Common Risk Factors in Currency Markets*

Hanno Lustig — UCLA Anderson and NBER

Nick Roussanov — Wharton

Adrien Verdelhan — Boston University

December 12, 2008

Abstract

Currency excess returns are highly predictable and strongly counter-cyclical. The average excess returns on low interest rate currencies are 4.8 percent per annum smaller than those on high interest rate currencies after accounting for transaction costs. A single return-based factor, the return on the highest minus the return on the lowest interest rate currency portfolios, explains the cross-sectional variation in average currency excess returns from low to high interest rate currencies. In a simple affine pricing model, we show that the high-minus-low currency return measures that component of the stochastic discount factor innovations that is common across countries. To match the carry trade returns in the data, low interest rate currencies need to load more on this common innovation when the market price of global risk is high.

JEL: G12, G15, F31.

*First Version September 2007. Lustig: Anderson School of Management, University of California at Los Angeles, Box 951477, Los Angeles, CA 90095; hlustig@anderson.ucla.edu; Tel: (310) 825-8018; `http://hlustig2001.squarespace.com/`. Roussanov: Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104; nroussan@wharton.upenn.edu; Tel: (215) 545-1020; `http://www.wharton.upenn.edu/faculty/roussanov.html`. Verdelhan: Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215; av@bu.edu; Tel: (617) 353-6324; `http://people.bu.edu/av`.

The authors thank Andy Atkeson, Alessandro Beber, Frederico Belo, Michael Brennan, Alain Chaboud, John Cochrane, Pierre Collin-Dufresne, Magnus Dahlquist, Kent Daniel, Darrell Duffie, Xavier Gabaix, Urban Jermann, Don Keim, Leonid Kogan, Olivier Jeanne, Karen Lewis, Fang Li, Francis Longstaff, Pascal Maenhout, Rob Martin, Anna Pavlova, Monika Piazzesi, Richard Roll, Geert Rouwenhorst, Clemens Sialm, Rob Stambaugh, Jessica Wachter, Amir Yaron, Hongjun Yan, Moto Yogo and seminar participants at the University of Amsterdam Business School, Dartmouth, Duke - Fuqua School of Business, the Federal Reserve Board of Governors, the Federal Reserve Bank of Texas, INSEAD, MIT Sloan, the University of North Carolina - Kenan-Flagler Business School, Purdue University, SUNY Buffalo, UCLA Finance Lunch, UCLA, the University of Minnesota - Carlson School of Business, the University of Illinois at Urbana-Champaign, Wharton, Yale School of Management, the 2008 CEPR and Swiss National Bank Conference, the 2008 Western Finance Association Annual Meeting, the 2008 CEPR European Summer Symposium in Financial Markets, the 2008 European Economic Association Annual Meeting, the 2008 European Finance Association Annual Meeting and the NBER IFM meeting for helpful comments.

In this paper, we demonstrate that currency risk premia are a robust feature of the data, even after accounting for transaction costs. We show that currency risk premia are determined by their exposure to a single, global risk factor, and that interest rates measure currency exposure to this factor. This global risk factor explains most of the cross-sectional variation in average excess returns between high and low interest rate currencies. We show that by investing in high interest rate currencies and borrowing in low interest rate currencies, US investors load up on global risk, especially during "bad times". After accounting for the covariance with this risk factor, there are no significant anomalous or unexplained excess returns in currency markets. In addition, we show that most of the time-series variation in currency risk premia is explained by the average interest rate difference between the US and foreign currencies, not the currency-specific interest rate difference. The average interest rate difference is highly counter-cyclical, and so are currency risk premia. We can replicate our main findings in a no-arbitrage model of exchange rates with two factors, a country-specific factor and a global factor, but only if low interest rate currencies are more exposed to global risk in bad times. Heterogeneity in exposure to country-specific risk cannot explain the carry trade returns.

We identify the common risk factor in the data by building portfolios of currencies. As in Lustig and Verdelhan (2007), we sort currencies on their forward discounts and allocate them to six portfolios. Forward discounts are the difference between log forward rates and log spot rates. Since covered interest rate parity typically holds, forward discounts equal the interest rate difference between two currencies. As a result, the first portfolio contains the lowest interest rate currencies while the last portfolio contains the highest interest rate currencies. Unlike Lustig and Verdelhan (2007), we only use spot and forward exchange rates to compute returns. These contracts are easily tradable, subject to minimal counterparty risk, and their transaction costs are easily available. As a consequence, our main sample comprises 37 currencies. We account for bid-ask spreads that investors incur when they trade these spot and forward contracts.

Risk premia in currency markets are large and time-varying. For each portfolio, we compute the monthly foreign currency excess returns realized by buying or selling one-month forward contracts for all currencies in the portfolio, net of transaction costs. Between the end of 1983 and the beginning of 2008, US investors earn an annualized log excess return of 4.8 percent by buying one-month forward contracts for currencies in the last portfolio and by selling forward contracts for currencies in the first portfolio. The annualized Sharpe ratio on such a strategy is .54. These findings are robust. We find similar results when we limit the sample to developed currencies, and when we take the perspective of investors in other countries. In this paper, we investigate the cross-sectional and time-series properties of these currency excess returns.

In the data, the first two principal components of the currency portfolio returns account for most of the time series variation in currency returns. The first principal component is essentially the average excess return on all foreign currency portfolios. We call this average excess return the dollar risk factor RX. The second component is very similar to the return on a zero-cost strategy that goes long in the last portfolio and short in the first portfolio. We label this excess return the carry trade risk factor HML_{FX}, for high interest rate minus low interest rate currencies. The carry trade risk factor HML_{FX} explains about 70 percent of the variation in average excess returns on our 6 currency portfolios. The risk price of this carry trade factor that we estimate from the cross-section of currency portfolio returns is roughly equal to its sample mean, consistent with a linear factor pricing model. Low interest rate currencies provide US investors with insurance against HML_{FX} risk, while high interest rate currencies expose investors to more HML_{FX} risk. By ranking currencies into portfolios based on their forward discounts, we find that forward discounts determine currencies' exposure to HML_{FX}, and hence their risk premia. As a check, we also rank currencies based on their HML_{FX}-betas, and we find that portfolios with high HML_{FX}-exposure do yield higher average returns and have higher forward discounts.

We show that the carry trade risk factor has explanatory power for the returns on momentum currency portfolios built by ranking currencies on past returns rather than on interest rates. This lends support to a risk-based rather than a characteristic-based explanation of our findings; a characteristic-based explanation would imply that our risk factor has no explanatory power for currency portfolios not constructed by sorting on interest rates.

To explain our findings, we use a standard no-arbitrage exponentially-affine asset pricing model. Our model features a large number of countries. The stochastic discount factor (SDF) that prices assets in the units of a given country's currency is composed of two risk factors: one is country-specific, the other is common for all countries. We show analytically that two conditions need to be satisfied in order to match the data. First, we need a common risk factor because it is the only source of cross-sectional variation in currency risk premia. Second, we need low interest rate currencies to be more exposed to the common risk factor in times when the price of common risk is high, i.e. in bad times. Using the model, we show analytically that by sorting currencies into portfolios and constructing HML_{FX}, we measure the common innovation to the SDF. Similarly, we show that the dollar risk factor RX measures the home-country-specific innovation to the SDF. Thus, we provide a theoretical foundation for building currency portfolios: by doing so, we recover the two factors that drive the pricing kernel.

In the model, currency risk premia are determined by a dollar risk premium and a carry trade risk premium. The size of the carry trade risk premium depends on the spread in the loadings on the common component between high and low interest rate currencies, and on the global risk price. As the global risk price increases, the spread increases endogenously and the carry trade risk premium goes up. If there is no spread, i.e. if low and high interest rate currencies share the same loadings on the common risk factor, then HML_{FX} cannot be a risk factor, because the global component does not affect exchange rates. The larger the spread, the riskier high interest rate currencies become relative to low interest rate currencies, because the latter appreciate relative to the former in case of a negative global shock and hence offer insurance. In a version of the model that is calibrated to match moments of exchange rates and interest rates in the data, we replicate the carry trade risk premium as well as the failure of the CAPM to explain average currency returns in the data.

Finally, there is far more predictability in currency portfolio returns than in the returns on individual currencies. As predicted by the no-arbitrage model, the average forward discount rate is a better predictor of portfolio returns than the forward discounts of individual currency portfolios. This result echoes the finding of Cochrane and Piazzesi (2005) that a linear combination of forward rates across maturities is a powerful predictor of excess returns on bonds. Expected excess returns on portfolios with medium to high interest rates co-move negatively with the US business cycle as measured by industrial production, payroll or help wanted indices, and they co-move positively with the term and default premia as well as the option-implied volatility index VIX. Forecasted excess returns on high interest rate portfolios are strongly counter-cyclical and increase in times of crisis, as predicted by our model. In fact, we find that US industrial production growth has predictive power for currency excess returns even when controlling for forward discounts. In recent work, Duffee (2008) and Ludvigson and Ng (2005) report similar findings for the bond market, and Piazzesi and Swanson (2008) document that payroll growth predicts excess returns on interest rate futures.

Related Literature There is a large literature that documents the failure of UIP in the time series, starting with the work of Hansen and Hodrick (1980a) and Fama (1984): higher than usual interest rates lead to further appreciation, and investors earns more by holding bonds from currencies with interest rates that are *higher than usual* (see e.g. Cochrane (2001)). Hence, this seems to imply that currency investors need to know what 'higher than usual' means for a specific currency. Bansal and Dahlquist (2000) survey the time series evidence for a large number of currencies and they conclude that country-specific attributes are critical to understanding the cross-sectional variation in currency risk premia.

By building portfolios of positions in currency forward contracts sorted on forward discounts (as Lustig and Verdelhan (2007) do with T-bills), we show that UIP also fails in the cross-section: currently high interest rate currencies depreciate 5.9 % per annum less than the interest rate difference, while currently low interest rate currencies appreciate 2.9 % per annum less than the interest rate difference. Hence, investors earns more simply by holding bonds from currencies with interest rates that are *currently high*. As a result, currency-specific attributes other than the interest rate cannot be the only explanation, because these currencies switch portfolios when their interest rates change, and these switches are frequent. Instead, we show that low and high interest rate currencies have different risk characteristics.

Papers that address the failure of UIP can be divided into two broad classes. The first class aims to understand exchange rate predictability within a standard asset pricing framework based on systematic risk.[1] Hollifield and Yaron (2001) are the first to provide empirical evidence that real factors account for the forward premium. The second class looks for non-risk-based explanations.[2] Recent contributions to the risk-based literature offer three types of fully-specified models of the forward premium puzzle: Verdelhan (2005) uses habit preferences in the vein of Campbell and Cochrane (1999), Bansal and Shaliastovich (2007) build on the long run risk model pioneered by Bansal and Yaron (2004), and Farhi and Gabaix (2007) augment the standard consumption-based model with disaster risk following Barro (2006). These three models have two elements in common: a persistent variable drives the volatility of the log stochastic discount factor, and this variable comoves negatively with the country's risk-free interest rate. Backus et al. (2001) show that the latter is a necessary condition for models with log-normal shocks to reproduce the forward premium puzzle. Our paper adds to this list of requirements. To reproduce our finding that a single global risk factor explains the cross-section of currency returns, the SDF in these models needs to have a common heteroscedastic component, and the SDF in low interest rate currencies needs to load more on the common component. This heterogeneity is critical for replicating our empirical findings; we show that heterogeneity in the loadings on the country-specific factor cannot explain the cross-sectional variation in currency returns, even though it can generate negative UIP slope coefficients. Finally, we also show that HML_{FX} is strongly related to macroeconomic risk; it has a US consumption growth beta between 1 and 1.5, consistent with the findings of Lustig and Verdelhan (2007) who use the Consumption-CAPM to explain currency returns. In recent related work, DeSantis and Fornari (2008) provide more evidence that currency returns compensate investors for systematic, business cycle risk. Finally, our paper is connected to work by Gorton, Hayashi and Rouwenhorst (2007), who rank commodities into portfolios based on their basis. They also find a connection between these 'carry' portfolios of commodities and momentum portfolios of commodities.

Our paper is organized as follows. We start by describing the data, how we build currency portfolios and the main characteristics of these portfolios in section 1. Section 2 shows that a

[1]This segment includes recent papers by Backus, Foresi and Telmer (2001), Harvey, Solnik and Zhou (2002), Alvarez, Atkeson and Kehoe (2005), Verdelhan (2005), Campbell, de Medeiros and Viceira (2006), Lustig and Verdelhan (2007), Graveline (2006), Bansal and Shaliastovich (2007), Brennan and Xia (2006), Farhi and Gabaix (2007) and Hau and Rey (2007), Colacito (2008) and Brunnermeier, Nagel and Pedersen (2008). Earlier work includes Hansen and Hodrick (1980a), Fama (1984), Korajczyk (1985), Bekaert and Hodrick (1992), Bekaert (1995) and Bekaert (1996).

[2]This segment includes papers by Froot and Thaler (1990), Lyons (2001), Gourinchas and Tornell (2004), Bacchetta and van Wincoop (2006), Frankel and Poonawala (2007), Sarno, Leon and Valente (2006), Plantin and Shin (2007), Burnside, Eichenbaum, Kleshchelski and Rebelo (2006), Burnside, Eichenbaum and Rebelo (2007), Burnside, Eichenbaum and Rebelo (2008a) and Burnside, Eichenbaum, Kleshchelski and Rebelo (2008b).

single factor, HML_{FX}, explains most of the cross-sectional variation in foreign currency excess returns. In section 3, we use a no-arbitrage model of exchange rates to interpret these findings. Section 4 describes the time variation in excess returns that investors demand on these currency portfolios. Section 5 considers a calibrated version of the model that replicates the key moments of the data. Finally, section 6 shows that the carry trade risk factor explains some of the variation in momentum currency portfolios that are sorted on past returns, lending additional support to our risk-based explanation. Section 7 concludes. All the tables and figures are in the appendix. The portfolio data can be downloaded from our web site and are regularly updated. We also posted a separate appendix on-line with some additional results.

1 Currency Portfolios and Risk Factors

We focus on investments in forward and spot currency markets. Compared to Treasury Bill markets, forward currency markets only exist for a limited set of currencies and shorter time-periods. However, forward currency markets offer two distinct advantages. First, the carry trade is easy to implement in these markets, and the data on bid-ask spreads for forward currency markets are readily available. This is not the case for most foreign fixed income markets. Second, these forward contracts are subject to minimal default and counterparty risks. This section describes the properties of monthly foreign currency excess returns from the perspective of a US investor. We consider currency portfolios that include developed and emerging market countries for which forward contracts are traded. We find that currency markets offer Sharpe ratios comparable to the ones measured in equity markets, even after controlling for bid-ask spreads.

1.1 Building Currency Portfolios

We start by setting up some notation. Then, we describe our portfolio building methodology, and we conclude by giving a summary of the currency portfolio returns.

Currency Excess Returns We use s to denote the log of the spot exchange rate in units of foreign currency per US dollar, and f for the log of the forward exchange rate, also in units of foreign currency per US 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: $f_t - s_t \approx i_t^\star - i_t$, where $i^\star$ and i denote the foreign and domestic nominal risk-free rates over the maturity of the contract. Akram, Rime and Sarno (2008) study high frequency deviations from covered interest rate parity (CIP). They conclude that CIP holds at daily and lower frequencies. Hence, the log currency excess return approximately equals the interest rate differential less the rate of depreciation:

$$rx_{t+1} \approx i_t^\star - i_t - \Delta s_{t+1}.$$

Transaction Costs Since we have bid-ask quotes for spot and forward contracts, we can compute the investor's actual realized excess return net of transaction costs. The *net* log currency excess return for an investor who goes long in foreign currency is:

$$rx_{t+1}^l = f_t^b - s_{t+1}^a.$$

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:

$$rx_{t+1}^s = -f_t^a + s_{t+1}^b.$$

Data We start from daily spot and forward exchange rates in US dollars. We build end-of-month series from November 1983 to March 2008. These data are collected by Barclays and Reuters and available on Datastream. Lyons (2001) reports that bid-ask spreads from Reuters are roughly twice the size of inter-dealer spreads (page 115). As a result, our estimates of the transaction costs are conservative. Lyons (2001) also notes that these indicative quotes track inter-dealer quotes closely, only lagging the inter-dealer market slightly at very high intra-day frequency. This is clearly not an issue here at monthly horizons. Our main data set contains 37 currencies: Australia, Austria, Belgium, Canada, Hong Kong, Czech Republic, Denmark, Euro area, Finland, France, Germany, Greece, Hungary, India, Indonesia, Ireland, Italy, Japan, Kuwait, Malaysia, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Saudi Arabia, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, United Kingdom. Some of these currencies have pegged their exchange rate partly or completely to the US dollar over the course of the sample. We keep them in our sample because forward contracts were easily accessible to investors. We leave out Turkey and United Arab Emirates, even if we have data for these countries, because their forward rates appear disconnected from their spot rates. As a robustness check, we also a study a smaller data set that contains only 15 developed countries: Australia, Belgium, Canada, Denmark, Euro area, France, Germany, Italy, Japan, Netherlands, New Zealand, Norway, Sweden, Switzerland and United Kingdom. We present all of our results on these two samples.

Currency Portfolios At the end of each period t, we allocate all currencies in the sample to six portfolios on the basis of their forward discounts $f-s$ observed at the end of period t. Portfolios are re-balanced at the end of every month. They are ranked from low to high interests rates; portfolio 1 contains the currencies with the lowest interest rate or smallest forward discounts, and portfolio 6 contains the currencies with the highest interest rates or largest forward discounts. We compute the log currency excess return rx^j_{t+1} for portfolio j by taking the average of the log currency excess returns in each portfolio j. For the purpose of computing returns net of bid-ask spreads we assume that investors *short* all the foreign currencies in the *first* portfolio and go *long* in all the other foreign currencies.

The total number of currencies in our portfolios varies over time. We have a total of 9 countries at the beginning of the sample in 1983 and 26 at the end in 2008. We only include currencies for which we have forward and spot rates in the current and subsequent period. The maximum number of currencies attained during the sample is 34; the launch of the euro accounts for the subsequent decrease in the number of currencies. The average number of portfolio switches per month is 6.01 for portfolios sorted on one-month forward rates. We define the average frequency as the time-average of the following ratio: the number of portfolio switches divided by the total number of currencies at each date. The average frequency is 29.32 percent, implying that currencies switch portfolios roughly every three months. When we break it down by portfolio, we get the following frequency of portfolio switches (in percentage points): 19.9 for the 1st, 33.8 for the 2nd, 40.7 for the 3rd, 43.4 for the 4th, 42.0 for the 5th, and 13.4 for the 6th. Overall, there is quite some variation in the composition of these portfolios, but there is more persistence in the composition of the corner portfolios. As an example, we consider the Japanese yen. The yen starts off in the fourth portfolio early on in the sample, then gradually ends up in the first portfolio as Japanese interest rates fall in the late eighties and it briefly climbs back up to the sixth portfolio in the early nineties. The yen stays in the first portfolio for the remainder of the sample.

1.2 Returns to Currency Speculation for a US investor

Table 1 provides an overview of the properties of the six currency portfolios from the perspective of a US investor. For each portfolio j, we report average changes in the spot rate Δs^j, the forward discounts $f^j - s^j$, the log currency excess returns $rx^j = -\Delta s^j + f^j - s^j$, and the log currency excess returns net of bid-ask spreads rx^j_{net}. Finally, we also report log currency excess returns on carry trades or high-minus-low investment strategies that go long in portfolio $j = 2, 3 \ldots, 6$, and short in the first portfolio: $rx^j_{net} - rx^1_{net}$. All exchange rates and returns are reported in US dollars and the moments of returns are annualized: we multiply the mean of the monthly data by 12 and the standard deviation by $\sqrt{12}$. The Sharpe ratio is the ratio of the annualized mean to the annualized standard deviation.

The first panel reports the average rate of depreciation for all currencies in portfolio j. According to the standard uncovered interest rate parity (UIP) condition, the average rate of depreciation $E_T(\Delta s^j)$ of currencies in portfolio j should equal the average forward discount on these currencies $E_T(f^j - s^j)$, reported in the second panel. Instead, currencies in the first portfolio trade at an average forward discount of -390 basis points, but they appreciate on average only by almost 100 basis points over this sample. This adds up to a log currency excess return of minus 290 basis points on average, which is reported in the third panel. Currencies in the last portfolio trade at an average discount of 778 basis points but they depreciate only by 188 basis points on average. This adds up to a log currency excess return of 590 basis points on average. A large body of empirical work starting with Hansen and Hodrick (1980b) and Fama (1984) reports violations of UIP. However, our results are different because our investment strategy only considers whether the currency's interest rate is currently high, not whether it is higher than usual.

The fourth panel reports average log currency excess returns net of transaction costs. Since we rebalance portfolios monthly, and transaction costs are incurred each month, these estimates of net returns to currency speculation are conservative. After taking into account bid-ask spreads, the average return on the first portfolio drops to minus 170 basis points. Note that the first column reports *minus* the actual log excess return for the first portfolio, because the investor is short in these currencies. The corresponding Sharpe ratio on this first portfolio is minus 0.21. The return on the sixth portfolio drops to 314 basis points. The corresponding Sharpe ratio on the last portfolio is 0.34.

The fifth panel reports returns on zero-cost strategies that go long in the high interest rate portfolios and short in the low interest rate portfolio. The spread between the net returns on the first and the last portfolio is 483 basis points. This high-minus-low strategy delivers a Sharpe ratio of 0.54, after taking into account bid-ask spreads. Equity returns provide a natural benchmark. Over the same sample, the (annualized) Fama-French monthly excess return on the US stock market is 7.11 percent, and the equity Sharpe ratio is 0.48. Note that this equity return does *not* reflect any transaction cost.

[Table 1 about here.]

We have documented that a US investor with access to forward currency markets can realize large excess returns with annualized Sharpe ratios that are comparable to those in the US stock market. Table 1 also reports results obtained on a smaller sample of developed countries. The

Sharpe ratio on a long-short strategy is 0.39. There is no evidence that time-varying bid-ask spreads can account for the failure of UIP in these data or that currency excess returns are small in developed countries, as suggested by Burnside et al. (2006). We turn now to cross-sectional asset pricing tests on these currency portfolios.

2 Common Factors in Currency Returns

We show that the sizeable currency excess returns described in the previous section are matched by covariances with risk factors. The riskiness of different currencies can be fully understood in terms of two currency factors that are essentially the first two principal components of the portfolio returns. All portfolios load equally on the first component, which is essentially the average currency excess return. We label it the *dollar risk factor*. The second principal component, which is very close to the difference in returns between the low and high interest rate currencies, explains a large share of the cross-section. We refer to this component as the *carry risk factor*. The risk premium on any currency is determined by the dollar risk premium and the carry risk premium. The carry risk premium depends on which portfolio a currency belongs to, i.e. whether the currency has high or low interest rates, but the dollar risk premium does not. To show that a currency's interest rate relative to that of other currencies truly measures its exposure to carry risk, we also sort all the currencies into portfolios based on their carry-betas, and we recover a similar pattern in the forward discounts and in the excess returns.

2.1 Methodology

Linear factor models predict that average returns on a cross-section of assets can be attributed to risk premia associated with their exposure to a small number of risk factors. In the arbitrage pricing theory of Ross (1976), these factors capture common variation in individual asset returns. A principal component analysis on our currency portfolios reveals that two factors explain more than 80 percent of the variation in returns on these six portfolios. The top panel in table 2 reports the loadings of our currency portfolios on each of the principal components as well as the fraction of the total variance of portfolio returns attributed to each principal component. The first principal component explains 70 percent of common variation in portfolio returns, and can be interpreted as a *level* factor, since all portfolios load equally on it. The second principal component, which is responsible for over 12 percent of common variation, can be interpreted as a *slope* factor, since portfolio loadings increase monotonically across portfolios. The first principal component is indistinguishable from the average portfolio return. The second principal component is essentially the difference between the return on the sixth portfolio and the return on the first portfolio. As a consequence, we consider two risk factors: the average currency excess return, denoted RX, and the difference between the return on the last portfolio and the one on the first portfolio, denoted HML_{FX}. The correlation of the first principal component with RX is .99. The correlation of the second principal component with HML_{FX} is .94. Both factors are computed from net returns, after taking into account bid-ask spreads. The bottom panel confirms that we obtain similar results even when we exclude developing countries from the sample.

These currency risk factors have a natural interpretation. HML_{FX} is the return in dollars on a zero-cost strategy that goes long in the highest interest rate currencies and short in the lowest interest rate currencies. This is the portfolio return of a US investor engaged in the usual currency carry trade. Hence, this is a natural candidate currency risk factor, and, as we are about to show, it explains much of the cross-sectional variation in average excess returns. RX is the average portfolio return of a US investor who buys all foreign currencies available in the forward market. This second factor is essentially the currency "market" return in dollars available to an US investor.

[Table 2 about here.]

Before turning to our main asset pricing estimates, we report on a simple experiment to build intuition for our results. Following Cochrane and Piazzesi (2008), we compute the covariance of each principal component with the currency portfolio returns, and we compare these covariances (indicated by triangles) with the average currency excess returns (indicated by squares) for each portfolio. Figure 1 illustrates that the second principal component plays a key role. Its covariance with currency excess returns increases monotonically as we go from portfolio 1 to 6.[3] This is not the case for any of the other principal components. As a result, in the space of portfolio returns, the second principal component seems crucial.

[Figure 1 about here.]

Cross-Sectional Asset Pricing We use Rx^j_{t+1} to denote the average excess return on portfolio j in period $t+1$. All asset pricing tests are run on excess returns and not log excess returns. In the absence of arbitrage opportunities, this excess return has a zero price and satisfies the following Euler equation:

$$E_t[M_{t+1}Rx^j_{t+1}] = 0.$$

We assume that the stochastic discount factor M is linear in the pricing factors f:

$$M_{t+1} = 1 - b(f_{t+1} - \mu),$$

[3]We thank John Cochrane for suggesting this figure. Figure 1 is the equivalent of figure 6 page 25 of Cochrane and Piazzesi (2008).

where b is the vector of factor loadings and μ denotes the factor means. This linear factor model implies a beta pricing model: the expected excess return is equal to the factor price λ times the beta of each portfolio β^j:

$$E[Rx^j] = \lambda' \beta^j,$$

where $\lambda = \Sigma_{ff} b$, $\Sigma_{ff} = E(f_t - \mu_f)(f_t - \mu_f)'$ is the variance-covariance matrix of the factor, and β^j denotes the regression coefficients of the return Rx^j on the factors. To estimate the factor prices λ and the portfolio betas β, we use two different procedures: a Generalized Method of Moments estimation (GMM) applied to linear factor models, following Hansen (1982), and a two-stage OLS estimation following Fama and MacBeth (1973), henceforth FMB. In the first step, we run a time series regression of returns on the factors. In the second step, we run a cross-sectional regression of average returns on the factors. We do not include a constant in the second step ($\lambda_0 = 0$) .

2.2 Results

Table 3 reports the asset pricing results obtained using GMM and FMB on currency portfolios sorted on forward discounts. The left hand side of the table corresponds to our large sample of developed and emerging countries, while the right hand side focuses on developed countries. We describe first results obtained on our large sample.

[Table 3 about here.]

Market Prices of Risk The top panel of the table reports estimates of the market prices of risk λ and the SDF factor loadings b, the adjusted R^2, the square-root of mean-squared errors $RMSE$ and the p-values of χ^2 tests (in percentage points). The market price of HML_{FX} risk is 546 basis points *per annum*. This means that an asset with a beta of one earns a risk premium of 5.46 percent per annum. Since the factors are returns, no arbitrage implies that the risk prices of these factors should equal their average excess returns. This condition stems from the fact that the Euler equation applies to the risk factor itself, which clearly has a regression coefficient β of one on itself. In our estimation, this no-arbitrage condition is satisfied. The average excess return on the high-minus-low strategy (last row in Table 3) is 537 basis points. This value differs slightly from the previously reported mean excess return because we use excess returns in *levels* in the asset pricing exercise, but table 1 reports *log* excess returns to illustrate their link to changes in exchange rates and interest rate differentials. So the estimated risk price is only 9 basis points removed from the point estimate implied by linear factor pricing. The GMM standard error of the risk price is 234 basis points. The FMB standard error is 183 basis points. In both cases, the risk price is more than two standard errors from zero, and thus highly statistically significant.

The second risk factor RX, the average currency excess return, has an estimated risk price of 135 basis points, compared to a sample mean for the factor of 136 basis points. This is not surprising, because all the portfolios have a beta close to one with respect to this second factor. As a result, the second factor explains none of the cross-sectional variation in portfolio returns, and the standard errors on the risk price estimates are large: for example, the GMM standard error is 168 basis points. When we drop the dollar factor, the RMSE rises from 95 to 168 basis points, but the adjusted R^2 is still 76 %. The dollar factor does not explain any of the cross-sectional variation in returns, but it is crucial to get the average returns right. When we include a constant in the 2nd step of the FMB procedure, the RMSE drops to 92 basis points with only HML_{FX} as the pricing factor. Including a constant and the dollar risk factor is a problem, because the dollar factor acts like a constant in the cross-sectional regression.

Overall, the pricing errors are small. The RMSE is around 95 basis points and the adjusted R^2 is 69 percent. The null that the pricing errors are zero cannot be rejected, regardless of the estimation procedure. Figure 2 plots predicted against realized excess returns for all six currency portfolios. Clearly, the model's predicted excess returns line up rather well with the average excess returns. Note that the predicted excess return is here simply the OLS estimate of the betas times the sample mean of the factors, not the estimated prices of risk. The latter would imply an even better fit by construction. These results are robust. They also hold in a smaller sample of developed countries, as shown in the right-hand side of Table 3.

Alphas in the Carry Trade? The bottom panel of Table 3 reports the constants (denoted α^j) and the slope coefficients (denoted β^j) obtained by running time-series regressions of each portfolio's currency excess returns Rx^j on a constant and risk factors. The returns and α's are in percentage points per annum. The first column reports α's estimates. The fourth portfolio has a large α of 162 basis points per annum, significant at the 10 percent level but not statistically significant at the 5 percent level. The other α estimates are much smaller and not significantly different from zero. The null that the α's are jointly zero cannot be rejected at the 5 or 10 % significance level.

The second column of the same panel reports the estimated βs for the HML_{FX} factor. These βs increase monotonically from -.39 for the first portfolio to .61 for the last currency portfolio, and they are estimated very precisely. The first three portfolios have betas that are negative and significantly different from zero. The last two have betas that are positive and significantly different from zero. The third column shows that betas for the second factor are essentially all equal to one. Obviously, this second factor does not explain any of the variation in average excess returns across portfolios, but it helps to explain the average level of excess returns. These results are robust and comparable to the ones obtained on a sample of developed countries (reported on the right hand side of the table).

[Figure 2 about here.]

A natural question is whether these unconditional betas the bottom panel of Table 3 are driven by the covariance between exchange rate changes and risk factors, or between interest rate changes and risk factors. This is important because the conditional covariance between the log currency returns and the carry trade risk factor obviously only depends on the spot exchange rate changes:

$$cov_t\left[rx^j_{t+1}, HML_{FX,t+1}\right] = -cov_t\left[\Delta s^j_{t+1}, HML_{FX,t+1}\right].$$

In Table 4, we report the regression results of the log changes in the spot exchange rate for each portfolio on the factors. These conditional betas are almost identical to the unconditional ones (with a minus sign), as expected. Low interest currencies offer a hedge against carry trade risk because they appreciate when the carry return is low, not because the interest rates on these currencies increase. High interest rate currencies expose investors to more carry risk, because they depreciate when the carry return is low, not because the interest rates on these currencies decline. This is exactly the pattern that our no-arbitrage model in section 3 delivers. Our analysis inside the model focusses on conditional betas.

[Table 4 about here.]

Principal Components as Factors Using a linear combination of the portfolio returns as factors entails linear restrictions on the α's. When the two factors HML_{FX} and RX_{FX} are orthogonal, it is easy to check that $\alpha^1 = \alpha^6$, because $\beta^6_{HML_{FX}} - \beta^1_{HML_{FX}} = 1$ by construction. In this case, the risk prices equal the factor means. This is roughly what we find in the data. Alternatively, we can use the two first principal components themselves as factors. We re-scaled these principal component coefficients to obtain zero cost investment strategies, and we use $w^c_j, j = 1, \ldots, 6$ to denote these weights. For the second component, these portfolio weights are:

$$w^c = \left[\begin{array}{cccccc} -0.757 & -0.472 & -0.479 & -0.100 & 0.203 & 1.501 \end{array}\right].$$

Since the factors are orthogonal, we know that $\sum_{j=1}^{6} w^i_j \beta^j_i = 1$ for each risk factor $i = c, d$, and hence we know that $\sum_{j=1}^{6} w^i_j \alpha^j_0 = 0$ by construction. These results are reported in Table 5. This investment strategy involves borrowing 75 cents in currencies in the first portfolio, 47 cents in the currencies in the second portfolio, etc, and finally investing \$1.50 in currencies in the last portfolio. This is a risky strategy. The risk price of the carry factor (the second principal component) is 7.42 % per annum and the risk price of the dollar factor (the first principal component) is 1.37 % per annum. The risk-adjusted return on HML_{FX} is only 35 basis points per annum. The only portfolio with a statistically significant positive risk-adjusted return is the fourth one. However, the null that the α's are jointly zero cannot be rejected. All of the statistics of fit are virtually identical to those that we obtained we when we used HML_{FX} and RX_{FX} as factors.

[Table 5 about here.]

2.3 Sorting on HML_{FX} exposure

To show that the sorting of forward discounts really does measure a currency's exposure to the risk factor, we build portfolios based on each currency's exposure to aggregate currency risk as measured by HML_{FX}. For each date t, we first regress each currency i log excess return rx^i on a constant and HML_{FX} using a 36-month rolling window that ends in period $t-1$. This gives us currency i's exposure to HML_{FX}, and we denote it $\beta_t^{i,HML}$. Note that it only uses information available at date t. We then sort currencies into six groups at time t based on these slope coefficients $\beta_t^{i,HML}$. Portfolio 1 contains currencies with the lowest βs. Portfolio 6 contains currencies with the highest βs. Table 6 reports summary statistics on these portfolios. We do not take into account bid-ask spreads here, because it is not obvious a priori when the investor wants to go long or short. The first panel reports average changes in exchange rates. The second panel shows that average forward discounts increase monotonically from portfolio 1 to portfolio 6. Thus, sorts based on forward discounts and sorts based on betas are clearly related, which implies that the forward discounts convey information about riskiness of individual currencies. The third panel reports the average log excess returns. They are monotonically increasing from the first to the last portfolio. Clearly, currencies that covary more with our risk factor - and are thus riskier - provide higher excess returns. The last panel reports the post-formation betas. They vary monotonically from $-.31$ to $.38$. This finding is quite robust. When we estimate betas using a 12-month rolling window, we also obtain a 300 basis point spread between the first and the last portfolio.

[Table 6 about here.]

2.4 Robustness

We conducted several other robustness checks that are not reported in the paper, but are avaialble in a separate appendix on the authors' web sites.

We check the Euler equation of foreign investors in the UK, Japan and Switzerland. We construct the new asset pricing factors (HML_{FX} and RX) in local currencies, and we use the local currency returns as test assets. Note that HML_{FX} is essentially the same risk factor in all currencies, if we abstract from bid-ask spreads. Our initial spot and forward rates are quoted in US dollars. In order to convert these quotes into pounds, yen and Swiss francs, we use the corresponding midpoint quotes of these currencies against the US dollar. The correlation of HML_{FX} across different base currencies is above .95 in all cases. In fact, without the bid-ask spreads, HML_{FX} is identical across base currencies. For all countries, the estimated market price of HML_{FX} risk is less than 70 basis points removed from the sample mean of the factor. The HML_{FX} risk price is estimated at 5.54 percent in the UK, 5.50 percent in Japan and 5.79 percent in Switzerland. These estimates are statistically different from zero in all three cases. The two currency factors explain between 47 and 71 percent of the variation (after adjusting for degrees of freedom). The mean squared pricing error is 95 basis points for the UK, 116 basis points for Japan and 81 basis points for Switzerland. The null that the underlying pricing errors are zero cannot be rejected except for Japan, for which the p-values are smaller than 10 percent.

We conducted several additional robustness checks that we describe here succinctly. First, we consider the sample proposed by Burnside et al. (2008b). Following the methodology of Lustig and Verdelhan (2007), Burnside et al. (2008b) build 5 currency portfolios and argue that these currency excess returns bear no relation to their riskiness. In their data, we show that the average excess returns on these portfolios are explained by the carry trade and dollar risk factors. The α's are smaller than 60 basis points per annum, but the high-minus-low return yields 6.3 percent per annum in their sample (without bid-ask spreads). Second, we checked our results on portfolio returns from the perspective of foreign investors. Third, we divided our main sample into two subsamples, starting in 1983 and in 1995. Fourth, we considered the longer sample of currency excess returns built using Treasury bills in Lustig and Verdelhan (2007). All these results confirm that currency excess returns are large and that they are well explained by the portfolios' covariances with these risk factors.

3 A No-Arbitrage Model of Exchange Rates

In order to interpret these findings, we use a standard no-arbitrage model of exchange rates. We show that HML_{FX}, the factor that we construct by building currency portfolios, measures the common innovation to the stochastic discount factors (henceforth SDFs). Similarly, RX measures the dollar-specific innovation to the SDF of U.S. investors. In addition, we show how sorting currencies based on interest rates is equivalent to sorting these currencies on their exposure to the global risk factor. We derive conditions on stochastic discount factors at home and abroad that need to be satisfied in order to produce a carry trade risk premium that is explained by HML_{FX}.

Our model falls in the essentially-affine class and therefore shares some features with the models proposed by Frachot (1996) and Brennan and Xia (2006), as well as Backus et al. (2001). Like these authors, we do not specify a full economy complete with preferences and technologies; instead we posit a law of motion for the SDFs directly. We consider a world with N countries and currencies. Following Backus et al. (2001), we assume that in each country i, the logarithm of the SDF m^i

follows a two-factor Cox, Ingersoll and Ross (1985)-type process:

$$-m_{t+1}^i = \lambda^i z_t^i + \sqrt{\gamma^i z_t^i} u_{t+1}^i + \tau^i z_t^w + \sqrt{\delta^i z_t^w} u_{t+1}^w.$$

There is a common global factor z_t^w and a country-specific factor z_t^i. The currency-specific innovations u_{t+1}^i and global innovations u_{t+1}^w are *i.i.d* gaussian, with zero mean and unit variance; u_{t+1}^w is a world shock, common across countries, while u_{t+1}^i is country-specific. The country-specific volatility component is governed by a square root process:

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

where the innovations v_{t+1}^i are uncorrelated across countries, *i.i.d* gaussian, with zero mean and unit variance. The world volatility component is also governed by a square root process:

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

where the innovations v_{t+1}^w are also *i.i.d* gaussian, with zero mean and unit variance. In this model, the conditional market price of risk has a domestic component $\sqrt{\gamma^i z_t^i}$ and a global component $\sqrt{\delta^i z_t^w}$.[4] Brandt, Cochrane and Santa-Clara (2006) and Colacito and Croce (2008) emphasize the importance of a large common component in stochastic discount factors to make sense of the high volatility of SDF's and the 'low' volatility of exchange rates. In addition, there is a lot evidence that much of the stock return predictability around the world is driven by variation in the global risk price (Ferson and Harvey (1993)).

A major difference between our model and that proposed by Backus et al. (2001) is that we allow the loadings δ^i on the common component to differ across currencies. This will turn out to be critically important.

Complete Markets 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:

$$\Delta q_{t+1}^i = m_{t+1} - m_{t+1}^i,$$

[4]The real interest rate investors earn on currency i is given by:

$$r_t^i = \left(\lambda - \frac{1}{2}\gamma\right) z_t^i + \left(\tau - \frac{1}{2}\delta^i\right) z_t^w.$$

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. For the home country (the US), we drop the superscript. The expected excess return in levels (i.e. corrected for the Jensen term) consists of two components:

$$E_t[rx^i_{t+1}] + \frac{1}{2}Var_t[rx^i_{t+1}] \quad = \quad \sqrt{\delta^i}\left(\sqrt{\delta} - \sqrt{\delta^i}\right) z^w_t + \gamma z_t.$$

The risk premium has a global and a dollar component. $\left(\sqrt{\delta} - \sqrt{\delta^i}\right)$ is the beta of the return on currency i w.r.t. the common shock, and z^w_t is the risk price. The beta w.r.t. the dollar shock is one for all currencies, and z_t is the risk price for dollar shocks. So, the expected return on currency i has a simple beta representation: $E_t[rx^i_{t+1}] + \frac{1}{2}Var_t[rx^i_{t+1}] = \beta^i\lambda_t$ with $\beta^i = \delta^i \left[\delta^i(\sqrt{\delta} - \sqrt{\delta^i}), 1\right]$ and $\lambda_t = [z^w_t, \gamma z_t]'$. The risk premium is *independent* of the foreign country-specific factor z^i_t and the foreign country-specific loading γ^i.[5] Hence, we need asymmetric loadings on the common component as a source of variation across currencies. While asymmetric loadings on the country-specific component can explain the negative UIP slope coefficients in time series regression (as Backus et al. (2001) show), these asymmetries cannot account for any variation in risk premia across different currencies. As a consequence, and in order to simplify the analysis, we impose more symmetry on the model with the following assumption:

Assumption. *All countries share the same loading on the domestic component γ. The home country has the average loading on the global component δ: $\sqrt{\delta} = \overline{\sqrt{\delta}}$.*

3.1 Building Currency Portfolios to Extract Factors

As in the data, we sort currencies into portfolios based on their forward discounts. We use H to denote the set of currencies in the last portfolio and L to denote the currencies in the first portfolio. The carry trade risk factor HML_{FX} and the dollar risk factor $\overline{rx}$ are defined as follows:

$$\begin{aligned} hml_{t+1} &= \frac{1}{N_H}\sum_{i\in H} rx^i_{t+1} - \frac{1}{N_L}\sum_{i\in L} rx^i_{t+1}, \\ \overline{rx}_{t+1} &= \frac{1}{N}\sum_{i} rx^i_{t+1}, \end{aligned}$$

[5] The expected log currency excess return does depend on the foreign factor; it equals the interest rate difference plus the expected rate of appreciation:

$$\begin{aligned} E_t[rx^i_{t+1}] &= -E_t[\Delta q^i_{t+1}] + r^i_t - r_t, \\ &= \frac{1}{2}[\gamma z_t - \gamma^i z^i_t + \left(\delta - \delta^i\right) z^w_t]. \end{aligned}$$

where lower letters denote logs. We let $\overline{\sqrt{\delta_t^j}}$ denote the average $\sqrt{\delta^i}$ of all currencies (indexed by i) in portfolio j. Note that the portfolio composition changes over time, and in particular, it depends on the global risk price z_t^w.

In this setting, the carry trade and dollar risk factors have a very natural interpretation. The first one measures the common innovation, while the second one measures the country-specific innovation. In order to show this result, we appeal to the law of large numbers, and we assume that the country-specific shocks average out within each portfolio.

Proposition. *The innovation to the HML_{FX} risk factor only measures exposure to the common factor u_{t+1}^w, and the innovation to the dollar risk factor only measures exposure to the country-specific factor u_{t+1}:*

$$
\begin{aligned}
hml_{t+1} - E_t[hml_{t+1}] &= \left(\overline{\sqrt{\delta_t^L}} - \overline{\sqrt{\delta_t^H}}\right)\sqrt{z_t^w}u_{t+1}^w, \\
\overline{rx}_{t+1} - E_t[\overline{rx}_{t+1}] &= \sqrt{\gamma}\sqrt{z_t}u_{t+1}.
\end{aligned}
$$

When currencies share the same loading on the common component, there is no HML_{FX} risk factor. This is the case considered by Backus et al. (2001). However, if lower interest rate currencies have different exposure to the common volatility factor - $\overline{\sqrt{\delta^L}} \neq \overline{\sqrt{\delta^H}}$ - then the innovation to HML_{FX} measures the common innovation to the SDF. As a result, the return on the zero-cost strategy HML_{FX} measures the stochastic discount factors' exposure to the common shock u_{t+1}^w.

Proposition. *The HML_{FX} betas and the RX_{FX} betas of the returns on currency portfolio j:*

$$
\begin{aligned}
\beta_{hml,t}^j &= \frac{\sqrt{\delta} - \overline{\sqrt{\delta_t^j}}}{\overline{\sqrt{\delta_t^L}} - \overline{\sqrt{\delta_t^H}}}, \\
\beta_{rx,t}^j &= 1.
\end{aligned}
$$

The betas for the dollar factor are all one. Not so for the carry trade risk factor. If the sorting of currencies on interest rate produces a monotonic ranking of δ , then the HML_{FX} betas will increase monotonically as we go from low to high interest rate portfolios. As it turns out the model with asymmetric loadings automatically delivers this if interest rates decrease when global risk decreases. This case is summarized in the following condition:

Condition.

$$
0 < \tau < \frac{1}{2}\delta^i.
$$

The real short rate depends both on country-specific factors and on a global factor. The only sources of cross-sectional variation in interest rates are the shocks to the country-specific factor z_t^i, and the heterogeneity in the SDF loadings δ^i on the world factor z^w. As a result, as z^w increases, on average, the currencies with the high loadings δ will tend to end up in the lowest interest rate portfolios, and the gap $\left(\sqrt{\overline{\delta_t^L}} - \sqrt{\overline{\delta_t^H}}\right)$ increases. This implies that in bad times the spread in the loadings increases. In section 5, we provide a calibrated version of the model that illustrates these effects.

As shown above, in our model economy, the currency portfolios recover the two factors that drive innovations in the pricing kernel. Therefore, these two factors together do span the mean-variance efficient portfolio, and it comes as no surprise that these two factors can explain the cross-sectional variation in average currency returns.

3.2 Risk Premia in No-Arbitrage Currency Model

In our model, the risk premium on individual currencies consists of two parts: a dollar risk premium and a carry trade risk premium. Our no-arbitrage model also delivers simple closed-form expression for these risk premia.

Proposition. *The carry trade risk premium and the dollar risk premium are:*

$$\begin{aligned} E_t[hml_{t+1}] &= \frac{1}{2}\left(\overline{\delta_t^L} - \overline{\delta_t^H}\right) z_t^w, \\ E_t[\overline{rx}_{t+1}] &= \frac{1}{2}\gamma\left(z_t - \overline{z_t}\right). \end{aligned} \tag{3.1}$$

The carry trade risk premium is driven by the global risk factor. The size of the carry trade risk premium is governed by the spread in the loadings (δ) on the common factor between low and high interest rate currencies, and by the global price of risk. When this spread doubles, the carry trade risk premium doubles. However, the spread itself also increases when the global Sharpe ratio is high. As a result, the carry trade risk premium increases non-linearly when global risk increases. The dollar risk premium is driven only by the US risk factor, if the home country's exposure to global risk factor equals to the average δ. When the home country's δ is lower than average, then the dollar risk premium also loads on the global factor:

$$rp_t^{\overline{rx}} = \frac{1}{2}\gamma\left(z_t - \overline{z_t}\right) + \frac{1}{2}\left(\delta - \overline{\delta}\right) z_t^w.$$

The risk premia on the currency portfolios have a dollar risk premium and a carry trade component:

$$rp_t^j = \frac{1}{2}\gamma\left(z_t - \overline{z_t^j}\right) + \frac{1}{2}\left(\delta - \overline{\delta^j}\right) z_t^w. \tag{3.2}$$

The first component is the dollar risk premium part. The second component is the carry trade part. The highest interest rate portfolios load more on the carry trade component, because their loadings are smaller than the home country's δ, while the lowest interest rate currencies have a negative loading on the carry trade premium, because their loadings exceed the home country's δ. Note that $\overline{z^j}$ is constant in the limit $N \to \infty$ by the law of large numbers. This means that there should be no role for portfolio-specific variables in forecasting currency excess returns. This is exactly what we find. We show in the next section that the average interest rate difference is a better predictor than the portfolio-specific one.

In addition, in a reasonably specified model, the US-specific component of the risk price, z_t, and hence the dollar risk premium, should be counter-cyclical -with respect to the US-specific component of the business cycle-, and the global component z_t^w, and hence the carry risk premium, should be counter-cyclical with respect to the global business cycle. In the next section, we show that the predicted excess returns on medium to high interest rate currencies are highly counter-cyclical, and that business cycle indices (like US industrial production growth) predict these excess returns, even after controlling for interest rate differences. We also show that the predicted excess returns on a long position in the sixth portfolio and a short position in the first portfolio are highly correlated with the VIX volatility index, one proxy of higher frequency variation in the global risk factor z_t^w.

4 Return Predictability in Currency Markets

In this section, we investigate the predictability of returns on these currency portfolios, and we show that the average forward discount across portfolios does a better job of describing the time variation in expected currency excess returns than the individual portfolio forward discounts, as implied by the no-arbitrage model. In addition, we show that these expected excess returns are closely tied to the US business cycle: expected currency returns increase in downturns and decrease in expansions, as is the case in stock and bond markets. Finally, we show that the variation in expected returns on long-short strategies are linked to higher frequency variation in global credit spreads and global market volatility.

4.1 Predictability in Portfolio Excess Returns

We first investigate the predictive power of the portfolio-specific forward discount, and then turn to the predictive power of the average forward discount.

Individual Forward Discounts For each portfolio j, we run a regression of each portfolio's average log currency excess returns on each portfolio's average log forward discounts:

$$rx_{t+1}^j = \kappa_0^j + \kappa_f^j (f_t^j - s_t^j) + \eta_t^j.$$

If UIP were an accurate description of the data, there would be no predictability in currency excess returns, and the slope coefficient κ_f would be zero. Table 7 reports regression results. We use net excess returns that take into account bid-ask spreads. Bid-ask spreads vary with time. For example, the average spread in the last portfolio increases with the volatility index VIX, but this time-variation is very small compared to the mean bid-ask spread and the mean excess return.

Portfolio forward discounts account for between 1.8 percent and 6.4 percent of the monthly variation in excess returns on these currency portfolios. There is strong evidence against UIP in these portfolio returns, more so than in individual currency returns. Looking across portfolios, from low to high interest rates, the slope coefficient κ_f^j (column 3) varies a lot: it increases from 108 basis points for currencies in the first portfolio to 357 basis points for currencies in the fourth portfolio. The slope coefficient decreases to 72 basis points for the sixth portfolio. Deviations from UIP are highest for currencies with medium to high forward discounts. However, forward rates are strongly autocorrelated. This complicates statistical inference about these slope coefficients. To deal with this issue, we use two asymptotically-valid corrections. The Newey-West standard errors (NW) are computed with the optimal number of lags following Andrews (1991). The Hansen-Hodrick standard errors (HH) are computed with one lag. Both of these methods correct for arbitrary error correlation and conditional heteroscedasticity. Bekaert, Hodrick and Marshall (1997) note that the small sample performance of these test statistics is also a source of concern. To address this problem, we also report small sample standard errors. These were generated by bootstrapping 10,000 samples of returns and forward discounts from a bivariate VAR with one lag. The null of no predictability is rejected at the 1 percent significance level for all of these portfolios except for the third. At the one-month horizon, the R^2 on these predictability regressions varies between 1.61 and 5.98 percent. In other words, when considering currency portfolios, up to 6 percent of the variation in spot rates is predictable at a one-month horizon.

Average Forward Discount There is even more predictability in these excess returns than the standard UIP regressions reveal, because forward discounts on the other currency portfolios also help to forecast returns. We found that a single return forecasting variable describes time variation in the dollar risk premium even better than the forward discount rates on the individual currency portfolios. This variable is the average of all the forward discounts across portfolios. We also examined the optimal linear combination of forward discounts along the lines of Cochrane and Piazzesi (2005). However, it does not really outperform the average forward discount as a predictor. We use ι to denote the 6×1 vector with all elements equal to $1/6$. For each portfolio j, we run the following regression of log excess returns after bid-ask spreads on the average forward rates:

$$rx^j_{net,t+1} = \kappa^j_0 + \kappa^j_{\mathbf{f}} \iota'(\mathbf{f}_t - \mathbf{s}_t) + \eta^j_t,$$

where $\mathbf{f}_t - \mathbf{s}_t$ bunches together all forward discounts. A summary of the results is reported in columns 3 and 4 of Table 7. This single factor explains between 2.68 and 7.85 percent of the variation in returns at the one-month horizon. The average forward discount outperforms the portfolio-specific forward discounts, except in portfolios 4 and 5. In this case, the slope coefficients are more stable across the different portfolios. Portfolio-specific time variation in expected exchange rate movements driven by the sorting variable (relative interest rates) does not appear to be the main driver of return predictability in currency markets. The average interest rate difference is the main driver.

[Table 7 about here.]

The right panel of Table 7 focuses on the predictability of carry trade returns: the returns on a high-minus-low strategy that goes long in high interest rate currencies and short in low interest rate currencies. We run the following predictability regression of the one-month high-minus-low return $rx^j - rx^1$ on the spread in the one-month forward discount between the j-th and the first portfolio:

$$rx^j_{t+1} - rx^1_{t+1} = \kappa_{sp,0} + \kappa_{sp,f}\left[(f^j_t - s^j_t) - (f^1_t - s^1_t)\right] + \eta^j_t.$$

There is some evidence that the high-minus-low returns are forecastable by the forward spreads, but the evidence is less strong

than on individual portfolio returns. Since the spread in forward discounts is much less persistent than the forward discount and there is no overlap in returns, there is less cause for concern about persistent regressor bias.

Longer Horizons At longer horizons, the fraction of changes in log spot rates explained by the forward discount is even greater than at short horizons. We use k-month maturity forward contracts to compute k-period horizon returns (where $k = 1,\ 2,\ 3,\ 6,\ 12$). The log excess return on the k-month contract is:

$$rx^k_{t+k} = -\Delta s_{t\to t+k} + f^k_t - s_t.$$

Then we sort the currencies into portfolios based on forward rates with the corresponding maturity, and we compute the average excess return for each portfolio. Table 8 provides a summary of the results: it lists the R^2s we obtained for each portfolio (rows) and for each forecasting horizon (columns). We only consider the corner portfolios.

At longer horizons, the returns on the first portfolio are most predictable; the returns on the last portfolio are least predictable. On the first portfolio, more than a quarter of the variation in excess returns is accounted for by the forward rate at the 12-month horizon. On the last portfolio, 10 percent is accounted for by the forward rate. One concern is that these measures of fit may be biased because we use overlapping returns and because the predictors are highly autocorrelated. In the bottom panel of Table 8 we also provide the same R^2 measures that we obtained for each forecasting horizon with non-overlapping data. To produce these measures, we simply used the first month of every period (quarter, year) to run the same regressions. Though there are some differences, these R^2s are not systematically lower. Even at longer horizons, the average forward discount seems to do a better job in describing the variation in expected excess returns. This single factor explains between 18 and 32 percent of the variation at the one-year horizon. This single factor mostly does as well and sometimes better than the forward discount of the specific portfolio in forecasting excess returns over the entire period.

[Table 8 about here.]

As a result, we conclude that the average forward discount contains information that is useful for forecasting excess returns on all currency portfolios, while little information is lost by aggregating all these forward discounts into a single predictor. The fact that the average forward discount is a *better predictor* of future excess returns on foreign currency than individual forward discount rates is consistent with the risk premium view: by using the average forward discount, we throw away all information related to country-specific inflation, and we do better in predicting future changes in exchange rates. In fact, if we take the residuals of the average forward discount forecasting regression and we project these on the individual portfolio forward discounts, there is no predictability left. In the right panel of Table 8, we also report the R^2s of these regressions. There is no information in the individual forward discounts left that helps to forecast currency returns. This finding is similar to results of Stambaugh (1988) and Cochrane and Piazzesi (2005) for the predictability of Treasury bill and bond returns. These studies show that linear combinations of forward rates across maturities outperform the forward rate of a particular maturity in forecasting returns. In particular, Cochrane and Piazzesi (2005) report R^2s of up to 40 percent on one-year holding period returns for zero coupon bonds using a single forecasting factor. Currency returns are *more predictable* than stock returns, and almost as predictable as bond returns.

Counter-Cyclical Dollar Risk Premium Our predictability results imply that expected excess returns on currency portfolios vary over time. We now show that this time variation has a large US business cycle component: expected excess returns go up in US recessions and go down in US expansions. The same counter-cyclical behavior has been documented for bond and stock excess returns. We use $\widehat{E}_t rx^j_{t+1}$ to denote the forecast of the one-month-ahead excess return based on the forward discount:

$$\widehat{E}_t rx^j_{t+1} = \kappa^j_0 + \kappa^j_f (f^j_t - s^j_t).$$

At high frequencies, forecasted returns on high interest rate currency portfolios – especially for the sixth portfolio – increase very strongly in response to events like the Asian crisis in 1997 and the LTCM crisis in 1998, but at lower frequencies, a big fraction of the variation in forecasted excess returns is driven by the US business cycle, especially for the third, fourth and fifth portfolios. To assess the cyclicality of these forecasted excess returns, we use three standard business cycle indicators and three financial variables: (i) the 12-month percentage change in US industrial production index, (ii) the 12-month percentage change in total US non-farm payroll index, (iii) the 12-month percentage change in the Help Wanted index, (iv) the default spread – the difference between the 20-Year Government Bond Yield and the $S\&P$ 15-year BBB Utility Bond Yield – (v) the slope of the yield curve – the difference between the 5-year and the 1-year zero coupon yield on Treasuries, and (vi) the $S\&P$ 500 VIX volatility index.[6] 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 to the ones reported in this paper.

[Table 9 about here.]

Table 9 reports the contemporaneous correlation of the month-ahead forecasted excess returns with these macroeconomic and financial variables. As expected, forecasted excess returns for high interest rate portfolios are strongly counter-cyclical.

On the one hand, the monthly contemporaneous correlation between predicted excess returns and percentage changes in industrial production (first column), the non-farm payroll (second column) and the help wanted index (third column) are negative for all portfolios except the first one. For payroll changes, the correlations range from -.70 for the second portfolio to -.09. for the sixth. Figure 3 plots the forecasted excess return on portfolio 2 against the 12-month change in US industrial production. Forecasted excess returns on the other portfolios have similar low frequency dynamics, but in the case of portfolios 5 and 6, they also respond to other events, like the Russian default and LTCM crisis, the Asian currency crisis and the Argentine default.

[Figure 3 about here.]

[6]Industrial production data are from the IMF International Financial Statistics. The payroll index is from the BEA. The Help Wanted Index is from the Conference Board. Zero coupon yields are computed from the Fama-Bliss series available from CRSP. These can be downloaded from `http://wrds.wharton.upenn.edu`. Payroll data can be downloaded from `http://www.bea.gov.` The VIX index, the corporate bond yield and the 20-year government bond yield are from `http://www.globalfinancialdata.com.`

On the other hand, monthly correlations of the high interest rate currency portfolio with the default spread (fourth column) and the term spread (fifth column) are, as expected, positive. Finally, the last column reports correlations with the implied volatility index (VIX). The VIX seems like a good proxy for the global risk factor. The VIX is highly correlated with similar volatility indices abroad.[7] The correlations in the last column reveal a clear difference between the low interest rate currencies with negative correlations, and the high interest rate currencies, with positive correlations. This is consistent with the predictions of our no-arbitrage model. Recall that the model predicts negative loadings on the common risk factor for the risk premia on low interest rate currencies and positive loadings for the risk premia on high interest rate currencies (see equation 3.2). In times of global market uncertainty, there is a flight to quality: investors demand a much higher risk premium for investing in high interest rate currencies, and they accept lower (or more negative) risk premia on low interest rate currencies.

Longer Horizons We find the same business cycle variation in expected returns over longer holding periods. The predictability is partly due to the counter-cyclical nature of the forward discount, but not entirely. Controlling for the forward discount reduces the IP slope coefficient by 50 basis points on portfolios 1-4, 20-30 basis points for portfolios 5-6, but the forward discount does not drive out the macroeconomic variable. Table 10 reports forecasting results for currency portfolios obtained using the 12-month change in industrial production and either the portfolio-specific forward discount or the average forward discount. The currency risk premium increase in response to a one percentage point drop in the growth rate of industrial production varies between 90 (portfolio 1) and 170 basis points (portfolio 5). The IP slope coefficients are still significantly different from zero for the high interest rate portfolios, but the slope coefficients on the (average) forward discounts are not. In recent work, Duffee (2008) and Ludvigson and Ng (2005) report a similar finding for the bond market, while Piazzesi and Swanson (2008) find that the annual growth rate of the non-farm payroll predicts excess returns on interest rate futures.

[Table 10 about here.]

We have documented in this section that returns in currency markets are highly predictable. The average forward discount rate accurately predicts up to 33 percent of the variation in annual excess returns. The time variation in expected returns has a clear business cycle pattern: US macroeconomic variables are powerful predictors of these returns, especially at longer holding periods, and expected currency returns are strongly counter-cyclical. We now turn to the behavior of the second moments of currency returns over time.

[7] The VIX starts in February 1990. The DAX equivalent starts in February 1992; the SMI in February 1999; the CAC, BEL and AEX indices start in January 2000. Using the longest sample available for each index, the correlation coefficients with the VIX are very high, respectively 0.85, 0.82, 0.88, 0.83 and 0.82 using monthly time-series.

4.2 Time-Varying Risk of the Carry Trade

In this section, we show that while the average beta of HML_{FX} with the US stock market return is too small to explain carry trade risk premia, the conditional market beta varies a lot through time, and is particularly high during episodes of global financial crises.

We run the same asset pricing experiment on the cross-section of currency excess returns using the US stock market excess return as the pricing factor, instead of the slope risk factor HML_{FX}. To measure the return on the market, we use the CRSP value-weighted return on the NYSE, AMEX and NASDAQ markets in excess of the one-month average Fama risk-free rate. The US stock market excess return and the level factor RX can explain 52 percent of the variation in returns. However, the estimated price of US market risk is 37 percent, while the actual annualized excess return on the market is only 7.1 percent over this sample. The risk price is 5 times too large. The CAPM betas vary monotonically from -.05 for the first portfolio to .08 for the last one. Low interest rate currencies provide a hedge, while high interest rate currencies expose US investors to more stock market risk. These betas increase almost monotonically from low to high interest rates, but they are too small to explain these excess returns. Therefore, the cross-sectional regression of currency returns on market betas implies market price of risk that are far too high. The null that that the α's are zero is rejected at the 5 % significance level.[8]

Despite the low unconditional market beta of the carry trade, the carry risk factor HML_{FX} is very highly correlated with the stock market during periods of increased market volatility. The recent subprime mortgage crisis offers a good example. A typical currency carry trade at the start of July 2007 was to borrow in yen - a low interest rate currency - and invest in Australian and New Zealand dollars - high interest rate currencies. Over the course of the summer, each large drop in the S&P 500 was accompanied by a large appreciation of the yen of up to 1.7 percent in one day and a large depreciation of the New Zealand and Australian dollar of up to 2.3 percent in one day. Figure 4 plots the monthly returns on HML_FX at daily frequencies against the US stock market return. Clearly, a US investor who was long in these high interest rate currencies and short in low interest rate currencies, was heavily exposed to US aggregate stock market risk during the subprime mortgage crisis, and thus should have been compensated by a risk premium ex ante.

[Figure 4 about here.]

This pattern is consistent with the model. In the two-factor affine model, the conditional correlation of HML_{FX} and the SDF in the home country is:

$$corr_t\left(hml_{t+1}, m_{t+1}\right) = \frac{\sqrt{\delta z_t^w}}{\sqrt{\delta z_t^w} + \sqrt{\gamma z_t}}. \tag{4.1}$$

[8]Detailed results available upon request

As the global component of the conditional market price of risk increases, the conditional correlation between the stochastic discount factor at home and the carry trade returns HML_{FX} increases. We find strong evidence for this type of time-varying correlation in the data.

As a first pass, we use the US stock market return as a proxy for the domestic SDF. We compute the correlation between one-month currency returns and the return on the value-weighted US stock market return using 12-month rolling windows on daily data. Figure 5 plots the difference between the correlation of the 6th and the 1st portfolio with the US stock market excess return. We denote it $Corr_\tau[R_t^m, rx_t^6] - Corr_\tau[R_t^m, rx_t^1]$, where $Corr_\tau$ is the sample correlation over the previous 12 months $[\tau - 12, \tau]$ and R^m, the stock market excess return. We also plot the stock market beta of HML_{FX}. These market correlations exhibit enormous variation. In times of crisis and during US recessions, the difference in market correlation between high and low currencies increases significantly. During the Mexican, Asian, Russian and Argentinean crises, the correlation difference jumps up by 50 to 90 basis points.

[Figure 5 about here.]

We now explore time-variation in market betas. There is evidence that, in times of financial crisis, the stock market beta of the high-minus-low strategy in currency markets increases dramatically. We start by examining the recent sub-prime mortgage crisis, and we then consider other crisis episodes. The last 4 columns of Table 11 reports the market betas of all the currency portfolios that we obtain on a 6-month window before 08/31/2007. To estimate the market betas, we use daily observations on monthly currency and stock market returns. The Newey-West standard error correction is computed with 20 lags. We estimate a market beta of HML_{FX} of up to 62 basis points. The estimated market betas increase monotonically as we move from low to high interest rate currency portfolios, as we would expect. We report the αs in the bottom panel of Table 11. Over this period, the estimated pricing errors α on the high-minus-low strategy dropped to 30 basis points over 6 months or 60 basis points per annum compared to an unconditional pricing error α_{HML} of more than 500 basis points per annum.

This is not an isolated event, as these results extend to other crises. In Table 11, we document similar increases in the US market beta of HML_{FX} during the LTCM crisis (column 1-4), the Mexican "Tequila" crisis (column 5-8) and the Brazilian/Argentine crisis (column 9-12). Again, the market betas increase monotonically in the forward discount rates. For example, $\beta^m_{\tau,HML}$ increases to 1.14 in the run-up to the Russian default in 1998, implying that high interest rate currencies depreciate on average by 1.14 percent relative to low interest rate currencies when the stock market goes down by one percent. Low interest rate currencies provide a hedge against market risk while high interest rate currencies expose US investors to more market risk in times of crisis. For the Tequila crisis, the market betas of all the currency portfolios are negative. This is consistent with our model, as the dollar risk premium component is counter-cyclical with respect to the US business cycle, and hence the expected returns on all portfolios can be negative (see equation 3.2). In two of these crisis, the α on the high-minus-low strategy is negative: minus 271 basis over the 6 months preceding the Russian default and minus 382 basis points during the Tequila crisis.[9] In the two other crisis, the αs are positive (96 and 29 basis points over 6 months respectively) but small, well below the average α of 4.46 percent per annum that we obtained over the entire sample. As we have shown, the market beta of the high-minus-low strategy increases dramatically in times when the price of global risk is high.

[Table 11 about here.]

5 Calibrated Model

A reasonably calibrated version of the no-arbitrage model in section 3 can match the key moments of currency returns in the data. We calibrate the model at monthly frequency by targeting annualized moments of monthly data. In this calibration, we focus on developed countries over the 1983-2008 sample. The calibration proceeds in two stages. In a first stage, we calibrate the real side of the model by targeting moments of the real variables. In the second stage, we turn to the nominal SDFs by matching some moments of inflation.

5.1 Calibration

We start by calibrating a completely symmetric version of the model, and then we introduce enough heterogeneity in the SDF loadings on the global shock across countries to match the carry trade risk premium. There are 7 parameters in the model: 4 parameters govern the countries' SDFs (λ, γ, τ and δ) and 3 parameters describe the evolution of the country-specific and global state variables. We choose these parameters to match 7 key moments in the data: the mean, standard deviation and autocorrelation of real risk-free rates, the average conditional variance of changes in real exchange rates, the mean and standard deviation of the maximal conditional Sharpe ratio and the UIP slope coefficient. Panel I of Table 12 lists all of these moments and panel II lists the parameter choices. These moments were generated by drawing $10,000$ observations from a model with 40 currencies. The simulated model produces a real risk-free rate with a mean of 1.2 percent, a standard deviation of 0.2 percent and an autocorrelation of 0.7 (on annual basis). The mean and autocorrelation of the real interest rate fall within the range of empirical estimates for the post-war U.S. data (e.g. Campbell (2003)). The average (annualized) standard deviation of real exchange rates is about 10 percent and the average regression coefficient of exchange rate changes on interest rate differentials is around -1, roughly consistent with our data. Our model produces

[9]These numbers need to be multiplied by 12 to be annualized.

an average conditional maximum Sharpe ratio for the domestic investor with a mean of 0.32 and a standard deviation of 0.04, in annual units.

Next, we set the heterogeneity in the loadings on the common risk factor by choosing the range of parameters δ^i to match the mean of the carry trade risk factor. Setting the range for countries' global risk loadings to be within 40 percent of the home country's loading leads to a mean carry factor HML_{FX} of 5 percent per annum, which is broadly consistent with our empirical results for developed countries. Expanding the range of global shock loadings allows us to match the higher average return obtained using all countries in our sample, but also increases the average exchange rate volatility.

[Table 12 about here.]

We add inflation to the model in order to match moments of nominal interest rates and exchange rates. The log of the nominal pricing kernel in country i is simply given by the real pricing kernel less the rate of inflation π^i:

$$m_{t+1}^{i,\$} = m_{t+1}^i - \pi_{t+1}^i.$$

We assume that inflation is composed of a country-specific component and a global component. The bottom panel of Table 12 lists the moments of inflation processes used in calibration; the details of the calibration are in the appendix.

The calibrated version of our multi-country model delivers reasonable interest rates and exchange rates. The annualized average real one-period yields are between -1 and 10 percent. The mean nominal one-period interest rate is between 2 and 6 percent, with an average 4.2 across countries, and average standard deviation of about 2 percent. The annualized standard deviations of changes in the real and nominal exchanges rates are between 9 and 15 percent.

Since we want to test the CAPM on model-generated data, we also need to add stocks to our model. We define country i's total stock market portfolio as a claim to the aggregate dividend stream of that country, D_t^i. We model each country's dividend process as a random walk with a drift for the logarithm $d_t^i = \log D_t^i$:

$$\Delta d_{t+1}^i = d_{t+1}^i - d_t^i = g^{Di} + \sigma^{Di} w_{t+1}^{Di},$$

where the w^D innovations are i.i.d. and normally distributed. In order to command a risk premium, the dividend growth innovations must be correlated with the SDF. In particular, we specify the conditional correlations of the dividend growth process with both the world and country-specific innovations to the SDF:

$$\rho^{Dw} = corr\left(w^{Di}, u^w\right) \text{ and } \rho^{Di} = corr\left(w^{Di}, u^i\right).$$

We choose the standard deviation of log dividend growth σ^{Di} to be 10 percent per annum, and we simply choose the correlations with the two SDF shocks $\rho^{Dw} = \rho^{Di} = 0.7$, since it is *a priori* reasonable that aggregate dividends are equally affected by global and country-specific shocks. The resulting stock market return process has an empirically plausible annualized monthly Sharpe ratio of 0.29, although both the equity premium and the stock return volatility are low, at about 3 percent and just over 10 percent, respectively. This is because the amount of variation in expected stock returns generated by the model is too small.

5.2 Simulated Currency Portfolios

We simulate a version of the model with $N = 40$ countries over 10,000 periods. We build currency portfolios starting from the simulated data in the same way as for the actual data. Table 13 reports summary statistics on these portfolios and estimates of the market prices of risk associated with the two factors, RX and HML_{FX}. The model delivers a sizable cross-section of currency excess returns. The spread between the first and last portfolio is 5 percent per annum, implying an annualized Sharpe ratio of 0.44. In the cross-sectional asset pricing tests, the market price of the carry trade factor HML_{FX} is 5 percent per annum, very close to the sample mean. The price of the aggregate market return RX is close to zero and not statistically significant. This is not surprising; with a large number of periods, the mean of RX should be zero according to equation (3.1) as long as the home country SDF has an "average" loading on the global risk factor. At the same time, due to the cross-sectional heterogeneity in the loadings on the world risk factor, our model is able to reproduce the variation in average returns on currency portfolios, and in particular the large average return on the carry trade factor.

[Table 13 about here.]

The simulated market price of carry risk varies for two reasons. First, it is high when the world risk factor z^w is high. Second, this effect is amplified by changes in portfolio composition: higher world risk price drives the selection of low-global risk countries into high interest rate portfolios, and vice versa. Thus, in "bad times," when z^w is high, the spread between the average δ in the first and the last portfolio increases. Figure 6 illustrates this second effect.

[Figure 6 about here.]

Using the simulated return on the stock market portfolio we can also show that the CAPM fails to explain currency return generated by our model, as in the data. In a sample of 5000 simulated periods, we run a time-series regression of HML_{FX} on the stock market return. We find that the CAPM α of HML_{FX} is large and statistically significantly different from zero: the CAPM understates the average return by over 3 percent (with very little statistical uncertainty given the large size of the simulated sample). This large CAPM α represents the bulk of the average HML_{FX} return. As a result, the CAPM cannot explain currency returns in this no-arbitrage model of exchange rates, even though the stock market wealth is priced using the same stochastic discount factor that prices currencies. The average stock market beta of the carry trade is somewhat higher than in the data, at about .6. This is in part because the model understates the stock market volatility.

Both the betas and the correlations of the currency portfolio returns with the stock market return exhibit a lot of variation over time, due to the fact that time-varying prices of risk imply time-varying conditional correlations of portfolio returns with the stochastic discount factor. Figure 7 plots the conditional betas and correlations of the carry factor returns with the stock market return (Panel A) as well as the realized volatility of the stock market return (Panels B), both computed using 12-month rolling windows, as used when estimation these quantities in the data. The periods of high global risk and, consequently, high stock market volatility correspond to a greater spread in correlations/betas of currency portfolios with the stock market return. Conditional market beta of HML_{FX} varies between close to zero in times of low volatility to well above one during episodes of spiking uncertainty. Thus, in our model the stock market risk of the carry trade varies over time in a manner consistent with the empirical evidence documented in Section 4.2.

[Figure 7 about here.]

6 Covariances or Characteristics

We conclude by providing additional evidence for the importance of common risk factors in currency returns. One potential concern is that by sorting currencies into portfolios based on interest rates, we might be picking up the effects of the characteristics of currencies rather than the true exposure to risk (Daniel and Titman (2005)). In fact, Bansal and Dahlquist (2000) argue that individual currency characteristics, not risk exposures, can account for the cross-sectional variation in currency returns. To address this concern, we exploit a different source of variation in currency returns: momentum. We show that the carry risk factor can account for at least 50 % of the cross-sectional variation in momentum-driven currency returns, even though the momentum portfolios are constructed on the basis of past returns, not interest rate differentials.

6.1 Momentum in Currency Returns

Table 14 reports the momentum returns. The momentum portfolios are constructed by sorting currencies at time t into portfolio based on one-month returns realized at the end of period $t-1$.

We chose not to double-sort on interest rates and past returns, because of the limited number of currencies we have, so there is some overlap between the carry and momentum portfolios, but it does not appear to be the sole driver of momentum. The low momentum currencies tend to depreciate at an annualized rate of 4.48 %, while the high momentum tend to appreciate at an annualized rate 1.92 %. The rate of appreciation varies monotonically from low to high momentum portfolios. For the carry portfolios, there was no such pattern in portfolio-by-portfolio exchange rates. The high momentum portfolios do tend to have higher interest rates, but the spread between the lowest and the highest momentum portfolio is less than 300 basis points on average, much smaller than the spread between low and high interest rate portfolios of more than 11 percentage points that we reported in Table 1. Since high momentum currencies tend to have higher interest rates on average, there is a concern that these are too similar to the carry portfolios to provide an independent sort. However, momentum and carry strategies are very different. In fact, the return correlations between corresponding (i.e. high/high or low/low) carry and momentum strategies are small and sometimes even negative. These correlations are reported in the separate appendix.

The momentum strategy in currencies produces an impressive 9.32 % return before transaction costs. However, high momentum currencies tend to have larger bid/ask spreads. The annual excess return drops to 5.42 % per annum after accounting for transaction costs. The annualized Sharpe ratio for this momentum strategy is .5. Both the average excess returns and the Sharpe ratios increase monotonically from low to high momentum portfolios.

[Table 14 about here.]

The first principal component of these 12 portfolios is clearly the dollar risk factor. It accounts for 67 % of the time-series variation in returns on all of the 12 currency portfolios. The second principal component is clearly a momentum factor. This represents an investment strategy that shorts low momentum and goes long in high momentum portfolios. The portfolio weights are given by:

$$w^m = \begin{bmatrix} 1.54 & 2.17 & 0.78 & 2.34 & 1.49 & -4.31 & -11.33 & -1.37 & -0.11 & 2.12 & 3.32 & 4.34 \end{bmatrix}.$$

However, the most interesting one from an asset pricing perspective is the third one. This component represents an investment strategy that goes long in high momentum and long in high interest rate currencies, while shorting low momentum and low interest rate currencies. The portfolio weights are given by:

$$w^c = \begin{bmatrix} 6.01 & 3.55 & 4.76 & 0.05 & -2.80 & -10.35 & 0.46 & 5.89 & 5.15 & 0.45 & -4.32 & -7.84 \end{bmatrix}.$$

The portfolio weights increase monotonically from low to high interest rates and from low to high momentum (except for portfolio 7). Not surprisingly, this third principal component is highly correlated with HML_{FX} (.75) and with the second principal component of the carry portfolios (.80). For each of 12 principal components, Figure 8 plots the covariance of excess returns with that principal component against the average excess returns. The first principal component is a flat line. The second principal component does not covary with the carry portfolios in the right way. However, the third principal component clearly does, and it is the only one, as is apparent from the other 11 subplots. This suggests that the carry risk factor that we identify has explanatory power for other currency portfolios not sorted on interest rates. We confirm this by estimating a linear factor model on the cross-section of currency returns.

[Figure 8 about here.]

In Table 15, we estimate a linear factor model. In the first subpanel, we use all 12 portfolios as test assets. In the second subpanel, we use only the 6 carry portfolios as test assets. In the third subpanel, we use the 6 momentum portfolios as test assets. In the baseline version, the three factors are the first three principal components: the dollar factor (denoted d, the first principal component), the momentum factor (denoted m, the second principal component) and the carry factor (denoted c, the third principal component). These results are reported in the left panel. We rescaled the principal component weights so they to sum to one, hence these three factors represent excess returns on zero cost investment strategies in these 12 currency portfolios. The panel on the right uses only two factors, dropping momentum.

With all 12 test assets and 3 risk factors, the risk prices of the factors equal their means. The risk price of the carry factor is precisely estimated and highly statistically significant. The adjusted R^2 of this three-factor model is .83 and the $RMSE$ is 70 basis points. The momentum portfolios load significantly on the carry factor, and the betas vary from -.18 on portfolio 8 to .24 on portfolio 12 (high momentum). Most of these betas (not reported) are highly statistically significant. The variation in these carry risk betas is enough to account for most of the variation in returns on the momentum portfolios. In the right panel of Table 15, we report estimates of a two-factor model, without the momentum factor, on the same 12 test assets (6 carry and 6 momentum currency portfolios). The carry factor can actually account for most of the cross-sectional variation in returns on the momentum portfolios, except for the lowest momentum portfolio (portfolio 7): in the two-factor model, the $RMSE$ increases to 1.02 basis points and the adjusted R^2 drops to .68. The null that the α's are jointly zero in the two-factor model cannot be rejected at standard significance levels.

[Table 15 about here.]

We also report estimates of the three-factor and the two-factor model that are obtained by using the carry portfolios and the momentum portfolios separately. The second subpanel in Table

15 uses only the carry portfolios as test assets. The bottom subpanel uses only the momentum portfolios as test assets. So, these estimates only use 6 test assets. On the carry portfolios, the carry factor that we construct from these 12 currency portfolios actually does slightly *better* than the carry factor constructed from the 6 carry portfolios. The results in the right panel can be directly compared to Table 5. The RMSE drops from 96 to 84 basis points and the adjusted R^2 increases from .7 to .76. So, bringing information from momentum portfolios to bear actually improves the fit. Also, the estimated price of carry risk is 9.96 % per annum, very close to its mean of 10 % per annum. However, the key finding is in the bottom panel, when we only use the momentum portfolios as test assets: the carry risk factor is statistically significant even when controlling for the momentum factor. In fact, when we drop the momentum factor, the model still explains half of the cross-sectional variation in momentum returns (in adjusted R^2). The risk prices that we estimate in the two-factor model are almost invariant to the test assets: in both cases, the price of carry risk is estimated to be 10 % per annum. This evidence is a major challenge for the characteristics-based explanation because it shows that covariances with this carry trade risk factor line up with returns on portfolios that are sorted on an different characteristic.

7 Conclusion

In this paper, we show that currency markets offer large and time-varying risk premia. Currency excess returns are highly predictable. In addition, these predicted returns are strongly countercyclical. The average excess returns on low interest rate currencies are about 5 percent per annum smaller than those on high interest rate currencies after accounting for transaction costs. We show that a single return-based factor explains the cross-section of average currency excess returns. These findings are consistent with the notion that carry trade profits are compensation for systematic risk.

Using a no-arbitrage model of exchange rates, we show that a single risk factor, obtained as the return on the highest minus the return on the lowest interest rate currency portfolio, measures exposure to common or global shocks to investors' marginal utilities/stochastic discount factors. We can replicate our main empirical findings in a reasonably calibrated version of this model, provided that low interest rate currencies are more exposed to global risk in bad times, when the price of global risk is high. This heterogeneity in the loadings on the global risk factor is critical for explaining the cross section of currency returns.

We conclude that currency excess returns reflect risk premia for exposure to aggregate risk. Identifying the economic mechanism that drives the relationship between macroeconomic risk and asset prices is therefore key to understanding the dynamics of currency markets. Heterogeneity in a country's exposure to global risk factors can be driven by the differences in preferences (risk aversion) across investors in different countries or by the cross-sectional variation in the goods market integration.

References

Akram, Q. Farooq, Dagfinn Rime, and Lucio Sarno, "Arbitrage in the Foreign Exchange Market: Turning on the Microscope," *Journal of International Economics*, 2008, *forthcoming.*

Alvarez, Fernando, Andy Atkeson, and Patrick Kehoe, "Time-Varying Risk, Interest Rates and Exchange Rates in General Equilibrium," 2005. Working paper No 627 Federal Reserve Bank of Minneapolis Research Department.

Andrews, Donald W.K., "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 1991, *59* (1), 817–858.

Bacchetta, Philippe and Eric van Wincoop, "Incomplete Information Processing: A Solution to the Forward Discount Puzzle," September 2006. Working Paper University of Virginia.

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.

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 2007. 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.

Barro, Robert, "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 2006, *121*, 823–866.

Bekaert, Geert, "The Time Variation of Expected Returns and Volatility in Foreign-Exchange Markets," *Journal of Business and Economic Statistics*, 1995, *13* (4), 397–408.

___ , "The Time Variation of Risk and Return in Foreign Exchange Markets: A General Equilibrium Perspective," *The Review of Financial Studies*, 1996, *9* (2), 427–470.

___ **and Robert J. Hodrick**, "Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets," *The Journal of Finance*, 1992, *47*, 467–509.

——— **, Robert Hodrick, and David Marshall**, "The Implications of First-Order Risk Aversion for Asset Market Risk Premiums," *Journal of Monetary Economics*, 1997, *40*, 3–39.

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(4)*, 671–698.

Brennan, Michael J. and Yihong Xia, "International Capital Markets and Foreign Exchange Risk," *Review of Financial Studies*, 2006, *19* (3), 753–795.

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen, "Carry Trades and Currency Crashes," 2008. forthcoming NBER Macroannual.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, "The Returns to Currency Speculation in Emerging Markets," *American Economic Review Papers and Proceedings*, May 2007, *97* (2), 333–338.

——— **,** ——— **, and** ——— , "Understanding the forward premium puzzle: A microstructure approach," *American Economic Journal: Macroeconomics*, 2008, *forthcoming.*

——— **,** ——— **, Isaac Kleshchelski, and Sergio Rebelo**, "The Returns to Currency Speculation," November 2006. Working Paper NBER No. 12489.

——— **,** ——— **,** ——— **, and** ——— , "Can Peso Problems Explain the Returns to the Carry Trade?," May 2008. NBER Working Paper 14054.

Campbell, John Y., "Consumption-Based Asset Pricing," in George Constantinides, Milton Harris, and René Stultz, eds., *Handbook of the Economics of Finance*, Vol. 1, Amsterdam: Elsevier - North Holland, 2003, chapter 13, p. 801885.

——— **and John H. Cochrane**, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior.," *Journal of Political Economy*, 1999, *107* (2), 205–251.

——— **, Karine Serfaty de Medeiros, and Luis M. Viceira**, "Global Currency Hedging," 2006. Working Paper NBER No. 13088.

Cochrane, John H., *Asset Pricing*, Princeton, N.J.: Princeton University Press, 2001.

——— **and Monika Piazzesi**, "Bond Risk Premia," *The American Economic Review*, March 2005, *95* (1), 138–160.

——— **and** ——— , "Decomposing the Yield Curve," 2008. working paper University of Chicago.

Colacito, Riccardo, “Six Anomalies looking for a model. A consumption based explanation of International Finance Puzzles,” 2008. Mimeo.

—— **and Mariano Massimiliano Croce**, “Risks for the Long-Run the Real Exchange Rate,” September 2008. Available at SSRN: http://ssrn.com/abstract=1105496.

Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross, “A Theory of the Term Structure of Interest Rates,” *Econometrica*, 1985, *53* (2), 385–408.

Daniel, Kent and Sheridan Titman, “Testing Factor-Model Explanations of Market Anomalies,” 2005.

DeSantis, Roberto A. and Fabio Fornari, “Does Business Cycle Risk Account for Systematic Returns from Currency Positioning? The International Perspective,” June 2008. European Central Bank.

Duffee, Gregory R., “Information in (and not in) the term structure,” 2008. Working Paper.

Fama, Eugene F., “Forward and Spot Exchange Rates,” *Journal of Monetary Economics*, 1984, *14*, 319–338.

—— **and James D. MacBeth**, “Risk, Return, and Equilibrium: Empirical Tests,” *The Journal of Political Economy*, 3 1973, *81*, 607–636.

Farhi, Emmanuel and Xavier Gabaix, “Rare Disasters and Exchange Rates: A Theory of the Forward Premium Puzzle,” October 2007. Working Paper Harvard University.

Ferson, Wayne E. and Campbell R. Harvey, “The Risk and Predictability of International Equity Returns,” *Review of Financial Studies*, 1993, *6(3)*, 527–566.

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,” 2007. Working paper NBER No. 12496.

Froot, Kenneth and Richard Thaler, “Anomalies: Foreign Exchange,” *The Journal of Economic Perspectives*, 3 1990, *4*, 179–192.

Gorton, Gary B., Fumio Hayashi, and K. Geert Rouwenhorst, “The Fundamentals of Commodity Futures Returns,” June 2007. Yale ICF Working Paper No. 07-08.

Gourinchas, Pierre-Olivier and Aaron Tornell, "Exchange Rate Puzzle and Distorted Beliefs," *Journal of International Economics*, 2004, *64* (2), 303–333.

Graveline, Jeremy J., "Exchange Rate Volatility and the Forward Premium Anomaly," 2006. Working Paper.

Hansen, Lars Peter, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 1982, *50*, 1029–54.

___ and Robert J. Hodrick, "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, 1980, *88*, 829–853.

___ and ___, "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., Bruno Solnik, and Guofu Zhou, "What Determines Expected International Asset Returns?," *Annals of Economics and Finance*, 2002, *3* (2), 249–298.

Hau, Harald and Helene Rey, "Global Portfolio Rebalancing Under the Microscope," 2007. Working Paper.

Hollifield, Burton and Amir Yaron, "The Foreign Exchange Risk Premium: Real and Nominal Factors," 2001. Working Paper Wharton School of Business, University of Pennsylvania.

Korajczyk, Robert A., "The Pricing of Forward Contracts for Foreign Exchange," *Journal of Political Economy*, 1985, *93* (2), 346 – 368.

Ludvigson, Sydney C. and Serena Ng, "Macro Factors in Bond Risk Premia," 2005. NBER Working Paper 11703.

Lustig, Hanno and Adrien Verdelhan, "The Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk," *American Economic Review*, March 2007, *97* (1), 89–117.

Lyons, Richard K., *The Microstructure Approach to Exchange Rates*, M.I.T Press, 2001.

Newey, Whitney K. and Kenneth D. West, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 1987, *55* (3), 703–708.

Piazzesi, Monika and Eric Swanson, "Futures Prices as Risk-Adjusted Forecasts of Monetary Policy," *forthcoming Journal of Monetary Economics*, 2008.

Plantin, Guillaume and Hyun Song Shin, "Carry Trades and Speculative Dynamics," July 2007. Working Paper Princeton University.

Ross, Stephen A., "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, December 1976, *13*, 341–360.

Sarno, Lucio, Gene Leon, and Giorgio Valente, "Nonlinearity in Deviations from Uncovered Interest Parity: An Explanation of the Forward Bias Puzzle," *Review of Finance*, 2006, pp. 443–482.

Shanken, Jay, "On the Estimation of Beta-Pricing Models," *The Review of Financial Studies*, 1992, *5* (2), 1–33.

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.

Verdelhan, Adrien, "A Habit Based Explanation of the Exchange Rate Risk Premium," 2005. Working Paper Boston University.

Appendix A Model

Appendix A1 Inflation

We assume that inflation is composed of a country-specific component and a global component. Both components follow AR(1) processes:

$$\begin{aligned} \pi^w_{t+1} &= (1-\rho^w)\overline{\pi}^w + \rho^w \pi^w_t + \sigma^{w\$}\epsilon^w_{t+1}, \\ \pi^{ci}_{t+1} &= (1-\rho^i)\overline{\pi}^i + \rho^i \pi^i_t + \sigma^{i\$}\epsilon^i_{t+1}, \end{aligned}$$

where the innovations ϵ^w_t and ϵ^i_t are also $i.i.d$ gaussian, with zero mean and unit variance. Inflation in country i is a weighted average of these two components:

$$\pi^i_{t+1} = \mu^i \pi^{ci}_{t+1} + (1-\mu^i)\pi^w_{t+1}.$$

We define world inflation as the cross-sectional, unweighted average of all annual inflation rates, and we measure the moments of the average world inflation rate for the countries in our sample. The average global inflation is calibrated to be 3 percent annually, autocorrelation is equal to 0.87, and standard deviation is 2.1%. The relative weight μ on domestic versus world inflation set equal to 0.16; it is determined by the share of the total variance explained by the first principal component. We subtract the world component from each country inflation rate to obtain the autocorrelation and the shocks' standard deviation in each country. We use the average of these moments. This yields an average for the country-specific component equal to 3 percent, an autocorrelation of 0.5 and standard deviation equal to 10 percent, for the annualized series. Inflation moments are reported in panel III of table 12. We use monthly values corresponding to these annual quantities in calibrating the model parameters. Table 12 (panel IV) reports the calibrated parameters.

Figure 9 documents the statistical properties of the interest rates and exchange rates simulated from the model for a range of currencies.

[Figure 9 about here.]

Appendix A2 Stock Market Return

The ex-dividend price of the stock market portfolio at time t in the units of domestic currency is given by

$$P^i_t = E_t \sum_{s=1}^{\infty} D^i_{t+s} \exp\left[\sum_j^s m^i_{t+j}\right].$$

Since all the relevant information at time t is summarized by the state vector $[z_t^i, z_t^w]$, we can write the price-dividend ratio as

$$\frac{P_t^i}{D_t^i} = E\left\{ \sum_{s=1}^{\infty} \exp\left[\sum_{j=1}^{s} \left(\Delta d_{t+j}^i + m_{t+j}^i \right) \right] \middle| z_t^i, z_t^w \right\}$$

We compute the price-dividend ratios that correspond to the simulated values of the state vector using Monte Carlo simulation and interpolate them using a kernel regression.

The stock market return is then calculated using the identity

We compute the stock market returns using

$$R_{t+1}^{Di} = \frac{P_{t+1}^i + D_{t+1}^i}{P_t^i} = \frac{P_{t+1}^i / D_{t+1}^i + 1}{P_t^i / D_t^i} \exp\left(\Delta d_{t+1}^i \right)$$

by simulating the dividend process jointly with the state variables and SDF innovations and using the kernel projection to interpolate the price-dividend ratios. In order to calculate the conditional correlations and betas of this return with the currency portfolio returns, as well as its conditional volatility, we consider 12-month rolling windows and estimate these moments in the same way as in the data.

Table 1: Currency Portfolios - US Investor

Portfolio	1	2	3	4	5	6	1	2	3	4	5
	Panel I: All Countries						Panel II: Developed Countries				
	Spot change: Δs^j						Δs^j				
Mean	−0.97	−1.33	−1.55	−2.73	−0.99	1.88	−1.86	−2.54	−4.05	−2.11	−1.11
Std	8.04	7.29	7.41	7.42	7.74	9.16	10.12	9.71	9.24	8.92	9.20
	Forward Discount: $f^j - s^j$						$f^j - s^j$				
Mean	−3.90	−1.30	−0.15	0.94	2.55	7.78	−3.09	−1.02	0.07	1.13	3.94
Std	1.57	0.49	0.48	0.53	0.59	2.09	0.78	0.63	0.65	0.67	0.76
	Excess Return: rx^j (without b-a)						rx^j (without b-a)				
Mean	−2.92	0.02	1.40	3.66	3.54	5.90	−1.24	1.52	4.11	3.24	5.06
Std	8.22	7.36	7.46	7.53	7.85	9.26	10.20	9.75	9.35	9.01	9.30
SR	−0.36	0.00	0.19	0.49	0.45	0.64	−0.12	0.16	0.44	0.36	0.54
	Net Excess Return: rx^j_{net} (with b-a)						rx^j_{net} (with b-a)				
Mean	−1.70	−0.95	0.12	2.31	2.04	3.14	−0.11	0.46	2.71	1.98	3.35
Std	8.21	7.35	7.43	7.48	7.85	9.25	10.20	9.75	9.32	9.02	9.30
SR	−0.21	−0.13	0.02	0.31	0.26	0.34	−0.01	0.05	0.29	0.22	0.36
	High-minus-Low: $rx^j - rx^1$ (without b-a)						$rx^j - rx^1$ (without b-a)				
Mean		2.95	4.33	6.59	6.46	8.83		2.75	5.35	4.47	6.29
Std		5.36	5.54	6.65	6.34	8.95		6.42	6.44	7.38	8.70
SR		0.55	0.78	0.99	1.02	0.99		0.43	0.83	0.61	0.72
	High-minus-Low: $rx^j_{net} - rx^1_{net}$ (with b-a)						$rx^j_{net} - rx^1_{net}$ (with b-a)				
Mean		0.75	1.82	4.00	3.73	4.83		0.57	2.82	2.09	3.46
Std		5.36	5.56	6.63	6.35	8.98		6.45	6.44	7.41	8.73
SR		0.14	0.33	0.60	0.59	0.54		0.09	0.44	0.28	0.40

Notes: This table reports, for each portfolio j, the average change in log spot exchange rates Δs^j, the average log forward discount $f^j - s^j$, the average log excess return rx^j without bid-ask spreads, the average log excess return rx^j_{net} with bid-ask spreads, and the average return on the long short strategy $rx^j_{net} - rx^1_{net}$ and $rx^j - rx^1$ (with and without bid-ask spreads). Log currency excess returns are computed as $rx^j_{t+1} = -\Delta s^j_{t+1} + f^j_t - s^j_t$. All moments are annualized and reported in percentage points. For excess returns, the table also reports Sharpe ratios, computed as ratios of annualized means to annualized standard deviations. The portfolios are constructed by sorting currencies into six groups at time t based on the one-month forward discount (i.e nominal interest rate differential) at the end of period $t-1$. Portfolio 1 contains currencies with the lowest interest rates. Portfolio 6 contains currencies with the highest interest rates. Panel I uses all countries, panel II focuses on developed countries. Data are monthly, from Barclays and Reuters (Datastream). The sample period is 11/1983 - 03/2008.

Table 2: Principal Components

	Panel I: All Countries					
Portfolio	1	2	3	4	5	6
1	0.43	0.41	−0.18	0.31	0.72	0.03
2	0.39	0.26	−0.14	−0.02	−0.44	0.75
3	0.39	0.26	−0.46	−0.38	−0.31	−0.57
4	0.38	0.05	0.72	−0.56	0.16	−0.01
5	0.42	−0.11	0.38	0.66	−0.37	−0.31
6	0.43	−0.82	−0.28	−0.10	0.18	0.11
% Var.	70.07	12.25	6.18	4.51	3.76	3.23
	Panel II: Developed Countries					
Portfolio	1	2	3	4	5	
1	0.48	0.56	0.60	0.23	0.20	
2	0.47	0.29	−0.66	−0.32	0.40	
3	0.46	0.05	−0.30	0.36	−0.76	
4	0.42	−0.34	0.34	−0.72	−0.25	
5	0.41	−0.69	0.02	0.44	0.40	
% Var	79.06	9.33	4.73	3.58	3.30	

Notes: This table reports the principal component coefficients of the currency portfolios. In each panel, the last row reports (in %) the share of the total variance explained by each common factor. Data are monthly, from Barclays and Reuters (Datastream). The sample period is 11/1983 - 03/2008.

Table 3: Asset Pricing - US Investor

Panel I: Factor Prices and Loadings

	All Countries							Developed Countries						
	$\lambda_{HML_{FX}}$	λ_{RX}	$b_{HML_{FX}}$	b_{RX}	R^2	$RMSE$	χ^2	$\lambda_{HML_{FX}}$	λ_{RX}	$b_{HML_{FX}}$	b_{RX}	R^2	$RMSE$	χ^2
GMM_1	5.46	1.35	0.59	0.26	69.28	0.95		3.56	2.24	0.43	0.32	71.06	0.61	
	[2.34]	[1.68]	[0.25]	[0.32]			13.83	[2.19]	[2.02]	[0.24]	[0.24]			41.06
GMM_2	4.88	0.58	0.52	0.12	47.89	1.24		3.78	3.03	0.46	0.42	20.41	1.00	
	[2.23]	[1.63]	[0.24]	[0.31]			15.42	[2.14]	[1.95]	[0.23]	[0.23]			44.36
FMB	5.46	1.35	0.58	0.26	69.28	0.95		3.56	2.24	0.42	0.32	71.06	0.61	
	[1.82]	[1.34]	[0.19]	[0.25]			13.02	[1.80]	[1.71]	[0.20]	[0.20]			41.34
	(1.83)	(1.34)	(0.20)	(0.25)			14.32	(1.80)	(1.71)	(0.20)	(0.20)			42.35
$Mean$	**5.37**	**1.36**						**3.44**	**2.24**					

Panel II: Factor Betas

	All Countries						Developed Countries					
$Portfolio$	$\alpha_0^j(\%)$	$\beta^j_{HML_{FX}}$	β^j_{RX}	$R^2(\%)$	$\chi^2(\alpha)$	$p-value$	$\alpha_0^j(\%)$	$\beta^j_{HML_{FX}}$	β^j_{RX}	$R^2(\%)$	$\chi^2(\alpha)$	$p-value$
1	−0.56	−0.39	1.06	91.36			0.00	−0.50	1.00	94.95		
	[0.52]	[0.02]	[0.03]				[0.48]	[0.02]	[0.02]			
2	−1.21	−0.13	0.97	78.54			−0.90	−0.11	1.02	82.38		
	[0.76]	[0.03]	[0.05]				[0.81]	[0.04]	[0.04]			
3	−0.13	−0.12	0.95	73.73			1.01	−0.02	1.02	85.22		
	[0.82]	[0.03]	[0.04]				[0.83]	[0.03]	[0.03]			
4	1.62	−0.02	0.93	68.86			−0.12	0.13	0.97	81.43		
	[0.86]	[0.04]	[0.06]				[0.85]	[0.04]	[0.04]			
5	0.84	0.05	1.03	76.37			0.00	0.50	1.00	93.87		
	[0.80]	[0.04]	[0.05]				[0.48]	[0.02]	[0.02]			
6	−0.56	0.61	1.06	93.03								
	[0.52]	[0.02]	[0.03]									
All					10.11	0.12					2.61	0.76

Notes: The panel on the left reports results for all countries. The panel on the right reports results for the developed countries. Panel I reports results from GMM and Fama-McBeth asset pricing procedures. Market prices of risk λ, the adjusted R^2, the square-root of mean-squared errors $RMSE$ and the p-values of χ^2 tests on pricing errors are reported in percentage points. b denotes the vector of factor loadings. Excess returns used as test assets and risk factors take into account bid-ask spreads. All excess returns are multiplied by 12 (annualized). The standard errors in brackets are Newey and West (1987) standard errors computed with the optimal number of lags according to Andrews (1991). Shanken (1992)-corrected standard errors are reported in parentheses. We do not include a constant in the second step of the FMB procedure. Panel II reports OLS estimates of the factor betas. R^2s and p-values are reported in percentage points. The χ^2 test statistic $\alpha' V_\alpha^{-1} \alpha$ tests the null that all intercepts are jointly zero. This statistic is constructed from the Newey-West variance-covariance matrix (1 lag) for the system of equations (see Cochrane (2001), p. 234). Data are monthly, from Barclays and Reuters in Datastream. The sample period is 11/1983 - 03/2008. The alphas are annualized and in percentage points.

Table 4: Conditional Betas - US Investor

Factor Betas of Exchange Rate Changes Δs^j_{t+1}								
	All Countries				Developed Countries			
$Portfolio$	$\alpha^j_0(\%)$	$\beta^j_{HML_{FX}}$	β^j_{RX}	$R^2(\%)$	$\alpha^j_0(\%)$	$\beta^j_{HML_{FX}}$	β^j_{RX}	$R^2(\%)$
1	-1.59	0.37	-1.03	89.73	-1.25	0.49	-0.98	94.59
	[0.56]	[0.02]	[0.03]		[0.50]	[0.02]	[0.02]	
2	-0.66	0.13	-0.96	78.08	-0.71	0.11	-1.01	81.54
	[0.76]	[0.03]	[0.05]		[0.83]	[0.05]	[0.04]	
3	-0.95	0.12	-0.94	72.78	-1.88	0.02	-1.00	84.89
	[0.85]	[0.03]	[0.04]		[0.80]	[0.03]	[0.03]	
4	-1.51	0.02	-0.92	67.25	0.41	-0.12	-0.95	80.10
	[0.89]	[0.04]	[0.06]		[0.94]	[0.04]	[0.04]	
5	0.71	-0.05	-1.02	75.81	2.90	-0.50	-0.98	92.87
	[0.79]	[0.04]	[0.05]		[0.53]	[0.02]	[0.02]	
6	6.56	-0.62	-1.02	89.60				
	[0.65]	[0.02]	[0.04]					

Notes: The panel on the left reports results for all countries. The panel on the right reports results for the developed countries. The table reports OLS estimates of the factor betas obtained by regressing changes in log spot exchange rates Δs^j_{t+1} on the factors. R^2s are reported in percentage points. Data are monthly, from Barclays and Reuters in Datastream. The sample period is 11/1983 - 03/2008. The alphas are annualized and in percentage points.

Table 5: Asset Pricing - US Investor - Principal Components

Panel I: Factor Prices and Loadings

	All Countries							Developed Countries						
	λ_c	λ_d	b_c	b_d	R^2	$RMSE$	χ^2	λ_2	λ_1	b_c	b_d	R^2	$RMSE$	χ^2
GMM_1	7.42	1.37	0.40	0.26	68.69	0.96		2.20	2.17	0.72	0.25	70.75	0.61	
	[3.12]	[1.65]	[0.17]	[0.31]			12.92	[1.22]	[2.02]	[0.40]	[0.23]			51.15
GMM_2	6.23	0.54	0.34	0.10	43.12	1.30		2.63	2.90	0.86	0.34	24.08	0.98	
	[2.86]	[1.60]	[0.15]	[0.30]			15.21	[1.17]	[1.94]	[0.38]	[0.23]			57.14
FMB	7.42	1.37	0.40	0.26	68.72	0.96		2.20	2.17	0.72	0.25	70.75	0.61	
	[2.52]	[1.35]	[0.14]	[0.25]			11.25	[1.02]	[1.72]	[0.33]	[0.20]			41.67
	[2.52]	[1.35]	[0.14]	[0.25]			12.37	[1.02]	[1.72]	[0.33]	[0.20]			42.64
$Mean$	**7.42**	**1.37**						**2.20**	**2.17**					

Panel II: Factor Betas

	All Countries						Developed Countries					
$Portfolio$	$\alpha_0^j(\%)$	β_c^j	β_d^j	$R^2(\%)$	$\chi^2(\alpha)$	$p-value$	$\alpha_0^j(\%)$	β_d^j	β_c^j	$R^2(\%)$	$\chi^2(\alpha)$	$p-value$
1	-0.99	-0.23	1.06	85.69			-0.21	-0.72	1.07	91.14		
	[0.72]	[0.02]	[0.04]				[0.64]	[0.05]	[0.02]			
2	-0.85	-0.14	0.96	81.38			-0.43	-0.38	1.04	85.94		
	[0.69]	[0.02]	[0.04]				[0.72]	[0.07]	[0.03]			
3	0.31	-0.14	0.94	76.89			1.15	-0.07	1.02	85.59		
	[0.84]	[0.02]	[0.04]				[0.81]	[0.07]	[0.03]			
4	1.72	-0.03	0.92	68.16			-0.54	0.44	0.94	85.14		
	[0.86]	[0.03]	[0.06]				[0.77]	[0.06]	[0.03]			
5	0.64	0.06	1.03	77.41			0.01	0.89	0.92	93.64		
	[0.80]	[0.03]	[0.04]				[0.49]	[0.04]	[0.02]			
6	-0.64	0.45	1.06	96.83								
	[0.34]	[0.01]	[0.02]									
All					6.90	0.33					2.40	0.79

Notes: The factors are the first and the second principal components (denoted d, for the "dollar" factor, and c, for the "carry" factor, respectively). The panel on the left reports results for all countries. The panel on the right reports results for the developed countries. Panel I reports results from GMM and Fama-McBeth asset pricing procedures. Market prices of risk λ, the adjusted R^2, the square-root of mean-squared errors $RMSE$ and the p-values of χ^2 tests on pricing errors are reported in percentage points. b denotes the vector of factor loadings. Excess returns used as test assets and risk factors take into account bid-ask spreads. All excess returns are multiplied by 12 (annualized). The standard errors in brackets are Newey and West (1987) standard errors computed with the optimal number of lags according to Andrews (1991). Shanken (1992)-corrected standard errors are reported in parentheses. We do not include a constant in the second step of the FMB procedure. Panel II reports OLS estimates of the factor betas. R^2s and p-values are reported in percentage points. The χ^2 test statistic $\alpha' V_\alpha^{-1} \alpha$ tests the null that all intercepts are jointly zero. This statistic is constructed from the Newey-West variance-covariance matrix (1 lag) for the system of equations (see Cochrane (2001), p. 234). Data are monthly, from Barclays and Reuters in Datastream. The sample period is 11/1983 - 03/2008. The alphas are annualized and in percentage points.

Table 6: Beta-Sorted Currency Portfolios - US Investor

Portfolio	1	2	3	4	5	6	1	2	3	4	5
	Panel I: Developed and Emerging Countries						Panel II: Developed Countries				
	Spot change: Δs^j						Spot change: Δs^j				
Mean	−2.11	−1.80	−1.25	−1.97	−1.80	−0.14	−1.95	−2.33	−1.88	−2.20	0.28
Std	8.74	7.86	7.28	6.75	8.06	7.45	8.79	8.20	8.15	7.83	7.58
	Discount: $f^j - s^j$						Discount: $f^j - s^j$				
Mean	−1.45	−0.38	0.75	0.93	1.48	3.18	−1.46	−0.51	0.98	1.28	4.15
Std	0.77	0.56	1.23	0.64	0.80	1.26	0.69	0.60	0.71	0.82	1.65
	Excess Return: rx^j (without b-a)						Excess Return: rx^j (without b-a)				
Mean	0.66	1.42	2.00	2.90	3.29	3.32	0.48	1.82	2.86	3.48	3.87
Std	8.88	7.87	7.33	6.71	8.07	7.48	8.87	8.24	8.20	7.79	7.97
SR	0.07	0.18	0.27	0.43	0.41	0.44	0.05	0.22	0.35	0.45	0.49
	High-minus-Low: $rx^j - rx^1$ (without b-a)						Excess Return: rx^j (without b-a)				
Mean		0.76	1.34	2.24	2.63	2.66		1.34	2.38	2.99	3.38
Std		5.24	6.34	7.43	8.88	9.23		5.34	5.96	7.96	9.02
SR		0.15	0.21	0.30	0.30	0.29		0.25	0.40	0.38	0.38
	Pre-formation β's						Pre-formation β's				
Mean	−0.40	−0.24	−0.15	0.01	0.21	0.57	−0.39	−0.23	−0.04	0.15	0.46
Std	0.29	0.23	0.24	0.26	0.43	0.41	0.26	0.25	0.35	0.45	0.41
	Post-formation β's						Post-formation β's				
Estimate	−0.31	−0.20	−0.14	0.01	0.13	0.28	−0.26	−0.15	0.04	0.08	0.30
s.e	[0.04]	[0.05]	[0.05]	[0.05]	[0.06]	[0.06]	[0.05]	[0.05]	[0.05]	[0.05]	[0.04]

Notes: This table reports, for each portfolio j, the average change in the log spot exchange rate Δs^j, the average log forward discount $f^j - s^j$, the average log excess return rx^j without bid-ask spreads and the average returns on the long short strategy $rx^j - rx^1$. The left panel uses our sample of developed and emerging countries. The right panel uses our sample of developed countries. Log currency excess returns are computed as $rx^j_{t+1} = -\Delta s^j_{t+1} + f^j_t - s^j_t$. All moments are annualized and reported in percentage points. For excess returns, the table also reports Sharpe ratios, computed as ratios of annualized means to annualized standard deviations. Portfolios are constructed by sorting currencies into six groups at time t based on slope coefficients β^i_t. Each β^i_t is obtained by regressing currency i log excess return rx^i on HML_{FX} on a 36-period moving window that ends in period $t-1$. The first portfolio contains currencies with the lowest βs. The last portfolio contains currencies with the highest βs. We report the average pre-formation beta for each portfolio. The last panel reports the post-formation betas obtained by regressing realized log excess returns on portfolio j on HML_{FX} and RX_{FX}. We only report the HML_{FX} betas. The standard errors are reported in brackets. Data are monthly, from Barclays and Reuters (Datastream). The sample period is 11/1983 - 03/2008.

Table 7: One-Month Ahead Return Predictability

Portfolio	$\kappa_{\mathbf{f}}$	W	R^2	κ_f	W	R^2	*Portfolio*	$\kappa_{\mathbf{sp,f}}$	W	R^2	$\kappa_{sp,f}$	W	R^2
			Panel A: Returns							Panel B: Spreads			
1	3.65		7.85	1.08		4.30							
NW	[0.64]	32.10		[0.33]	11.03								
HH	[0.57]	40.36		[0.23]	21.92								
VAR	[0.73]	37.57		[0.36]	17.28								
2	2.29		3.86	2.44		2.65	*2 minus 1*	8.31		3.05	9.08		4.81
NW	[0.70]	10.76		[0.97]	6.28		*NW*	[3.02]	7.57		[2.66]	11.64	
HH	[0.69]	11.13		[0.92]	6.98		*HH*	[0.43]	368.58		[2.44]	13.86	
VAR	[0.72]	16.49		[1.02]	8.79		*VAR*	[3.58]	8.51		[3.43]	12.66	
3	1.93		2.68	1.96		1.61	*3 minus 1*	7.10		2.09	7.28		2.89
NW	[0.65]	8.92		[1.04]	3.56		*NW*	[3.01]	5.58		[2.27]	10.26	
HH	[0.63]	9.48		[1.02]	3.67		*HH*	[3.03]	5.49		[2.40]	9.23	
VAR	[0.66]	12.88		[0.97]	5.94		*VAR*	[4.01]	5.74		[3.72]	7.75	
4	2.22		3.47	3.47		5.98	*4 minus 1*	8.33		1.99	9.27		3.28
NW	[0.65]	11.61		[0.87]	16.03		*NW*	[2.99]	7.75		[2.38]	15.22	
HH	[0.64]	12.16		[0.82]	18.02		*HH*	[2.85]	8.53		[2.38]	15.22	
VAR	[0.72]	14.28		[0.92]	18.32		*VAR*	[4.34]	6.69		[4.03]	10.97	
5	2.68		4.63	3.02		5.10	*5 minus 1*	7.13		1.61	6.83		2.15
NW	[0.74]	13.01		[0.91]	11.11		*NW*	[3.49]	4.17		[2.82]	5.86	
HH	[0.76]	12.44		[0.93]	10.61		*HH*	[1.68]	17.94		[0.96]	50.74	
VAR	[0.77]	19.80		[0.83]	16.33		*VAR*	[4.00]	6.18		[3.32]	8.31	
6	3.09		4.44	0.71		2.56	*6 minus 1*	9.93		1.57	3.73		0.80
NW	[0.84]	13.61		[0.21]	11.40		*NW*	[4.20]	5.59		[3.10]	1.45	
HH	[0.85]	13.27		[0.21]	11.48		*HH*	[3.73]	7.09		[3.08]	1.47	
VAR	[0.94]	16.80		[0.32]	12.78		*VAR*	[5.30]	7.36		[2.99]	3.60	

Notes: Panel A reports summary statistics for return predictability regressions at a one-month horizon. For each portfolio j, we report the R^2, and the slope coefficient in the time-series regression of the log currency excess return on the average log forward discount ($\kappa_{\mathbf{f}}$) in the left panel and the portfolio-specific log forward discount (κ_f) in the right panel. Panel B reports summary statistics for return predictability regressions of the spread at a one-month horizon. The left panel reports the statistics in the regression of one-month excess returns on the average one-month forward discount spread ($\kappa_{\mathbf{sp,f}}$). The right panel reports the statistics in the regression of one-month excess returns on that portfolio's one-month forward discount spread ($\kappa_{sp,f}$). W is the Wald-test χ^2 statistic for the slope coefficient. The Newey and West (1987) NW standard errors are computed with the optimal number of lags following Andrews (1991). The Hansen and Hodrick (1980b) HH standard error are computed with one lag. The bootstrapped standard errors VAR are computed by drawing from the residuals of a VAR with one lag. All the returns are annualized and reported in percentage points. Data are monthly, from Barclays and Reuters (Datastream). The returns take into account bid-ask spreads. The sample period is 11/1983 - 03/2008.

Table 8: Return Predictability: Longer Horizons

Horizon	1	2	3	6	12	1	2	3	6	12
Panel I: All Countries										
Overlapping Data										
Portfolio	Forward Discount					Residual Predictability				
1	4.30	4.64	8.03	25.30	25.93	0.23	0.00	0.01	1.18	0.20
6	2.56	3.07	3.82	5.72	10.03	0.01	0.03	0.06	0.03	0.05
	Average Forward Discount									
1	7.85	12.58	17.16	28.32	32.57					
6	4.44	6.13	8.46	12.70	17.54					
No Overlapping Data										
Portfolio	Forward Discount					Residual Predictability				
1	4.30	2.52	8.84	24.62	28.18	0.23	0.23	0.05	0.54	0.61
6	2.56	3.59	4.19	4.67	14.50	0.01	0.01	0.00	0.01	0.04
	Average Forward Discount									
1	7.85	13.41	17.87	31.74	30.22					
6	4.44	6.49	7.58	12.58	25.55					
Panel II: Developed Countries										
Overlapping Data										
Portfolio	Forward Discount					Residual Predictability				
1	1.95	3.51	6.86	14.41	17.23	0.01	0.25	0.17	0.12	0.06
5	3.29	5.74	7.67	12.26	13.55	0.24	0.24	0.21	0.42	1.22
	Average Forward Discount									
1	3.02	6.31	10.08	18.39	20.51					
5	2.85	5.34	7.80	12.27	10.43					
No Overlapping Data										
Portfolio	Forward Discount					Residual Predictability				
1	1.95	1.90	7.54	16.67	17.17	0.01	1.04	0.12	0.33	0.04
5	3.29	6.21	8.29	19.22	19.14	0.34	0.83	0.36	1.95	1.87
Portfolio	Average Forward Discount									
1	3.02	6.37	10.56	22.74	20.12					
5	2.85	4.19	7.79	15.81	14.19					

Notes: In the left panel, we report the R^2 in the time-series regressions of the log k-period currency excess return on the log forward discount for each portfolio j: $rx^{j,k}_{net,t+k} = \kappa^j_0 + \kappa^j_1(f^{j,k}_t - s^j_t) + \eta^j_t$. In the left panel, we also report the R^2 in the time-series regression the log k-period currency excess return on the linear combination of log forward discounts for each portfolio j: $rx^{j,k}_{net,t+k} = \kappa^j_0 + \kappa^j_1\iota'(\mathbf{f}^k_t - \mathbf{s}^k_t) + \eta^j_t$ for each portfolio j. In the right panel, we report the residual predictability: In a first step, we regress the log k-period currency excess return on the average log forward discount for each portfolio j: $rx^{j,k}_{net,t+k} = \kappa^j_0 + \kappa^j_1\iota'(\mathbf{f}^k_t - \mathbf{s}^k_t) + \eta^j_t$. We report the R^2 in the time-series regression of the residuals η^j_t from the first step on the log forward discounts for each portfolio j: $rx^{j,k}_{net,t+k} = \kappa^j_0 + \kappa^j_1(f^k_t - s^k_t) + \epsilon^j_t$ for each portfolio j. Data are monthly, from Barclays and Reuters (Datastream). The sample period is 11/1983 - 03/2008. Panel I uses developed and emerging countries. Panel II focuses on developed countries. In both cases, the top panel uses overlapping data and the bottom panel does not.

Table 9: Contemporaneous Correlations Between Expected Excess Returns or Forward Discounts and Macroeconomic and Financial Variables

	IP	Pay	$Help$	$spread$	$slope$	vol
Portfolio	Panel I: Expected Excess Returns					
1	0.18	0.02	0.19	−0.21	0.04	−0.17
	[0.04]	[0.02]	[0.11]	[0.03]	[0.04]	[0.02]
2	−0.57	−0.70	−0.41	0.34	0.42	−0.14
	[0.04]	[0.04]	[0.05]	[0.02]	[0.04]	[0.02]
3	−0.61	−0.64	−0.37	0.33	0.47	−0.04
	[0.05]	[0.05]	[0.06]	[0.02]	[0.04]	[0.02]
4	−0.57	−0.51	−0.30	0.26	0.42	0.09
	[0.06]	[0.05]	[0.06]	[0.02]	[0.04]	[0.02]
5	−0.51	−0.39	−0.24	0.28	0.38	0.28
	[0.05]	[0.05]	[0.05]	[0.02]	[0.03]	[0.02]
6	−0.14	−0.09	−0.05	0.17	0.15	0.52
	[0.05]	[0.05]	[0.05]	[0.02]	[0.05]	[0.02]
Maturity	Panel II: Average Forward Discount					
1	−0.31	−0.34	−0.13	0.17	0.33	0.18
	[0.12]	[0.04]	[0.14]	[0.04]	[0.08]	[0.05]
2	−0.46	−0.47	−0.24	0.26	0.40	0.24
	[0.15]	[0.05]	[0.15]	[0.04]	[0.09]	[0.05]
3	−0.51	−0.52	−0.30	0.30	0.41	0.27
	[0.16]	[0.05]	[0.15]	[0.04]	[0.09]	[0.05]
6	−0.54	−0.57	−0.38	0.35	0.40	0.32
	[0.18]	[0.05]	[0.15]	[0.05]	[0.10]	[0.07]
12	−0.50	−0.60	−0.37	0.29	0.41	0.24
	[0.18]	[0.05]	[0.17]	[0.06]	[0.12]	[0.08]

Notes: Panel I reports the contemporaneous correlation $Corr\left[\widehat{E}_t rx^j_{t+1}, x_t\right]$ of forecasted excess returns using the portfolio forward discount with different variables x_t: the 12-month percentage change in industrial production ($\Delta \log IP_t$), the 12-month percentage change in the total US non-farm payroll ($\Delta \log Pay_t$), and the 12-month percentage change of the Help-Wanted index ($\Delta \log Help_t$), the default spread ($spread_t$), the slope of the yield curve ($slope_t$) and the CBOE $S\&P$ 500 volatility index (vol_t). Panel II reports the contemporaneous correlation of the average forward discount with these variables. Data are monthly, from Datastream and Global Financial Data. The sample period is 11/1983 - 03/2008.

Table 10: Forecasting 12-month ahead Excess Returns with Industrial Production and Forward Discounts

	κ_{IP}	κ_f	W	R^2	κ_{IP}	$\kappa_{\mathbf{f}}$	W	R^2	κ_{IP}	κ_f	W	R^2	κ_{IP}	$\kappa_{\mathbf{f}}$	W	R^2
	All Countries								Developed Countries							
1	−0.92	2.23		30.20	−0.89	3.09		37.37	−1.30	1.27		23.45	−1.13	1.79		25.03
NW	[0.60]	[1.21]	37.13		[0.28]	[0.80]	41.77		[0.72]	[1.16]	19.66		[0.55]	[0.93]	21.24	
HH	[0.67]	[1.38]	38.95		[0.29]	[0.83]	47.75		[0.78]	[1.31]	17.37		[0.59]	[1.02]	19.39	
VAR	[0.71]	[1.31]	38.13		[0.61]	[1.10]	41.20		[0.91]	[1.55]	33.55		[0.89]	[1.49]	33.92	
No overlap	[0.78]	[1.60]	22.31		[0.51]	[1.37]	24.37		[0.91]	[1.48]	12.23		[0.78]	[1.38]	13.71	
2	−0.98	0.69		18.68	−0.94	0.98		20.13	−1.91	−0.21		21.25	−1.42	1.03		22.58
NW	[0.52]	[1.00]	15.30		[0.36]	[0.70]	15.11		[0.83]	[1.41]	16.63		[0.60]	[1.25]	17.89	
HH	[0.58]	[1.11]	16.11		[0.40]	[0.71]	16.36		[0.92]	[1.59]	14.45		[0.66]	[1.40]	15.64	
VAR	[0.54]	[1.08]	21.93		[0.51]	[0.92]	41.20		[0.88]	[1.56]	44.24		[0.89]	[1.49]	33.92	
No overlap	[0.68]	[1.61]	8.12		[0.48]	[1.25]	9.65		[0.96]	[1.94]	11.50		[0.79]	[1.98]	12.55	
3	−1.18	1.18		29.42	−1.15	1.51		31.75	−1.71	0.61		29.92	−1.68	0.71		30.02
NW	[0.36]	[0.92]	26.76		[0.30]	[0.82]	28.02		[0.43]	[0.86]	39.90		[0.46]	[0.99]	40.18	
HH	[0.40]	[0.99]	23.17		[0.33]	[0.90]	24.16		[0.46]	[0.93]	35.58		[0.48]	[1.09]	36.04	
VAR	[0.54]	[0.93]	62.73		[0.49]	[0.89]	56.88		[0.66]	[0.92]	52.70		[0.69]	[1.09]	48.97	
No overlap	[0.71]	[1.50]	14.59		[0.56]	[1.42]	16.13		[0.61]	[1.48]	92.52		[0.58]	[1.43]	92.46	
4	−1.19	1.02		31.66	−1.19	1.20		32.38	−1.48	0.84		32.46	−1.42	1.08		33.01
NW	[0.28]	[0.69]	32.51		[0.27]	[0.74]	31.14		[0.46]	[0.97]	51.55		[0.49]	[1.18]	49.47	
HH	[0.30]	[0.72]	29.88		[0.29]	[0.79]	28.37		[0.50]	[1.05]	49.98		[0.54]	[1.30]	47.69	
VAR	[0.46]	[0.64]	61.11		[0.44]	[0.77]	63.26		[0.57]	[0.85]	50.78		[0.58]	[1.02]	61.71	
No overlap	[0.39]	[1.44]	24.95		[0.31]	[1.48]	21.21		[0.62]	[1.54]	45.16		[0.57]	[1.82]	69.50	
5	−1.71	1.20		39.97	−1.72	0.97		37.90	−1.76	0.64		32.75	−2.14	−0.45		32.03
NW	[0.31]	[0.66]	43.03		[0.35]	[0.79]	38.81		[0.39]	[1.22]	41.94		[0.52]	[1.43]	48.03	
HH	[0.32]	[0.69]	47.98		[0.38]	[0.79]	43.60		[0.41]	[1.37]	38.25		[0.56]	[1.60]	44.46	
VAR	[0.41]	[0.71]	68.34		[0.46]	[0.81]	53.27		[0.68]	[1.10]	48.86		[0.73]	[1.25]	51.50	
No overlap	[0.54]	[0.98]	33.12		[0.70]	[1.51]	22.11		[0.45]	[1.46]	37.95		[0.67]	[1.86]	40.11	
6	−1.50	1.08		26.64	−1.08	1.95		24.09								
NW	[0.42]	[0.45]	23.97		[0.50]	[1.38]	17.97									
HH	[0.45]	[0.46]	20.20		[0.53]	[1.51]	15.68									
VAR	[0.52]	[0.57]	53.36		[0.65]	[1.13]	33.36									
No overlap	[0.45]	[0.50]	20.01		[0.50]	[1.40]	14.78									

Notes: This table reports forecasting results obtained on currency portfolios using the 12-month change in Industrial Production and either the portfolio 12-month forward discount or the average 12-month forward discount. We report the R^2 in the time-series regressions of the log 12-month currency excess return on the log forward discount for each portfolio j: $rx^{j,12}_{net,t+12} = \kappa^j_0 + \kappa^j_1(f^{j,12}_t - s^j_t) + \kappa^j_1 \Delta IP_{t-12,t} + \eta^j_t$. The left panel uses our sample of developed and emerging countries. The right panel uses our sample of developed countries. The Newey and West (1987) (NW) standard errors are computed with the optimal number of lags. W is the Wald-test χ^2 statistic for the slope coefficients. The Hansen and Hodrick (1980b) (HH) standard errors are computed with 12 lags for the 12-month returns. For the bootstrapped standard errors, the VAR uses 12 lags for the 12-month returns. All the returns are annualized and reported in percentage points. Data are monthly, from Datastream and Global Financial Data. The sample period is 11/1983 - 03/2008.

Table 11: CAPM in Crisis

$Portfolio$	α_m^i	β_m^i	$p(\%)$	R^2	α_m^i	β_m^i	$p(\%)$	R^2	α_m^i	β_m^i	$p(\%)$	R^2	α_m^i	β_m^i	$p(\%)$	R^2
$Sample$	26-May-1998				02-Aug-1995				10-Oct-1999				31-Aug-2007			
1	−1.13	0.02	86.16	0.10	4.24	−1.22	0.09	18.20	−0.16	−0.13	16.91	7.33	0.15	−0.13	1.38	11.85
	[0.62]	[0.14]			[1.57]	[0.37]			[0.57]	[0.09]			[0.38]	[0.05]		
2	−0.64	−0.05	75.70	0.59	3.48	−0.90	8.76	8.52	−0.45	−0.11	5.19	9.30	0.17	0.21	0.04	27.84
	[0.92]	[0.16]			[1.90]	[0.53]			[0.35]	[0.05]			[0.37]	[0.06]		
3	−1.45	0.21	11.09	10.97	3.51	−0.89	7.88	11.97	0.85	−0.05	34.63	1.93	0.74	0.18	0.02	28.38
	[0.71]	[0.13]			[1.80]	[0.50]			[0.34]	[0.05]			[0.27]	[0.05]		
4	−1.43	0.28	2.50	13.55	2.21	−0.48	5.52	11.88	−0.24	−0.23	3.95	29.24	0.31	0.21	0.00	40.08
	[0.59]	[0.12]			[0.83]	[0.25]			[0.22]	[0.11]			[0.25]	[0.03]		
5	−1.81	0.50	0.00	23.41	2.14	−0.55	5.20	10.14	−0.40	0.06	22.28	4.82	0.51	0.25	0.00	45.52
	[0.47]	[0.11]			[0.92]	[0.28]			[0.30]	[0.05]			[0.23]	[0.04]		
6	−3.84	1.14	0.00	23.41	0.42	−0.00	98.46	10.14	0.80	0.25	0.00	4.82	0.44	0.50	0.00	45.52
	[1.53]	[0.27]			[0.43]	[0.14]			[0.48]	[0.05]			[0.43]	[0.10]		
HML_{FX}	−2.71	1.11	0.00	20.15	−3.82	1.22	0.02	11.24	0.96	0.37	0.03	20.87	0.29	0.62	0.00	56.12
	0.60	0.16			1.38	0.33			0.75	0.10			[0.38]	[0.08]		

Notes: This table reports results OLS estimates of the factor betas. The sample period is 129 days (6 months) before and including the mentioned date. The intercept α_0 β, and the R^2 are reported in percentage points. The standard errors in brackets are Newey-West standard errors computed with the optimal number of lags. The p-value is for a t-test on the slope coefficient. The portfolios are constructed by sorting currencies into six groups at time t based on the one-month forward discount (i.e nominal interest rate differential) at the end of period $t-1$. The returns are 1-month returns, and take into account bid-ask spreads. Portfolio 1 contains currencies with the lowest interest rates. Portfolio 6 contains currencies with the highest interest rates. Data are daily, from Barclays and Reuters in Datastream. We use the value-weighted return on the US stock market (CRSP).

Table 12: Calibration

Panel I: Moments		
Moment		*Value*
Mean real interest rate	$E[r]$	1.2%
Std real interest rate	$Std[r]$	.2%
Autocorr. real interest rate	$\rho[r]$	.66
Mean volatility log SDF	$E[\sigma_t(m_{t+1})]$	.35
Std volatility log SDF	$Std\,[\sigma_t(m_{t+1})]$	.04
Std changes in real exchange rates	$Std[\Delta q_{t+1}]$	10%
UIP slope coefficient	β_{UIP}	-1

Panel II: Real SDF Parameters						
λ	γ	τ	δ	ϕ	θ	$\sigma(\%)$
1.01	0.68	8.17	14.75	0.96	0.00	0.19

Panel III: Inflation Moments	
Moment	*Value*
Mean World inflation rate	3%
Std World inflation rate	2.1%
Autocor. World Inflation	0.87
Mean Country Inflation Rate	3%
Std Country Inflation Rate	2.4%
Autocor. Country Inflation	0.7
Std Country-Specific component	10%
Autocor. Country-Specific component	0.5

Panel IV: Inflation Parameters						
$\sigma^{w\$}(\%)$	ρ^w	$\overline{\pi^w}(\%)$	$\sigma^{\$}(\%)$	$\rho^{\$}$	$\overline{\pi}(\%)$	μ
0.03	0.98	0.25	0.43	0.90	0.25	0.16

This table reports the annualized moments of the real variables (Panel I), as well as the corresponding parameters used in calibration (Panel II). The moments in Panel I are: mean, standard deviation and autocorrelation of the (nominal) risk-free rate, mean and variance of the conditional variance of the real SDF, average real exchange rate volatility and the coefficient from the regression of the exchange rate change on the forward discount, both real (the latter two moments are averages across all foreign countries). All countries share the same parameters except for δ. The parameters δ^i are linearly spaced on the interval $[0.5\delta, 1.5\delta]$. Panel III reports the moments of the common and country-specific inflation processes, and panel IV - the corresponding inflation process parameters (see appendix for details).

Table 13: Currency Portfolios - Simulated data

Portfolio	1	2	3	4	5	6
			Spot change: Δs^j			
Mean	0.87	0.69	0.37	0.31	0.20	-0.09
Std	9.62	8.74	7.89	7.01	7.70	8.81
			Forward Discount: $f^j - s^j$			
Mean	-2.23	-1.38	-0.65	0.08	0.82	1.86
Std	0.54	0.39	0.23	0.17	0.29	0.51
			Excess Return: rx^j			
Mean	-3.10	-2.08	-1.02	-0.23	0.62	1.95
Std	9.65	8.79	7.92	7.03	7.72	8.84
SR	-0.32	-0.24	-0.13	-0.03	0.08	0.22
			High-minus-Low: $rx^j - rx^1$			
Mean		1.02	2.08	2.87	3.72	5.05
Std		4.63	5.72	7.20	9.09	11.42
SR		0.22	0.36	0.40	0.41	0.44

	λ_{RX}	$\lambda_{HML_{FX}}$	b_{RX}	$b_{HML_{FX}}$	R^2	$RMSE$	χ^2
GMM_1	-0.64	5.05	-0.08	0.32	99.73	0.08	
	[0.34]	[0.56]	[0.06]	[0.04]			79.06
GMM_2	-0.64	5.03	-0.08	0.32	99.73	0.08	
	[0.34]	[0.55]	[0.06]	[0.04]			79.07
FMB	-0.64	5.05	-0.08	0.32	99.65	0.08	
	[0.35]	[0.57]	[0.06]	[0.04]			77.45
	[0.35]	[0.57]	[0.06]	[0.04]			77.99
Mean	**-0.6**	**5.05**					

Notes: This table reports, for each portfolio j, the average change in log spot exchange rates Δs^j, the average log forward discount $f^j - s^j$, the average log excess return rx^j and the average return on the long short strategy $rx^j - rx^1$. Log currency excess returns are computed as $rx^j_{t+1} = -\Delta s^j_{t+1} + f^j_t - s^j_t$. All moments are annualized monthly values and reported in percentage points. For excess returns, the table also reports annualized Sharpe ratios. The portfolios are constructed by sorting currencies into six groups at time t based on the one-year forward discount (i.e nominal interest rate differential) at the end of period $t-1$. Portfolio 1 contains currencies with the lowest interest rates. Portfolio 6 contains currencies with the highest interest rates. All data are simulated from the model.

Table 14: Currency Momentum Portfolios - US Investor

Portfolio	1	2	3	4	5	6	1	2	3	4	5
	Panel I: All Countries						Panel II: Developed Countries				
	Spot change: Δs^j						Δs^j				
Mean	4.48	0.17	−0.49	−2.12	−2.22	−1.92	−0.37	−1.02	−2.27	−3.83	−1.69
Std	10.28	8.34	8.24	7.67	8.20	8.35	9.47	9.85	9.92	9.66	8.79
	Forward Discount: $f^j - s^j$						$f^j - s^j$				
Mean	0.47	0.70	1.39	1.32	1.86	3.39	0.16	0.42	0.71	0.86	1.54
Std	1.92	0.80	1.69	0.79	0.83	1.26	1.04	0.98	0.94	0.78	0.78
	Excess Return: rx^j (without b-a)						rx^j (without b-a)				
Mean	−4.01	0.53	1.88	3.44	4.08	5.31	0.53	1.44	2.98	4.69	3.23
Std	10.30	8.38	8.25	7.77	8.31	8.42	9.55	9.92	10.03	9.74	8.93
SR	−0.39	0.06	0.23	0.44	0.49	0.63	0.06	0.15	0.30	0.48	0.36
	Net Excess Return: rx^j_{net} (with b-a)						rx^j_{net} (with b-a)				
Mean	−1.98	−0.83	0.43	2.05	2.73	3.44	1.88	0.11	1.64	3.37	1.77
Std	10.26	8.36	8.21	7.74	8.28	8.42	9.56	9.90	10.04	9.73	8.92
SR	−0.19	−0.10	0.05	0.26	0.33	0.41	0.20	0.01	0.16	0.35	0.20
	High-minus-Low: $rx^j - rx^1$ (without b-a)						$rx^j - rx^1$ (without b-a)				
Mean		4.54	5.90	7.45	8.09	9.32		0.91	2.45	4.16	2.70
Std		8.69	8.97	9.36	10.00	10.79		7.14	7.53	8.32	8.63
SR		0.52	0.66	0.80	0.81	0.86		0.13	0.33	0.50	0.31
	High-minus-Low: $rx^j_{net} - rx^1_{net}$ (with b-a)						$rx^j_{net} - rx^1_{net}$ (with b-a)				
Mean		1.15	2.41	4.03	4.71	5.42		−1.76	−0.24	1.49	−0.11
Std		8.66	8.91	9.34	9.98	10.75		7.13	7.53	8.31	8.61
SR		0.13	0.27	0.43	0.47	0.50		−0.25	−0.03	0.18	−0.01

Notes: This table reports, for each portfolio j, the average change in log spot exchange rates Δs^j, the average log forward discount $f^j - s^j$, the average log excess return rx^j without bid-ask spreads, the average log excess return rx^j_{net} with bid-ask spreads, and the average return on the long short strategy $rx^j_{net} - rx^1_{net}$ and $rx^j - rx^1$ (with and without bid-ask spreads). Log currency excess returns are computed as $rx^j_{t+1} = -\Delta s^j_{t+1} + f^j_t - s^j_t$. All moments are annualized and reported in percentage points. For excess returns, the table also reports Sharpe ratios, computed as ratios of annualized means to annualized standard deviations. The portfolios are constructed by sorting currencies into six groups at time t based on the return realized at the end of period $t-1$. Portfolio 1 contains currencies with the lowest returns. Portfolio 6 contains currencies with the highest returns. Panel I uses all countries, panel II uses developed countries only. Data are monthly, from Barclays and Reuters (Datastream). The sample period is 11/1983 - 03/2008.

Table 15: Asset Pricing - US Investor - Carry and Momentum Currency Portfolios

	All Countries															
	3 Factors									2 Factors						
	λ_c	λ_m	λ_d	b_c	b_m	b_d	R^2	$RMSE$	χ^2	λ_c	λ_d	b_c	b_d	R^2	$RMSE$	χ^2
	All 12 Portfolios															
FMB	10.00	3.62	1.51	0.47	0.21	0.28	83.23	0.70		10.00	1.51	0.47	0.28	68.13	1.02	
	[2.70]	[2.44]	[1.37]	[0.13]	[0.14]	[0.25]			49.82	[2.70]	[1.37]	[0.13]	[0.25]			39.19
	[2.70]	[2.44]	[1.37]	[0.13]	[0.14]	[0.25]			54.45	[2.70]	[1.37]	[0.13]	[0.25]			43.55
$Mean$	**10.00**	**3.62**	**1.51**							**10.00**	**1.51**					
	6 Carry Portfolios															
FMB	13.02	5.63	1.34	0.61	0.32	0.25	75.39	0.74		9.96	1.51	0.47	0.28	76.13	0.84	
	[4.22]	[5.54]	[1.38]	[0.20]	[0.32]	[0.25]			14.00	[3.30]	[1.38]	[0.16]	[0.25]			17.89
	[4.34]	[5.76]	[1.38]	[0.20]	[0.33]	[0.25]			17.38	[3.33]	[1.38]	[0.16]	0.25			20.06
$Mean$	**10.00**	**3.62**	**1.51**							**10.00**	**1.51**					
	6 Momentum Portfolios															
FMB	6.65	4.76	1.63	0.31	0.27	0.30	96.01	0.29		10.05	1.50	0.47	0.27	50.98	1.18	
	[3.61]	[2.61]	[1.37]	[0.17]	[0.15]	[0.25]			85.93	[3.68]	[1.38]	[0.17]	[0.25]			38.76
	[3.64]	[2.61]	[1.38]	[0.17]	[0.15]	[0.25]			86.60	[3.72]	[1.38]	[0.18]	[0.25]			41.49
$Mean$	**10.00**	**3.62**	**1.51**							**10.00**	**1.51**					

Notes: The momentum portfolios are constructed by sorting currencies into six groups at time t based on the return realized at the end of period $t-1$. Portfolio 1 contains currencies with the lowest returns. Portfolio 6 contains currencies with the highest returns. The risk factors are the third (the carry factor denoted c), the second (the momentum factor denoted m) and the first principal component (the dollar factor denoted d) of the 12 currency portfolios. The test assets are the six carry and the six momentum currency portfolios. The first subpanel reports results Fama-McBeth asset pricing procedures using all 12 test assets. The second subpanel uses only the 6 carry trade portfolios as test assets. The third subpanel uses only the 6 momentum portfolios as test assets. Market prices of risk λ, the adjusted R^2, the square-root of mean-squared errors $RMSE$ and the p-values of χ^2 tests on pricing errors are reported in percentage points. b denotes the vector of factor loadings. Excess returns used as test assets and risk factors take into account bid-ask spreads. All excess returns are multiplied by 12 (annualized). The standard errors in brackets are Newey and West (1987) standard errors computed with the optimal number of lags according to Andrews (1991). Shanken (1992)-corrected standard errors are reported in parentheses. We do not include a constant in the second step of the FMB procedure. R^2s and p-values are reported in percentage points. The χ^2 test statistic $\alpha' V_\alpha^{-1} \alpha$ tests the null that all intercepts are jointly zero. This statistic is constructed from the Newey-West variance-covariance matrix (1 lag) for the system of equations (see Cochrane (2001), p. 234). Data are monthly, from Barclays and Reuters in Datastream. The sample period is 11/1983 - 03/2008.

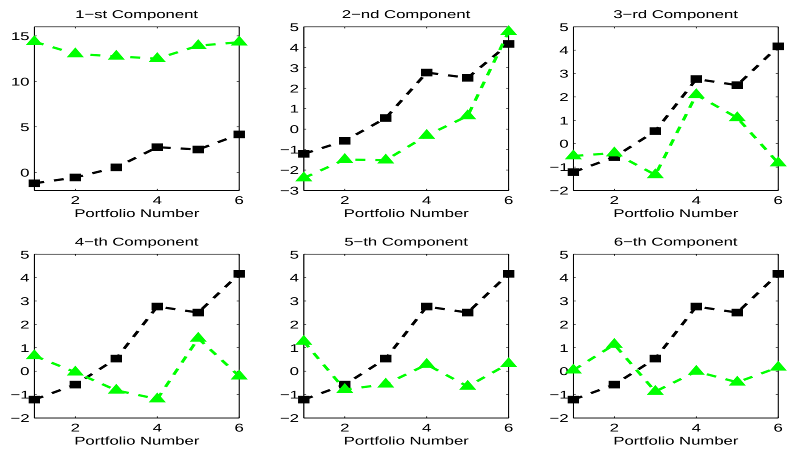

Figure 1: Mean Excess Returns and Covariances between Excess Returns and Principal Components - Developed and Emerging Countries

Each panel corresponds to a principal component. The upper left panel uses the first principal component. The black squares represent the average currency excess returns for the six portfolios. Each green triangle represents a covariance between a given principal component and a given currency portfolio. The covariances are rescaled (multiplied by 15,000). The average excess returns are annualized (multiplied by 12) and reported in percentage points. The sample is 11/1983 - 03/2008.

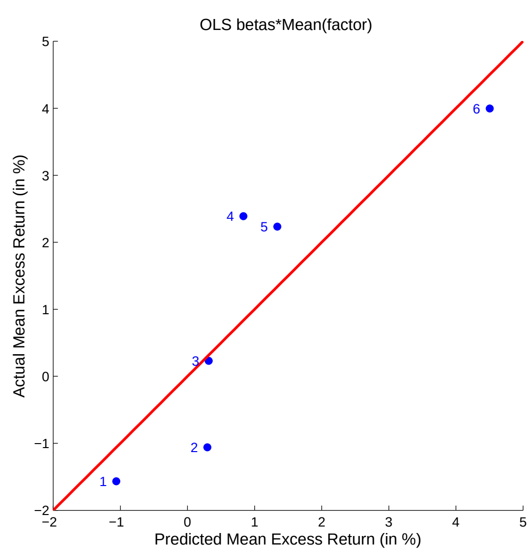

Figure 2: Predicted against Actual Excess Returns.

This figure plots realized average excess returns on the vertical axis against predicted average excess returns on the horizontal axis. We regress each actual excess return on a constant and the risk factors RX and HML_{FX} to obtain the slope coefficient β^j. Each predicted excess returns is obtained using the OLS estimate of β^j times the sample mean of the factors. All returns are annualized. The date are monthly. The sample is 11/1983 - 03/2008.

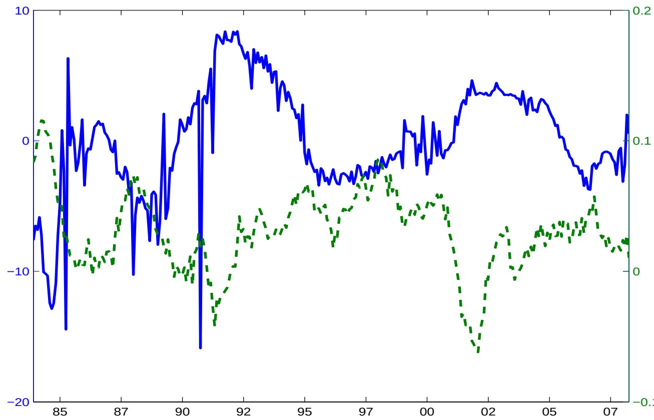

Figure 3: Forecasted Excess Return in Currency Markets and US Business Cycle.

This figure plots the one-month ahead forecasted excess returns on portfolio 2 ($\widehat{E}_t rx^2_{t+1}$). All returns are annualized. The dashed line is the year-on-year log change in US Industrial Production Index.

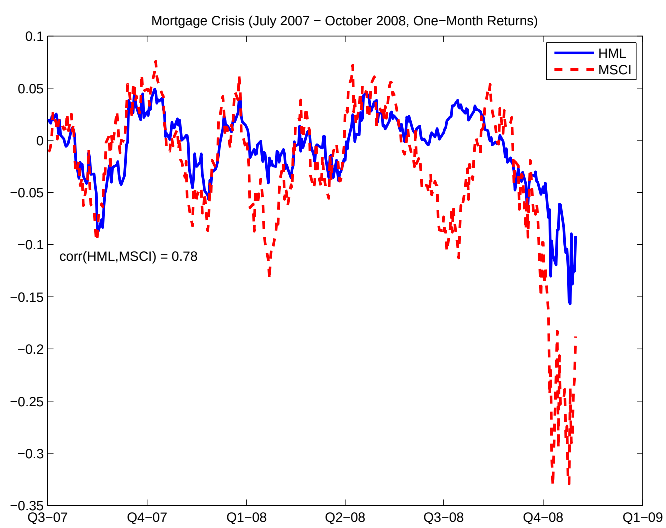

Figure 4: Carry Trade and US Stock Market Returns during the Mortgage Crisis - July 2007 to March 2008.

This figure plots the one-month HML_{FX} return at daily frequency against the one-month return on the US MSCI stock market index at daily frequency. The sample is 07/02/07-03/31/08.

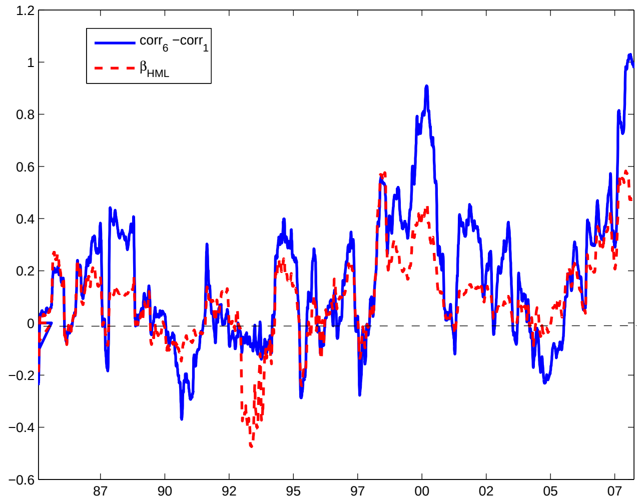

Figure 5: Market Correlation Spread of Currency Returns

This figure plots $Corr_\tau[R_t^m, rx_t^6] - Corr_\tau[R_t^m, rx_t^1]$, where $Corr_\tau$ is the sample correlation over the previous 12 months $[\tau - 253, \tau]$. We use monthly returns at daily frequency. We also plot the stock market beta of HML_{FX}, β_{HML} . The stock market return is the return on the value-weighted US index (CRSP).

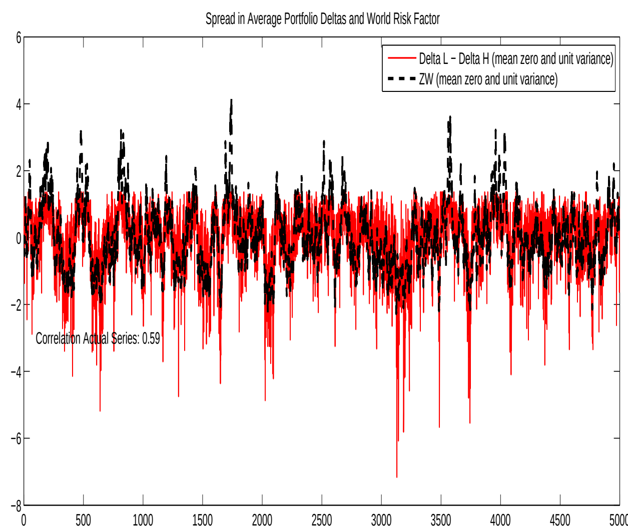

Figure 6: Spreads in Portfolio Deltas and World Risk Factor - Simulated Data.

This figure plots the difference between the average delta in the first portfolio and the average delta in the last portfolio, along with the world risk factor ZW. Both series are centered and scaled by their standard deviations.

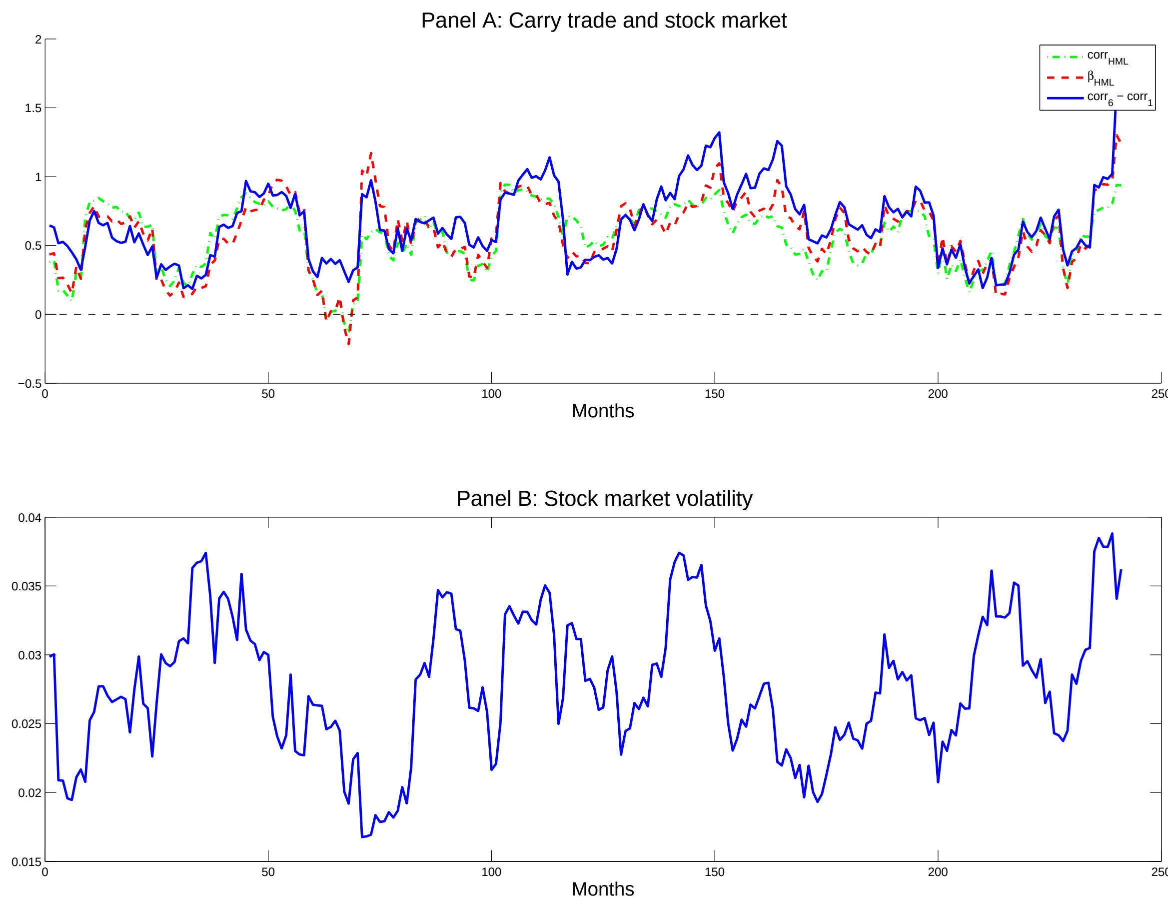

Figure 7: Stock market risk of the carry trade

This figure plots the conditional risk measures implied by the calibrated model. Panel A displays the conditional correlations and betas of the carry trade factor HML_{FX} with the stock market return simulated from the model. $corr_6 - corr_1$ denotes the difference in conditional correlations with the stock market return between the highest interest rate portfolio and the lowest interest rate portfolio, for a 20-year period (using monthly data). These quantities are estimated from simulated data using rolling 12-month windows. Panel B plots the standard deviation of the stock market return using the same rolling windows as the estimated betas.

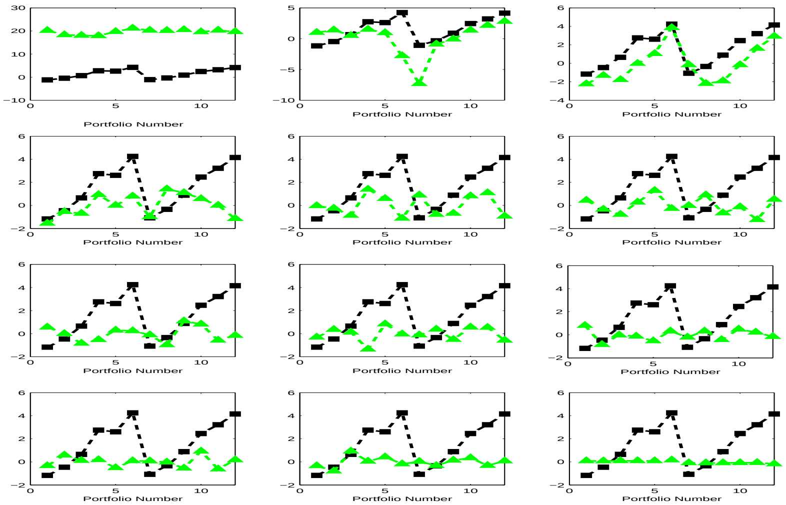

Figure 8: Mean Excess Returns and Covariances between Excess Returns and Principal Components - Carry and Momentum Currency Portfolios

Each panel corresponds to a principal component of the 6 carry trade portfolios (1-6) and the 6 momentum portfolios (7-12). The upper left panel uses the first principal component. The lower right panel uses the 12-th principal component. The black squares represent the average currency excess returns for the twelve portfolios. Each green triangle represents a covariance between a given principal component and a given currency portfolio. The covariances are rescaled (multiplied by 15,000). The average excess returns are annualized (multiplied by 12) and reported in percentage points. The sample is 11/1983 - 03/2008.

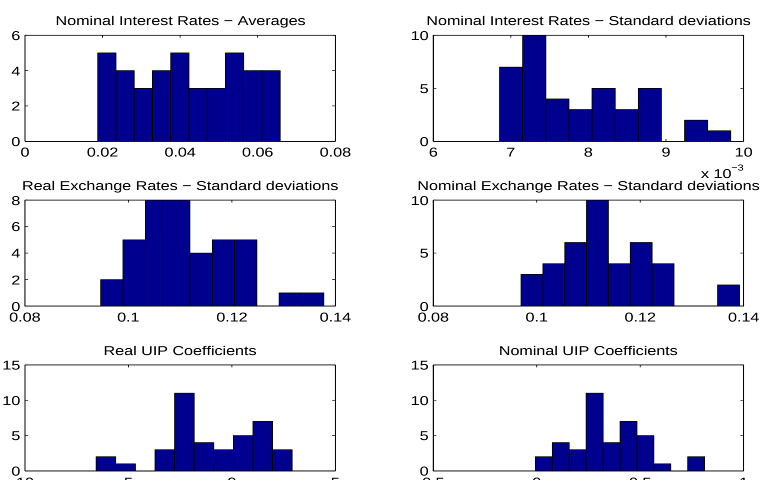

Figure 9: Interest Rates, Exchange Rates, and UIP Slope Coefficients - Simulated Data.

This figure plots several histograms summarizing our simulated data. We report the distributions of the interest rates' first two moments, the volatility of real and nominal exchange rates and the UIP slope coefficients.