

NOTES

A NOTE ON THE FORWARD AND THE EQUITY PREMIUM PUZZLES: TWO SYMPTOMS OF THE SAME ILLNESS?

CARLOS E. DA COSTA AND JOÃO V. ISSLER

FGV/EPGE

PAULO F. MATOS

UFC/CAEN

We build a stochastic discount factor—SDF—using U.S. domestic financial data only, and provide evidence that it accounts for stylized facts about foreign markets that escape SDFs generated by consumption-based models. When our SDF is interpreted as the projection of the pricing kernel from a fully specified model in the space of returns, our results indicate that a model that accounts for the behavior of domestic assets goes a long way toward accounting for the behavior of foreign asset prices. In our tests, we address predictability, a defining feature of the forward premium puzzle—FPP—by using instruments that are known to forecast excess returns in the moment restrictions associated with Euler equations both in the equity and in the foreign markets.

Keywords: Equity Premium Puzzle, Forward Premium Puzzle, Return-Based Pricing Kernel

1. INTRODUCTION

The forward premium puzzle—henceforth, FPP—is how one calls the systematic departure from the intuitive proposition that, conditional on available information, the expected return to speculation in the forward foreign exchange market should be zero.

Because rational expectations alone do not restrict the behavior of forward rates, as emphasized by Fama (1984), one may always include a risk-premium term that reconciles the time series behavior of the associated data. Therefore, what makes

This paper has greatly benefited from comments and suggestions by Caio Almeida, Marco Bonomo, Luis Braidó, Ricardo Brito, Marcelo Fernandes, Jaime Filho, Marcio Garcia, Luiz Lima, Celina Ozawa, Rafael Santos, Andre Silva, and seminar participants at EPGE, PUC-Rio, CAEN-UFC, the XXVII SBE Meeting, Econometrics in Rio, the 2006 Meeting of the Brazilian Finance Society, the 61th ESEM, and the 1st Luso-Brazilian Finance Meeting. Matos gratefully acknowledges financial support from FUNCAP and CNPq-Brazil. Issler and da Costa gratefully acknowledge financial support from CNPq-Brazil and FAPERJ. Address correspondence to: João V. Issler, Escola de Pós-Graduação em Economia da Fundação Getúlio Vargas, Praia de Botafogo, 190 11o andar, Rio de Janeiro, RJ 22250-900, Brazil; e-mail: jissler@fgv.br.

the departure from an intuitive proposition puzzling is the fact that no sensible model for the risk premium seems to account for the forward premium behavior.

The natural candidate for a sensible model is, of course, the consumption capital asset-pricing model—CCAPM—whose use to investigate the FPP was pioneered by Mark (1985). Using a nonlinear generalized method of moments (GMM) approach, he estimated extremely high values for the representative agents' risk aversion parameter and offered evidence of the canonical consumption model's inability to account for its overidentifying restrictions. Similar findings in equity markets by Hansen and Singleton (1982) and Mehra and Prescott (1985) led to the recognition of the so-called equity premium puzzle—henceforth, EPP.

Despite these similar findings, the literature on the FPP and the EPP followed very different paths. We may wonder why this happened but the existence of an early specificity for the FPP with no parallel in the case of the EPP—the predictability of returns based on interest rate differentials—may be to blame.¹ Moreover, the absence of a successful model for the risk premium meant that one could not formally assess whether the two puzzles had the same root.

The purpose of this paper is to offer suggestive evidence that the solution to the puzzle in domestic assets is to account for the analogous puzzle in foreign-asset pricing. What we lack, in our opinion, is a proper model for asset pricing, which, if found, would solve both puzzles at once. Because a direct answer to our question cannot be given without actually writing down such a model, we devise an indirect strategy. We build a return-based pricing-kernel estimate using only U.S. domestic assets and provide evidence that it successfully prices the forward premium.

The return-based kernel is a stochastic discount factor—henceforth, SDF—mimicking portfolio: the projection of any fully specified economic model's SDF (yet to be written) on the space of payoffs. An advantage of concentrating on the projection is that we can approximate it arbitrarily well in sample using statistical methods and asset returns alone. Using such a projection not only circumvents the nonexistence of a proper consumption model but is also guaranteed not to underperform such an ideal model in sample—see Hansen and Jagannathan (1991) and Hansen and Jagannathan (1997).²

We estimate the SDF mimicking employing the unconditional linear multifactor model using standard techniques. To assess whether this SDF prices the forward premium, we run tests that are based on Euler equations. We exploit the theoretical lack of correlation between discounted risk premia and variables in the conditioning set, or between discounted returns and their respective theoretical means.

After failing to reject the null that average discounted returns of representative domestic assets are equal to unity and that excess returns are equal to zero, we take our estimated pricing kernel to try to price foreign assets for the widest possible group of developed countries with a long enough span of future exchange-rate data: Canada, Germany, Japan, Switzerland, and the United Kingdom. For all countries, we fail to reject the null that average discounted excess returns on currency speculation are zero and that average discounted returns are equal to

one, using instruments known to having strong forecasting power for currency movements—one of the defining features of the FPP. Identical results are reported for all the overidentifying restriction tests, except for the Japanese case. In a joint test of the return-based pricing kernels, results confirm the previous evidence, with isolated rejections only in Wald tests for the Canadian and German cases.

We finally run robustness exercises to show that our empirical results are based on insignificant pricing deviations that come from reasonable values for these deviations, rather than higher variance, i.e., lower power of the test.

The remainder of the paper is organized as follows. Section 2 gives an account of the literature that tries to explain the FPP and is related to our current effort. Section 3 discusses the techniques used to estimate the SDF and the pricing tests implemented in this paper. Section 4 presents the empirical results obtained in this paper. Concluding remarks are offered in Section 5.

2. CRITICAL APPRAISAL OF CURRENT DEBATE ON THE FORWARD PREMIUM PUZZLE

Let R^i denote the return on asset i , and recall that, given free portfolio formation, the law of one price guarantees, through the Riesz representation theorem, the existence of an SDF, M_{t+1} , satisfying³

$$1 = \mathbf{E}_t [M_{t+1} R_{t+1}^i], \quad \forall i = 1, 2, \dots, N. \tag{1}$$

The reason that we find the behavior of the equity premium puzzling is that the pricing kernel generated by the CCAPM, $M_{t+1} = \beta U'(C_{t+1})/U'(C_t)$, does not satisfy

$$0 = \mathbf{E}_t [M_{t+1} (R_{t+1}^i - R_{t+1}^j)], \quad \forall i, j, \tag{2}$$

when $R_{t+1}^i = (1 + i_{t+1}^{\text{SP}})P_t/P_{t+1}$, where i_{t+1}^{SP} is the nominal return on S&P500, $R_{t+1}^j = (1 + i_{t+1}^b)P_t/P_{t+1}$, where i_{t+1}^b is the nominal return on the U.S. Treasury Bill, and P_t is the dollar price level when empirically sound parameters are chosen for the representative agent’s preferences.

Although seldom stated along the same lines, the FPP is also associated with the fact that no theoretically sound economic model seems to offer a definition of risk capable of correctly pricing the forward premium.⁴ Let the covered, R^C , and the uncovered return, R^U , on foreign government bonds trade be represented by

$$R_{t+1}^C = \frac{{}_t F_{t+1}(1 + i_{t+1}^*)P_t}{S_t P_{t+1}} \quad \text{and} \quad R_{t+1}^U = \frac{S_{t+1}(1 + i_{t+1}^*)P_t}{S_t P_{t+1}}, \tag{3}$$

respectively, where ${}_t F_{t+1}$ and S_t are the forward and spot prices of foreign currency in units of domestic currency and i_{t+1}^* is the nominal net return on a foreign asset in terms of the foreign investor’s currency. Assuming $u(C) = C^{1-\alpha}(1 - \alpha)^{-1}$,

Mark (1985) applied Hansen's (1982) GMM to estimate and test

$$0 = \mathbf{E}_t \left[\frac{u'(C_{t+1})}{u'(C_t)} \frac{P_t(1 + i_{t+1}^*) [F_{t+1} - S_{t+1}]}{S_t P_{t+1}} \right]. \quad (4)$$

He estimated a coefficient of relative risk aversion, $\hat{\alpha}$, above 40, and tested and rejected the overidentifying restrictions when the forward premium and its lags were used as instruments.⁵ Analogous findings for the equity premium define the EPP.

Even though important progress has been made in building more successful consumption models, there is no consensus as of this moment on whether any of the current models derived from the primitives of the economy accounts for either puzzle.⁶ The current state of the art thus precludes a direct answer to the question in the title of this paper. To resolve this issue, we devise an indirect approach under which we extract a pricing kernel from U.S. return data alone and show that it prices both the domestic and the foreign-exchange returns and excess returns.

Although the law of one price guarantees that an SDF does exist, uniqueness only obtains if markets are complete. If markets are incomplete, there exists a continuum of SDFs, M_{t+1} , that price all traded securities correctly. If, however, one fixes the set of assets, then there is a unique projection of all those SDFs in the space of returns associated with this set of assets—see Hansen and Jagannathan (1991, 1997).

We rely on this projection argument to combine statistical methods with the asset pricing equation (1) to devise SDF-mimicking portfolio estimates. Our exercise consists of exploring a large cross section of U.S. time-series stock returns to construct return-based pricing kernel estimates satisfying the pricing equation (1) for that group of assets. We then take this “domestic” SDF estimate and use it to price assets not used in constructing it: foreign assets. Therefore, we perform a genuine out-of-sample forecasting exercise using SDF-mimicking portfolio estimates, avoiding the in-sample overfitting critique discussed in Cochrane (2001), for example.

Our empirical strategy exploits the fact that the mimicking portfolio has identical in-sample-pricing properties vis-à-vis any other proper SDF. In this sense, we are guaranteed not to underperform any model when the set of assets is the one from which the mimicking portfolio is built. Next, we take a larger set of assets—that is, we include the foreign assets for which the FPP is defined, and use the SDF projection to try to price these larger set of assets. As we will discuss later, we cannot reject the null that the mimicking portfolio extracted from the restricted set of assets prices the assets of the larger set.

Of course we cannot guarantee that any SDF whose projection is the return on the SDF-mimicking portfolio will price the assets in the larger set correctly. The component that is orthogonal to the space of returns of the smaller set of assets may worsen the pricing properties of the mimicking portfolio. As we have said, a complete answer to the question in the title of this paper requires one to actually

write down the asset-pricing model. Still, our results represent a strong indication that the SDF will price the assets correctly.

We cannot overemphasize the importance of out-of-sample forecasting for our purposes. The main point of this paper is to show that the forward and the equity premium puzzle are intertwined. Under the law of one price, an SDF that prices all assets necessarily exists. Thus, an in-sample exercise would only provide evidence that the forward-premium puzzle is not simply a consequence of violations of the law of one price. We aim at showing more: an SDF can be constructed using only domestic assets, i.e., using the same source of information that guides research regarding the equity premium puzzle, and still price foreign assets. It is our view that this SDF is to capture the growth of the marginal utility of consumption in a model yet to be written.

3. ECONOMETRIC TESTS

3.1. Return-Based Pricing Kernels

To see how we build \hat{M}_{t+1} , start with the unconditional version of the standard asset-pricing model, letting R be a vector of returns including the returns of a risk-free asset. Then

$$\iota = \mathbf{E}[RM], \tag{5}$$

where M is the SDF of a proper model pricing all assets whose returns are recorded, and ι is the unit vector. Despite the fact that we do not observe M , Hansen and Jagannathan (1991) propose a projection of M onto the space of returns that preserves linear pricing. They label it by M^* :

$$M^* = R'\beta, \tag{6}$$

where β are the weights in the linear projection. Consider now the direct sum

$$M = M^* + \nu, \tag{7}$$

where ν is orthogonal to the space of returns R because of standard projection arguments. This implies that M and M^* have identical pricing properties concerning (5), leading to⁷

$$\iota = \mathbf{E}[RR'\beta] + \mathbf{E}[R\nu] = \mathbf{E}[RR']\beta \tag{8}$$

or

$$\beta = (\mathbf{E}[RR'])^{-1}\iota. \tag{9}$$

Thus, despite not observing M , we can construct a mimicking portfolio with identical pricing properties:

$$M^* = R'(\mathbf{E}[RR'])^{-1}\iota,$$

which is only a function of observables.⁸

3.2. Prediction and Pricing Tests

Because the two puzzles are present in logs and in levels, we may work on an Euler-equation framework. This avoids the imposition of stringent auxiliary restrictions in hypothesis testing, while keeping the possibility of testing the conditional moments through the use of lagged instruments along the lines of Hansen and Singleton (1982, 1984) and Mark (1985). Euler equations (1) and (2) must hold for all assets and portfolios.

If we had observations on M^* , return data could be used to test directly whether they held. Of course, M^* is a latent variable, and the best we can do is to find a consistent estimator for M^* . With a large enough sample of return data, so that M^* and their estimators are “close enough,” we may still directly test the validity of these Euler equations, (1) and (2).

We employ standard econometric techniques to construct an estimate of M_{t+1}^* relying on standard methods: principal-component and factor analyses, coupled with the asset-pricing equation; see, e.g., Cochrane (2001, Ch. 6). These methods are asymptotic: either $N \rightarrow \infty$ or $N, T \rightarrow \infty$, relying on a weak law of large numbers to provide consistent estimators of the SDF mimicking portfolio—the unique systematic portion of asset returns. The assets used to estimate the multifactor version of the SDF, labeled $\widetilde{M}_t^{\text{US\$}}$, are described in Subsection 4.1. Only domestic (U.S.) assets are used in estimating $\widetilde{M}_t^{\text{US\$}}$, which is an important part of our empirical strategy.

In-sample exercise: The equity premium puzzle. Let \mathbf{z}_t be a vector of instrumental variables, all measurable with respect to $\mathbf{E}_t(\cdot)$. Employing *scaled returns* and *scaled excess returns*—defined as $R_{t+1}^i \times z_t$ and $(R_{t+1}^i - R_{t+1}^j) \times z_t$, respectively—we are able to test the conditional moment restrictions associated with (1) and (2). This is particularly important in the case of the FPP, because the overidentifying restriction associated with having the current forward premium itself as an instrument was usually rejected both when the CCAPM was employed and when returns were not discounted: a manifestation of its predictive power.

The first exercise consists of using $\widetilde{M}_t^{\text{US\$}}$ to price the S&P500 and the 90-day T-bill. Its purpose is to investigate if the anomaly related to the EPP is present with our pricing-kernel estimate.

Under a consumption-based approach, it is usual to estimate the utility function parameters and then to test the associated system of orthogonality restrictions using GMM. In our two-step procedure, we need not estimate the utility specification parameters, for we have already constructed a time series for the pricing kernel. The only parameters to be estimated are the pricing deviations, μ_1 and μ_2 , in the following system:

$$\mathbf{0} = \mathbf{E} \left\{ \begin{bmatrix} \widetilde{M}_t^{\text{US\$}} \frac{(i_{t+1}^{\text{SP}} - i_{t+1}^b) P_t}{P_{t+1}} - \mu_1 \\ \widetilde{M}_t^{\text{US\$}} \frac{(1+i_{t+1}^{\text{SP}}) P_t}{P_{t+1}} - (1 + \mu_2) \end{bmatrix} \otimes \mathbf{z}_t \right\}, \tag{10}$$

where i_{t+1}^{SP} and i_{t+1}^b are respectively the nominal returns on the S&P500 and on a U.S. government short-term bond. The deviations, μ_i , are estimated by minimizing any quadratic form of the sample mean of the errors for some weighting matrix. The estimation of a nonsignificant deviation may be interpreted as a first evidence that the pricing kernel that we use is able to approximate the discounted payoff on trading foreign government bonds and their price. Once we have estimated the deviations that produces the best fit, testing the system of overidentifying restrictions allows us to assess the overall fit of the model.⁹

Reassured that our pricing kernel does a good job in pricing the relevant U.S. domestic assets, we proceed to our main exercise. We test the pricing properties of the same pricing kernel for foreign assets: now, an out-of-sample exercise.

Out-of-sample exercise: The forward premium puzzle. The out-of-sample test for the foreign exchange market parallels the in-sample exercise with $(1 + i_{t+1}^*)[{}_tF_{t+1}]/S_t$ and $(1 + i_{t+1}^*)[{}_tS_{t+1}]/S_t$ substituting for i_{t+1}^{SP} and i_{t+1}^b , respectively:

$$\mathbf{0} = \mathbf{E} \left\{ \begin{bmatrix} \widetilde{M}_t^{USS} \frac{{}_tP_t(1+i_{t+1}^*)[{}_tF_{t+1} - S_{t+1}]}{S_t P_{t+1}} - \mu_1 \\ \widetilde{M}_t^{USS} \frac{S_{t+1}(1+i_{t+1}^*)P_t}{S_t P_{t+1}} - (1 + \mu_2) \end{bmatrix} \otimes \mathbf{z}_t \right\}. \tag{11}$$

Once again, we estimate the deviations μ_1 of the null price for the excess return and μ_2 of the unitary price for the uncovered return. For (11) to hold, we must have $\mu_1 = 0$ and $\mu_2 = 0$, which can be jointly tested using a Wald test. As in the in-sample exercise, we also check whether the overidentifying restrictions are not rejected using the J -test.¹⁰ Finally, we include a joint test of the return-based pricing kernels to assess whether both EPP and FPP hold jointly. We test pricing deviations individually and jointly along with the restriction J -test.

3.3. GMM Estimation Setup and Instruments

In choosing how to weight the different moments used for estimation of the parameters of interest, one faces a trade-off between attaining full efficiency and correctly specifying the structure of the variance–covariance matrix; see the discussion in a GMM framework in Cochrane (2001, Ch. 11). As is well known, both OLS and GLS are consistent under correct specification of the variance–covariance matrix. However, in trying to achieve full efficiency, one can render GLS estimates inconsistent if the structure chosen is incorrectly specified. OLS is a robust estimate in the sense that it does not rely on a correct choice for the weighting matrix. For that reason, most applied econometric studies use OLS in estimation and properly estimate its standard errors using the methods advanced by White (1980). Similarly, using the identity matrix instead of optimal GMM weights (\hat{S}^{-1}) may be better under some circumstances considered here, especially when the number of restrictions is large vis-à-vis the number of time-series observations used in estimating optimal GMM weights.¹¹

For our purposes, there are two relevant tests that can be performed in a GMM setup. The first is a test of whether the pricing errors in (10) and (11) are statistically zero, which uses a robust Wald test on individual and joint significance of the pricing errors in these systems. The second is a standard overidentifying restrictions test (*J*-test), which has the usual interpretation of a test of orthogonality between the errors in each moment restriction and the instruments used in GMM estimation, thus being a specification test for the validity of instruments. Note, however, that if an optimal method of GMM estimation is not employed, and an identity matrix is used in weighting the moments, we should not place too much emphasis on the results of the *J*-test.

Pricing assets in domestic and foreign markets individually. In this paper, $T = 112$, whereas the number of instruments $k = 2, 3$, or 4 , and the number of asset-pricing equations $N = 2$ or 4 , when assets in domestic and in foreign markets are priced. To implement the in-sample test considering moment restrictions (10) and to account for the FPP based on the system of orthogonality restrictions (11), we will use optimal GMM, because we have a reasonable number of moment conditions, 6 and 4, vis-à-vis our sample size. Following Hansen and Singleton (1982), our estimates are produced by an iterative procedure.

As instruments we use specific financial variables carefully chosen according to their forecasting potential. For the EPP, we employ the dividend–price ratio and the investment–capital ratio.¹² For the FPP, we use the current value of the respective forward premium, $({}_t F_{t+1} - S_t)/S_t$, because its forecasting ability is a defining feature of this puzzle and this series is measurable with respect to the information set used by the representative consumer. In addition to these variables, we used lagged values of returns for the assets being tested as a robustness check. Taking into account the fact that expected returns and business cycles are correlated, e.g., Fama and French (1989), we also use macroeconomic variables with documented forecasting ability regarding financial returns—e.g., real consumption and GDP instantaneous growth rates, consumption-GDP ratio—as instruments.

Pricing assets in both markets jointly. In the joint tests for domestic and foreign markets, we have the number of instruments $k = 4$ and the number of asset-pricing equations $N = 4$. This leaves us with 16 moment conditions, a relatively large number for our sample of 112 time-series observations. Therefore, we will use the identity matrix in weighting orthogonality conditions, but will still use robust estimates for the variance–covariance matrix of our estimates using the techniques in Newey and West (1987), and primarily use the Wald test discussed previously to verify whether pricing errors are statistically zero.

4. EMPIRICAL RESULTS

4.1. Data and Summary Statistics

In terms of sample size, our main limitation for the time-series span used here regards the FPP tests. Foreign exchange only floated post–Bretton Woods.

The Chicago Mercantile Exchange—CME, the pioneer of the financial-futures market—only launched currency futures in 1972. In addition to that, only futures data for a few developed countries are available since then. In order to have a common sample for the largest possible set of countries, we used U.S. foreign-exchange data for Canada, Germany, Japan, Switzerland, and the United Kingdom, covering the period from 1977:1 to 2004:4, at a quarterly frequency, comprising 112 observations.

Returns from covered and uncovered trading with foreign government bonds were transformed into real returns using the U.S. Consumers Price Index—the CPI. The forward-rate series were extracted from the CME database, whereas the spot-rate series were extracted from the Bank of England and the IFS databases. Returns on the S&P500 and on 90-day T-bills were extracted in nominal terms from the CRSP and the IFS database, respectively. Real returns were obtained using the U.S. CPI as deflator.

$\widetilde{M}_t^{\text{US\$}}$ is constructed as a linear function of factors. Using a principal-components technique, seven factors are extracted from (i) the U.S.\$ real returns on all U.S. stocks traded on NYSE along the period in question, in a totality of 464 stocks, according to the CRSP database, and (ii) the real returns on one hundred size-*BE/ME* Fama–French portfolios. In addition to these seven factors, we also use U.S. real returns on real estate, bonds on AAA U.S. corporations, S&P500, AMEX, NASDAQ, NYSE, and Dow Jones as factors. All these factors are used in constructing our estimate of $\widetilde{M}_t^{\text{US\$}}$ following standard techniques; see the discussion in Cochrane (2001, Ch. 6).

All macroeconomic variables used in econometric tests were extracted from FED's FRED database. We also employed additional forecasting financial variables that are specific to each test performed, and are listed in the appropriate tables of results.

4.2. SDF Estimates

Under GMM, we need stationarity to obtain consistent estimates. We tested all returns used either in estimating SDFs or in the pricing test procedures for the presence of unit roots using the robust test in Phillips and Perron (1988), including an intercept in the test regression. For almost all cases, we rejected the null of a unit root at the 5% level and rejected the null for all cases at the 10% level.¹³

As for the number of factors used in constructing $\widetilde{M}_t^{\text{US\$}}$, we examined the time plot of the eigenvalues ordered from the largest to the smallest, and found that seven factors accounted for most of the variation in the returns on stocks and the Fama–French portfolios. This choice was close to the one in Connor and Korajczyk (1993), who examined returns from stocks listed on NYSE and AMEX.

The estimate of M_t^* is plotted in Figure 1.

In Figure 2, we plot mean–variance frontiers for the discount factors, following Hansen and Jagannathan (1997). The gray line restricts the moments of a pricing




FIGURE 1. Multifactor pricing kernel. $\widetilde{M}_t^{\text{US\$}}$ is a return-based pricing kernel constructed as a linear function of factors. Seven factors are extracted from the U.S.\$ real returns on 464 stocks traded on the NYSE and real returns on 100 size-*BE/ME* Fama–French portfolios. In addition to these factors, we also use U.S.\$ real returns on gold, U.S. real estate, bonds on AAA U.S. corporations, S&P500, AMEX, NASDAQ, NYSE, and Dow Jones as factors. Data from 1977:1 to 2004:4, 112 observations.

kernel needed to price unconditionally the set of U.S.\$ real returns on gold, U.S. real estate, bonds on AAA U.S. corporations, S&P500, AMEX, NASDAQ, NYSE, and Dow Jones. The black line restricts this set of assets and also the set comprising returns on uncovered and covered trading with British, Canadian, German, Japanese, and Swiss government bonds. Comparing the Hansen and Jagannathan (1997) bounds for the pricing kernels, for periods for which the U.S. real risk-free rates remained below 2.64% per year, shows that a more volatile pricing kernel is needed to account for the EPP than to accommodate the FPP anomaly. As this real rate decreases, foreign assets provide less information on valid SDFs in the extended bounds for U.S. and foreign assets, and the difference between the gray and black line representing the bounds is reduced.

4.3. Pricing-Test Results

In-sample exercise: The equity premium puzzle. Table 1 displays the performance of our return-based pricing kernel $\widetilde{M}_t^{\text{US\$}}$ in pricing returns related to the EPP. We label these *in-sample* tests, in which we test the system of orthogonality restrictions in (10), which accounts for the equity premium, $R_{t+1}^{\text{SP}} - R_{t+1}^{\text{b}}$, and for the return on the S&P500, R_{t+1}^{SP} . We use a constant as instrument as well as $(D/P)_t$ and $(I/C)_t$. We fail to reject the null that pricing errors are statistically zero individually and jointly statistically zero. Also, the *J*-test shows no signs




FIGURE 2. Hansen and Jagannathan (1997) bounds for pricing kernels. In this figure, we plot a mean–variance frontier of all discount factors that price two sets of assets. The gray line restricts the moments of a pricing kernel in order to price unconditionally the set of U.S. real returns on gold, U.S. real estate, bonds on AAA U.S. corporations, S&P500, AMEX, NASDAQ, NYSE, and Dow Jones. The gray line (dashed) restricts the set of pricing kernels given the returns on uncovered and covered trading with British, Canadian, German, Japanes, and Swiss government bonds. The black line restricts both sets of assets. The horizontal axis contains the unconditional mean and the vertical one contains the unconditional standard deviation. Data from 1977:1 to 2004:4, 112 observations.

of misspecification for our estimates. We also use $\widetilde{M}_t^{US\$}$ to price the six Fama and French (1993) benchmark portfolios, dynamic portfolios extracted from the Fama–French library. Once again, the performance of $\widetilde{M}_t^{US\$}$ is reflected in the fact that there is no rejection of the null.¹⁴

Out-of-sample exercise: The forward premium puzzle. Table 2 presents results for what we have labelled *out-of-sample* pricing tests, for which $\widetilde{M}_t^{US\$}$ is used to price foreign assets related to the FPP. We report the results of tests for the system of orthogonality restrictions (11). There are two Euler equations: one with the excess return on uncovered over covered trade with foreign government bonds, $(R_{t+1}^U - R_{t+1}^C)$, and one with the uncovered return, R_{t+1}^U . The instrument set includes a constant and the variable $({}_t F_{t+1}^i - S_t^j) / S_t^j$ that generates the FPP.

At the 5% significance level, all individual pricing deviations are statistically zero for all foreign currencies studied here. Also, Wald-test results show no sign of joint significance for currency-pricing errors. Moreover, the size of the pricing deviations is relatively modest: $\widehat{\mu}_2$ is comparable to returns on uncovered trading, whereas $\widehat{\mu}_1$ is lower than the respective excess returns. Except in the

TABLE 1. Equity premium puzzle test: In-sample asset-pricing exercise for S&P500 and American 90-day T-bill^a

System of conditional moment restrictions:

$$\mathbf{0} = \mathbf{E} \left\{ \begin{bmatrix} \widetilde{M}_t^{\text{US\$}} \frac{(i_{t+1}^{\text{SP}} - i_{t+1}^{\text{b}}) P_t}{P_{t+1}} - \mu_1 \\ \widetilde{M}_t^{\text{US\$}} \frac{(1+i_{t+1}^{\text{SP}}) P_t}{P_{t+1}} - (1 + \mu_2) \end{bmatrix} \otimes \mathbf{z}_t \right\} \quad \mathbf{z}_t = (D/P)_t, (I/C)_t$$

Wald test (deviations)	Asset pricing and deviations			Overall fit				
$H_0: \mu_1 = \mu_2 = 0$	$\widehat{\mu}_1$	0.0052	(0.758)	[0.492]	J_T Test	0.0377		
χ^2	0.5242	[0.769]	$1 + \widehat{\mu}_2$	0.9898	(5.871)	[0.862]	P-value	[0.377]

*Indicates the rejection of the null hypothesis (individually insignificant pricing deviation) at 5% level.
 **Indicates the rejection of the null hypothesis (jointly insignificant pricing deviations) in Wald test at 5% level.
 ***Indicates the rejection of the validity of the overidentifying restrictions at 5% significance level.
^aTechnique: Hansen’s (1982) Generalized Method of Moments is used to test Euler equations and to estimate the model parameters, over the period from 1977:1 to 2004:4, 112 observations. The GMM procedure forms a consistent estimate of the weighting matrix and then uses it to iterate the coefficient estimates until convergence. Respective standard errors are reported in the parentheses, whereas *p*-values are in the brackets. The standard errors are reported multiplied by 100.

Japanese case, results for all overidentifying restriction tests show no signs of misspecification.

Wrapping up, a domestically constructed pricing kernel $\widetilde{M}_t^{\text{US\$}}$, which prices domestic markets correctly, also prices foreign markets correctly.

Forward and equity premium puzzles jointly. In Table 3 we present joint tests of the FPP and of the EPP for Canada, Germany, Japan, Switzerland, and the United Kingdom. These pricing restrictions stack the Euler equations that characterize the EPP tested in Table 1 and those that characterize the FPP tested in Table 2. The instrument set is the union of the sets previously considered: a constant, $({}_t F_{t+1}^j - S_t^j) / S_t^j$, $(D/P)_t$, and $(I/C)_t$. In this strongest test of the pricing kernels, the results confirm that return-based kernels generate no puzzles whatsoever in pricing returns and excess returns for the United Kingdom, Japan, and Switzerland: at 5% significance, all pricing errors are statistically zero individually and jointly. However, for Canada and Germany, joint tests (Wald) reject the null of zero pricing errors for domestic and foreign markets. Regarding the *J*-test, we cannot reject the null of no misspecification for any of those countries. However, because here we employ an identity matrix in weighting moment restrictions, we should not put much emphasis on the latter.

Additional tests and robustness. To account for other domestic and international market stylized facts that escape consumption-based models, we also use U.S.\$ real returns on the Fama and French (1993) benchmark portfolios formed on size and book-to-market, extracted from the French data library, as well as eight foreign currency portfolios, constructed by da Costa et al. (2010) using quarterly

TABLE 2. Forward premium puzzle test: Out-of-sample asset-pricing exercise for the foreign currencies^a

$$\text{System of conditional moment restrictions: } \mathbf{0} = \mathbf{E} \left\{ \left[\begin{array}{l} \widehat{M}_t^{\text{USS}} \frac{P_t(1+i_{t+1}^*)P_{t+1} - S_{t+1}^j}{S_t^j P_{t+1}} - \mu_1 \\ \widehat{M}_t^{\text{USS}} \frac{S_{t+1}^j(1+i_{t+1}^*)P_t}{S_t^j P_{t+1}} - (1 + \mu_2) \end{array} \right] \otimes \mathbf{z}_t \right\} \quad \mathbf{z}_t = ({}_tF_{t+1}^j - S_t^j)/S_t^j$$

Wald test (deviations)	Asset pricing and deviations	Overall fit	Wald test (deviations)	Asset pricing and deviations	Overall fit
Results for British government bonds:			Results for Canadian government bonds:		
$H_0: \mu_1 = \mu_2 = 0$ χ^2 5.1523 [0.076]	$\widehat{\mu}_1$ 0.0123 (0.637) [0.056] $1 + \widehat{\mu}_2$ 0.9362 (7.360) [0.387]	J_T Test 0.0503 P-value [0.060]	$H_0: \mu_1 = \mu_2 = 0$ χ^2 1.4302 [0.489]	$\widehat{\mu}_1$ 0.0033 (0.334) [0.319] $1 + \widehat{\mu}_2$ 0.9577 (7.679) [0.582]	J_T Test 0.0530 P-value [0.051]
Results for German government bonds:			Results for Japanese government bonds:		
$H_0: \mu_1 = \mu_2 = 0$ χ^2 2.2808 [0.320]	$\widehat{\mu}_1$ 0.0100 (0.776) [0.199] $1 + \widehat{\mu}_2$ 0.9740 (6.666) [0.697]	J_T Test 0.0473 P-value [0.071]	$H_0: \mu_1 = \mu_2 = 0$ χ^2 2.2866 [0.319]	$\widehat{\mu}_1$ 0.0103 (0.853) [0.229] $1 + \widehat{\mu}_2$ 0.9583 (7.799) [0.594]	J_T Test 0.0747*** P-value [0.015]
Results for Swiss government bonds:					
$H_0: \mu_1 = \mu_2 = 0$ χ^2 0.8176 [0.665]	$\widehat{\mu}_1$ 0.0081 (0.901) [0.367] $1 + \widehat{\mu}_2$ 1.0001 (7.578) [0.999]	J_T Test 0.0247 P-value [0.252]			

*Indicates the rejection of the null hypothesis (insignificant pricing deviation) at 5% level. **Indicates the rejection of the null hypothesis (jointly insignificant pricing deviations) in Wald test at 5% level. ***Indicates the rejection of the validity of the overidentifying restrictions at 5% significance level.

^aTechnique: Hansen's (1982) Generalized Method of Moments is used to test Euler equations and to estimate the model parameters, over the period from 1977:1 to 2004:4, 112 observations. The GMM procedure forms a consistent estimate of the weighting matrix and then uses it to iterate the coefficient estimates until convergence. We use the superscript *j* for forward, spot and interest rates in order to associate these variables to the country *j*. Standard errors are reported in the parentheses, whereas *p*-values are in the brackets. The standard errors are reported multiplied by 100.

TABLE 3. The forward and equity premium puzzles joint test: Out-of-sample asset-pricing exercise^a

$$\text{System of conditional moment restrictions: } \mathbf{0} = \mathbf{E} \left\{ \begin{bmatrix} \widehat{M}_t^{\text{USS}} \frac{P_t (1+i_t^{*j}) (F_{t+1}^j - S_{t+1}^j)}{S_t^j P_{t+1}} - \mu_1 \\ \widehat{M}_t^{\text{USS}} \frac{S_{t+1}^j (1+i_t^{*j}) P_t}{S_t^j P_{t+1}} - (1 + \mu_2) \\ \widehat{M}_t^{\text{USS}} \frac{(F_{t+1}^{\text{SP}} - F_{t+1}^{\text{b}}) P_t}{P_{t+1}} - \mu_3 \\ \widehat{M}_t^{\text{USS}} \frac{(1+i_t^{\text{SP}}) P_t}{P_{t+1}} - (1 + \mu_4) \end{bmatrix} \otimes \mathbf{z}_t \right\}$$

$$\mathbf{z}_t = (F_{t+1}^j - S_t^j) / S_t^j, (D/P)_t, (I/C)_t$$

Wald test (deviations)	Asset pricing and deviations	Overall fit	Wald test (deviations)	Asset pricing and deviations	Overall fit
Results for British government bonds:			Results for Canadian government bonds:		
$H_0: \mu_i = 0, i = 1, \dots, 4$	$\widehat{\mu}_1$ 0.0089 (0.701) [0.205]	J_T Test 0.0524	$H_0: \mu_i = 0, i = 1, \dots, 4$	$\widehat{\mu}_1$ 0.0028 (0.338) [0.401]	J_T Test 0.0469
χ^2 2.6653 [0.615]	$1 + \widehat{\mu}_2$ 1.0074 (8.007) [0.926]	P-value [0.923]	χ^2 13.5724** [0.008]	$1 + \widehat{\mu}_2$ 1.0030 (8.026) [0.971]	P-value [0.949]
	$\widehat{\mu}_3$ 0.0072 (0.899) [0.421]			$\widehat{\mu}_3$ 0.0072 (0.899) [0.421]	
	$1 + \widehat{\mu}_4$ 1.0059 (8.085) [0.942]			$1 + \widehat{\mu}_4$ 1.0059 (8.085) [0.942]	
Results for German government bonds:			Results for Japanese government bonds:		
$H_0: \mu_i = 0, i = 1, \dots, 4$	$\widehat{\mu}_1$ 0.0074 (0.824) [0.372]	J_T Test 0.0477	$H_0: \mu_i = 0, i = 1, \dots, 4$	$\widehat{\mu}_1$ 0.0070 (0.874) [0.423]	J_T Test 0.0484
χ^2 12.1051** [0.017]	$1 + \widehat{\mu}_2$ 1.0077 (8.013) [0.923]	P-value [0.946]	χ^2 2.0286 [0.731]	$1 + \widehat{\mu}_2$ 1.0049 (8.137) [0.953]	P-value [0.943]
	$\widehat{\mu}_3$ 0.0072 (0.899) [0.421]			$\widehat{\mu}_3$ 0.0072 (0.899) [0.421]	
	$1 + \widehat{\mu}_4$ 1.0059 (8.085) [0.942]			$1 + \widehat{\mu}_4$ 1.0059 (8.085) [0.942]	
Results for Swiss government bonds:					
$H_0: \mu_i = 0, i = 1, \dots, 4$	$\widehat{\mu}_1$ -0.0006 (1.264) [0.961]	J_T Test 0.0454			
χ^2 4.8852 [0.299]	$1 + \widehat{\mu}_2$ 1.0098 (7.994) [0.819]	P-value [0.955]			
	$\widehat{\mu}_3$ 0.0072 (0.899) [0.421]				
	$1 + \widehat{\mu}_4$ 1.0059 (8.085) [0.942]				

^{*}Indicates the rejection of the null hypothesis (insignificant pricing deviation) at 5% level. ^{**}Indicates the rejection of the null hypothesis (jointly insignificant pricing deviations) in Wald test at 5% level. ^{***}Indicates the rejection of the validity of the overidentifying restrictions at 5% significance level.
^aTechnique: Hansen's (1982) Generalized Method of Moments is used to test Euler equations and to estimate the model parameters, over the period from 1977:1 to 2004:4, 112 observations. The GMM procedure minimizes a quadratic form of the sample mean errors for the identity matrix. We use the superscript *j* for forward, spot and interest rates in order to associate these variables to the country *j*. Standard errors are reported in the parentheses, whereas *p*-values are in the brackets. The standard errors are reported multiplied by 100.

series of spot exchange rates and short-term foreign government bonds available in IFS database. They follow Lustig and Verdelhan's (2007) procedure, using, as the foreign interest rate, the interest rate on a 3-month government security (e.g., a U.S. T-bill) or an equivalent instrument, taking into account that as data became available, new countries were added (or subtracted) to these portfolios. We cannot reject the null that they price correctly all the dynamic Fama–French portfolios and all but portfolio 1 of Lustig and Verdelhan (2007). These results are all available upon request.

Our claims in this paper depend fundamentally on how strong the unconditional fitting and the out-of-sample Euler-based pricing exercises are. We have compared our results with those results found using a canonical consumption-based pricing kernel calculated with reasonable values for the curvature parameters, according to Mehra and Prescott (1985). At the 5% significance level, we reject the hypothesis that all deviations of the unitary price for the returns are statistically zero individually.

Most of the deviations of the null price for the excess return are also significant, and the Wald test confirms these rejections. The failures in Hansen's moments test are also systematic. The root mean-squared error as prediction performance measure takes on higher values: 3.25% and 3.28% in predicting U.S. and foreign assets, respectively. It is, therefore, *not* true that any, albeit theoretically sound, pricing kernel could be successful in the tests proposed here!

Along the same lines, a crucial element in our analysis was showing that the returns on investments in foreign assets are accounted for by an adequate pricing kernel. This is indeed what we found when the ratio ${}_t F_{t+1}/S_t$ was used as an instrument and the theoretical restrictions tested were not rejected. This suggests that ${}_t F_{t+1}/S_t$ has no predictive power for $M_{t+1}^* \frac{P_t(1+i_{t+1}^*)}{P_{t+1}} \frac{{}_t F_{t+1} - S_{t+1}}{S_t}$, despite having predictive power for $\frac{P_t(1+i_{t+1}^*)}{P_{t+1}} \frac{{}_t F_{t+1} - S_{t+1}}{S_t}$. Although excess returns on uncovered over covered trading with foreign bonds are predictable, "risk adjusted" excess returns are not.

Therefore, notwithstanding the two joint rejections in Wald tests, the out-of-sample test results obtained in Tables 2 and 3 are in sharp contrast with those formerly obtained in log-linear tests of the FPP and/or those obtained with consumption-based kernels. As in any empirical exercise, it is important to assess whether results are robust to changes in the environment used in testing. We changed the conditioning set used previously. Basic results remained unchanged and are available upon request.

5. CONCLUSION

Previous research has cast doubt on whether a single asset-pricing model was capable of correctly pricing the equity and the forward premium, which led to the emergence of two separate literatures. We challenge this position and propose a fresh look into the relationship between the equity and the forward premium

puzzle. Our take on the FPP is that it is not only an international-economics issue, but an asset-pricing problem with important repercussions for monetary economics.¹⁵

We do so by extracting SDF estimates using U.S. return data and showing them to be able to properly price returns and excess returns of assets that compose the equity premium and the forward premium puzzles. We thus establish the common nature of the FPP and the EPP. By using consistent estimates of the SDF mimicking portfolio, we have shown that the projection of a proper CCAPM model—yet to be written—onto the space of returns is able to properly price assets that compose the EPP and the FPP. We have not found this proper model yet. Nevertheless, we rely on the fact that its pricing properties cannot be better than those of its projection onto the space of returns.

By working with a U.S.-based version of SDF estimates, we were able to show that the factors contained in domestic (U.S.) market returns yield neither evidence of the EPP nor evidence of the FPP for most cases. Our starting point is the asset-pricing equation, coupled with the use of consistent estimators of the SDF mimicking portfolio. We first show that our return-based kernels constructed using domestic (U.S.) returns alone price correctly the equity premium, the 90-day T-bill and the Fama-French portfolios. When discounted by our pricing kernels, excess returns are shown to be orthogonal to past information that is usually known to forecast undiscounted excess returns. Based on these results, we go one step further and ask whether the EPP and the FPP are but two symptoms of the same illness—the inability of standard (and augmented) consumption-based pricing kernels to price asset returns or excess -returns. In our tests, we found that the ex ante forward premium is not a predictor of discounted excess returns, despite their being so for undiscounted excess returns.

We believe we have offered evidence that the answer to the question posed in the title of this paper is in the affirmative. A different and interesting issue is whether the forward premium should be regarded as a reward for risk taking. If we take the covariance with M_{t+1}^* as the relevant measure of risk, then our answer is yes. This position is not without controversy, however. Citing Engel (1996, p. 162): “If the [CAPM] model were found to provide a good description of excess returns in foreign exchange markets, there would be some ambiguity about whether these predicted excess returns actually represent premiums.” We do, however, side with the position implicit in Brandt et al. (2006) for whom the behavior of the SDF is equal to that of the intertemporal marginal rate of substitution for a model of preferences and/or market structure yet to be written. In this sense, we have provided grounds to believe that the striking similarity in the results found in trying to apply the CCAPM for the two markets is not accidental.

NOTES

1. We do not claim that returns on equity are not predictable. On the contrary, there is now strong evidence of predictability of returns on equity, accumulated through the last decades, which has caused

the discount factor to take center stage in the discussion of asset pricing—e.g., Cochrane (2008). The point is that predictability was not first seen as a defining feature of the EPP.

2. Pure nonarbitrage arguments are also used by Iwata and Wu (2006) to address foreign exchange risk premia using a VAR system.

3. E.g., Harrison and Kreps (1979), Hansen and Richard (1987), and Hansen and Jagannathan (1991).

4. This view, championed by Frankel (1979), among others, is based on the idea that most exchange rate risks are diversifiable, there being no grounds for agents to be rewarded for holding foreign assets. The relationship between exchange rate shocks and nondiversifiable aggregate risks is, however, considered in Kim and Kose (2003).

5. Similar results were reported later by Modjtahedi (1991). Using a different, larger data set, Hodrick (1989) reported estimated values of $\hat{\alpha}$ above 60, but did not reject the overidentifying restrictions, whereas Engel (1996) reported some estimated $\hat{\alpha}$'s in excess of 100. A more recent attempt to use Euler equations to account for the FPP is Lustig and Verdelhan (2006), where risk aversion in excess of 100 is needed to price the forward premium on portfolios of foreign currency. Finally, using an international real business cycle model, Thoenissen (2011), shows that trade elasticities defined on a very narrow range are needed to account for some of the international markets puzzles, and how these constraints lead to other anomalies.

6. In the case of the FPP, macroeconomic models do not do a good job in explaining exchange rate movements. See Rossi (2006) for a summary of the evidence and a contrasting view.

7. We regress “ones” against returns following Hansen and Richard (1987).

8. By construction, $\mathbf{E}[RM^*] = \mathbf{E}[RR'(\mathbf{E}[RR'])^{-1}1] = \iota$. Moreover, the variance of M^* can serve as a lower bound for any SDF proxy because, as a result of (7), $\text{VAR}(M) = \text{VAR}(M^*) + \text{VAR}(v)$.

9. To account for the cross-section behavior of domestic assets, we also perform in-sample pricing tests for the six Fama and French (1993) benchmark portfolios: dynamic portfolios extracted from the Fama–French library. These results are available upon request.

10. We have also addressed predictability in foreign government bond markets by following Lustig and Verdelhan (2007) in building eight different zero-cost foreign-currency portfolios and applying the orthogonality tests for these portfolios. These results are also available upon request.

11. Suppose that we use N pricing equations with k instruments to form the $N \times k$ orthogonality conditions used in GMM estimation. The variance–covariance matrix of sample moments associated with these orthogonality conditions is of order $(N \times k) \times (N \times k)$ and has $(N \times k) \times (N \times k + 1)/2$ parameters. These must be estimated using T time-series observations. If $N \times k$ is large vis-à-vis T , it may be infeasible to estimate the variance–covariance matrix of sample moments. Even if estimation is feasible, the estimate of this variance–covariance may be far from its asymptotic probability limit, which is a problem for asymptotic tests.

12. Campbell and Shiller (1988), Fama and French (1988), and Cochrane (1991) provide evidence that these variables are good predictors of stock-market returns.

13. Results available upon request.

14. Results not shown here, but available upon request.

15. The uncovered interest parity is key to many influential models in monetary economics, e.g., Mundell (1963).

REFERENCES

- Brandt, Michael W., John H. Cochrane, and Pedro Santa-Clara (2006) International risk sharing is better than you think, or exchange rates are too smooth. *Journal of Monetary Economics* 53, 671–698.
- Campbell, John Y. and Robert J. Shiller (1988) The dividend–price ratio and the expectations of future dividends and discount factors. *Review of Financial Studies* 1, 195–228.
- Cochrane, John H. (1991) Production-based asset pricing and the link between stock returns and economic fluctuations. *Econometrica* 46, 207–234.
- Cochrane, John H. (2001) *Asset Pricing*. Princeton, NJ: Princeton University Press.

- Cochrane, John H. (2008) The dog that did not bark: A defense of return predictability. *Review of Financial Studies* 21(4), 1533–1575.
- Connor, Gregory and Robert Korajczyk (1993) A test for the number of factors in an approximate factor structure. *Journal of Finance* 48, 1263–1291.
- da Costa, Carlos E., Jaime J. Filho, and Paulo F. Matos (2010) Forward-Premium Puzzle: Is It Time to Abandon the Usual Regression? VI Portuguese Financial Network, Azores Island.
- Engel, Charles (1996) The forward discount anomaly and the risk premium: A survey of recent evidence. *Journal of Empirical Finance* 3(2), 123–192.
- Fama, Eugene F. (1984) Term premiums in bond returns. *Journal of Financial Economics* 13(4), 529–546.
- Fama, Eugene F. and Kenneth R. French (1988) Dividend yields and expected stock returns. *Journal of Financial Economics* 22(1), 3–25.
- Fama, Eugene F. and Kenneth R. French (1989) Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25(1), 23–49.
- Fama, Eugene F. and Kenneth R. French (1993) Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33(1), 3–56.
- Frankel, Jeffrey A. (1979) The diversifiability of exchange risk. *Journal of International Economics* 9(3), 379–393.
- Hansen, Lars P. (1982) Large sample properties of generalized method of moments estimators. *Econometrica* 50(5), 1029–1054.
- Hansen, Lars P. and Ravi Jagannathan (1991) Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99(2), 225–262.
- Hansen, Lars P. and Ravi Jagannathan (1997) Assessing specification errors in stochastic discount factor models. *Journal of Finance* 52(2), 557–590.
- Hansen, Lars P. and Scott F. Richard (1987) The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica* 55(3), 587–613.
- Hansen, Lars P. and Kenneth Singleton (1982) Generalized instrumental variables estimation of nonlinear expectations models. *Econometrica* 50(5), 1269–1286.
- Hansen, Lars P. and Kenneth Singleton (1984) Erratum of the article: Generalized instrumental variables estimation of nonlinear expectations models. *Econometrica* 52(1), 267–268.
- Harrison, J. Michael and David M. Kreps (1979) Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20(3), 381–408.
- Hodrick, Robert J. (1989) U.S. international capital flows: Perspectives from rational maximizing models. *Carnegie-Rochester Conference Series on Public Policy* 30(1), 231–288.
- Iwata, Shigeru and Shu Wu (2006) Macroeconomic shocks and the foreign exchange risk premia. *Macroeconomic Dynamics* 10(4), 439–466.
- Kim, Sunghyun Henry and M. Ayhan Kose (2003) Dynamics of open-economy business-cycle models: Role of the discount factor. *Macroeconomic Dynamics* 7(2), 263–290.
- Lustig, Hanno and Adrien Verdelhan (2006) Investing in foreign currency is like betting on your intertemporal marginal rate of substitution. *Journal of the European Economic Association* 4(2–3), 644–655.
- Lustig, Hanno and Adrien Verdelhan (2007) The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review* 97(1), 89–117.
- Mark, Nelson C. (1985) On time varying risk premiums in the foreign exchange market: An econometric analysis. *Journal of Monetary Economics* 16(2), 3–18.
- Mehra, Rajnish and Edward C. Prescott (1985) The equity risk premium: A puzzle. *Journal of Monetary Economics* 15(2), 145–161.
- Modjtahedi, Bagher (1991) Multiple maturities and time-varying risk premia in forward exchange markets: An econometric analysis. *Journal of International Economics* 30(1–2), 69–86.
- Mundell, Robert A. (1963) Capital mobility and stabilization policy under fixed and flexible exchange rates. *Canadian Journal of Economics and Political Science* 29(4), 475–485.

- Newey, Whitney K. and Kenneth D. West (1987) A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55(3), 703–708.
- Phillips, Peter C. B. and Pierre Perron (1988) Testing for a unit root in time series regression. *Biometrika* 75(2), 335–346.
- Rossi, Barbara (2006) Are exchange rates really random walks? Some evidence robust to parameter instability. *Macroeconomic Dynamics* 10(1), 20–38.
- Thoenissen, Christoph (2011) Exchange rate dynamics, asset market structure and the role of the trade elasticity. *Macroeconomic Dynamics* 15(1), 119–143.
- White, Halbert (1980) A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48(4), 817–838.