The Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk: A Comment*

Craig Burnside†

February 2009

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. The fact that such contrasting conclusions can be reached from the same set of data reflects the statistically weak identification of their model.

J.E.L. Classification: F31, G12.

*I thank John Cochrane, Martin Eichenbaum, Ravi Jagannathan, Sergio Rebelo and an anonymous referee for helpful comments and useful conversations, and the NSF for research support. Responsibility for the opinions expressed and possible errors contained within this comment is entirely my own.

†Duke University and NBER.

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 simple 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. They argue that the R^2 statistic corresponding to their benchmark estimates means that their model explains about 87 percent of the cross-sectional variation in expected returns.

To the contrary, I argue that their evidence is consistent with consumption risk explaining very little of the cross-sectional variation in the expected returns of their portfolios. Theory predicts that the expected excess return on an asset, $E(R_t^e)$, is given by $-\operatorname{cov}(R_t^e, m_t)$, 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, one cannot reject the null hypothesis that the covariances in question are all zero. Consequently, one cannot reject the null hypothesis that consumption risk explains none of the cross-sectional variation in the expected excess returns in their data set.

The lack of statistical spread, across portfolios, in the covariance between the returns and the SDF reflects that the model is, in a statistical sense, very poorly identified. LV build SDFs that are linear in a vector risk factors. They implement a widely-used two-pass procedure to estimate these SDFs. The first pass is a time series regression of each portfolio's return on the risk factors. This regression determines the factor betas, $\boldsymbol{\beta}$. The second pass is a cross-sectional regression of average portfolio returns on these betas. This regression determines the lambdas, $\boldsymbol{\lambda}$, or factor risk premia. When there are n portfolios and k risk factors, $\boldsymbol{\beta}$ is an $n \times k$ matrix and $\boldsymbol{\lambda}$ is a $k \times 1$ vector. For $\boldsymbol{\lambda}$ to be identified, $\boldsymbol{\beta}$ must have full column rank. Statistical tests indicate that for LV's model, in which $k = 3$, the rank of $\boldsymbol{\beta}$ is very low, perhaps as low as 0. It is not surprising, in this situation, that the covariances between the returns and the SDF are close to zero, because these covariances are given by a linear combination of the columns of $\boldsymbol{\beta}$.

The identification problem raises three 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, 2007). 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 $\boldsymbol{\lambda}$. Third, LV are able to draw such favorable inference about their model because they focus mainly on standard errors for $\boldsymbol{\lambda}$ that treat the betas, $\boldsymbol{\beta}$, as known regressors, rather than generated regressors. This effectively abstracts away the identification problem since $\boldsymbol{\beta}$ always has full column rank in a finite sample. But even if one accepts that $\boldsymbol{\beta}$ has full rank, treating $\boldsymbol{\beta}$ as known leads to a misleading level of confidence in the model. With conventionally calculated standard errors 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.

Even if we discount all of these econometric considerations, however, there is still a problem with LV's findings. As is common in the finance literature, LV include a constant in the second-pass regression. If the model is true this constant should be zero, so, correctly interpreted, the constant is really part of the model's pricing error. It is not a component of the expected return "explained" by the SDF. I show that without the constant the R^2 of the model drops from 0.87 to a maximum of 0.34 across a range of model estimates. The mean absolute pricing error increases from 0.4 percent to between 1.2 and 2.3 percent. The economic interpretation of this finding is that there is a large unexplained common component of the expected excess returns of LV's currency portfolios.

In their reply, LV defend their findings on two main grounds. First, 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. Here, I show that the inclusion of these test assets 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 $E(R^e_t) = 0$ 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.

Second, LV refer to evidence not in their original paper. 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. They show that the fit of their model to these seven portfolios is insensitive to the inclusion of a constant in the second pass regression. I discuss why one should not be in any way surprised by the latter finding in Section 5. However, the evidence regarding the seven "differenced" portfolios affects none of my other 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 $E(R_t^e) = 0$ for all of the "differenced" portfolios. Also, because the "differenced" portfolios are linear combinations of the original portfolios, working with these portfolios can only make the identification problem worse.

I conclude that, taken as a whole, the evidence for LV's consumption-based model is weak. I cannot reject that the model predicted expected returns of the currency portfolios they study are all zero. It is also true that it is impossible to reject their model using formal statistical tests. But this reflects lack of information in the data about the factor betas of currency portfolios. It is not, I would argue, a reflection of the model's "success".

In Section 1 I briefly review their 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 identification problem and its implications. In Section 4, 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 5. Section 6 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

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

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, ..., P8 with the order running from low interest rate currencies to high interest rate currencies.[1]

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

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

I rewrite (1) generically as

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

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

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

where $\boldsymbol{\beta}$ is a $n \times k$ matrix of factor betas, and $\boldsymbol{\lambda}$ is a $k \times 1$ vector of factor risk premia.

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

[1] Further details of the model, portfolio formation, and data sources can be found in LV's article and in an appendix to this comment available from the author.

on the vector of risk factors:

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

Here $\boldsymbol{\beta}_i'$ represents the ith 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:

$$\bar{R}_i^e = \hat{\boldsymbol{\beta}}_i'\boldsymbol{\lambda} + \alpha_i, \qquad i = 1, \ldots, n, \tag{6}$$

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

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

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

where $\ddot{R}^e = \frac{1}{n}\sum_{i=1}^n \bar{R}_i^e$ 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 $C_{\hat{\boldsymbol{\alpha}}} = T\hat{\boldsymbol{\alpha}}'\hat{\boldsymbol{\Omega}}_{\hat{\alpha}}^{-1}\hat{\boldsymbol{\alpha}}$, where $\hat{\boldsymbol{\Omega}}_{\hat{\alpha}}$ is a consistent estimator for the asymptotic covariance matrix of $\sqrt{T}\hat{\boldsymbol{\alpha}}$ and the inverse is generalized. John H. Cochrane (2005) discusses how to form $\hat{\boldsymbol{\Omega}}_{\hat{\alpha}}$ and shows that $C_{\hat{\boldsymbol{\alpha}}} \xrightarrow{d} \chi_{n-k}^2$.

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

$$\bar{R}_i^e = \gamma + \hat{\boldsymbol{\beta}}_i'\boldsymbol{\lambda} + u_i, \qquad i = 1, \ldots, n. \tag{8}$$

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.

Including the constant in the regression does not bias estimates of $\boldsymbol{\lambda}$, since, if the model is true and identified, $\text{plim}\,\hat{\gamma} = 0$. At the very least, however, the economic and statistical significance of $\hat{\gamma}$ should be considered before a model is deemed reasonable. We should also be skeptical if inclusion of the constant significantly boosts the R^2 of the model since, if the model is true, the probability limit of the R^2 statistic is 1, whether or not a constant is included in the second-pass regression.

1.2 GMM Estimation

Cochrane (2005) describes a generalized method of moments (GMM, Lars P. Hansen 1982) 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

$$E(R^e_{it} - a_i - \boldsymbol{\beta}_i'\mathbf{f}_t) = 0, \qquad i = 1, \ldots, n. \tag{9}$$

$$E[(R^e_{it} - a_i - \boldsymbol{\beta}_i'\mathbf{f}_t)\mathbf{f}_t'] = 0, \qquad i = 1, \ldots, n. \tag{10}$$

$$E(R^e_{it} - \gamma - \boldsymbol{\beta}_i'\boldsymbol{\lambda}) = 0, \qquad i = 1, \ldots, n. \tag{11}$$

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

$$E(R^e_{it} - \boldsymbol{\beta}_i'\boldsymbol{\lambda}) = 0, \qquad i = 1, \ldots, n. \tag{12}$$

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:

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

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

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:

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

When an identity matrix is used to weight the moments in (15), and the GMM procedure is set up in such a way that $\hat{\boldsymbol{\mu}}$ equals the sample mean of $\mathbf{f}_t$, the GMM procedure is numerically identical to the two-pass regression method in terms of pricing errors.[2]

2 First-Pass Estimates of Betas

Table 1 presents first-pass estimates of the betas obtained by running the least squares regressions described by (5). Standard errors are computed using either the standard system

[2] If 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 an appendix available from the author.

OLS formulas, or a GMM-based VARHAC procedure.[3] The individual standard errors are similar across the two procedures. None of the betas reported in Table 1 are *individually* statistically significant.

We can also test whether for each *portfolio* the three factor betas are jointly significant. As Table 1 indicates, at conventional significance levels one cannot reject that $\beta_{ij} = 0 \ \forall j$, for each $i = 1, \ldots, n$. Since the model-predicted expected return for each portfolio is given by $\boldsymbol{\beta}_i' \boldsymbol{\lambda}$ the hypothesis tests in Table 1 imply that one cannot reject the null that the model-predicted expected return for each portfolio is zero. Of course, the success of an asset pricing model does not rest on the statistical significance of *all* the betas. And we would not expect a portfolio with a zero average excess return to have significant betas. When a portfolio, however, has a non-zero and statistically significant average excess return, presumably *some* of its betas should be statistically significant.

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 indicates, at conventional significance levels one cannot reject that $\beta_{ij} = 0 \ \forall i$, for each $j = 1, \ldots, k$. Since the contribution of factor j to the vector of model-predicted expected returns is $\boldsymbol{\beta}_j \lambda_j$ these hypothesis tests imply that one cannot reject the null that each factor's risk premium contributes nothing to the model-predicted expected returns.

Finally, if we test whether literally all the betas are jointly zero, the χ^2_{24} statistic has a p-value of 0.64, although with VARHAC standard errors the p-value is very small. It turns out that this result is entirely driven by a single beta. If we test whether every beta except the durables beta of P7 is zero, the χ^2_{23} statistic has an OLS p-value of 0.78 and a VARHAC p-value of 0.06.

Another way to assess the model is to look for covariance between the SDF and the portfolio returns. By forming SDFs that are linear combinations of the factors, we may induce significant cross-sectional spread in $\text{cov}(m_t, R^e_{it})$ across i that is hard to detect in the factor betas. One way to capture such spread is to measure the SDF betas. Using (3), (2) can be rewritten as

$$E(\mathbf{R}^e_t) = -\,\text{cov}(\mathbf{R}^e_t, m_t)/E(m_t). \tag{16}$$

[3] 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). Details are provided in an available appendix.

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

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

To measure m_t, using (3), we need values for the elements of $\mathbf{b}$. Here I use three versions of $\mathbf{b}$ taken directly from LV's article. The SDF betas for these different values of $\mathbf{b}$ are shown in Table 2.

In Table 2(a) I use the $\mathbf{b}$ vector corresponding to LV's two pass estimates of $\boldsymbol{\lambda}$: $b_c = -21$, $b_d = 130$ and $b_r = 4.5$. The beta for P7 is significant at the 5 percent level. However, as Table 2 indicates, the null hypotheses that $\beta_{im} = 0$ or that $\beta_{im} = \beta_m$ (a common constant) for all i cannot be rejected.

Table 2(b) uses the $\mathbf{b}$ vector corresponding to LV's GMM estimates of $\mathbf{b}$: $b_c = 37$, $b_d = 75$ and $b_r = 4.7$. None of the estimated SDF betas is individually significantly different from zero at the 5 percent level. As Table 2 indicates, the null hypotheses that $\beta_{im} = 0$ or that $\beta_{im} = \beta_m$ for all i cannot be rejected.

Table 2(c) repeats the exercise using the $\mathbf{b}$ vector corresponding to the calibrated model discussed in section I.E of LV's paper: $b_c = 6.7$, $b_d = 23$ and $b_r = 0.31$. As in case (a) only the beta for P7 is significant at the 5 percent level. However, as Table 2 indicates, the null hypotheses that $\beta_{im} = 0$ for all i or that $\beta_{im} = \beta_m$ for all i cannot be rejected.

Linear factor models rely on there being significant spread in the covariance between the risk factors and the returns. The general lack of statistical significance in the factor and SDF betas and the lack of significant spread of these betas across assets cast reasonable doubt on the hypothesis that LV's model explains the cross-section of the expected returns of their portfolios.

3 Weak Identification

In this section I argue that tests of the pricing errors fail to reject LV's model at low levels of significance due to an identification problem. In the second-pass regression with the constant, the parameters γ and $\boldsymbol{\lambda}$ are identified under the assumption that $\boldsymbol{\beta}^+ = (\ \boldsymbol{\iota} \quad \boldsymbol{\beta}\ )$ has full column rank, where $\boldsymbol{\iota}$ is an $n \times 1$ vector of ones. In the second-pass regression with no constant, $\boldsymbol{\lambda}$ is identified if $\boldsymbol{\beta}$ has full column rank. The same conditions must hold for the GMM procedure used to estimate $\mathbf{b}$. 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 (2007) discusses, t-statistics for $\hat{\boldsymbol{\lambda}}$ and $\hat{\mathbf{b}}$ have non-standard distributions, and, most importantly, pricing-error tests cannot reliably detect model misspecification.

It is straightforward to test whether LV's model is identified using a rank test from Jonathan H. Wright (2003). Table 3 presents tests of the null hypothesis that $\boldsymbol{\beta}$ has reduced rank. Since there are three risk factors in the model, $\boldsymbol{\beta}$ has reduced rank if rank($\boldsymbol{\beta}$) < 3. As Table 3 indicates, it is not possible to reject that rank($\boldsymbol{\beta}$) $= 1$ at conventional significance levels. In fact, it is only when VARHAC standard errors are used that we can reject the null hypothesis that rank($\boldsymbol{\beta}$) $= 0$, which is equivalent to the null that *every* element of $\boldsymbol{\beta}$ is zero. Similar results are obtained when I test the rank of $\boldsymbol{\beta}^+$.

To see, intuitively, why lack of identification limits the power of pricing-error tests, imagine that a researcher was confronted with the m_t series constructed using the $\mathbf{b}$ corresponding to LV's benchmark two-pass estimates ($b_c = -21.0$, $b_d = 130$, $b_r = 4.46$). Suppose this researcher used the two-pass procedure to estimate a linear factor model, treating $-m_t$ as a risk factor. In the second pass, the researcher would regress the sample expected returns, $\bar{R}_i^e$, against the SDF betas, $\hat{\beta}_{mi}$, presented in Table 2(a). A scatter plot of $\bar{R}_i^e$ against $\hat{\beta}_{mi}$ is presented in Figure 1, with the estimated regression line The horizontal bar around each estimate, $\hat{\beta}_{mi}$, represents a 95 percent confidence interval for $\hat{\beta}_{mi}$ treating $\mathbf{b}$ as known. The width of the bars illustrates the enormous degree of uncertainty about the betas that we saw in Section 2. The vertical bar around each estimate, $\bar{R}_i^e$, represents a 95 percent confidence interval for $\bar{R}_i^e$.

The estimated regression line corresponds to $\hat{\gamma}_m = -2.9$ and $\hat{\lambda}_m \cong 5.8$, and, as Figure 1 indicates, it fits the scatter plot quite well. However, once the degree of uncertainty associated with $\bar{R}^e$ and $\hat{\beta}_m$ is taken into account, it is abundantly clear that a very wide range of parameter pairs (γ_m, λ_m) would also fit the scatter plot reasonably well. This conjecture is easy to verify formally. 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.[4] Roughly speaking, the confidence set is constructed as follows. At each point, (γ_m, λ_m), in the parameter space the CU-GMM objective function is evaluated, and the corresponding p-value is calculated. Those pairs, (γ_m, λ_m), for which the p-value exceeds 0.05 lie in the 95 percent confidence

[4] See James H. Stock and Wright (2000), Stock, Wright and Yogo (2002) and Yogo (2004) for details.

set for (γ_m, λ_m). The difference between this approach and standard GMM is that the weighting matrix is recalculated at each (γ_m, λ_m), and no degrees of freedom are lost because γ_m and λ_m are not estimated. Figure 2 shows the robust 95 percent confidence set for (γ_m, λ_m) constructed using the m_t series described above. Points outside the confidence set are indicated by grey shading. The two-pass estimate, $(-2.9, 5.8)$, is indicated by the solid dot. The CU-GMM estimate, $(-4.2, 8.0)$, is indicated by the open dot. While the rank tests are the definitive indication that the data are uninformative about $\boldsymbol{\lambda}$, the fact that the confidence set in Figure 2 is a vast, disjoint, and unbounded subset of the parameter space provides a graphical illustration of the problem.

It is not surprising that points around the horizontal axis ($\lambda_m = 0$) lie outside the confidence set. The value of the CU-GMM objective function at $\lambda_m = 0$ and any value of γ_m is equivalent to a test statistic for the null hypothesis that $E(R_i^e) = \gamma_m$ for all i. As long as there is statistically significant variation in the mean excess returns, this hypothesis will be rejected for *any* proposed SDF.[5]

When confidence sets for parameters are constructed using the CU-GMM objective function, lack of information about the parameters goes hand-in-hand with inability to reject the model. Given the confidence set in Figure 2, of course, we cannot formally reject the model. It is easy to find parameter pairs where the test of the over-identifying restrictions fails to reject. But this hardly seems like a signal of the model's success.

4 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 regressions, which include the constant γ, are reproduced in Table 4(a). When presenting their findings they compute standard errors—shown in the 'OLS' column—under the assumption 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

[5] In LV's case, as Table 1 and Figure 1 indicate, the mean excess returns of P1 and P7 are significantly different from zero.

pricing errors is 0.48. These results are the main basis of LV's positive assessment of the model.

There are three reasons our assessment should be less sanguine. 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 correctly, 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 doubtful about the model estimates is the identification problem discussed earlier. Under non-identification or weak identification even appropriately calculated asymptotic standard errors are likely to understate the degree of uncertainty about the model parameters. The third 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.

4.1 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 is true even asymptotically. There are two standard ways to deal with this problem. One is to use the correction of the standard errors suggested by Jay Shanken (1992). The other is to compute the standard errors using the first of the two GMM procedures described above that replicates the point estimates. By construction, the alternative approaches to calculating standard errors do not affect the point estimates of the factor risk premia. Also, the Shanken standard errors are a special case of the GMM standard errors when the ϵ_{it}, in (5), are i.i.d. and homoskedastic. The GMM procedure is more general, but as I show here, the two procedures deliver quite similar results. I also show that correcting the errors matters quantitatively.

The Shanken and GMM-corrected standard errors for the model with the constant [Table 4(a)] 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 $\boldsymbol{\Sigma}_{\mathbf{f}} = E[(\mathbf{f}_t - \boldsymbol{\mu})(\mathbf{f}_t - \boldsymbol{\mu})']$, 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

$$\boldsymbol{\Omega}_{\hat{\boldsymbol{\theta}}} = (\boldsymbol{\beta}^{+\prime}\boldsymbol{\beta}^{+})^{-1}\boldsymbol{\beta}^{+\prime}\boldsymbol{\Sigma}\boldsymbol{\beta}^{+}(\boldsymbol{\beta}^{+\prime}\boldsymbol{\beta}^{+})^{-1} + \tilde{\boldsymbol{\Sigma}}_{\mathbf{f}}. \tag{18}$$

With the Shanken correction the covariance matrix is

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

In some finance applications the Shanken correction is small. For example, for the CAPM estimated using annual returns of the Fama and Kenneth R. 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 1.035. In LV's case the estimate of $1 + \boldsymbol{\lambda}'\boldsymbol{\Sigma}_{\mathbf{f}}^{-1}\boldsymbol{\lambda}$ is 6.79. Although the individual λs 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 4(a) are statistically significant at the 5 percent level, though the risk premium corresponding to durables growth is significant at the 10 percent level. 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 relatively uninformative about the price of consumption risk.[6]

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 **b**. GMM estimates of **b** for the model with the constant are found in Table 5. 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 **b** are statistically insignificant for every risk factor. 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.

[6] 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 use inappropriate GMM standard errors, as detailed in the appendix. Third, they argue that standard errors from a bootstrap procedure are small enough to make the estimated risk premia significant. These standard errors are also calculated inappropriately as detailed in the available appendix.

The lack of identification we saw in Section 3, however, suggests that even appropriately calculated asymptotic standard errors understate the degree of uncertainty we have about $\boldsymbol{\lambda}$ and $\mathbf{b}$. Unfortunately, it is not possible to plot robust confidence sets for these parameter vectors, since they have three elements. We can, however, repeat the exercise of Section 3, in which we imagined a researcher confronted with the m_t series constructed using the benchmark two-pass estimates. Suppose this researcher used the two-pass procedure to estimate a linear factor model, treating $-m_t$ as a risk factor. Suppose this researcher used his estimates to construct confidence sets for (γ_m, λ_m) using OLS, Shanken and GMM standard errors. These confidence sets are plotted in Figure 3, along with the robust confidence set calculated in Section 3. The GMM and Shanken procedures produce confidence sets with roughly nine to ten times the area of the OLS procedure. The confidence sets certainly cast doubt on the significance of the parameter estimates. However, none of the confidence sets capture the uncertainty revealed by the robust confidence set.

4.2 Imposing a Zero Intercept

Given the estimates in Table 4(a) 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 4(a) 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.

Table 4(b) 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.

We can also estimate the model without the constant by GMM. In Table 6(a) I present the results from the first stage of GMM. Here, none of the **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 [Table 6(b)] 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 λs associated with consumption and durables are statistically significant, the **b** parameter is only significant for durables. Further iterations on the weighting matrix [Table 6(c)] 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 **b** parameters are significant. In summary, the model fits very poorly absent the constant.

Unfortunately, excluding the constant from the model does not alleviate the identification problem alluded to earlier. The model with the constant is identified if the $\boldsymbol{\beta}^+$ matrix has full rank. The model without the constant is identified if the $\boldsymbol{\beta}$ matrix has full rank. If the problem with the model were simply that there was lack of spread in the betas for one factor, then getting rid of the constant would solve the problem. As Table 3 indicates, however, the rank of the $\boldsymbol{\beta}$ matrix appears to be 1 or less. The identification problem is severe. If we construct confidence sets for λ_m with γ_m set equal to zero, we obtain $(-0.33, 2.28)$ with OLS standard errors, $(-1.32, 3.27)$ with Shanken standard errors and $(-0.87, 2.82)$ with GMM standard errors. The robust confidence set constructed using the CU-GMM objective function, $(-\infty, -4.92) \cup (1.78, +\infty)$, is much bigger and has little overlap with the other confidence sets. This is a classic indication of non-identification.

A scatter plot of expected returns against factor betas would provide insight into the role of the constant in the fit of the model, but constructing such a scatter plot is impossible for a three factor model. I consider, instead, a scatter plot of expected returns against SDF betas. Figure 4 shows a scatter plot of $\bar{R}_i^e$ against $\hat{\beta}_{mi}$, $i = 1, \ldots, n$, with m_t being constructed using the **b** vector corresponding to the calibrated model discussed in section I.E of LV's paper: $b_c = 6.74$, $b_d = 23.3$ and $b_r = 0.31$. Equation (17) implies that a scatter plot of $E(R_{it}^e)$ against the SDF betas, β_{mi}, should lie on a line through the origin with slope $\lambda_m = \sigma_m^2$. The constructed m_t series has sample variance $\hat{\sigma}_m^2 = 0.29$, so I indicate an estimated version of this theoretical line, $\beta_m \hat{\sigma}_m^2$, in bold black in Figure 4. This line correctly prices the risk free asset, since the intercept is zero. It also correctly prices SDF risk because an SDF mimicking portfolio has a beta of 1, and a risk premium σ_m^2. The theoretical line, however, does not correctly price LV's portfolio returns because the values of $(\hat{\beta}_{mi}, \bar{R}_i^e)$, indicated by the small circles in Figure 4, do not fit closely around the black line. Better fit can be obtained by running a regression of $\bar{R}_i^e$ on a constant and $\hat{\beta}_{mi}$. The regression line (indicated in bold grey) has an intercept of $\hat{\gamma}_m = -3.0$ percent and a slope of $\hat{\lambda}_m = 1.3$. This means that although the regression line fits the scatter plot reasonably well, it misprices the risk free rate by -3 percent. It also misprices SDF risk because an SDF mimicking portfolio has a beta of 1, and an implied expected return of $\hat{\gamma}_m + \hat{\lambda}_m = -1.7$, rather than the theoretically predicted expected return, $\hat{\sigma}_m^2 = 0.29$. When LV report high R^2 statistics for the calibrated model, it is because they measure fit with respect to the grey line, not the black line.

Figure 4 also sheds light on the degree to which we should be surprised by the good fit of LV's model with the constant. There are effectively three data points: the P1 and P7 observations, plus the cluster of points formed by P2–P6 and P8. It is not too surprising that a regression line with a free constant can do a good job of fitting three data points. When the SDF is estimated (as opposed to being calibrated) the problem becomes worse because, statistically, the P2–P6 and P8 portfolios have the same expected returns and betas as each other and we have four free parameters in γ and $\boldsymbol{\lambda}$.

By comparing Figure 5 (constructed using the estimate of $\mathbf{b}$ in Table 5a), with Figure 4 (constructed using the calibrated $\mathbf{b}$) we can also see what estimation of the model accomplishes. By re-weighting the factors in the SDF, estimation of the model with the constant forces the black line and the grey line to be parallel. That is, by construction, the regression line has slope $\hat{\lambda}_m = \hat{\sigma}_m^2$. Estimation of the model also, of course, unambiguously improves the fit of the model. The important point, again, is that LV report high R^2 statistics for the estimated model because they measure fit with respect to the grey line, not the black line.

A further problem with the estimated version of the model is that the implied value of $\hat{\sigma}_m$ is 2.4, while the mean of m_t is 1, by assumption. Consequently, in sample, the constructed m_t is negative in 38 percent of the observations. Furthermore, the implied structural parameters are theoretically implausible. In the original nonlinear model, the parameter values corresponding to LV's estimates of $\mathbf{b}$ imply that the marginal utility of nondurables is always negative, the intertemporal elasticity of substitution is negative, and the coefficient of relative risk aversion exceeds 100. This means that even if we accept the point estimates we still have an asset pricing puzzle.

4.3 Summary Discussion

When a constant is included in the model it fits well, as measured by the cross-sectional R^2 and the mean absolute pricing error, and it passes formal tests of the pricing errors. But this "success" is due to the convention which excludes the constant γ from the pricing errors. Furthermore, not one of the estimated factor risk premia, $\hat{\boldsymbol{\lambda}}$, and model parameters, $\hat{\mathbf{b}}$, is statistically significant using appropriate approaches to inference. Robust confidence sets for (γ_m, λ_m) indicate that the data are even less informative about the model than appropriate standard errors indicate.

If the constant is excluded from the model, it is still badly identified, and its fit deteriorates markedly. The R^2 is less than or equal to 0.34, while $\hat{\boldsymbol{\lambda}}$ and $\hat{\mathbf{b}}$ are statistically insignificant, with one exception. In two-stage and iterated GMM, some of the parameters associated with nondurable consumption growth and durables growth are statistically significant, but the R^2 of the model is negative and the mean absolute price error is large.

5 Robustness

In this section I consider the robustness of my findings to five different considerations. 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. Next, I check whether considering a different set of currency portfolios affects my conclusions. Third, I consider the post-Bretton Woods subsample of the data set. Fourth, I discuss LV's focus on the returns to a particular portfolio. Finally, I discuss additional evidence from higher frequency data in the post-Bretton Woods period.

5.1 Equity and Bond Portfolios

LV, in their reply, argue that the results in Section IV.C of their original paper 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.[7] Indeed, additional test assets could be useful because they might alleviate the identification problem alluded to in Section 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.

Of course, adding extra portfolios does not change the factor betas for the currency portfolios. The tests in Section 2 still apply. Adding extra portfolios changes the estimates of $\boldsymbol{\lambda}$, and therefore does change the estimates of the SDF betas. I discuss parameter estimates, below, for two different cases. In both cases, the SDF betas of the currency portfolios are jointly statistically insignificantly different from zero. Therefore, adding the equity and bond portfolios does not change my conclusion, based on betas alone, that we cannot reject the null that the predicted expected returns of the currency portfolios are all zero.

Adding the six Fama-French equity portfolios *slightly* alleviates the identification problem because equities have statistically significant betas with respect to the market return factor, r_W. However, the rank tests in Table 7 indicate that the $\boldsymbol{\beta}$ 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 Table 8. As the table indicates, 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 -0.82 (for iterated GMM). The mean absolute pricing error for the currency portfolios ranges from 1.43 (at the first stage of GMM) to 1.93 (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. Why do some currency portfolios (like P7) earn high returns, and why do 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 compared to the equity portfolios. That is not a puzzle, given that currency portfolios are only weakly correlated with risk factors that price equity portfolios. To see how this point affects the interpretation of R^2 statistics, consider the following example.

[7] These data are described in detail in an appendix available from the author.

Suppose we have a factor model that predicts that all six of the Fama-French portfolios have a mean return of 9.4 percent, and, because the factor is uncorrelated with the currency portfolios, predicts that all of them have a mean return of zero. This factor explains none of the cross-section of the Fama-French portfolios, nor does it explain the cross-section of the currency portfolios, but the R^2 across all assets is 0.82. If the factor explains the cross-section of the equities, but none of the cross-section of the currencies the R^2 across all assets is 0.94. 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.

Adding the five Fama bond portfolios leads to the results presented in Table 9. 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 -1.33 (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).

5.2 Different Currency Portfolios

In their reply, LV argue that they can explain the excess returns to the strategy of holding Pi and shorting P1, for $i = 2, \ldots, 8$. They claim that their paper is really about these seven "differenced" portfolios, which I refer to as D2, D3, . . . , 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:

$$\bar{R}_i^e = \hat{\gamma} + \hat{\boldsymbol{\beta}}_i' \hat{\boldsymbol{\lambda}} + \hat{u}_i, \qquad i = 1, \ldots, n. \tag{20}$$

This is equation (8) with γ and $\boldsymbol{\lambda}$ replaced by $\hat{\gamma}$ and $\hat{\boldsymbol{\lambda}}$ (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 (20), it follows that the sample mean of R_i^d is given by:

$$\bar{R}_i^d = (\hat{\boldsymbol{\beta}}_i - \boldsymbol{\beta}_1)' \hat{\boldsymbol{\lambda}} + \hat{u}_i - \hat{u}_1. \tag{21}$$

P1 is not just any asset. As can be seen in Figure 5, it happens to be an asset for which the SDF beta is roughly zero ($\boldsymbol{\beta}_1' \hat{\boldsymbol{\lambda}} \cong 0$) and the idiosyncratic pricing error is very small ($\hat{u}_1 \cong 0$). Thus

$$\bar{R}_i^d \cong \hat{\boldsymbol{\beta}}_i' \hat{\boldsymbol{\lambda}} + \hat{u}_i. \tag{22}$$

If the game is to fit the cross-sectional distribution of $\bar{R}_i^d$ it will make little difference whether the model includes a constant, or not.

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: $\boldsymbol{\beta}^{\mathrm{D}} = \boldsymbol{\Phi}\boldsymbol{\beta}^{\mathrm{P}}$ where $\boldsymbol{\Phi} = \begin{pmatrix} -\boldsymbol{\iota} & \mathbf{I}_{n-1} \end{pmatrix}$ and $\boldsymbol{\iota}$ is an $(n-1) \times 1$ unit vector. Not surprisingly, since we could not reject the null hypotheses that $\beta_{ij}^{\mathrm{P}} = 0$ for all i, and each j, we also cannot reject the null hypotheses that $\beta_{ij}^{\mathrm{D}} = \beta_{ij}^{\mathrm{P}} - \beta_{1j}^{\mathrm{P}} = 0$ for all i, and each j. 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. If we re-estimate the model using the D-portfolios, the SDF betas are also statistically insignificant.

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 $\boldsymbol{\beta}^{\mathrm{P}}\mathbf{x} \cong \mathbf{0}$, statistically. Given the definition of $\boldsymbol{\beta}^{\mathrm{D}}$ it follows that $\boldsymbol{\beta}^{\mathrm{D}}\mathbf{x} \cong \mathbf{0}$, statistically. In fact, the identification problem gets worse, because the transformation $\boldsymbol{\Phi}$ is not invertible. Any $\mathbf{x}$ such that $\boldsymbol{\beta}^{\mathrm{P}}\mathbf{x} \cong \mathbf{0}$ implies that $\boldsymbol{\beta}^{\mathrm{D}}\mathbf{x} \cong \mathbf{0}$. But there may be additional $\mathbf{x}$ such that $\boldsymbol{\beta}^{\mathrm{D}}\mathbf{x} \cong \mathbf{0}$ for which $\boldsymbol{\beta}^{\mathrm{P}}\mathbf{x} \cong \mathbf{0}$. This analysis is hardly surprising. Throwing away information is never likely to improve identification. Formal test statistics verifying this are provided in Table 7.[8]

5.3 Post-Bretton Woods Sample

LV present results for two sample periods, 1953–2002 and 1971–2002. The shorter sample may be more relevant given that the Bretton Woods system was in place until 1971. Choosing the shorter sample, however, does not greatly affect my conclusions.

As in Section 2 we can test whether for each factor the betas (across portfolios) are jointly significantly different from zero. The p-values associated with these tests are 0.22 for consumption growth, 0.12 for durables and 0.05 for the market return. While this suggests that there may be some amount of spread across the betas in the shorter sample, this does not translate into significant spread across the SDF betas. For the SDF associated with the two-pass regression, the p-value for a test that the SDF betas are all zero is 0.23. For the SDF associated with the 2nd-stage GMM estimate the p-value is 0.32. The model is still poorly identified, as the rank tests in Table 7 indicates. This is not too surprising given the extremely short sample. If the model is re-estimated without the constant over this period, the R^2 is never greater than 0.38 and the MAE is never less than 1.41.[9]

5.4 The D7 Portfolio

In their reply, LV emphasize a single portfolio, the D7 portfolio, which is formed by going long in P7 and short in P1. I agree with LV that if one considers the D7 portfolio in isolation, and if one considers individual factor betas, the consumption beta and durables beta are individually statistically significant. However, if the factor betas are estimated jointly (as they are in the first-pass regressions) the consumption beta and durables beta of the D7 portfolio are not individually statistically significant. Obviously this reflects the degree of

[8] Detailed empirical results using the D-portfolios are available upon request.

[9] Detailed empirical results using the post-Bretton Woods subsample are available upon request.

correlation between consumption growth and durables growth. Additionally, although the SDF betas of the D-portfolios are jointly insignificantly different from zero, the SDF beta of the D7 portfolio is statistically significant at the 5 percent level.

The appropriate statistics indicate that the average return of the D7 portfolio is positive, and its SDF beta is positive and statistically significant. But there are six other D-portfolios. Because they are formed by going short in P1 and because P1 has a significantly negative average return, several of these portfolios (D3, D4, D5, and, marginally, D6) also have a significantly positive average excess return. But their SDF betas are statistically insignificant. I conclude, from this evidence, that the cross-section of returns has not been explained.

5.5 Other Evidence from the Literature

My comment mainly discusses the conclusions we should draw from LV's evidence. Burnside, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo (2008) provide additional evidence by studying carry-trade portfolios formed monthly over the period 1976 to 2007. They show that the consumption, durables and market return betas of these carry-trade portfolios are statistically insignificant. When consumption-based models are estimated using carry-trade and equity portfolios as test assets the models, in many cases, cannot be formally rejected using the test of the pricing errors. But this, as in LV's case, reflects the weak identification problem stemming from the imprecise estimation of the betas. In almost every case, however, the pricing errors of the currency portfolios, alone, are statistically significant.

6 Conclusion

To explain cross-sectional variation in expected returns, a risk-based story requires that at least some of the returns be correlated with the risk factors. As the first-pass regressions reported in Section 2 demonstrate, however, LV's risk factors are very close to being uncorrelated with the returns they study. A symptom of this is that there is no statistically significant spread in the factor betas. 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 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, and 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. I conclude, however, that there are no grounds for drawing the sharp conclusion that consumption risk does explain the currency returns in LV's data set.

References

Burnside, Craig "Empirical Asset Pricing and Statistical Power in the Presence of Weak Risk Factors," NBER Working Paper No. 13357, August 2007.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo (2008) "Do Peso Problems Explain the Returns to the Carry Trade?" NBER Working Paper No. 14054, June 2008.

Cochrane, John H. *Asset Pricing*, Revised edition. Princeton: Princeton University Press, 2005.

Center for Research in Security Prices "CRSP Monthly Treasury US Database Guide," http://www.crsp.com/documentation/pdfs/monthly_treasury.pdf, 2007.

den Haan, Wouter J. and Andrew T. Levin "Robust Covariance Matrix Estimation with Data-Dependent VAR Prewhitening Order," NBER Technical Working Paper No. 255, June 2000.

Fama, Eugene and Kenneth R. French "Common Risk Factors in the Returns on Stocks and Bonds," February 1993, *Journal of Financial Economics*, *33*(1), pp. 3–56.

Fama, Eugene and James D. MacBeth "Risk Return and Equilibrium: Empirical Tests," May/June 1973, *Journal of Political Economy*, *81*(3), pp. 607–36.

Hansen, Lars P. "Large Sample Properties of Generalized Method of Moments Estimators," July 1982, *Econometrica*, *50*(4), pp. 1029–54.

Jagannathan, Ravi and Zhenyu Wang "An Asymptotic Theory for Estimating Beta-Pricing Models Using Cross-Sectional Regressions," August 1998, *Journal of Finance*, *53*(4), pp. 1285–309.

Kan, Raymond and Chu Zhang "GMM Tests of Stochastic Discount Factor Models with Useless Factors," October 1999a, *Journal of Financial Economics* *54*(1), pp. 103–27.

Kan, Raymond and Chu Zhang "Two-Pass Tests of Asset Pricing Models with Useless Factors," February 1999b, *Journal of Finance*, *54*(1), pp. 203–35.

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

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

Stock, James H. and Jonathan H. Wright "GMM with Weak Identification," September 2000, *Econometrica*, *68*(5), pp. 1055–96.

Stock, James H., Jonathan H. Wright, and Motohiro Yogo "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," October 2002, *Journal of Business and Economic Statistics*, *20*(4), pp. 518–29.

Wright, Jonathan H. "Detecting Lack of Identification in GMM," April 2003, *Econometric Theory*, *19*(2), pp. 322–30.

Yogo, Motohiro "Estimating the Elasticity of Intertemporal Substitution when Instruments are Weak," August 2004, *Review of Economics and Statistics*, *86*(3), pp. 797–810.

Yogo, Motohiro "A Consumption-Based Explanation of Expected Stock Returns," April 2006, *Journal of Finance*, *61*(2), pp. 539–80.

TABLE 1: First-Pass Estimates of the Betas

Portfolio (i)	$\bar{R}^e_i$	Factor (j) Δc	Δd	r_W	Test of $\beta_{ij} = 0\ \forall j$
P1	−2.336	0.201	0.028	−0.068	
	(0.897)	(0.852)	(0.612)	(0.055)	(0.600)
	[0.906]	[0.643]	[0.563]	[0.049]	[0.483]
P2	−0.873	0.740	0.091	−0.034	
	(0.937)	(0.889)	(0.638)	(0.058)	(0.579)
	[0.947]	[0.620]	[0.466]	[0.062]	[0.157]
P3	−0.747	−0.639	0.962	0.019	
	(0.935)	(0.882)	(0.633)	(0.057)	(0.464)
	[1.361]	[1.026]	[0.682]	[0.047]	[0.513]
P4	0.329	−0.546	0.982	−0.089	
	(1.190)	(1.095)	(0.786)	(0.071)	(0.156)
	[1.202]	[1.075]	[0.760]	[0.068]	[0.068]
P5	−0.151	0.180	0.485	0.009	
	(1.053)	(1.006)	(0.722)	(0.065)	(0.740)
	[1.063]	[0.754]	[0.708]	[0.065]	[0.740]
P6	−0.213	−0.755	1.079	0.023	
	(1.148)	(1.089)	(0.781)	(0.071)	(0.556)
	[1.160]	[0.958]	[0.830]	[0.068]	[0.592]
P7	2.988	0.036	1.234	−0.027	
	(1.144)	(1.044)	(0.749)	(0.068)	(0.101)
	[1.155]	[0.797]	[0.755]	[0.062]	[0.129]
P8	2.031	−1.342	1.426	0.079	
	(1.756)	(1.674)	(1.201)	(0.108)	(0.684)
	[2.586]	[1.646]	[0.762]	[0.116]	[0.251]
Test of		(0.813)	(0.623)	(0.365)	
$\beta_{ij} = 0\ \forall i$		[0.799]	[0.668]	[0.511]	

Notes: Annual data, 1953–2002. The regression equation is $R^e_{it} = a_i + \mathbf{f}_t' \boldsymbol{\beta}_i + \epsilon_{it}$, where R^e_{it} 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 β_{ij} and the sample mean of each portfolio return, $\bar{R}^e_i$. The table also reports p-values for tests of the hypotheses that $\beta_{ij} = 0\ \forall j$ for each portfolio i, and $\beta_{ij} = 0\ \forall i$ for each factor j. For estimates of betas, OLS standard errors are in parentheses and GMM-VARHAC standard errors are in square brackets. For test statistics, corresponding p-values are presented.

TABLE 2: ESTIMATES OF THE SDF BETAS FOR SPECIFIC VALUES OF **b**

Portfolio	Estimates of β_{im}		
(i)	Model (a)	Model (b)	Model (c)
P1	0.059	−0.030	0.638
	(0.373)	(0.452)	(1.665)
	[0.544]	[0.658]	[2.594]
P2	0.335	0.437	1.925
	(0.387)	(0.468)	(1.720)
	[0.241]	[0.285]	[1.064]
P3	0.549	0.513	2.229
	(0.381)	(0.465)	(1.709)
	[0.542]	[0.526]	[2.617]
P4	0.539	0.308	2.692
	(0.489)	(0.597)	(2.178)
	[0.702]	[1.031]	[2.288]
P5	0.467	0.559	2.158
	(0.433)	(0.524)	(1.932)
	[0.470]	[0.553]	[2.031]
P6	0.603	0.556	2.422
	(0.470)	(0.573)	(2.106)
	[0.705]	[0.684]	[3.233]
P7	0.999	1.038	4.714
	(0.454)	(0.557)	(2.018)
	[0.465]	[0.566]	[2.043]
P8	0.694	0.653	2.403
	(0.723)	(0.879)	(3.245)
	[0.696]	[0.844]	[3.481]
Hypothesis Tests			
$\beta_{im} = 0\ \forall i$	(0.469)	(0.443)	(0.505)
	[0.306]	[0.170]	(0.329)
$\beta_{im} = \beta_m\ \forall i$	(0.508)	(0.477)	(0.575)
	[0.688]	[0.589]	[0.602]

Notes: Annual data, 1953–2002. The regression equation is $R^e_{it} = a_i + m_t\beta_{im} + \epsilon_{it}$, where R^e_{it} is the excess return of portfolio i at time t, $m_t = 1 - (\mathbf{f}_t - \bar{\mathbf{f}})'\mathbf{b}$, $\mathbf{f}_t = (\ \Delta c_t \quad \Delta d_t \quad r_{Wt}\ )'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return, $\bar{\mathbf{f}}$ is the sample mean of $\mathbf{f}_t$ and the vector $\mathbf{b}$ takes on one of the following three values: (a) $\mathbf{b} = (\ -21.0 \quad 129.9 \quad 4.46\ )'$, (b) $\mathbf{b} = (\ 37.0 \quad 74.7 \quad 4.65\ )'$ and (c) $\mathbf{b} = (\ 6.74 \quad 23.3 \quad 0.31\ )'$. The portfolios are described in Table 1. This table reports estimates of β_{im} and p-values for tests of restrictions on the β_{im}s. For estimates of betas, OLS standard errors are in parentheses and GMM-VARHAC standard errors are in square brackets. For test statistics, corresponding p-values are presented.

TABLE 3: TESTS OF THE RANK OF THE FACTOR BETA MATRIX

Test of H_0: rank($\boldsymbol{\beta}$) $= r$		Test of H_0: rank($\boldsymbol{\beta}^+$) $= r+1$	
r	p-value	$r+1$	p-value
2	(0.615) [0.657]	3	(0.529) [0.699]
1	(0.643) [0.444]	2	(0.618) [0.326]
0	(0.639) [0.000]	1	(0.562) [0.003]

Notes: Annual data, 1953–2002. The matrix $\boldsymbol{\beta}$ is obtained by running the regressions described in Table 1. The matrix $\boldsymbol{\beta}$ must have full column rank (3) for $\boldsymbol{\lambda}$ to be identified in the two-pass procedure without the constant. The matrix $\boldsymbol{\beta}^+ = (\boldsymbol{\iota} \quad \boldsymbol{\beta})$ must have full column rank (4) for γ and $\boldsymbol{\lambda}$ to be identified in the two-pass procedure with the constant. The table presents tests of the null hypothesis that these rank conditions fail. The p-values for the tests are presented in parentheses (OLS standard errors) and square brackets (GMM-VARHAC standard errors).

TABLE 4: Second-Pass Estimates of the Factor Risk Premia

	(a) Model With a Constant				(b) Model Without a Constant			
		Standard Error of $\hat{\boldsymbol{\lambda}}$				Standard Error of $\hat{\boldsymbol{\lambda}}$		
Factor	$\hat{\boldsymbol{\lambda}}$	OLS	Shanken	GMM	$\hat{\boldsymbol{\lambda}}$	OLS	Shanken	GMM
Constant (γ)	−2.94	(0.86)	[2.23]	{2.66}				
Δc	2.19	(0.83)	[2.11]	{2.48}	0.59	(0.73)	[1.01]	{1.17}
Δd	4.70	(0.97)	[2.42]	{2.41}	1.10	(1.02)	[1.40]	{1.69}
r_W	3.33	(7.59)	[18.8]	{23.1}	11.7	(7.40)	[10.1]	{10.6}
R^2		0.87				0.34		
MAE		0.44				1.17		

	P-value for $C_{\hat{\mathbf{u}}}$			P-value for $C_{\hat{\boldsymbol{\alpha}}}$		
	OLS	Shanken	GMM	OLS	Shanken	GMM
Pricing Error Test	0.483	0.972	0.994	0.001	0.059	0.173

Notes: Annual data, 1953–2002. Part (a) reports results from running the cross-sectional regression $\bar{R}_i^e = \gamma + \hat{\boldsymbol{\beta}}_i' \boldsymbol{\lambda} + u_i$.where $\bar{R}_i^e$ is the mean excess return of portfolio i and $\hat{\boldsymbol{\beta}}_i$ is the vector of factor betas of portfolio i estimated in the first pass regression. The portfolios are described in Table 1. Part (b) reports results from the a cross-sectional regression without the constant: $\bar{R}_i^e = \hat{\boldsymbol{\beta}}_i' \boldsymbol{\lambda} + \alpha_i$. For the factor risk premia ($\hat{\boldsymbol{\lambda}}$) OLS standard errors are in parentheses, Shanken standard errors are in square brackets, and GMM-VARHAC standard errors are in braces. For the tests of the pricing errors I compute the test statistic for each of the three methods of computing the covariance matrix of $\hat{\mathbf{u}}$ or $\hat{\boldsymbol{\alpha}}$, 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 5: GMM ESTIMATES OF THE MODEL WITH THE CONSTANT

	(a) 1st Stage		(b) 2nd Stage		(c) Iterated GMM	
Factor	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$
Constant (γ)		−2.94		−3.02		−3.56
		(2.63)		(2.23)		(2.01)
Δc	−21.0	2.19	27.2	2.73	59.5	2.94
	(87.7)	(1.88)	(76.2)	(2.81)	(74.2)	(1.59)
Δd	129.9	4.70	108.3	4.85	88.4	4.81
	(97.4)	(3.16)	(97.3)	(3.15)	(76.4)	(2.62)
r_W	4.46	3.33	2.53	0.78	0.15	−3.87
	(4.83)	(13.0)	(4.38)	(10.9)	(4.03)	(11.8)
R^2	0.87		0.81		0.50	
MAE	0.44		0.42		0.66	
Pricing Error Test			0.703		0.915	

Notes: Annual data, 1953–2002. The table reports GMM estimates of $\mathbf{b}$, γ and $\boldsymbol{\lambda}$ 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 a vector of currency portfolio returns described in Table 1, $\mathbf{f}_t=(\ \Delta c_t \quad \Delta d_t \quad r_{Wt}\ )'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return. GMM-VARHAC standard errors are in parentheses. For the test of the pricing errors I report the p-value associated with the test-statistic.

TABLE 6: GMM Estimates of the Model with no Constant

	(a) 1st Stage		(b) 2nd Stage		(c) Iterated GMM	
Factor	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$
Δc	−22.0	0.59	37.0	2.37	25.8	1.72
	(62.9)	(1.07)	(45.2)	(1.00)	(40.6)	(0.93)
Δd	45.5	1.10	74.7	3.48	66.0	3.41
	(50.3)	(1.64)	(33.0)	(1.13)	(47.4)	(1.52)
r_W	5.16	11.7	4.65	10.2	−2.14	−10.6
	(2.88)	(8.26)	(2.70)	(7.37)	(2.94)	(8.80)
R^2	0.34		−0.66		−1.45	
MAE	1.17		1.86		2.28	
Pricing Error Test			0.068		0.281	

Notes: Annual data, 1953–2002. The table reports GMM estimates of $\mathbf{b}$ and $\boldsymbol{\lambda}$ 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 a vector of currency portfolio returns described in Table 1, $\mathbf{f}_t=(\ \Delta c_t \quad \Delta d_t \quad r_{Wt}\ )'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return. GMM-VARHAC standard errors are in parentheses. For the test of the pricing errors I report the p-value associated with the test-statistic. The appendix provides details of the weighting matrices at each stage, and explains the equivalence of the GMM approach to the two-pass method. It also explains why the test of the pricing errors is the same at both stages of GMM.

TABLE 7: Tests of the Rank of the Factor Beta Matrix with Different Test Assets

Tests of H_0: rank($\boldsymbol{\beta}$) $= r$					Tests of H_0: rank($\boldsymbol{\beta}^+$) $= r+1$				
r	Currencies & Equities	Currencies, Equities & Bonds	Differenced Currencies	Currencies 1971-2002	$r+1$	Currencies & Equities	Currencies, Equities & Bonds	Differenced Currencies	Currencies 1971-2002
2	(0.440) [0.324]	(0.751) –	(0.529) [0.706]	(0.528) –	3	(0.476) [0.376]	(0.757) –	(0.406) [0.626]	(0.438) –
1	(0.435) [0.000]	(0.012) –	(0.618) [0.134]	(0.250)	2	(0.519) [0.000]	(0.168) –	(0.612) [0.339]	(0.228) –
0	(0.000) [0.000]	(0.000) –	(0.562) [0.001]	(0.108) –	1	(0.000) [0.000]	(0.000) –	(0.603) [0.005]	(0.088) –

Notes: See the note to Table 3. P-values for rank tests are presented in parentheses (OLS standard errors) and square brackets (GMM-VARHAC standard errors). The currency portfolios are described in Table 1. The equity portfolios are described in Table 8. The bond portfolios are described in Table 9. The differenced currency portfolios are described in the main text. The currency portfolios in the post-Bretton Woods period (1971–2002) are also considered. Dashes indicate cases where the GMM-VARHAC standard errors cannot be calculated because the number of parameters in the $\boldsymbol{\beta}$ matrix exceeds the sample size.

TABLE 8: GMM ESTIMATES OF THE MODEL WITH NO CONSTANT CURRENCY AND EQUITY PORTFOLIOS AS TEST ASSETS

	(a) 1st Stage		(b) 2nd Stage		(c) Iterated GMM	
Factor	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$
Δc	112.6	2.37	83.9	2.59	56.5	2.50
	(79.7)	(1.22)	(54.2)	(0.62)	(50.7)	(0.99)
Δd	−4.39	1.80	35.0	2.74	63.8	3.55
	(65.6)	(2.04)	(28.1)	(0.89)	(46.8)	(1.64)
r_W	2.09	11.0	4.24	13.6	2.62	5.47
	(3.34)	(3.82)	(1.70)	(6.2)	(2.16)	(4.91)
R^2	0.03		−0.09		−0.82	
MAE	1.43		1.39		1.93	

Notes: Annual data, 1953–2002. The table reports GMM estimates of $\mathbf{b}$ and $\boldsymbol{\lambda}$ 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 a vector of excess returns that includes the currency portfolios described in Table 1 as well as Fama and French's (1993) six equity portfolios created by sorting stocks on the basis of size and value, $\mathbf{f}_t = (\ \Delta c_t \quad \Delta d_t \quad r_{Wt}\ )'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return. GMM-VARHAC standard errors are in parentheses. The R^2 statistic and mean absolute pricing error (MAE) are presented for currency portfolios only, and are comparable to the statistics in Table 6.

TABLE 9: GMM Estimates of the Model with no Constant
Currency, Equity and Bond Portfolios as Test Assets

	(a) 1st Stage		(b) 2nd Stage		(c) Iterated GMM	
Factor	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$	$\hat{\mathbf{b}}$	$\hat{\boldsymbol{\lambda}}$
Δc	88.0	1.65	91.5	1.79	14.3	−0.020
	(59.3)	(1.06)	(32.5)	(0.45)	(24.0)	(0.82)
Δd	−14.6	0.92	−13.0	0.97	−28.6	−1.36
	(52.1)	(1.74)	(20.6)	(0.67)	(23.5)	(0.75)
r_W	1.90	10.3	3.00	13.7	5.85	20.7
	(2.61)	(3.19)	(1.24)	(5.5)	(1.50)	(10.6)
R^2	0.03		0.04		−1.33	
MAE	1.40		1.42		1.64	

Notes: Annual data, 1953–2002. The table reports GMM estimates of $\mathbf{b}$ and $\boldsymbol{\lambda}$ obtained by exploiting the moment restrictions $E\{\mathbf{R}^e_t[1-(\mathbf{f}_t-\boldsymbol{\mu})'\,\mathbf{b}]\}=0$, $E(\mathbf{f}_t-\boldsymbol{\mu})=\mathbf{0}$ and $E[(\mathbf{f}_t-\boldsymbol{\mu})(\mathbf{f}_t-\boldsymbol{\mu})'-\boldsymbol{\Sigma}_f]=0$, where $\mathbf{R}^e_t$ is a vector of excess returns that includes the currency portfolios described in Table 1, the equity portfolios described in Table 8, and five Fama bonds portfolios sorted by maturity (from the Center for Research in Securities Prices, 2007), $\mathbf{f}_t = (\ \Delta c_t \quad \Delta d_t \quad r_{Wt}\ )'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return. GMM-VARHAC standard errors are in parentheses. The R^2 statistic and mean absolute pricing error (MAE) are presented for currency portfolios only, and are comparable to the statistics in Table 6.

FIGURE 1: SDF Betas, Expected Returns and Two Standard Error Bars for the Benchmark Estimated Model

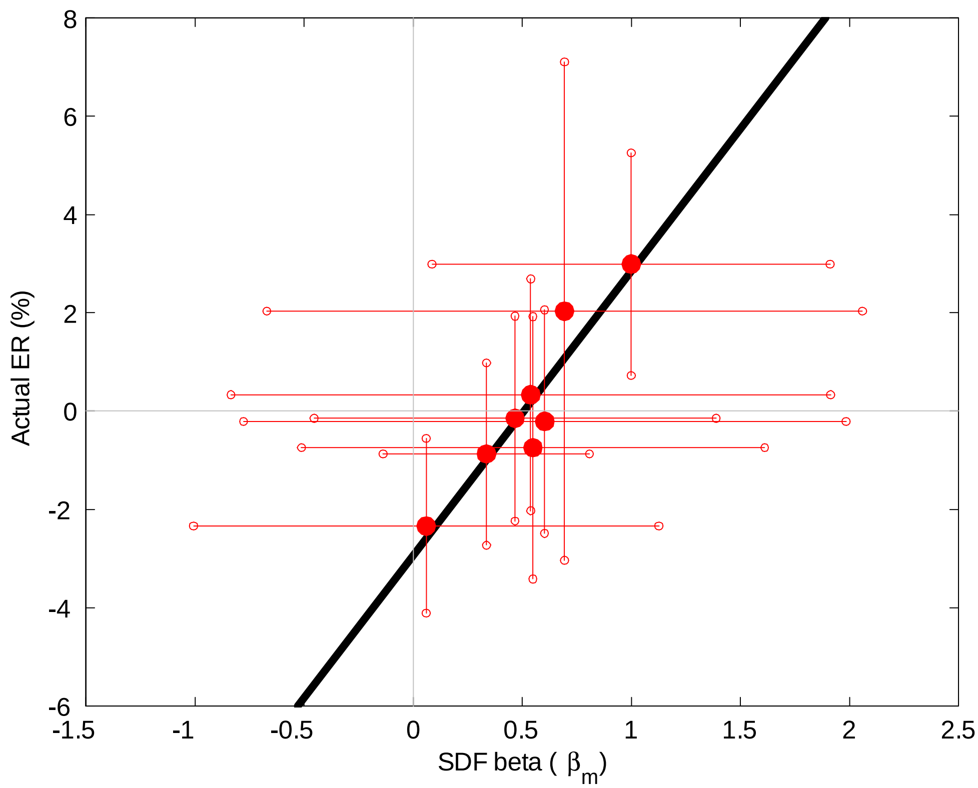

Note: Filled circles represent (SDF beta, mean excess return) pairs for LV's eight currency portfolios, described in Table 1. The "SDF beta", $\hat{\beta}_{mi}$, for portfolio i is the slope coefficient from a regression of the portfolio excess return, R^e_{it}, on the SDF, $m_t = 1 - \left(\mathbf{f}_t - \bar{\mathbf{f}}\right)' \mathbf{b}$, where $\mathbf{f}_t = (\ \Delta c_t \quad \Delta d_t \quad r_{Wt}\ )'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return, $\bar{\mathbf{f}}$ is the sample mean of $\mathbf{f}_t$ and $\mathbf{b}$ corresponds to LV's two pass estimates of $\boldsymbol{\lambda}$: $b_c = -21.0$, $b_d = 130$ and $b_r = 4.46$. The horizontal lines at each circle are two standard error bands around each $\hat{\beta}_{mi}$. "Actual ER" is the sample mean of the portfolio return, $\bar{R}^e_i$. The vertical lines are two standard error bands around $\bar{R}^e_i$. The black line corresponds to $\beta_m \hat{\sigma}^2_m$ where $\hat{\sigma}^2_m$ is the variance of the constructed SDF. The grey line is the estimated regression line $\bar{R}^e_i = \hat{\gamma}_m + \hat{\lambda}_m \beta_{mi}$.

FIGURE 2: Robust Confidence Sets for the Constant and the Price of SDF Risk

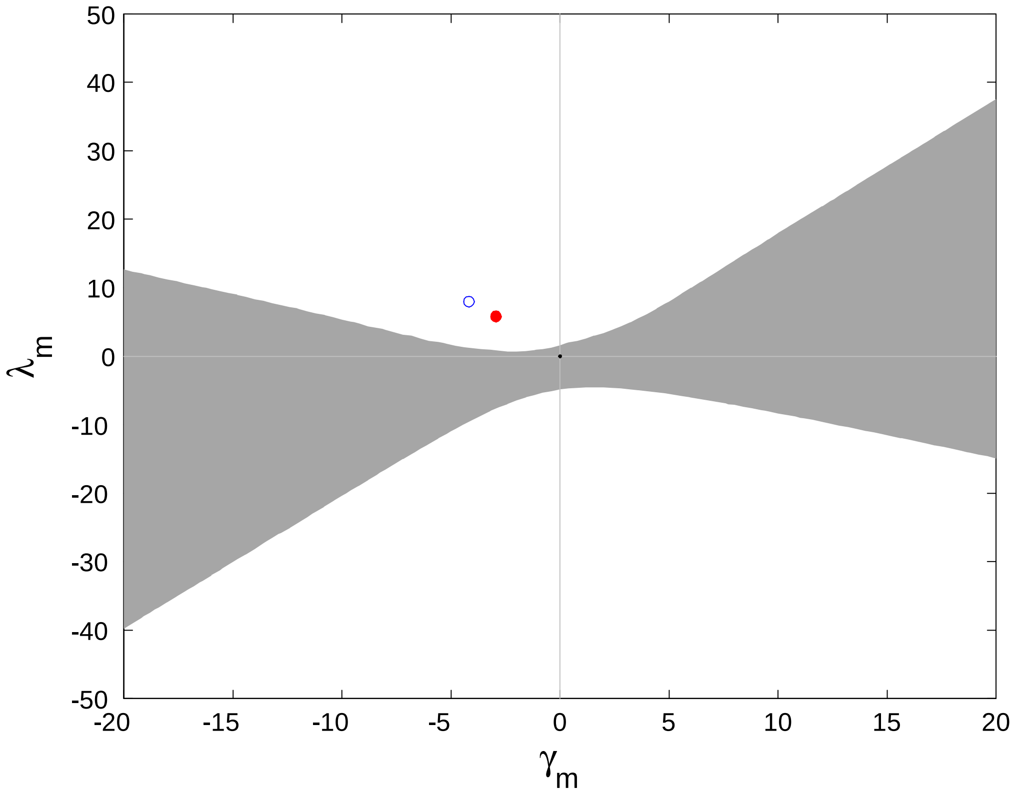

Note: The white area in the figure represents a robust 95 percent confidence set for (γ_m, λ_m) when the SDF is specified by $m_t = 1 - \left(\mathbf{f}_t - \bar{\mathbf{f}}\right)' \mathbf{b}$, $\mathbf{f}_t = \begin{pmatrix} \Delta c_t & \Delta d_t & r_{Wt} \end{pmatrix}'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return, $\bar{\mathbf{f}}$ is the sample mean of $\mathbf{f}_t$ and $\mathbf{b}$ corresponds to LV's two pass estimates of $\boldsymbol{\lambda}$: $b_c = -21.0$, $b_d = 130$ and $b_r = 4.46$. The confidence set is constructed using the CU-GMM objective function. The filled dot in the figure is the two-pass estimate of (γ_m, λ_m). The open dot in the figure is the CU-GMM estimate of (γ_m, λ_m).

FIGURE 3: Various Confidence Sets for the Constant and the Price of SDF Risk

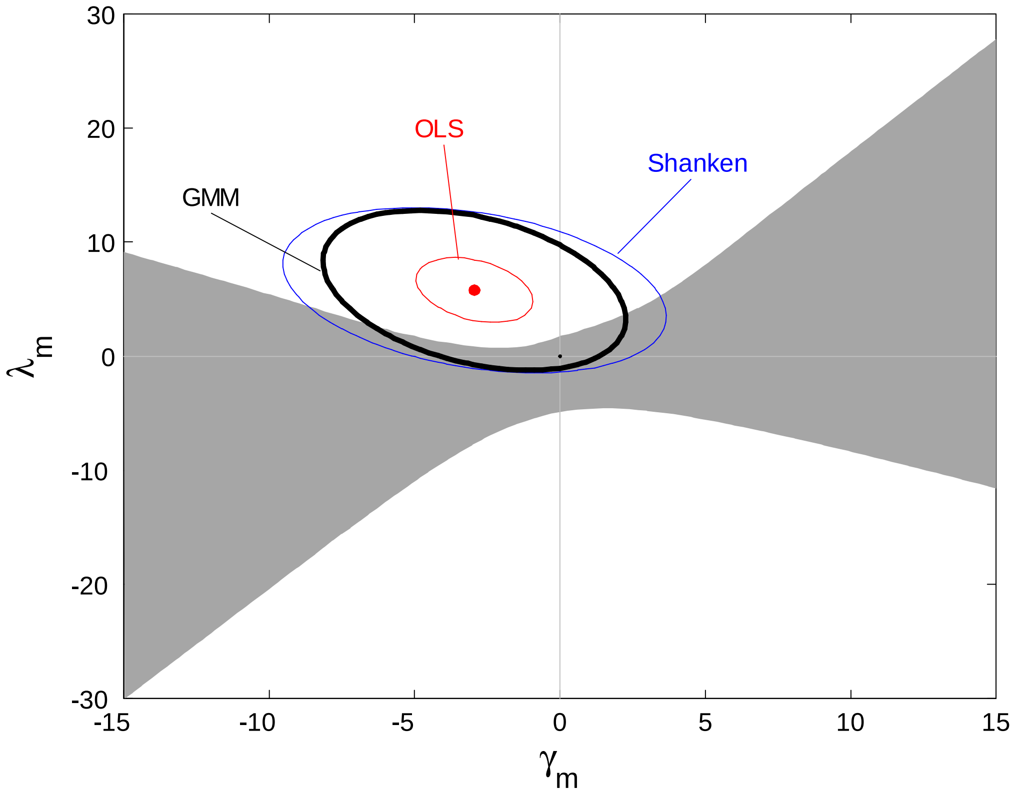

Note: The white area in the figure represents a robust 95 percent confidence set for (γ_m, λ_m) when the SDF is specified by $m_t = 1 - \left(\mathbf{f}_t - \bar{\mathbf{f}}\right)' \mathbf{b}$, $\mathbf{f}_t = (\ \Delta c_t \quad \Delta d_t \quad r_{Wt}\ )'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return, $\bar{\mathbf{f}}$ is the sample mean of $\mathbf{f}_t$ and $\mathbf{b}$ corresponds to LV's two pass estimates of $\boldsymbol{\lambda}$: $b_c = -21.0$, $b_d = 130$ and $b_r = 4.46$. The confidence set is constructed using the CU-GMM objective function. The filled dot in the figure is the two-pass estimate of (γ_m, λ_m). The 95 percent confidence sets based on asymptotic standard errors are indicated by their respective labels, OLS, Shanken, and GMM.

FIGURE 4: SDF Betas and Expected Returns for the Calibrated Model

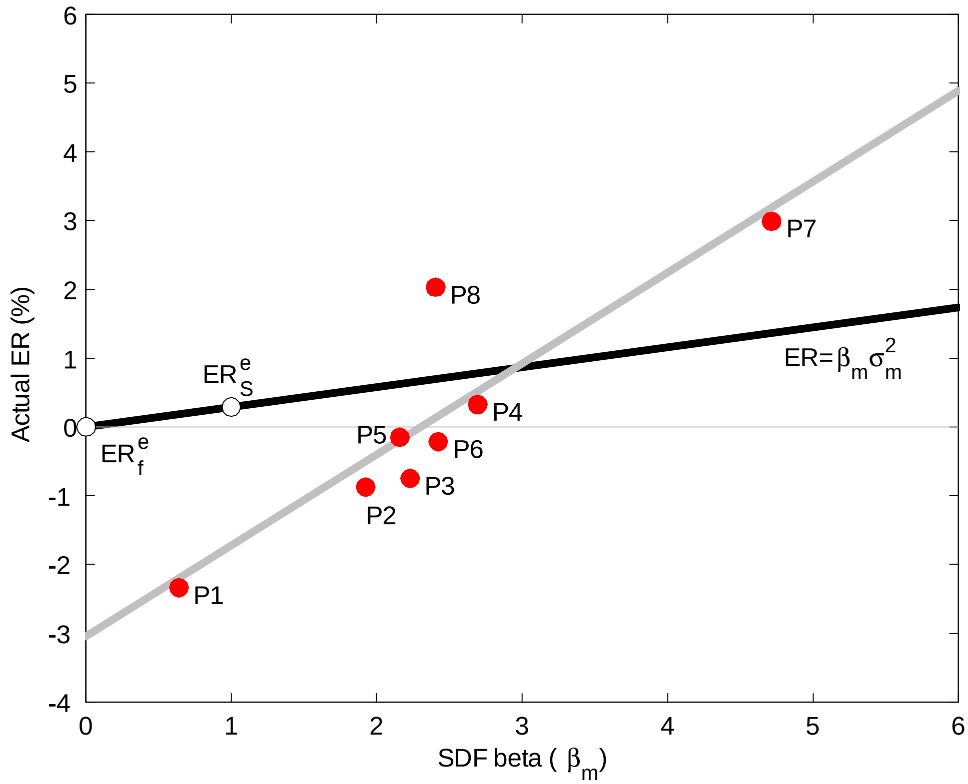

Note: Filled circles represent (SDF beta, mean excess return) pairs for LV's eight currency portfolios described in Table 1. The "SDF beta", $\hat{\beta}_{mi}$, for portfolio i is the slope coefficient from a regression of the portfolio excess return, R^e_{it}, on the SDF, $m_t = 1 - \left(\mathbf{f}_t - \bar{\mathbf{f}}\right)' \mathbf{b}$, where $\mathbf{f}_t = (\; \Delta c_t \quad \Delta d_t \quad r_{Wt} \;)'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return, $\bar{\mathbf{f}}$ is the sample mean of $\mathbf{f}_t$ and $\mathbf{b}$ corresponds to the calibrated model with $b_c = 6.74$, $b_d = 23.3$ and $b_r = 0.31$. "Actual ER" is the sample mean of the portfolio return, $\bar{R}^e_i$. The black line corresponds to $\beta_m \hat{\sigma}^2_m$ where $\hat{\sigma}^2_m$ is the variance of the constructed SDF. The grey line is the estimated regression line $\bar{R}^e_i = \hat{\gamma}_m + \hat{\lambda}_m \beta_{mi}$. The empty circle marked ER^e_f signifies that a risk free asset has a zero beta, and a zero excess return. The empty circle marked ER^e_S signifies that an SDF mimicking portfolio has a beta of 1 and expected excess return of σ^2_m.

FIGURE 5: SDF Betas and Expected Returns for the Benchmark Estimated Model

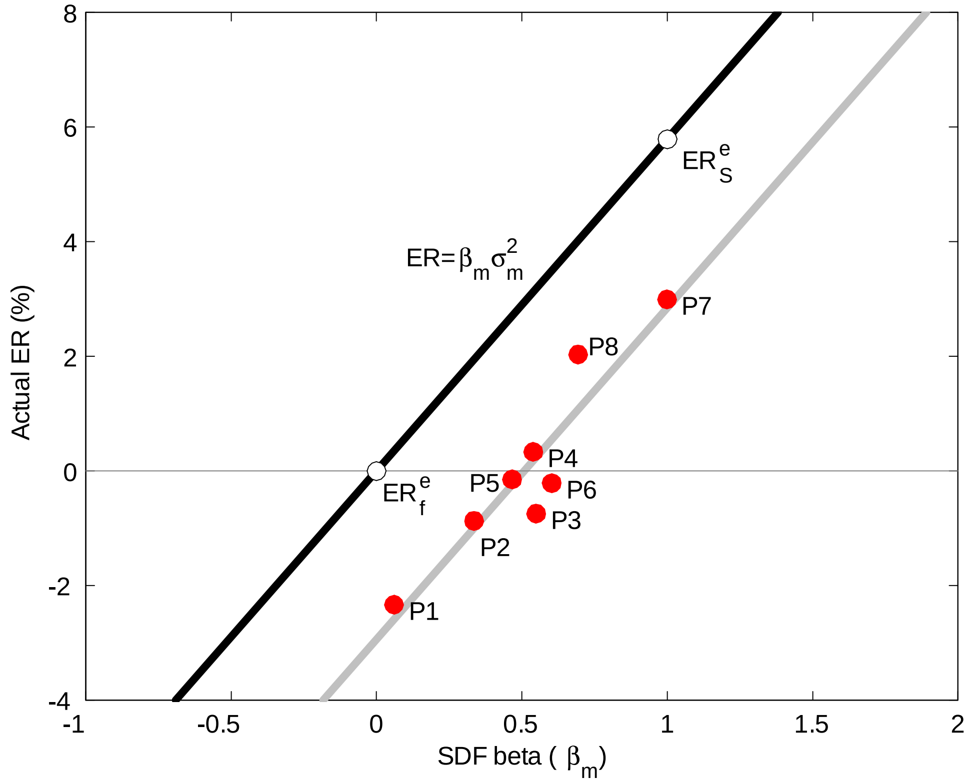

Note: Filled circles represent (SDF beta, mean excess return) pairs for LV's eight portfolios described in Table 1. The "SDF beta", $\hat{\beta}_{mi}$, for portfolio i is the slope coefficient from a regression of the portfolio excess return, R^e_{it}, on the SDF, $m_t = 1 - \left(\mathbf{f}_t - \bar{\mathbf{f}}\right)' \mathbf{b}$, where $\mathbf{f}_t = (\ \Delta c_t \quad \Delta d_t \quad r_{Wt}\ )'$, Δc is real per household consumption (nondurables & services) growth, Δd is real per household durable consumption growth, r_W is the value weighted US stock market return, $\bar{\mathbf{f}}$ is the sample mean of $\mathbf{f}_t$ and $\mathbf{b}$ corresponds to LV's two pass estimates of $\boldsymbol{\lambda}$: $b_c = -21.0$, $b_d = 130$ and $b_r = 4.46$. "Actual ER" is the sample mean of the portfolio return, $\bar{R}^e_i$. The black line corresponds to $\beta_m \hat{\sigma}^2_m$ where $\hat{\sigma}^2_m$ is the variance of the constructed SDF. The grey line is the estimated regression line $\bar{R}^e_i = \hat{\gamma}_m + \hat{\lambda}_m \beta_{mi}$. The empty circle marked ER^e_f signifies that a risk free asset has a zero beta, and a zero excess return. The empty circle marked ER^e_S signifies that an SDF mimicking portfolio has a beta of 1 and expected excess return of σ^2_m.