Institute for Empirical Research in Economics
University of Zurich

Working Paper Series
ISSN 1424-0459

Working Paper No. 331

The Consumption - Real Exchange Rate Anomaly: an Asset Pricing Perspective

Mathias Hoffmann and Thomas Nitschka

September 2007

The Consumption - Real Exchange Rate Anomaly: an Asset Pricing Perspective[1]

Mathias Hoffmann[2]
University of Zurich, Institute for Empirical Research in Economics and CESifo

Thomas Nitschka[3]
University of Zurich, Institute for Empirical Research in Economics

This version: September 2007
http://www.iew.uzh.ch/itf

[1]The authors have benefited from comments by seminar participants at UZH and at the MAREM 2007. In particular, we would like to thank Bent Sørensen, Josef Falkinger and our discussant, Christian Conrad.

[2]e-mail: mathias.hoffmann@iew.unizh.ch, phone: ++41-(0)44-63-45258, address: Institute for Empirical Research in Economics, Chair of International Trade and Finance, Zuerichbergstrasse 14, CH-8032 Zürich

[3]e-mail: thomas.nitschka@iew.unizh.ch

Abstract

Idiosyncratic consumption risk explains more than 60 percent of the cross-sectional variation in quarterly exchange rate changes and currency returns. Our results are obtained from data of 13 industrialized countries and are based on an international version of the consumption capital asset pricing model (CCAPM) in which we account for international consumption heterogeneity. We use this framework to dissect the consumption-exchange rate anomaly, the empirical fact that international variation in purchasing power alone does not appear to account for differences in consumption growth rates across countries. As an explanation for this phenomenon, we explore the presence of currency risk premia that also lead to departures from uncovered interest parity (UIP). We decompose the cross-sectional variation in consumption into one component that is due to cross-country differences in inflation rates and a second component that is due to international variation in nominal interest rates. We interpret these factors as indicators of goods and financial market segmentation respectively. We find that both help account to virtually equal parts for the cross-section of exchange rate changes. Interestingly, the price of aggregate consumption risk has declined over the 1990s, in line with a growing literature that documents a growing internationalisation of country portfolios over this period.

JEL classification: E21, F30, G12

Keywords: uncovered interest rate parity, consumption CAPM, international financial integration, consumption risk sharing

1 Introduction

According to the consumption capital asset pricing model (CCAPM), currencies – like any asset – should be priced for the exposure to world aggregate consumption growth risk they deliver. Lustig and Verdelhan (2007) have just recently argued that a version of the CCAPM can account for departures from uncovered interest parity (UIP). High interest rate currencies therefore depreciate less than is implied by UIP since they expose the investor to lots of aggregate consumption growth risk. Conversely, low interest rate currencies depreciate by more than is implied by UIP because they provide a hedge against aggregate consumption growth risk.

As emphasized by Brandt et al. (2006), exchange rate changes should reflect relative intertemporal marginal rates of substitution (IMRS). In a consumption-based model this implies that countries with relatively high consumption growth will tend to have depreciating currencies, those with low consumption growth appreciating ones.

The international dispersion of consumption growth rates should therefore be a key determinant of the cross-section of exchange rate changes. In particular, if consumption growth rates differ across countries – and low international consumption correlations are a salient feature of the data (see Backus, Kehoe and Kydland (1992)) – this heterogeneity should be reflected in standard pricing relations for currency risk. This insight forms our point of departure for this paper.

Based on a CCAPM framework augmented to include international consumption dispersion as an additional pricing factor, we find that idiosyncratic consumption risk explains more than 60 percent of the cross-sectional variation in exchange rate changes. These results are obtained from quarterly data of 13 industrialized countries for the period 1971-2003.

Our analysis places itself between two strands of the recent literature. First, a number of recent papers has employed empirical asset pricing models with heterogeneous consumption.[1] However, to our knowledge, Sarkissian (2003) is to date the only study that applies such a model to price currency returns. Here, we employ a framework similar to Sarkissian's to dissect the mechanism to which a second strand of recent papers, notably Lustig and Verdelhan (2007)[2] and Brandt et al. (2006), have drawn the profession's attention and that emphasizes the role of variation in relative intertemporal marginal rates of substitution for the determination of exchange rates:

In a world with complete financial markets, frictionless trade in goods and homothetic preferences separable in consumption and leisure, consumption growth rates should equalize and consumption dispersion is zero. Introducing non-traded goods in a model with complete financial markets, Backus and Smith (1993) show that even under complete financial markets, deviations from purchasing power parity can drive a wedge between consumption growth rates so that relative consumption growth should move in lockstep with (real) exchange rate changes. This theoretical link between consumption and real exchange rates has so far mainly been examined in the time-series dimension (Backus and Smith (1993), Kollmann (1995) and more recently Ravn (2001)) where it has found virtually no empirical support. Here, we argue that the use of an asset pricing model in which consumption growth rates differ across countries is a natural starting point to investigate the consumption-exchange rate anomaly further.

Once we allow for incompleteness of financial markets, the Backus-Smith

[1] Constantinides and Duffie (1996) was the first paper to demonstrate theoretically how idiosyncratic consumption risk, captured through consumption dispersion, can come to figure in aggregate pricing relations when markets are incomplete. This seminal paper has motivated a number of empirical studies, including Brav et al. (2002), Jacobs and Wang (2004) and Sarkissian (2003).

[2] We will often refer to this latter paper as 'LV'.

relation should still hold in expectation, but – as we show – only to the extent that uncovered interest parity holds. In theory, while consumption and real exchange rates do not have to be perfectly correlated if financial markets are incomplete, departures from UIP and departures from the Backus-Smith condition should still be highly correlated in the cross-section, though not necessarily for individual countries over time. Our empirical results suggest that this is the case, which may help explain why earlier studies that have focused on the time series dimension of the data have not identified this relation.

Our results support the view that the cross-section of international differences in consumption growth rates should be related to the cross-section of real interest rate differentials. This in turn suggests that consumption heterogeneity and its explanatory power for exchange rate changes can have two conceptually different sources:

The first source is heterogeneity in *nominal* interest rates. The importance of the international spread of nominal interest rates has recently been emphasized by LV who find that forming currency portfolios by the size of the interest rate differential is key for the successful application of the CCAPM to currency returns. We argue that using portfolios sorted by interest rate differentials as test assets in a representative-agent model is conceptually very similar to using the cross-sectional dispersion of interest rates as an additional pricing factor in a heterogeneous-agent framework, as we do here.[3] In fact, in our model, the cross-sectional dispersion of interest rates helps price the cross-section of currency returns using quarterly data and without prior portfolio formation. We therefore call this interest-rate

[3] One interpretation of our approach is that we use conditioning information in a different way. While the LV approach can be interpeted as using scaled returns, our approach relies on scaled betas (see Cochrane (2005), ch. 8)

dispersion factor the LV factor and we interpret it as an indicator of financial market integration: when nominal interest rates equalize across countries, then *ceteris paribus*, consumption growth rates should equalize, reflecting the same cost of intertemporal substitution across countries.

The second important driver of international consumption dispersion to which we draw attention here is *heterogeneity in inflation rates*: as emphasized by Backus-Smith (1993), even in complete and seamlessly integrated financial markets, intertemporal marginal rates of substitution and therefore consumption growth rates should equalize only to the extent that purchasing power equalizes. We find that the inclusion of the cross-section of inflation rates contributes significantly to the empirical success of our model. We therefore call this factor the Backus-Smith factor and we interpret it as an indicator of goods market integration:[4] if goods markets are perfectly integrated and an ever larger fraction of goods is traded, consumption price levels denominated in a common currency should equalize and the international dispersion of inflation rates should decrease.

Our decomposition of the cross-sectional variance of consumption into one component that is due to cross-sectional heterogeneity in inflation rates – the Backus-Smith (or 'BS') component – and a second component that is due to international variation in nominal interest rates – the LV component – shows that both factors account to virtually equal parts for the pricing power that the cross-sectional variance of consumption has for exchange rates.[5] Our results not only suggest that the failure of uncovered interest

[4]A third determinant of the cross-sectional dispersion of consumption should be the cross-sectional covariation between interest rate and inflation rates. We also explore the role of this 'Fisher' factor in our analysis. It is, however, empirically not very relevant and that's why we omit it from our discussion here.

[5]Adler and Dumas (1983) were among the first to argue that idiosyncratic inflation risk could be a powerful factor in explaining currency returns. Our results support this view.

rate parity and the consumption real exchange rate anomaly are more closely linked than may have commonly been believed. They can also be read as an extension of the results by Lustig and Verdelhan (2007) to a setting with potentially incomplete markets. In addition, they provide a taxonomy of the extent to which the consumption-real exchange rate link is determined by financial market incompleteness on the one hand and segmented goods markets on the other.

As a policy application of our framework, we ask how the price of idiosyncratic consumption risk – as an indicator of the marginal value of buying an additional unit of consumption insurance – has changed over time, as world financial markets have become more integrated. We find that both the quantity of idiosyncratic risk as well as its marginal disutility have decreased substantially, in particular since 1990. This finding ties in with an important and by now well-established literature that dates the onset of the recent major wave in financial globalization in the first half of the 1990s (Lane and Milesi-Ferretti (2002,2003,2004)). Our results also line up with those of Brandt et al. (2006) who argue that given the volatility of exchange rate changes international risk sharing must be much better than is commonly implied by consumption based tests. Our analysis here puts a price on consumption dispersion, shows that it is relatively small and that it has been declining over time.

The remainder of this paper is organized as follows. Section two presents our theoretical framework and relates it to the Backus-Smith puzzle. We present our main results in section three and finally conclude in section four.

2 A simple framework

In complete currency markets, the change in the exchange rate can be expressed as the relative intertemporal marginal rate of substitution:

$$\Delta q_{t+1}^{k} = m_{t+1}^{k} - m_{t+1} \tag{1}$$

where $\Delta q_{t+1}^{k}$is the percentage change in the real exchange rate of country k and $m_{t+1}$and m_{t+1}^{k} are the logarithmic home and foreign discount factors respectively. In a setting with constant relative risk aversion, (1) can be specialized to

$$\Delta q_{t+1}^{k} = \gamma(\Delta c_{t+1}^{k} - \Delta c_{t+1}) \tag{2}$$

where γ is the coefficient of risk aversion, Δc denotes consumption growth and superscript k denotes a country k variable. Variables without a superscript k pertain to the home country. Since Backus and Smith (1993), this relationship has been the focus of much empirical research. Most studies find the link between real exchange rates and consumption tenuous at best, the correlation generally insignificant and often wrongly signed. In our analysis here, we build on Lustig and Verdelhan (2007) to explore this link in the cross-section. In so doing, we use and extend a simple version of the consumption capital asset pricing model (CCAPM).

The standard version of the CCAPM assumes the existence of a representative agent which implies that consumption growth rates should be equalized internationally. However, the very notion of equation (2) is that international consumption growth rates differ across countries and since Backus, Kehoe and Kydland (1992) low international consumption correlations are a prominently documented feature of the data. We therefore adapt our

CCAPM to account for this international consumption heterogeneity. We start from the first-order condition of a country k investor with respect to holdings of foreign currency:

$$\mathbf{E}_t(M_{t+1}^k R_{t+1}^j) = 1 \tag{3}$$

where $R_{t+1}^j$is the return of holding one unit of country j currency relative to an arbitrary base currency (We will generally use U.S. dollars as the base denomination). We discuss the exact specification of this return below.
For now we average condition (3) over all countries to obtain

$$\mathbf{E}_t\left(\frac{1}{K}\sum_{k=1}^{K} M_{t+1}^k R_{t+1}^j\right) = 1 \tag{4}$$

where K is the number of countries. We denote the average pricing kernel $\frac{1}{K}\sum_{k=1}^{K} M_{t+1}^k$with $\overline{M}_{t+1}$. Under CRRA-utility, we can approximate

$$M_{t+1}^k = \beta(1+\Delta c_{t+1}^k)^{-\gamma} \tag{5}$$

where, in addition to the notation already introduced above, β is the subjective discount factor. We then expand $\overline{M}_{t+1}$ around world average consumption growth, $\Delta\overline{c}_{t+1}$to obtain

$$\begin{aligned}\overline{M}_{t+1} \approx{}& \beta(1+\Delta\overline{c}_{t+1})^{-\gamma} - \frac{\gamma}{K}\sum_{k=1}^{K}\beta(1+\Delta\overline{c}_{t+1})^{-\gamma-1}(\Delta c_{t+1}^k - \Delta\overline{c}_{t+1}) \\ &+\frac{\gamma(\gamma+1)}{2K}\sum_{k=1}^{K}\beta(1+\Delta\overline{c}_{t+1})^{-\gamma-2}(\Delta c_{t+1}^k - \Delta\overline{c}_{t+1})^2\end{aligned} \tag{6}$$

The first order term is zero by construction whereas for large K, the second term is a multiple of the cross-sectional sample variance of Δc_{t+1}^k.

Denoting this consumption dispersion term with

$$\sigma^2_{K,t+1} = \frac{1}{K}\sum_{k=1}^{K}(\Delta c^k_{t+1} - \Delta\bar{c}_{t+1})^2$$

we obtain

$$\overline{M}_{t+1} \approx \beta(1+\Delta\bar{c}_{t+1})^{-\gamma}\left[1+\frac{\gamma(\gamma+1)\sigma^2_{K,t+1}}{2(1+\Delta\bar{c}_{t+1})^2}\right] \tag{7}$$

In our empirical implementation, we use the logarithm of $\overline{M}_{t+1}$that we approximate as

$$\overline{m}_{t+1} \approx \kappa - \gamma\Delta\bar{c}_{t+1} + \delta\frac{\sigma^2_{K,t+1}}{(1+\Delta\bar{c}_{t+1})^2} \tag{8}$$

where

$$\delta = \gamma(\gamma+1)/2$$

κ is a constant term and $m_{t+1} = \ln M_{t+1}$. Hence, the international heterogeneity of consumption growth rates gives rise to a second pricing factor, the cross-sectional variance of consumption, scaled by the square of the aggregate consumption growth rate.[6]

According to the Backus-Smith relation (1), international differences in consumption growth rates across countries are a key determinant of exchange rate changes. Holding foreign currency therefore necessarily exposes the investor to idiosyncratic consumption risk. It would therefore seem that differences in the exposure to such risk are a natural determinant of the cross-section of currency returns. In fact, our results suggest exactly this.

[6] In our empirical implementation we report results for $\sigma^2_{K,t+1}$ rather than $\sigma^2_{K,t+1}(1+\overline{\Delta c}_{t+1})^{-1}$ as pricing factor. This does not affect any of our conclusions.

As we will argue, the sorting of individual currencies into portfolios according to the size of the interest rate differential (as done in Lustig and Verdelhan (2007)) implicitly amounts to using the cross-sectional variance term as a pricing factor: using the cross-sectional variance term directly, we price quarterly currency returns in industrialized countries without prior portfolio formation. We also show that differences in the exposure to to risk are directly related to interest rate differentials and, therefore, to the currency risk premium.

2.1 Interest parity and the Backus-Smith condition

If uncovered interest parity holds, regressions of exchange rate changes on the nominal interest rate differential of the form

$$\Delta e_{t+1}^{k} = \alpha_k (i_t^k - i_t) + \epsilon_t \tag{9}$$

where Δe_{t+1}^{k} is the change in the nominal log exchange rate (measured in terms of foreign (country k) currency to home currency) and i_t denotes the nominal interest rate, should yield a coefficient of unity. The UIP puzzle, first stated by Fama (1984), is the empirical regularity that α_k is typically much smaller than unity and often even negative in the data. The excess return of investing into foreign over domestic bonds is given by $i_t^k - i_t - \Delta e_{t+1}^{k}$.Hence, a value of $\alpha < 1$ implies that the expected return on investing in foreign currency is predictable.

$$\mathbf{E}_t (i_t^k - i_t - \Delta e_{t+1}^{k}) = (1 - \alpha_k)(i_t^k - i_t) \tag{10}$$

so that high interest rate currencies generate a positive excess return. To highlight the link between the UIP puzzle and the Backus-Smith condition, we find it useful to start from the first-order condition of the country k investor for whom, in analogy to (3) above, the risk-free rate is determined by

$$\mathbf{E}_t(M^k_{t+1}R^f_{t+1}) = 1 \tag{11}$$

Note that the country k bond is a safe asset only for country k investors[7] – investors in all other countries would necessarily face exchange rate risk when investing into bond k. Hence, there are as many different risk-free rates as there are countries and under our maintained assumption of CRRA-utility, a second-order expansion of this condition (and the assumption of log-normality of consumption growth) yields

$$i^k_{t+1} = \mathbf{E}_t(\pi^k_{t+1}) + \kappa + \gamma\mathbf{E}_t(\Delta c^k_{t+1}) - \frac{\gamma^2}{2}\operatorname*{var}_t(\Delta c^k_{t+1}) \tag{12}$$

where $\kappa = \log(\beta)$ and π^k_{t+1} is the inflation rate in country k. We assume the variance term that reflects the precautionary savings motive to be constant over time and we normalize it to zero. Plugging in for $i^k_t - i_t$, we can then write the excess return of investing into foreign currency bonds as

$$i^k_t - i_t - \mathbf{E}_t(\Delta e^k_{t+1}) = \mathbf{E}_t(\pi^k_{t+1} - \pi_{t+1}) + \gamma\mathbf{E}_t(\Delta c^k_{t+1} - \Delta c_{t+1}) - \mathbf{E}_t(\Delta e^k_{t+1}) \tag{13}$$

Under UIP, this excess return should be zero, so that we obtain

$$\mathbf{E}_t(\Delta e^k_{t+1} + \pi_{t+1} - \pi^k_{t+1}) = \gamma\mathbf{E}_t(\Delta c^k_{t+1} - \Delta c_{t+1}) \tag{14}$$

This equation is almost identical with the Backus-Smith condition (2). It

[7] Clearly, country k investors may still face domestic inflation risk if the bonds are nominal.

is however, important to note that (14) holds in expectations only, whereas (2) holds in all states of nature; whereas the latter condition will only be satisfied in complete markets, the former is much more general. However, even tests that have focused on the much weaker version of the condition, e.g. Kollmann (1995), have generally found little evidence to support it. Our derivation here suggests that the failure to identify (2) – or even the weaker condition (14) – in the data is intimately linked to the failure of UIP. Clearly, condition (14) will hold only if UIP holds. Since UIP is grossly violated in the data, maybe we should not hope to identify a relation such as (2) from the data directly. But, according to (13), expected excess returns on a currency, i.e. the deviations from UIP, should be related to expected deviations from the Backus-Smith condition. Our empirical results suggest that they are. Whereas (14) has found little empirical support in the data, relation (13) has to our knowledge not been examined directly. Investigating (13) in the time series would involve proxying for the dynamics of unobserved expectations. Our approach is simpler – and arguably more successful: we focus on the cross-sectional implications of (13) by taking time series means of realized excess returns and of actual deviations of relative consumption from real exchange rates. We then compare the cross-section of the country-specific means of these two time series. Our empirical results below provide strong evidence in support of relation (13).

We therefore suggest to explore the link between exchange rates, consumption, inflation and interest rates further by exploiting the asset pricing implications of (14) using the pricing kernel expansion (8).

We first rewrite condition (13) as

$$\Delta c_{t+1}^{k} - \Delta c_{t+1} = \frac{1}{\gamma}\left(i_t^k - i_t - (\pi_{t+1}^k - \pi_{t+1})\right) + v_{t+1}^k$$

where v_{t+1}^k is an i.i.d. disturbance term. This implies that the cross-sectional variance of consumption should equal the cross-sectional variance of (ex post) real interest rates, up to the cross-sectional variance of the expectation error v_{t+1}^k so that

$$\sigma_{K,t}^2 = \text{var}_K(\Delta c_{t+1}^k) = \frac{\text{var}_K(i_t^k) - 2\,\text{cov}_K(i_t^k, \pi_{t+1}^k) + \text{var}_K(\pi_{t+1}^k)}{\gamma^2} + var_K\left(v_{t+1}^k\right) \tag{15}$$

where $\text{var}_K(.)$ denotes the sample cross-sectional variance operator. Plugging this relation into (8) and dropping the residual term $\text{var}_K(v_{t+1}^k)$, we obtain

$$\overline{m}_{t+1} \approx \kappa - \gamma\Delta\overline{c}_{t+1} + \frac{(\gamma+1)}{2\gamma}\frac{\text{var}_K(i_t^k) - 2\,\text{cov}_K(i_t^k, \pi_{t+1}^k) + \text{var}_K(\pi_{t+1}^k)}{(1+\Delta\overline{c}_{t+1})^2} \tag{16}$$

Note that it is an empirical question whether we can neglect $var_K(v_{t+1}^k)$. To the extent that we can, the approximation (16) should price exchange rates almost as well as the consumption-based representation (8) above. The approximation here implies a model with four pricing factors: the first is world average consumption growth. The second factor is the cross-sectional variance of nominal interest rates. This factor is key in understanding the link between our results here and the ones obtained in the Lustig and Verdelhan study. As Lustig and Verdelhan themselves note, their representative agent CCAPM successfully prices excess returns on a set of currency portfolios that are formed by the size of the interest rate differential, but not the individual currencies. This is because the portfolio formation emphasizes the cross-sectional spread of interest rates and this spread is needed to explain the cross-section of currency returns. In our setting, where we set out by explicitly acknowledging the heterogeneity in consumption, the cross-sectional dispersion of interest rates arises naturally as a pricing factor. We believe that it is the feature of our model that is at the root of one of our main empirical results: our ability to price currencies in a CCAPM context without prior portfolio formation. We call $var_K(i_t^k)$ the Lustig-Verdelhan (LV) factor.

Our framework gives rise to two more pricing factors. Most importantly, the term $var_K(\pi_{t+1}^k)$ captures idiosyncratic inflation risk. It is this risk which is at the heart of the Backus-Smith condition: international consumption correlations should be perfect only to the extent that relative purchasing powers are equalized. Consumption growth should be high in countries with high domestic purchasing power (low prices) and low in places with high prices (low purchasing power). To the extent that only inflation differentials account for departures of consumption correlations from unity – as the complete markets version of the Backus-Smith condition would suggest — the cross-sectional variance of consumption and the cross-sectional dispersion of inflation rates should comove perfectly. We therefore call $var_K(\pi_{t+1}^k)$ the Backus-Smith (BS) factor. We interpret this factor as an indication of goods market segmentation: if goods markets are perfectly integrated, international inflation differentials should at least be small or even vanish altogether.

Finally, there is the covariance between nominal interest rates and expected inflation that measures the cross-sectional comovement between idiosyncratic interest rate and inflation risk. We call this term a Fisher-parity factor. However, as we will see in our empirical implementation, this factor is only of very limited importance and we therefore do not discuss it any further.

We start our empirical analysis by using the purely consumption based representation (8) of the discount factor $\overline{m}_{t+1}$ to price currency returns. We then move on to the decomposition (16) of $\overline{m}_{t+1}$ that allows us to ask to what extent cross-sectional variation in interest and inflation rates contributes to the success of $var_K(\Delta c^k_{t+1})$ as a pricing factor.

A direct implication of our derivations above is that the cross-sectional dispersion of consumption growth rates should fully capture the conditioning information contained in interest rates. In our analysis we therefore focus on cash currency returns – exchange rate changes – directly.[8] Below, we show that consumption dispersion does indeed capture the international dispersion of interest rates and we demonstrate that our model therefore also helps explain departures from UIP.

3 Data and Empirical Results

3.1 Data

Our full sample contains quarterly data on exchange rates as well as real, p.c. consumption of Australia, Austria, Canada, France, Germany, Italy, Japan, Norway, Spain, Sweden, Switzerland, United Kingdom and the United States for the time period from 1971Q1 to 2003Q2. We calculate quarterly exchange rate changes from end-of-quarter MSCI stock market returns denominated in U.S. dollar (or the respective currency of interest) and local currency freely available on www.mscibarra.com. We update the Campbell (1999) consumption data set with consumption, cpi and population data from the IFS January 2004 tape to construct world average consumption growth as

[8] In this respect, our approach differs from Sarkissian (2003) who prices forward currency premia.

well as the cross-sectional variance of world consumption growth at each point in time. Since consumption growth is unitless, average world consumption growth is an equal-weighted average of growth rates from real, p.c. consumption in local currency of the 13 countries under consideration.

In order to illustrate the Lustig and Verdelhan (2007) mechanism in our framework we obtain short-term interest rates (treasury bills or call money market rates) from the IFS January 2004 CD following the recommendation by Campbell (1999).

3.2 Empirical Results

3.2.1 Departures from UIP and the Backus-Smith condition

According to (13), departures from UIP should be directly related to the violation of the Backus-Smith relation. A time-series test of this relation would require us to proxy for unobserved expectations. Rather, we provide evidence on (13) by focussing on its cross-sectional dimension. Figure (1) plots the time series means for each country in our sample of $i_t^k - i_t^* - \Delta e_{t+1}^k$ against that of $\gamma \left(\Delta c_{t+1}^k - \Delta c_{t+1}^*\right) - \Delta e_{t+1}^k - (\pi_{t+1}^k - \pi_{t+1})$. The different figure panels pertain to different values of the risk aversion parameters ($\gamma = 1, 5, 25, 50$). As becomes apparent, irrespective of a particular value of γ, there is a very close relation between these variables, suggesting that the cross-section of deviations from UIP is indeed closely related to the cross-section of departures from the incomplete markets version of the Backus-Smith relation. A regression of the mean deviations from UIP on the deviations from the Backus-Smith condition corroborates the visual impression. We find highly significant estimates and high R^2 throughout.

3.2.2 Pricing currency returns

We follow two different asset pricing approaches to explore the relation between our world consumption-based CAPM with heterogeneity and returns on currencies. First, we ask if one of the two factors in question, world consumption growth and the cross-sectional variance of consumption growth, helps to price currency returns given the presence of the other. Therefore, we use the generalized methods of moments (GMM) framework of Hansen and Singleton (1982) to exploit the testable restrictions imposed by

$$E_t(\overline{m}_{t+1} r^j_{t+1}) = 0$$

with $\overline{m}_{t+1}$ the natural logarithm of the average pricing kernel and r^j_{t+1} the log return on currency j. We simplify the log-linear discount factor to

$$m_{t+1} = 1 - b_1 \Delta \bar{c}_{t+1} - b_2 \sigma^2_{K,t+1}$$

The GMM estimator chooses the parameters b_1 and b_2, such that the model fulfils the restrictions in (3) best. This procedure amounts to minimizing the errors of our pricing model. Therefore, we use a two-stage procedure. We employ the identity matrix as weighting matrix in the first step to obtain estimates of b_1 and b_2 and the optimal weighting matrix suggested by Hansen and Singleton (1982) in the second step to compute standard errors. All of our results remain qualitatively the same if we use first-stage GMM estimates only. [9]

Secondly, we assess what risk factor is actually priced in currency returns. We estimate risk prices via the Fama – MacBeth (Fama and MacBeth

[9] See Cochrane (2005) for an excellent introduction to the GMM framework in asset pricing.

(1973)) cross-sectional regression. In the first stage we run the following time series regressions

$$\Delta e^k_{t+1} = \mu + \beta^k_c \Delta \bar{c}_{t+1} + \beta^k_\sigma \sigma^2_{K,t+1} + \xi^k_{t+1} \tag{17}$$

with $\beta^j_{\Delta \bar{c}}$ the sensitivity of currency return k to world consumption growth and $\beta^j_{\sigma^2_K}$ the exposure of the return on currency k to idiosyncratic consumption risk, i.e. the cross-sectional variance of consumption.
In the second step we run the following cross-sectional regressions

$$\Delta e^k_t = \mu + \lambda_c \hat{\beta}^k_c + \lambda_\sigma \hat{\beta}^k_\sigma +, \forall t \tag{18}$$

at each point in time to obtain estimates of the risk prices, λ_c and λ_σ.

Table 1 presents GMM estimates and risk prices when we consider nominal returns on U.S. dollar exchange rate changes of twelve developed countries at quarterly frequency.[10] We start with pricing nominal exchange rate changes because the "true" rational expectations risk premium on foreign currency should be predominantly determined by the covariance of exchange rate changes (currency returns) with consumption growth or other systematic sources of risk as shown by Engel (1996) and Lustig and Verdelhan (2007).

The row 'GMM' provides estimates of the parameters b_1 and b_2 with Newey-West (Newey and West (1987)) corrected t-statistics in parenthesis below the estimates. The column 'J-Test' gives the p-value of a test of the null that all pricing errors are jointly zero. Apparently we cannot reject this null at conventional confidence levels. Furthermore, the parameter estimates are both statistically significantly different from zero. Both of the factors

[10] Australia, Austria, Canada, France, Germany, Italy, Japan, Norway, Spain, Sweden, Switzerland, United Kingdom.

seem to help pricing currency returns. Note also that b_1 mirrors the coefficient of constant relative risk aversion (CRRA) since we assume investors maximize a power utility function.

The estimate of the CRRA coefficient of about 22 is considerably lower than in other studies that focused on the pricing of currency excess returns with a consumption-based CAPM not allowing for consumer heterogeneity (see e.g. Mark (1985), Lustig and Verdelhan (2007)). This result is due to the cross-sectional variance term that derives immediately from our heterogeneous agent framework. On the other hand, the coefficient on the cross-sectional variance term turns out to be incorrectly signed.[11] This pattern seems to be quite common in models with heterogeneous consumption and it lines up with the findings by Jacobs and Wang (2004) who introduce heterogeneity in domestic consumption into a CCAPM to price domestic stock returns. It is interesting that our results corroborate this pattern on a completely different sample period using currency returns instead of stock returns.

The GMM estimates reflect that both of the risk factors are helpful to price currency returns. However, we are particularly interested if these factors are actually priced and how large the respective price of risk is. The row 'Fama-MacBeth' of table 1 presents Fama-MacBeth cross-sectional regression estimates of the risk prices of our two factors in question. The t-statistics in parenthesis are adjusted for the fact that the regressors are generated (Shanken (1992)). R^2 denotes the cross-sectional R^2 proposed by Jagannathan and Wang (1996).

The estimated price of world consumption growth risk is statistically indistinguishable from zero. Though average consumption growth risk seems to be helpful in pricing currencies, it is not priced in average currency returns.

[11] However, Mankiw (1986) argues that positive and negative estimates are possible.

Quite in contrast to this finding, the cross-sectional variation in consumption growth is a significant determinant of returns on currencies. We estimate its price to be around 0.14 percentage points per quarter, i.e. approximately 0.56 percentage points annually. In addition, the heterogeneous agent world consumption CCAPM explains about sixty percent of the cross-sectional dispersion in average returns on currencies. We are thus able to show that international consumption dispersion is an important determinant of the cross-section of exchange rate changes. Figure 2 plots mean actual currency returns against their counterparts predicted by the model to give a visual impression of the fit of the model.

Our results suggest that idiosyncratic consumption risk is crucial in order to explain the cross-sectional differences in average exchange rate changes and hence risk premia on foreign currencies. In addition, taking account of international consumption heterogeneity allows to price individual currency returns directly which is interesting given that the seminal results of Lustig and Verdelhan (2007) rely on a set of currency portfolios that are formed with respect to interest rate differentials. As we argue next, one way to interpret this sorting of currencies by the size of their current interest rate differential implicitly amounts to doing what we do here – using the cross-sectional variance of consumption as a pricing factor.

3.2.3 Inside the mechanism: consumption heterogeneity and the UIP puzzle

The idea that an extra risk premium for holding foreign assets explains the empirical failure of UIP has attracted considerable attention (see e.g. the survey in Engel (1996)). As argued by Bansal and Dahlquist (2000), basic asset pricing theory should apply to the UIP puzzle and Lustig and

Verdelhan provide an empirical illustration of this insight in a representative agent CCAPM. The excess return on foreign bonds, ϕ^k_{t+1}, i.e. the departures from UIP, are given by

$$\phi^k_{t+1} = i^k_t - i_t - \Delta e^k_{t+1} \tag{19}$$

In our framework we price cash currency returns directly. Since interest rate differentials are known as of time t, this is tantamount to saying that we should also be able to explain departures from UIP. Our new pricing factor, the cross-sectional dispersion of consumption, should therefore capture the conditioning information embedded in interest rate differentials.

To explore how and to what extent the consumption-dispersion factor captures this information, we first split exchange rate changes into the part that is perfectly correlated with interest rate differentials and into the orthogonal complement unrelated to interest rate differentials by running the Fama (1984) type regression[12]

$$\Delta e^k_{t+1} = \mu_k + \alpha_k(i^k_t - i_t) + \epsilon^k_{t+1} \tag{20}$$

We do not report our estimates of α_k since they largely corroborate the estimates provided by Bansal and Dahlquist (2000) for the countries under consideration here. But one should note that the estimates are virtually all smaller than one, and many of them negative– there is clear evidence of a UIP puzzle in our data set.

In the mould of the large literature on the UIP puzzle, we then proxy for the conditional expectation of Δe_{t+1} using the fitted value of regression

[12] The results in this subsection (3.2.3) are based only on the G7 economies (Canada, France, Germany, Japan, United Kingdom and the United States) since we we have uninterrupted time series for the entire sample period only for these countries.

(20) as[13]

$$\widehat{\Delta e}^k_{t+1} = \alpha_k(i^k_t - i_t)$$

With some abuse of language, we call $\widehat{\Delta e}^k_{t+1}$ the UIP-related component of the exchange rate change, based on which we construct a measure of the conditional expectation of the currency risk premium, ϕ^k_{t+1} as

$$\widehat{\phi^k_{t+1}} = i^k_t - i_t - \widehat{\Delta e}^k_{t+1} = (1 - \alpha_k)(i^k_t - i_t)$$

We then ask if the cross-sectional dispersion in $\widehat{\phi^k_{t+1}}$ reflects exposure to the two price factors, i.e. average consumption growth and consumption dispersion. To this end, we obtain a set of conditional betas by regressing

$$\epsilon^k_{t+1} = \mu + \beta^k_c \Delta\bar{c}_{t+1} + \beta^k_\sigma \sigma^2_{K,t+1} + \xi^k_{t+1}$$

where ϵ^k_{t+1} is the residual from (20). We then regress the time means of $\widehat{\phi}^k_{t+1}$ on the consumption dispersion betas β^k_σ. The results, reported in table 2, are quite clear: the model explains roundabout 90 percent of the cross-sectional variation in currency excess returns. Both world consumption growth as well as consumption dispersion are significantly priced – and this is at the level of the individual currency.

The results here also shed light on *why* we are able to price currency returns without prior portfolio formation. As Lustig and Verdelhan (2007) note, high interest rate currencies generate positive excess returns because they deliver lots of exposure to aggregate consumption growth risk, low interest rate currencies generate low expected returns because they are a hedge. Since individual interest differentials are mean-reverting so that individual

[13]Note that the coefficient μ_k is generally insignificant in our sample, so that we neglect this term in obtaining the fitted values.

countries shift from low to high interest rate differentials and vice versa over time, the excess returns on individual countries' currencies cannot have a time invariant beta with respect to consumption growth risk. One interpretation of our approach is that consumption dispersion proxies for exactly this time-variation in country's exposure to aggregate consumption growth risk. Specifically, our setup can be seen as the approximation to a model in which conditional betas with respect to consumption growth get scaled with the cross-sectional dispersion of consumption. To see this, note that for small $\sigma^2_{K,t+1}\Delta c_{t+1}$ we can expand the beta-representation for currency k's return

$$
\begin{aligned}
\epsilon^k_{t+1} &= \mu^k + \beta_c \Delta c_{t+1} + \beta_\sigma \sigma^2_{K,t+1} + \xi^k_{t+1} \\
&\approx \mu^k + \beta_c \Delta c_{t+1} + \beta_\sigma \sigma^2_{K,t+1} + \beta_\sigma \sigma^2_{K,t+1} \Delta c_{t+1} \\
&= const + (\beta_c + \beta_\sigma \sigma^2_{K,t+1})(1 + \Delta c_{t+1}) + \xi^k_{t+1} \\
&= const + \beta_{t+1}(1 + \Delta c_{t+1}) + \xi^k_{t+1}
\end{aligned}
$$

with $\beta_{t+1} = \beta_c + \beta_\sigma \sigma^2_{K,t+1}$ and $const = \beta_c + \mu^k$.

Hence, instead of a model in which conditioning information is effectively used to scale returns, as in LV, our approach uses consumption dispersion as conditioning information to obtain scaled betas. We illustrate this intuition in figure (3), that further underscores the importance of consumption dispersion in obtaining our results. The figure plots the mean interest rate differentials $i^k_t - i_t$ against the conditional β_σ, i.e. the beta of ϵ^k_{t+1} with respect to idiosyncratic consumption risk. There is a clear positive relation: currencies that have low β_σ, appreciate when idiosyncratic consumption risk is high. This in turn is associated with low interest rate differentials. Our scaled-factor interpretation provides an explanation for this pattern: the currency's low exposure to idiosyncratic consumption risk also implies a low conditional beta with respect to aggregate consumption risk. The currency is a hedge and it will be a better hedge, the higher idiosyncratic risk is. According to the standard CCAPM, the currency should therefore yield low expected returns, and the results by Lustig and Verdelhan imply that it should be a low interest differential currency.

Our results extend those of Lustig and Verdelhan (2007) to an empirical setup with (potentially) incomplete markets. While the mechanism highlighted by Lustig and Verdelhan seems to be key in reconciling cross-sectional differences in returns on currencies with the CCAPM, the intuition behind this mechanism explicitly relies on heterogeneity in consumption across countries. In the data, international consumption correlations are far from perfect. To the extent that this is the case, idiosyncratic consumption risk should be priced into currency returns. Indeed, we find that idiosyncratic consumption risk is priced and the results in this subsection suggest that consumption dispersion captures the information implicit in interest rate differentials. Next, we show that there is another important reason for why consumption heterogeneity is priced: the risk associated with international fluctuations in purchasing power. It is this heterogeneity in inflation rates that is at the heart of the Backus-Smith relation. We show that idiosyncratic inflation risk is an equally important determinant of exchange rate changes.

3.2.4 The cross-section of currency returns and the Backus-Smith puzzle

In section 2.1, we related our asset pricing framework to the relation between real exchange rates and relative consumption growth as proposed in Backus and Smith (1993). We derive a similar but weaker condition, equation (14), that earlier research focusing on the time-series dimension has found to be an empirical challenge (e.g. Kollmann (1995)). As we have demonstrated, deviations from UIP and deviations from the BS condition are highly correlated which implies that the international dispersion in consumption growth rates is indeed related to variation in (nominal) interest rates and inflation rates. To further dissect the role of consumption heterogeneity, we now decompose international consumption dispersion along the lines of the decomposition in (16) into a Lustig-Verdelhan factor (the cross-sectional variance of nominal interest rates), the Backus-Smith factor (the cross-sectional variance of inflation) and a Fisher parity term (the covariance of interest rates with inflation). We assess the importance of each of the factors for the pricing of currency returns by estimating their respective risk prices.[14]

Table 3 displays the results from Fama-MacBeth cross-sectional regressions when the full cross-section of currency returns is considered. World average consumption growth still is not a significant pricing factor - real exchange rates seem to be determined primarily by idiosyncratic factors. The cross-sectional variance of interest rates as well as inflation rates explain average exchange rate changes. Their risk price estimates are comfortably statistically significant whereas the Fisher parity term does not seem to be priced. These results show clearly that both the Backus-Smith (BS) and the Lustig-Verdelhan (LV) factor are conceptually different sources of risk each of which is individually important. Considering them jointly hugely increases the ability of our model to explain the cross-section of currency returns.

[14] Note that our results from here on are again based on all 13 currencies. We obtain the cross-sectional dispersion of interest rates over all currencies for which interest rates are available at a given point in time. Results based on the G7 economies alone are qualitatively very similar.

Moreover, note that the explanatory power of this specification of the heterogeneous agent CCAPM is approximately the same as when we consider the cross-sectional consumption dispersion as pricing factor. The importance of both factors for capturing the cross-sectional dispersion in currency returns is highlighted by the results displayed in table 4. Here we assess the performance of our asset pricing model when either the LV or BS factor are left out. Panel A gives the risk price estimates as well as mean squared and mean absolute pricing errors when we exclude the BS factor. Panel B shows the corresponding estimates when the LV factor is missing.

Compared to the baseline regression with all four factors in table 3, it turns out that consumption growth and the LV factor alone capture a substantial fraction of the cross-sectional variation in exchange rate changes. However, the results in panel B suggest that, when considered on its own, the BS factor is even more successful in explaining the cross-section of currency returns. The cross-sectional R^2 is substantially higher and at the same time pricing errors are lower than those reported in panel A. While the LV factor does a remarkable job in explaining the cross-section of currency returns, we seem to miss a lot of explanatory power if we do not consider the BS factor – idiosyncratic inflation risk seems to be an important determinant of currency returns.

3.2.5 The price of idiosyncratic consumption risk over time

The impact of idiosyncratic consumption risk on currency returns is interesting from a macroeconomic and policy perspective: a growing literature documents that – in spite of two decades of continued financial globalization, international consumption correlations do not seem to have increased (see e.g. Heathcote and Perri (2004), Artis and Hoffmann (2004)). Does this mean that risk sharing has not increased? Our framework suggests a simple way to get at this issue: if idiosyncratic consumption risk is bad for consumers, it should have a price. Then improvements in risk sharing should be reflected either in a decline in idiosyncratic risk or in a decline of the price of this risk or in both.

Figure (4) presents the cross-sectional standard deviation of world consumption growth (solid line) as well as the Hodrick-Prescott filtered trend (dashed line) from 1971Q1 to 2003Q2. This graph leaves the impression that the cross-sectional dispersion in consumption growth rates decreases over time. Hence, idiosyncratic consumption risk in the 1990s might have been of less importance than in the 1970s since its *level* has decreased over time. This observation inevitably raises the question if the *price* of idiosyncratic consumption risk has declined over time as well.

We therefore estimate the factor risk prices when we regard currency returns on the full cross-section of 13 countries over two subsample periods. The first one ranges from 1971Q1 to 1989Q4 because the drop in the cross-sectional volatility in figure (4) at the end of 1980s or beginning of 1990s seems to be especially pronounced. Moreover, Lane and Milesi-Ferretti (2002, 2003, 2004) highlight the tremendous increase in cross-border capital flows since the early 1990s which is suggestive of deeper market integration in the course of financial globalization. The second period thus runs from 1990Q1 to 2003Q2. Table 5 summarizes our results.

Panel A shows that for the first sample period, the 1970s and 1980s, the price of idiosyncratic consumption growth is significant and about as high as in the full sample. It drastically decreases in the 1990s and is not distinguishable from zero in the latter subperiod (panel B). All in all, these findings leave the impression that the lower the level of idiosyncratic consumption risk, the lower is also its price. Thus, even though low international consumption correlations remain a salient feature of the data, they are not necessarily associated with a lack of international risk sharing. Rather, our results suggest that the price of idiosyncratic consumption risk has fallen considerably in the course of the last 15-20 years. This ties in with the findings in Brandt et al. (2006) who argue that – in spite of low international consumption correlations, international risk sharing is much better than is commonly thought.

As a last exercise, we explore to what extent goods and financial market integration respectively can account for the decline in the price of idiosyncratic consumption risk. Table 6 reports subsample estimates of the LV and BS factor respectively. The LV and the BS factor are statistically significant determinants of average currency returns in the period from 1971Q1 to 1989Q4 but not in the more recent period from 1990Q1 to 2003Q2. In the light of our earlier interpretation of these factors, this finding suggests that both financial and goods market integration are key drivers of the increase in international consumption risk sharing.

3.2.6 Robustness Checks

Before we conclude, we briefly report on a range of robustness checks with respect to alternative data specifications: we tested our CCAPM with heterogeneity using U.S. dollar exchange rate changes deflated with realized inflation. The results in panel A of table 7 reveal that the qualitative results remain unaltered.

The same is true if we change the base currency of our returns and regard Yen (panel B) as well as Deutschmark/Euro (panel C) returns. Furthermore, the estimate of the coefficient of relative risk aversion as well as the price of idiosyncratic consumption growth risk have the same order of magnitude as in the U.S. dollar case.

4 Conclusions

In this paper we have provided an asset pricing perspective on the link between real exchange rates and consumption. This link, first explored by Backus and Smith (1993), suggests that relative consumption growth rates should be inversely related to fluctuations in real exchange rates: a country's consumption growth should be high when local prices are low and vice versa. While this link is central in most modern macro-models, it has found little support in the data. Here we have argued that the empirical failure of this link and the failure of uncovered interest parity may be two sides of the same medal.

In a first step we recognize that if consumption growth rates differ across countries, then this heterogeneity should be reflected in standard pricing relations for currency risk. We find that idiosyncratic consumption risk – measured by the cross-sectional variance of real consumption growth – explains more than 60 percent of the cross-sectional variation in exchange rates in a standard consumption CAPM adjusted for heterogeneous consumers. In a second step we further decompose idiosyncratic consumption risk into one component that is due to heterogeneity in inflation rates – a Backus-Smith factor – and another that reflects international heterogeneity in the price of intertemporal substitution and that is measured by the cross-sectional variance of interest rate differentials – the Lustig-Verdelhan factor. We find both factors significant and successful in pricing the cross-section of currency returns.

As a policy application of our framework, we have examined how the price of idiosyncratic consumption risk has changed over time as financial globalization has accelerated in pace. We find that, even though international consumption correlations do not at first sight seem to have increased since the beginning of the 1990s, the cross-sectional dispersion of consumption risk has clearly been trending downward as has the world price of idiosyncratic consumption risk. This drop in the price of idiosyncratic consumption risk seems to be driven to equal parts by financial and goods market integration – we find that the risk prices of both the interest rate dispersion factor as well as of the idiosyncratic inflation risk factor have fallen considerably.

References

[1] Adler, Michael and Bernard Dumas (1983), "International Portfolio Choice and Corporation Finance: A Synthesis", *Journal of Finance* 38, 925-984.

[2] Artis, Michael J. and Mathias Hoffmann (2004), "Financial Globalization, International Business Cycles and Consumption Risk Sharing", CEPR discussion paper 5697.

[3] Backus, David K. and Gregory W. Smith (1993), "Consumption and real exchange rates in dynamic economies with non-traded goods", *Journal of International Economics*, 35, 297-316.

[4] Backus, David K, Kehoe, Patrick J. and Fynn E. Kydland (1992), "International real business cycles", *Journal of Political Economy*, 100, 745-775.

[5] Bansal, Ravi and Magnus Dahlquist (2000), "The Forward Premium Puzzle: Different Tales from Developed and Emerging Markets", *Journal of International Economics* 51(1), 115-144.

[6] Brandt, Michael W., Cochrane, John H. and Pedro Santa-Clara (2006), "International risk sharing is better than you think, or exchange rates are too smooth", *Journal of Monetary Economics* 53, 671-698.

[7] Brav, Alon, Constantinides, George M. and Christopher C. Geczy (2002), "Asset Pricing with Heterogeneous Consumers and Limited

Participation: Empirical Evidence”, *Journal of Political Economy* 110, 793-824

[8] Campbell, John Y. (1999), “Asset Prices, Consumption and the Business Cycle”, in John Taylor and Mark Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: North Holland.

[9] Cochrane, John H. (2005), *Asset Pricing*, 2nd edition, Princeton University Press, Princeton and Oxford.

[10] Constantinides, George M. and Darrell Duffie (1996). ‘Asset Pricing with Heterogeneous Consumers’, *Journal of Political Economy* 104(2), 219-40.

[11] Engel, Charles (1996), "The forward discount anomaly and the risk premium: A survey of recent evidence", *Journal of Empirical Finance*, 3, 123-192.

[12] Fama, Eugene F. (1984), “Forward and Spot Exchange Rates”, *Journal of Monetary Economics*, 14, 319-338.

[13] Fama, Eugene F. and James D. MacBeth (1973), "Risk, Return and Equilibrium: Empirical Tests", *Journal of Political Economy*, 81, 607-631.

[14] Hansen, Lars Peter and Kenneth J. Singleton (1982), "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models", *Econometrica*, 50, 1269-1288.

[15] Heathcote, John and Fabrizio Perri (2004), “Financial globalization and real regionalization”, *Journal of Economic Theory*, 119, 207-243.

[16] Jacobs, Chris and Kevin Q. Wang (2004), “Idiosyncratic Consumption Risk and the Cross Section of Asset Returns”, *Journal of Finance*, 59, 2211-2242.

[17] Jagannathan, Ravi and Zhengyu Wang (1996), “The Conditional CCAPM and the Cross-Section of Expected Returns”, *Journal of Finance*, 51, 3-53.

[18] Kollmann, Robert (1995), “Consumption, real exchange rates and the structure of international asset markets”, *Journal of International Money and Finance*, 14, 191-211.

[19] Lane, Philip R. and Gian Maria Milesi-Ferretti (2002), “Long-Term Capital Movements”, *NBER Macroeconomics Annual*, 16, 73-116.

[20] Lane, Philip R. and Gian Maria Milesi-Ferretti (2003), “International Financial Integration”, *IMF Staff Papers*, 50, 82-113.

[21] Lane, Philip R. and Gian Maria Milesi-Ferretti (2004), “Financial Globalization and Exchange Rates”, IIIS discussion paper No. 44

[22] Lustig, Hanno and Adrien Verdelhan (2007), "The Cross-Section Of Foreign Currency Risk Premia And Consumption Growth Risk", *American Economic Review*, 97, 89-117.

[23] Mankiw, N. Gregory (1986), “The Equity Premium and the Concentration of Aggregate Shocks”, *Journal of Financial Economics*, 17, 211-219.

[24] Mark, Nelson C. (1985), "On Time Varying Risk Premia in the Foreign Exchange Market", *Journal of Monetary Economics*, 16, 3-18.

[25] Newey, Whitney K. and Kenneth D. West (1987), "A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix", *Econometrica*, 55, 703-708.

[26] Ravn, Morten O. (2001), “Consumption Dynamics and Real Exchange Rates”, CEPR discussion paper 2940.

[27] Sarkissian, Sergei (2003), “Incomplete Consumption Risk Sharing and Currency Risk Premiums”, *Review of Financial Studies*,16, 983-1005.

[28] Shanken, Jay (1992), "On the Estimation of Beta-Pricing Models", *Review of Financial Studies*, 5, 1-33.

Tables

Table 1: GMM estimates and risk prices with Δe_{t+1}^{k} as test assets

	b_1	b_2	J-Test (p-value)
GMM	22.43 (4.31)	357.35 (7.54)	0.44
	λ_c	λ_σ	R^2
Fama-MacBeth	0.87 (1.05)	0.14 (3.35)	0.62

Notes: The table provides estimates from two-stage GMM estimations of $1 = E_t(M_{t+1}\Delta e_{t+1}^{k})$ with $m_t = 1 - b_1\Delta\overline{c_t} - b_2\sigma^2{}_{K,t+1}$
where Δe_{t+1}^{k} represents U.S. dollar exchange rate changes defined as foreign (country k) currency over U.S. dollar, $\Delta\overline{c}_t$ average world consumption growth at time t and $\sigma^2{}_{K,t+1} = \underset{K}{\mathrm{var}}\ \Delta c_{t+1}^{k}$ the cross-sectional consumption dispersion at time t+1. T-statistics occur in parenthesis. In the first stage we use the identity matrix as weighting matrix and in the second stage the Newey-West (Newey and West, 1987) corrected optimal weighting matrix as in Hansen and Singleton (1982). The column J-Test gives the p-value of a test of overidentifying restrictions that tests the null of all pricing errors being zero.

Furthermore, we present estimates of risk prices of the factors under consideration using the Fama-MacBeth cross-sectional regression which takes the following form:

$$\Delta e_t^{k} = \mu + \lambda_c\beta_c^{k} + \lambda_\sigma\beta_\sigma^{k} + \varepsilon^{k}{}_t, \forall t$$

The betas are obtained from time series regressions of exchange rate changes on average world consumption growth and the cross-sectional variance of consumption growth. T-statistics of the risk price estimates, λ, are corrected for this generated regressor bias (Shanken, 1992).
The sample spans the period from first quarter 1971 to second quarter 2003 and uses data on exchange rates and consumption from the following countries: Australia, Austria, Canada, France, Germany, Italy, Japan, Norway, Spain, Sweden, Switzerland, United Kingdom and the United States.

Table 2: Conditional risk price estimates for excess currency returns

	λ_c	λ_σ	R^2
Fama-MacBeth	-9.98 (-10.53)	0.94 (10.24)	0.89

Notes: The table reports risk price estimates (in % per quarter) from cross-sectional Fama-MacBeth regressions of $(1-\alpha^k)(i_t^k - i_t)$ on consumption growth and consumption dispersion betas of the UIP residuals ε_{t+1}^k. The estimate equation takes the form

$$(1-\alpha^k)(i_t^k - i_t) = \lambda_c \hat{\beta}_c^{\varepsilon,k} + \lambda_\sigma \hat{\beta}_\sigma^{\varepsilon,k} + v_t, \forall t$$

where $\hat{\beta}_x^{\varepsilon,k}$ (x=c,σ) is the exposure of the UIP residuals ε_{t+1}^k with respect to the pricing factors world average consumption growth and consumption dispersion respectively. The $\hat{\beta}_x^{\varepsilon,k}$ are estimated from a first-stage time series regression. The sample spans the period from 1971Q1 to 2003Q2. Countries included are CND, FRA, GER, ITA, JPN, UK.

Table 3: Risk prices of LV and BS factor

	λ_c	λ_{LV}	λ_{BS}	λ_{FP}	R^2
Fama-MacBeth	0.56 (0.76)	0.01 (2.69)	0.01 (3.66)	-0.06 (-1.60)	0.69

Notes: Here, we present estimates of risk prices if we split the cross-sectional variance of consumption growth into its components suggested by the relationship between real interest rates and consumption growth implied by the Backus and Smith (1993) framework.
Our cross-sectional regression thus takes the following form:

$$\Delta e_{t+1}^{k} = \mu + \lambda_c \beta_c^k + \lambda_{LV} \beta_{LV_t}^k + \lambda_{BS} \beta_{BS}^k + \lambda_{FP} \beta_{FP}{}^k + v_{t+1}^k, \forall t$$

The betas are obtained from time series regressions of exchange rate changes on average world consumption growth and the cross-sectional variance of short-term interest rates (LV), inflation (BS) and the cross-sectional covariance of interest rates with inflation (FP). T-statistics of the risk price estimates, λ, are corrected for this generated regressor bias (Shanken, 1992).
The sample spans the period from first quarter 1971 to second quarter 2003 and uses data on exchange rates and consumption from the following countries: Australia, Austria, Canada, France, Germany, Italy, Japan, Norway, Spain, Sweden, Switzerland, United Kingdom and the United States.

Table 4: Risk prices of LV and BS factor separately

Panel A: without BS factor						
	λ_c	λ_{LV}		R^2	mspe	mape
Fama-MacBeth	0.48 (0.91)	0.01 (3.61)		0.38	0.12	0.25
Panel B: without LV factor						
	λ_c		λ_{BS}	R^2	mspe	mape
Fama-MacBeth	0.48 (0.70)		0.01 (359)	0.59	0.06	0.20

Notes: This table shows estimates of cross-sectional regressions of exchange rate changes on either world consumption growth and the LV factor or consumption growth and the BS factor. For further details consider the notes to table 4. The two right columns give mean squared prediction errors (mespe) and mean absolute prediction errors (mape) to judge the goodness of fit in terms of pricing errors.

Table 5:Evolution of risk prices over subsamples

λ_c	λ_σ	R^2
Panel A: 1970Q1-1989Q4		
λ_c	λ_σ	R^2
0.16 (0.20)	0.13 (2.43)	0.35
Panel B: 1990Q1-2003Q2		
λ_c	λ_σ	R^2
-0.63 (-0.83)	0.01 (0.40)	0.23

Notes: This table presents estimates risk prices of average world consumption growth and the cross-sectional variance of consumption growth over sub-samples. The risk prices are obtained with Fama-MacBeth cross-sectional regressions using exchange rate changes of the countries mentioned in table 1 as test assets.

Table 6: Risk prices of LV and BS factors (subsamples)

	λ_c	λ_{LV}	λ_{BS}	λ_{FP}	R^2
Panel A: 1971Q1 – 1989Q4					
Fama-MacBeth	0.33 (1.10)	0.00 (2.07)	0.01 (2.62)	−0.01 (−1.57)	0.56
Panel B: 1990Q1 – 2003Q2					
	λ_c	λ_{LV}	λ_{BS}	λ_{FP}	R^2
Fama-MacBeth	−0.31 (−1.25)	0.00 (1.12)	0.00 (1.20)	−0.00 (−0.31)	0.65

Notes: Here, we present estimates of risk prices if we split the cross-sectional variance of consumption growth into its components suggested by the relationship between real interest rates and consumption growth implied by the Backus and Smith (1993) framework.
Our cross-sectional regression thus takes the following form:

$$\Delta e_{t+1}^{k} = \mu + \lambda_c \beta^{k}{}_{c} + \lambda_{LV} \beta_{LV_t}^{k} + \lambda_{BS} \beta_{BS}^{k} + \lambda_{FP} \beta_{FP}^{k} + v_{t+1}^{k}, \forall t$$

The betas are obtained from time series regressions of exchange rate changes on average world consumption growth and the cross-sectional variance of short-term interest rates, inflation and the cross-sectional covariance of interest rates with inflation. T-statistics of the risk price estimates, λ, are corrected for this generated regressor bias (Shanken, 1992).
Panel A reports the results for the sample period from first quarter 1971 to fourth quarter 1989, panel B the estimates for the period from 1990Q1 to second quarter 2003. We use data on exchange rates and consumption from the following countries: Australia, Austria, Canada, France, Germany, Italy, Japan, Norway, Spain, Sweden, Switzerland, United Kingdom and the United States.

Table 7: robustness checks

Panel A: $\Delta e_t - \Delta p_t$ from U.S. perspective			
	b_1	b_2	J-Test
GMM	25.04 (4.70)	409.27 (10.72)	0.17
	λ_1	λ_2	R^2
Fama-MacBeth	0.88 (1.01)	0.16 (3.37)	0.77
Panel B: Δe_t (Yen)			
	b_1	b_2	J-Test
GMM	23.34 (4.64)	366.31 (8.41)	0.33
	λ_1	λ_2	R^2
Fama-MacBeth	0.86 (0.99)	0.13 (2.24)	0.73
Panel C: Δe_t (Deutschmark/Euro)			
	b_1	b_2	J-Test
GMM	23.08 (4.63)	362.42 (8.28)	0.36
	λ_1	λ_2	R^2
Fama-MacBeth	0.42 (0.46)	0.13 (2.99)	0.71

Notes: The table provides estimates from two-stage GMM estimations of $m_t = 1 - b_1 \Delta \overline{c_t} - b_2 \sigma^2{}_{K,t+1}$

where Δe_{t+1}^k represents U.S. dollar exchange rate changes, $\Delta \overline{c}_t$ average world consumption growth at time t and $\sigma^2{}_{K,t+1} = \underset{K}{\text{var}}\ \Delta c_{t+1}^k$ and the cross-sectional consumption dispersion at time t. T-statistics occur in parenthesis.

Furthermore, we present estimates of risk prices of the factors under consideration using the Fama-MacBeth cross-sectional regression.

Figures

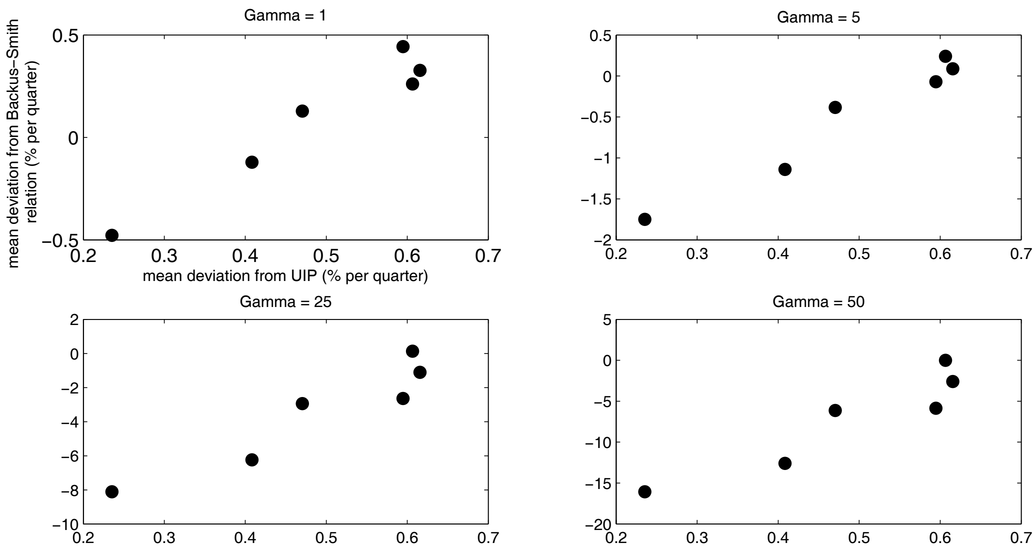

Figure 1: Time series means of deviations from UIP vs. deviations from Backus and Smith condition for different values of γ, the coefficient of relative risk aversion (Sample restricted to G7 countries due to availability of short-term interest rates).

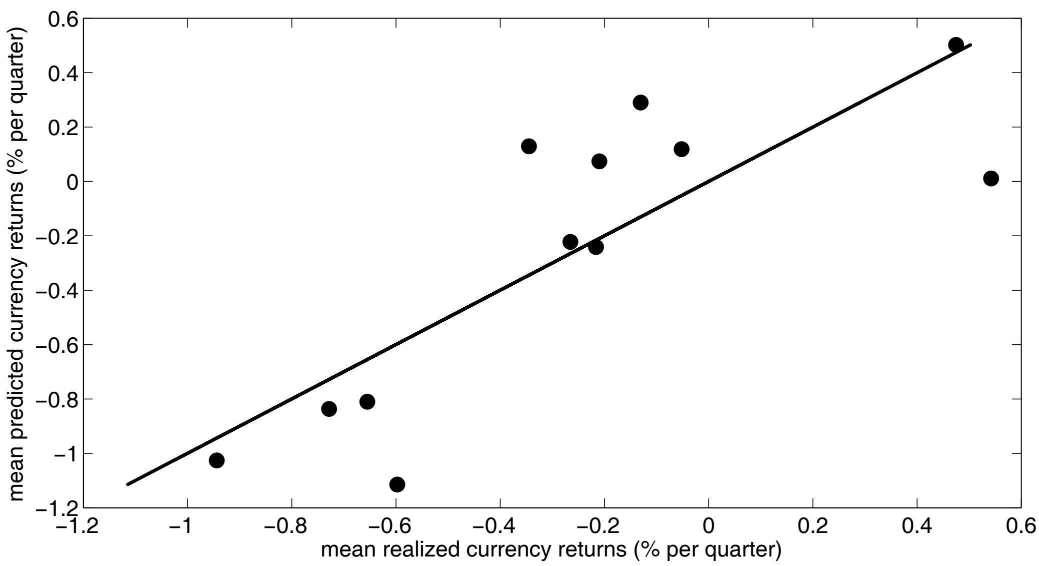

Figure 2: Fit of the CCAPM with heterogeneity (mean actual vs. mean predicted currency returns)

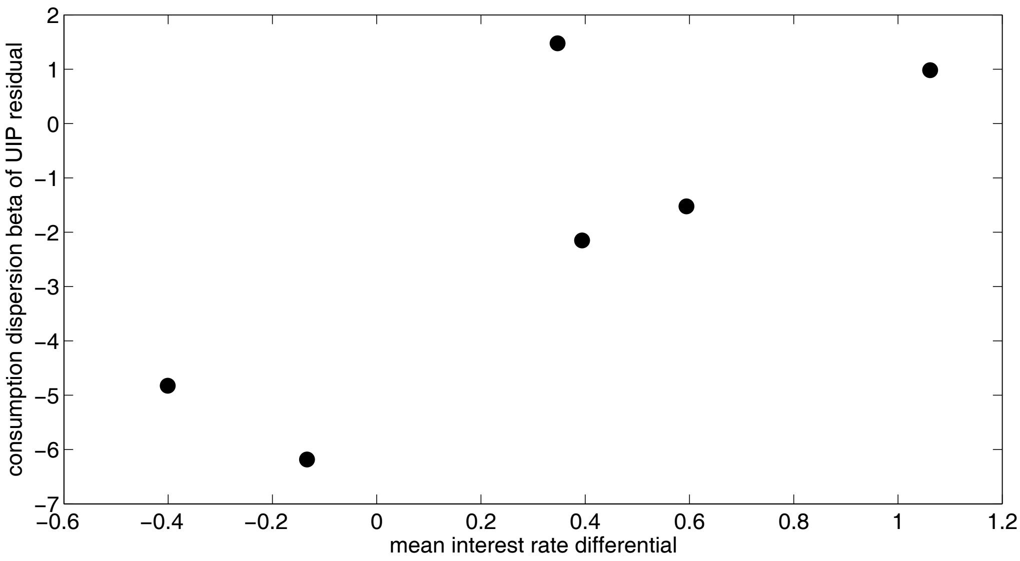

Figure 3: Mean interest rate differentials vs. idiosyncratic consumption risk betas of residuals from UIP regression (sample restricted to G7 countries due to availability of short-term interest rates)

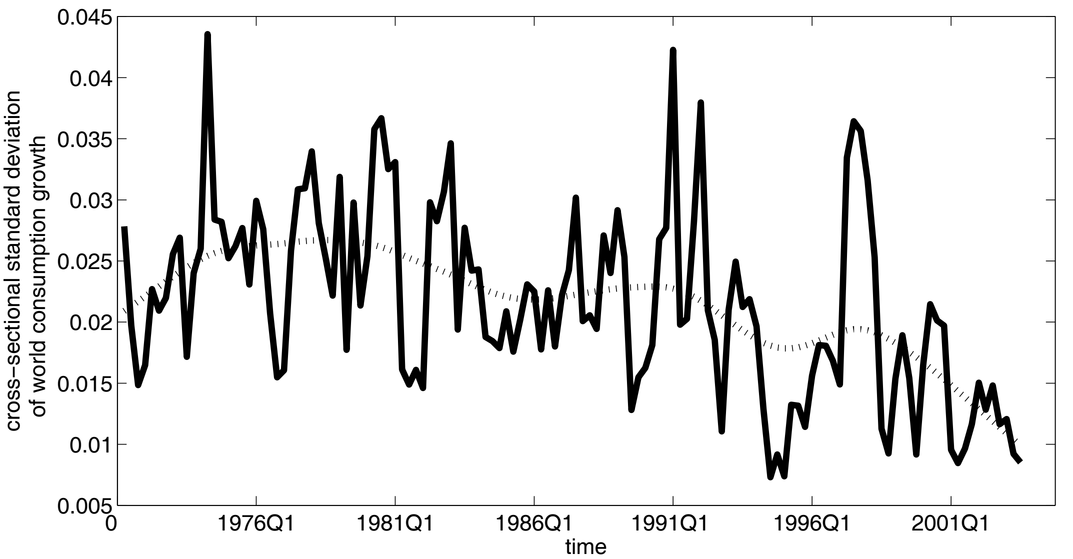

Figure 4: Cross-sectional standard deviation of world consumption growth