

Countercyclical Currency Risk Premia*

Hanno Lustig Nikolai Roussanov Adrien Verdelhan
UCLA Anderson Wharton MIT Sloan

March 30, 2012

Abstract

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 price of risk is high. The counter-cyclical 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.

JEL: G12, G15, F31.

Keywords: Exchange Rates, Forecasting, Risk.

*Lustig: Anderson School of Management, University of California at Los Angeles, Box 951477, Los Angeles, CA 90095, and NBER; hlustig@anderson.ucla.edu; Tel: (310) 825-8018. Roussanov: Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104, and NBER; nroussan@wharton.upenn.edu; Tel: (215) 746-0004. Verdelhan: MIT Sloan, 100 Main Street, Cambridge, MA 02142, and NBER; adrienv@mit.edu. The authors thank seminar participants at many institutions and conferences for helpful comments.

The high-minus-low currency carry trade, which goes long in baskets of currencies with high interest rates and short in baskets of currencies with low interest rates, is a dollar-neutral investment strategy that has been shown to deliver high Sharpe ratios. This currency investment strategy, which has been thoroughly studied by researchers, ignores all information in the level of short-term U.S. interest rates, and uses only the ranking of foreign interest rates to build portfolios. 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.

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. 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 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 case of a bad innovation to the U.S. pricing kernel. In both cases, U.S. investors collect a positive currency risk premium because they are betting against their own intertemporal marginal rates of substitution.

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 for the U.S. 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 underlies the profitability of the dollar carry trade strategy.

The key predictor is the average forward discount, which is the difference between the U.S. short-term interest rate and the average short-term interest in all the other developed countries. The one-month ahead average forward discount explains 1%-5% of the variation in the foreign currency excess returns on a basket of developed country foreign 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 15% 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, and they earn an additional 150 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 250 basis points per annum and it leads to an annualized depreciation of the dollar by 150 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 of the 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 25% at the one-year horizon. 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 up to 190 basis points per annum, after controlling for the average forward discount. We also show the U.S. consumption growth volatility forecasts dollar returns. As in the model, these macro-economic 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 (2008), 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 OECD countries, a proxy for world-wide macroeconomic risk, and the rate of U.S.-specific component of industrial production growth. We propose a no-arbitrage model that decomposes pricing kernels and returns into country-specific and global components. A version of our model calibrated to match the dynamics and the cross-section of interest rates and exchange rates generates a large dollar carry trade risk premium, but cannot match the one observed in the data without imputing too much volatility to changes in exchange rates.

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 1 discusses the relation of our paper to existing literature. Section 2 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 3 presents a no-arbitrage model of exchange rates, which belongs to the essentially-affine class that is popular in the term-structure literature. 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 4 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 5 concludes.

1 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. 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. Asness, Moskowitz and Pedersen (2009) show that 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. 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 (2011) 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 (2011) show that the futures open market interest has strong predictive power for returns on a portfolio of currency futures. Bakshi and Panayotov (forthcoming) present evidence that long-short carry trade returns are predictable using a measure of foreign exchange volatility and commodity prices. 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

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

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 falls within 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), and Lustig, Roussanov and Verdelhan (2011). In our model the bulk of the stochastic discount factor variation is common across countries, consistently 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 procyclical effect on the short-term interest rate, such as Verdelhan (2010), Bansal and Shaliastovich (2008), Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2009), and Backus, Gavazzoni, Telmer and Zin (2010). Atkeson and Kehoe (2008) argue that this effect is important for understanding the impact of monetary policy on interest rates. We show that distinguishing country-specific variation from global variation is key for understanding risk compensation in foreign exchange markets.

2 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 of exchange rate depreciation. 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.

2.1 Preliminaries

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 home currency. 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 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 is equal to the interest rate differential: $f_t - s_t \approx i_t^* - i_t$, where i^* and i denote the foreign and domestic nominal risk-free rates over the maturity of the contract.² Hence, the log currency excess return equals the interest rate differential less the rate of depreciation: $rx_{t+1} = i_t^* - i_t - \Delta s_{t+1}$.

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 $rx_{t+k}^i = -\Delta s_{t \rightarrow t+k}^i + f_{t \rightarrow t+k}^i - s_t^i$, where $f_{t \rightarrow t+k}^i$ is the k -month forward exchange rate, and the k -month change in the log exchange rate is $\Delta s_{t \rightarrow t+k}^i = s_{t+k}^i - s_t^i$. For horizons above one month, our series consist of overlapping k -month returns computed at a monthly frequency.

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

²Akram, 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, including or excluding those observations does not have a major effect on our results.

for an investor who goes long in foreign currency is: $rx_{t+1}^l = f_t^b - s_{t+1}^a$. 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_{t+1}^a) 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: $rx_{t+1}^s = -f_t^a + s_{t+1}^b$. 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.

Data We start from daily spot and forward exchange rates in U.S. dollars. 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 at most 38 different currencies of the following countries/currencies: Australia, Austria, Belgium, Canada, Hong Kong, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, India, Indonesia, Ireland, 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.

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 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 $\overline{rx}_{t \rightarrow t+k}^j = \frac{1}{N_t^j} \sum_{i=1}^{N_t^j} rx_{t+k}^i$, 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 $\overline{\Delta s}_{t \rightarrow t+k}^j = \frac{1}{N_t^j} \sum_{i=1}^{N_t^j} \Delta s_{t \rightarrow t+k}^i$, and the average forward discount (AFD) for maturity k is $\overline{f}_{t \rightarrow t+k}^j - \overline{s}_t^j = \frac{1}{N_t^j} \sum_{i=1}^{N_t^j} f_{t \rightarrow t+k}^i - s_t^i$.

[Figure 1 about here.]

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% (3-month Treasury bill rate) to 0% because of the secular decline in (global) inflation over the sample. The AFD differences out common variation in inflation expectations across countries and picks up real interest rate differences. 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 averages. This disparity suggests that one should expect different patterns of predictability for the different baskets.

The mean AFD for the U.S. is between 68 and 100 basis points per annum. In the sample of developed countries, the unconditional average annualized dollar premium varies between 2.25% and 2.43% per annum depending on the horizon. The AFDs are highly persistent, especially at longer horizons, with monthly autocorrelations between 0.83 and 0.98. Hence, the annualized autocorrelations vary between 0.11 and 0.78. Therefore, they 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.

2.2 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 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 Figure 1, the dollar carry trade shorts the dollar just before the peak (start of NBER recessions), and goes long in the dollar after the trough (end of the NBER recession).

Figure 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 et al. (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 lever the currency positions so that all currency returns are

equally volatile as equity returns. 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 predictable changes in the dollar exchange rate. On the other hand, the HML currency carry trade delivers \$356 dollars, while the country-level strategy barely breaks even.

Table I 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 bootstrapping under the assumption that excess returns are *i.i.d.* The sample average of dollar carry returns, our estimate of the dollar premium, are 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% (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% 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 if 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 right panel of Table I 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 5%. 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 Figure 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.

[Table 1 about here.]

[Figure 2 about here.]

2.3 Predictability in Currency Markets

The dollar carry trade is highly profitable because the U.S. average forward discount forecasts basket-level exchange rate changes and returns, even in a horse race with the individual currency pairs' forward discounts.

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:

$$\overline{r\bar{x}}_{t \rightarrow t+k} = \psi_0 + \psi_{\mathbf{f}}(\overline{f}_{t \rightarrow t+k} - \overline{s}_t) + \eta_{t+k}, \quad (1)$$

$$-\overline{\Delta s}_{t \rightarrow t+k} = \zeta_0 + \zeta_{\mathbf{f}}(\overline{f}_{t \rightarrow t+k} - \overline{s}_t) + \epsilon_{t+k}. \quad (2)$$

We report t -statistics for the slope coefficients $\psi_{\mathbf{f}}$ and $\zeta_{\mathbf{f}}$ for both asymptotic and finite-sample tests. The AFD are strongly autocorrelated, albeit less so than individual countries' interest rates. We use Newey-West standard errors (NW) computed with the optimal number of lags following Andrews (1991) in order to correct for error correlation and conditional heteroscedasticity. We

also verify that our results are robust by computing Hansen-Hodrick standard errors as well as Newey-West standard errors using non-overlapping data. To save space, we do not report the additional standard errors, but they are available on demand.

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 computed bias-adjusted small sample t -statistics, generated by bootstrapping 10,000 samples of returns and forward discounts from a corresponding VAR under the null of no predictability.³

The regression Equations (1) and (2) test different hypotheses. In the regression for excess returns in Equation (1), the null states that the log expected excess currency returns are constant. In the regression for log exchange rates changes in Equation (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).

Predictability Results Table II 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.⁴ 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,

³Our 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. 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.

⁴The UIP condition states that expected changes in exchange rates are equal to interest rate differentials. With our notation, UIP implies that $\zeta_{\mathbf{f}} = -1$ and $\psi_{\mathbf{f}} = 0$. The UIP condition is flatly rejected by the data. 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.

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. Since the log excess returns are the difference between changes in spot rates at $t + 1$ and the AFD at t , these two regressions are equivalent and $\psi_f = \zeta_f + 1$. The R^2 s for the exchange rate regressions are lower, ranging from just over 1% for monthly to almost 5% for annual horizon.

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 250 basis points increase in the expected return, and it increases the expected appreciation of the foreign currency basket by 150 basis points. The estimates are very similar for all maturities, except the 12-month estimate, which is 34 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 II 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 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 II pertains to the sample of all countries. Not surprisingly, excess return predictability is very strong there as well.

[Table 2 about here.]

2.4 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 for excess returns:

$$rx_{t \rightarrow t+k}^i = \psi_0^i + \tilde{\psi}_f(\bar{f}_{t \rightarrow t+k} - \bar{s}_t) + \tilde{\psi}_f(f_t^i - s_t^i) + \tilde{\eta}_{t+k}^i,$$

using the AFD as well as the currency-specific forward discount, and a similar regression for spot exchange rate changes:

$$-\Delta s_{t \rightarrow t+k}^i = \zeta_0^i + \tilde{\zeta}_f(\bar{f}_{t \rightarrow t+k} - \bar{s}_t) + \tilde{\zeta}_f(f_t^i - s_t^i) + \tilde{\eta}_{t+k}^i,$$

where ψ_0^i and ζ_0^i are currency fixed effects, so that only the slope coefficients are constrained to be equal across currencies.

Table III 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 robustly 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 equally important as the AFD for predicting excess returns, but not for exchange rate changes.

[Table 3 about here.]

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 6-month horizons yields an average R^2 (across all developed country currencies) of 10% and 5%, compared to the average R^2 using the bilateral forward discount of 7.3% and 2.3%, respectively. Similarly, at 12-month horizons the average R^2 using the AFD are 15% and 9% for excess returns and spot rate changes, respectively, compared to the average R^2 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 to the HML currency carry trades. We turn now to a no-arbitrage model that offers an interpretation of our empirical findings.

3 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.

3.1 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 $\Delta q_{t+1}^i = m_{t+1} - m_{t+1}^i$, where 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. The log stochastic discount factor (SDF) is denoted m ; it is also known as a pricing kernel or intertemporal marginal rate of substitution (IMRS). 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:

$$E_t[r_{t+1}^{basket}] = -\frac{1}{N} \sum_i E_t[\Delta q_{t+1}^i] + \frac{1}{N} \sum_i r_t^i - r_t = \frac{1}{2} [Var_t(m_{t+1}) - \frac{1}{N} \sum_i Var_t(m_{t+1}^i)].$$

The dollar carry trade goes 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. The signal to go long or short is the average forward discount ($\frac{1}{N} \sum_i r_t^i - r_t$), which tends to increase when the volatility of the U.S. pricing kernel is high. These expected returns are driven by the U.S. business cycle. By contrast, the real expected log currency excess return on the HML currency carry trade

is given by:

$$E_t[rx_{t+1}^{basket}] = \frac{1}{2} \left[\frac{1}{N_L} \sum_{j \in L} Var_t(m_{t+1}^j) - \frac{1}{N_H} \sum_{j \in H} Var_t(m_{t+1}^j) \right],$$

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 volatility of low and high interest rate currencies increases. These expected returns are driven by global volatility in financial markets.

We use a no-arbitrage asset pricing model in the tradition of the affine models of the term structure of interest rates 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 for the cross-sectional properties of average returns on currency baskets formed from the perspective of different base currencies, which we verify in the data.

3.2 Setup

We extend the no-arbitrage model developed by Lustig et al. (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 et al. (2006), Bakshi et al. (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 et al. (2011) show that, in order to reproduce cross-sectional evidence on currency excess returns, risk prices must load differently on this common component.

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. The SDFs assume that each type of risk is priced differently. The risk prices of 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_t^w shared by all countries and a country-specific state variable z_t^i . The country-specific and world state variables follow autoregressive square root processes:

$$\begin{aligned} z_{t+1}^i &= (1 - \phi)\theta + \phi z_t^i - \sigma \sqrt{z_t^i} u_{t+1}^i, \\ z_{t+1}^w &= (1 - \phi^w)\theta^w + \phi^w z_t^w - \sigma^w \sqrt{z_t^w} u_{t+1}^w. \end{aligned}$$

Intuitively, z_t^i captures variation in the risk price due the business cycle of country i , while z_t^w 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:

$$-m_{t+1}^i = \alpha + \chi z_t^i + \sqrt{\gamma z_t^i} u_{t+1}^i + \tau z_t^w + \sqrt{\delta^i z_t^w} u_{t+1}^w + \sqrt{\kappa z_t^i} u_{t+1}^g, \quad (3)$$

where u_{t+1}^i is a country-specific SDF shock while u_{t+1}^w and u_{t+1}^g are common to all countries SDFs. All of these three innovations are *i.i.d* Gaussian, with zero mean and unit variance, independent of one another. There are two types of common shocks. The first type u_{t+1}^w is priced identically in all countries with the same exposure δ . 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_{t+1}^g , is not; it is, for example, a productivity shock that affects some economies more than others.⁵ 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 IMRS to global shocks, and therefore the price of global risk, to vary over time with that country's economic and financial conditions (z_t^i).

To be parsimonious, we limit the heterogeneity in the SDF parameters to the different loadings, denoted δ^i , on the world shock u_{t+1}^w ; all the other parameters are identical for all countries. The only qualitative departure of our model from Lustig et al. (2011) is the separation of the world

⁵Gourio, Siemer and Verdelhan (forthcoming) 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.

shock into two independent shocks, u_{t+1}^w and u_{t+1}^g , that is dictated by the correlation properties of different carry trade strategies.

Currency Risk Premia for Individual Currencies In our model, the forward discount between currency i and the U.S. is equal to: $r_t^i - r_t = (\chi - \frac{1}{2}(\gamma + \kappa)) (z_t^i - z_t) - \frac{1}{2} (\delta^i - \delta) z_t^w$. The three same variables, z_t , z_t^i , and z_t^w , 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: $E_t[r x_{t+1}^i] = \frac{1}{2}[(\gamma + \kappa) (z_t - z_t^i) + (\delta - \delta^i) z_t^w]$.

Accounting for the variation in expected currency excess returns across different currencies requires variation in the SDFs' exposures to the common innovation. Lustig et al. (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 to not to depreciate on average), whereas the *transitory* heterogeneity in loadings, captured by the κz_t^i term in Equation (3), is necessary to match the variation in *conditional* expected returns (why currencies with currently high interest rate 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 et al. (2010)] and account for between a third and a half of the carry trade profits, as reported in Lustig et al. (2011).

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 (\bar{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:

$$E_t[\bar{r x}_{t+1}] = \frac{1}{2}(\gamma + \kappa) (z_t - \bar{z}_t) + \frac{1}{2} (\delta - \bar{\delta}) z_t^w. \quad (4)$$

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, \bar{z} is approximately constant (and exactly in the limit $N \rightarrow \infty$ 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_t^w). The real expected excess return of this basket depends only on z and z^w . These are the same variables that drive the AFD: $\bar{r}_t - r_t = (\chi - \frac{1}{2}(\gamma + \kappa))(\bar{z}_t - z_t) + \frac{1}{2}(\delta - \bar{\delta})z_t^w$. Thus, the AFD contains information about average excess returns on a basket of currencies.

HML Unconditional Carry Trades The model has strong predictions on HML currency carry trades; Lustig et al. (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 et al. (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:

$$hml_{t+1}^{unc} - E[hml_{t+1}^{unc}] = \left(\frac{1}{N_L} \sum_{i \in L} \sqrt{\delta^i z_t^w} - \frac{1}{N_H} \sum_{i \in H} \sqrt{\delta^i z_t^w} \right) u_{t+1}^w.$$

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

3.3 The Dollar Carry Risk Premium in the Model

We turn now to the dollar carry strategy. In order to build intuition on the dollar carry risk, it is useful to first consider a special case: assume that the global state variable (z^w in the model, e.g. a measure of global volatility in financial markets) has the same impact on the U.S. pricing kernel as on the average foreign pricing kernel. Recall that exchange rates correspond to differences in log pricing kernels. Then 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. The dollar, however, responds to U.S.-specific shocks (u in the model) and to global shocks (u^g) that are priced differently in each country. In this special case, a short position in the dollar is risky because the dollar appreciates following negative U.S.-specific shocks and negative global shocks. The compensation for bearing this risk depends only on the U.S. economic conditions, captured in the model by the U.S.-specific state variable z . U.S. investors expect to be compensated more for bearing this risk during recessions when U.S. interest rates are low and the average forward discount is high. 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.

Second, assume now that the U.S. pricing kernel does not always exhibit exactly the same exposure to common shocks as the average developed country. The dollar now responds to a second set of global shocks (u^w). Again, a short position in the dollar is risky because the dollar appreciates following negative global shocks. The compensation for bearing those shocks evolve according to the global state variable z^w . We refer to the corresponding risk premium as the global component of the dollar premium. By implementing the dollar carry trade, the U.S. investor is pocketing both components of the dollar premium when the price of risk is high.

With this intuition in mind, we now describe more precisely expected excess returns and predictability regressions in the model.

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 (e.g., the U.S.): $\bar{\delta} = \delta$. 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:

$$E_t[\bar{r}\bar{x}_{t+1}] = \frac{1}{2}(\gamma + \kappa)(z_t - \bar{z}_t). \quad (5)$$

Hence, the currency risk premium on this basket is the *dollar risk premium*. It compensates U.S. investors proportionally to their exposures to the local risk governed by γ and 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). Similarly, given the average exposure assumption, the AFD only depends on the U.S. factor z_t :

$$\bar{r}_t - r_t = \left(\chi - \frac{1}{2}(\gamma + \kappa) \right) (\bar{z}_t - z_t). \quad (6)$$

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 $\psi_{\mathbf{f}} = -\frac{1}{2}(\gamma + \kappa) / (\chi - \frac{1}{2}(\gamma + \kappa))$. Correspondingly, the slope coefficient in a regression of average real exchange rate changes on the real forward discount is $\zeta_{\mathbf{f}} = -\chi / (\chi - \frac{1}{2}(\gamma + \kappa))$. 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.⁶

Under the empirically relevant case $\chi < \frac{1}{2}(\gamma + \kappa)$ and the assumption that $\bar{\delta} = \delta$, 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:

$$E_t[\text{Dollar Carry}_{t+1}] = \frac{1}{2}(\gamma + \kappa) (z_t - \bar{z}_t) \text{sign}(z_t - \bar{z}_t) > 0. \quad (7)$$

The dollar premium is driven entirely by the domestic state variable z_t . This state variable influ-

⁶If $\chi = 0$, 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., $\zeta_{\mathbf{f}} = -1$

ences 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_{t+1}^g).

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, given our assumptions, only on the global state variable z_t^w . 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 $\bar{\delta} \approx \delta$. This prediction is confirmed in the data: there is no evidence of predictability of HML carry trade returns using the AFD.

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

$$Dollar\ Carry_{t+1} - E_t[Dollar\ Carry_{t+1}] = \left[\begin{array}{c} \sqrt{\gamma z_t} u_{t+1} \\ + \left(\sqrt{\delta z_t^w} - \frac{1}{N} \sum_i \sqrt{\delta^i z_t^w} \right) u_{t+1}^w \\ + \left(\sqrt{\kappa z_t} - \frac{1}{N} \sum_i \sqrt{\kappa z_t^i} \right) u_{t+1}^g \end{array} \right] sign(\bar{r}_t - r_t).$$

Therefore, if the domestic investor's SDF is more exposed to the global risk than the average ($\delta > \bar{\delta}$) 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 ($\delta < \bar{\delta}$). In either case, the global component of the dollar carry trade is driven by both of the global shocks, u^w and u^g .

Interpreting Shocks Is there a potential economic interpretation of these shocks? Lustig et al. (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. Interestingly, this variable is 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 OECD countries (the correlation is 0.19 and is statistically significant), as is 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 4.

Predictability and Heterogeneity When the home country exposure to the global shock u^w differs from that of the average foreign country ($\delta \neq \bar{\delta}$), 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, $\delta = \bar{\delta}$, 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

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

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:

$$\zeta_f = -\frac{1}{2}(\gamma + \kappa) / \left(\chi - \frac{1}{2}(\gamma + \kappa)\right).$$

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

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.

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:

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

To derive this result, we assume that the dispersion in δ^i is sufficiently small so that $\frac{1}{N} \sum_i \sqrt{\delta^i} \approx \sqrt{\bar{\delta}} = \sqrt{\delta}$. 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 trade returns have a positive correlation with the unconditional HML carry trade, and hence a higher correlation with the conditional one 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.⁷ In Section 3.6, we trace out this U-shaped relation between the correlation and the average forward discount, determined in the model by δ^i .

3.4 Calibration

We calibrate the no-arbitrage model in three steps. We first target a set of real moments that pin down the domestic pricing kernel under the assumption that the home country (U.S.) has the same δ as the average country in a basket used to compute the AFD and the dollar risk premium. Then we calibrate the process for inflation and verify that the implied nominal moments are consistent with those observed in the data. Finally, we introduce heterogeneity in the model by calibrating the range of δ s.

Real Moments There are 12 parameters in the real part of the model: 6 parameters govern the dynamics of the real stochastic discount factors (α , χ , τ , γ , κ , and δ) and 6 parameters describe the evolution of the country-specific and global factors (ϕ , θ , and σ for z and ϕ^w , θ^w , and σ^w for z^w). We choose these parameters to match the following 12 moments in the data: the mean, standard deviation, and autocorrelation of the U.S. real short-term interest rates, as well as the standard deviation and the autocorrelation of the average forward discount (its mean is set to zero by the average δ assumption), the standard deviation of changes in real exchange rates, the slope coefficients for the regression of average exchange rate changes on the average forward discount, the R^2 of this regression, which is closely related to the Sharpe ratio on the dollar carry strategy, the average excess return on the developed country currency basket, the cross-country correlation

⁷The 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:

$$\text{Corr}_t(\text{Dollar Carry}_{t+1}, \text{hml}_{t+1}) = \frac{\kappa \left| \sqrt{z_t} - \frac{1}{N} \sum_i \sqrt{z_t^i} \right| \left(\frac{1}{N_L} \sum_{i \in L} \sqrt{z_t^i} - \frac{1}{N_H} \sum_{i \in H} \sqrt{z_t^i} \right)}{\sqrt{\text{Var}_t(\text{Dollar Carry}_{t+1})} \sqrt{\text{Var}_t(\text{hml}_{t+1})}},$$

where $\text{Var}_t(\text{Dollar Carry}_{t+1}) = \gamma z_t + \left| \sqrt{\kappa z_t} - \frac{1}{N} \sum_i \sqrt{\kappa z_t^i} \right|^2$, and $\text{Var}_t(\text{hml}_{t+1}) = \left(\frac{1}{N_L} \sum_{i \in L} \sqrt{\delta^i z_t^w} - \frac{1}{N_H} \sum_{i \in H} \sqrt{\delta^i z_t^w} \right)^2 + \left(\frac{1}{N_L} \sum_{i \in L} \sqrt{\kappa z_t^i} - \frac{1}{N_H} \sum_{i \in H} \sqrt{\kappa z_t^i} \right)^2$.

of real interest rates, and two Feller parameters (equal to $2(1 - \phi)\theta/\sigma^2$ and $2(1 - \phi^w)\theta^w/\sigma^{w2}$). These 12 moments, as well as the targets in the data that we match, are listed in Panel A of Table IV. The table reports the annualized versions of means and standard deviations (the latter scaled by $\sqrt{12}$) and monthly autocorrelations. The model analogues of the empirical moments are closed form approximations (whenever available, otherwise simulated data are used).

Inflation The nominal pricing kernel is the real pricing kernel minus the rate of inflation: $m_{t+1}^{i,\$} = m_{t+1}^i - \pi_{t+1}^i$. To keep the analysis simple, we assume that inflation innovations are not priced. Hence, the expected nominal excess returns in levels on the individual currencies and currency portfolios are identical to the expected real excess returns we have derived, but in logs they are slightly different, because of Jensen's inequality. However, these differences are of second order.

We simply assume that the same factors driving the real pricing kernel also drive expected inflation. Thus, country i 's inflation process is given by:

$$\pi_{t+1}^i = \pi_0 + \eta z_t^i + \eta^w z_t^w + \sigma_\pi \epsilon_{t+1}^i,$$

where the inflation innovations ϵ_{t+1}^i are *i.i.d.* Gaussian. The nominal risk-free interest rate (in logarithms) is given by $r_t^{i,\$} = \pi_0 + \alpha + (\chi + \eta - \frac{1}{2}(\gamma + \kappa)) z_t^i + (\chi + \eta^w - \frac{1}{2}\delta^i) z_t^w - \frac{1}{2}\sigma_\pi^2$.

We set $\eta = 0$, so expected inflation does not respond to the country-specific factor. This assumption is reasonable as we consider only the developed countries currencies in this calibration. As a result, there is no difference in UIP slope coefficients between the nominal and the real model. We obtain the 3 inflation parameters $(\eta^w, \sigma^\pi, \pi_0)$ by targeting the first two moments of inflation, as well as the fraction of inflation that is explained by the common component.

In Panel B of Table IV, we list the expression for the variance of inflation and the fraction explained by the common component. We target an annualized standard deviation for inflation of 1.09% and an average inflation rate of 2.90%; 28% of inflation is accounted by the common component. Finally, for completeness, Panel C also shows the implied moments of nominal interest rates and exchange rates in this symmetric version of the model, including the yield on a long-term

nominal bond that was not targeted in the calibration. Panel D reports the implied moments of the domestic stochastic discount factor such as the the average conditional maximum Sharpe ratio (e.g., the average standard deviation of the log SDF), the average correlation between the domestic and foreign SDFs, and the standard deviation of the conditional volatility of the log SDF.

[Table 4 about here.]

Targets The specific values we consider are listed in Table IV. We target a basket exchange rate predictive regression slope coefficient of 1.5 (which implies a coefficient of 2.5 for the regression of excess return on the forward discount), and an R^2 in the UIP regression of 1%, a standard deviation of the AFD of 0.74% per annum and a monthly autocorrelation of 0.83, an average real interest rate of 1.72% per annum, an annualized standard deviation of the real interest rate of 0.57% per annum, and an autocorrelation (in monthly data) of 0.92. The annual standard deviation of real exchange rate changes is 10%. The annual dollar risk premium is 0.5% per annum, which is equal to Jensen’s inequality term (since unconditional average log return is equal to zero due to the average delta assumption). The average pairwise correlation of real interest rates is 0.3. 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).

Heterogeneity Finally, we introduce a single source of heterogeneity in the countries’ exposure to the global shocks. The parameters (δ^i) are assumed to be distributed uniformly on the interval $[\delta_h, \delta_l]$. The mean forward discount on currency i is given by $E(r_t^i - r_t) = -\frac{1}{2}(\delta^i - \delta)\theta^w$, so we can use the range of unconditional mean average forward discounts to calibrate the range of δ s: $\delta^i = \delta - \frac{2E(r_t^i - r_t)}{\theta^w}$. We use $E(r_t^h - r_t) = 5\%$ and $E(r_t^l - r_t) = -5\%$ (annualized), which is broadly consistent with our sample. All the model parameters are reported in Table V.

Match We do not match all of the moments exactly due to the fact that some of the moments are nonlinear in the parameters. In particular, the model overstates the volatility of the nominal risk-free rate and understates the volatility of the AFD, since in the model the latter must be smaller than the former, which does not appear to be the case in the data. Most of the other moments are matched very closely. The maximum Sharpe ratio is fairly high but still below unity on average at 0.85, and fairly volatile, with an annualized standard deviation of 5.69%. As a consequence, in our calibration the pricing kernels are also highly correlated across countries at 0.99 on average as exchange rates are fairly smooth relative to the SDFs (cf. Brandt et al., 2006). The model produces reasonable magnitudes of term premia on nominal bonds even though bond yields were not targeted in the calibration (n -year nominal bond yields $y^{s,n}$ are computed via a version of standard affine formulas).

[Table 5 about here.]

3.5 Simulated Dollar and HML Carry Trades

Table VI reports the moments of the dollar and HML carry trade returns generated by the calibrated version of our model. The unlevered average dollar carry trade return is only 2.08% in the model, because the Sharpe ratio is half of that in the data (0.24 vs. 0.56). Levering up the return as we do in the data to match the stock market volatility of 15.5% produces an average dollar carry excess return of 3.75% 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 more profitable in the model than in the data, with a Sharpe ratio of 0.77. The correlation between the dollar and HML currency carry trade is 0.15, which is higher than in the data but still close to zero. Overall, the model can match the HML carry trade returns, but it has more difficulty matching the dollar carry trade returns.

[Table 6 about here.]

Return Predictability The model generates high 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 VII 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 $T = 33600$ that is meant to closely approximate population values. Panel B presents results using a large number ($N = 1000$) 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 II.

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 2.5 for returns and 1.5 for exchange rate changes at the one-month horizon, declining only slightly at the twelve-month horizon. The “population” R^2 s appear 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 somewhat less than 1% of the variation in one-month excess returns, and over 10% of variation in twelve-month returns, almost as much as in the data; similarly, it explains up to 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 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 1.5 for returns and of 1 for exchange rate changes. This fact suggests that even if our model is true the evidence of predictability of excess returns and, especially, exchange rate changes may be hard to ascertain empirically. Even though exchange rates are predictable in the model, since $\chi > 0$, the statistical evidence for exchange rate predictability in small samples is weak, as it is in the data. From the perspective of the no-arbitrage model it is no surprise that

exchange rates are so close to the random walk that their changes are almost impossible to predict.

[Table 7 about here.]

3.6 Base Carry vs. HML Carry: Heterogeneity

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 a basket of all currencies when the AFD 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.

Figure 3 compares the correlations simulated from the calibrated model (solid line) to those observed in the data for the so called G10 currencies, i.e. the 10 most traded currencies: U.S. Dollar, U.K. Pound Sterling, Euro, Australian, Canadian, and New Zealand Dollars, Japanese Yen, Swiss Franc, Norwegian Krone and Swedish Krona. Consistently 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., Euro) 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 δ s. Note that this result is not mechanical, since the HML strategy formed from a perspective of a given base currency does not include the base currency itself, where as 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.

[Figure 3 about here.]

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 are much lower than average as they typically have high 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.

4 Countercyclical Currency Risk Premia

Our empirical results imply that expected excess returns on currency portfolios vary over time. The no-arbitrage model in Section 3 suggests that this variation is driven by time-variation in the U.S.-specific prices of domestic and global risk. Voluminous 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 counter-cyclical 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 counter-cyclical 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) and 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. Consistently 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.

4.1 Cyclical Properties of the Average Forward Discount

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

$$\widehat{E}_t r x_{t+1} = \psi_0 + \psi_{\mathbf{f}}(\bar{f}_{t \rightarrow t+k} - \bar{s}_t).$$

Therefore, expected excess returns on currency baskets inherit the cyclical properties of the AFDs. To assess the cyclicity 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 1-year Treasury zero coupon yields, $Term$), the default spread (the difference between the BBB and AAA bond yields, Def), and the CBOE VIX index of S&P 500 index-option implied volatility.⁸ 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 8 about here.]

Table VIII reports the contemporaneous correlations of the AFDs (across horizons) with these macroeconomic and financial variables. Developed countries are in the first panel and emerging countries in the second. As expected, the AFD for the developed country basket (and, therefore, forecasted excess returns) is counter-cyclical: 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 counter-cyclical. Since excess returns load positively on the AFD, they are also counter-cyclical, 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 counter-cyclical. The AFD, however, is not correlated to the VIX index (point estimates are negative but not significant), pointing again to the difference between dollar carry and HML carry trades.

⁸Industrial production data are from the IMF 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, the corporate and Treasury bond yields are from Datastream.

4.2 Macro Factors and Currency Return Predictability

So far we have focused on the predictive power of the AFD, but the counter-cyclical 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 rate, 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), and Joslin et al. (2010)). In our context, if we are looking to identify those components of the domestic state variable \mathbf{z}_t that are not captured by interest rate differentials, we expect a U.S.-specific macroeconomic variable to have forecasting power for currency excess returns, as well as spot exchange rate changes.

4.2.1 Industrial Production Growth

We use $\overline{r\mathbf{x}}_{t \rightarrow t+k}$ to denote the k -month ahead excess return between time t and $t+k$, as well as the corresponding regression for exchange rate changes. Table IX reports two sets of regression results:

$$\begin{aligned}\overline{r\mathbf{x}}_{t \rightarrow t+k} &= \psi_0 + \psi_{IP} \Delta \log IP_t + \psi_{\mathbf{f}} (\overline{f}_{t \rightarrow t+k} - \overline{s}_t) + \eta_{t+k}, \\ -\overline{\Delta s}_{t \rightarrow t+k} &= \zeta_0 + \zeta_{IP} \Delta \log IP_t + \zeta_{\mathbf{f}} (\overline{f}_{t \rightarrow t+k} - \overline{s}_t) + \varepsilon_{t+k}.\end{aligned}$$

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 3 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 6 and 12 months. The R^2 for average excess returns at 12-month horizon are between 24% and 32% across different baskets, and between 15% and 32% for average exchange rate changes.

Since we are controlling for the average forward 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 counter-cyclical 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 to 100 basis points per annum at the monthly horizon and by as much as 90 to 135 basis points at the annual horizon, all coming from the expected appreciation of the foreign currencies against the dollar. Since the AFD itself is counter-cyclical, the total effect is even greater, implying an increase in expected returns of up to 120 basis points for annual data.

[Table 9 about here.]

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 0.5.⁹ 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 for the projection of these 12-month changes on the average foreign IP indices:

$$\begin{aligned}\Delta \log IP_t &= \alpha + \beta \overline{\Delta \log IP_t} + IP_{res,t}, \\ \overline{r\bar{x}_{t \rightarrow t+k}} &= \psi_0 + \psi_{IP_{res}} IP_{res,t} + \psi_{\mathbf{f}}(\bar{f}_{t \rightarrow t+k} - \bar{s}_t) + \eta_{t+k},\end{aligned}$$

⁹The G7 countries are Canada, France, Germany, Italy, Japan, and the U.K., as well as the U.S.

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. 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 to 190 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 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.

4.2.2 Consumption Volatility

In consumption-based asset pricing models, time-varying risk premia can arise due to heteroscedasticity of aggregate consumption growth [e.g. as in Kandel and Stambaugh (1991) and 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 et al. (2010)]. 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 $\sigma_t(\pi_{t+1})$:

$$\begin{aligned}\overline{r\bar{x}}_{t \rightarrow t+k} &= \psi_0 + \psi_{\sigma_c} \sigma_t(\Delta c_{t+1}) + \psi_{\sigma_p} \sigma_t(\Delta \pi_{t+1}) + \eta_{t+k}, \\ -\overline{\Delta \bar{s}}_{t \rightarrow t+k} &= \zeta_0 + \zeta_{\sigma_c} \sigma_t(\Delta c_{t+1}) + \zeta_{\sigma_p} \sigma_t(\Delta \pi_{t+1}) + \varepsilon_{t+k}.\end{aligned}$$

Results of these regressions are reported in Table X. 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.

[Table 10 about here.]

4.2.3 Out-of-Sample Evidence

We conclude by looking at the out-of-sample predictability. We check whether our predictors outperform the random walk in forecasting exchange rates, as well as excess returns out-of-sample. For brevity, we simply summarize the main results.

On the one hand, for excess returns, the hypothesis of equal prediction is rejected for all specifications at most horizons using both the ratio of mean squared errors and the Clark and McCracken (2001) ENC test, but not for the Diebold and Mariano (1995) MSE_t tests. This suggests that the MSE_t test has low power in this setting. The AFD and IP growth predict currency excess returns out-of-sample better than a simple random walk.

On the other hand, for exchange rate changes, the evidence is more mixed, which is not surprising given the results in the literature and the weak evidence for exchange rate predictability inside our model. At the one-month horizon, the standard result of Meese and Rogoff (1983) stands. At longer horizons, however, changes in IP predict changes in exchange rates much better than a simple constant. While the random walk is hard to beat as the best predictor of these changes in

exchange rates, our results indicate that using business-cycle variables such as IP allows for some improvement in the forecasting power. In our calibrated model, the exchange rate is close to a random walk and the statistical evidence for exchange rate predictability is weak.

Overall, we find that the expected returns on shorting the U.S. dollar are counter-cyclical: 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 counter-cyclical. 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 counter-cyclical currency risk premia.

5 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 counter-cyclical.

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 (e.g., time-varying U.S. exposure to world aggregate consumption growth) 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, Tobias, Erkki Etula, and Hyun Song Shin**, “Risk Appetite and Exchange Rates,” May 2010. Federal Reserve Bank of New York Staff Reports 361.
- Akram, Q. Farooq, Dagfinn Rime, and Lucio Sarno**, “Arbitrage in the Foreign Exchange Market: Turning on the Microscope,” *Journal of International Economics*, 2008, 76 (2), 237–253.
- Andrews, Donald W.K.**, “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica*, 1991, 59 (1), 817–858.
- Ang, Andrew and Geert Bekaert**, “Stock Return Predictability: Is it There?,” *Review of Financial Studies*, 2007, 20, 651–707.
- and **Joseph S. Chen**, “Yield Curve Predictors of Foreign Exchange Returns,” 2010. Working Paper Columbia University.
- Asness, Cliff S., Tobias J. Moskowitz, and Lasse H. Pedersen**, “Value and Momentum Everywhere,” *NBER Working Paper*, 2009.
- Atkeson, Andrew and Patrick Kehoe**, “On the Need for a New Approach to Analyzing Monetary Policy,” in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., *NBER Macroeconomics Annual 2008*, University of Chicago Press, 2008.
- Backus, David K., Federico Gavazzoni, Christopher Telmer, and Stanley E. Zin**, “Monetary Policy and the Uncovered Interest Parity Puzzle,” 2010. NBER Working Paper.
- Backus, David, Silverio Foresi, and Chris Telmer**, “Affine Models of Currency Pricing: Accounting for the Forward Premium Anomaly,” *Journal of Finance*, 2001, 56, 279–304.
- Bakshi, Gurdip and George Panayotov**, “Predictability of Currency Carry Trades and Asset Pricing Implications,” *Journal of Financial Economics*, forthcoming.
- , — , and **George Skoulakis**, “Improving the predictability of real economic activity and asset returns with forward variances inferred from option portfolios,” *Journal of Financial Economics*, 2011, 100, 475–495.
- , **Peter Carr, and Liuren Wu**, “Stochastic Risk Premiums, Stochastic Skewness in Currency Options, and Stochastic Discount Factors in International Economies,” *Journal of Financial Economics*, January 2008, 87 (1), 132–156.

- Bansal, Ravi and Amir Yaron**, “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 2004, 59 (4), 1481 – 1509.
- and **Ivan Shaliastovich**, “Long-Run Risks Explanation of Forward Premium Puzzle,” April 2008. Working Paper Duke University.
- and **Magnus Dahlquist**, “The Forward Premium Puzzle: Different Tales from Developed and Emerging Economies,” *Journal of International Economics*, 2000, 51, 115–144.
- Bekaert, Geert, Robert J. Hodrick, and David A. Marshall**, “On Biases in Tests of the Expectations Hypothesis of the Term Structure of Interest Rates,” *Journal of Financial Economics*, 1997, 44(3), 309–48.
- Brandt, Michael W., John H. Cochrane, and Pedro Santa-Clara**, “International Risk-Sharing is Better Than You Think (or Exchange Rates are Much Too Smooth),” *Journal of Monetary Economics*, 2006, 53, 671 – 698.
- Brennan, Michael J. and Yihong Xia**, “International Capital Markets and Foreign Exchange Risk,” *Review of Financial Studies*, 2006, 19 (3), 753–795.
- Campbell, John Y and Yasushi Hamao**, “Predictable Stock Returns in the United States and Japan: A Study of Long-Term Capital Market Integration,” *Journal of Finance*, March 1992, 47 (1), 43–69.
- Campbell, John Y., Karine Serfaty De Medeiros, and Luis M. Viceira**, “Global Currency Hedging,” *Journal of Finance*, 02 2010, 65 (1), 87–121.
- Campbell, J.Y. and J. H. Cochrane**, “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 1999, 107, 205–251.
- Chen, Yu-Chin and Kwok Ping Tsang**, “What Does the Yield Curve Tell Us About Exchange Rate Predictability?,” 2011. Working Paper, University of Washington.
- Clark, Todd E. and Michael W. McCracken**, “Tests of Equal Forecast Accuracy and Encompassing for Nested Models,” *Journal of Econometrics*, 2001, 105, 85–110.
- Cochrane, John H. and Monika Piazzesi**, “Bond Risk Premia,” *The American Economic Review*, March 2005, 95 (1), 138–160.
- Colacito, Riccardo**, “Six Anomalies Looking For a Model. A Consumption-Based Explanation of International Finance Puzzles,” 2008. Working Paper.

- and **Mariano M. Croce**, “Risks For The Long Run And The Real Exchange Rate,” *Journal of Political Economy*, 2011, 119 (1), 153–182.
- Cooper, Ilan and Richard Priestley**, “Time-Varying Risk Premiums and the Output Gap,” *Review of Financial Studies*, 2009, 22 (7), 2801–2833.
- Diebold, F.X. and R.S. Mariano**, “Comparing Predictive Accuracy,” *Journal of Business and Economic Statistics*, 1995, 13, 253–263.
- Duffee, Gregory R.**, “Information in (and not in) the Term Structure,” *Review of Financial Studies*, forthcoming, 2011.
- Fama, Eugene F.**, “Forward and Spot Exchange Rates,” *Journal of Monetary Economics*, 1984, 14, 319–338.
- Farhi, Emmanuel, Samuel P. Fraiburger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan**, “Crash Risk in Currency Markets,” 2009. Working Paper.
- Frachot, Antoine**, “A Reexamination of the Uncovered Interest Rate Parity Hypothesis,” *Journal of International Money and Finance*, 1996, 15 (3), 419–437.
- Frankel, Jeffrey and Jumana Poonawala**, “The Forward Market in Emerging Currencies: Less Biased than in Major Currencies,” *Journal of International Money and Finance*, 2010, 29 (3), 585–598.
- Gourio, Francois, Michael Siemer, and Adrien Verdelhan**, “International Risk Cycles,” *Journal of International Economics*, forthcoming.
- Goyal, Amit and Ivo Welch**, “A Comprehensive Look at The Empirical Performance of Equity Premium Prediction,” *Review of Financial Studies*, 2005, 21 (4), 1455–1508.
- Hansen, Lars Peter and Robert J. Hodrick**, “Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis,” *Journal of Political Economy*, October 1980, 88 (5), 829–853.
- Harvey, Campbell R.**, “The World Price of Covariance Risk,” *Journal of Finance*, 1991, (1), 111–157.
- Hodrick, Robert**, *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*, Chur, Switzerland: Harwood Academic Publishers, 1987.
- Hong, Harrison and Motohiro Yogo**, “What Does Futures Market Interest Tell Us about the Macroeconomy and Asset Prices?,” *Journal of Financial Economics*, forthcoming, 2011.

- Joslin, Scott, Marcel Pribsch, and Kenneth J. Singleton**, “Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks,” May 2010. Stanford University.
- Kandel, Shmuel and Robert F. Stambaugh**, “Asset Returns and Intertemporal Preferences,” *Journal of Monetary Economics*, February 1991, *27* (1), 39–71.
- Kilian, Lutz**, “Exchange Rates and Monetary Fundamentals: What Do We Learn from Long-Horizon Regressions?,” *Journal of Applied Econometrics*, 1999, *14* (5), 491–510.
- Lewis, Karen K.**, “Puzzles in International Financial Markets,” in G. Grossman and K. Rogoff, eds., *Handbook of International Economics*, Amsterdam: Elsevier Science B.V., 1995, pp. 1913–1971.
- Ludvigson, Sydney C. and Serena Ng**, “Macro Factors in Bond Risk Premia,” *Review of Financial Studies*, 2009, *22*(12), 5027–5067.
- Lustig, Hanno, Nick Roussanov, and Adrien Verdelhan**, “Common Risk Factors in Currency Returns,” 2008. Working Paper NBER No 14082.
- , **Nikolai Roussanov, and Adrien Verdelhan**, “Common Risk Factors in Currency Returns,” *Review of Financial Studies*, November 2011, *24* (11), 3731–3777.
- Mark, Nelson C.**, “Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability,” *American Economic Review*, 1995, *85* (1), 201–218.
- Meese, Richard A. and Kenneth Rogoff**, “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?,” *Journal of International Economics*, 1983, *14*, 3–24.
- Newey, Whitney K. and Kenneth D. West**, “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 1987, *55* (3), 703–708.
- Stambaugh, Robert F.**, “The information in forward rates : Implications for models of the term structure,” *Journal of Financial Economics*, May 1988, *21* (1), 41–70.
- , “Predictive Regressions,” *Journal of Financial Economics*, 1999, *54* (1), 375–421.
- Verdelhan, Adrien**, “A Habit-Based Explanation of the Exchange Rate Risk Premium,” *Journal of Finance*, 2010, *65* (1), 123–145.

Table I: Currency Carry Trades and Equity Market Excess Returns

	Raw Returns				Scaled Returns			
	USD	FX	HML	Equity	USD	FX	HML	Equity
Panel A: Developed Countries								
Mean	5.60	0.43	3.00	6.26	10.18	0.88	4.77	6.26
Std. Error	[1.64]	[1.52]	[1.82]	[3.02]	[3.16]	[3.10]	[2.91]	[3.06]
Std. Dev.	8.53	7.67	9.73	15.49	15.49	15.49	15.49	15.49
Sharpe Ratio	0.66	0.06	0.31	0.40	0.66	0.06	0.31	0.40
$Corr(USD, .)$		0.32	-0.03	0.01				
Panel B: All Countries								
Mean	3.80	0.57	4.41	6.26	8.43	1.25	7.58	6.26
Std. Error	[1.38]	[1.40]	[1.70]	[3.13]	[2.96]	[3.13]	[3.03]	[2.88]
Std. Dev.	6.99	7.12	9.02	15.49	15.49	15.49	15.49	15.49
Sharpe Ratio	0.54	0.08	0.49	0.40	0.54	0.08	0.49	0.40
$Corr(USD, .)$		0.31	0.02	-0.00				

Notes: 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 et al. (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 (available on WRDS). 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 on the means (in brackets). These standard errors are obtained by bootstrapping under the assumption that excess returns are *i.i.d.* The sample period is 11/1983–6/2010.

Table II: Forecasting Currency Excess Returns and Exchange Rates with the Average Forward Discount

<i>Horizon</i>	Developed Countries				Emerging Countries				All Countries			
	Excess Returns		Exchange Rates		Excess Returns		Exchange Rates		Excess Returns		Exchange Rates	
	ψ_f	R^2	ζ_f	R^2	ψ_f	R^2	ζ_f	R^2	ψ_f	R^2	ζ_f	R^2
1	2.45	2.91	1.45	1.04	2.06	2.21	2.04	2.24	2.19	2.93	1.50	1.43
<i>NW</i>	[2.44]		[1.44]		[2.09]		[1.93]		[2.40]		[1.65]	
<i>VAR</i>	[2.91]		[1.78]		[2.40]		[2.35]		[2.84]		[2.04]	
2	2.49	5.00	1.50	1.86	2.09	3.96	2.14	4.05	2.25	5.08	1.57	2.57
<i>NW</i>	[2.34]		[1.40]		[1.71]		[1.58]		[2.27]		[1.57]	
<i>VAR</i>	[2.39]		[1.51]		[2.02]		[2.17]		[2.42]		[1.82]	
3	2.46	6.52	1.46	2.40	2.04	4.94	2.08	4.99	2.21	6.49	1.53	3.28
<i>NW</i>	[2.22]		[1.32]		[1.65]		[1.52]		[2.14]		[1.48]	
<i>VAR</i>	[2.38]		[1.43]		[1.90]		[1.96]		[2.16]		[1.64]	
6	2.45	10.23	1.45	3.84	2.02	6.96	2.10	7.28	2.19	9.95	1.53	5.20
<i>NW</i>	[2.14]		[1.27]		[1.64]		[1.53]		[2.06]		[1.43]	
<i>VAR</i>	[2.55]		[1.44]		[1.97]		[2.01]		[2.40]		[1.71]	
12	2.12	13.14	1.12	4.08	2.27	12.94	2.34	13.52	1.90	12.37	1.24	5.91
<i>NW</i>	[2.00]		[1.06]		[1.81]		[1.68]		[1.95]		[1.26]	
<i>VAR</i>	[2.05]		[1.14]		[2.58]		[2.98]		[2.12]		[1.42]	

Notes: 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 twelve months. For each basket we report the R^2 , and the slope coefficient ψ_f in the time-series regression of the log currency excess return on the average log forward discount, and similarly the slope coefficient ζ_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. *NW* denotes Newey and West (1987) standard errors computed with the optimal number of lags following Andrews (1991). 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.

Table III: Predictability Using Bilateral and Average Forward Discounts: Panel Regressions

<i>Horizon</i>	Developed Countries				Emerging Countries				All countries			
	Excess Returns		Exchange Rates		Excess Returns		Exchange Rates		Excess Returns		Exchange Rates	
	$\tilde{\psi}_{\mathbf{f}}$	$\tilde{\psi}_f$	$\tilde{\zeta}_{\mathbf{f}}$	$\tilde{\zeta}_f$	$\tilde{\psi}_{\mathbf{f}}$	$\tilde{\psi}_f$	$\tilde{\zeta}_{\mathbf{f}}$	$\tilde{\zeta}_f$	$\tilde{\psi}_{\mathbf{f}}$	$\tilde{\psi}_f$	$\tilde{\zeta}_{\mathbf{f}}$	$\tilde{\zeta}_f$
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
<i>Robust</i>	[2.13]	[1.52]	[2.13]	[-0.99]	[1.91]	[2.30]	[1.91]	[0.25]	[2.02]	[2.47]	[2.02]	[-0.12]
<i>NW</i>	[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	2.10	0.51	2.10	-0.49	1.35	1.19	1.35	0.19	1.52	1.01	1.52	0.01
<i>Robust</i>	[2.74]	[1.15]	[2.74]	[-1.12]	[1.81]	[2.15]	[1.81]	[0.34]	[2.22]	[2.26]	[2.22]	[0.01]
<i>NW</i>	[2.10]	[0.51]	[2.10]	[-0.49]	[1.35]	[1.19]	[1.35]	[0.19]	[1.52]	[1.01]	[1.52]	[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
<i>Robust</i>	[3.02]	[0.88]	[3.02]	[-1.36]	[1.66]	[2.47]	[1.66]	[0.57]	[2.18]	[2.44]	[2.18]	[0.16]
<i>NW</i>	[2.15]	[0.39]	[2.15]	[-0.61]	[1.15]	[1.30]	[1.15]	[0.30]	[1.38]	[1.07]	[1.38]	[0.07]
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
<i>Robust</i>	[3.45]	[0.77]	[3.45]	[-1.53]	[1.53]	[2.77]	[1.53]	[0.66]	[2.54]	[2.75]	[2.54]	[0.23]
<i>NW</i>	[2.23]	[0.33]	[2.23]	[-0.67]	[1.02]	[1.31]	[1.02]	[0.31]	[1.34]	[1.09]	[1.34]	[0.09]
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
<i>Robust</i>	[4.12]	[0.99]	[4.12]	[-2.13]	[1.52]	[3.00]	[1.52]	[1.08]	[2.38]	[2.83]	[2.38]	[0.52]
<i>NW</i>	[1.89]	[0.32]	[1.89]	[-0.68]	[1.00]	[1.56]	[1.00]	[0.56]	[1.10]	[1.22]	[1.10]	[0.22]

Notes: 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 twelve 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 (ψ_f), and similarly the slope coefficient $\zeta_{\mathbf{f}}$ and ζ_f for the exchange rate changes. The t -statistics for the slope coefficients in brackets are computed using the following methods. *Robust* denotes the robust standard errors clustered by month and country; *NW* denotes Newey and West (1987) standard errors computed with the number of lags equal to the horizon of forward discount plus one month. 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.

Table IV: Calibration

Panel A: Targets: Moments of Real Variables			
Moment	Closed Form	Target	Model
$\zeta_{\mathbf{r}}$	$-\frac{\chi}{(\chi - \frac{1}{2}(\gamma + \kappa))}$	1.50	1.50
$R_{\zeta_{\mathbf{r}}}^2$	$\frac{2\chi^2 \text{var}(z)}{\text{var}(\Delta q)}$	1%	0.40%
$Std(\overline{r_t^i} - r_t)$	$(\chi - \frac{1}{2}(\gamma + \kappa)) Std(z)$	0.74%	0.31%
$E(r)$	$\alpha + \theta(\chi - \frac{1}{2}(\gamma + \kappa)) + \theta^w(\tau - \frac{1}{2}\delta^i)$	1.72%	1.72%
$Std(r)$	$\sqrt{(\chi - \frac{1}{2}(\gamma + \kappa))^2 \text{var}(z) + (\tau - \frac{1}{2}\delta^i)^2 \text{var}(z^w)}$	0.57%	1.56%
$Corr(r_t, r_{t+1})$	$\frac{\phi(\chi - \frac{1}{2}(\gamma + \kappa))^2 \text{Var}(z) + (\tau - \frac{1}{2}\delta^i)^2 \phi^w \text{Var}(z^w)}{\text{Var}(r)}$	0.92	0.86
$Corr(\overline{r_t^i} - r_t, \overline{r_{t+1}^i} - r_{t+1})$	ϕ	0.83	0.93
$Std(\Delta q)$	$\sqrt{2\gamma\theta + 2\chi^2 \text{var}(z) + o}$	10.00%	10.05%
$Corr(r_t, r_t^i)$	$(\tau - \frac{1}{2}\delta^i)^2 \frac{\text{Var}(z^w)}{\text{Var}(r)}$	0.30	0.28
$E(rx_t^{\text{dollar}})$	$\gamma\theta$	0.50%	0.50%
<i>Feller</i>	$2(1 - \phi) \frac{\theta}{\text{Var}(z)}, 2(1 - \phi^w) \frac{\theta^w}{\text{Var}(z^w)}$	30.00	30.00
$E(r_t^l - r_t)$	$-\frac{1}{2}(\delta^l - \delta)\theta^w$	5%	5%
$E(r_t^h - r_t)$	$-\frac{1}{2}(\delta^h - \delta)\theta^w$	-5%	-5%
Panel B: Targets: Moments of Inflation			
$Std(\text{inflation})$	$\sqrt{(\eta^w)^2 \text{var}(z^w) + \sigma_\pi^2}$	1.10%	1.10%
R_π^2	$\frac{(\eta^w)^2 \text{var}(z^w)}{\text{var}(\text{inflation})}$	0.28	0.28
$E(\text{inflation})$	$\pi_0 + \eta^w \theta$	2.91%	2.94%
Panel C: Implied Moments of Nominal Variables			
$E(r^{\mathbb{S}})$	$\alpha + \theta(\chi - \frac{1}{2}(\gamma + \kappa)) + \theta^w(\tau + \eta^w - \frac{1}{2}\delta^i) - \frac{1}{2}\sigma_\pi^2$	4.69%	4.38%
$Std(r^{\mathbb{S}})$	$\sqrt{(\chi - \frac{1}{2}(\gamma + \kappa))^2 \text{var}(z^i) + (\tau + \eta^w - \frac{1}{2}\delta^i)^2 \text{var}(z^w)}$	0.65%	1.04%
$Std(\Delta s)$	$\sqrt{2\gamma\theta + 2\chi^2 \text{var}(z^i) + 2\sigma_\pi^2 + o}$	11.07%	10.14%
$Corr(r_t^{\mathbb{S}}, r_t^{\mathbb{S},i})$	$(\tau + \eta^w - \frac{1}{2}\delta^i)^2 \frac{\text{Var}(z^w)}{\text{Var}(r)}$	0.78	0.91
$E(y^{\mathbb{S},10})$		5.90%	6.60%
Panel D: Implied Moments of Real SDF			
$E(Std_t(m))$	$\sqrt{(\gamma + \kappa)\theta + \delta\theta^w + \chi^2 \text{var}(z^i) + \tau^2 \text{var}(z^w)}$		0.85
$Std(Std_t(m_{t+1}))$			5.69%
$E(Corr(m_{t+1}, m_{t+1}^i))$			0.99

Note: $\text{var}(z^w) = \frac{\sigma_w^2 \theta}{1 - \phi^2}$ and $\text{var}(z^i) = \frac{\sigma_i^2 \theta}{1 - \phi^2}$. $o = 2(\delta + \kappa)\theta - 2E\left(\sqrt{\delta^i z_t^w + \kappa^i z_t}\right)\left(\sqrt{\delta^i z_t^w + \kappa^i z_t^i}\right)$ is an order of magnitude smaller than the other terms.

Table V: Parameter Values

Stochastic Discount Factor					
α (%)	χ	τ	γ	κ	δ
2.95	0.61	-1.45	0.04	2.00	9.92
State Variable Dynamics					
ϕ	θ (%)	σ (%)	ϕ^w	θ^w (%)	σ^w (%)
0.93	11.69	0.74	0.85	3.63	0.60
Inflation Dynamics			Heterogeneity		
η^w	π_0 (%)	σ^π (%)	δ_h	δ_l	
2.30	-0.59	0.27	7.71	12.29	

This table reports the parameter values for the calibrated version of the model. These 17 parameters were chosen to match the moments in Table IV under the assumption that all countries share the same parameter values except for δ^i , which is distributed uniformly on $[\delta_h, \delta_l]$.

Table VI: Excess Returns on Carry Strategies: Simulated Data

	Raw Returns		Scaled Returns	
	Dollar	HML	Dollar	HML
<i>Mean</i>	2.08 [1.47]	10.22 [2.28]	3.75 [2.67]	11.87 [2.86]
<i>Std.Dev.</i>	8.76	13.42	15.50	15.50
<i>Sharpe Ratio</i>	0.24	0.77	0.24	0.77

Notes: This table reports the mean, standard deviation and Sharpe ratios of three investment strategies using returns simulated from the model. The first strategy, *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 - *HML carry* - corresponds to currency carry trades (long high interest rate currencies, short low interest currencies). 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. We report (in brackets) standard errors on the means calculated by simulating 1000 samples of length 336 periods and taking the standard deviations of the means across simulations.

Table VII: Forecasting Basket Returns and Exchange Rates: Simulated Data

<i>Horizon</i>	Large sample				Small sample			
	Excess Returns		Exchange Rates		Excess Returns		Exchange Rates	
	$\psi_{\mathbf{f}}$	R^2	$\zeta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\zeta_{\mathbf{f}}$	R^2
1	2.48	0.77	1.48	0.27	3.19	1.43	2.19	0.90
	[13.30]		[7.94]		[1.47]		[1.01]	
2	2.43	1.40	1.43	0.49	3.19	2.70	2.19	1.71
	[12.92]		[7.61]		[1.48]		[1.02]	
3	2.38	1.92	1.38	0.65	3.21	3.87	2.21	2.47
	[12.61]		[7.31]		[1.50]		[1.03]	
6	2.41	3.25	1.41	1.13	3.26	6.80	2.26	4.45
	[12.66]		[7.40]		[1.51]		[1.05]	
12	2.41	4.61	1.41	1.63	3.34	10.59	2.34	7.15
	[12.22]		[7.16]		[1.53]		[1.07]	

Notes: 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 Newey-West (*NW*) 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 brackets use standard deviations of the point estimates across simulations as standard errors.

Table VIII: Contemporaneous Correlations Between Average Forward Discounts and Macroeconomic and Financial Variables

<i>Horizon</i>	<i>IP</i>	<i>Pay</i>	<i>Help</i>	<i>Term</i>	<i>Def</i>	<i>VIX</i>
1	-0.29	-0.20	-0.12	0.48	0.22	-0.09
2	-0.31	-0.21	-0.14	0.49	0.23	-0.09
3	-0.31	-0.21	-0.15	0.49	0.23	-0.08
6	-0.33	-0.23	-0.20	0.49	0.20	-0.06
12	-0.37	-0.28	-0.28	0.48	0.15	-0.05

Notes: 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 *S&P* 500 volatility index (*VIX*). Data are monthly, from Datastream and Global Financial Data. The sample period is 11/1983–6/2010.

Table IX: Forecasting Excess Returns and Exchange Rates with Industrial Production and the Average Forward Discount

Horizon	Developed Countries												Emerging Countries												All Countries											
	Excess Returns				Exchange Rates				Excess Returns				Exchange Rates				Excess Returns				Exchange Rates				Excess Returns				Exchange Rates							
	ψ_{IP}	ψ_f	W	R^2	ζ_{IP}	ζ_f	W	R^2	ψ_{IP}	ψ_f	W	R^2	ζ_{IP}	ζ_f	W	R^2	ψ_{IP}	ψ_f	W	R^2	ζ_{IP}	ζ_f	W	R^2	ψ_{IP}	ψ_f	W	R^2	ζ_{IP}	ζ_f	W	R^2				
1	-0.54	2.14	7.00	3.40	-0.55	1.14	3.18	1.54	-0.95	1.52	5.93	3.79	-1.03	1.46	4.95	4.18	-0.57	1.87	6.67	3.60	-0.55	1.19	3.69	2.08												
NW	[-0.92]	[1.87]	[4.00]		[-0.93]	[1.00]	[39.48]		[-1.60]	[1.76]	[8.41]		[-1.83]	[1.53]	[15.57]		[-1.06]	[1.83]	[5.08]		[-1.02]	[1.16]	[30.96]													
VAR	[-0.96]	[2.44]	[0.00]		[-0.96]	[1.30]	[0.20]		[-1.86]	[1.63]	[0.00]		[-2.08]	[1.74]	[0.00]		[-1.14]	[2.32]	[0.10]		[-1.09]	[1.50]	[0.00]													
2	-0.65	2.09	10.35	6.25	-0.65	1.09	6.72	3.16	-1.02	1.46	6.02	7.43	-1.10	1.46	4.91	7.98	-0.66	1.84	8.88	6.68	-0.64	1.18	5.94	4.14												
NW	[-2.14]	[1.93]	[0.27]		[-2.14]	[1.01]	[4.89]		[-2.44]	[1.52]	[7.94]		[-2.20]	[1.29]	[15.91]		[-2.17]	[1.94]	[0.94]		[-2.01]	[1.20]	[8.34]													
VAR	[-1.23]	[1.92]	[0.00]		[-1.21]	[1.06]	[0.20]		[-1.90]	[1.26]	[0.00]		[-2.25]	[1.30]	[0.00]		[-1.36]	[1.99]	[0.00]		[-1.45]	[1.26]	[0.00]													
3	-0.72	1.99	23.67	8.68	-0.72	0.99	19.74	4.67	-1.14	1.30	8.46	10.86	-1.20	1.30	6.23	11.42	-0.74	1.73	13.96	9.28	-0.71	1.07	10.44	6.01												
NW	[-4.05]	[1.66]	[0.00]		[-4.05]	[0.83]	[0.00]		[-2.86]	[1.42]	[1.32]		[-2.49]	[1.20]	[6.87]		[-3.38]	[1.66]	[0.01]		[-3.03]	[1.01]	[0.25]													
VAR	[-1.41]	[1.77]	[0.00]		[-1.58]	[0.96]	[0.00]		[-2.40]	[1.17]	[0.00]		[-2.56]	[1.16]	[0.00]		[-1.69]	[1.82]	[0.00]		[-1.64]	[1.19]	[0.00]													
6	-0.87	1.84	38.02	15.58	-0.87	0.84	32.05	9.57	-1.34	1.08	7.12	19.75	-1.39	1.11	5.87	20.67	-0.90	1.56	16.35	16.96	-0.88	0.92	13.78	12.28												
NW	[-4.83]	[1.47]	[0.00]		[-4.83]	[0.67]	[0.00]		[-2.66]	[1.31]	[3.68]		[-2.41]	[1.13]	[8.78]		[-3.61]	[1.46]	[0.00]		[-3.42]	[0.84]	[0.01]													
VAR	[-1.79]	[1.72]	[0.00]		[-1.87]	[0.82]	[0.00]		[-2.97]	[0.96]	[0.00]		[-3.12]	[1.03]	[0.00]		[-2.18]	[1.66]	[0.00]		[-2.16]	[1.00]	[0.00]													
12	-0.91	1.37	16.75	23.20	-0.92	0.37	13.01	15.23	-1.33	1.18	8.65	31.37	-1.36	1.22	7.49	32.53	-0.94	1.13	15.46	24.86	-0.90	0.50	12.99	18.87												
NW	[-3.27]	[1.19]	[0.00]		[-3.27]	[0.33]	[0.02]		[-2.88]	[1.34]	[1.13]		[-2.67]	[1.13]	[2.81]		[-3.35]	[1.17]	[0.00]		[-3.26]	[0.50]	[0.02]													
VAR	[-2.27]	[1.41]	[0.00]		[-2.32]	[0.34]	[0.00]		[-4.22]	[1.48]	[0.00]		[-4.36]	[1.51]	[0.00]		[-2.69]	[1.28]	[0.00]		[-2.68]	[0.57]	[0.00]													


Notes: 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 twelve 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 (ψ_f), and similarly the slope coefficients ζ_{IP} , ζ_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. NW denotes Newey and West (1987) standard errors computed with the optimal number of lags following Andrews (1991). 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. 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 (NW) 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.

Table X: Forecasting Returns and Exchange Rates with Realized Consumption Volatility and Inflation Volatility

Horizon	Developed Countries										Emerging Countries										All Countries									
	Excess Returns			Exchange Rates			Excess Returns			Exchange Rates			Excess Returns			Exchange Rates			Excess Returns			Exchange Rates								
	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	2.27	0.66	1.26	0.75	1.13	0.31	-1.66	1.90	1.90	0.62	-1.55	2.69	2.50	0.82	1.03	1.09	1.00	0.35	0.74	1.21	0.87	0.28						
NW	[1.24]	[0.42]	[57.45]		[0.83]	[0.34]	[81.76]		[-1.04]	[0.91]	[65.47]		[-0.90]	[1.27]	[52.62]		[0.69]	[0.55]	[84.27]		[0.53]	[0.61]	[86.80]							
VAR	[0.85]	[0.26]	[0.20]		[0.50]	[0.24]	[0.30]		[-0.87]	[0.70]	[0.20]		[-0.75]	[1.01]	[0.10]		[0.41]	[0.40]	[0.20]		[0.32]	[0.40]	[0.30]							
2	2.19	0.91	2.28	1.48	1.52	0.72	1.33	0.76	-1.57	2.12	1.81	1.20	-1.47	2.92	2.41	1.58	1.20	1.18	1.13	0.82	0.92	1.30	1.02	0.67						
NW	[1.22]	[0.38]	[57.26]		[0.89]	[0.30]	[77.64]		[-0.83]	[1.02]	[67.29]		[-0.72]	[1.33]	[54.41]		[0.72]	[0.56]	[81.65]		[0.58]	[0.62]	[83.96]							
VAR	[1.06]	[0.32]	[0.20]		[0.74]	[0.30]	[0.30]		[-0.85]	[0.79]	[0.10]		[-0.85]	[1.23]	[0.00]		[0.69]	[0.47]	[0.30]		[0.49]	[0.51]	[0.10]							
3	2.38	1.28	3.09	2.65	1.71	1.08	2.00	1.52	-1.49	2.58	2.44	1.93	-1.40	3.38	3.25	2.56	1.36	1.57	1.72	1.63	1.08	1.68	1.66	1.42						
NW	[1.34]	[0.57]	[41.10]		[1.02]	[0.48]	[63.31]		[-0.79]	[1.30]	[53.79]		[-0.68]	[1.60]	[38.27]		[0.82]	[0.79]	[69.18]		[0.68]	[0.85]	[70.63]							
VAR	[1.19]	[0.45]	[0.10]		[0.80]	[0.41]	[0.00]		[-0.83]	[1.01]	[0.00]		[-0.83]	[1.37]	[0.00]		[0.75]	[0.60]	[0.20]		[0.62]	[0.73]	[0.20]							
6	2.95	1.36	5.16	6.73	2.33	1.12	3.78	4.54	-1.08	2.77	2.97	2.59	-0.96	3.59	4.01	3.68	1.86	1.62	2.72	4.23	1.63	1.75	2.88	3.95						
NW	[1.88]	[0.70]	[13.66]		[1.55]	[0.58]	[29.52]		[-0.62]	[1.47]	[43.32]		[-0.52]	[1.79]	[26.25]		[1.25]	[0.89]	[48.16]		[1.15]	[0.97]	[45.07]							
VAR	[1.38]	[0.51]	[0.00]		[1.15]	[0.44]	[0.00]		[-0.69]	[1.11]	[0.00]		[-0.57]	[1.53]	[0.00]		[1.02]	[0.75]	[0.20]		[0.95]	[0.78]	[0.00]							
12	3.75	1.00	13.94	19.57	3.23	0.71	11.11	15.69	-0.35	1.86	1.08	1.45	-0.20	2.81	1.96	3.09	2.67	1.05	6.99	12.45	2.52	1.24	7.65	12.67						
NW	[3.34]	[0.59]	[0.01]		[2.90]	[0.42]	[0.14]		[-0.22]	[0.94]	[82.66]		[-0.12]	[1.32]	[64.02]		[2.47]	[0.61]	[4.03]		[2.40]	[0.72]	[2.49]							
VAR	[1.86]	[1.57]	[0.00]		[1.62]	[0.91]	[0.00]		[-0.26]	[3.85]	[0.20]		[-0.19]	[6.97]	[0.20]		[1.58]	[1.49]	[0.10]		[1.48]	[2.38]	[0.00]							


Notes: 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 twelve 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. *NW* denotes Newey and West (1987) standard errors computed with the optimal number of lags following Andrews (1991). 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.

Figure 1: Average 12-month Forward Discounts on Three Currency Baskets



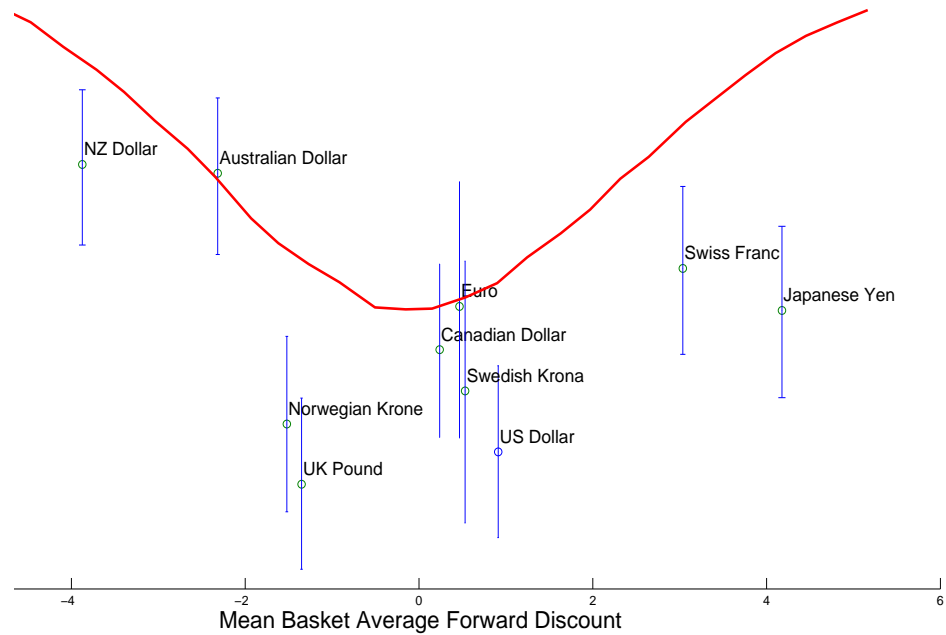
This figure presents the average 12-month forward discounts on three currency baskets (developing countries, developed countries, and all countries). The shaded areas are U.S. recessions according to NBER. The sample period is 11/1983–6/2010.

Figure 2: Carry Trade Excess Return Indexes



The 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 CRSP value-weighted U.S. stock market portfolio. All strategies are levered to match the volatility of the stock market.

Figure 3: 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 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.