

NBER WORKING PAPER SERIES

THE FORWARD PREMIUM IS STILL A PUZZLE

Craig Burnside

Working Paper 13129

<http://www.nber.org/papers/w13129>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

May 2007

I thank John Cochrane, Martin Eichenbaum, Ravi Jagannathan, Sergio Rebelo, Michael Weber and an anonymous referee for helpful comments, and the National Science Foundation for financial support (SES-0516697). The usual disclaimer applies. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Craig Burnside. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Forward Premium Is Still a Puzzle
Craig Burnside
NBER Working Paper No. 13129
May 2007, Revised July 2007, Revised June 2011
JEL No. F31,G12

ABSTRACT

Lustig and Verdelhan (2007) argue that the excess returns to borrowing US dollars and lending in foreign currency "compensate US investors for taking on more US consumption growth risk," yet the stochastic discount factor corresponding to their benchmark model is approximately uncorrelated with the returns they study. Hence, one cannot reject the null hypothesis that their model explains none of the cross-sectional variation of the expected returns. Given this finding, and other evidence, I argue that the forward premium puzzle remains a puzzle.

Craig Burnside
Department of Economics
Duke University
213 Social Sciences Building
Durham, NC 27708-0097
and NBER
burnside@econ.duke.edu

Hanno Lustig and Adrien Verdelhan (2007) claim that aggregate consumption growth risk explains the excess returns to borrowing U.S. dollars to finance lending in other currencies. They reach this conclusion after estimating a consumption-based asset pricing model using data on the returns of portfolios of short-term foreign-currency denominated money market securities sorted according to their interest differential with the U.S. Based on their evidence and additional U.S. data, I argue that consumption risk explains none of the cross-sectional variation in the expected returns of their portfolios.

Standard theory predicts that the expected excess return of an asset, $E(R_t^e)$, is given by $-\text{cov}(R_t^e, m_t)$, where m_t denotes some proposed stochastic discount factor (SDF). Therefore, any risk-based explanation of the cross-section of returns relies on significant spread, across portfolios, in the covariance between the returns and the SDF. For the SDFs that Lustig and Verdelhan (henceforth, LV) calibrate and estimate in their 2007 article, it is impossible to reject that there is no spread in these covariances. In fact, it is impossible to reject that these covariances are all zero.

LV's SDF is linear in a vector risk factors, so they implement a widely-used two-pass procedure to estimate its parameters. The first pass is a series of time series regressions of each portfolio's excess return on the risk factors. These regressions determine the factor betas, β . When there are n portfolios and k risk factors, β is an $n \times k$ matrix. In LV's case $n = 8$ and $k = 3$. None of the individual elements of LV's estimate, $\hat{\beta}$, is statistically different from zero. For each of the three factors, we also cannot reject the hypothesis that all eight of the relevant elements of $\hat{\beta}$ are jointly zero. Confronted with this evidence, alone, it would be reasonable to conclude that LV's model does not explain currency portfolios sorted on interest rates.

The statistical insignificance of the factor betas implies that LV's measure of the SDF is also uncorrelated with the excess returns that they study. To demonstrate this, I consider three calibrations of the parameters of the SDF in order to construct time series for m_t : (i) the SDF parameters corresponding to LV's two-pass estimates, (ii) LV's Generalized Method of Moments (GMM, Lars P. Hansen 1982) estimates of the SDF parameters, and (iii) Motohiro Yogo's (2006) estimates of the SDF parameters based on stock returns. I then run a series of time series regressions of each portfolio's excess return on the resulting m_t series. In each case, I find that estimated SDF betas are jointly statistically zero. This, again, suggests that LV's model does not explain the returns to their currency portfolios.

The second pass component of LV's estimation procedure is a cross-sectional regression of average portfolio returns on the betas. This regression determines the lambdas, $\boldsymbol{\lambda}$, a $k \times 1$ vector of factor risk premia. There are two problems with LV's estimates of $\boldsymbol{\lambda}$. First, they focus almost entirely on standard errors for $\hat{\boldsymbol{\lambda}}$ that treat the betas, $\boldsymbol{\beta}$, as known regressors, rather than generated regressors. With these standard errors $\hat{\boldsymbol{\lambda}}$ appears to be statistically significant, so LV draw favorable inference about their model.¹ But treating $\boldsymbol{\beta}$ as known leads to a misleading level of confidence in the model. With conventionally calculated standard errors (Shanken, GMM) none of the estimated factor risk premia in LV's benchmark model, and none of the parameters of the corresponding estimated SDF are statistically significant, except in cases where the model has very poor fit. Bootstrapped 95 percent confidence regions for these parameters always encompass zero. Consequently, I draw unfavorable inference where they do not.

Second, for $\boldsymbol{\lambda}$ to be identified, $\boldsymbol{\beta}$ must have full column rank. Because most of the elements of $\hat{\boldsymbol{\beta}}$ are statistically close to zero, statistical tests indicate that the rank of $\boldsymbol{\beta}$ is very low, perhaps as low as 0. The identification problem raises two important issues. First, and most importantly, it weakens inference in the sense that tests of the pricing errors based on the second pass regressions have little power to reject misspecified models (Raymond Kan and Chu Zhang, 1999a; Craig Burnside, 2010). Second, confidence regions for estimates of the factor risk premia, $\boldsymbol{\lambda}$, generated using asymptotic standard errors become unreliable (Kan and Zhang, 1999b). Using methods that are robust to weak identification, I show that LV's data contain almost no information about $\boldsymbol{\lambda}$. This reinforces the unfavorable inference I draw regarding their model.

In their reply, LV defend their findings on four main grounds. First, they discard most of my comment as an obscure discussion of sampling uncertainty as opposed to point estimates. It is true that I do not dispute their point estimates; this comment is not a trivial report on errors in LV's code for ordinary least squares (OLS). Unfortunately, however, until data sets are infinitely large, inference will involve both point estimates and standard errors. In their original article, LV clearly recognize the importance of statistical significance for inference. They repeatedly refer to the statistical significance of their estimates and to the results of

¹There are seven tables of parameter estimates in the original article. The standard errors computed in the first six tables treat the betas as known. The standard errors computed in the last table effectively treat the betas as unknown, but they cannot be compared to the standard errors in the rest of the article, because they are calculated for a different model (one without a constant). They are also calculated incorrectly, as I explain below.

statistical tests. Once inference is conducted properly, however, there is little support for LV's model.

Second, they appeal to a robustness check, described in their article, in which additional tests assets (six equity portfolios and five bond portfolios) are included in the model estimation. The inclusion of these test assets, however, has little effect on my conclusions. One still cannot reject the null hypothesis that the covariances between the excess returns of LV's currency portfolios and the SDF are all zero. Thus, regardless of the statistical significance of the parameters that determine the factor loadings in the SDF, I cannot reject that the model predicts that $E(R_t^e) = 0$ for all of the currency portfolios. Including more test assets leads to modest improvement on the identification front, largely because the equity portfolios are correlated with one of the model's risk factors: the return to the aggregate US stock market. Not surprisingly, this leads to some estimates of the model parameters being statistically significant. However, the model still does not explain the cross-section of foreign currency risk premia. The R^2 for the currency portfolios alone is at best roughly zero, indicating that the model cannot explain why some currency portfolios have significantly positive returns while others have significantly negative returns.

Third, LV refer to empirical evidence not in their original article. As in their more recent paper with Nick Roussanov (2009), they construct a new set of seven portfolios from their original set of eight portfolios by considering strategies whereby the investor short sells the low interest rate portfolio while going long in one of the seven higher interest rate portfolios. Evidence regarding the seven "differenced" portfolios affects none of my conclusions. Most importantly, since these portfolios are linear combinations of the original portfolios one cannot reject the null hypothesis that the covariances between the excess returns of the "differenced" portfolios and the SDF are all zero, and, therefore, one cannot reject the null hypothesis that the model predicts $E(R_t^e) = 0$ for all of the "differenced" portfolios. Also, because the "differenced" portfolios are smaller in number, and are formed as linear combinations of the original portfolios, working with these portfolios can only make the identification problem worse.

Finally, LV bring to bear additional evidence based on the recent financial crisis. They argue that the financial crisis, alone, is sufficient evidence that their consumption-based model works. Indeed, the financial crisis is a single observation that suits their hypothesis. Consumption growth fell, and currency returns were negative, in late 2008. However, I show

that carry trade returns and consumption growth are uncorrelated over the full post-Bretton Woods period. Carry trade returns are correlated with stock returns in the post-Bretton Woods period, but the market beta of the carry trade is far too small to explain its high average return. I also establish that there is only a very weak tendency of the market beta of carry trade returns to increase during US recessions, and periods of stock market turmoil. This casts doubt on any simple explanation of the returns to the carry trade based on market risk.

I conclude that, taken as a whole, the evidence for LV’s consumption-based model is extremely weak. I cannot reject that the model-predicted expected returns of the currency portfolios they study are all zero. In their reply, LV conclude by arguing that had the researchers of 25 years ago been confronted with their results, there would never have been a “forward premium puzzle.” I am alive now, and I have read their article. The forward premium is still a puzzle.

In Section 1 I briefly review LV’s model, data and methodological approach. In Section 2, I present the first-pass estimates of the betas that underlie their estimates of the factor risk premia and demonstrate that there is little evidence of significant covariance between the portfolio returns and the risk factors. In Section 3, I discuss the second-pass estimates of the factor risk premia and the interpretation of the pricing errors, and calculate standard errors for factor risk premia that correctly account for estimation of the betas. I discuss robustness of my negative findings in Section 4. Section 5 concludes.

1 Model, Data, Estimation and Inference

LV work primarily with a log-linearized version of Motohiro Yogo’s (2006) model, in which the stochastic discount factor is given by

$$m_t = \xi[1 - b_c(\Delta c_t - \mu_c) - b_d(\Delta d_t - \mu_d) - b_r(r_{Wt} - \mu_r)]. \quad (1)$$

Here c_t represents the logarithm of a representative household’s consumption of nondurable goods, d_t is the logarithm of the household’s durable consumption, r_{Wt} is the logarithm of the gross aggregate return to wealth, $\mu_c = E(\Delta c_t)$, $\mu_d = E(\Delta d_t)$ and $\mu_r = E(r_{Wt})$.

LV study the returns to borrowing U.S. dollars in the money market to finance short-term securities denominated in foreign currency. They form eight portfolios of such positions, which are created by sorting the currencies according to their interest differential versus the

U.S. I refer to these portfolios as P1, P2, \dots , P8 with the order running from low interest rate currencies to high interest rate currencies.²

LV estimate the model by exploiting the null hypothesis that the approximated stochastic discount factor (SDF), m_t , prices the $n \times 1$ vector of portfolio excess returns, \mathbf{R}_t^e . The pricing equation is

$$E(\mathbf{R}_t^e m_t) = 0. \quad (2)$$

I rewrite (1) generically as

$$m_t = \xi[1 - (\mathbf{f}_t - \boldsymbol{\mu})' \mathbf{b}], \quad (3)$$

where \mathbf{f}_t is a $k \times 1$ vector of risk factors, $\boldsymbol{\mu} = E(\mathbf{f}_t)$, \mathbf{b} is a $k \times 1$ vector of coefficients, and ξ is a scalar representing the mean of the SDF.

1.1 The Beta Representation and Two-Pass Regressions

It follows from (3) and (2) that

$$E(\mathbf{R}_t^e) = \text{cov}(\mathbf{R}_t^e, \mathbf{f}_t) \mathbf{b} = \underbrace{\text{cov}(\mathbf{R}_t^e, \mathbf{f}_t)}_{\boldsymbol{\beta}} \underbrace{\boldsymbol{\Sigma}_f^{-1} \boldsymbol{\Sigma}_f \mathbf{b}}_{\boldsymbol{\lambda}}. \quad (4)$$

where $\boldsymbol{\beta}$ is a $n \times k$ matrix of factor betas, $\boldsymbol{\lambda}$ is a $k \times 1$ vector of factor risk premia, and $\boldsymbol{\Sigma}_f$ is the covariance matrix of \mathbf{f}_t .

LV estimate $\boldsymbol{\beta}$ and $\boldsymbol{\lambda}$ using a two-pass procedure associated with Eugene Fama and James D. MacBeth (1973). The first pass is a time series regression of each portfolio's excess return on the vector of risk factors:

$$eR_{it}^e = a_i + \mathbf{f}_t' \boldsymbol{\beta}_i + \epsilon_{it}, \quad t = 1, \dots, T, \text{ for each } i = 1, \dots, n. \quad (5)$$

Here $\boldsymbol{\beta}_i'$ represents the i th row in $\boldsymbol{\beta}$. LV estimate the system of equations represented by (5) using equation-by-equation OLS. Given (4), the second pass is a cross-sectional regression of average portfolio returns on the estimated betas:

$$\bar{R}_i^e = \hat{\boldsymbol{\beta}}_i' \boldsymbol{\lambda} + \alpha_i, \quad i = 1, \dots, n, \quad (6)$$

where $\bar{R}_i^e = \frac{1}{T} \sum_{t=1}^T R_{it}^e$, $\hat{\boldsymbol{\beta}}_i$ is the OLS estimate of $\boldsymbol{\beta}_i$ obtained in the first stage, and α_i is a pricing error. Let the OLS estimator of $\boldsymbol{\lambda}$ be $\hat{\boldsymbol{\lambda}} = (\hat{\boldsymbol{\beta}}' \hat{\boldsymbol{\beta}})^{-1} \hat{\boldsymbol{\beta}}' \bar{\mathbf{R}}^e$, where $\bar{\mathbf{R}}^e$ is an $n \times 1$ vector

²Further details of the model, portfolio formation, and data sources can be found in LV's article.

formed from the individual mean returns. The model's *predicted mean returns* are $\hat{\beta}\hat{\lambda}$ and the *pricing errors* are the residuals, $\hat{\alpha} = \bar{\mathbf{R}}^e - \hat{\beta}\hat{\lambda}$.

The model's fit is assessed using the following statistic:

$$R^2 = 1 - \frac{(\bar{\mathbf{R}}^e - \hat{\beta}\hat{\lambda})'(\bar{\mathbf{R}}^e - \hat{\beta}\hat{\lambda})}{(\bar{\mathbf{R}}^e - \bar{R}^e)'(\bar{\mathbf{R}}^e - \bar{R}^e)}, \quad (7)$$

where $\bar{R}^e = \frac{1}{n} \sum_{i=1}^n \bar{R}_i^e$ is the cross-sectional average of the mean returns in the data.

The model is tested on the basis of the estimated pricing errors using the statistic $C_{\hat{\alpha}} = T\hat{\alpha}'\hat{\Omega}_{\hat{\alpha}}^{-1}\hat{\alpha}$, where $\hat{\Omega}_{\hat{\alpha}}$ is a consistent estimator for the asymptotic covariance matrix of $\sqrt{T}\hat{\alpha}$ and the inverse is generalized. John H. Cochrane (2005) discusses how to form $\hat{\Omega}_{\hat{\alpha}}$ and shows that $C_{\hat{\alpha}} \xrightarrow{d} \chi_{n-k}^2$.

It is common to include a constant in the second-pass regression as follows:

$$\bar{R}_i^e = \gamma + \hat{\beta}'_i \boldsymbol{\lambda} + u_i, \quad i = 1, \dots, n. \quad (8)$$

The constant, γ , is often interpreted as the model's pricing error for the risk free rate, but this error is shared by all assets. The statistical argument for running the regression without the constant is that we know with certainty that the excess return to a risk free asset, or any other zero-beta asset, is zero. One argument for including the constant is the notion that the risk free rate is imperfectly measured as the real return on T-bills.

1.2 GMM Estimation

Cochrane (2005) describes a GMM procedure that produces the same point estimates as the two pass regression method, but allows for heteroskedasticity-robust inference. When the constant is included in the model the moment restrictions are

$$E(R_{it}^e - a_i - \boldsymbol{\beta}'_i \mathbf{f}_t) = 0, \quad i = 1, \dots, n. \quad (9)$$

$$E[(R_{it}^e - a_i - \boldsymbol{\beta}'_i \mathbf{f}_t) \mathbf{f}'_t] = 0, \quad i = 1, \dots, n. \quad (10)$$

$$E(R_{it}^e - \gamma - \boldsymbol{\beta}'_i \boldsymbol{\lambda}) = 0, \quad i = 1, \dots, n. \quad (11)$$

When the constant is excluded from the model, the last set of moment restrictions is replaced by

$$E(R_{it}^e - \boldsymbol{\beta}'_i \boldsymbol{\lambda}) = 0, \quad i = 1, \dots, n. \quad (12)$$

In both cases, an identity matrix is used to weight the moment conditions.

The model can also be estimated using a GMM procedure that treats the SDF as the primary object of interest. This procedure, described in more detail in Cochrane (2005), estimates the model, (3), using the moment conditions:

$$E\{\mathbf{R}_t^e[1 - (\mathbf{f}_t - \boldsymbol{\mu})' \mathbf{b}]\} = 0 \quad (13)$$

$$E(\mathbf{f}_t - \boldsymbol{\mu}) = 0 \quad (14)$$

The parameter ξ is unidentified and is set equal to 1. The moment condition (13) can also be modified to allow for a common pricing error across assets:

$$E\{\mathbf{R}_t^e[1 - (\mathbf{f}_t - \boldsymbol{\mu})' \mathbf{b}] - \gamma\} = 0. \quad (15)$$

As described in the online appendix, the GMM procedure based on (14) and (15) can be set up so that it is numerically identical to the two-pass regression method in terms of pricing errors.³

2 First-Pass Estimates of Betas

Like LV, I compute first-pass estimates of the betas by running the least squares regressions described by (5). I compute standard errors using standard system OLS formulas, as well as GMM-based procedures. I also calculate 95 percent confidence regions using a bootstrap procedure. Using any of these procedures, none of the 24 estimated betas is *individually* statistically significant at the 5 percent level.⁴ More importantly, when there is spread in the expected returns across portfolios, there should also be statistically significant spread in the betas across portfolios. With this in mind, we can test whether for each *factor* the eight factor betas are jointly significantly different from zero. As Table 1(a) indicates, at conventional significance levels one cannot reject the hypotheses that $\beta_{ij} = \beta_j \forall i$, and $\beta_{ij} = 0 \forall i$, for each factor $j = 1, \dots, k$. Since the contribution of factor j to the vector of model-predicted expected returns is $\boldsymbol{\beta}_j \lambda_j$ the latter hypothesis tests imply that one cannot reject the null that each factor's risk premium contributes nothing to the model-predicted expected returns.

³If an estimate of $\boldsymbol{\lambda}$ is computed as $\hat{\boldsymbol{\lambda}} = \hat{\boldsymbol{\Sigma}}_f \hat{\mathbf{b}}$, where $\hat{\boldsymbol{\Sigma}}_f$ sample covariance matrix of \mathbf{f}_t , this estimate is identical to the two-pass estimate of $\boldsymbol{\lambda}$. The equivalence of the GMM and two-pass procedures is demonstrated in the online appendix.

⁴Due to space limitations I report full tables of betas in the online appendix. The GMM-based standard errors I present are computed using a variant of the VARHAC procedure described by Wouter J. den Haan and Andrew T. Levin (2000). I use VARHAC standard errors to take into account possible serial correlation in GMM errors.

In their reply, LV mistakenly argue that I have only looked at individual betas, when, in fact, in every version of my comment, including this one, I have tested for spread in the betas. Second, they argue that I should have looked at univariate betas rather than multivariate betas. This is puzzling, given that multivariate betas are what appears in the beta representation, (4), and what enters into the second-pass regression. Certainly, one can define the matrix of univariate betas, $\boldsymbol{\beta}^u = \text{cov}(\mathbf{R}_t^e, \mathbf{f}_t) \mathbf{D}_f^{-1}$, where \mathbf{D}_f is a matrix with the variances of the factors on the diagonal, and zeros off-diagonal. This leads to an alternative beta representation in which $E(\mathbf{R}_t^e) = \boldsymbol{\beta}^u \boldsymbol{\lambda}^u$ and $\boldsymbol{\lambda}^u = \mathbf{D}_f \mathbf{b}$. It is perhaps more straightforward, however, to work with the SDF representation $E(\mathbf{R}_t^e) = \text{cov}(\mathbf{R}_t^e, \mathbf{f}_t) \mathbf{b}$. The columns of $\text{cov}(\mathbf{R}_t^e, \mathbf{f}_t)$ are proportional to the columns of $\boldsymbol{\beta}^u$. We can directly estimate the elements of $\text{cov}(\mathbf{R}_t^e, \mathbf{f}_t)$ by GMM and then test whether $\text{cov}(\mathbf{R}_{it}^e, \mathbf{f}_{jt}) = c_j \forall i$, and $\text{cov}(\mathbf{R}_{it}^e, \mathbf{f}_{jt}) = 0 \forall i$. As Table 1(b) indicates, these hypotheses cannot be rejected at conventional significance levels.

It makes little difference which betas we use, however, because what matters in the end is how these betas are reflected in the covariance between the SDF and the portfolio returns. Using (3), (2) can be rewritten as

$$E(\mathbf{R}_t^e) = -\text{cov}(\mathbf{R}_t^e, m_t) / E(m_t). \quad (16)$$

With the normalization $\xi = 1$, (3) implies that $E(m_t) = 1$, so we can rewrite (16) as

$$E(\mathbf{R}_t^e) = -\text{cov}(\mathbf{R}_t^e, m_t) = \boldsymbol{\beta}_m \lambda_m, \quad (17)$$

where $\boldsymbol{\beta}_m \equiv -\text{cov}(\mathbf{R}_t^e, m_t) / \sigma_m^2$, $\lambda_m \equiv \sigma_m^2$ and σ_m^2 is the variance of m_t . Given the definition of m_t ,

$$\boldsymbol{\beta}_m = \frac{\text{cov}(\mathbf{R}_t^e, \mathbf{f}_t) \mathbf{b}}{\sigma_m^2} = \frac{\boldsymbol{\beta} \boldsymbol{\lambda}}{\sigma_m^2} = \frac{\boldsymbol{\beta}^u \boldsymbol{\lambda}^u}{\sigma_m^2}. \quad (18)$$

Which set of factor betas we consider is immaterial, since the SDF betas, $\boldsymbol{\beta}_m$, derived from either are identical, and what matters is whether there is sufficient spread in the SDF betas.

To measure the SDF betas, we need data for m_t , which can be constructed using (3), and values for the elements of the vector \mathbf{b} . Here I use three versions of \mathbf{b} taken directly from LV's article.

In Table 2(i) I use the \mathbf{b} vector corresponding to LV's two pass estimates of $\boldsymbol{\lambda}$: $b_c = -21$, $b_d = 130$ and $b_r = 4.5$. Table 2(ii) uses the \mathbf{b} vector corresponding to LV's GMM estimates of \mathbf{b} : $b_c = 37$, $b_d = 75$ and $b_r = 4.7$. Table 2(iii) repeats the exercise using the \mathbf{b} vector corresponding to the calibrated model discussed in section I.E of LV's article: $b_c = 6.7$,

$b_d = 23$ and $b_r = 0.31$. As Table 2 indicates, in all of these cases the null hypotheses that $\beta_{im} = \beta_m$ for all i and $\beta_{im} = 0$ for all i cannot be rejected. In other words, there is no spread in the betas, and they are jointly zero. Tests based directly on $\text{cov}(\mathbf{R}_t^e, m_t)$ rather than β_m reach the same conclusion but are not reported in the table.

Given that the SDF betas are jointly statistically insignificant, I conclude that LV’s model does not explain the cross-section of the expected returns of their portfolios. In Section 4 I show that this finding is robust to (i) estimates of the SDF based on an expanded set of test assets including equities and bonds, (ii) using the seven “differenced” portfolios emphasized in their reply, and (iii) a higher frequency, post-Bretton Woods, developed-country, database that extends through the recent financial crisis. Thus, the message of this comment is immune to the points emphasized by LV in their reply.

3 Second Pass and GMM Estimates of the Model

The second pass and GMM estimates of the model provide us another opportunity to assess LV’s proposed explanation of the cross-section of returns to foreign currency portfolios. Of particular interest are the point estimates of $\boldsymbol{\lambda}$ and \mathbf{b} , the R^2 measure of fit and the tests of the pricing errors.

LV’s second pass regression, which includes the constant γ , is reproduced in Table 3. When presenting their findings they show OLS standard errors, which assume that the first-pass betas are known. Given these standard errors, the factor risk premia for consumption and durables are both positive and highly statistically significant. The R^2 of the model is 0.87 and the p-value for the test for significance of the pricing errors is 0.48. These results are a key basis of LV’s positive assessment of the model.

There are three reasons our assessment should be less sanguine. The main one is that OLS standard errors are inappropriate given the estimation of the betas, and this turns out to matter a great deal for inference. Once standard errors are computed appropriately, estimates of $\boldsymbol{\lambda}$ and \mathbf{b} are statistically insignificant. The latter finding is especially important because it suggests that the consumption factors do not help price currency returns. The second reason to be skeptical is that the model performs much more poorly when we impose the restriction that the constant is equal to zero. The \mathbf{b} parameters remain insignificant, and the fit of the model deteriorates substantially. The third reason to be doubtful about the model estimates is that there is a severe identification problem. Under non-identification or

weak identification, asymptotic standard errors (even arguably appropriate ones) are likely to understate the degree of uncertainty about the model parameters.

3.1 Inference About Model Parameters

As Cochrane (2005) points out, the fact that the betas are estimated in the first pass matters for inference about the factor risk premia, and this remains true asymptotically. There are three standard ways to deal with this problem. One is to use the correction of the standard errors suggested by Jay Shanken (1992). Another is to compute standard errors using the first of the two GMM procedures described above, because it produces the same point estimates. A third is to construct confidence regions for the parameters using bootstrap methods. By construction, the alternative approaches to calculating standard errors do not affect the point estimates of the factor risk premia. The three procedures lead to similar inference regarding the model, and using them, rather than OLS standard errors, matters both qualitatively and quantitatively.

The Shanken and GMM-corrected standard errors for the model with the constant (Table 3) are roughly two to three times larger than the OLS standard errors that ignore estimation of the betas. Why is the Shanken correction so big? Let $\boldsymbol{\theta} = (\gamma \quad \boldsymbol{\lambda}')'$, $\boldsymbol{\Sigma} = E(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t')$ and let $\tilde{\boldsymbol{\Sigma}}_{\mathbf{f}}$ be a matrix with a leading column and row of zeros, and $\boldsymbol{\Sigma}_{\mathbf{f}}$ in the lower right corner. When the betas are treated as known the covariance matrix of $\sqrt{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$ is

$$\boldsymbol{\Omega}_{\hat{\boldsymbol{\theta}}} = (\boldsymbol{\beta}^{+'} \boldsymbol{\beta}^+)^{-1} \boldsymbol{\beta}^{+'} \boldsymbol{\Sigma} \boldsymbol{\beta}^+ (\boldsymbol{\beta}^{+'} \boldsymbol{\beta}^+)^{-1} + \tilde{\boldsymbol{\Sigma}}_{\mathbf{f}}. \quad (19)$$

Here $\boldsymbol{\beta}^+ = (\boldsymbol{\iota} \quad \boldsymbol{\beta})$ and $\boldsymbol{\iota}$ is an $n \times 1$ vector of ones. With the Shanken correction the covariance matrix is

$$\boldsymbol{\Omega}_{\hat{\boldsymbol{\theta}}} = (1 + \boldsymbol{\lambda}' \boldsymbol{\Sigma}_{\mathbf{f}}^{-1} \boldsymbol{\lambda}) (\boldsymbol{\beta}^{+'} \boldsymbol{\beta}^+)^{-1} \boldsymbol{\beta}^{+'} \boldsymbol{\Sigma} \boldsymbol{\beta}^+ (\boldsymbol{\beta}^{+'} \boldsymbol{\beta}^+)^{-1} + \tilde{\boldsymbol{\Sigma}}_{\mathbf{f}}. \quad (20)$$

In some finance applications the Shanken correction is small. For example, for the CAPM estimated using the annual returns of Fama and French's (2003) 25 portfolios sorted on size and book-to-market value over the period 1953–2002, the Shanken-correction term, $1 + \lambda^2 / \sigma_f^2$, is estimated to be 1.03. In LV's case the estimate of $1 + \boldsymbol{\lambda}' \boldsymbol{\Sigma}_{\mathbf{f}}^{-1} \boldsymbol{\lambda}$ is 6.79. Although the individual λ s in LV's model are of the same order of magnitude as for the CAPM, the consumption factors have much smaller variance than the market return. This blows up the size of the Shanken correction substantially.

Using either the Shanken or GMM standard errors, none of the estimated factor risk premia in Table 3 are statistically significant at the 5 percent level. The bootstrap-based 95 percent confidence regions for the parameters also encompass 0. These results do not imply that the price of consumption risk is zero. Instead, they indicate that the joint behavior of the currency returns and consumption factors is uninformative about the price of consumption risk.

LV defend the statistical significance of their findings on three grounds. First, they appeal to Ravi Jagannathan and Zhenyu Wang (1998) to defend the use of OLS standard errors rather than the Shanken correction. This is inappropriate. Jagannathan and Wang’s point is that under heteroskedasticity, the Shanken correction is inappropriate, and that more general GMM errors are appropriate. Shanken’s proof that corrected standard errors are necessarily bigger than OLS standard errors does not work for GMM standard errors. GMM errors *could* be smaller than OLS standard errors, but in LV’s case they are not. Second, in their footnote 11, they argue that the OLS standard errors are close in magnitude to the GMM standard errors. This is because they inappropriately compare GMM standard errors for the model without a constant, to the OLS standard errors for a model with a constant.⁵ The appropriate GMM standard errors are actually close in magnitude to the Shanken standard errors. Third, they argue that standard errors from a bootstrap procedure are small enough to make the estimated risk premia significant. Rather than focus on bootstrap standard errors, I use the entire distribution of bootstrapped estimates to show that 95 percent confidence regions easily encompass 0.

It is especially important to know whether the consumption factors help to price the currency returns. This requires us to focus on the parameter vector \mathbf{b} . GMM estimates of \mathbf{b} for the model with the constant are found in Table 4. The first stage of GMM is equivalent to the second-pass regression in terms of point estimates. I also show GMM estimates from the second stage, and after iterating over the weighting matrix to convergence. At all stages of GMM, the estimates of \mathbf{b} are statistically insignificant for every risk factor. Bootstrapped confidence regions also easily encompass zero. Thus we cannot reject the null hypothesis that consumption factors do not help price the cross-section of currency returns. The statistical insignificance of the estimates of $\boldsymbol{\lambda}$ is also robust to using this GMM procedure to estimate the model (Table 5).

⁵When LV report Shanken and HAC standard errors for the model without the constant, they appear to use incorrect formulas, as detailed in the online appendix.

3.2 Imposing a Zero Intercept

Given the estimates in Table 3 the R^2 of the model is 0.87, the cross-sectional mean absolute pricing error (MAE) is 0.44 percent, and the test of the pricing errors fails to reject the model regardless of how standard errors are calculated. The model's high R^2 is attributable to the inclusion of the common pricing error parameter $\hat{\gamma}$ that, by convention, is treated as part of the model's predicted expected returns. The theoretical model does not include a constant and predicts that the expected returns depend only on the covariance between the factors and the returns. So, whenever a constant is included in the second-pass regression it is important to consider its economic and statistical significance. In fact, as Table 3 indicates, the constant is big, implying a -3 percent per annum pricing error for the risk free rate. Measurement error in the estimated betas, and resulting downward bias in the estimated factor risk premia, can explain a positive pricing error for the risk free rate. So can a liquidity premium in T-bills. But a large *negative* pricing error for the risk free rate is bad news for the model. Like the other model parameters, however, the constant is statistically insignificant when standard errors are calculated appropriately.

Table 6 presents two-pass estimates of the model obtained by imposing the restriction that the constant is zero. In this case the factor risk premia for the consumption factors are much smaller, none of factor risk premia are statistically significant, the R^2 is only 0.34, and the mean absolute pricing error of the model is 1.17 percent.

Tables 7 and 8 present GMM estimates of the model without the constant. In the first stage of GMM, none of the \mathbf{b} coefficients is individually significant, nor are any of the elements of $\boldsymbol{\lambda}$. The R^2 of the model is 0.34 and the MAE is 1.17 percent. Turning to the second stage of GMM, I reproduce LV's point estimates (their Table 14, Column C). Importantly, the R^2 of the model is now negative (-0.66), and the MAE is 1.86 percent. This is very bad news, because it indicates that a constant would do a better job explaining the cross-sectional distribution of the returns than the model does. While the individual λ s associated with consumption and durables are statistically significant, the \mathbf{b} parameter is only significant for durables, and none of the parameters appears to be significant based on the bootstrapped 95 percent confidence regions. Further iterations on the weighting matrix lead to a further deterioration of the model's performance. The R^2 of the model drops to -1.45 and the MAE rises to 2.28 percent. The factor risk premium for durables is statistically significant, but none of the \mathbf{b} parameters are significant, and none of the parameters appears to be significant

based on the bootstrapped 95 percent confidence regions. In summary, the model fits very poorly absent the constant.

3.3 Weak Identification

In this section I argue that we should be skeptical of LV's model because it is very weakly identified as a consequence of the sampling uncertainty associated with $\hat{\beta}$. In the second-pass regression with the constant, the parameters γ and λ are identified under the assumption that β^+ has full column rank. In the second-pass regression with no constant, λ is identified if β has full column rank. When the rank conditions fail, conventional inference drawn from second pass regressions and GMM is unreliable because standard asymptotic theory does not apply. As Burnside (2010) discusses, t -statistics for $\hat{\lambda}$ and $\hat{\mathbf{b}}$ have non-standard distributions, and, most importantly, pricing-error tests cannot reliably detect model misspecification.

The online appendix contains an extended discussion of the identification problem. In it I show that it is straightforward to test whether LV's model is identified using a rank test from Jonathan H. Wright (2003). I show that it is not possible to reject that $\text{rank}(\beta) = 1$ and $\text{rank}(\beta^+) = 2$ at conventional significance levels. This implies that both the model without the constant, which needs $\text{rank}(\beta) = 3$, and the model with the constant, which needs $\text{rank}(\beta^+) = 4$, are grossly underidentified.

A standard tool for conducting inference under weak identification is to construct confidence sets for the weakly identified parameters using the objective function corresponding to the continuously updated (CU) GMM estimator. In the online appendix, I construct CU-GMM based confidence sets for the price of SDF risk, λ_m . The robust 95 percent confidence set for λ_m encompasses almost the entire real line. Thus currency portfolios are virtually uninformative about the price of risk associated with LV's SDF.

4 Robustness

In this section I consider the robustness of my findings to three considerations emphasized by LV in their reply. First, I ask whether adding information from equity and bond portfolios improves identification and sheds additional light on whether the SDF prices the currency portfolios. Second, I check whether considering a set of "differenced" currency portfolios affects my conclusions. Third, I discuss additional evidence from higher frequency data in the post-Bretton Woods period, including the relevance of the recent global financial crisis.

4.1 Equity and Bond Portfolios

LV, in their reply, argue that the results in Section IV.C of their original article are important, because these show that the same SDF that prices currency portfolios also prices other test assets, such as Fama and French's (1993) six equity portfolios created by sorting stocks on the basis of size and value, and five Fama bond portfolios (Center for Research in Security Prices, 2007) created by sorting bonds on the basis of maturity. Indeed, additional test assets could be useful because they might alleviate the identification problem alluded to in Section 3.3. However, I find that these additional test assets have no effect on my main conclusion, that currency portfolios do not appear to be priced by consumption factors.

The main reason that adding additional portfolios does not change my conclusions is that it does not change the factor betas for the currency portfolios. The tests in Section 2 still apply. So, adding the equity and bond portfolios does not change my conclusion, based on betas alone, that the model cannot explain differences in expected returns across currency portfolios.

Adding the six Fama-French equity portfolios to the set of test assets *slightly* alleviates the identification problem because equities have statistically significant betas with respect to the market return factor, r_W . However, the rank tests presented in the online appendix indicate that the β matrix still appears to have reduced rank. The identification problem does not go away with the further addition of the five Fama bond portfolios.

Estimates of the model without the constant using the currency and equity portfolios as test assets are presented in the online appendix. As indicated there, with sufficient iterations over the weighting matrix the factor risk premia for consumption growth and durables growth are statistically significant. However, the fit of the model with respect to currency portfolios is very poor. When the R^2 statistic is calculated just for currency portfolios it ranges from 0.03 (at the first stage of GMM) to -0.76 (for iterated GMM). The mean absolute pricing error for the currency portfolios ranges from 1.44 (at the first stage of GMM) to 1.88 (for iterated GMM).

Why do I compute these statistics just for currency portfolios? First, the goal is to explain the cross-section of returns of currency portfolios. The R^2 across all assets does not tell us whether the model explains why some currency portfolios earn high returns and others earn low returns. Instead, the R^2 across just the currency portfolios tell us whether the model explains why some currency portfolios (like P7) earn high returns, and other

currency portfolios (like P1) earn low returns, on average. Second, we are not trying to explain why the currency portfolios all have relatively low returns (the average excess return across LV’s eight currency portfolios is 0.1 percent) compared to the equity portfolios (the average excess returns of the six Fama-French portfolios are all above 6 percent, and they are centered near 9 percent). That is not a puzzle, given that currency portfolios are only weakly correlated with risk factors that price equity portfolios. The puzzle is why the low interest rate currency portfolio, P1, has an average return of -2.3 percent, and why the high interest rate currency portfolios, P7 and P8, have average returns in excess of 2 percent.

Adding the five Fama bond portfolios does not improve the situation. At the first two stages of GMM the results are quite similar to those obtained using only the currency and equity portfolios, although further iterations over the weighting matrix eventually drive out consumption growth and durables growth as significant risk factors. Once again, the fit of the model with respect to currency portfolios is very poor. If the R^2 statistic is calculated just for currency portfolios it ranges from 0.03 (at the first stage of GMM) to -1.35 (for iterated GMM). The mean absolute pricing error for the currency portfolios ranges from 1.40 (at the first stage of GMM) to 1.64 (for iterated GMM).

4.2 Differenced Currency Portfolios

In their reply, LV argue that they can explain the excess returns to the strategy of holding P_i and shorting P1, for $i = 2, \dots, 8$. They claim that their article is really about these seven “differenced” portfolios, which I refer to as D2, D3, \dots , D8, and not really about the original eight portfolios. Given that the entire article is about the P-portfolios this comment is surprising, especially given the following statement in the original article: “consumption-based models can explain the cross-section of currency excess returns if and only if high interest rate currencies typically depreciate when real US consumption growth is low, while low interest rate currencies appreciate”. Notice that these statements are not about whether there is a difference between the rates of return of the portfolios, it is a statement about the rates of return themselves.

LV also argue for working with the D-portfolios on the basis that “large swings in the dollar make it hard to accurately estimate the constant”, the constant being γ in the model for the P-portfolios. This argument is not persuasive because the intercept can be “estimated” with perfect accuracy. We *know* that the mean excess return of a zero beta asset is

zero, so we can set γ equal to zero without even having to estimate it.

Nonetheless, what of LV's point that the constant no longer plays an important role in explaining the cross-section once we consider the seven D portfolios? LV are right, but this point can easily be made without new tables of point estimates. Consider the model with the constant. The estimates of the second pass regression satisfy:

$$\bar{R}_i^e = \hat{\gamma} + \hat{\beta}'_i \hat{\lambda} + \hat{u}_i, \quad i = 1, \dots, n. \quad (21)$$

This is equation (8) with γ and λ replaced by $\hat{\gamma}$ and $\hat{\lambda}$ (the two-pass estimates) and u_i replaced by \hat{u}_i (the idiosyncratic pricing error or residual). Now suppose we consider the new set of excess returns, $R_i^d = R_i^e - R_1^e$, for $i \geq 2$. Given the definition of R_i^d and equation (21), it follows that the sample mean of R_i^d is given by:

$$\bar{R}_i^d = (\hat{\beta}_i - \beta_1)' \hat{\lambda} + \hat{u}_i - \hat{u}_1, \quad i = 2, \dots, n. \quad (22)$$

P1 is not just any asset. It happens to be an asset for which the SDF beta is roughly zero ($\beta_1' \hat{\lambda} \cong 0$) and the idiosyncratic pricing error is very small ($\hat{u}_1 \cong 0$). Using these facts in equation (22), we have

$$\bar{R}_i^d \cong \hat{\beta}'_i \hat{\lambda} + \hat{u}_i, \quad i = 2, \dots, n. \quad (23)$$

Now, since $\hat{u}_1 \cong 0$ it also follows that the average value of \hat{u}_i across $i = 2, \dots, n$, is roughly zero. Thus, the same $\hat{\lambda}$ obtained by running the cross-sectional regression for the P-portfolios, can approximately fit the cross-sectional distribution of \bar{R}_i^d without a constant.

In their reply, it seems that LV concede that their model does not price the original portfolios. But this means they have not identified the true SDF. All of the portfolios (the P and D-portfolios) should be priced by the same SDF. Effectively this means there must be a missing factor that prices P1. Since this factor is responsible for the fit of the original portfolios we are back to square one.

Are the D-portfolios priced by consumption growth? On the basis of factor betas the answer is clearly "no". The beta matrix for the D-portfolios is the same transformation of the P-portfolios used to test whether they are equal to a common constant: i.e., $\beta^D = \Phi \beta^P$ where $\Phi = (-\mathbf{1} \quad \mathbf{I}_{n-1})$ and $\mathbf{1}$ is an $(n-1) \times 1$ unit vector. Not surprisingly, since we could not reject the null hypotheses that $\beta_{ij}^P = 0$ for all i , and each j , we also cannot reject the null hypotheses that $\beta_{ij}^D = \beta_{ij}^P - \beta_{1j}^P = 0$ for all i , and each j . As Table 9 indicates, the p-values associated with the test are 0.818 for consumption growth, 0.471 for durables growth and

0.186 for the market return. The SDF betas associated with the D-portfolios are also jointly statistically insignificant (Table 10).

Working with the D-portfolios also does not alleviate the identification problem. The identification problem arises because there exists at least one non-zero $k \times 1$ vector \mathbf{x} such that $\beta^P \mathbf{x} \cong \mathbf{0}$, statistically. Given the definition of β^D it follows that $\beta^D \mathbf{x} \cong \mathbf{0}$, statistically. In fact, the identification problem gets worse, because the transformation Φ is not invertible. Any \mathbf{x} such that $\beta^P \mathbf{x} \cong \mathbf{0}$ implies that $\beta^D \mathbf{x} \cong \mathbf{0}$. But there may be additional \mathbf{x} such that $\beta^D \mathbf{x} \cong \mathbf{0}$ for which $\beta^P \mathbf{x} \neq \mathbf{0}$. This is hardly surprising, since throwing away information is never likely to improve identification. Formal test statistics verifying this are provided in the appendix.

4.3 The Post-Bretton Woods Era

My comment mainly discusses the conclusions we should draw from LV's evidence. Additional evidence from the post-Bretton Woods era, similar to the evidence introduced by LV in their reply, casts further doubt on a consumption-based explanation of carry trade returns. Here I discuss evidence gleaned from a sample of 21 developed-country currencies over the period 1976–2010. The same sample of currencies is used by Burnside, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo (2011) for the period 1976–2009, and similar samples are used by Lustig, Roussanov and Verdelhan (2009) and Lukas Menkhoff, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf (forthcoming), for the period after 1983, to study carry trade portfolios.

My data set consists of spot and forward exchange rates from Reuters/WMR and Barclays, available on Datastream. The raw data are daily observations of spot and one-month forward exchange rates. I use end of month values of these data to create monthly observations. The data span the period January 1976 to December 2010, with the sample varying by currency. As in Lustig, Roussanov and Verdelhan (2009), in each period, the available currencies in my sample are sorted into six bins according to their forward discount against the US dollar. The first bin includes those currencies with the smallest forward discounts (the lowest interest rates), the second bin the next smallest, etc., with the sixth bin consisting of those currencies with the largest forward discounts (and, therefore, the highest interest rates). I then compute the payoff associated with borrowing one dollar in order to invest equally in the riskless securities of the constituent currencies of each bin. This procedure

produces six currency portfolios that I refer to as Q1, Q2, . . . , Q6. Then, I follow the procedure advocated by LV in their reply. I construct the five differenced portfolios, DQ2, DQ3, . . . , DQ6, which involve holding Q_i and shorting Q1, for $i = 2, \dots, 6$. I measure payoffs to these portfolios on a monthly basis, and then compute quarterly excess returns. To assess the model I use data for consumption growth, durables growth and the market return that are available from 1976Q2 to 2010Q1 (sources for these data are described in Burnside et al., 2011, and in the online data archive).

Using these portfolios I find even stronger evidence against a consumption-based explanation of currency returns. As the detailed results in the appendix indicate, none of the betas of the differenced portfolios with respect to consumption growth and durables growth are individually significant. The point estimates also do not display any pattern of increasing with the interest differential. Not surprisingly, as Table 11 indicates, we cannot reject the null hypotheses that $\beta_{ij} = 0$ for all i , for the consumption and durables factors. In this sample, the currency portfolios are correlated with the market return, but the market betas are small. They are roughly 0.08 for DQ2, DQ3 and DQ4, and 0.16 for DQ5 and DQ6. The average excess returns of these five portfolios in the period 1976–2010 were, however, about 1.75, 2.25, 3, 5.25 and 6 percent per annum, as compared to about 7 percent for the stock market. The market betas would need to be five times larger for the CAPM to explain currency returns. Finally, the SDF betas associated with these portfolios are also jointly statistically insignificant (Table 12).

In their reply, LV make much of the fact that carry trades lost money during the recent global financial crisis and other times of market turmoil. But I have never claimed that financial turmoil and currency returns are unrelated. My comment centers on the fact that currency returns and consumption growth are approximately uncorrelated with one another. The global financial crisis is an observation that suits LV’s hypothesis, because it is also an episode in which US consumption growth was low. The other crises (Mexican, Asian, Russian, and Argentinian) that they mention in their reply are not, because they did not coincide with periods of low US consumption growth. According to LV’s data, durables growth was well above average through all of the latter episodes, while nondurable consumption growth was well above average during two and slightly below average in two.

Overall, however, there is little evidence that average carry trade returns can be explained by increased exposure to stock market risk during periods of recession, market downturns,

or market volatility. To demonstrate this, I first regress the monthly returns to the DQ6 portfolio (equivalent to LV's HML carry trade portfolio) on the monthly excess returns of the value-weighted US stock market (the CAPM factor). The market beta of DQ6 in the period 1976M2–2010M12 is 0.18 and is statistically significant. As in the quarterly sample, this beta is much too small (by a factor of more than 5) to explain the average return of the carry trade.

I then divide my data into recessions and expansions as defined by the NBER, periods of high and low stock market volatility, and periods of high and low stock market returns. To measure volatility, I use the daily standard deviation of the market excess return in each month. Months in which this measure of volatility is more than one standard deviation above its mean are denoted as “high volatility”. To divide the data into periods of high and low stock returns, I denote months in which the market excess return is more than one standard deviation below its mean as “low return” months. The market beta of DQ6 is 0.14 during expansions and 0.26 during recessions. It is 0.25 in periods of high volatility and 0.11 in other periods. It is 0.33 in periods of low stock returns and 0.15 in other periods. In each case, however, the difference between the two betas is not statistically significant, and the betas certainly remain insufficiently large to explain the average returns of the carry trade.

5 Conclusion

To explain cross-sectional variation in expected returns, a risk-based story requires that assets with non-zero returns be correlated with the proposed SDF. As Section 2 demonstrates, however, LV's consumption-based SDF is approximately uncorrelated with all of the returns they study. Given this fact, one cannot reject that LV's estimates are consistent with consumption risk explaining none of the cross-sectional variation in the expected returns they studied.

I have argued that the statistical insignificance of the SDF betas leads to two additional problems with LV's conclusions. First, it implies that one cannot ignore sampling uncertainty in the betas when conducting inference about factor risk premia. The statistical significance LV point to largely vanishes when standard errors appropriately reflect this uncertainty. Second, the degree of uncertainty about the betas implies that the factor risk premia are very weakly identified. This makes asymptotic inference less reliable. Confidence sets that are robust to weak identification suggest that LV's data are approximately uninformative

about consumption risk.

Finally, I have argued that LV are able to report strikingly high R^2 measures of fit because their model includes a constant pricing error, which is treated as part of the model's predicted expected returns. When this constant is excluded from the model, the R^2 statistics are much smaller and, in many cases, negative.

A central point of my discussion is that the betas of consumption factors are very poorly estimated. This is why a consumption-based model is difficult to reject using a formal test of the pricing errors. If a great deal more data were collected, one might obtain sufficiently precise estimates of the betas to enable sharper conclusions about the model. But with the data we have, both LV's data and high frequency data from the post-Bretton Woods period, there is no justification for drawing a strong conclusion in favor of a consumption-based model. Burnside et al. (2011) have argued that out-of-sample risk is a potential explanation for the returns to the carry trade. If these currently out-of-sample events occur in the future, it may well turn out that a consumption-based model works. At the moment, however, a model based on in-sample variation in consumption does not work.

REFERENCES

- Burnside, Craig. 2010. "Identification and Inference in Linear Stochastic Discount Factor Models." National Bureau of Economic Research Working Paper 16634.
- Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo. 2011. "Do Peso Problems Explain the Returns to the Carry Trade?" *Review of Financial Studies*, 24(3): 853–91.
- Cochrane, John H. 2005. *Asset Pricing*, Revised edition. Princeton: Princeton University Press.
- Center for Research in Security Prices. 2007. "CRSP Monthly Treasury US Database Guide." http://www.crsp.com/documentation/pdfs/monthly_treasury.pdf.
- den Haan, Wouter J. and Andrew T. Levin. 2000. "Robust Covariance Matrix Estimation with Data-Dependent VAR Prewhitening Order." National Bureau of Economic Research Technical Working Paper 255.
- Fama, Eugene and Kenneth R. French. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, 33(1): 3–56.
- Fama, Eugene and James D. MacBeth. 1973. "Risk Return and Equilibrium: Empirical Tests." *Journal of Political Economy*, 81(3): 607–36.
- Hansen, Lars P. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, 50(4): 1029–54.
- Jagannathan, Ravi and Zhenyu Wang. 1998. "An Asymptotic Theory for Estimating Beta-Pricing Models Using Cross-Sectional Regressions." *Journal of Finance*, 53(4): 1285–309.
- Kan, Raymond and Chu Zhang. 1999a. "GMM Tests of Stochastic Discount Factor Models with Useless Factors." *Journal of Financial Economics*, 54(1): 103–27.
- Kan, Raymond and Chu Zhang. 1999b. "Two-Pass Tests of Asset Pricing Models with Useless Factors." *Journal of Finance*, 54(1): 203–35.
- Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan. 2009. "Common Risk Factors in Currency Markets." Social Science Research Network Paper 1139447.
- 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.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. Forthcoming. "Carry Trades and Global Foreign Exchange Volatility." *Journal of Finance*.

- Shanken, Jay. 1992. "On the Estimation of Beta-Pricing Models." *Review of Financial Studies*, 5(1): 1–33.
- Wright, Jonathan H. 2003. "Detecting Lack of Identification in GMM." *Econometric Theory*, 19(2): 322–30.
- Yogo, Motohiro. 2006. "A Consumption-Based Explanation of Expected Stock Returns." *Journal of Finance*, 61(2): 539–80.

TABLE 1: **Factor Betas and Covariances: Tests for Spread and Against Zero**

	(a) Betas			(b) Covariances		
	Δc	Δd	r_W	Δc	Δd	r_W
Standard error type	Tests for no spread (p-values)					
System-OLS	0.738	0.563	0.273	–	–	–
GMM-VARHAC	0.838	0.596	0.437	0.611	0.901	0.405
	Joint tests vs. zero (p-values)					
System-OLS	0.813	0.623	0.365	–	–	–
GMM-VARHAC	0.799	0.668	0.511	0.447	0.682	0.510

Annual data, 1953–2002. In part (a) the regression equation is $R_{it}^e = a_i + \mathbf{f}_t' \boldsymbol{\beta}_i + \epsilon_{it}$, where R_{it}^e is the excess return of portfolio i at time t , $\mathbf{f}_t = (\Delta c_t \quad \Delta d_t \quad r_{Wt})'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, and r_W is the value weighted US stock market return. The portfolios are equally-weighted groups of short-term foreign-currency denominated money market securities sorted according to their interest differential with the United States, where P1 and P8 are the portfolios with, respectively, the smallest and largest interest differentials. The table reports p-values for tests of the hypotheses that $\beta_{ij} = \beta_j \forall i$ and $\beta_{ij} = 0 \forall i$ for each factor j . In part (b) the covariances between the excess returns of the portfolios and the factors, $\text{cov}(R_i^e, f_j)$, are estimated by GMM. The table reports p-values for tests of the hypotheses that $\text{cov}(R_i^e, f_j) = c_j \forall i$ and $\text{cov}(R_i^e, f_j) = 0 \forall i$ for each factor j .

TABLE 2: **SDF Betas: Tests for Spread and Against Zero**

	Model (i)	Model (ii)	Model (iii)
Standard error type	Tests for no spread (p-values)		
System-OLS	0.508	0.477	0.575
GMM-VARHAC	0.688	0.589	0.602
	Joint tests vs. zero (p-values)		
System-OLS	0.469	0.443	0.505
GMM-VARHAC	0.306	0.170	0.329

Annual data, 1953–2002. The regression equation is $R_{it}^e = a_i - m_t \beta_{im} + \epsilon_{it}$, where $m_t = 1 - (\mathbf{f}_t - \bar{\mathbf{f}})' \mathbf{b}$, $\bar{\mathbf{f}}$ is the sample mean of \mathbf{f}_t , and the vector \mathbf{b} takes on one of the following three values: (i) $\mathbf{b} = (-21.0 \ 129.9 \ 4.46)'$ (corresponds to LV's two pass estimate with a constant), (ii) $\mathbf{b} = (37.0 \ 74.7 \ 4.65)'$ (corresponds to LV's GMM estimate with no constant) and (iii) $\mathbf{b} = (6.74 \ 23.3 \ 0.31)'$ (the calibrated model). Here R_{it}^e is a portfolio return, and \mathbf{f}_t is the vector of factors, described in Table 1. The table reports p-values for tests of the hypotheses that $\beta_{im} = \beta_m \ \forall i$ and $\beta_{im} = 0 \ \forall i$.

TABLE 3: Second-Pass Regression with a Constant

Factor Prices (λ)				R^2	p-value	MAE
Constant	Nondurables (Δc)	Durables (Δd)	Market (r_W)			
-2.94	2.19	4.70	3.33	0.87		0.44
(0.86)	(0.83)	(0.97)	(7.59)		(0.483)	
[2.23]	[2.11]	[2.42]	[18.8]		[0.972]	
{2.66}	{2.48}	{2.41}	{23.1}		{0.994}	
<-4.9,2.5>	<-1.8,4.7>	<-3.8,7.8>	<-32,28>			

Annual data, 1953–2002. The table reports results from running the cross-sectional regression $\bar{R}_i^e = \gamma + \hat{\beta}_i' \lambda + u_i$, where \bar{R}_i^e is the mean excess return of portfolio i and $\hat{\beta}_i$ is the vector of factor betas of portfolio i estimated in the first pass regression. The portfolios and factors are described in Table 1. For the factor risk premia ($\hat{\lambda}$) OLS standard errors are in parentheses, Shanken standard errors are in square brackets, and GMM-VARHAC standard errors are in braces. Bootstrapped 95 percent confidence regions are in angled brackets. For the tests of the pricing errors I compute the test statistic for each of the three methods of computing the covariance matrix of $\hat{\mathbf{u}}$ (OLS, Shanken and GMM-VARHAC), and report the associated p-value. The R^2 statistic from the second-pass regression is reported along with the mean absolute pricing error (MAE).

TABLE 4: GMM Estimates of the Model with the Constant

Stage	Model Parameters (\mathbf{b})			R^2	p-value	MAE
	Nondurables (Δc)	Durables (Δd)	Market (r_W)			
1st	-21.0 (87.7) <-291,304>	129.9 (97.4) <-199,296>	4.46 (4.83) <-11,13>	0.87		0.44
2nd	27.2 (76.2) <-201,253>	108.3 (97.3) <-147,227>	2.53 (4.38) <-8.7,10>	0.81	0.703	0.42
Iterated	59.5 (74.2) <-206,268>	88.4 (76.4) <-155,228>	0.15 (4.03) <-9.5,11>	0.50	0.915	0.66

Annual data, 1953–2002. The table reports GMM estimates of \mathbf{b} obtained by exploiting the moment restrictions $E\{\mathbf{R}_t^e[1 - (\mathbf{f}_t - \boldsymbol{\mu})' \mathbf{b}] - \gamma\} = 0$ and $E(\mathbf{f}_t - \boldsymbol{\mu}) = \mathbf{0}$, where \mathbf{R}_t^e is the vector of currency portfolio returns, and \mathbf{f}_t is the vector of risk factors, described in Table 1. GMM-VARHAC standard errors are in parentheses. Bootstrapped 95 percent confidence regions are in angled brackets. For the test of the pricing errors I report the p-value associated with the test-statistic.

TABLE 5: GMM Estimates of the Model with the Constant

Stage	Factor Prices (λ)			
	Constant	Nondurables (Δc)	Durables (Δd)	Market (r_W)
1st	-2.94 (2.63) <-4.9,2.5>	2.19 (1.88) <-1.8,4.7>	4.70 (3.16) <-3.8,7.8>	3.33 (13.0) <-32,28>
2nd	-3.02 (2.23) <-4.2,1.4>	2.73 (1.81) <-1.1,4.2>	4.85 (3.15) <-2.6,6.4>	0.78 (10.9) <-25,24>
Iterated	-3.56 (2.01) <-4.3,1.5>	2.94 (1.59) <-1.3,4.6>	4.81 (2.62) <-2.9,6.6>	-3.87 (11.8) <-28,31>

Annual data, 1953–2002. The table reports GMM estimates of γ and λ obtained by exploiting the moment restrictions $E\{\mathbf{R}_t^e[1 - (\mathbf{f}_t - \boldsymbol{\mu})' \mathbf{b}] - \gamma\} = 0$, $E(\mathbf{f}_t - \boldsymbol{\mu}) = \mathbf{0}$ and $E[(\mathbf{f}_t - \boldsymbol{\mu})(\mathbf{f}_t - \boldsymbol{\mu})' - \boldsymbol{\Sigma}_f] = 0$, where \mathbf{R}_t^e is the vector of currency portfolio returns, and \mathbf{f}_t is the vector of risk factors, described in Table 1. GMM-VARHAC standard errors are in parentheses. Bootstrapped 95 percent confidence regions are in angled brackets.

TABLE 6: **Second-Pass Regression without a Constant**

Factor Prices (λ)			R^2	p-value	MAE
Nondurables (Δc)	Durables (Δd)	Market (r_W)			
0.59 (0.73) [1.01] {1.17} <-1.7,3.4>	1.10 (1.02) [1.40] {1.69} <-3.0,5.9>	11.7 (7.40) [10.1] {10.6} <-24,24>	0.34	(0.001) [0.059] {0.173}	1.17

Annual data, 1953–2002. The table reports results from running the cross-sectional regression $\bar{R}_i^e = \hat{\beta}_i' \lambda + \alpha_i$, where \bar{R}_i^e is the mean excess return of portfolio i and $\hat{\beta}_i$ is the vector of factor betas of portfolio i estimated in the first pass regression. The portfolios and factors are described in Table 1. For the factor risk premia ($\hat{\lambda}$) OLS standard errors are in parentheses, Shanken standard errors are in square brackets, and GMM-VARHAC standard errors are in braces. Bootstrapped 95 percent confidence regions are in angled brackets. For the tests of the pricing errors I compute the test statistic for each of the three methods of computing the covariance matrix of $\hat{\alpha}$ (OLS, Shanken and GMM-VARHAC), and report the p-value associated with the test-statistic. The R^2 statistic from the second-pass regression is reported along with the mean absolute pricing error (MAE).

TABLE 7: GMM Estimates of the Model with no Constant

Stage	Model Parameters (\mathbf{b})			R^2	p-value	MAE
	Nondurables (Δc)	Durables (Δd)	Market (r_W)			
1st	-22.0 (62.9) <-237,235>	45.5 (50.3) <-155,231>	5.16 (2.88) <-8.4,11>	0.34		1.17
2nd	37.0 (45.2) <-183,213>	74.7 (33.0) <-129,187>	4.65 (2.70) <-7.0,9.4>	-0.66	0.068	1.86
Iterated	25.8 (40.6) <-191,235>	66.0 (47.4) <-138,193>	-2.14 (2.94) <-8.0,11>	-1.45	0.281	2.28

Annual data, 1953–2002. The table reports GMM estimates of \mathbf{b} obtained by exploiting the moment restrictions $E\{\mathbf{R}_t^e[1 - (\mathbf{f}_t - \boldsymbol{\mu})' \mathbf{b}]\} = 0$ and $E(\mathbf{f}_t - \boldsymbol{\mu}) = \mathbf{0}$, where \mathbf{R}_t^e is the vector of currency portfolio returns, and \mathbf{f}_t is the vector of risk factors, described in Table 1. GMM-VARHAC standard errors are in parentheses. Bootstrapped 95 percent confidence regions are in angled brackets. For the test of the pricing errors I report the p-value associated with the test-statistic.

TABLE 8: GMM Estimates of the Model with no Constant

Factor Prices (λ)			
Stage	Nondurables (Δc)	Durables (Δd)	Market (r_W)
1st	0.59 (1.07) <-1.7,3.4>	1.10 (1.64) <-3.0,5.9>	11.7 (8.26) <-24,24>
2nd	2.37 (1.00) <-1.2,3.4>	3.48 (1.13) <-2.4,5.0>	10.2 (7.37) <-21,25>
Iterated	1.72 (0.93) <-1.4,3.7>	3.41 (1.52) <-2.6,5.2>	-10.6 (8.80) <-24,31>

Annual data, 1953–2002. The table reports GMM estimates of λ obtained by exploiting the moment restrictions $E\{\mathbf{R}_t^e[1 - (\mathbf{f}_t - \boldsymbol{\mu})' \mathbf{b}]\} = 0$, $E(\mathbf{f}_t - \boldsymbol{\mu}) = \mathbf{0}$ and $E[(\mathbf{f}_t - \boldsymbol{\mu})(\mathbf{f}_t - \boldsymbol{\mu})' - \boldsymbol{\Sigma}_f] = 0$, where \mathbf{R}_t^e is the vector of currency portfolio returns, and \mathbf{f}_t is the vector of risk factors, described in Table 1. GMM-VARHAC standard errors are in parentheses. Bootstrapped 95 percent confidence regions are in angled brackets.

TABLE 9: Factor Betas and Covariances for Annual Differenced Portfolios, Tests for Spread and Against Zero

	(a) Betas			(b) Covariances		
	Δc	Δd	r_W	Δc	Δd	r_W
Standard error type	Tests for no spread (p-values)					
System-OLS	0.646	0.644	0.371
GMM-VARHAC	0.728	0.567	0.350	0.517	0.845	0.503
	Joint tests vs. zero (p-values)					
System-OLS	0.738	0.563	0.273
GMM-VARHAC	0.818	0.471	0.186	0.560	0.895	0.402

Annual Data, 1953–2002. In part (a) the regression equation is $R_{it}^e = a_i + \mathbf{f}'_t \beta_i + \epsilon_{it}$, where \mathbf{f}_t is the vector of factors described in Table 1, and R_{it}^e is a portfolio return. The portfolios are D2, D3, ..., D8, the returns to holding long positions in LV's portfolios P2, P3, ..., P8 (described in Table 1) while holding a short position in portfolio P1. The table reports p-values for tests of the hypotheses that $\beta_{ij} = \beta_j \forall i$ and $\beta_{ij} = 0 \forall i$ for each factor j . In part (b) the covariances between the excess returns and the factors, $\text{cov}(R_i^e, f_j)$, are estimated by GMM. The table reports p-values for tests of the hypotheses that $\text{cov}(R_i^e, f_j) = c_j \forall i$ and $\text{cov}(R_i^e, f_j) = 0 \forall i$ for each factor j .

TABLE 10: **SDF Betas for Annual Differenced Portfolios, Tests for Spread and Against Zero**

	Model (i)	Model (ii)	Model (iii)
Standard error type	Tests for no spread (p-values)		
System-OLS	0.672	0.664	0.660
GMM-VARHAC	0.668	0.475	0.485
	Joint tests vs. zero (p-values)		
System-OLS	0.508	0.477	0.575
GMM-VARHAC	0.404	0.162	0.321

Annual Data, 1953–2002. The regression equation is $R_{it}^e = a_i - m_t \beta_{im} + \epsilon_{it}$, where $m_t = 1 - (\mathbf{f}_t - \bar{\mathbf{f}})' \mathbf{b}$, $\bar{\mathbf{f}}$ is the sample mean of \mathbf{f}_t , and the vector \mathbf{b} takes on one of the following three values: (i) $\mathbf{b} = (-21.0 \ 129.9 \ 4.46)'$ (corresponds to LV's two pass estimate with a constant), (ii) $\mathbf{b} = (37.0 \ 74.7 \ 4.65)'$ (corresponds to LV's GMM estimate with no constant) and (iii) $\mathbf{b} = (6.74 \ 23.3 \ 0.31)'$ (the calibrated model). Here, R_{it}^e and \mathbf{f}_t are the returns and factors described in the note to Table 10. The table reports p-values for tests of the hypotheses that $\beta_{im} = \beta_m \ \forall i$ and $\beta_{im} = 0 \ \forall i$.

TABLE 11: Factor Betas and Covariances for Quarterly Differenced Portfolios, Tests for Spread and Against Zero

	(a) Betas			(b) Covariances		
	Δc	Δd	r_W	Δc	Δd	r_W
Standard error type	Tests for no spread (p-values)					
System-OLS	0.629	0.499	0.085
GMM-VARHAC	0.542	0.615	0.044	0.494	0.402	0.072
	Joint tests vs. zero (p-values)					
System-OLS	0.763	0.618	0.016
GMM-VARHAC	0.685	0.681	0.049	0.640	0.470	0.113

Quarterly Data, 1976Q2–2010Q1. In part (a) the regression equation is $R_{it}^e = a_i + \mathbf{f}_t' \beta_i + \epsilon_{it}$, where R_{it}^e is the excess return of portfolio i at time t , and \mathbf{f}_t is the vector of risk factor described in Table 1. The portfolios are DQ2, DQ3, ..., DQ5, the returns to holding long positions in the five quarterly portfolios Q2, Q3, ..., Q6 while holding a short position in the quarterly portfolio Q1. The table reports p-values for tests of the hypotheses that $\beta_{ij} = \beta_j \forall i$ and $\beta_{ij} = 0 \forall i$ for each factor j . In part (b) the covariances between the excess returns and the factors, $\text{cov}(R_i^e, f_j)$, are estimated by GMM. The table reports p-values for tests of the hypotheses that $\text{cov}(R_i^e, f_j) = c_j \forall i$ and $\text{cov}(R_i^e, f_j) = 0 \forall i$ for each factor j .

TABLE 12: SDF Betas for Quarterly Differenced Portfolios, Tests for Spread and Against Zero

	Model (i)	Model (ii)	Model (iii)
Standard error type	Tests for no spread (p-values)		
System-OLS	0.333	0.185	0.273
GMM-VARHAC	0.280	0.103	0.301
	Joint tests vs. zero (p-values)		
System-OLS	0.409	0.168	0.399
GMM-VARHAC	0.332	0.139	0.431

Quarterly Data, 1976Q2–2010Q1. The regression equation is $R_{it}^e = a_i - m_t \beta_{im} + \epsilon_{it}$, where $m_t = 1 - (\mathbf{f}_t - \bar{\mathbf{f}})' \mathbf{b}$, $\bar{\mathbf{f}}$ is the sample mean of \mathbf{f}_t , and the vector \mathbf{b} takes on one of the following three values: (i) $\mathbf{b} = (-21.0 \ 129.9 \ 4.46)'$ (corresponds to LV's two pass estimate with a constant), (ii) $\mathbf{b} = (37.0 \ 74.7 \ 4.65)'$ (corresponds to LV's GMM estimate with no constant) and (iii) $\mathbf{b} = (6.74 \ 23.3 \ 0.31)'$ (the calibrated model). Here, R_{it}^e and \mathbf{f}_t are the returns and factors described in the note to Table 12. The table reports p-values for tests of the hypotheses that $\beta_{im} = \beta_m \ \forall i$ and $\beta_{im} = 0 \ \forall i$.