Incomplete Consumption Risk Sharing and Currency Risk Premiums

Sergei Sarkissian
McGill University

This article studies the impact of imperfect consumption risk sharing across countries on the formation of time-varying risk premiums in the foreign exchange market and on their cross-sectional differences. These issues are addressed within the framework of the Constantinides and Duffie (1996) model applied to a multicountry world. The article shows that the cross-country variance of consumption growth rates is counter-cyclical and that this feature of consumption data is mildly helpful for currency pricing. In particular, unlike the standard CCAPM, the new model is able to generate currency risk premiums at lower values of risk aversion and provide certain explanatory power for cross-sectional differences in currency returns.

The forward rate premium puzzle — the empirical observation of a negative relation between future changes in the spot rates and the forward premium — has been a long-standing phenomenon in international finance. Earlier studies in this area include articles by Bilson (1981), Cumby (1988), Fama (1984), Gregory and McCurdy (1984), and Hodrick and Srivastava (1984), among others.[1] The most widely accepted interpretation of the apparent predictability of currency returns is that there exists a time-varying risk premium in the foreign exchange market. However, researchers have had very limited success in reconciling intertemporal asset pricing models with these results.

The equilibrium asset pricing models based on complete market assumption, such as the representative agent consumption capital asset pricing models (CCAPMs), are able to explain with economically reasonable parameters of the risk aversion neither the time variation in risk premiums nor the large cross-sectional differences in ex post currency

This article is based in part on the author's dissertation at the University of Washington. I am indebted to Wayne Ferson for discussions and invaluable suggestions. I would like to thank George Allayannis, Jonathan Berk, Laurence Booth, Michael Cooper, Charles Engel, Steve Foerster, Avraham Kamara, Charles Nelson, Ed Rice, Joshua Rosenberg, Andrew Siegel, Kenneth Singleton (the editor), Raman Uppal, and especially the anonymous referee for valuable comments. The earlier version of this article circulated as "Heterogeneous Consumption and Asset Pricing in Global Financial Markets" was presented at the 1998 Conference on Financial Economics and Accounting at New York University and the 1999 European Finance Association Meetings in Helsinki. I also thank Richard Levich for providing me with the Harris Bank data. All remaining errors are my own. Financial support from IFM_2 is gratefully acknowledged. Address correspondence to Sergei Sarkissian, Faculty of Management, 1001 Sherbrooke St. W, Montreal H3A1G5, Canada, or e-mail: sergei.sarkissian@mcgill.ca.

[1] See Engel (1996) for a comprehensive literature review.

The Review of Financial Studies Fall 2003 Vol. 16, No. 3, pp. 983–1005, DOI: 10.1093/rfs/hhg019

returns. In these models the conditional covariances between the marginal rate of substitution in consumption and currency returns are not sufficient to generate the observed data properties. Evidence of the difficulties with the representative agent CCAPMs in the explanation of the forward premium puzzle is presented in Mark (1985), Backus, Gregory and Telmer (1993), Bekaert (1996), and Bekaert, Hodrick and Marshall (1997).

This article studies whether incomplete consumption risk sharing can provide any explanation for the forward premium puzzle. I examine the impact of imperfect consumption risk sharing across countries on the formation of the time-varying risk premium in the foreign exchange market and on cross-country differences in expected excess returns from holding foreign deposits. To address these issues I use the framework of the general equilibrium incomplete markets asset pricing model of Constantinides and Duffie (1996; hereafter CD) and data from eight industrialized countries: G7 and Switzerland. While the CD model requires data on the cross-sectional variance of individual consumption growth, the time-series of this measure is difficult to obtain. The advantage of my approach is that I am able to observe consumption dispersion directly, by virtue of the fact that it is dispersion in *aggregate* consumption across countries.

First, I show in the data that the cross-country consumption dispersion covaries negatively with the world consumption growth and currency returns. I then estimate the international version of the CD model and find that it is able to generate risk premiums at substantially lower values of risk aversion than the standard CCAPM. The new model also shortens the Hansen and Jagannathan (1997) distance. In this model, the expected currency returns are related to two risk factors: the global, systematic risk proxied by the world consumption growth, and the idiosyncratic but not diversifiable risk proxied by the cross-country variation in consumption growth. Next, using the beta-pricing framework, I show that cross-country consumption variability is useful in explaining cross-sectional differences in currency returns and, consequently, is also helpful in reducing the extent of the forward premium puzzle.

Most of the literature on imperfect risk sharing is limited to U.S. data. Earlier articles such as Telmer (1993) and Heaton and Lucas (1996) conclude that incomplete markets and individual consumption growth rates are not important for asset pricing. These articles model idiosyncratic shocks as transitory events. By contrast, CD show that if economic agents experience persistent consumption shocks then consumption dispersion across investors can have important implications for asset pricing. Storesletten, Telmer, and Yaron (2001, 2003) find that idiosyncratic risk in income not only has a transitory component but also a persistent one, which significantly decreases the agents' ability to diversify. They emphasize that the idiosyncratic risk should be measured at low frequencies, which is very convenient in tests with consumption data.[2]

The impact of cross-country consumption dispersion on currency pricing is likely to be of considerable significance if global markets are incomplete. Atkeson and Bayoumi (1993) document that risk insurance is less efficient across countries than in a single country. Bansal and Dahlquist (2000) find that country-specific attributes are more important than the systematic (i.e., world portfolio) risk in characterizing cross-country differences in currency risk premiums.[3] Indeed, investors in a given country are less subject to idiosyncratic consumption shocks because of the availability of certain common hedging mechanisms, such as unemployment insurance. This type of risk hedging is not readily available across countries, and its absence can lead to cross-country differences in optimal consumption growth rates and in expected returns on investments, including currency deposits.[4] Thus the forward premium puzzle may be better explained with a nonrepresentative agent CCAPM rather than with a representative agent model.

In a recent article, Ramchand (1999) extends the CD model to the two-country, two-good framework and, using a calibration exercise, shows that the resulting model can generate a more volatile pricing kernel. My study differs substantially from that article. First, my international version of the CD model maintains a one-good paradigm. Second, I analyze directly the performance of the CD model across countries and examine to what extent it can ease the forward premium puzzle. Third, I explicitly account for consumption heterogeneity only across countries, but not within them. Fourth, I study the impact of incomplete consumption risk sharing on cross-sectional differences in excess currency returns. Thus the main contribution of my article is that I present a more comprehensive empirical analysis of the implications of the CD model for currency pricing.

The remainder of the article is organized as follows. Section 1 extends the CD model to the international setting. Section 2 describes the data and construction of consumption-based variables, and presents summary statistics. Section 3 relates the model to the forward premium puzzle, outlines

[2] Jacobs (1999) finds no support for the CCAPM with disaggregate consumption, but Cogley (2002) and Brav, Constantinides, and Geczy (2002) provide some weak evidence regarding the importance of investors' heterogeneity for U.S. equity pricing.

[3] Lewis (1996) argues that the apparent lack of consumption risk sharing across countries can be explained by the combination of capital market restrictions and nonseparabilities in investors' utility function. However, investors need only have unrestricted access to a single asset to circumvent these restrictions and achieve consumption smoothing. Telmer (1993) and Heaton and Lucas (1996) provide more discussion on this issue.

[4] In contrast to my article, Bansal (1997) and Bansal and Dahlquist (2000) primarily emphasize the importance of different monetary uncertainties across countries for the observed differences in excess currency returns.

the econometric methodology, and presents the test results based on the Euler equation estimation. Section 4 focuses on cross-sectional characteristics of currency returns within the beta-pricing framework of the model. It also directly addresses the extent to which the forward premium puzzle is resolved by the new model. Section 5 concludes.

1. The World CCAPM with Heterogeneity

The standard canonical asset pricing relation has the following form:

$$E_t\left[m_{t+1}R_{j,\,t+1}\right]=1, \tag{1}$$

where $R_{j,t+1}$ is the gross real rate of return to an investor from holding an asset j (e.g., in country j) one period and m_{t+1} is the pricing kernel. In the CCAPM framework, Equation (1) is the first-order or Euler condition for an investor's optimization problem, while m_{t+1} is the intertemporal marginal rate of substitution in consumption (IMRS).

Constantinides and Duffie (1996) account for individual investor heterogeneity in consumption and, assuming a single-good economy, perfect markets, and the existence of permanent income shocks, derive the Euler equation for a nonrepresentative agent CCAPM:

$$\rho E_t\left[\left(\frac{C_{j,t+1}}{C_{j,t}}\right)^{-\gamma}\exp\left[\frac{\gamma(\gamma+1)}{2}d^{w}_{j,t+1}\right]R_{j,t+1}\right]=1, \tag{2}$$

where $C_{j,t}$ is the aggregate consumption in country j at time t, $d^{w}_{j,t}=\mathrm{var}_i[\ln(C_{ij,t}/C_{ij,t}C_{ij,t-1})]$ is the cross-sectional variance (dispersion) of consumption growth, C_{ij} is the consumption of investor i in country j, γ is the relative risk aversion, and ρ, $\rho\in(0,1)$ is the time preference.[5]

When investors are unable to perfectly hedge themselves against single-country consumption shocks, then consumption heterogeneity will exist not only within but also across countries. I assume that financial markets of all countries are open but not complete. In other words, investors in each country can freely trade in securities of other countries, but the set of all available assets is not sufficient to ensure full consumption insurance. I also assume that all investors are identical in their preferences.

Let C_t denote the world consumption at time t. I express investor i's consumption in country j as $C_{ij,t}=\delta_{ij,t}\delta_{j,t}C_t$ ($\delta_{ij,t}$ and $\delta_{j,t}$ are the portions of country j consumption by investor i and world consumption by country j, respectively). Similar to CD, I assume that the law of large numbers holds across investors and countries, and I specify $\delta_{ij,t}$ and $\delta_{j,t}$ as

$$\delta_{ij,t}=\delta_{ij,t-1}\exp\left(\eta_{ij,t}\sqrt{d^{w}_{j,t}}-\frac{d^{w}_{j,t}}{2}\right)\quad\text{and}\quad\delta_{j,t}=\delta_{j,t-1}\exp\left(\eta_{j,t}\sqrt{d_t}-\frac{d_t}{2}\right)$$

[5] Note that $\mathrm{var}_i[\ln(C_{ij,t}/C_{ij,t-1})]=\mathrm{var}_i[\ln((C_{ij,t}/C_{j,t})/(C_{ij,t-1}/C_{j,t-1}))]$.

Here, $\eta_{ij,t}$ and $\eta_{j,t}$ are the standard normals denoting, respectively, investor i's and country j's consumption shocks at time t, $d_t = \text{var}_j[\ln(C_{j,t}/C_{j,t}C_{j,t-1})]$ is the cross-country variation of consumption growth, and $\eta_{ij,t}$, $d^w_{j,t}$, $\eta_{j,t}$, and d_t are independent across all investors, countries, and time. Thus the Euler equation of investor i in country j is

$$\rho E_t\left[\left(\frac{C_{ij,t+1}}{C_{ij,t}}\right)^{-\gamma} R_{j,t+1}\right] = \rho E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} \exp\left(-\gamma\left(\eta_{ij,t+1}\sqrt{d^w_{j,t+1}} - \frac{d^w_{j,t+1}}{2} + \eta_{j,t+1}\sqrt{d_{t+1}} - \frac{d_{t+1}}{2}\right)\right) R_{j,t+1}\right].$$

Using the method of iterated expectations, the new aggregate Euler equation for country j is

$$\rho E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} \exp\left(\frac{\gamma(\gamma+1)}{2}(d_{t+1} + d^w_{j,t+1})\right) R_{j,t+1}\right] = 1. \tag{3}$$

The Euler equation [Equation (3)] relates asset returns to the world consumption growth as well as within- and cross-country consumption variations. In this article I will not identify within-country consumption dispersion, due to practical difficulties in obtaining this measure for each country. This implies that d cannot be identified separately from d^w_j. In terms of excess returns, Equation (3) is rewritten as

$$E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} \exp\left(k\frac{\gamma(\gamma+1)}{2} d_{t+1}\right) r_{j,t+1}\right] = 0, \tag{4}$$

where $r_{j,t+1}$ is the excess return on asset j at time $t+1$, while k is a scale factor which represents cross-sectional consumption variation above and beyond cross-country dispersion (or mismeasured cross-country dispersion).[6] Note that the magnitude of consumption dispersion affects only the level of returns. Thus the ad hoc parameter k cannot impact the risk premium induced by consumption dispersion, since the premium is determined solely by the covariance of dispersion with returns and aggregate consumption growth.

2. Data and Summary Statistics

The CD model and hence its extension [Equation (4)] apply for any assets, but in this article we focus exclusively on currencies. My data are quarterly and cover a period from 1973:2 to 1995:4, or 91 observations. The

[6] Equation (4) implies that when $k = 0$, we have the standard CCAPM of Lucas (1978); when $k > 1$, we have a representative agent in each country; when $k > 1$, we implicitly account for missing within- or cross-country consumption dispersion.

foreign exchange markets considered in this article are comprised of eight developed countries: Canada, France, Germany, Italy, Japan, Switzerland, the United Kingdom, and the United States, with the latter being the domestic (numeraire) country.[7]

The exchange rates used are Canadian dollars, French francs, German marks, Italian lira, Japanese yen, Swiss francs, and British pounds, all of which are relative to the U.S. dollar. The spot and one-month forward exchange rates are for the last Friday of the month and correspond to the *Harris Bank Weekly Review* quotes. The monthly one-month forward market return at time $t+1$ for country j is computed as $(S_{j,t+1}-F_{j,t})/S_{j,t}$. Here $S_{j,t}$ and $S_{j,t+1}$ are the spot prices in U.S. dollars of one unit of foreign currency j at times t and $t+1$, respectively, while $F_{j,t}$ is the forward price of one unit of currency j at time t to be delivered at time $t+1$. The forward premium on currency j is defined as $(F_{j,t}-S_{j,t})/S_{j,t}$. The quarterly forward premiums and currency returns are obtained by compounding the corresponding monthly values over the quarter. I do not use the end-of-quarter rates because all available consumption data are averaged over the quarter. Since foreign exchange market returns are in U.S. dollars, I use the U.S. quarterly consumer price index (CPI) from *Ibbotson Associates* to obtain real returns. I also construct the world forward premium (WFP), which is an equally weighted average of forward premiums in individual currency markets.

The seasonally adjusted real aggregate consumption data for all countries are from *National Accounts*. To arrive at the per capita consumption, the aggregate consumption for each country is divided by the quarterly population estimates. These estimates are obtained by linearly interpolating annual midyear population figures as reported in *Datastream*. I construct the world per capita consumption growth (WCG) as the gross domestic product (GDP)-weighted average of countries' real per capita consumption growth rates, which are denominated in local currency units. The motivation for this construction is as follows. If national consumption data are expressed in U.S. dollars, then, since the volatility of consumption growth rates is much lower than that of exchange rates, the time-series properties of the changes in consumption will be dominated by the changes in exchange rates. However, it is impossible to meaningfully aggregate consumption data expressed in different currency units. The aggregation of consumption growth rates rather than consumption levels helps resolve these problems. The consumption growth rates, unlike consumption levels, are unitless. Nevertheless, the aggregation must somehow reflect the relative wealth distribution across countries. The GDP weights

[7] I must point out that the CD model only holds for a continuum of heterogeneous agents. This seems natural for individuals within a country, but it is a bit more of a stretch for the eight countries examined in this article. The problem is that without the cross-country law of large numbers, C_t in the equation $C_{jt} = \delta_{jt} C_t$ is not the aggregate world consumption.

Table 1
Summary statistics (quarterly data 1973–1995)

	Mean	SD	Max	Min	Autocorrelations AC1	AC2	AC3	AC4	AC8	JB
WCG	0.0049	0.0057	0.0128	−0.0164	0.22	0.16	0.41	0.11	−0.05	85.57
lnWCD	−9.9014	0.8886	−7.1215	−11.859	0.39	0.12	0.18	0.24	0.10	6.62
C$	0.0004	0.0220	0.0566	−0.0465	0.11	−0.04	0.24	0.03	−0.03	0.46
FF	0.0076	0.0596	0.1406	−0.1240	0.15	−0.06	0.06	0.17	0.11	0.93
DM	0.0047	0.0638	0.1487	−0.1191	0.08	−0.16	0.18	0.23	0.07	0.99
L	0.0074	0.0583	0.1419	−0.1848	0.14	−0.12	0.02	0.19	0.09	5.76
Y	0.0076	0.0633	0.1753	−0.1502	0.19	−0.08	0.07	0.15	0.02	1.03
SF	0.0060	0.0736	0.1749	−0.1445	0.05	−0.07	0.11	0.15	−0.01	1.69
£	0.0036	0.0557	0.1532	−0.1343	0.19	−0.11	0.17	0.04	0.05	0.56
WFP	−0.0023	0.0067	0.0150	−0.0262	0.70	0.46	0.44	0.49	0.18	7.01

The reported values are the sample mean, standard deviation (SD), maximum, minimum, autocorrelations (AC1, ..., AC8), and Jarque–Bera statistic (JB) for consumption and currency returns data for the G-7 countries and Switzerland. The sample consists of 91 observations. WCG is the world per capita consumption growth, lnWCD is the log of the cross-country consumption dispersion, WFP is the world forward premium. The quarterly currency returns and forward premiums are obtained by compounding the corresponding monthly values over the quarter. The monthly forward market return on currency j is computed as $(S_{j,t+1} - F_{j,t})/S_{j,t}$, where $S_{j,t}$ and $S_{j,t+1}$ are the spot prices in U.S. dollars of one unit of foreign currency j at times t and $t+1$ respectively, while $F_{j,t}$ is the forward price of one unit of currency j at time t to be delivered at time $t+1$. The monthly forward premium on currency j is defined as $(F_{j,t} - S_{j,t})/S_{j,t}$. The WFP is the equally weighted average of all forward premiums.

denominated in U.S. dollars attach more value to the consumption growth of the countries with higher GDP and less value to those with lower GDP. Thus the exchange rate fluctuations affect the WCG only indirectly — through the U.S. dollar-denominated GDP.[8] The cross-country (world) consumption dispersion (WCD) is calculated as the variance of the log changes in the countries' real per capita consumption growth rates in local currency units.

Table 1 reports the summary statistics of the data, including the means, standard deviations, minimum and maximum values, autocorrelations, and the Jarque–Bera statistics for normality. It shows that the first and second moments of the WCG, that is, the mean growth rate of 0.0049 and the standard deviation of 0.0057, are close to the corresponding values for the U.S. data, 0.0041 and 0.0069, respectively. The WCG has a relatively large third-order autocorrelation. Meanwhile, the largest autocorrelation for the WCD is first order. All excess currency returns are small but positive, on average, and, while the Canadian dollar returns have the lowest mean, the Japanese yen and French franc returns have the highest. Returns on the Italian lira exhibit the largest deviation from normality with the Jarque–Bera statistic of 5.76. The autocorrelation of returns is between 0.05 and 0.2, but exceeds 0.7 for the WFP.

[8] The method of constructing some world growth measure as a GDP-weighted average of the countries' growth rates expressed in local currency units has been used in the literature before [e.g., see Harvey (1990)].

Table 2
Unconditional cross-correlations

	Contemporaneous cross-correlations								Lagged cross-correlations		
	WCD	C$	FF	DM	L	Y	SF	£	WCG_{-1}	WCD_{-1}	WFP_{-1}
WCG	−0.31	−0.11	−0.02	−0.09	0.04	−0.05	−0.14	−0.05	0.22	−0.06	0.03
WCD	1	0.13	−0.17	−0.15	−0.16	0.01	−0.18	−0.08	−0.18	0.40	−0.02
C$		1	0.04	0.06	0.03	0.08	0.08	0.13	0.03	0.15	0.13
FF			1	0.93	0.81	0.62	0.85	0.66	0.08	−0.11	−0.18
DM				1	0.72	0.63	0.88	0.66	0.09	−0.13	−0.20
L					1	0.51	0.67	0.60	0.07	0.01	−0.14
Y						1	0.64	0.52	−0.03	−0.06	−0.28
SF							1	0.63	0.13	−0.15	−0.18
£								1	0.07	−0.07	−0.14

WCG_{-1} is the lagged world consumption growth. WCD_{-1} is the lagged world consumption dispersion. WFP_{-1} is the lagged world forward premium.

Table 2 reports cross-correlations. Its left-hand side shows the contemporaneous correlations. The correlation of the WCG with currency returns is extremely weak. More importantly, the correlation of the WCD with currency returns is moderately negative for five of the seven currencies. Cross-correlations among closely linked currencies are, as expected, very high: for example, as high as 0.93 between the German mark and the French franc. The returns on the Canadian dollar have the lowest correlations with all the other currency returns. Finally, the correlation between WCG and WCD is −0.3.

The right-hand side of Table 2 shows the correlations with the lagged values of the three major variables. These variables are the one-quarter lags of the WCG, WCD, and WFP. The correlation of the WCG with WFP and all currency returns except for the yen is positive. The correlation of consumption dispersion with the WFP is positive, whereas it is negative with almost all currency returns, except for the Canadian dollar. The correlation of the lagged WCD and excess returns on the Italian lira is close to zero. Thus the WCD, partially because of its moderate level of persistence, shows a similar correlation pattern with currency returns at both contemporaneous and lagged levels. The other interesting feature of the table is that the lagged WFP, which is usually negatively related to returns in currency markets, exhibits sizable positive correlation again with the excess returns on the Canadian dollar.

Figure 1 shows the time series of the world consumption growth and dispersion. For comparison, it also depicts the U.S. per capita consumption growth. Consistent with the sizable negative correlation between WCG and WCD pointed out above, these plots illustrate that dispersion tends to be higher when growth is low, and vice versa. For example, the three largest peaks of the WCD occur within the periods of worldwide recessions of 1973–1975, 1980–1981, and the early 1990s. This relation between the dispersion and business cycles coincides with the intuition of

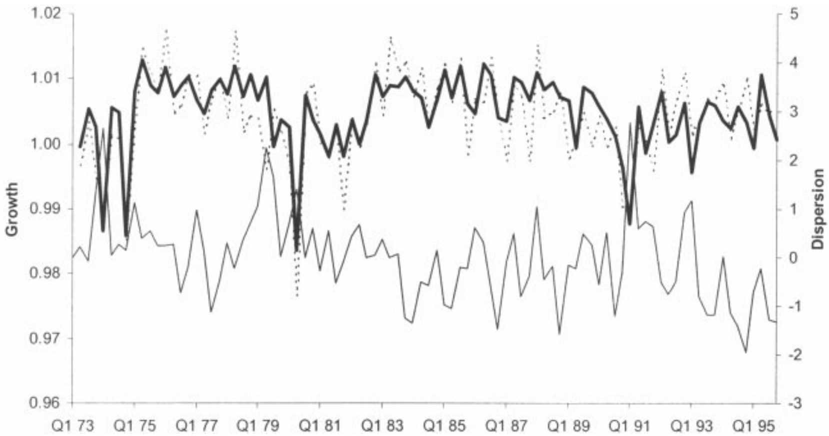

Figure 1
The world consumption growth and consumption dispersion
This figure shows the world and the U.S. quarterly real per capita consumption growth and the cross-country dispersion of the real per capita consumption growth rates. The world consumption growth rate (bold solid line) is the GDP-weighted average of the real per capita consumption growth rates from period t to period $t+1$ for Canada, France, Germany, Italy, Japan, Switzerland, the United Kingdom, and the United States, all of which are expressed in local currency units. The U.S. consumption growth rate is the dashed line. The demeaned log transformation of the cross-country dispersion of the real per capita consumption growth rates from period t to period $t+1$ expressed in local currency units is shown with the thin solid line.

Mankiw (1986), Constantinides and Duffie (1996), and Storesletten, Telmer, and Yaron (2001), who point out that uninsurable risk may have a measurable impact on asset returns only if it is inversely related to aggregate shocks. Thus the patterns in the data, in particular a negative relation between consumption dispersion and growth, as well as excess returns, suggest that empirical work may meet with some success.

3. Tests of the Euler Equations

3.1 Relation to the forward premium puzzle

In its essence, the forward premium puzzle is similar to a conditional version of the equity premium puzzle — the finding of large and variable conditional risk premiums.[9] In the case of excess currency returns, $r_{j,t+1} = [(S_{j,t+1} - F_{j,t})/(S_{j,t}\pi_{t+1})]$ is the real return at time $t+1$ from the forward position in currency j (π_{t+1} is the one-period gross inflation rate in a home country). However, unlike equity returns, the average excess returns (unconditional risk premiums) in currency markets are approximately zero, that is, $E[r_{j,t+1}] \approx 0$. To produce higher conditional risk premiums,

[9] See Mehra and Prescott (1985) for a discussion of the equity premium puzzle.

Equation (4) must generate a more variable IMRS or a higher covariance between IMRS and currency returns, or both, than the standard CCAPM. For simplicity, let $K = 0.5k\gamma(\gamma+1)$. Since $\exp(Kd_{t+1}) \geq 1$ for any positive values of γ and k, and as data suggest, $(C_{t+1}/C_t)^{-\gamma}$ and $\exp(Kd_{t+1})$ are not negatively correlated, Equation (4) will produce a higher conditional variation of IMRS than the standard CCAPM. Furthermore, consumption dispersion has a sizable correlation with excess currency returns. This should also produce a higher conditional covariance between IMRS and those returns. As a result, the time variation in currency risk premiums can now be attributed to the time variation in both WCG and WCD.

Thus we can expect that the new CCAPM will be able to generate time-varying currency risk premiums with much lower values of the risk aversion. However, to resolve the forward premium puzzle, it must also be able to explain why forward premium is such a good predictor of excess currency returns. This means that the new model must be able to generate risk premiums with specific characteristics. I discuss this issue in detail in the next section.

3.2 Estimation procedure

In the Euler equation [Equation (4)], there is one parameter of interest — the relative risk aversion, γ. I use Hansen's (1982) GMM methodology and estimate the model together with mean pricing errors. The transformed Euler equation takes the form

$$E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\exp\left(k\frac{\gamma(\gamma+1)}{2}d_{t+1}\right)(r_{j,t+1}-\alpha_j)\right]=0, \qquad (5)$$

where α_j is the mean pricing error for currency j, that is, the average deviation between the excess currency return and the excess currency return predicted by the model. These asset-specific coefficients are similar to the Jensen's alpha measures in the excess return formulations of beta pricing models.[10] Therefore an alternative way to examine the validity of our Euler equation is to test whether the α_j's equal zero. It is important to note that even though the average return on forward positions is zero, Table 1 shows that in our sample the averages of all currency returns are greater than zero.

In effect, the Euler equation [Equation (5)] defines the error term

$$u_{j,t+1} = (C_{t+1}/C_t)^{-\gamma}\exp\left(k\frac{\gamma(\gamma+1)}{2}d_{t+1}\right)(r_{j,t+1}-\alpha_j),$$

[10] To see this, we can rewrite Equation (5) in a shorthand form as $E_t[m_{t+1}(r_{j,t+1}-\alpha_j)]=0$. This is equivalent to: $E_t[r_{j,t+1}] = \alpha_j - \text{cov}_t(m_{t+1}r_{j,t+1})/E_t[m_{t+1}]$, which is a 'beta representation," and so α_j is like a "Jensen's alpha."

for each currency return j. Thus $E[u_{j,t+1}] = 0$ and $E[u_{j,t+1}\boldsymbol{Z}_t'] = 0$ for all returns, where $\boldsymbol{Z}_t$ is the set of L instruments which are assumed to be known to the market at time t. Since the number of returns, N, is seven, the model is overidentified when L exceeds two. To derive consistent and asymptotically efficient GMM estimators, I assume that all explanatory variables in Euler equations are strictly covariance stationary. Since consumption data are time averaged, in the GMM tests I use the Newey and West (1987) hetero- and autoconsistent matrix with the first-order moving average term.

As in any GMM-based test, a careful selection of the instrument vector is important. Due to the small sample size, having too many instruments is not desirable: a large number of instruments can lead to finite sample biases [e.g., see Andersen and Sorensen (1996)]. Since my work focuses on CCAPMs, it is natural to include the lagged WCG in the instrument set. For example, Hall (1978) and Hansen and Singleton (1983) show that lagged consumption growth is useful in predicting future U.S. consumption growth. The lagged WCD is another choice for an instrument, given the results of Table 2. However, to simplify the comparison of the test results between the new and standard CCAPMs, I only report outcomes with no lagged WCD in the information set. The final instrument is also natural: the lagged WFP. With seven moment conditions and three instruments, there are 21 orthogonality conditions.[11]

After the estimation of the model, I compute the mean and standard deviation of the implied pricing kernel and use the Hansen and Jagannathan distance [see Hansen and Jagannathan (1997)] to compare the performance of the model for various values of k. The goal here is to find the shortest distance between the estimated pricing kernel $\hat{m}$ and the true pricing kernel m when $E(m\boldsymbol{R}-\boldsymbol{1}|\boldsymbol{Z})$ is indeed equal to zero. The distance between $\hat{m}$ and m can be measured as: $[E(\hat{m}\tilde{\boldsymbol{r}}) - 1]'E(\tilde{\boldsymbol{r}}\tilde{\boldsymbol{r}}')^{-1}[E(\hat{m}\tilde{\boldsymbol{r}}) - 1]$, where $\boldsymbol{1}$ is the $1 \times NL$ vector of ones, $\tilde{\boldsymbol{r}} = \boldsymbol{R} \otimes (\boldsymbol{Z}/E(\boldsymbol{Z}))$ is the $N \times L$ vector of augmented returns, $\otimes$ is the row-wise Kroneker product, and is the $E(\boldsymbol{Z})$ unconditional mean of the instrument vector.

3.3 Test results

To give the first perspective on the behavior of the pricing kernel in the new CCAPM, Figure 2 shows the Hansen and Jagannathan (1991) volatility bounds for gross currency returns. It also shows the same bounds adjusted for the small sample size [see Ferson and Siegel (2003)]. The bounds are quite tight: the standard deviation of unity is found approximately within

[11] The inclusion of all individual forward premiums in the instrument set leads to a singular weighting matrix due to high auto- and cross-correlation of these variables.

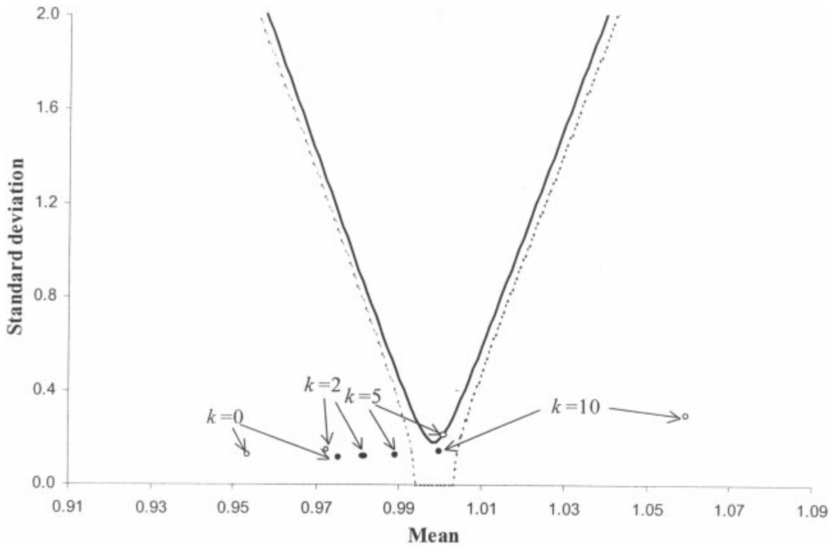

Figure 2
The Hansen and Jagannathan volatility bounds with bias correction
The plot shows the Hansen and Jagannathan lower bounds for volatility for forward currency market returns (solid curve). It also shows the bias-adjusted bounds (dashed curve). The filled and unfilled dots depict the sample mean and standard deviation of the implied pricing kernel when the risk aversion parameter, γ, equals 5 and 10, respectively. The results are shown for the values of $k = 0$, 2, 5, and 10.

the mean values of 0.98 and 1.02.[12] The implied mean-standard deviation pairs of the pricing kernel are shown for the values of $\gamma = 5$ or 10 and for $k = 0, 2, 5$, and 10. Note that when $\gamma = 10$, the implied pricing kernel, while still outside the standard mean-standard deviation "cup," appears to be inside the bias-adjusted bounds when k is approximately between 2 and 5. This implies that accounting for only cross-country consumption dispersion from the eight countries in our sample is not sufficient to satisfy the bound. The additional variation must come from within-country or perhaps the omitted cross-country dispersion. Nevertheless, the plot provides initial evidence that the new CCAPM is likely to outperform the standard model.

I estimate the Euler equation [Equation (5)] for the standard CCAPM, by setting $k = 0$, and for the new CCAPM by setting $k = 1$, 2, 5, and 10. The instrument set is composed of a constant and the lagged values of WCG and WFP. Table 3 shows the results. Note that while the absolute numbers of parameter estimates are of certain interest to us, we are more concerned with the incremental effect of consumption dispersion on the

[12] Similarly, Backus, Gregory, and Telmer (1993) and Cecchetti, Lam, and Mark (1994) report that the Hansen and Jagannathan volatility bound for currency returns is tighter than that for equity returns.

Table 3
Tests of the Euler equations

		Average pricing errors									Pricing kernel		
k	γ	$\alpha_{C\$}$	α_{FF}	α_{DM}	α_{L}	α_{Y}	α_{SF}	$\alpha_{£}$	J-statistic	p	Mean	SD	HJD
0	118.89	0.0015	0.0219	0.0203	0.0132	0.0265	0.0211	0.0170	14.10	0.367	0.7805	1.0847	0.6078
	(33.65)	(0.0016)	(0.0075)	(0.0084)	(0.0053)	(0.0051)	(0.0126)	(0.0049)					
1	23.33	−0.0005	0.0160	0.0123	0.0098	0.0209	0.0150	0.0098	15.74	0.263	0.9235	0.1673	0.4562
	(34.80)	(0.0023)	(0.0058)	(0.0062)	(0.0058)	(0.0061)	(0.0076)	(0.0058)					
2	13.28	−0.0007	0.0153	0.0115	0.0091	0.0206	0.0142	0.0091	16.03	0.247	0.9549	0.0937	0.4501
	(37.34)	(0.0023)	(0.0056)	(0.0060)	(0.0058)	(0.0061)	(0.0074)	(0.0058)					
5	5.70	−0.0007	0.0149	0.0109	0.0087	0.0201	0.0136	0.0087	16.26	0.235	0.9805	0.0401	0.4475
	(39.52)	(0.0023)	(0.0055)	(0.0058)	(0.0057)	(0.0061)	(0.0072)	(0.0058)					
10	2.75	−0.0009	0.0148	0.0108	0.0085	0.0204	0.0134	0.0086	16.32	0.232	0.9908	0.0196	0.4470
	(40.41)	(0.0024)	(0.0055)	(0.0058)	(0.0057)	(0.0060)	(0.0071)	(0.0058)					

The GMM estimates of the following model are reported:

$$E\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\exp\left[k\frac{\gamma(\gamma+1)}{2}d_{t+1}\right]\left(\frac{S_{j,t+1}-F_{j,t}}{S_{j,t}\pi_{t+1}}-\alpha_j\right)\middle|Z_t\right]=0.$$

Here C_{t+1}/C_t is the WCG, d_t is the WCD, Z_t is the instrument vector composed of a constant and the lagged values of WCG, WCD, and WFP, γ is the risk aversion parameter, α_j is the average pricing error for currency j, k is the multiplicative factor, π_{t+1} is the gross inflation rate. The description of $S_{j,t}$, $S_{j,t+1}$, and $F_{j,t}$ is in Table 1. The case $k=0$ defines the standard CCAPM framework. For each test, the table also reports the standard errors of parameters (in parentheses), the goodness-of-fit J-statistic with its p-value (p) as well as the mean, standard deviation (SD), and the Hansen and Jagannathan distance (HJD) measure of the estimated pricing kernel. The HJD measure is multiplied by 100.

resolution of the forward premium puzzle. The estimated risk aversion is lower when $k > 0$ than when $k = 0$: γ is about 118.9 for the standard CCAPM, but decreases to 5.7 at $k = 5$, although its precision does not decrease.[13] The estimates of the average pricing errors, and most of their own standard errors, decrease once we account for the WCD. Moreover, for the standard CCAPM, the α_j's are insignificant only for the Canadian dollar, but for the new CCAPM, they are insignificant also for the mark, the lira, the Swiss franc, and the pound. We can also observe that the standard CCAPM produces a pricing kernel with an unrealistically low mean, 0.7805, but with a high standard deviation, 1.0847. The mean of the implied pricing kernel is more reasonable at $k = 5$. Indeed, from Table 1, we can see that the average mean currency return is 0.0053, implying the mean of 0.9950 for the pricing kernel, while the average standard deviation is 0.0566. At $k = 5$, the corresponding values, 0.9805 and 0.0401, are much closer to the real data than when $k = 0$.

It appears that the improvement in parameter estimates for the new CCAPM negatively impacts the overall fit of the model: the *J*-statistic somewhat increases with k and the variability of the pricing kernel decreases. This observation is misleading however: a higher *p*-value of the *J*-statistic for the standard CCAPM is solely the result of a relatively higher volatility of the pricing kernel. The Hansen and Jagannathan distance shows that the overall fit of the new model is in fact much better than that of the standard CCAPM. The difference between the implied pricing kernel and the true one decreases in k.[14] The above results provide certain support for the relevance of both incomplete consumption risk sharing and investors' heterogeneity across countries in explaining the time variation of currency risk premiums. This also implies that the new CCAPM is likely to explain the cross-sectional differences in currency returns across time and on average better than the standard model.

4. The Beta-Pricing Framework

4.1 Formulation and implications

The above results show that the CCAPM with consumption heterogeneity is able to reconcile the lower values of the risk aversion with the estimated

[13] With the U.S. data, the estimate of γ for the standard CCAPM in Mark (1985) and Backus, Gregory, and Telmer (1993) is around 50. However, since the world consumption growth is less variable than the U.S. growth due to averaging, having a larger estimate of the risk aversion for the standard CCAPM is not surprising.

[14] The main distinction between the Hansen and Jagannathan distance and chi-square statistic (*J*-statistic) is that the former does not merely reward the variability of the pricing kernel. In other words, ceteris paribus, a pricing kernel with higher standard deviation is less likely rejected by the chi-square test, though it may still have the same Hansen and Jagannathan distance as the one with lower standard deviation.

risk premiums on currency returns. In this section I analyze whether the two risk factors — the world consumption growth and dispersion — can account for differences in currency risk premiums across countries in our sample.[15] I also examine whether the new model can shed light on the apparent predictability of currency returns — the widely observed negative relation between changes in exchange rates and respective forward premiums.

Following Ferson (1983), Hansen and Singleton (1983), Campbell (1993), and others, I derive an approximate beta pricing relation by assuming that the joint conditional distribution of consumption growth, dispersion, and asset returns is lognormal but not necessarily homoscedastic. For simplicity, I do not distinguish here between the distributional properties of d_t and $\exp(d_t)$. The difference in lognormal approximations of the Euler equation [Equation (4)] expressed for uncovered and covered positions in the foreign exchange market yields the following linear conditional asset pricing relation:

$$E_t[s_{j,t+1}] - f_{j,t} \approx -0.5\sigma^2_{j,t} + \gamma\sigma_{jc,t} - K\sigma_{jd,t}. \tag{6}$$

Here $s_{j,t+1} = \ln(S_{j,t+1}) - \ln(S_{j,t})$ and $f_{j,t} = \ln(F_{j,t}) - \ln(S_{j,t})$ are the log returns at time $t+1$ for taking, respectively, uncovered and covered positions in currency j at time t, $\sigma_{jc,t} = \mathrm{cov}_t((s_{j,t+1} - \ln\pi_{t+1}), c_{t+1})$ and $\sigma_{jd,t} = \mathrm{cov}_t((s_{j,t+1} - \ln\pi_{t+1}), d_{t+1})$ are the conditional covariances of returns with WCG and WCD, respectively, $\sigma^2_{j,t} = \mathrm{var}_t(s_{j,t+1} - \ln\pi_{t+1})$ is the conditional variance of return, and $c_{t+1} = \ln(C_{t+1}) - \ln(C_t)$. Equation (6) is equivalent to the following beta pricing formulation:

$$E_t[s_{j,t+1}] - f_{j,t} \approx \lambda_{0,t} + \lambda_{c,t}\beta_{jc,t} + \lambda_{d,t}\beta_{jd,t}, \tag{7}$$

where $\lambda_{0t} = -0.5\sigma^2_{j,t}$ is a Jensen's inequality term, $\lambda_{c,t} = \gamma\sigma^2_{c,t}$ and $\lambda_{d,t} = -K\sigma^2_{d,t}$ are time-varying coefficients that can be thought of as prices of risk for the WCG and WCD, respectively; $\sigma^2_{c,t}$ and $\sigma^2_{d,t}$ are the conditional variances of the WCG and WCD. In this equation, $\beta_{jc,t} = \sigma_{jc,t}/\sigma^2_{c,t}$ is the well-known consumption growth beta similar to that in Breeden (1979), while the new ratio, $\beta_{jd,t} = \sigma_{jd,t}/\sigma^2_{d,t}$, will be called the consumption dispersion beta. Thus Equation (7) is in the spirit of asset pricing models of Merton (1973) and Ross (1976). In this specification, the return on currency j speculation is determined by its covariance with the two state variables: WCG and WCD.

Equation (7) allows us to make two observations. First, it implies that each currency risk premium, at any time t, should be proportional to $\lambda_{c,t}$ and $\lambda_{d,t}$. The averages of $\lambda_{c,t}$ and $\lambda_{d,t}$ provide information on the average

[15] A similar exercise can be found, for example, in Bansal and Dahlquist (2000), who examine cross-sectional differences in expected excess currency returns with respect to the set of country-specific risk factors.

sensitivities of expected currency returns to the world consumption growth and dispersion risk within the given sample of data. Bansal and Dahlquist (2000) find that systematic risk does not explain cross-sectional differences in currency returns. In my setup, this is equivalent to saying that differences in excess currency returns should be largely related to the world consumption dispersion risk.

Second, Fama (1984) interprets the forward premium on any currency j, $f_{j,t} - s_{j,t}$, as the sum of the risk premium, $f_{j,t} - E_t[s_{j,t+1}]$, and the expected depreciation rate, $E_t[s_{j,t+1}] - s_{j,t}$. He shows that the negative slope from the regression of the depreciation rate of currency j, $s_{j,t+1} - s_{j,t}$, on the forward premium, $f_{j,t} - s_{j,t}$, has two implications: first, the risk premium must have a negative covariance with the expected depreciation rate, and second, that

$$\begin{aligned} \mathrm{var}(f_{j,t} - E_t[s_{j,t+1}] &> \mathrm{cov}|(f_{j,t} - E_t[s_{j,t+1}], E_t[s_{j,t+1}] - s_{j,t})| \\ &> \mathrm{var}(E_t[s_{j,t+1}] - s_{j,t}). \end{aligned} \tag{8}$$

I refer to these requirements as Fama's necessary conditions. In my setup, they imply that to resolve the forward premium puzzle, the right-hand side of Equation (7) taken with the opposite sign must have negative covariance with and be more volatile than the expected depreciation rate. Since the first two terms on the right-hand side of Equation (7) are the same in the representative agent model, the primary contribution for resolving or lessening the puzzle should come from the third term, $\lambda_{d,t}\beta_{jd,t}$.

4.2 Estimation procedure and test results

I first test whether $\beta_{jc,t}$ and $\beta_{jd,t}$ are constant for each country j, adapting the approach of Ferson and Harvey (1993) to our purposes. Assuming that the conditional expectation of risk factors is linear in $\boldsymbol{Z}$, that is, $E_t[\boldsymbol{F}_{t+1}] = \boldsymbol{Z}_t\eta$, where $\boldsymbol{F}_{t+1} = [\mathrm{WCG}_{t+1}, \mathrm{WCD}_{t+1}]$, and η is the $L \times 2$ coefficient matrix, I define a disturbance vector, $\boldsymbol{u}1_{j,t+1} = \boldsymbol{F}_{t+1} - \boldsymbol{Z}_t\eta$, for each asset j at time t. The additional error term can be obtained from the definition of conditional beta,

$$\begin{aligned} \beta_{j,t} &= \mathrm{cov}(r_{j,t+1}, \boldsymbol{F}_{t+1}|\boldsymbol{Z}_t)\mathrm{var}(\boldsymbol{F}_{t+1}|\boldsymbol{Z}_t)^{-1} \\ &= \mathrm{cov}(r_{j,t+1}, \boldsymbol{u}1_{j,t+1}|\boldsymbol{Z}_t)\mathrm{var}(\boldsymbol{u}1_{j,t+1}|\boldsymbol{Z}_t)^{-1}. \end{aligned}$$

Since my null hypothesis is that $\beta_{j,t} = [\beta_{jc,t}, \beta_{jd,t}]$ is time invariant, the second error term can be written as $\boldsymbol{u2}_{j,t+1} = \beta_j(\boldsymbol{u1}'_{j,t+1}\boldsymbol{u1}_{j,t+1}) - \boldsymbol{u1}_{j,t+1}r_{j,t+1}$. Thus we have a system of equations:

$$\begin{cases} \boldsymbol{u}1_{j,t+1} = \boldsymbol{F}_{t+1} - \boldsymbol{Z}_t\eta \\ \boldsymbol{u}2_{j,t+1} = \beta_j(\boldsymbol{u}1'_{j,t+1}\boldsymbol{u}1_{j,t+1}) - \boldsymbol{u}1_{j,t+1}r_{j,t+1}. \end{cases} \tag{9}$$

Table 4
Tests for the time-variation in consumption growth and dispersion betas

	C$	FF	DM	L	Y	SF	£
β_c	−0.3749	−1.4232	−2.1567	−0.1715	−0.3175	−3.6912	−1.2283
	(−1.12)	(−1.39)	(−1.97)	(−0.16)	(−0.25)	(−2.64)	(1.20)
β_d	0.0030	−0.0114	−0.0122	−0.0070	−0.0044	−0.0182	−0.0066
	(1.17)	(−1.99)	(−1.89)	(−1.15)	(−0.68)	(−2.49)	(−1.07)
J-statistic	0.53	1.48	1.47	2.70	3.95	2.28	3.45

The GMM estimates of the following model are reported:

$$\begin{cases} \boldsymbol{u1}_{j,t+1} = \boldsymbol{F}_{t+1} - \boldsymbol{Z}_t \eta \\ \boldsymbol{u2}_{j,t+1} = \beta_j(\boldsymbol{u1}'_{i,t+1}\boldsymbol{u1}_{j,t+1}) - \boldsymbol{u1}_{j,t+1} r_{j,t+1}. \end{cases}$$

Here $\beta_j = [\beta_{jc}, \beta_{jd}]$ is a set composed of the world consumption growth and dispersion betas, respectively, for each currency j, J-statistic is the goodness-of-fit J-statistic. The t-statistics are in parentheses. $\boldsymbol{F}$ is a vector composed of the WCG and WCD, $\boldsymbol{Z}$ is the instrument set, η is the coefficient vector.

As usual, $E[\boldsymbol{u}_{j,t+1}] = 0$ and $E[\boldsymbol{u}_{j,t+1}\boldsymbol{Z}'_t] = \boldsymbol{0}$, where $\boldsymbol{u}_{j,t+1} = [\boldsymbol{u1}_{j,t+1}, \boldsymbol{u2}_{j,t+1}]$. Equation (9), which is overidentified as long as the vector $\boldsymbol{Z}$ has at least one time-varying component, is estimated by the GMM separately for each currency return j.

Table 4 shows the consumption growth and dispersion betas, their t-statistics, and the Hansen's goodness-of-fit J-statistics from estimating Equation (9) for every excess currency return. The instrument vector $\boldsymbol{Z}$ is again composed of a constant as well as the lagged values of the WCG and WFP. This implies 8 parameters, 12 orthogonality conditions, and 4 degrees of freedom in the GMM estimation. The reported J-statistics show that the model cannot be rejected for any currency since the 5% critical value for the chi-square test at four degrees of freedom is 9.49. Thus we cannot reject the hypothesis that the betas, with respect to the WCG and WCD, are conditionally constant.

Given the assumption that $\beta_{jc,t}$ and $\beta_{jd,t}$ are constant, we can estimate the risk premiums $\lambda_{c,t}$ and $\lambda_{d,t}$ using the following cross-sectional regression:

$$s_{j,t+1} - f_{j,t} = \lambda_{0,t} + \lambda_{c,t}\hat{\beta}_{jc} + \lambda_{d,t}\hat{\beta}_{jd} + e_{j,t}, \tag{10}$$

where $e_{j,t}$ is the residual for currency j. To examine the marginal importance of the third term, $\lambda_{d,t}\beta_{jd,t}$, in the beta pricing relation of Equation (7), I also estimate Equation (10) based on the standard CCAPM, that is, when the consumption dispersion risk $\beta_{jd,t} = 0$ for all j. In this case I assign the variable $\boldsymbol{F}$ to be the WCG and compute β_{jc}'s by estimating Equation (9) jointly across all seven currencies. When both risk factors are considered, the betas from Table 4 are used.[16]

[16] It is preferable to estimate the betas and risk premiums simultaneously to avoid the error-in-variables problem. Using mimicking portfolios, I have also estimated these coefficients simultaneously but found no qualitative difference from the results reported in the article. These results, along with those based on the joint estimation of betas across all currencies using only the world consumption growth, are available on request.

Table 5
Tests for the cross-sectional differences in excess currency returns

	Risk premiums			
λ_0	λ_c	λ_d	R^2	F-statistic
0.0056	0.0004		0.02	
(1.28)	(0.17)			
0.0042	0.0016	−0.3113	0.20	4.29
(1.26)	(0.62)	(−0.80)		[0.08]

The time-series averages of the risk premiums, the Fama and MacBeth t-statistics (in parentheses), and adjusted R^2 are from the following cross-sectional regression model:

$$s_{j,t+1} - f_{j,t} = \lambda_{0,t} + \lambda_{c,t}\hat{\beta}_{jc} + \lambda_{d,t}\hat{\beta}_{jd} + e_{j,t}.$$

Here $s_{j,t+1}$ and $f_{j,t}$ are the log returns at time $t+1$ for taking, respectively, uncovered and covered positions in currency j at time t, $\hat{\beta}_{jc}$ and $\hat{\beta}_{jd}$ are the estimates of the world consumption growth and dispersion betas, respectively, for each currency j. When both risk factors, the WCG and WCD, are considered, the estimated betas are from Table 4. When only the WCG is considered, the estimated $\hat{\beta}_{jc}$ are from the joint estimation of Equation (9) across all currencies (see the text for details). F-statistic is the mean value of the F-statistics (its p-value is in the square brackets).

Table 5 shows the time-series averages of the cross-sectional coefficients, their Fama and MacBeth (1973) t-statistics, and adjusted R^2. When both the WCG and WCD are used as risk factors, the table reports the average estimated F-statistic as well. The average adjusted R^2 is barely 2% when we account for only the WCG risk, but it is much higher, 20%, when both risk factors are present. Moreover, while the t-statistics are not significant at any conventional level for either consumption growth or dispersion, the average F-statistic of 4.29 is marginally significant: its p-value is less than 0.1. Thus it appears that the WCD adds a substantial explanatory power for the cross-section of excess currency returns. Given this result, we now turn our attention to the extent to which the new model is helpful in resolving the forward premium puzzle.

Table 6 compares the performance of the new CCAPM to the standard one in light of Fama's conditions. The first two columns show how deep the forward premium puzzle is for different currencies: that is, how negative the slope coefficient is from the regression of the depreciation rate on the corresponding forward premium. In these regressions, to be consistent with our construction of the forward premium, we compute the depreciation rate on each currency as the compounded change in the monthly exchange rate changes over the quarter. The regression results show that, similar to other articles, the estimates of all slopes except for the Italian lira are negative, although insignificant.[17] I report the variances of the risk premium and the expected depreciation rate, as well

[17] Due to compounding, my estimates of the slopes are generally less negative than those found in other studies; however, qualitatively they are similar. For example, Backus, Foresi, and Telmer (2001) also find that the slope for the Italian lira is relatively more positive than for other major currencies.

Table 6
The extent of resolving the forward premium puzzle

			Standard CCAPM			New CCAPM		
	θ	$se(\theta)$	var(p)	cov(p,q)	var(q)	var(p)	cov(p,q)	var(q)
C$	−0.36	0.54	0.0020	−0.0018	0.0034	0.0583	−0.0578	0.0591
FF	−0.45	0.75	0.0073	−0.0054	0.0106	0.1750	−0.1659	0.1638
DM	−0.53	0.79	0.0639	−0.0596	0.0623	0.1691	−0.1629	0.1637
L	0.47	0.54	0.0178	−0.0184	0.0329	0.1270	−0.1194	0.1257
Y	−0.41	0.42	0.0003	−0.0001	0.0208	0.0364	−0.0360	0.0565
SF	−0.60	0.58	0.1852	−0.1696	0.1700	0.3859	−0.3630	0.3560
£	−1.18	0.75	0.0023	−0.0021	0.0077	0.0490	−0.0466	0.0501

The coefficient θ is the slope estimate from the regression of the currency depreciation rate, $s_{t+1} - s_t$, on the corresponding forward premium, $f_t - s_t$, $se(\theta)$ is the standard error of θ. The depreciation rate on each currency is computed as the compounded change in the monthly exchange rate changes over the quarter. The other notation is as follows: var(p) is the variance of the currency risk premium, var(q) is the variance of the expected depreciation rate, cov(p,q) is the covariance between the risk premium and the expected depreciation rate. The risk premium, $p_t = f_t - E_t[s_{t+1}]$, is the corresponding expected excess return, taken with the negative sign, generated using model Equation (10), as described in Table 5. The expected depreciation rate, $q_t = E_t\,[s_{t+1}] - s_t$, is computed as the difference between the forward premium and the estimated risk premium. The variances and covariances are multiplied by 100.

as the covariance of these two series within both the standard and new CCAPM frameworks in columns three through eight. The risk premium on each currency is the corresponding expected excess return, taken with the negative sign, generated using Equation (12), as described above. The expected depreciation rate is the difference between the forward premium and the estimated risk premium.

First of all, the variances of risk premiums are markedly higher and the covariances of risk premiums with depreciation rates are substantially more negative for the new model than for the standard one. Second, under the standard CCAPM, the variance of the risk premium is larger than the corresponding depreciated rate for only two currencies, while under the new model it is for four. More importantly, the variance of the expected depreciation rate is greater than the absolute value of its covariance with the risk premium across all currencies for the standard CCAPM. This implies that Fama's second condition, that is, Equation (8), fails for all currencies. However, for the new model, the disparity in these two measures is smaller in magnitude and, in the cases of the French and Swiss francs, the picture is even reversed, implying that in these two instances Equation (8) holds.[18] Thus, even though Fama's conditions are not fully satisfied for all currencies, the new CCAPM is able to generate currency risk premiums with properties that are partially aligned with those observed in the data.

[18] Note that even when Equation (8) holds for a given risk premium, it only implies that this risk premium can explain the negative sign of the slope coefficient in the regression of the depreciation rate on the forward premium. It does not necessarily imply that this risk premium fully explains the *magnitude* of the slope [e.g., see Fama (1984), p. 327].

What is the reason that the new model is unable to fully resolve the forward premium puzzle? In the framework of affine models, Backus, Foresi, and Telmer (2001) show that any solution to the puzzle requires that the state variables have an asymmetric effect on interest rates in different countries, or that the assumption of strictly positive interest rates is relaxed. First, since the pricing kernel defined by Equation (4) is exactly the same in any currency, the two state variables, the world consumption growth and dispersion, have an identical effect on interest rates in all countries. A possible remedy here is to assume that time preference and relative risk aversion can differ across investors in different countries. Second, if one allows a small probability for negative interest rates, then the state variables will have a different impact on the pricing kernels and the interest rates in different countries, if at least one of those variables is country specific. A possible solution to this problem is to use Equation (3) instead of Equation (4). In this case, currency returns in country j will depend on the world consumption and dispersion factors, which are common to all countries, as well as consumption dispersion in country j. A detailed empirical analysis of these issues is beyond the scope of this article.

5. Conclusion

In this article I study whether incomplete consumption risk sharing can be responsible for the forward premium puzzle, that is, the appearance of a large conditional bias in the prediction of the future spot exchange rate from the current forward rate. The existence of this puzzle has implications for expected excess returns from currency speculation. I argue that if cross-country differences in currency risk premiums are driven by local, country-specific factors that are not diversifiable at the global level, then there must be a relation between currency returns and the idiosyncratic consumption risk. I model the uninsurable consumption risk across countries using the multicountry extension of the Constantinides and Duffie (1996) CCAPM, which accounts for investors' heterogeneity and incomplete markets.

I find that theoretical implications of the new CCAPM are partially supported in the data: the cross-country consumption dispersion is countercyclical. This characteristic of consumption data has a positive impact on the usefulness of the model for currency pricing. In particular, the new CCAPM, unlike the standard model, is able to generate foreign exchange risk premiums at substantially lower values of risk aversion. In addition, the tests of the approximate beta pricing relation derived from the model reveal that consumption dispersion provides some explanatory power for the differences in expected excess currency returns and, consequently, helps reduce the extent of the forward premium puzzle.

In spite of certain encouraging results, there are several outstanding issues as well. For example, I obtain the best empirical results in the estimation of the Euler equations when I multiply the world consumption dispersion series by some positive factor. What remains to be seen is the true value of the cross-country consumption dispersion. Second, the power of our tests is generally low, due to the small sample. Therefore it is appealing to construct and estimate an asset pricing model with wealth-based measures for world growth and dispersion instead of consumption-based ones. Such a formulation will allow us to substantially increase the sample size, both across time and cross-sectionally. Finally, explicit accounting for the within-country consumption variation may be useful for further reduction of the forward premium puzzle. All these issues are left for future research.

References

Andersen, T. G., and B. E. Sorensen, 1996, "GMM Estimation of a Stochastic Volatility Model: A Monte Carlo Study," *Journal of Business & Economic Statistics*, 14, 328–352.

Atkeson, A., and T. Bayoumi, 1993, "Do Private Capital Markets Insure Regional Risk? Evidence from the United States and Europe," *Open Economies Review*, 4, 303–324.

Backus, D. K., S. Foresi, and C. I. Telmer, 2001, "Affine Term Structure Models and Forward Premium Anomaly," *Journal of Finance*, 56, 279–304.

Backus, D. K., A. Gregory, and C. I. Telmer, 1993, "Accounting for Forward Rates in Markets for Foreign Currency," *Journal of Finance*, 48, 1887–1908.

Bansal, R., 1997, "An Exploration of the Forward Premium Puzzle in Currency Markets," *Review of Financial Studies*, 10, 369–404.

Bansal, R., and M. Dahlquist, 2000, "The Forward Premium Puzzle: Different Tales from Developed and Emerging Economies," *Journal of International Economics*, 51, 115–144.

Bekaert, G., 1996, "The Time-Variation of Risk and Return in Foreign Exchange Markets: A General Equilibrium Perspective," *Review of Financial Studies*, 9, 427–470.

Bekaert, G., R. J. Hodrick, and D. A. Marshall, 1997, "The Implications of First-Order Risk Aversion for Asset Market Risk Premiums," *Journal of Monetary Economics*, 40, 3–39.

Bilson, J. F. O., 1981, "The 'Speculative Efficiency' Hypothesis," *Journal of Business*, 54, 435–452.

Brav, A., G. Constantinides, and C. Geczy, 2002, "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," *Journal of Political Economy*, 110, 793–824.

Breeden, D., 1979, "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–296.

Campbell, J. Y., 1993, "Intertemporal Asset Pricing without Consumption Data," *American Economic Review*, 83, 487–512.

Cecchetti S., P.-S. Lam, and N. Mark, 1994, "Testing Volatility Restrictions on Intertemporal Marginal Rates of Substitution Implied by Euler Equations and Asset Returns," *Journal of Finance*, 49, 123–152.

Cogley, T., 2002, "Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey," *Journal of Monetary Economics*, 49, 309–334.

Constantinides, G., and D. Duffie, 1996, "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104, 219–240.

Cumby, R. E., 1988, "Is it Risk? Explaining Deviations from Uncovered Interest Parity," *Journal of Monetary Economics*, 22, 279–299.

Engel, C. M., 1996, "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence," *Journal of Empirical Finance*, 3, 123–192.

Fama, E., 1984, "Forward and Spot Exchange Rates," *Journal of Monetary Economics*, 14, 319–338.

Fama, E., and J. MacBeth, 1973, "Risk, Return and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81, 607–636.

Ferson W. E., 1983, "Expectations of Real Interest Rates and Aggregate Consumption: Empirical Tests," *Journal of Financial and Quantitative Analysis*, 18, 477–497.

Ferson, W. E., and C. R. Harvey, 1993, "The Risk and Predictability of International Equity Returns," *Review of Financial Studies*, 6, 773–816.

Ferson, W. E., and A. F. Siegel, 2003, "Stochastic Discount Factor Bounds with Conditioning Information," *Review of Financial Studies*, 16, 567–595.

Gregory, A., and T. McCurdy, 1984, "Testing the Unbiasedness Hypothesis in the Forward Foreign Exchange Market," *Journal of International Money and Finance*, 3, 357–368.

Hall, R. E., 1978, "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86, 971–988.

Hansen, L., 1982, "Large Sample Properties of Generalized Method of Moment Estimators," *Econometrica*, 50, 1029–1054.

Hansen, L., and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262.

Hansen, L., and R. Jagannathan, 1997, "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52, 557–590.

Hansen, L., and K. J. Singleton, 1983, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 91, 249–265.

Harvey, C. R., 1990, "The Term Structure and the World Economic Growth," *Journal of Fixed Income*, 1, 7–19.

Heaton, J., and D. Lucas, 1996, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104, 443–487.

Hodrick, R. J., and S. Srivastava, 1984, "An Investigation of Risk and Return in Forward Foreign Exchange," *Journal of International Money and Finance*, 3, 1–29.

Jacobs, K., 1999, "Incomplete Markets and Security Prices: Do Asset Pricing Puzzles Result from Aggregation Problems," *Journal of Finance*, 54, 123–163.

Lewis, K. K., 1996, "What Can Explain the Apparent Lack of International Consumption Risk Sharing?" *Journal of Political Economy*, 104, 267–297.

Lucas, R. E., Jr., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

Mankiw, N. C., 1986, "The Equity Premium and the Concentration of Aggregate Shocks," *Journal of Financial Economics*, 17, 221–219.

Mark, N., 1985, "On Time-Varying Risk Premiums in the Foreign Exchange Market," *Journal of Monetary Economics*, 15, 145–162.

Mehra, R., and E. Prescott, 1985, "The Equity Premium Puzzle," *Journal of Monetary Economics*, 15, 145–161.

Merton, R. C., 1973, "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, 867–887.

Newey, W., and K. D. West, 1987, "A Simple, Positive Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Ramchand, L., 1999, "Asset Pricing in International Markets in the Context of Agent Heterogeneity and Market Incompleteness," *Journal of International Money and Finance*, 18, 871–890.

Ross, S. A., 1976, "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, 13, 341–360.

Storesletten, K. C., C. I. Telmer, and A. Yaron, 2001, "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," working paper, Carnegie-Mellon University.

Telmer, C. I., 1993, "Asset Pricing Puzzles and Incomplete Markets," *Journal of Finance*, 48, 1803–1832.