Pricing currency risks*

Mikhail Chernov,† Magnus Dahlquist,‡ and Lars A. Lochstoer§

First Draft: May, 2020

This Draft: January 26, 2021

Comments are welcome!

Abstract

The currency market features a relatively small cross-section and conditional expected returns can be characterized by only a few signals – interest differentials, trend and mean-reversion. We exploit these properties to construct a conditional projection of the stochastic discount factor onto excess returns of individual currencies. Our approach is implementable in real time and prices all currencies and prominent strategies conditionally as well as unconditionally. We document that the fraction of unpriced risk in these assets is at least 85%. Extant explanations of carry strategies based on intermediary capital or global volatility are related to these unpriced components, while consumption growth is related to the priced component of returns.

JEL Classification Codes: F31, G12, G15.

Keywords: currency risk premiums, stochastic discount factor, factor models.

---

*We thank Valentin Haddad, Serhiy Kozak, Francis Longstaff, Hanno Lustig, Tyler Muir, Adrien Verdelhan and Irina Zviadadze for comments on earlier drafts, as well as participants in the seminar at UCLA. We thank Felix Wilke for excellent research assistance. Dahlquist gratefully acknowledges support from the Jan Wallander and Tom Hedelius Foundation. The latest version is available here.

†UCLA, NBER, and CEPR; mikhail.chernov@anderson.ucla.edu.

‡Stockholm School of Economics, and CEPR; magnus.dahlquist@hhs.se.

§UCLA; lars.lochstoer@anderson.ucla.edu.

1 Introduction

In this paper we argue that research on cross-sectional currency pricing can depart fruitfully from the factor mining path established by the equity literature. That is because direct *conditional* projection of the stochastic discount factor (SDF), as advocated by Hansen and Richard (1987), is feasible in the exchange-rate setting. The value of this approach, if it is empirically successful, is obvious. One obtains a measure of the SDF that prices, both conditionally and unconditionally, the full cross-section of currencies and trading strategies associated with them. This measure can help with directing research that tries to develop an understanding of currency risk pricing.

Specifically, the projection approach requires two critical ingredients: estimates of the conditional mean and covariance matrix of currency returns. This is where exchange rates have an advantage over equities. First, the size of the cross-section of exchange rates is small, not exceeding 40 as compared to thousands of stocks. That feature dramatically affects the precision of the covariance matrix regardless of the specifics of the estimation method. Second, exchange rates have a very particular economic structure that helps with hypothesizing the functional form of the conditional means. Research that goes back decades suggests the importance of the interest rate differential (Covered Interest Parity, Keynes, 1923; Uncovered Interest Parity, Porter, 1971; Random Walk Hypothesis, Meese and Rogoff, 1983), real exchange rate (Purchasing Power Parity, Cassel, 1918), and trend in nominal exchange rate (weak form of market efficiency, Cornell and Dietrich, 1978). These three drivers remain the key ingredients in the modern currency trading strategies under the names of Carry, Value, and Momentum, respectively.

We work with a sample of G10 currencies, the most commonly used data in the literature, with monthly returns based on forward rates from January 1985 through May 2020. We construct monthly conditional expected excess returns using the three aforementioned signals. The loading on the first signal is one as motivated by the random walk hypothesis, while the loadings on the other signals are estimated on a rolling basis via panel regressions. We construct a conditional covariance matrix of currency returns using daily data within each month. Importantly, the overall approach uses data that are available to investors in real time.

We consider nine leading trading strategies. The Dollar strategy represents the currency version of "the market" (Lustig, Roussanov, and Verdelhan, 2011). The Dollar Carry goes long or short the Dollar factor depending on the average carry signal (Lustig, Roussanov, and Verdelhan, 2014). Next, we evaluate the cross-sectional and time-series Carry strategies (Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2011, Daniel, Hodrick, and Lu, 2017, Koijen, Moskowitz, Pedersen, and Vrugt, 2018, Lustig, Roussanov, and Verdelhan, 2011). These strategies are followed by two versions of cross-sectional and time-series Momentum depending on whether one uses the past month or year as the measure of past performance (Asness, Moskowitz, and Pedersen, 2013, Burnside, Eichenbaum, and Rebelo, 2011, Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2011, Menkhoff, Sarno, Schmeling, and Schrimpf, 2012b, Moskowitz, Ooi, and Pedersen, 2012). Lastly, we study cross-sectional Value that sorts currencies based on the real exchange rate (Asness, Moskowitz, and Pedersen, 2013). As emphasized by Asness, Moskowitz, and Pedersen (2013) and Koijen, Moskowitz, Pedersen, and Vrugt (2018), besides being the most common strategies in FX, they are more broadly considered in other asset classes.

The estimated SDF can be represented as a linear function of the unconditional

mean-variance efficient portfolio (UMVE). The UMVE is tradeable in real time, has a Sharpe ratio in excess of one, and, conveniently, allows us to compare directly to existing trading strategies using standard regression-based tests. We verify that it prices both excess returns on nine individual exchange rates vs. the U.S. dollar and the nine above-mentioned trading strategies via the unconditional and conditional pricing tests of Gibbons, Ross, and Shanken (1989) and Chernov, Lochstoer, and Lundeby (2018), respectively.

Having tested the SDF, we proceed with evaluating its properties. The UMVE portfolio, which is perfectly negatively correlated with the SDF, has approximately zero skewness of -0.04 and modest excess kurtosis of 3.1. Thus, there is little prima facie evidence of crash risk or extreme events as being key for risk pricing in our sample. Further, the UMVE explains only a small fraction of the time-series variation in strategy returns. For instance, less than 1% of the variation in the Dollar factor returns (the currency "market factor") are priced, while the same for Carry strategies is at most 15%. Thus, there are large unpriced components in strategy returns and we show that hedging out these unpriced components in an out-of-sample manner dramatically increases Sharpe Ratios of the hedged returns. As a model-check, we verify that the average returns to the (out-of-sample) unpriced components are indeed not significantly different from zero.

Next, we document that the factors that are the most important for time-series co-variation among currencies and the currency trading strategies have little relation to the UMVE portfolio. In other words, the main factors that drive the covariance matrix of returns are not an important source of risk for the marginal currency investor. For instance, the first two principal components of individual currency returns capture 75% of the return variation and are strongly associated with the

Dollar and cross-sectional Carry strategies, consistent with Verdelhan (2018).[1] However, these principal components explain only around 5% of the variation in UMVE returns. In fact, all nine of these principal components only explain around 10% of the UMVE returns indicating that conditional currency risk premiums vary strongly and underscoring the importance of correct conditioning.

The currency strategies involve a lot of timing and therefore potentially accounts for these conditional dynamics. We investigate whether their principal components are capable of explaining variation in the UMVE. The answer is no, although there is some improvement. The first three principal components of the currency trading strategies explain around 20% of the variation in the UMVE returns. Using all the principal components, this increases to 50%. This evidence is consistent with the uncovered large idiosyncratic component of strategy returns. Taken together, this suggests a large role of optimal timing implicit in the construction of the UMVE above and beyond the timing already implicit in existing trading strategies.

To give a sense of the magnitudes, consider the standard deviation of a strategy's conditional expected returns. The average of this number across strategies, annualized, is 3.4%. For comparison, the average unconditional return across strategies is 3.2%. The large variation in conditional returns suggests a substantial benefit also to individual factor timing. We do so optimally in an out-of-sample fashion by scaling a position in each strategy up or down depending on our estimates of the individual currency conditional expected returns and covariance matrix. In six cases alphas of the timed strategies with respect to their original counterparts are significant with large Information ratios. Thus, there are large gains in terms of the unconditional

[1] Aloosh and Bekaert (2019) and Greenaway-McGrevy, Mark, Sul, and Wu (2018) offer alternative approaches to analysis of time-series behavior of exchange rates.

Sharpe ratio to timing existing factors proposed in the literature.

General equilibrium models have struggled with explaining FX moments. The current generation of models based on habit (Verdelhan, 2010), long-run risks (Bansal and Shaliastovich, 2013; Colacito, Croce, Gavazzoni, and Ready, 2018), and rare disasters (Farhi and Gabaix, 2016) tend to aim at explaining the forward premium puzzle (and to some extent the carry trade). However, as our analysis shows, the most popular carry factors together explain only about 20% of the variation in our pricing factor. Thus, there is much left to explain and our empirical model can help guide the specification and tests of future models on currency risk and return.

To that end, we first regress the UMVE returns on various candidate explanatory factors. We find that quarterly and 3-year consumption growth, a proxy for long-run consumption risk, both are significantly positively related to the UMVE returns, consistent with Lustig and Verdelhan (2007) and Zviadadze (2017). We next find that the factors in the Fama-French 5-factor model are only weakly related to the currency UMVE portfolio, with only the HML and CMA portfolios having a significant relation. However, here the R^2s are only about 1%, so economically there appears to be effectively no relation between priced risks in the equity and currency markets. These results extend findings in Burnside (2012), who shows a similar weak relation between the Carry trade returns and equity factors, to the UMVE portfolio. We also find that intermediary capital factors and shocks to equity and currency variance are unrelated to the UMVE returns. This may appear puzzling given that previous literature has highlighted these factors as potential explanations of the carry trade, but we show that this correlation comes from the large unpriced component in carry factor returns.

Finally, we show that the conditional price of risk, as given by the volatility of our currency market SDF, has a downward trend over the sample, largely caused by a trend decline in interest differentials across countries. The cyclical variation in the price of risk is *negatively* related to measures of the conditional variance of consumption and equity returns, and currency depreciation. In sum, our paper provides a rich set of facts that can help guide future developments of equilibrium models of currency market risk and return.

Related literature. A number of papers consider conditional mean-variance efficient portfolios (CMVE) using Carry alone as a signal (Ackermann, Pohl, and Schmedders, 2017, Baz, Breedon, Naik, and Peress, 2001, Daniel, Hodrick, and Lu, 2017, Mauer, To, and Tran, 2018, 2020). Further, with the exception of Mauer, To, and Tran (2018), these papers do not test if CMVE explains other versions of Carry or individual currency returns. We reject this version of our model with respect to our full model, which includes Value and Momentum, using the Barillas and Shanken (2017) test. We also reject it using the conditional pricing test, even when using Carry-based strategies only as test assets. Della Corte, Sarno, and Tsiakas (2009) depart from the lone Carry signal by considering three monetary fundamentals variables. They perform an out-of-sample analysis of exchange rate predictability in a mean-variance framework where the allocation choice is between the risk-free asset and one of the three exchange rates they consider.

A large literature implements model-free SDF projections in the context of unconditional pricing of assets. Many authors use Hansen and Jagannathan (1991) bounds as a diagnostic of their models, but not as a standalone tool for SDF estimation. A variation on the unconditional version of the Hansen and Richard (1987) approach is that of entropy minimization advocated by Stutzer (1996) in the context of deriva-

tives pricing. Ghosh, Julliard, and Taylor (2019) develop and apply this framework to cross-sectional asset pricing. Korsaye, Trojani, and Vedolin (2020), Orlowski, Sokolowski, and Sverdrup (2019), and Sandulescu, Trojani, and Vedolin (2020) are examples of international applications of this framework. We critically depart from this work by constructing real-time conditional SDF projections. In particular, that allows valuation of any trading strategies in the set of assets involved in the projection.

2 Linear factor models and exchange rates

Linear factor models are popular in empirical asset pricing, in general, and in currency pricing in particular. The pursuit of such a model has two ultimate objectives. The first objective, common across all of asset pricing, is to achieve correct conditional pricing. That is, find a stochastic discount factor (SDF), $M_{t,t+1}$, such that

$$E_t(M_{t,t+1}R^{ei}_{t,t+1}) = 0, \tag{1}$$

where $R^{ei}_{t,t+1}$ is a one-period excess return on asset i.

The second objective is to find a parsimonious factor representation of the risk-return trade-off. To make an extreme argument, having as many factors as assets does not help with understanding risk pricing in the economy. Further, a parsimonious representation is helpful in developing economic explanations of risk pricing. That is because it offers a succinct set of quantitative benchmarks, which an equilibrium model should match.

2.1 Conditional MVE portfolio

The ultimate parsimonious factor model is a single-factor one. The projection approach of Hansen and Richard (1987) offers a way to construct such a factor via a CMVE portfolio (also, see Cochrane and Hansen, 1992 and Cochrane, 2004). Stack excess returns $R^{ei}_{t,t+1}$ on all assets into an $N_t \times 1$ vector $R^{e}_{t,t+1}$. Then

$$\begin{aligned} R^{C}_{t,t+1} &= (w^{C}_{t})^{\top} R^{e}_{t,t+1}, & (2) \\ w^{C}_{t} &= k_t^{-1} V_t^{-1}\left(R^{e}_{t,t+1}\right) E_t\left(R^{e}_{t,t+1}\right) & (3) \end{aligned}$$

is a CMVE portfolio. Here k_t is any positive constant known at time t, governing the leverage of the portfolio. In particular, given w^C_t and $R^C_{t,t+1}$ above, this constant governs the ratio of the conditional mean to the conditional variance,

$$k_t = V_t^{-1}\left(R^{C}_{t,t+1}\right) E_t\left(R^{C}_{t,t+1}\right).$$

The CMVE portfolio "prices" any combination of these assets. That is, a projected SDF

$$M_{t,t+1} = 1 - k_t(R^{C}_{t,t+1} - E_t(R^{C}_{t,t+1})) \quad (4)$$

satisfies Equation (1). That there is a single factor driving this SDF does not mean that it prices only one source of risk in the economy. The factor is capturing the correct combination of the various risks that is priced.

The issue with this approach is practical. Implementation requires estimation of the conditional covariance matrix and conditional mean of a potentially large number

of assets. The literature has responded by searching for a small number of factors whose linear combination with constant loadings in the SDF price assets unconditionally. That is, researchers use a low-dimensional factor approach to approximate the *unconditional* MVE portfolio. See Chernov, Lochstoer, and Lundeby (2018) for a recent distillation of the approach.

2.2 Exchange rates

Exchange rates differ from equities in two very important dimensions. First, the cross-section of "assets" is limited. Indeed, the largest data panel considered in the literature goes up to 39 currencies at any given point in time (e.g., Hassan and Mano, 2019, Lustig, Roussanov, and Verdelhan, 2011, Menkhoff, Sarno, Schmeling, and Schrimpf, 2012a), while most studies limit their attention to ten currencies representing developed economies.

Second, a voluminous literature on the forecasting of exchange rates suggests that there are few variables outperform the random walk hypothesis (RWH) out of sample. That observation, perhaps paradoxically, simplifies the modeling effort required to construct conditional expectations. We exploit these differences from the equity literature and construct a projection of the SDF on currency excess returns directly via Equation (4). Before we show how we do it, we introduce notation and basic concepts pertaining to exchange rates.

Let the USD be the base and measurement currency. Let S_t^i and F_t^i denote the spot exchange rate and the one-month forward exchange rate of country i, respectively. All exchange rates are expressed in USD per unit of foreign currency. $M_{t,t+1}$ is the nominal USD-denominated SDF. Superscript i refers to a foreign country i.

The payoff of a forward contract (when buying one unit of the foreign currency) is $S^i_{t+1} - F^i_t$. One common way to scale this payoff to define excess return is to divide by F^i_t :

$$R^{ei}_{t,t+1} = (S^i_{t+1} - F^i_t)/F^i_t. \tag{5}$$

This definition implies that the amount of foreign currency bought is one "forward" dollar. Thus, this is an excess return to a trading strategy regardless of whether Covered Interest Parity (CIP) holds.

Next, we proceed with estimation of $E_t(R^e_{t,t+1})$ and $V_t(R^e_{t,t+1})$. We use daily data within each month to construct our estimate of $V_t(R^e_{t,t+1})$ by combining the shrinkage method of Ledoit and Wolf (2020) along with an exponential moving average. See Appendix A.1 for details.

Our starting point for $E_t(R^e_{t,t+1})$ is the Random Walk Hypothesis (RWH) for spot rates. The reason for this is simple: the hypothesis leads to an exceptionally robust forecaster of exchange rates (e.g., Meese and Rogoff, 1983). This is also a natural straw man given the research in finance, starting with Baz, Breedon, Naik, and Peress (2001), that uses RWH to develop trading stratgies.

The RWH implies that expected excess currency returns are given by:

$$E_t(R^{ei}_{t,t+1}) = \gamma \cdot \left(S^i_t/F^i_t - 1\right)$$

with $\gamma = 1$. This is a particular violation of Uncovered Interest Parity (UIP), which posits $\gamma = 0$. Next, we add mean-reversion and trend signals for exchange rate forecasting. See Bilson (1984) and Sweeney (1986), respectively, for early contributions

in the context of trading strategies. Thus, the use of the RWH, as well as mean-reversion and trend signals, were all concepts that were in the public domain at the start of our sample in 1985.

Our mean-reversion signal is motivated by the literature on the role of real exchange rates (RER) in forecasting and capturing risk premiums. The RER is defined as

$$Q_t^i = S_t^i \cdot P_t^i / P_t, \tag{6}$$

where P_t and P_t^i are the U.S. and foreign price levels, respectively. The weak form of Purchasing Power Parity (PPP) implies mean-reversion in the RER. Thus, when the RER is far from its long-run mean it should forecast the currency depreciation.

Combining the RER-based signal with the RWH goes as far back as Bilson (1984). More recently, Chernov and Creal (2020) and Dahlquist and Penasse (2016) show that PPP implies that RER forecasts nominal depreciation rate and that currency risk premium depends on the RER. Eichenbaum, Johannsen, and Rebelo (2020) demonstrate that RER outperforms random walk in forecasting of exchange rates at horizons beyond one year.

As Jorda and Taylor (2012) emphasize, the RER's long-run mean is not a clearly defined object empirically. We divide each RER by its 5-year smoothed lag (specifically the average RER from 4.5 to 5.5 years ago) as a way to remove the dependence on the long-run mean while still preserving the long-run nature of mean-reversion signals:

$$\widetilde{Q}_t^i \equiv Q_t^i \cdot \left(\frac{1}{13} \sum_{j=-6}^{6} Q_{t-60+j}^i \right)^{-1}.$$

Lastly, we cross-sectionally demean the signal at each time t to create a cross-sectional

ranking of "cheap" and "expensive" currencies. That is, our signal is

$$z_{it}^{Q} \equiv \widetilde{Q}_t^i - \frac{1}{N}\sum_{i=1}^{N} \widetilde{Q}_t^i. \tag{7}$$

This definition has the virtue of removing any time and currency fixed-effects. The specific form of the reference point is inspired by the cross-sectional model of Asness, Moskowitz, and Pedersen (2013), who, in their turn, select the specifics to be comparable with the equity and commodity literatures.

Our trend signal is motivated by the academic and practitioner literature on using moving averages in trading and forecasting exchange rates. In contrast to the macro literature, the finance literature suggests that using past performance could be fruitful in improving trading performance. See, e.g., Kho (1996), Okunev and White (2003), Sweeney (1986), and references therein. Motivated by this work, we add a simple trend signal, which is a one-year depreciation rate.

In summary, we forecast excess returns out of sample via:

$$E_t(R_{t,t+1}^{ei}) = \gamma_t \cdot \left(S_t^i/F_t^i - 1\right) + \delta_t \cdot z_{it}^{Q} + \phi_t \cdot S_t^i/S_{t-12}^i.$$

We set $\gamma_t = 1$ to match the RWH baseline. The coefficients δ_t and ϕ_t are re-estimated every month using a panel of N_t exchange rates and an expanding sample. Estimating γ_t in a similar fashion does not change the results. This specification departs from the traditional forecasting exercises in the macro literature where all the coefficients are presumed constant.

2.3 Testing factor models

Tests

Factor models of expected returns are commonly tested using standard regression-based methods. The conditional SDF projection in Equation (4) implies that $\alpha = 0$ in the regression:

$$R^e_{t,t+1} = \alpha + \beta \cdot R^U_{t,t+1} + \varepsilon_{t,t+1}, \tag{8}$$

where, $R^U_{t,t+1}$ is the UMVE portfolio, $\varepsilon_{t,t+1}$ is an error term, and where the UMVE portfolio is derived from the CMVE portfolio in the SDF projection as follows:

$$R^U_{t,t+1} = (1 + \theta_t)^{-1} k_t R^C_{t,t+1}, \tag{9}$$

where θ_t is the maximal conditional squared Sharpe Ratio:

$$\theta_t = E_t(R^e_{t,t+1})^\top V_t(R^e_{t,t+1})^{-1} E_t(R^e_{t,t+1}).$$

Alternatively, since $R^C_{t,t+1} = (w^C_t)^\top R^e_{t,t+1}$, we can write the UMVE portfolio as a function of the conditional mean and variance of the underlying set of currencies:

$$R^U_{t,t+1} = (1 + \theta_t)^{-1} E_t \left(R^e_{t,t+1}\right)^\top V_t \left(R^e_{t,t+1}\right)^{-1} R^e_{t,t+1}. \tag{10}$$

See Chernov, Lochstoer, and Lundeby (2018) for derivation of these relationships.

The Gibbons, Ross, and Shanken (1989) test is a standard time-series joint test of $\alpha = 0$ for all test assets in this setting. It is asking if it is possible to improve

upon the Sharpe Ratio (SR) of the candidate UMVE portfolio by forming a portfolio consisting of both the factor and the test assets. Obviously, this is always possible in sample and so the test accounts for sampling uncertainty.

The typical testing challenge in the equity literature is to find a set of test portfolios that are sufficiently informative about a given model (e.g., Barillas and Shanken, 2017, Daniel and Titman, 2012, Lewellen, Nagel, and Shanken, 2010). In the case of currencies, major cross-sectional and time-series factors that were proposed in the literature are natural test assets. Intuitively, factors that reflect dynamic trading strategies are particularly informative about conditional properties of the SDF.

Furthermore, we follow Chernov, Lochstoer, and Lundeby (2018) and use multi-horizon returns (MHR) on the selected assets to generate additional test assets that are endogenous to the model of the SDF and pre-selected assets. Specifically, we test if

$$E(M_{t,t+h}R^i_{t,t+h}) = 1$$

for a range of horizons h.[2] The set of assets includes the same currency factors. In addition, we include MHR on the unconditional MVE. Chernov, Lochstoer, and Lundeby (2018) show that unconditional MHR pricing allows for testing most, if not all, aspects of conditional model misspecification.

[2]The multihorizon SDF is $M_{t,t+h} = \prod_{j=1}^{h} M_{t+j-1,t+j}$. The multihorizon gross return is $R^i_{t,t+h} = \prod_{j=1}^{h} R^i_{t+j-1,t+j}$, where $R^i_{t,t+1} = R^{ei}_{t,t+1} + R^f_{t,t+1}$ with the latter denoting the one-month U.S. gross funding rate.

Trading strategies

The portfolio excess return of a trading strategy is:

$$R^P_{t,t+1} = \sum_{i=1}^{N_t} w_{it} R^{ei}_{t,t+1}, \tag{11}$$

where w_{it} is the weight in currency i and N_t is the number of currencies at time t. The weight can be based on a signal, z_{it}, and chosen such that the portfolio has an exposure to the dollar or not. We use so-called rank weights for cross-sectional (CS) strategies and sign weights for time-series (TS) strategies. The CS strategies are dollar neutral; the TS strategies are dollar exposed. The rank and sign weights are used and discussed by Moskowitz, Ooi, and Pedersen (2012), Asness, Moskowitz, and Pedersen (2013), Koijen, Moskowitz, Pedersen, and Vrugt (2018), and others. An alternative to rank weights is "high-minus-low" portfolios with currencies being equal-, value-, or signal-weighted within the "high" and "low" portfolios (see, e.g., Lustig, Roussanov, and Verdelhan, 2011).

We consider nine leading strategies proposed in the literature. The Dollar strategy is an equal-weighted average of the individual currency (Lustig, Roussanov, and Verdelhan, 2011, Verdelhan, 2018). The Dollar Carry uses the average forward discount, S^i_t/F^i_t, as a signal. Specifically, it goes long (short) all currencies when the average discount is positive (negative) (Lustig, Roussanov, and Verdelhan, 2014).

The CS-Carry uses the currency's forward discount as a signal. Specifically, ranks are based on the signal and weights are based on the ranks according to

$$w_{it} = \kappa_t \left(\text{rank}(z_{it}) - N_t^{-1} \sum_{j=1}^{N_t} \text{rank}(z_{jt}) \right), \tag{12}$$

where the scaling factor κ_t makes the portfolio one dollar long and one dollar short (Koijen, Moskowitz, Pedersen, and Vrugt, 2018). The TS-Carry uses the sign of the currency's forward discount as a signal. It goes long (short) currencies with a positive (negative) discount (Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2011, Daniel, Hodrick, and Lu, 2017).

We consider two CS-Momentum strategies, which use the currency's performance as a signal. The first, CS-Mom 1, uses the performance in the most recent month as a signal (Menkhoff, Sarno, Schmeling, and Schrimpf, 2012b, Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2011) and the second, CS-Mom 12, uses the performance in the most recent year skipping the most recent month as a signal (Asness, Moskowitz, and Pedersen, 2013). Specifically, weights are based on the signal ranks according to Equation (12). We also consider two TS-Momentum strategies, which use the sign of the currency's recent performance as a signal. The first, TS-Mom 1, uses the currency's last month performance (Burnside, Eichenbaum, and Rebelo, 2011) and the second, TS-Mom 12, uses the currency's performance in the last twelve months as a signal (Moskowitz, Ooi, and Pedersen, 2012). They both go long (short) currencies with a positive (negative) performance. Similar to Moskowitz, Ooi, and Pedersen, 2012, the second strategy also scales the weights inversely with the conditional return volatilities (same estimates as those used in the construction of the UMVE portfolio, see below).

Lastly, CS-Value uses the Bilson (1984) signal in Equation (7), whereby a relatively low (high) real exchange rate today indicates that the currency is cheap (expensive). Specifically, weights are based on the signal ranks according to Equation (12) (Asness, Moskowitz, and Pedersen, 2013).

The role of currency denomination

Our main analysis pertains to a USD perspective. The question is whether the empirical results are affected by this choice. In fact, there are two separate choices with respect to currency that one may make.

The first choice is about the measuremnt currency. An investor funding herself in Japanese Yen (JPY), and receiving a payoff in JPY, faces a different excess return than a USD-funded investor in the same currency i. The one "forward" dollar strategy described in Equation (5) is equivalent to $1/(F_t^i/S_t^{¥})$ contracts purchased with JPY. Thus, JPY excess return on forward position in currency i is

$$R_{t,t+1}^{ei,¥} = R_{t,t+1}^{ei} \cdot S_t^{¥}/S_{t+1}^{¥}.$$

Unless financial markets are complete, the SDF projected on JPY-denominated excess returns is not equal to the SDF projected on USD-denominated excess returns, and then converted into JPY. As a result, the pricing of risks may be different depending on the measurement currency.

The second choice is about the base currency. All the CS strategies that we discuss are not exposed to the USD by construction. As regards the remaining strategies, one could contemplate a different base, e.g., exposure to the euro, even if the measurement currency continues to be USD. One of the big points in the literature is that the choice of USD as a base currency is intentional because of the USD specialness (e.g., Lustig, Roussanov, and Verdelhan, 2014, Maggiori, Neiman, and Schreger, 2020).

Given the specialness of USD as the base currency, we address only the choice of

the measurement currency. Daniel, Hodrick, and Lu (2017) show theoretically, in continuous time with normally distributed innovations, that the difference between expected excess returns from the two different measurement perspectives is related to the covariance between the depreciation rate between the two measurement currencies and the given strategy return. They argue that this covariance is small. We ask a related but different question whether our conclusions about the pricing of the individual currencies or strategies change when considering another measurement currency.

One approach to address this question is to construct a new SDF projection for each measurement currency and then test it as described above for the USD. We propose an alternative empirical strategy. We simply convert the USD UMVE into a new measurement currency, e.g., JPY, $R^{U,¥}_{t,t+1} = R^{U}_{t,t+1} \cdot S^{¥}_{t} / S^{¥}_{t+1}$, and test if it prices individual currencies, $R^{ei,¥}_{t,t+1}$, and strategies $R^{P,¥}_{t,t+1}$, expressed in JPY.

As noted earlier, if financial markets are complete, this strategy delivers the correct projection of the SDF onto JPY-denominated excess returns. Next, suppose markets are incomplete. That conversion is still useful to evaluate because it corresponds to an actual trading strategy where a JPY-funded investor pursues the USD UMVE. Further, if the JPY-denominated USD SDF prices JPY-denominated currencies and strategies correctly, it can serve as an estimate of the valid JPY SDF for the test assets we consider. We use JPY here as a specific example. In the empirical work we consider all currencies for measurement.

3 Evidence

3.1 Data

We construct a dataset of daily spot and one-month forward exchange rates expressed as USD per unit of the foreign currency for the G10 currencies (AUD, CAD, EUR spliced with DEM, JPY, NZD, NOK, SEK, CHF, GBP, USD) from January 1, 1976 to May 31, 2020. Daniel, Hodrick, and Lu (2017) offer compelling arguments for exclusion of emerging currencies and the European currencies other than the EUR. In particular, emerging currencies reflect credit risk of the respective sovereigns and, therefore, the economic composition of the associated currency risk premiums is different (e.g., Chernov, Creal, and Hördahl, 2020, Na, Schmitt-Grohe, Uribe, and Yue, 2018).

We use data from several providers through Datastream. Our monthly dataset includes the last day of every month from the daily dataset. Forward exchange rates for the AUD and NZD are available from December 1984, which effectively determines that January 1985 is the common starting month for currency excess returns.

We also collect monthly consumer price indices (CPIs) from the OECD for the period January 1976 to May 2020. Only quarterly CPIs are available for the AUD and NZD, so we forward fill the quarterly values to get monthly observations. As CPI is published with a lag and we want to ensure that all the variables are observable at time t, we construct the RER in Equation (6) using the price level lagged by three months.

The returns on the trading strategies are constructed using data going as far back as possible. All strategies use all nine currencies versus the USD from January 1985.

Note that since both the value and momentum signals rely on spot rates, for which we have data going back to 1976, we have ten years of data to estimate the first conditional means (per Equation (8)) and covariance matrix for January 1985 when the currency excess return sample starts. We then each month expand the sample by one month to update these estimates in an out-of-sample fashion.

3.2 Preliminary evidence

Before we proceed with testing the model, we check if the main objects that we use for constructing the UMVE, $E_t(R^e_{t,t+1})$ and $V_t(R^e_{t,t+1})$, are plausible. Specifically, we check if they predict $R^e_{t,t+1}$ and $(R^e_{t,t+1} - E_t(R^e_{t,t+1}))(R^e_{t,t+1} - E_t(R^e_{t,t+1}))^\top$ respectively.

Table 1A demonstrates results of regressing $R^e_{t,t+1}$ on $E_t(R^e_{t,t+1})$ and $(R^e_{t,t+1} - E_t(R^e_{t,t+1}))(R^e_{t,t+1} - E_t(R^e_{t,t+1}))^\top$ on $V_t(R^e_{t,t+1})$ for individual currencies in a panel setting. Slopes in both regressions are significantly different from 0 and insignificantly different from 1. Even point-wise these coefficients are pretty close to 1: 0.84 and 0.93, respectively. Panel B performs the same analysis for strategies with the same conclusion. The conditional expected return and variance of each strategy are computed using the expected return vector and covariance matrix for the underlying individual currencies, along with the conditional portfolio weights of each strategy.

Table 1C documents time variation in $E_t(R^e_{t,t+1})$ and $V_t(R^e_{t,t+1})$ for both individual currencies and strategies. In both cases, the variation in expected excess returns is substantial. For currencies, the volatility of conditional expectations is more than double than the average conditional expectation. For strategies, the two objects are about the same. As another perspective, conditional expectations are as variable as conditional variances.

Lastly, Panel D explores the role of the three signals that we used to construct conditional expectation. In the first row we report the unconditional SR for the full model, which equals 1.075. We emphasize that UMVE portfolio returns correspond to a real-time implementable trading strategy. In the subsequent rows, we consider UMVEs formed on the basis of using the Carry signal alone, or Carry and Momentum, or Carry and Value. All of these combinations result in lower SR, albeit still high.

In order to assess the economic significance of using all the signals, we implement the Barillas and Shanken (2017) test and check if these three alternative UMVEs have an alpha with respect to the UMVE formed using all three signals. All of the alphas are significant. The corresponding adjusted R^2s vary between 50% and 80%. In unreported results, we reverse the direction of the regressions and find no significant alphas. Thus, the UMVE portfolio from the first row can explain UMVE returns from other versions of our model.

3.3 Testing the model's implications

Headline results

Table 2 presents our initial testing results. Panel A displays summary statistics for individual currencies (average excess returns and sample SR). SR for the UMVE reported in Table 1D exceeds the largest currency SR, NZD, by 2.5 times. The panel also shows alphas from individual regressions (8). The largest $t-$statistic is 1.67 for NZD. The largest adjusted R^2 from regressing a return on the UMVE is 2.5% for AUD.

Panel B displays the same information for strategies. The UMVE SR exceeds the largest strategy SR, TS-Carry, by 1.6 times. None of the strategies have a significant alpha relative to the UMVE, with the highest individual t-statistic 1.75 for the Dollar Carry strategy. On average, the alphas are about 70% smaller than the mean excess return on the strategies. Lastly, adjusted R^2s in these regressions are low, which implies that only a small component of the variation in these portfolios' returns is being priced. The Carry strategies and CS-Value exhibit the largest exposures to the UMVE with R^2 around 15%, followed by Momentum around 9%, and the rest ranging between 0 and 6%.

The first row of Table 3, columns labeled "All strategies" display the GRS and MHR tests applied to the nine strategies. The MHR test uses strategy returns at the 1, 3, 6, 12, 24 and 48 month horizons. The $p-$values are 0.47 and 0.28, respectively. These results imply that the candidate UMVE prices single-horizon returns to the strategies both unconditionally and conditionally.

This failure to reject is meaningful. Consider a simpler model in which conditional expectations are computed on the basis of RWH/Carry alone and test it using the first five strategy portfolios, which do not rely on Momentum or Value signals. See the second row of Table 3, columns labeled "Dollar + All Carry". GRS fails to reject with a $p-$value of 0.31, but MHR does reject with a $p-$value of 0.03. In other words, the UMVE based only the conditional mean returns implied by the random walk hypothesis does not price long-run returns to the Dollar and Carry-based strategies.

Table 4 explores AUD and JPY as alternative measurement currencies. Results are qualitatively the same: strategy alphas are insignificant, adjusted R^2 are low, and GRS $p-$values are high. The results for the remaining seven currencies are qualita-

tively similar and available upon request. Thus, despite the potential market incompleteness, we can conclude that our approach and conclusions are not dependent on the choice of the measurement currency.

Unpriced components of returns

The low R^2 of strategy returns in Table 2 suggests a large unpriced component of these returns. Daniel, Mota, Rottke, and Santos (2020), in the context of equity strategies, make a strong case for hedging out the unpriced component to enhance strategy performance. It is easy to construct a conditional hedge in our setting because we have an explicit UMVE portfolio.

In particular, we use the conditional covariance matrix of the underlying currencies to construct the conditional beta of each strategy on the UMVE:

$$\beta_t = \frac{w_t^\top V_t(R^e_{t,t+1}) w_t^U}{(w_t^U)^\top V_t(R^e_{t,t+1}) w_t^U},$$

where w_t is a vector of strategy weights w_{it} as used in Equation (11). A portfolio with systematic exposure, a.k.a. the hedged portfolio, is simply $\beta_t R^U_{t,t+1}$. We refer to the residual as the hedging portfolio return:

$$R^{HP}_{t,t+1} = R^P_{t,t+1} - \beta_t R^U_{t,t+1}.$$

We note that this hedging strategy is implementable in real time.

Figure 1 compares the SR of the original strategies to those of the hedging and hedged components. The SR for the hedging components are close to zero as should be the

case for the unpriced piece if our model is well-specified. The SR for the hedged components are much larger than those of the original strategies as anticipated. For instance, for CS Carry, the Sharpe ratio goes from about 0.5 to slightly higher than 1. The one exception is the Dollar strategy where hedging does not seem to make much of a difference.

Table 3 reports GRS tests applied to the hedged portfolios in columns labeled "GRS hedged". Specifically, we jointly test whether the average hedged returns to the strategies are zero, as implied by the model. The use of the conditional betas means that this is an unconditional test of the models' conditional implications. The approach is similar in spirit to that undertaken in Lewellen and Nagel (2006) but where the pricing factor is the out-of-sample UMVE portfolio. Overall, this conditional test yields results similar to the unconditional GRS test: only our full model is not rejected in any of these tests.

A natural question related to the low R^2 of returns is whether there is some common source of variation in the unpriced components of strategy returns. We implement principal component (PC) analysis of the hedging components of currency returns and those of strategy returns. We find, in both cases, that there is no dominant component as the first PC explains 30% and 20%, respectively.

Sources of portfolio timing

To delineate the role of timing based on the conditional mean vs the conditional covariance, we also report in Table 3 test results for these sub-optimal models in rows three and four. The row labeled "$V-$weighted" corresponds to setting the conditional mean to a constant (a column of ones), while the row labeled "$E-$weighted"

corresponds to setting the conditional covariance matrix to a constant (the identity matrix). The SDF projection is not fully specified in these cases, so we cannot implement MHR tests. The column labeled "UMVE SR" documents a large drop in the SR as compared to the optimal portfolio weights. GRS strongly rejects both models.

3.4 Time series implications

We analyze the factor structure implied by our model by performing PC analysis on the covariance matrix of currency returns. In this section, we consider the unconditional covariance matrix of individual currencies.[3] Figure 2A displays the contributions of each of the nine PCs to the overall variation. The first PC explains 58%, the second 18%, and the third 9% of the variation.

Panel B shows loadings on individual returns on the first three PCs. We see that the first principal component loads similarly on each currency suggesting that it is related to the Dollar strategy. Indeed, the correlation between this first PC and the Dollar factor is 0.998. The second PC goes long the high interest currencies and goes short the low interest currency and is thus related to Carry – their correlation is 0.718. The third PC is harder to associate with an extant strategy, though a possible interpretation is geographical with largely positive loadings for European countries and largely negative for the other. These results are consistent with Verdelhan (2018) who argues that Dollar and Carry capture exchange rate movements.

Panel C explores how much of the UMVE portfolio's variation can be explained by the PCs. We start by repeating the information of Panel A. Rather than reporting

[3]Conditional analysis produces qualitatively similar results, which are available upon request.

the individual contribution of each PC, we now report their cumulative contribution. Next, starting with the first PC and gradually adding them one after another, we compute how much of the variance of UMVE returns can be explained by a given number of PCs.

The answer is not much. It is close to zero if we use the first PC alone. Recall that this PC explains almost 60% of the variation in individual currency returns. If we use all nine PCs, we progress to about 10% of the variation. This is consistent with low R^2 reported in Table 2B. The flipside of this result is that UMVE has a significant alpha with respect to the PCs as evidenced in Figure 2D.

We conclude that time-series variation in currency returns can be summarized by three factors, two of which are close to the Dollar and CS-Carry strategies. However, these factors are weakly related to the UMVE, which prices the cross-section of both currencies and strategies. That suggests that factor timing is an important ingredient of the pricing success. We proceed with fleshing that out in the next section.

3.5 Factor timing

The leading currency strategies, which we study in this paper involve a substantive amount of currency timing. Thus, a natural starting point is to investigate how well these trading strategies can explain UMVE portfolio returns. Figure 3 addresses this issue via PC analysis of strategies. This analysis parallels that in Figure 2.

Panel A shows individual contributions of PCs to the overall variation in strategy returns. Here three PCs explain 65% of the variation. Panel B displays the loadings of these three PCs on strategy, but no clear interpretation of what these "strategies on

strategies" might entail emerges. Panel C shows how much of the variance of UMVE returns can be explain by strategy PCs. The first three PCs explain only about 20% of the variation in UMVE returns. Thus, again the main drivers of the time series variation asset returns are not the main drivers of priced risk, as summarized by the UMVE. The first five PCs can explain about 50% of the variation and this number stays the same after adding the remaining PCs.

The improvement from 10% to 50% is a testament to the importance of currency timing implicit in the strategies. Yet, there is a lot left on the table. UMVE alphas with respect to these PCs continue to be significant as documented in Panel D.

As one manifestation of conditional dynamics not accounted for in existing strategies, we implement optimal timing of each strategy. That is, every month we determine the weight on a given (zero-investment) strategy that theoretically should yield the highest unconditional SR. In particular, we apply Equation (9) to a strategy return as opposed to the CMVE return. We can implement this equation because our model produces real-time conditional expectations and variance for each currency and the portfolio weights for each strategy are known.

Table 5 compares performance of the strategy returns to the returns of the corresponding timing strategy. The first column reports an alpha from regressing the timed strategy on the original strategy. Panel A reports the case of timing on both conditional expectations and variances. Panel B reports timing on conditional variance alone to highlight the role of conditional expectations.

Panel A of Table 2 shows that six timing alphas are significant and economically meaningful relative to the original strategy returns. The annualized information ra-

tios (IR) are comparable or larger than strategy SR, also reported in Table 2A.[4] Adjusted R^2s are small, in general. All these statistics are highlighting the importance of optimal conditioning that was foreshadowed in Figure 2. In Panel B, the alphas and IRs are much smaller and R^2s are higher underscoring the importance of variation in conditional expected returns for currency dynamics.

3.6 Properties of the projected SDF

The estimated UMVE corresponds to the linear projection of the SDF on the set of currency excess returns. Figure 4 plots time series of both the conditional annualized price of risk (the conditional annualized volatility of the SDF) and returns to the UMVE portfolio, which is perfectly negatively correlated with this SDF. The price of risk various strongly over the sample, from a high of about 4 to a low of about 0.5.

We see that the price of risk experiences its maximum value in a dramatic move during the 1992 Exchange Rate Mechanism crisis. Outside of this event, the main take-away is a substantial decline through the sample, mainly caused by a trend decrease in interest differentials across countries. In particular, the same decline is shared by the price of risk of the model that only uses the RWH when estimating currency expected returns (red dashed line). Notably, recessions, showed as yellow bars on the plot, are associated with a *decline* in the price of risk.

The bottom panel of Figure 4 shows the time series of returns to the UMVE portfolio. First, the returns are relatively smooth, with little evidence of a disaster event. For

[4]We define the Information ratio here as $\sqrt{12}\alpha/\sigma(\eta)$, where η is the residual in the regression. This definition is sometimes also referred to as the Appraisal Ratio.

instance, the Global Financial Crisis is associated with small negative returns and one big *positive* spike. Second, we can see a slow decrease in the volatility in returns over the sample, consistent with the pattern in the price of risk.

As a next step, we make an attempt to relate these two objects to observable variables that have the potential to capture currency risk premiums. As the rest of the literature we focus on macro and financial variables. Among macro variables, we have a particular interest in consumption as it has a rich history and connection to equilibrium models. Among financial variables, we focus on the Fama and French (2015) model as the de-facto SDF projection on on the space of equity returns. Finally, we consider measures of intermediary leverage. The full list of candidate covariates and motivation for considering them are offered in Appendix A.2.

Table 6 reports the results. Panel A displays univariate regressions for the significant covariates. We see that quarterly consumption growth and 3-year consumption growth, a proxy for the long-run consumption component, are significant. This result is consistent with Lustig and Verdelhan (2007) and Zviadadze (2017), who argue that these variables are important for explaining the CS-Carry premium. Our results generalize theirs because we find significant relation to the UMVE, and not just Carry returns. The R^2s of the regressions using consumption growth on the right hand side are only 3% and 7% indicating that other elements than consumption risk may be important for understanding the currency market conditional risk-return trade-off.

Panel A also shows that only two out of the five Fama-French factors are significant, value (HML) and investment (CMA). Even though they are significant, the relation to the UMVE returns is weak as betas are very small. R^2s are small as well, at 0.6% and 1.3%, respectively.

Panel B focuses on the properties of the CS-Carry, a strategy that has received a lot of attention both in the empirical and equilibrium literatures. The middle section of the Panel, labeled 'CS-Carry', demonstrates that consistent with the literature returns to this strategy are significantly related to intermediary leverage (He, Kelly, and Manela, 2017), realized variance of equity (Lustig, Roussanov, and Verdelhan, 2011), and that of exchange rates (Menkhoff, Sarno, Schmeling, and Schrimpf, 2012b). However, as foreshadowed by Panel A, UMVE returns are not related to these variables. Indeed, Panel B makes it explicit in the left panel labeled 'UMVE'. Results in Table 2B suggests reconciliation of that evidence. Only 15% of CS-Carry is explained by the UMVE. Thus, it must be the idiosyncratic component of the strategy that is related to these variables. Indeed, hedging out the systematic component just like we did in Figure 1, we find that the idiosyncratic component of CS-Carry is related to the three variables (see the right section of Table 6B).

Lastly, Panel C relates prices of risk to various measures of aggregate uncertainty: conditional consumption variance, realized equity variance, and realized FX variance. The conditional variance of consumption growth is estimated using a GARCH(1,1) model. Due to the downward trend in the price of risk in the sample, we run these regressions in 3-year differences. The conditional price of risk, that is, conditional volatility of the SDF, is *negatively* related to consumption variance and FX return variance. Thus, periods of high uncertainty are associated with lower Sharpe ratios in the currency market. The sign is negative also for equity variance, though it is insignificant in this case. Figure 5 plots the price of risk vs. the conditional variance of consumption growth. The common trend and the negatively correlated cyclical relation are both clearly visible.

4 Conclusion

We construct a conditional linear projection of the SDF on the space of excess currency returns. That is feasible as most of the literature focuses on the G10 currencies and the corresponding nine exchange rates versus the U.S. dollar. A small cross-section makes estimation of the conditional covariance matrix relatively precise. We then use standard rolling panel regressions and a limited set of longstanding currency characteristics to obtain conditional expected returns. These inputs are all that is needed to construct the projection.

The benefit of this approach is that it is a theoretically motivated empirical characterization of the risk-return trade-off. The use of the conditional covariance matrix means we also characterize the factor structure in realized returns, which allow us to discuss the priced versus unpriced sources of common variation in currency returns. Our analysis leads to a number of new insights.

First, the UMVE portfolio has a sample Sharpe ratio in excess of one and is not explained by any existing factors. At the same time, the UMVE portfolio accounts for known currency factors risk premiums and is not rejected as the pricing factor in standard model tests. Thus, our proposed estimation methodology is validated in the data and uncovers heretofore undiscovered sources of risk.

Second, we show that most of the common variation in currency returns is due to unpriced risks – that is, factors that do not command a risk premium. We show that existing currency strategies indeed have large uncompensated components that can be hedged away to obtain much higher factor Sharpe ratios.

Third, we document large gains to timing many of the existing factors, owing to

the fact that the conditional dynamics of these factors are strongly time-varying. In fact, the standard deviation of conditional risk premiums is as large as the level of unconditional risk premiums.

Fourth, the priced currency risks are effectively unrelated to priced risks in the equity market, as well as shocks to equity and currency market return variance and intermediary capital. Consumption growth, however, is related to priced risks, though the relation is not economically very strong. That said, our results support a role for consumption risk in understanding currency risk and return.

Fifth, the conditional price of risk is strongly downward trending through the sample, with a cyclical component that is *negatively* related with measures of the conditional variance of macroeconomic aggregates and market returns.

References

Ackermann, Fabian, Walt Pohl, and Karl Schmedders, 2017, Optimal and naive diversification in currency markets, *Management Science* 63, 3147–3529.

Adrian, Tobias, Erkko Etula, and Jan J.J. Groen, 2011, Financial amplification of foreign exchange risk premia, *European Economic Review* 55, 354–370.

Adrian, Tobias, Erkko Etula, and Tyler Muir, 2014, Financial intermediaries and the cross-section of asset returns, *Journal of Finance* 69, 2557–2596.

Aloosh, Arash, and Geert Bekaert, 2019, Currency factors, working paper.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen, 2013, Value and momentum everywhere, *Journal of Finance* 68, 929–985.

Bansal, Ravi, and Ivan Shaliastovich, 2013, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies* 26, 1–33.

Barillas, Fancisco, and Jay Shanken, 2017, Which alpha?, *Review of Financial Studies* 30, 1316–1338.

Baz, Jamil, Francis Breedon, Vasant Naik, and Joel Peress, 2001, Optimal portfolios of foreign currencies, *Journal of Portfolio Management* 28, 102–111.

Bilson, John, 1984, Purchasing power parity as a trading strategy, *Journal of Finance* 39, 715–724.

Burnside, Craig, 2012, Carry trades and risk, in Jessica James, Ian Marsh, and Lucio Sarno, ed.: *Handbook of Exchange Rates* (John Wiley and Sons, Ltd: Hoboken NJ).

——— , Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2011, Do peso problems explain the returns to the carry trade?, *Review of Financial Studies* 24, 853–891.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 2011, Carry trade and momentum in currency markets, *Annual Review of Financial Economics* 3, 511–535.

Cassel, Gustav, 1918, Abnormal deviations in international exchanges, *Economic Journal* 28, 413–415.

Chernov, Mikhail, and Drew Creal, 2020, The PPP view of multihorizon currency risk premiums, Review of Financial Studies, forthcoming.

——— , and Peter Hördahl, 2020, Sovereign credit and exchange rate risks: Evidence from Asia-Pacific local currency bonds, working paper.

Chernov, Mikhail, Lars Lochstoer, and Stig Lundeby, 2018, Conditional dynamics and the multi-horizon risk-return trade-off, working paper.

Cochrane, John, 2004, *Asset Pricing* (Princeton University Press: Princeton, NJ) revised edn.

——— , and Lars Hansen, 1992, Asset pricing explorations for macroeconomics, in Oliver Blanchard, and Stanley Fischer, ed.: *NBER Macroeconomics Annual* (MIT Press: Cambridge MA).

Colacito, Riccardo, and Mariano M. Croce, 2011, Risks for the long run and the real exchange rate, *Journal of Political Economy* 119, 153–181.

Colacito, Ric, Mariano M. Croce, Federico Gavazzoni, and Robert Ready, 2018, Currency risk factors in a recursive multicountry economy, *Journal of Finance* 73, 2719–2756.

Cornell, W.B., and J. K. Dietrich, 1978, The efficiency of the market for foreign exchange under floating exchange rates, *Review of Economics and Statistics* 60, 111–120.

Cumby, Robert E., 1988, Is it risk? Explaining deviations from uncovered interest parity, *Journal of Monetary Economics* 22, 279–299.

Dahlquist, Magnus, and Julien Penasse, 2016, The missing risk premium in exchange rates, working paper, SSE.

Daniel, Kent, Robert J. Hodrick, and Zhongjin Lu, 2017, The carry trade: Risks and drawdowns, *Critical Finance Review* 6, 211–262.

Daniel, Kent, Lira Mota, Simon Rottke, and Tano Santos, 2020, The cross-section of risk and return, *Review of Financial Studies* 33, 1927–1979.

Daniel, Kent, and Sheridan Titman, 2012, Testing factor-model explanations of market anomalies, *Critical Finance Review* 1, 103–139.

Della Corte, Pasquale, Lucio Sarno, and Ilias Tsiakas, 2009, An economic evaluation of empirical exchange rate models, *Review of Financial Studies* 22, 3491–3530.

Eichenbaum, Martin, Benjamin K. Johannsen, and Sergio Rebelo, 2020, Monetary policy and the predictability of nominal exchange rates, Review of Economic Studies, forthcoming.

Fama, Eugene, and Kenneth French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.

Farhi, Emmanuel, and Xavier Gabaix, 2016, Rare disasters and exchange rates, *Quarterly Journal of Economics* 131, 1–52.

Ghosh, Anisha, Christian Julliard, and Alex P. Taylor, 2019, An information-theoretic asset pricing model, working paper.

Gibbons, Michael, Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 59, 1121–1152.

Greenaway-McGrevy, Ryan, Nelson C. Mark, Donggyu Sul, and Juh-Lin Wu, 2018, Identifying exchange rate common factors, *International Economic Review* 59, 2193–2218.

Hansen, Lars, and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.

Hansen, Lars, and Scott Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.

Hassan, Tarek K., and Rui C. Mano, 2019, Forward and spot exchange rates in a multi-currency world, *Quarterly Journal of Economics* 134, 397–450.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

Jorda, Oscar, and Alan M. Taylor, 2012, The carry trade and fundamentals: Nothing to fear but FEER itself, *Journal of International Economics* 88, 74–90.

Keynes, John Maynard, 1923, *A Tract on Monetary Reform* (MacMillan and Co., Ltd., London).

Kho, Boing-Chan, 1996, Time varying risk premia, volatility, and technical reading profits: Evidence from foreign currency futures markets, *Journal of Financial Economics* 41, 249–290.

Koijen, Ralph S.J., Tobias J. Moskowitz, Lasse Heje Pedersen, and Evert B. Vrugt, 2018, Carry, *Journal of Financial Economics* 127, 197–225.

Korsaye, Sofonias Alemu, Fabio Trojani, and Andrea Vedolin, 2020, The global factor structure of exchange rates, working paper.

Ledoit, Olivier, and Michael Wolf, 2020, Analytical nonlinear shrinkage of large-dimensional covariance matrices, *Annals of Statistics* 48, 3043–3065.

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

Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional CAPM does not explain asset pricing anomalies, *Journal of Financial Economics* 79, 289–314.

——— , and Jay Shanken, 2010, A skeptical appraisal of asset-pricing tests, *Journal of Financial Economics* 96, 175–194.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2014, Countercyclical currency risk premia, *Journal of Financial Economics* 111, 527–553.

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

Lustig, Hanno N., Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

Maggiori, Matteo, Brent Neiman, and Jesse Schreger, 2020, International currencies and capital allocation, *Journal of Political Economy* 126, 2019–2066.

Mauer, Thomas, Thuy-Duong To, and Ngoc-Khanh Tran, 2018, Pricing implications of covariances and spreads in currency markets, working paper.

——— , 2020, Market timing and predictability in FX markets, working paper.

Meese, Richard A., and Kenneth Rogoff, 1983, Empirical exchange rate models of the seventies: Do they fit out of sample?, *Journal of International Economics* 14, 3–24.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012a, Carry trades and global foreign exchange volatility, *Journal of Finance* 67, 681–718.

——— , 2012b, Currency momentum strategies, *Journal of Financial Economics* 106, 660–684.

Moskowitz, Tobias J., Yao Hua Ooi, and Lasse Heje Pedersen, 2012, Time series momentum, *Journal of Financial Economics* 104, 228–250.

Na, Seunghoon, Stephanie Schmitt-Grohe, Martin Uribe, and Vivian Yue, 2018, The twin Ds: Optimal default and devaluation, *American Economic Review* 108, 1773–1819.

Okunev, John, and Derek White, 2003, Do momentum-based strategies still work in foreign currency markets?, *Journal of Financial and Quantitative Analysis* 38, 425–447.

Orlowski, Piotr, Valeri Sokolowski, and Erik Sverdrup, 2019, Benchmark currency stochastic discount factors, working paper.

Porter, Michael, 1971, A theoretical and empirical framework for analyzing the term structure of exchange rate expectations, *Staff Papers (International Monetary Fund)* 18, 613–645.

Sandulescu, Mirela, Fabio Trojani, and Andrea Vedolin, 2020, Model-free international stochastic discount factors, Journal of Finance, forthcoming.

Sarkissian, Sergei, 2003, Incomplete consumption risk sharing and currency risk premiums, *Review of Financial Studies* 16, 983–1005.

Stutzer, Michael, 1996, A simple nonparametric approach to derivative security valuation., *Journal of Finance* 60, 1633–1652.

Sweeney, Richard, 1986, Beating the foreign exchange market, *Journal of Finance* 41, 163–182.

Verdelhan, Adrien, 2010, A habit-based explanation of the exchange rate risk premium., *Journal of Finance* 65, 123–146.

——— , 2018, The share of systematic variation in bilateral exchange rates, *Journal of Finance* 73, 375–418.

Zviadadze, Irina, 2017, Term-structure of consumption risk premia in the cross-section of currency returns, *Journal of Finance* 72, 1529–1566.

Appendix A Appendix

Appendix A.1 Estimating covariance matrix

First, for each month t, we compute the realized covariance matrix based on daily depreciation rates. That is, define $X_{j,t}$ as the vector of percentage changes in spot currency rates over day j in month t. The realized covariance matrix is then

$$\hat{\Sigma}_t = \sum_{j=1}^{D_t} X_{j,t} X_{j,t}^\top$$

where D_t is the number of trading days in the month. Given a potentially large cross-section and short time-series, we apply the shrinkage method of Ledoit and Wolf (2020) to $\hat{\Sigma}_t$ to improve the mean-squared-error of this sample covariance estimate. Denote the resulting shrunk matrix as $\tilde{\Sigma}_t$.

Second, we apply a simple exponentially weighted average to $\tilde{\Sigma}_t$ to arrive at the estimate of the conditional covariance matrix for month $t+1$ currency percentage price changes:

$$V_t = (1-\lambda)\tilde{\Sigma}_t + \lambda V_{t-1},$$

where we set $\lambda = 0.94$ following, e.g., the RiskMetrics model. Finally, to get to the conditional covariance matrix of excess currency returns, we pre- and post-multiply this matrix by a diagonal matrix with $(i,i)^{th}$ element set to S_t^i/F_t^i.

Appendix A.2 Additional data sources

This appendix lists and motivates the macroeconomic and financial variables used as covariates in Section 3.6 to characterize the properties of the projected SDF. We consider several groups of covariates.

The first group of covariates are based on a per capita series of real consumption expenditure on nondurables and services that we construct from the NIPA Table 7.1. at the Bureau of Economic Analysis. The data and more detailed information about the series are available through the webpage:

https://apps.bea.gov/iTable/index_nipa.cfm.
We construct quarterly log consumption growth, three-year log consumption growth, and consumption growth volatility on a quarterly basis from 1985:Q1 to 2020:Q1. These consumption growth measures are appealing as they naturally relate to equilibrium models of asset prices. Several studies have attempted to explain currency returns or trading strategies to short-term consumption growth (Cumby, 1988; Sarkissian, 2003; Lustig and Verdelhan, 2007), long-term consumption growth (Colacito and Croce, 2011), and consumption growth volatility (Lustig, Roussanov, and Verdelhan, 2014).

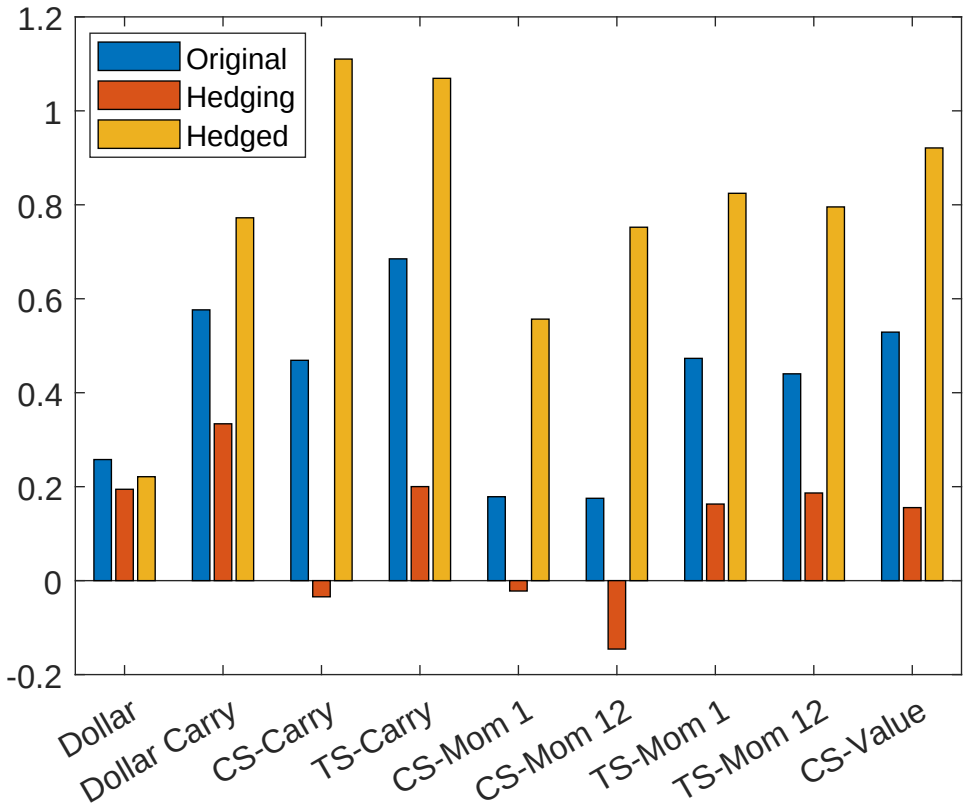
The second group of covariates are the equity factors in the Fama and French (2015) five-factor model. The data are available through the webpage:
http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/.
More specifically, the equity factors are: the return on a world market portfolio minus the U.S. one-month T-bill rate (market, MKT); the return on a small stock portfolio minus the return on a big stock portfolio (small minus big, SMB); the return on a value portfolio minus the return on a growth portfolio (high minus low, HML); the return on a robust-operating-profitability portfolio minus the return on a weak-operating-profitability portfolio (robust minus weak, RMW); and the return on a conservative investment portfolio minus the return on an aggressive investment portfolio (conservative minus aggressive, CMA). We use monthly observations from January 1985 to May 2020. Examples of studies relating carry trades to equity factors include Lettau, Maggiori, and Weber (2014) and Daniel, Hodrick, and Lu (2017).

The third group of covariates are based on the literature of financial intermediaries and their leverage (Adrian, Etula, and Groen, 2011; Adrian, Etula, and Muir, 2014; He, Kelly, and Manela, 2017). More specifically, we consider the monthly intermediary capital measures of He, Kelly, and Manela (2017) from January 1985 to May 2020. The data are available through the webpage:
http://apps.olin.wustl.edu/faculty/manela/data.html.
Adrian, Etula, and Groen, 2011 decompose currency risk premiums into components associated with macroeconomic and balance sheet variables. He, Kelly, and Manela (2017) explain currency trading strategies based on cross-sectional carry and momentum.

The fourth group of covariates are based on measures of global equity and foreign exchange rate volatility. Lustig, Roussanov, and Verdelhan (2011) and Menkhoff, Sarno, Schmeling, and Schrimpf (2012a) find that cross-sectional carry strategies

relate to global equity and foreign exchange rate volatility, respectively. We construct monthly realized variances from daily changes in equity prices and spot exchange rates similar to these authors. We consider the period January 1985 to May 2020.

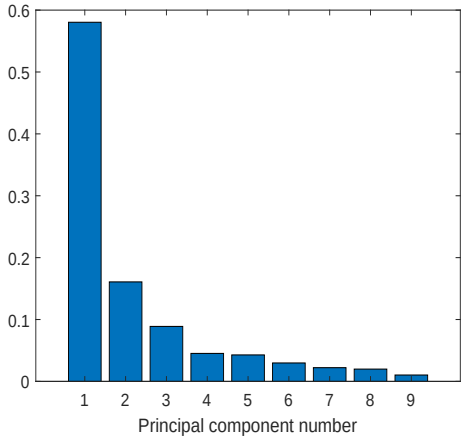
Figure 1
Sharpe ratio of original, hedging, and hedged strategies



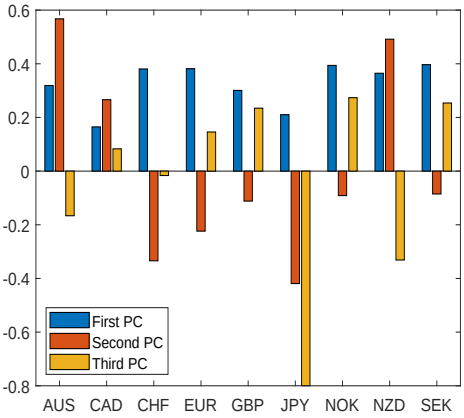
The blue bars show the sample annualized Sharpe ratios of each strategy. The red bars shows the same for each strategy's hedging portfolio, defined as the component of the strategy returns that is unpriced according to the model. The yellow bars show the Sharpe ratio of the factor returns when unpriced risks are hedged out. All portfolios are tradeable in real time. The sample is monthly from January 1985 to May 2020.

Figure 2
Principal component analysis of currencies

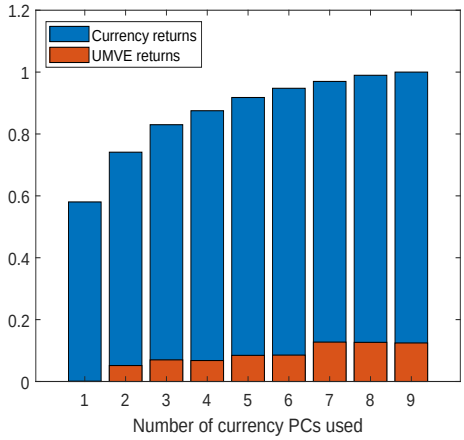
(A) Fraction of V explained



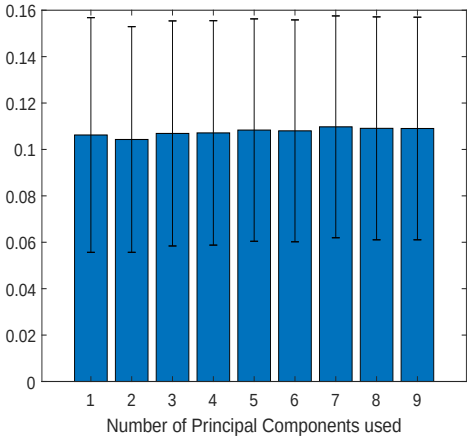
(B) Loadings of first three PCs



(C) Cumulative fraction of V explained



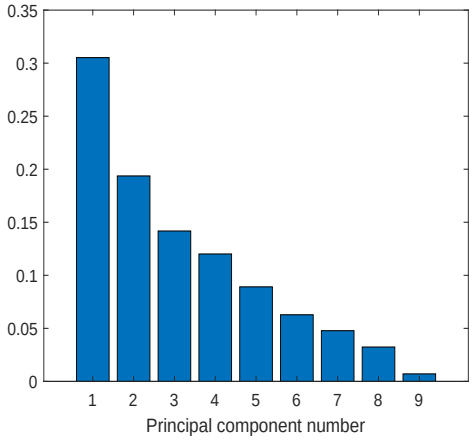
(D) UMVE α



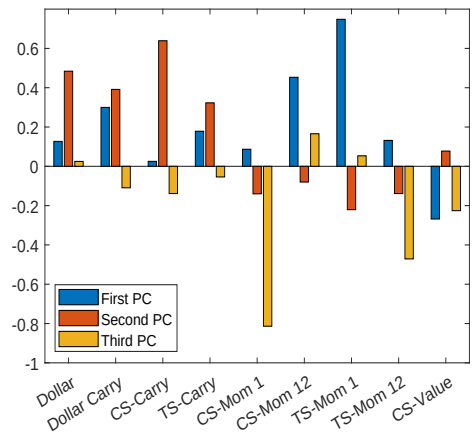
Panel A shows the fraction of variance across the nine individual currencies that each principal component (PC) explains. The PCs are obtained from the unconditional covariance matrix of individual currency returns. Panel B shows the loadings of the three first PCs on each currency. Panel C shows in blue the cumulative amount of currency variance explained as one goes from using one to all nine PCs. In red is the R^2 of a regression of the UMVE returns on an increasing number of PCs. Panel D shows the "alpha" of a regression of the UMVE returns on an increasing number of PCs. The error bars correspond to the 95% confidence interval computed using heteroskedasticity-adjusted standard errors. The sample is monthly from January 1985 to May 2020.

Figure 3
Principal component analysis of strategies

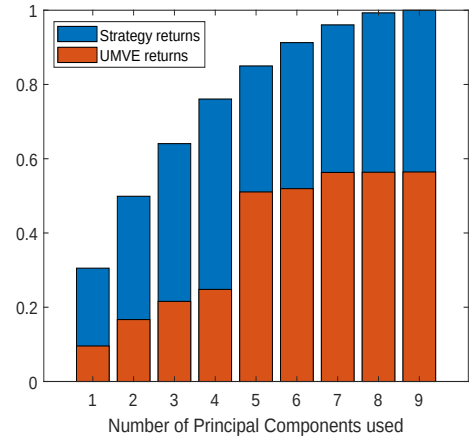
(A) Fraction of V explained



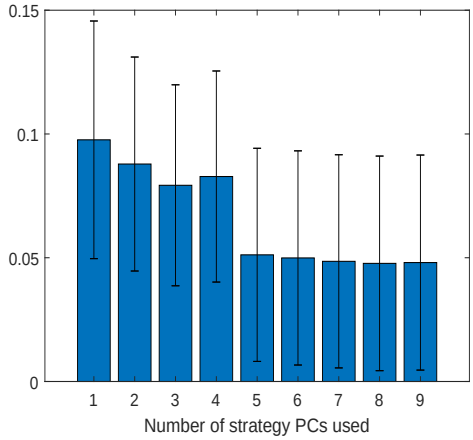
(B) Loadings of first three PCs



(C) Cumulative fraction of V explained

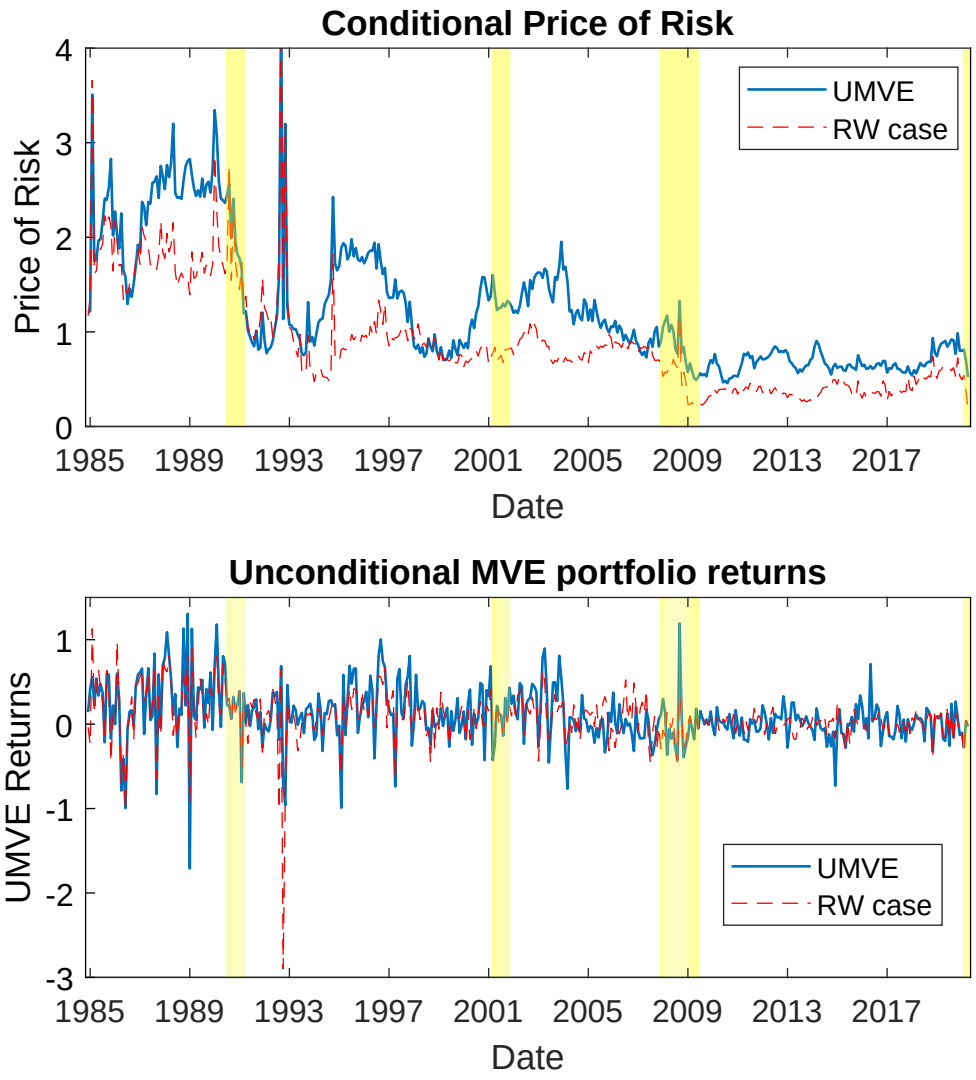


(D) UMVE α



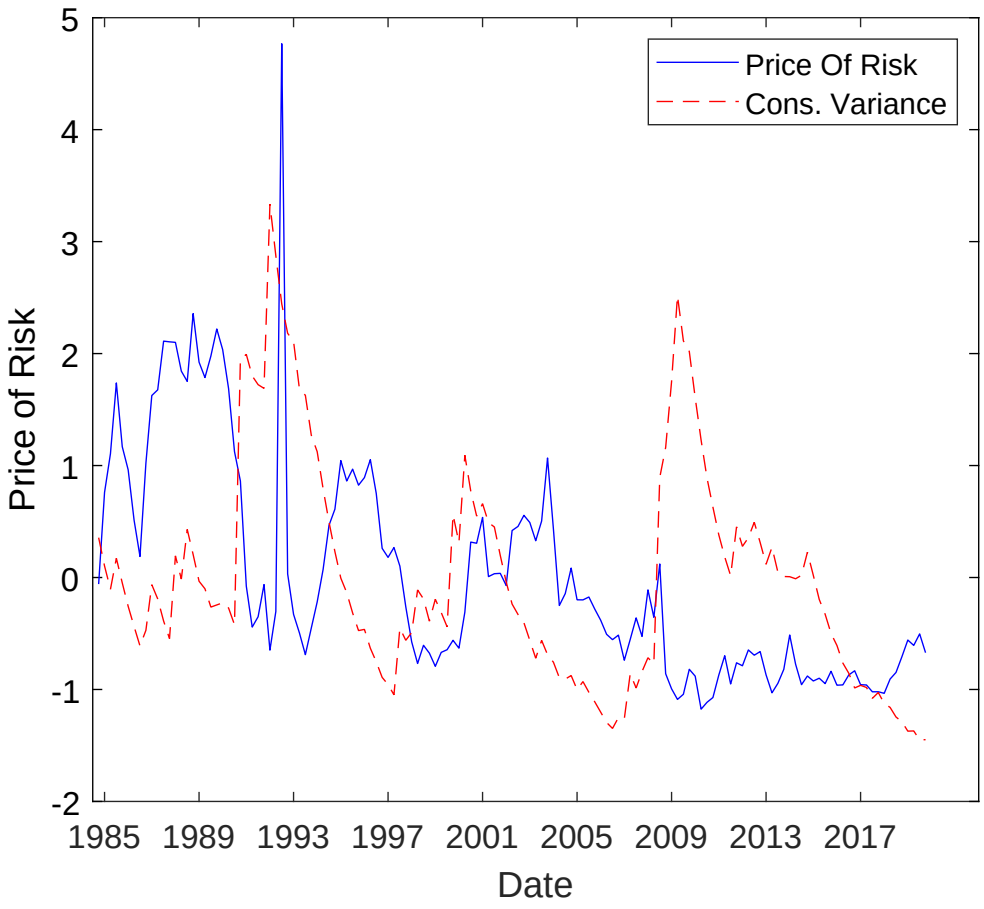
Panel A shows the fraction of variance across the nine currency trading strategies that each principal component (PC) explains. The PCs are obtained from the unconditional covariance matrix of these currency strategies. Panel B shows the loadings of the three first PCs on each strategy. Panel C shows in blue the cumulative amount of strategy variance explained as one goes from using one to all nine PCs. In red is the R^2 of a regression of the UMVE returns on an increasing number of these PCs. Panel D shows the "alpha" of a regression of the UMVE returns on an increasing number of PCs. The error bars correspond to the 95% confidence interval computed using heteroskedasticity-adjusted standard errors. The sample is monthly from January 1985 to May 2020.

Figure 4
Properties of UMVE



The top plot shows the annualized conditional price of risk as implied by the stochastic discount factor derived in the paper. The lower plot shows the realized returns to the UMVE portfolio implied from this SDF. The blue solid lines corresponds to the full model, whereas the dashed red lines shows the case where only the random walk assumption is used for estimating currency risk premiums. The sample is monthly from January 1985 to May 2020.

Figure 5
Price of risk and consumption's variance



The blue solid line is the conditional price of risk as implied by the stochastic discount factor derived in the paper. The dashed red line shows the conditional variance of quarterly, real per-capita nondurable+services log consumption growth, as estimated by a GARCH(1,1). Both series are standardized to be mean zero and unit variance to facilitate comparison. The sample is quarterly from end of December 1984 to end of April 2020.

Table 1: Predictive ability of conditional expectations and variance

Panel A Currency returns

$R^e_{t,t+1}$		$(R^e_{t,t+1} - E_t R^e_{t,t+1})(R^e_{t,t+1} - E_t(R^e_{t,t+1}))^\top$	
$E_t(R^e_{t,t+1})$	0.838	$V_t(R^e_{t,t+1})$	0.931
s.e.	0.253	s.e.	0.068
R^2_{adj}	0.009	R^2_{adj}	0.053

Panel B Strategy returns

$R^e_{t,t+1}$		$(R^e_{t,t+1} - E_t R^e_{t,t+1})(R^e_{t,t+1} - E_t(R^e_{t,t+1}))^\top$	
$E_t(R^e_{t,t+1})$	0.791	$V_t(R^e_{t,t+1})$	1.055
s.e.	0.237	s.e.	0.087
R^2_{adj}	0.003	R^2_{adj}	0.072

Panel C Summary statistics of conditional means and variances

	$E[E_t(R^e_{t,t+1})]$	$E[V_t^{1/2}(R^e_{t,t+1})]$	$V^{1/2}[E_t(R^e_{t,t+1})]$	$V^{1/2}[V_t^{1/2}(R^e_{t,t+1})]$
Currencies	0.021	0.107	0.044	0.048
Strategies	0.032	0.079	0.034	0.036

Panel D UMVE portfolios

Signals used	SR	alpha	$t-$stat	R^2_{adj}
Carry+Mom+Val	1.075	n/a	n/a	n/a
Carry	0.871	0.0440	2.5369	0.515
Carry+Mom	0.871	0.0400	2.6327	0.586
Carry+Val	0.998	0.0181	2.1292	0.796

Panel A shows panel regressions of currency excess returns $R^e_{t,t+1}$ on the conditional expected returns as estimated from our model, $E_t R^e_{t,t+1}$, and $(R^e_{t,t+1} - E_t R^e_{t,t+1})^2$ on the conditional variance as estimated by our model, $V_t(R^e_{t,t+1})$. Panel B shows the same for the strategy returns. Panel C shows summary statistics for estimated conditional means and variances. $E[E_t(R^e_{t,t+1})]$ denotes the grand average risk premium, $E[V_t^{1/2}(R^e_{t,t+1})]$ is the grand average conditional standard deviation, $V^{1/2}[E_t(R^e_{t,t+1})]$ is the standard deviation of conditional expected returns, and $V^{1/2}[V_t^{1/2}(R^e_{t,t+1})]$ is the standard deviation of conditional standard deviations. Panel D shows the Sharpe ratio (SR) if the UMVE portfolio as a function of the signals used in its construction, along with the "alpha" of alternative UMVE portfolios regressed on the full model UMVE portfolio.

Table 2: Testing the UMVE

Panel A Currency returns

Currency	SR	$E[R^e]$	$t-$stat	α	$t-$stat	R^2_{adj}
AUD	0.241	2.83	1.44	0.73	0.34	0.025
CAD	0.099	0.74	0.59	-0.03	-0.02	0.007
CHF	0.149	1.68	0.89	3.22	1.56	0.014
EUR	0.145	1.52	0.86	1.83	0.95	-0.002
GBP	0.217	2.17	1.29	1.08	0.65	0.008
JPY	0.054	0.60	0.32	2.45	1.20	0.022
NOK	0.205	2.26	1.22	1.27	0.63	0.005
NZD	0.427	5.29	2.54	3.72	1.67	0.012
SEK	0.134	1.48	0.80	0.28	0.14	0.008

Panel B Strategy returns

Strategy	SR	$E[R^e]$	$t-$stat	α	$t-$stat	R^2_{adj}
Dollar	0.258	2.06	1.54	1.61	1.14	0.000
Dollar Carry	0.576	4.56	3.43	2.40	1.75	0.062
CS-Carry	0.469	4.14	2.79	0.43	0.28	0.151
TS-Carry	0.685	3.44	4.08	1.25	1.47	0.163
CS-Mom 1	0.179	1.46	1.06	0.61	0.36	0.007
CS-Mom 12	0.175	1.47	1.04	-1.22	-0.82	0.087
TS-Mom 1	0.440	2.73	2.62	1.45	1.30	0.034
TS-Mom 12	0.473	5.19	2.82	1.65	0.86	0.088
CS-Value	0.529	4.02	3.15	0.98	0.69	0.136

Panel A gives the annualized Sharpe ratio (SR), average return, and $t-$statistic of excess returns to each currency, along with its "alpha" and R^2 with respect to the UMVE portfolio. The $t-$statistics are heteroskedasticity-adjusted. Panel B shows the same for the strategy returns. The sample is monthly from January 1985 to May 2020.

Table 3: Asset pricing tests

	UMVE	All strategies			Dollar + All Carry		
Model	SR	GRS	MHR	GRS hedged	GRS	MHR	GRS hedged
Optimal	1.075	0.470	0.277	0.475	0.467	0.568	0.434
Carry only	0.871	0.010	0.001	0.010	0.309	0.030	0.156
$V-$weighted	0.165	0.000	–	0.001	0.002	–	0.023
$E-$weighted	0.565	0.000	–	0.000	0.084	–	0.002

The GRS $p-$value is for the joint test of all "alpha" equal to zero. The MHR $p-$value is for the conditional test of the model as implied by multi-horizon returns. The horizons are 1, 3, 6, 12, 24, and 48 months. See the text for details. The "GRS hedged" $p-$value is for the joint test of all average hedged strategy returns equal to zero. For each strategy, we compute the conditional beta on the UMVE portfolio out-of-sample and thus hedge out the priced component, according to our model, in real time. This test is also an unconditional test of each model's conditional implications. The sample is monthly from January 1985 to May 2020.

Table 4: Testing the UMVE: Alternative measurement currencies

Panel A Returns in AUD

Strategy	SR	$E[R^e]$	$t-$stat	α	$t-$stat	R^2_{adj}
Dollar	0.177	1.42	1.05	1.11	0.78	-0.001
Dollar Carry	0.550	4.39	3.27	2.22	1.61	0.065
CS-Carry	0.397	3.58	2.36	-0.02	-0.02	0.145
TS-Carry	0.644	3.26	3.84	1.12	1.31	0.162
CS-Mom 1	0.185	1.55	1.10	0.58	0.33	0.010
CS-Mom 12	0.174	1.47	1.04	-1.23	-0.80	0.091
TS-Mom 1	0.475	5.28	2.83	1.70	0.88	0.093
TS-Mom 12	0.445	2.81	2.65	1.49	1.31	0.038
CS-Value	0.516	3.99	3.07	0.91	0.63	0.144
UMVE SR	1.042	GRS	$p-$value	0.565		

Panel B Returns in JPY

Strategy	SR	$E[R^e]$	$t-$stat	α	$t-$stat	R^2_{adj}
Dollar	0.204	1.62	1.22	1.17	0.83	0.000
Dollar-carry	0.556	4.37	3.31	2.26	1.65	0.058
CS-Carry	0.523	4.56	3.12	0.85	0.57	0.150
TS-Carry	0.703	3.52	4.19	1.33	1.56	0.158
CS-Mom 1	0.186	1.51	1.11	0.65	0.40	0.007
CS-Mom 12	0.179	1.49	1.06	-1.25	-0.85	0.088
TS-Mom 1	0.460	5.00	2.74	1.38	0.73	0.091
TS-Mom 12	0.433	2.64	2.58	1.42	1.28	0.031
CS-Value	0.547	4.11	3.26	1.10	0.79	0.132
UMVE SR	1.092	GRS	$p-$value	0.455		

The table gives the annualized Sharpe ratio (SR), average return, and $t-$statistic of excess returns to each currency strategy, along with its "alpha" and R^2 with respect to the UMVE portfolio. The $t-$statistics are heteroskedasticity-adjusted. Results for other measurement currency are similar and available upon reuqest. The sample is monthly from January 1985 to May 2020.

Table 5: Factor timing

Panel A Optimal timing

Strategy	α timing	$t-$stat	IR timing	R^2_{adj}
Dollar	3.75	2.72	0.482	0.055
Dollar-Carry	1.49	1.54	0.234	0.348
CS-Carry	0.48	0.46	0.083	0.568
TS-Carry	1.40	1.95	0.384	0.472
CS-Mom 1	2.40	2.14	0.304	0.070
CS Mom 12	2.69	2.42	0.384	0.296
TS-Mom 1	2.66	2.92	0.471	0.164
TS-Mom 12	2.18	2.06	0.319	0.609
CS-Value	-0.53	-0.62	-0.097	0.478

Panel B Volatility timing

Strategy	α timing	$t-$stat	IR timing	R^2_{adj}
Dollar	-1.31	-2.13	-0.446	0.861
Dollar-Carry	0.44	0.92	0.146	0.855
CS-Carry	-0.13	-0.20	-0.039	0.815
TS-Carry	1.14	1.29	0.257	0.687
CS-Mom 1	0.00	0.00	0.000	0.801
CS Mom 12	0.97	1.81	0.295	0.828
TS-Mom 1	1.13	1.20	0.219	0.571
TS-Mom 12	1.05	1.32	0.244	0.703
CS-Value	-0.76	-1.93	-0.291	0.890

Panel A gives the annualized "alpha" and $t-$statistic, Information Ratio (IR), and adjusted R^2 for all timed strategy returns regressed on the original strategy returns. A positive and significant alpha implies that the timing provides a significant Sharpe ratio that cannot be explained by the original strategy. The timing weights are calculated in an out-of-sample fashion using the estimated conditional means and covariance matrix of the underlying currencies. Panel B shows the same when only the conditional strategy variance is used in the timing strategy (volatility timing). The $t-$statistics are heteroskedasticity-adjusted. The sample is monthly from January 1985 to May 2020.

Table 6: Covariates

Panel A UMVE returns vs. observable factors

	Quarterly Δc	3-year Δc	HML	CMA
Beta	32.452	6.740	0.011	0.021
$t-$stat	2.516	3.167	2.039	2.724
R^2_{adj}	0.031	0.073	0.006	0.013
N	141	131	425	425

Panel B UMVE, CS-Carry, its idiosyncratic component and covariates

	UMVE			CS-Carry			CS-Carry idiosyncratic		
	HKM	Equity RV	FX RV	HKM	Equity RV	FX RV	HKM	Equity RV	FX RV
Beta	0.309	-3.633	-28.499	0.111	-2.066	-6.023	0.099	-1.814	-4.858
(t-stat)	1.066	-0.565	-1.494	3.482	-6.743	-6.628	3.572	-4.954	-4.752
R2adj	0.001	0	0.017	0.084	0.117	0.151	0.086	0.116	0.126
N	425	425	425	425	425	425	425	425	425

Panel C Price of risk vs. aggregate variance

	Δ Cons. Variance	Δ Equity RV	Δ FX RV
Beta	-22508.189	-1.478	-12.372
$t-$stat	-3.928	-0.717	-2.806
R^2_{adj}	0.140	0.000	0.032
N	129	399	399

Panel A gives the results from regressing the UMVE portfolio returns on various factors in univariate regressions. Δc denote log real, per capita nondurable+services consumption growth. HML is the high-minus-low book-to-market factor from the Fama-French 5-factor model, whereas CMA is the Conservative minus Aggressive investment factor from the same. Panel B shows regressions of UMVE, the CS Carry strategy, and the unpriced component of CS carry returns on the intermediary capital factor from He, Kelly, and Manela (2018), HKM, as well as shocks to realized variance in the equity and currency markets. Panel C displays a regression of the 3-year change in the conditional price of risk on the 3-year change in conditional consumption variance (obtained from a GARCH(1,1)), equity variance, and currency (FX) variance, respectively. We use 3-year changes instead of levels due to a downward drift in the price of risk over the sample. The $t-$statistics are heteroskedasticity-adjusted. The sample is monthly from January 1985 to May 2020.