[−0.55]	[1.55]	[0.00]		[1.02]	[0.62]	[0.00]		[0.85]	[0.79]	[0.20]	
12	3.75	1.00	18.18	19.57	3.23	0.71	12.99	15.69	−0.35	1.86	3.29	1.45	−0.20	2.81	3.60	3.09	2.67	1.05	8.89	12.45	2.52	1.24	8.28	12.67
HH	[3.19]	[0.56]	[0.01]		[2.79]	[0.39]	[0.13]		[−0.23]	[0.88]	[86.32]		[−0.13]	[1.29]	[68.42]		[2.38]	[0.58]	[4.11]		[2.33]	[0.69]	[2.06]	
VAR	[1.97]	[0.52]	[0.00]		[1.83]	[0.36]	[0.00]		[−0.31]	[1.26]	[0.20]		[−0.15]	[1.94]	[0.10]		[1.71]	[0.66]	[0.00]		[1.62]	[0.69]	[0.10]	

covariation of the AFD with the global macroeconomic conditions. For the basket of developed currencies, the slope coefficients are only 20 basis points lower across different maturities, which could be explained by the noise introduced in estimation of the residual. For the other baskets, the results are similar.

6.2.2. Consumption volatility

In consumption-based asset pricing models, time-varying risk premia can arise due to heteroskedasticity of aggregate consumption growth (e.g. as in Kandel and Stambaugh, 1991; Bansal and Yaron, 2004); in models with a nontrivial nominal side, the conditional volatility of inflation can also play an important role in generating the forward premium (e.g., as in Backus, Gavazzoni, Telmer, and Zin, 2011). Indeed, the realized volatility of U.S. aggregate consumption growth (estimated as a rolling 36-month standard deviation of monthly growth rates) is highly countercyclical, as is the realized volatility of inflation. They both increase during recessions (when IP falls), as does the average forward discount for the developed countries. If the time-variation in the currency risk premia is due to the time-varying consumption uncertainty, we should be able to detect a relation between the two statistically, as long as a good empirical measure of consumption volatility can be constructed.

We run regressions of average excess currency returns and exchange rate changes on these two volatility measures: the estimated conditional volatilities of consumption growth $\sigma_t(\Delta c_{t+1})$ and of inflation $\sigma_t(\pi_{t+1})$:

$$\overline{rx}_{t\to t+k} = \psi_0 + \psi_{\sigma_c}\sigma_t(\Delta c_{t+1}) + \psi_{\sigma_p}\sigma_t(\Delta\pi_{t+1}) + \eta_{t+k}, \tag{31}$$

$$-\overline{\Delta s}_{t\to t+k} = \zeta_0 + \zeta_{\sigma_c}\sigma_t(\Delta c_{t+1}) + \zeta_{\sigma_p}\sigma_t(\Delta\pi_{t+1}) + \varepsilon_{t+k}. \tag{32}$$

Results of these regressions are reported in Table 11. Consumption growth volatility appears to have substantial predictive power for currency excess returns and exchange rate changes, with R^2's as high as 19%. However, the finite-sample bias may potentially be severe, as the bootstrapped t-statistics are large enough for the coefficients to be statistically significant only at the 12-month horizon, and at all horizons are markedly different from the large asymptotic t-statistics.

Overall, we find that the expected returns on shorting the U.S. dollar are countercyclical: they increase when U.S. output declines (in particular, relative to the world average), and U.S. consumption growth volatility increases. This suggests that the market price of U.S.-specific risk, and thus z^i in our model, is countercyclical. Our model thus provides a potential explanation for our empirical findings—both the large excess returns on a novel trading strategy and the strong predictability results on currency excess returns—and this explanation is clearly based on countercyclical currency risk premia.

7. Conclusion

We document in this paper that aggregate returns in currency markets are highly predictable. This predictability manifests itself in the form of high Sharpe ratios on the dollar carry trade. The average forward discount and the change in the U.S. IP index explain up to 25% of the subsequent variation in average annual excess returns realized by shorting the dollar and going long in large baskets of currencies. The time-variation in expected returns has a clear business cycle pattern: U.S. macroeconomic variables are powerful predictors of these returns, especially at longer holding periods, and expected currency returns are strongly countercyclical.

We provide a simple, no-arbitrage model that reproduces our main findings and shows that the source of predictability is U.S.-specific variation in the price of global as well as U.S. specific risk that is unrelated to another global source of risk identified by the usual high-minus-low carry trades (e.g., global financial market volatility). However, the no-arbitrage model cannot fully match the dollar carry trade risk premium without imputing too much volatility to exchange rates. Perhaps, because the dollar is a reserve currency, investors are willing to forgo some return in exchange for a long position in dollars, especially in bad times. We leave this question for future research.

References

Adrian, T., Etula, E., Shin, H.S., 2010. Risk appetite and exchange rates. Federal Reserve Bank of New York Staff Reports 361, May.
Akram, Q.F., Rime, D., Sarno, L., 2008. Arbitrage in the foreign exchange market: turning on the microscope. J. Int. Econ. 76 (2), 237–253.
Andrews, D.W., 1991. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. Econometrica 59 (1), 817–858.
Ang, A., Bekaert, G., 2007. Stock return predictability: Is it there?. Rev. Financ. Stud. 20, 651–707.
Ang, A., Chen, J.S., 2010. Yield curve predictors of foreign exchange returns. unpublished working paper. Columbia University, University of California, Davis.
Asness, C.S., Moscowitz, T.J., Pedersen, L.H., 2013. Value and momentum everywhere. J. Finance 68 (3), 929–985.
Atkeson, A., Kehoe, P., 2009. On the need for a new approach to analyzing monetary policy. In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), NBER Macroeconomics Annual 2008, University of Chicago Press, Chicago, IL, pp. 389–425.
Baba, N., Packer, F., 2009. From turmoil to crisis: dislocations in the FX swap market before and after the failure of Lehman Brothers. J. Int. Money Finance 28, 1350–1374.
Backus, D., Foresi, S., Telmer, C., 2001. Affine models of currency pricing: accounting for the forward premium anomaly. J. Finance 56, 279–304.
Backus, D.K., Gavazzoni, F., Telmer, C., Zin, S.E., 2011. Monetary policy and the uncovered interest parity puzzle. NBER Working Paper No. 16218, July.
Bakshi, G., Carr, P., Wu, L., 2008. Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies. J. Financ. Econ. 87, 132–156.
Bakshi, G., Panayotov, G., 2013. Predictability of currency carry trades and asset pricing implications. J. Financ. Econ. 110 (1), 139–163.
Bakshi, G., Panayotov, G., Skoulakis, G., 2011. Improving the predictability of real economic activity and asset returns with forward variances inferred from option portfolios. J. Financ. Econ. 100, 475–495.
Bansal, R., Dahlquist, M., 2000. The forward premium puzzle: different tales from developed and emerging economies. J. Int. Econ. 51, 115–144.
Bansal, R., Shaliastovich, I., 2012. A long-run risks explanation of predictability puzzles in bond and currency markets. Rev. Financ. Stud. 26 (1), 1–33.
Bansal, R., Yaron, A., 2004. Risks for the long run: a potential resolution of asset pricing puzzles. J. Finance 59 (4), 1481–1509.
Bekaert, G., Hodrick, R.J., Marshall, D.A., 1997. On biases in tests of the expectations hypothesis of the term structure of interest rates. J. Financ. Econ. 44, 309–348.
Bhansali, V., 2007. Volatility and the carry trade. J. Fixed Income, 72–84.
Brandt, M.W., Cochrane, J.H., Santa-Clara, P., 2006. International risk-sharing is better than you think (or exchange rates are much too smooth). J. Monet. Econ. 53, 671–698.

Brennan, M.J., Xia, Y., 2006. International capital markets and foreign exchange risk. Rev. Financ. Stud. 19 (3), 753–795.
Brunnermeier, M.K., Nagel, S., Pedersen, L.H., 2009. Carry trades and currency crashes. In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), NBER Macroeconomics Annual, University of Chicago Press, Chicago, IL, pp. 313–347.
Burnside, C., Eichenbaum, M., Kleshchelski, I., Rebelo, S., 2011. Do peso problems explain the returns to the carry trade? Rev. Financ. Stud. 24 (3), 853–891.
Campbell, J., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. J. Polit. Econ. 107, 205–251.
Campbell, J.Y., Hamao, Y., 1992. Predictable stock returns in the United States and Japan: a study of long-term capital market integration. J. Finance 47, 43–69.
Campbell, J.Y., Medeiros, K.S.D., Viceira, L.M., 2010. Global currency hedging. J. Finance 65, 87–121.
Chen, Y.-c., Tsang, K.P., 2013. What does the yield curve tell us about exchange rate predictability? Rev. Econ. Stat. 95 (1), 185–205.
Chernov, M., Graveline, J., Zviadadze, I., 2012. Crash risk in currency returns. Unpublished working paper. London School of Economics, London Business School, University of Minnesota.
Cochrane, J.H., Piazzesi, M., 2005. Bond risk premia. Am. Econ. Rev. 95 (March (1)), 138–160.
Colacito, R., 2008. Six anomalies looking for a model. A consumption-based explanation of international finance puzzles. Unpublished working Paper. University of North Carolina at Chapel Hill.
Colacito, R., Croce, M.M., 2011. Risks for the long run and the real exchange rate. J. Polit. Econ. 119 (1), 153–182.
Cooper, I., Priestley, R., 2009. Time-varying risk premiums and the output gap. Rev. Financ. Stud. 22 (7), 2801–2833.
Duffee, G.R., 2011. Information in (and not in) the term structure. Rev. Financ. Stud. 24, 2895–2934.
Evans, M.D.D., Lewis, K.K., 1995. Do long-term swings in the dollar affect estimates of the risk premia? Rev. Financ. Stud. 8 (3), 709–742.
Fama, E.F., 1984. Forward and spot exchange rates. J. Monet. Econ. 14, 319–338.
Farhi, E., Fraiberger, S.P., Gabaix, X., Ranciere, R., Verdelhan, A., 2012. Crash risk in currency markets. NBER Working Paper No. 15062.
Frachot, A., 1996. A reexamination of the uncovered interest rate parity hypothesis. J. Int. Money Finance 15 (3), 419–437.
Frankel, J., Poonawala, J., 2010. The forward market in emerging currencies: less biased than in major currencies. J. Int. Money Finance 29 (3), 585–598.
Froot, K.A., Ramadorai, T., 2005. Currency returns, intrinsic value, and institutional-investor flows. J. Finance 60 (3), 1535–1566.
Gavazzoni, F., Sambalaibat, B., Telmer, C., 2012. Currency risk and pricing kernel volatility. Working Paper. Carnegie Mellon University.
Gourio, F., Siemer, M., Verdelhan, A., 2013. International risk cycles. J. Int. Econ. 89, 471–484.
Goyal, A., Welch, I., 2005. A comprehensive look at the empirical performance of equity premium prediction. Rev. Financ. Stud. 21 (4), 1455–1508.
Hansen, L.P., Hodrick, R.J., 1980. Forward exchange rates as optimal predictors of future spot rates: an econometric analysis. J. Polit. Econ. 88 (October (5)), 829–853.
Harvey, C.R., 1991. The world price of covariance risk. J. Finance 46 (1), 111–157.
Hodrick, R., 1987. The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets. Harwood Academic Publishers, Chur, Switzerland.
Hollifield, B., Yaron, A., 2001. The foreign exchange risk premium: Real and nominal factors. Working Paper Wharton School. University of Pennsylvania.
Hong, H., Yogo, M., 2012. What does futures market interest tell us about the macroeconomy and asset prices? J. Financ. Econ. 105 (3), 473–490.
Joslin, S., Priebsch, M., Singleton, K.J., 2010. Risk premiums in dynamic term structure models with unspanned macro risks. Working paper. Stanford University, May.
Jurek, J.W., 2009. Crash-neutral currency carry trades. Working paper.
Kandel, S., Stambaugh, R.F., 1991. Asset returns and intertemporal preferences. J. Monet. Econ. 27, 39–71.
Kilian, L., 1999. Exchange rates and monetary fundamentals: What do we learn from long-horizon regressions?. J. Appl. Econom. 14 (5), 491–510.
Lewis, K.K., 1995. Puzzles in international financial markets. In: Grossman, G., Rogoff, K. (Eds.), Handbook of International Economics, Elsevier Science, Amsterdam, pp. 1913–1971.
Ludvigson, S.C., Ng, S., 2009. Macro factors in bond risk premia. Rev. Financ. Stud. 22 (12), 5027–5067.
Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common risk factors in currency returns. Rev. Financ. Stud. 24 (November (11)), 3731–3777.
Mark, N.C., 1995. Exchange rates and fundamentals: evidence on long-horizon predictability. Am. Econ. Rev. 85 (1), 201–218.
Meese, R.A., Rogoff, K., 1983. Empirical exchange rate models of the seventies: Do they fit out of sample?. J. Int. Econ. 14, 3–24.
Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012. Carry trades and global foreign exchange rate volatility. J. Finance 67 (2), 681–718.
Politis, D.N., Romano, J.P., 1994. The stationary bootstrap. J. Am. Stat. Assoc. 89, 1303–1313.
Sarno, L., Schneider, P., Wagner, C., 2012. Properties of foreign exchange risk premiums. J. Financ. Econ. 105 (2), 279–310.
Stambaugh, R.F., 1988. The information in forward rates: implications for models of the term structure. J. Financ. Econ. 21, 41–70.
Stambaugh, R.F., 1999. Predictive regressions. J. Financ. Econ. 54 (1), 375–421.
Verdelhan, A., 2010. A habit-based explanation of the exchange rate risk premium. J. Finance 65 (1), 123–145.