The cross section of foreign currency risk premia and consumption growth risk: Reply

Citation: Lustig, Hanno, and Adrien Verdelhan. "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: Reply." American Economic Review 101.7 (2011): 3477–3500.

As Published: http://dx.doi.org/10.1257/aer.101.7.3477

Publisher: American Economic Association

Persistent URL: http://hdl.handle.net/1721.1/76222

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

American Economic Review 101 (December 2011): 3477–3500
http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.7.3477

The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: Reply†

By Hanno Lustig and Adrien Verdelhan*

Generations of economists have been trained to think that investors ex ante should expect to earn a zero return on investments in high interest rate currencies that are funded in low interest rate currencies. In this view, large carry trade returns are an anomaly. Our work shows that this is a reasonable view of the world only if you are willing to believe that investors in currency markets take on aggregate risk without being compensated for it.[1]

Investments are risky if they offer low returns in bad times, when the typical investor experiences higher marginal utility growth than average. As an example, let us examine what transpired in currency markets during the current crisis.

The current recession, which started in December 2007, provides an interesting out-of-sample test case for our claims. We start with the real side of the economy. In the fourth quarter of 2008, the United States recorded a 4.9 percent (annualized) drop in real personal consumption expenditures on nondurable goods, following a 5.6 percent (annualized) decrease in the third quarter. These growth rates are three standard deviations below the mean US consumption growth rate in postwar data. Table 1 summarizes the evidence for 2008: −0.8 percent in nondurable consumption growth, −4.5 percent in durable expenditures growth, and −38.4 percent in the US stock market return. The drop in durable consumption expenditures translates into a weak increase in the stock of durable goods (our measure of durable consumption growth) of 2.7 percent, much lower than the post–WWII average of 3.6 percent. It seems safe to say that the average investor in the United States experienced a higher than usual growth rate in marginal utility, regardless of which model is employed. A similar conclusion holds for foreign investors: total consumption growth in OECD countries decreased by 2.1 percent in the fall of 2008.

* Lustig: Anderson School of Management, University of California at Los Angeles, Box 951477, Los Angeles, CA 90095 (e-mail: hlustig@anderson.ucla.edu); Verdelhan: Department of Finance, MIT Sloan School of Management, E62-621, 77 Massachusetts Avenue, Cambridge, MA 02139 (e-mail: adrienv@mit.edu). The authors would like to thank Andy Atkeson, Dave Backus, John Cochrane, Xavier Gabaix, François Gourio, Christian Hellwig, Bob King, Nikolai Roussanov, and Stijn VanNieuwerburgh for comments.

† To view additional materials, visit the article page at
http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.7.3477.

[1] Other authors have pursued a risk-based explanation of the forward premium puzzle. Our work is closest to Burton Hollifield and Amir Yaron (2001) and Campbell Harvey, Bruno Solnik and Guofu Zhou (2002). Hollifield and Yaron (2001) find some evidence that real factors, not nominal ones, drive most of the predictable variation in currency risk premia. Using a latent factor technique on a sample of international bonds, Harvey, Solnik, and Zhou (2002) find empirical evidence of a factor premium that is related to foreign exchange risk. More recently, Michael J. Brennan and Yihong Xia (2006) show that their estimates of currency risk premia derived in affine term structure models satisfy the Eugene F. Fama (1984) necessary conditions for explaining the forward premium puzzle. Their paper builds on David K. Backus, Silverio Foresi, and Chris I. Telmer (2001), who delineate the class of affine models that satisfy the Fama's (1984) conditions.

Table 1—Currency Excess Returns and Risk Factors—Subprime Crisis

	2007–2008	2008:III–2008:IV	2009:I–2009:II
Panel A. Carry trade returns (percent)			
Deutsche Bank returns	−28.8	−31.4	16.8
LRV returns—all countries	−10.2	−6.1	3.6
LRV returns—developed countries	−31.1	−20.1	8.6
Panel B. Risk factors (percent)			
Expenditure—nondurables	−0.8	−1.6	4.1
Expenditure—durables	−4.5	−10.2	7.4
Consumption—durables	2.7	0.4	0.9
US stock market	−38.4	−29.9	6.06
Total consumption—OECD	−0.6	2.1	−0.6
Panel C. Other zero-cost investment strategies (percent)			
SMB	4.2	5.6	10.3
HML	1.0	10.3	−10.7
Momentum	13.4	0.4	−45.4

Notes: The annual sample starts on December 31, 2007 and ends on December 31, 2008 (first column). The quarterly data correspond to the sum of the 2008 third- and fourth-quarter (second column) growth rates and to the sum of the first two quarters of 2009 (third column). The Deutsche Bank Carry Harvest Index is available online at http://www.dbfunds.db.com/Dbv/index.aspx. The LRV returns were computed by Lustig, Roussanov, and Verdelhan (forthcoming), updated through June 2009. Section A in the separate Appendix contains a detailed description of our series. *SMB*, *HML,* and the momentum factor were taken from Kenneth French's website.

What happened in currency markets during the same period? High interest rate currencies depreciated and low interest rate currencies appreciated. As a result, returns on currency carry trades were low exactly in bad times. No computation needed here; market data are readily available. For example, the Deutsche Bank G10 Carry Harvest Index consists of long futures contracts on the three G10 currencies associated with the highest interest rates and of short futures contracts on the three G10 currencies associated with the lowest interest rates. We use this index as one measure of carry trade returns because the corresponding exchange traded fund is easily available to any investor. The evolution of this index is clear. The current crisis has erased almost all of the carry trade gains made since the end of 2002. Figure 1 plots the evolution of this carry index (we normalized the index to 100 at the end of 2002). Carry traders first enjoyed a long period of steadily high returns. The index peaked at 155 in June 2007, but by the end of 2008 it was back down to 105. Thus, a 55 percent cumulative gain was followed by a like decrease. During the last two quarters of 2008, we witnessed a decrease of more than 31 percent of the Deutsche Bank carry trade index, a negative return equivalent to three standard deviations. The currency portfolios constructed by Lustig, Nikolai Roussanov, and Verdelhan (forthcoming, hereafter denoted LRV) cover more contracts than the Deutsche Bank Index. Even so, applying the long-short strategy still yielded a decrease of more than 10 percent if one invested in both developed and emerging countries.

In a reversal, the Deutsche Bank carry trade index recovered 16.8 percent during the first two quarters of 2009. Nondurable expenditures increased by 4.11 percent, and expenditures on durables increased by 7.4 percent. The US stock market recovered 4.9 percent. If carry trades were not risky, then the opposite pattern would have been just as likely: investors would have fled from the dollar and yen directly to

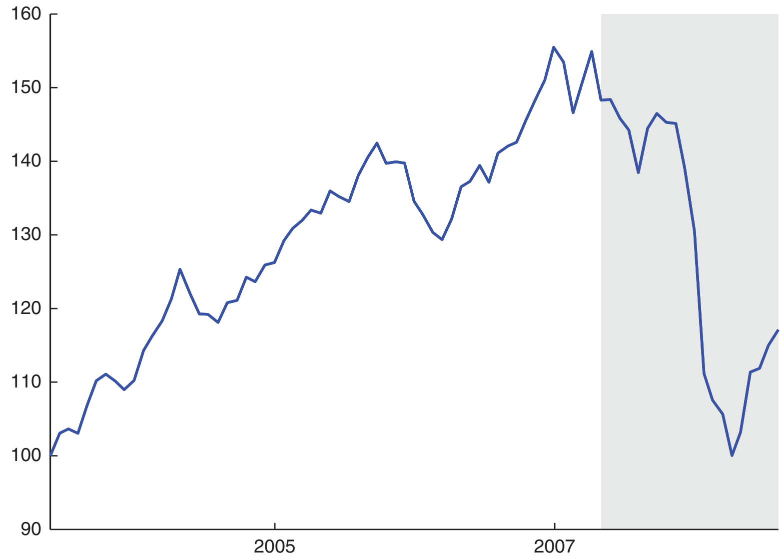

Figure 1. Currency Carry Trade Excess Returns

Notes: This figure presents an index of monthly excess returns on currency carry trades. Carry excess returns correspond to the Deutsche Bank Carry Harvest Index, which is available online at http://www.dbfunds.db.com/Dbv/index.aspx and starts in 1993. In Lustig and Verdelhan (2007), the sample ends in 2002. We focus here on the recent sample from December 2002 to June 2009. The index is equal to 100 at the end of 2002. The gray area, which starts in December 2007, corresponds to the latest recession.

the Australian dollar and the Icelandic krone during the fall of 2008, ignoring the higher-yield currencies when the crisis abated. This is clearly not what happened.

One might argue this is not surprising; maybe all risky zero-cost investment strategies had similar returns during this episode. Far from it. Panel C shows the returns for small-minus-big (long in small, short in large stocks), high-minus-low (long in high book-to-market stocks, short in low book-to-market stocks) and momentum (long in winners over the past 12 months, short in losers over the past 12 months) equity investment strategies. All of these are zero-cost strategies that historically have produced large and positive average excess returns. All three did well during the last two quarters of 2008, while HML and momentum actually did very poorly during the first two quarters of 2009, when US consumption rebounded. The exact opposite from what we see in currency markets. Clearly, US investors with positions in the carry trade have incurred larger portfolio losses and consumption drops than those experienced by the average US investor. Remarkably, this was not the case for other popular zero-cost investment strategies that have proven profitable: size, value, and momentum stock investment strategies all posted positive returns during 2008.

Our Previous Paper.—In a previous paper entitled "The Cross-Section of Currency Risk Premia and Consumption Growth" (Lustig and Verdelhan 2007) we showed that, on average, high interest rate currencies are more exposed to aggregate consumption growth risk than low interest rate currencies in a sample with 81 currencies spanning 50 years of data. High interest rate currencies do not depreciate as much as the interest rate difference, and, as a result, US investors can generally earn positive excess returns by investing in these currencies. However, these high interest rate currencies tend to depreciate relative to low interest rate currencies in bad times for US investors.

Furthermore, the average risk factor loadings of currency portfolios that we reported in Lustig and Verdelhan (2007) tend to understate the true risk inherent in the carry trade because the exposure of carry trade returns to aggregate risk factors increases dramatically during crisis episodes and recessions, exactly when the price of risk should increase in currency markets. To anyone who kept track of recent developments in currency markets, this may now seem obvious, but some economists still insist that "there is no relation between risk factors and currency returns" (see Burnside et al. 2008).

In Lustig and Verdelhan (2007), we departed from the literature by examining the cross-sectional relation between average returns on foreign currency investments and interest rates rather than examining the time-series relation.[2] To do so, we sorted currencies into eight portfolios based on their current interest rate. This approach (developed in Lustig and Verdelhan 2005) is helpful because it averages out changes in exchange rates that are purely idiosyncratic and, hence, are not priced in currency markets. We found that investors on average earn large excess returns simply by taking long positions in baskets of currencies with currently high interest rates and taking short positions in baskets of currencies with currently low interest rates, regardless of the history of interest rate differences for individual currency pairs. The average excess returns increase from the first portfolio, with currently low interest rate currencies, to the last portfolio, with currently high interest rate currencies. Moreover, we established that currencies sorted by interest rates share a lot of common variation, a necessary condition for a risk-based explanation. Finally, we tied this common variation to aggregate risk exposure—more specifically, to durable consumption growth, nondurable consumption growth, and the market return. Our paper was the first to make this point. The literature that precedes it focusses almost exclusively on currency-specific variation in bilateral exchange rates.

On average, the high interest rate currency portfolio produces a return that is 5 percentage points larger per annum than the return on the low interest rate currency portfolio. We find that US aggregate consumption growth risk explains a large share of the variation in average returns for these currency portfolios, because the consumption betas for low interest rate currencies are smaller than the consumption betas for high interest rate currencies. In other words, high interest rate currencies do not depreciate as much as the interest gap on average, but these currencies tend to depreciate in bad times for a US investor, who in turn receives a positive excess return in compensation for taking on this risk.

Our model is a standard representative agent model that allows for nonseparable utility from nondurable and durable consumption and also for nonseparable utility over time. In Lustig and Verdelhan (2007), our analysis proceeds in two steps. First, as is standard in modern macroeconomics, we calibrate the actual model. We adopted the structural parameters from Motohiro Yogo (2006), who estimates these

[2] See Lars Peter Hansen and Robert J. Hodrick (1980) and Fama (1984) for earlier examples of time-series tests. Geert Bekaert and Hodrick (1993) investigate biases as an explanation of the forward premium puzzle. Hodrick (1987) and Karen K. Lewis (1995) provide extensive surveys and updated regression results.

parameters based on stock returns and macroeconomic data. We compute the pricing errors implied by the representative agent's Euler equation evaluated over the sample of the eight currency portfolios. These results are shown in Table 4 (see Section IE). When confronted with the postwar sample of foreign currency returns and US aggregate consumption growth, the representative agent demands a much higher risk premium on the high interest rate currency portfolio than on the low interest rate portfolio. The benchmark model explains 68 percent of the variation in returns. This finding alone disproves the common claim that the forward premium puzzle *cannot* have a risk-based explanation (see, for example, Kenneth A. Froot and Richard H. Thaler 1990).

Second, as is standard in empirical finance, we linearize the model (in Section II of the paper). We then estimate the factor betas for this linearized model by regressing the currency portfolio returns on the three factors (nondurable and durable consumption growth and the market return). Finally, we regress average returns on these betas in order to estimate the risk prices. This exercise confirms our earlier results. The risk prices of nondurable and durable consumption are large, and they are in line with what we and others have found using different test assets (like stocks and bonds). Our paper concludes by explaining why low interest rate currencies tend to appreciate when US consumption growth is lower than average.

Burnside's Comments.—In his comment on our article, Burnside (2007) replicates our point estimates for the risk prices in the linear model using currency portfolios as test assets. He agrees that the consumption betas are aligned with the returns on these currency portfolios. In other words, there is no question that consumption risk is priced if you accept the consumption betas in our sample. However, Burnside questions how accurately these betas are measured. As a result, the debate has shifted away from the claim that risk premia cannot explain the forward premium puzzle—after all, we have shown that the sample moments of consumption growth and currency returns do support a risk-based explanation— to a debate about how accurately these sample moments are measured.

More specifically, Burnside questions the conclusion of our article by claiming (i) that there is no statistical evidence that aggregate consumption growth risk is priced in currency markets and (ii) that our definition of the measure of fit overstates our results.

Our Reply.—In this article, we first address these two claims. We show that they have no merit with respect to our initial sample that ends in 2002. Furthermore, extending the sample to 2009 actually reinforces our points.

First Claim: Burnside claims there is no statistical evidence that aggregate consumption growth risk is priced in currency markets or that currency excess returns do not covary with US consumption growth. This is his most important claim, and it is false.

Let us define HML_{FX} as the difference in returns between the high interest rate portfolio and the low interest rate portfolio. We focus on the seventh portfolio minus the first portfolio because this produces the largest spread (5.3 percent per annum). By construction, the consumption β of HML_{FX} is the difference between the consumption β of the seventh and the first portfolio ($\beta^{HML} = \beta^7 - \beta^1$). Hence, we can simply test Burnside's claim by regressing HML_{FX} on consumption growth.

The consumption growth beta of HML_{FX} is 1 for nondurable and durable consumption growth in a long sample starting in 1953. As a result, the consumption Capital Asset Pricing Model (CAPM) can account for the average return on this investment strategy of 5.3 percent per annum given a market price of consumption risk of around 5 percent per annum. This spread in betas is economically significant. As a benchmark, the consumption beta of the return on the US stock market (the return on the value-weighted CRSP index) is 0.97 for the same sample. In order to explain the average annual stock market excess return of almost 7 percent in the standard consumption CAPM, the price of consumption risk has to be 7.1 percent per annum. This implies a substantial carry trade premium of $7.1 = 1.0 \times 7.1$ percent on the *HML* strategy, compared to 5.3 percent in the data. We obtain similar results in the post–Bretton Woods sample. As a result, if we simply use risk prices from stocks then the model already predicts a sizable carry trade risk premium. Moreover, in Section IVC of our 2007 article, we show that the risk prices we found for currency excess returns are similar to those obtained when estimating the same model on other test assets such as equity and bonds, even though these currency returns are not spanned by the usual factors of value and size. Burnside does not discuss this evidence.

In addition, the spread in consumption betas is statistically significant. In the simple univariate regression case, the *p*-values for a *t*-test are smaller than 2.5 percent in all of the four cases that we consider: nondurables in the 1953–2002 and the 1971–2002 samples, durables in the 1953–2002 and the 1971–2002 samples. The multivariate regressions lead to the same conclusion. The Wald test statistic's *p*-values are both below 1 percent.[3]

Finally, it is not the case that *all* of the consumption betas should be statistically different from zero. The interesting economic question is whether betas are different from each other, not different from zero. Since, for example, the average excess returns on the fifth and sixth currency portfolios are very close to zero, we should expect to see betas close to zero for these portfolios. This is why we focus on the "corner portfolios."

Burnside then argues that the price of consumption risk estimated for currency portfolios is not significantly different from zero once you correct for the fact that the betas are estimated. In our previous article, we report the standard errors obtained by bootstrapping samples from the observed consumption and return data (see Section IVC). These standard errors take into account the two steps and the small sample size. Using these bootstrapped standard errors, the price of durable consumption growth risk is significant at the 5 percent level. The separate Appendix presents additional evidence from generalized least squares (GLS) and generalized method of moments (GMM) estimates that were omitted from the published version. All the evidence indicates that the price of consumption risk is statistically significant.

[3] Why does Burnside reach a different conclusion? In the multivariate case, the only case he considers, Burnside mistakenly focuses on the *t*-statistics of the individual betas; the strong correlation of the consumption factors renders the individual coefficient estimates imprecise. This inference problem is commonly referred to as multicollinearity in textbooks. Obviously, two low *t*-stats on the consumption growth betas in the multivariate regression do not imply that consumption growth does not covary with currency returns.

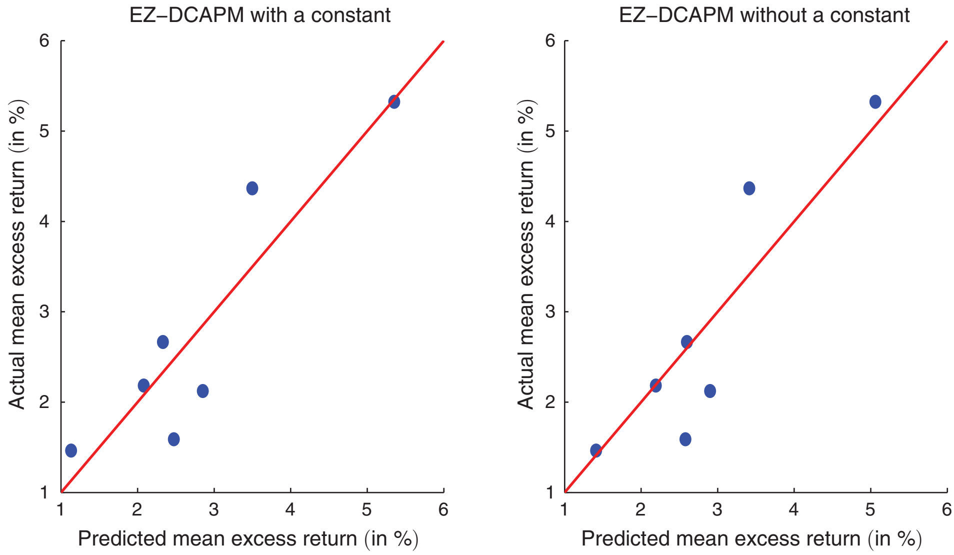

FIGURE 2. SHORT IN LOW AND LONG IN HIGH INTEREST RATE CURRENCIES

Notes: This figure plots actual versus predicted excess returns for seven test assets. Currencies are sorted into eight portfolios according to their interest rates. The seven test assets are obtained by subtracting the returns on the first portfolio from the returns on the other portfolios. These test assets correspond to an investment strategy of going long in the high interest rate currency portfolios and short in the low interest rate currency portfolio. The risk factors are nondurable consumption growth, durable consumption growth, and stock market returns. Risk prices are estimated following Fama and James D. MacBeth (1973). The data are annual from Lustig and Verdelhan (2007), and the sample is 1953–2002.

Second Claim: Burnside points out that the constant in the second stage of our regression is large and negative, and he argues that a risk-based explanation can be discounted because our model overpredicts the returns on the eight currency portfolios. The constant is large (about 300 basis points), but it is not precisely estimated and is not significantly different from zero. Since the rest of Burnside's comment is exclusively about estimation uncertainty, we are puzzled by his emphasis on the point estimate for the constant without emphasizing the large standard error.

This constant measures the price of that part of dollar risk that is not explained by our risk factors. All of our currency portfolios share the same loading on dollar fluctuations, so the cross-section of currency returns cannot be informative about the price of dollar risk. In other terms, the first principal component of currency returns is a dollar factor; all currency portfolios have essentially the same loadings on this factor. We never claim that consumption growth risk explains the returns on investing in a basket of all foreign currencies (i.e., the dollar risk premium), and no such claim is supported by the data.

If instead we use test assets that go long in high interest rate portfolios and short in low interest rate portfolios, we eliminate the dollar risk factor. These test assets are now dollar neutral. In this case the estimated constant is much smaller and insignificant, as expected, and the model does even better on these test assets. Figure 2 plots the benchmark model's predicted excess returns against the realized excess returns for these seven test assets. The model's predicted excess returns are a linear combination of the factor betas. The left panel reflects the inclusion of a constant, and the right panel does not, but there is hardly any difference in the fit. The consumption CAPM model explains 80 percent of the variation in currency excess returns regardless of whether we include a constant. Even though we agree that it overpredicts the average (dollar) excess return on foreign currency investments, the model has no trouble explaining the spread between high and low interest currency returns, and this is the essence of the forward premium puzzle. We could have written our entire article about these zero-cost investment strategies that go long in high and short in low interest rate currencies without changing a single line of the conclusion.

Our article is not about dollar risk. We agree with Burnside that consumption risk does not explain the average returns earned by US investors on a basket of all foreign currencies, and we have never claimed that it did. Our focus is on the returns obtained by going long in high interest rate currencies and short in low interest rate currencies, for this is how the carry trade is defined.

Additional Evidence.—To address Burnside's concerns about estimation uncertainty, we also bring new evidence to bear on the relation between aggregate risk and currency returns.

First, the statistical link between asset returns and macroeconomic factors is always weaker than the link between asset returns and return-based factors. That is why John H. Cochrane (2001) warns against pointless horse races between models with macroeconomic factors and those with return-based factors. Statistical uncertainty is not a license to ignore the link with the macroeconomy. However, an alternative is to use return-based risk factors. This reduces statistical uncertainty at the cost of economic content. Return-based risk factors are more precisely measured, and their loadings and prices are precisely estimated. Lustig, Roussanov, and Verdelhan (forthcoming) use the first two principal components of the currency excess returns as risk factors to explain the cross-sectional variation in monthly currency returns. The second principal component captures the common time-series variation in exchange rates of currencies sorted by their interest rates. As a risk factor, it can account for 65 percent of the cross-sectional variation in average excess returns across our eight portfolios. This factor's market price of risk is precisely estimated and has t-stats well in excess of the 1 percent significance level. Clearly, average currency excess returns can be attributed to covariances between returns and common risk factors. Going back to the sample of countries we used in Lustig and Verdelhan (2007), we obtain similar results on long series of quarterly returns.

Second, by looking at monthly and daily return data, we show that the risk factor loadings of carry returns vary over time. We emphasize that these loadings tend to increase during recessions and other crisis episodes, when the price of risk tends to be higher. As a result, the average factor loadings reported in our previous article tend to understate, not overstate, the true riskiness of these currency portfolios. For example, consider what happened during the recent crisis. At daily frequency, the correlation between one-month returns on the US stock market and carry trade returns increased to 0.70 during the subprime crisis. We report similar findings for the LTCM crisis, the Mexican "Tequila" crisis, and the Brazilian/Argentine crisis.

Third, we extend our findings to foreign investors. Currency carry trades, defined as long-short investment strategies, are dollar-neutral and broadly similar across foreign investors. We show that a consumption-based explanation holds from the perspective of many foreign investors because currency carry trades comove significantly with the *world* component of US consumption growth.

Taking Stock.—Burnside replicates our point estimates of both the quantity and price of consumption growth risk, and he concludes that "it is impossible to reject [our] model using formal statistical tests." He ignores the current crisis, which provides a striking counterexample to his own claims. In the face of all this evidence, what is the point of a lengthy digression on standard errors? In a series of papers, Burnside has consistently argued against risk-based explanations of the carry trade, and he has explored such other avenues as adverse selection and price pressure in attempting to explain the forward premium puzzle.[4] Most recently, however, Burnside et al. (2008) find some support for the disaster risk model of Robert J. Barro (2006). They conclude that "the same value of the stochastic discount factor that rationalizes the average payoffs to the carry trade also rationalizes the equity premium." Thus it seems as though they finally ended up where we started off. We agree that neither price pressure nor adverse selection is a plausible explanation of the forward premium puzzle. Remarkably, Burnside et al. (2008) still insist that carry excess returns are "uncorrelated to traditional risk factors."

The evidence presented in our original article, and in this reply, presents a serious challenge to the view that risk is not priced in currency markets. All the data used in Lustig and Verdelhan (2007) and in this reply are available on the *AER* website.[5] As a result, all tables in the 2007 article and in this reply can be easily replicated.

Outline.—The rest of this article is structured as follows. Section I addresses Burnside's first claim in detail. We have already addressed his second claim. To save space, we present additional results and robustness checks in the separate Appendix. Section II studies the low frequency changes in carry trade risk premia and risk exposure. Section III documents the high frequency time variation in risk prices and quantities on currency markets. Section IV shows that currency carry trades offer low returns when world consumption growth is low. Section V concludes. The separate Appendix, available online, details our sources and the construction of each variable and provides additional evidence.

[4] Burnside et al. (2006) argue that the average payoffs to currency-speculation strategies are low and not risky, and that spot and forward prices move against currency traders in response to "price pressure," thus explaining the forward premium puzzle. However, in a reversal, Burnside, Martin Eichenbaum, and Sergio Rebelo (2007) consider emerging market currencies and do obtain large excess returns, but they still argue against a risk-based explanation. The same claim appears in Burnside, Martin Eichenbaum, and Sergio Rebelo (2008) and Burnside et al. (2008). Burnside, Eichenbaum, and Rebelo (2009) argue that the forward premium puzzle is explained by adverse selection, but not by risk. In the first version of their paper, Burnside et al. (2008) conclude that excess returns in currency markets are not risk compensation, arguing that "the peso problem cannot be a major determinant of the payoff to the carry trade." The second version reaches a different conclusion. Finally, Burnside et al. (2010) argue that investors' overconfidence explains the forward premium puzzle.

[5] Datasets are also available at http://hlustig2001.squarespace.com/ and at http://web.mit.edu/adrienv/www/.

Table 2—Estimation of Consumption Betas for HML_{FX}

Panel I. Simple regression

	β_c^{HML}	$p(\%)$	R^2	β_d^{HML}	$p(\%)$	R^2
	Panel A. Nondurables			*Panel B. Durables*		
1953–2002	1.00		4.04	1.06		9.07
	[0.44]	2.23		[0.40]	0.89	
1971–2002	1.54		8.72	1.65		14.02
	[0.52]	0.28		[0.60]	0.63	

Panel II. Multivariate regression

	β_c^{HML}	β_d^{HML}	χ^2	R^2
1953–2002	0.07	1.03		9.07
	[0.68]	[0.62]	9.40	
1971–2002	0.28	1.48		14.90
	[1.20]	[1.24]	14.15	

Notes: In panel I, each entry reports OLS estimates of β_1 in the following time-series regression of the spread on the factor: $HML_{FX,t+1} = \beta_0 + \beta_1^{HML} f_t + \epsilon_{t+1}$, where $HML_{FX,t+1}$ is the return on the seventh portfolio minus the return on the first portfolio. The estimates are based on annual data; standard errors are reported in brackets. Following Donald W. K. Andrews (1991), we use Newey-West heteroskedasticity-consistent standard errors with an optimal number of lags to estimate the spectral density matrix. The p-values (reported in percent) are for a t-test on the slope coefficient. The factor f_t is nondurable consumption growth (Δc) in the left panel and durable consumption growth (Δd) in the right panel. In panel II, we report the multivariate regressions $HML_{FX,t+1} = \beta_0 + \beta_1^{HML} \mathbf{f}_t + \epsilon_{t+1}$, where $\mathbf{f}_t = [\Delta c_t, \Delta d_t]$. The χ^2 are for a Wald test that the slope coefficients are zero. The data are annual, and the samples cover 1953–2002 and 1971–2002.

I. The Exposure to Consumption Risk across Currencies

Burnside argues that the estimated market prices of risk are not significant once one considers the sampling uncertainty introduced by the first-stage estimation of the betas. In addition, he argues that the consumption betas are all indistinguishable from zero. This is incorrect. We start with the consumption beta estimates.

We first refute Burnside's claim with respect to our previous sample, which ends in 2002. We then check the robustness of our results by extending our annual sample through 2008.

A. *Consumption Growth Betas in Previous Sample*

Let us first recall our previous results. Table 6 of our previous article reports the univariate consumption betas and standard errors. The (nondurable and durable) consumption betas for the seventh currency portfolios are significantly different from zero, but this does not hold for the other portfolios. We obviously agree with Burnside's comment that consumption betas are not estimated as precisely as return-based betas.

Are the consumption betas all indistinguishable from zero and not different across portfolios, as Burnside claims? We show in Table 2 above that the consumption growth betas on a simple currency carry trade strategy (borrowing in low interest rate currencies and lending in high interest rate currencies) vary between 1 over the entire sample and 1.5 in the post–Bretton Woods sample. As already reported, this spread in betas over the entire sample easily accounts for the carry trade excess return. The same is true on the post–Bretton Woods sample. The consumption beta of the return on the US stock market is 1.2 over the 1971–2002 sample. To explain the average annual stock market excess return of 5.75 percent over the same post–Bretton Woods sample in the standard consumption CAPM, the price of consumption risk has to be 4.9 percent. This implies a substantial spread of $7.4 = 1.5 \times 4.9$ percent on the *HML* strategy, compared with 6.9 percent in the data for the 1971–2002 sample.

All these betas are statistically significant at the 5 percent confidence level and are economically meaningful as well.

B. *Durable Consumption Growth and Currency Returns in Longer Sample*

We now extend our sample through 2008. In order to conserve space, we consider only one risk factor: durable consumption growth. This risk factor matters for asset prices if preferences are nonseparable. In addition, durable consumption growth has intuitive appeal as an asset pricing factor because it is highly cyclical. Our baseline measure of durable consumption growth, denoted Δd_1, corresponds to the change in the stock of consumer durables, as in Yogo (2006). As a robustness check we also use a second measure, Δd_2, which is the log change in the quantity index for consumer durable goods from the US Bureau of Economic Analysis (BEA) fixed asset tables. Durable consumption growth is strongly procyclical, more so than nondurable consumption growth. In fact, Joao Gomes, Leonid Kogan, and Yogo (2009) find that an investment strategy that is long on the durable-good producers portfolio and short on the service industry portfolio earns a risk premium exceeding 4 percent annually.

In Figure 3, we plot the average currency excess returns against the durable consumption growth betas, for the first measure Δd_1, estimated on the 1953–2008 sample; we also add standard-error bands. The univariate durable consumption growth betas are reported in Table 7 in Section B of the online Appendix, where Newey-West standard errors are reported in brackets. The durable consumption beta of low interest rate currencies (portfolio 1) differs from the consumption beta of high interest rate currencies (portfolio 7). Clearly, portfolios with higher interest rates tend to have higher consumption betas. What does this mean? It means that investing in low interest rate currencies does not carry the same risk as investing in high interest rate currencies. Low interest rate currencies tend to appreciate in bad times while high interest rate currencies tend to depreciate. The betas on the last portfolios (7 and 8) are statistically different from zero. The betas of the intermediate portfolios are lower, but so are the excess returns on these portfolios. Recall that we are only using one risk factor here; hence Burnside's discussion about the rank of the beta matrix is meaningless.

Durable consumption growth betas are higher for higher interest rate currencies. For the most comprehensive sample of currencies, which includes developing countries, we find that the durable consumption growth betas (i.e., the factor loadings on Δd_1) increase from 0.37 for the first portfolio to 1.57 for the seventh portfolio (and 1.12 for the last portfolio). We find similar results using the second measure; the loadings on Δd_2 increase from 0.33 to 1.24. The variation in betas increases in the post–Bretton Woods sample: the loadings on Δd_1 increase from 0.67 on the first portfolio to 2.38 on the seventh portfolio (and 1.22 on the last portfolio); the loadings on Δd_2 range from 0.57 to 1.53. For the sample of developed currencies, the spreads in loadings are even larger. The loading increases

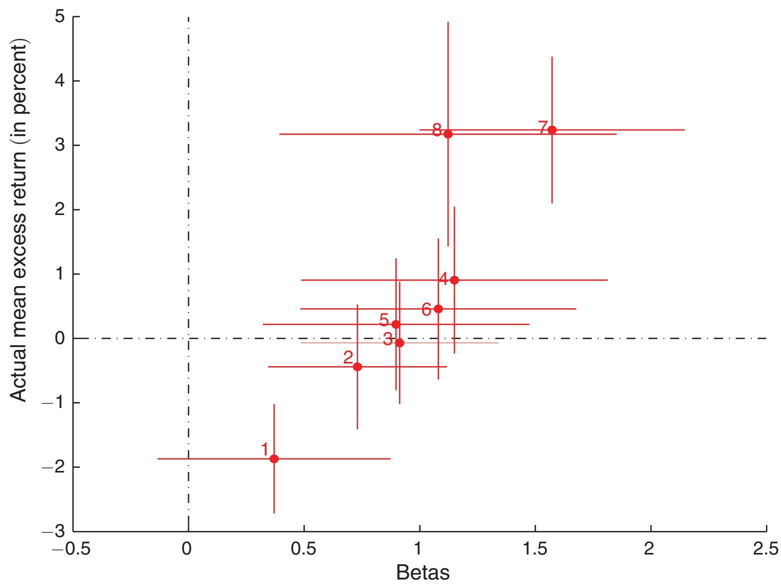

Figure 3. Durable Consumption Growth Betas and Average Currency Excess Returns

Notes: The dots represent point estimates; the lines represent one standard deviation above and below the point estimates. The sample is 1953–2008, and data are annual. Durable consumption growth (Δd_1) is defined in Section A in the online Appendix.

from 0.23 on the first portfolio to 1.85 on the last portfolio for our baseline measure, Δd_1, and from 0.31 to 1.41 for Δd_2. In the post–Bretton Woods sample, the loadings increase from 0.37 to 2.85 for Δd_1 and from 0.47 to 2.06 for Δd_2.

We now examine Burnside's (2007) claim that there is no statistically significant relation between durable consumption growth and the excess returns on the currency portfolios. For Δd_1, this null hypothesis is rejected at the 1 percent significance level for portfolio 7, at the 5 percent level for portfolio 3, and at the 10 percent level for portfolios 2 and 6. For Δd_2, this null is rejected at the 5 percent significance level for five portfolios out of eight and at the 10 percent level for an additional portfolio. In the post–Bretton Woods sample, this claim is rejected at the 1 percent significance level for three portfolios, at the 5 percent level for two more, and at 10 percent for a last one. In sum, six of the eight portfolios have significant consumption betas, and the results are similar for the second measure of durable consumption growth (Δd_2). The results are even stronger if we consider only developed currencies. When using the first measure the null is rejected at the 5 percent level four times and once at the 10 percent level. In the post–Bretton Woods sample, the null is rejected at the 1 percent level for three portfolios and at the 5 percent level for two more. Durable consumption growth betas are statistically different from zero.

Are the loadings on the other currency portfolios statistically different from those on the first portfolio? To answer this question, we look again at loadings on returns on zero-cost portfolios that short the first basket of currencies and go long in the other currencies. These results are reported in the bottom panel of Table 7 in the online Appendix (Section B). Without exception, the loadings are positive. For the sample of developed currencies, the loadings are significantly different from zero at the 10 percent level in half the cases, when we use Δd_1, the first measure of durable consumption growth. In the post–Bretton Woods sample, all of the differences (excepting the fourth and fifth) are statistically different from zero. Clearly, the null that the betas on the first portfolio of funding currencies are identical to those of the other currency portfolios is overwhelmingly rejected at conventional levels of significance. The loadings for the $7 - 1$ portfolio vary between 0.77 and 2.38 in the sample of all currencies, and between 1.28 and 2.98 in the sample of developed currencies (depending on the measure of durable consumption and the time window). Durable consumption growth betas of the carry trade are both economically large and statisically significant. Extending our sample reinforces our previous results.

We have established that the average high interest rate currency *depreciates* relative to the average low interest rate currency in case of negative durable consumption growth innovations and *appreciates* in the case of positive durable consumption growth innovations. This effect is statistically significant, and also economically meaningful; for the sample of developed currencies, the beta of the zero-cost portfolio is around 2, which indicates a 2 percent depreciation for every 1 percent drop in durable consumption growth below its mean.

We have shown in this section that average excess returns correspond to covariances between returns and risk factors. The quantities and prices of risk are precisely estimated when one uses return-based risk factors. Consequently, it is not possible to ignore a risk-based approach to exchange rates merely on the basis of the large standard errors obtained with consumption-based risk factors.

We now rapidly turn to Burnside's second claim. Burnside stresses that the constant in the second stage of our regression is large and negative. We have already largely addressed this claim in the introduction of this article. The constant is not significant. It simply represents the dollar factor. For the sake of completeness, we reestimate our model on a set of seven assets, built from the original eight portfolios by going short in first portfolio (low interest currencies) and long in the other ones. We use Fama and MacBeth (1973) and GMM procedures. To save space, we report these additional results in Section C of the online Appendix. Figure 2 in this article already shows the main point: estimating the model on dollar-neutral portfolios gives similar results with or without a constant.

In the asset pricing literature, standard errors on the estimates of macroeconomic factor loadings are typically large. We report in our previous article and in this reply estimates that are nonetheless significant. To reduce the estimation uncertainty, an alternative strategy is to construct return-based factors. These are much more precisely measured and thus deliver better estimates of the loadings and the risk prices. In Lustig, Roussanov, and Verdelhan (forthcoming), we pursue this approach using monthly currency return data constructed from one-month forward contracts (not T-bills) over the period from 1983:1 to 2008:12. Our results extend to a longer sample of quarterly returns starting either in 1953:I or in 1971:I. The second principal component is a slope factor that explains a large share of the cross-sectional variation in average returns on the currency portfolios. There is much less estimation uncertainty when we use return-based factors. The market price of risk has a t-statistic of 3.4. Again to save space, we report these additional results in Section D in the online Appendix.

Table 3—Estimation of Factor Loadings on Monthly Data

Excess returns			Industry return betas		
Portfolios	All currencies 6 minus 1	Developed currencies 5 minus 1		All currencies 6 minus 1	Developed currencies 5 minus 1
Panel A. 1983:11–2009:6					
Mean	4.97%	2.94%	Durables	0.08***	0.10***
				[0.02]	[0.03]
SD	9.03%	9.74%	Nondurables	0.11**	0.14**
				[0.04]	[0.05]
Sharpe	0.55	0.30	Market	0.15***	0.17***
				[0.03]	[0.05]
Panel B. 1983:11–1995:12					
Mean	2.90%	2.16%	Durables	0.02	0.03
				[0.04]	[0.04]
SD	9.14%	8.72%	Nondurables	0.07	0.09
				[0.04]	[0.05]
Sharpe	0.42	0.25	Market	0.07	0.10*
				[0.05]	[0.05]
Panel C. 1996:1–2009:6					
Mean	6.84%	3.65%	Durables	0.12***	0.14**
				[0.03]	[0.05]
SD	8.92%	10.61%	Nondurables	0.16**	0.22**
				[0.06]	[0.09]
Sharpe	0.77	0.34	Market	0.20***	0.22***
				[0.04]	[0.07]
Panel D. 2000:1–2009:6					
Mean	6.52%	4.42%	Durables	0.11***	0.16**
				[0.03]	[0.05]
SD	7.84%	11.26%	Nondurables	0.08**	0.32**
				[0.09]	[0.12]
Sharpe	0.83	0.39	Market	0.17***	0.29***
				[0.05]	[0.08]

Notes: The first column reports the mean, standard deviation, and Sharpe ratios of carry trade excess returns. The mean and standard deviation are annualized. The next columns report nondurable and durable goods industry betas, along with market betas (ten industry portfolios downloadable from French's website). Monthly currency portfolio returns come from Lustig, Roussanov, and Verdelhan (forthcoming), updated through June 2009. Standard errors are reported in brackets. Following Andrews (1991), we use Newey-West heteroskedasticity consistent standard errors with an optimal number of lags to estimate the spectral density matrix.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

II. Low Frequency Variation in the Carry Trade Risk Premium and Risk Exposure

We have refuted Burnside's two claims; they have no bearing on our initial sample nor on updated series. Risk-based explanations account for the cross section of average currency excess returns estimated over long time windows. In this section we push further the risk-based approach and provide additional evidence on the low frequency variation in carry trade risk premiums and risk exposure.

It turns out that during the last ten years the carry trade risk premium has doubled. The left panel of Table 3 provides an overview of the returns on the monthly currency portfolios of Lustig, Roussanov, and Verdelhan (forthcoming), where the monthly returns are annualized. The carry risk premium, the average return on a long position in the last portfolio and a short position in the first portfolio, has increased from 2.9 percent (2.16 percent for developed currencies) in the first half of the sample (1983–1995) to 6.84 percent (3.65 percent) in the second half of the sample (1996–2009).

The right panel of Table 3 reports betas using industry returns for nondurables and durables from Kenneth French's website; he creates ten industry portfolios using all stocks traded on NYSE-AMEX-NASDAQ. It appears clearly that the large increase in the carry trade risk premium over the past 14 years has been matched by a similar increase in the consumption and market betas of carry trade returns. The durable beta increases from 0.02 to 0.12, the nondurable beta from 0.07 to 0.16. In the sample of developed currencies, the increases in consumption betas measured using industry returns are even larger. In the second half of the sample (1996:1–2009:6), where carry trade excess returns are the largest, industry betas are all significant.

Monthly consumption series are less precisely estimated than industry returns, but they paint a similar picture. We study the contemporaneous betas for durable expenditure growth, nondurable expenditure growth, and the market return, as well as the forward-looking betas, which are obtained by regressing returns at $t + 1$ on consumption growth and returns between t and $t + 3$.[6]

In the 1983:11–1995:12 sample, which is characterized by a small carry risk premium, a long position in the highest interest rate currencies and a short position in the first position did not expose investors to durable expenditures risk. The estimated loadings are close to zero and sometimes negative. Similarly, the market betas are small or negative. However, in the second half of the sample (1996:1–2009:6), the durable expenditure beta of the long short position is 0.16, increasing to 0.19 in the subsample that starts in 2000. The market beta is 0.20. The results for developed currencies are similar but not significant. In the forward-looking consumption betas, the differences are even more striking. The durable expenditure beta increases from -0.46 in the first half of the sample to 0.45 in the second half, and the market beta increases from -0.01 to 0.25. These differences are statistically significant. Overall, long periods of low carry trade returns are associated with low amounts of consumption growth risk, while long periods of high carry trade returns are associated with large consumption risk exposures.

III. High Frequency Variation in the Carry Trade Risk Premium and Risk Exposure

We turn now to high frequency variation in currency risk. This variation implies that unconditional currency betas understate the risk of the currency carry trades.

[6] Jonathan A. Parker and Christian Julliard (2005) show that future consumption growth improves the explanatory power of the standard consumption CAPM, presumably because of lags in household consumption adjustment, which are especially relevant at monthly frequencies. These are the time-series regressions we run:

$$R^{j,e}_{t+1} = \theta^{j,k}_0 + \theta^{j,k}_1 f^k_{t\to t+3} + \eta^{j,k}_{t+1},$$

where $R^{j,e}_{t+1}$ denotes the excess return on portfolio j and $f^k_{t\to t+3}$ denotes the risk factor.

Unconditional Betas Understate Risk.—In the first sections of this article, we have used only unconditional betas, obtained by regressing returns on the factors over the entire sample, to measure the quantity of risk inherent in each portfolio of currencies. However, as we have just seen, the factor betas of currencies vary over time. As demonstrated by Ravi Jagannathan and Zhenyu Wang (1996), this time variation matters for asset pricing if the betas co-vary with risk prices. Consider two assets with zero average market betas: one asset has a beta of 1.5 during recessions and crises, whereas the other asset has a beta of -1.5 during those times. Considering only average betas, one might conclude that investors should be indifferent between these two assets, but this conjecture is false unless the price of risk is constant (see Jagannathan and Wang 1996 for a detailed analysis). Yet it seems reasonable to assume that the price of risk *increases* during episodes like the financial crisis.[7] We show that a similar result holds for currency markets: the factor loadings of carry trades tend to increase dramatically during these episodes. This implies that the average betas tend to understate the true riskiness of these investments.

Stock Market Risk.—In Lustig and Verdelhan (2007) we start from Epstein-Zin preferences and use the implied three components of the stochastic discount factor as risk factors – namely, consumption growth in nondurables and durables and stock market returns. This last factor proxies for the return on wealth. In unconditional asset pricing tests, this return factor does not play much of a role in explaining the cross-section of currency excess returns. Using the annual returns of Lustig and Verdelhan (2007), we show in Section E of the online Appendix that unconditional equity market betas are too low, and thus lead to implausibly high market prices of risk. However, the stock market risk in carry trades increases during crisis episodes. To make this point we turn again to higher-frequency data, the daily and monthly returns of Lustig, Roussanov, and Verdelhan (forthcoming). We start with the current recession and then show that similar results obtain for previous crises.[8]

The recent subprime mortgage crisis offers a good example of the changing nature of the connection between currency and equity markets. Figure 4 plots the monthly returns on a carry trade at daily frequencies against the US stock market return. Clearly, a US investor who was long in these high interest rate currencies and short in low interest rate currencies was heavily exposed to US aggregate stock market risk during the subprime mortgage crisis, and therefore should have been compensated by a risk premium ex ante.

This increase in correlations is not specific to the recent mortgage crisis. We compute the correlation between one-month currency returns and the return on the value-weighted US stock market return using 12-month rolling windows on daily data over the entire 1983–2009 sample. Figure 5 plots the difference between the

[7] Time variation in the quantity and price of risk is indeed well established on equity and bond markets. It underlies the leading dynamic asset pricing models: the habit preferences of John Y. Campbell and John H. Cochrane (1999), the long-run risk model of Ravi Bansal and Amir Yaron (2004), and the time-varying disaster risk of Xavier Gabaix (2009).

[8] We could make a similar point about the subprime crisis and the CAPM using the Deutsche Bank currency index as in the introduction. However, because this index does not begin until 1993, we would not be able to present results for previous crises.

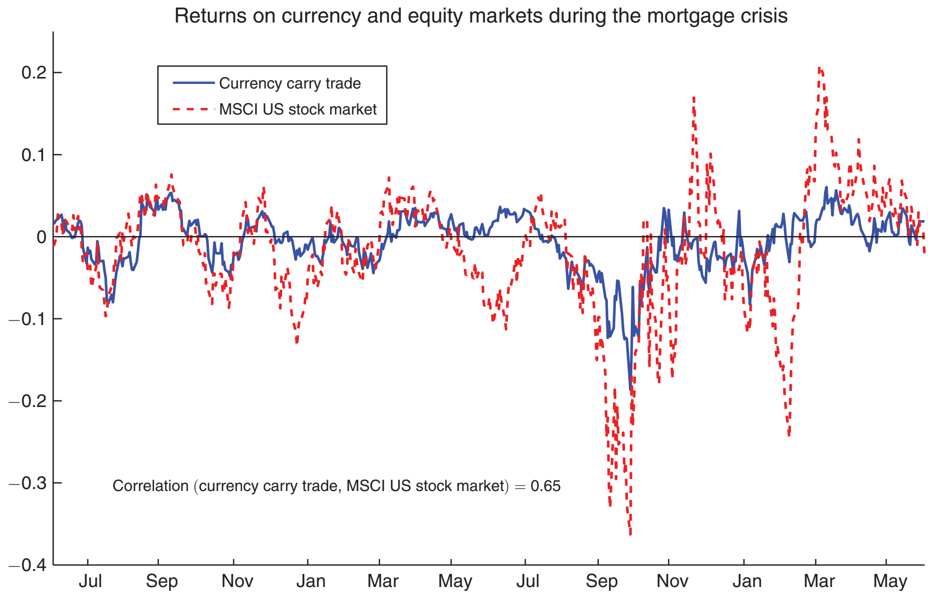

Figure 4. Currency Carry Trade and US Stock Market Returns during the Mortgage Crisis
(*July 2007 through June 2009*)

Notes: This figure plots one-month carry trade returns at daily frequency against one-month returns on the US MSCI stock market index at daily frequency. Currency carry trade returns come from Lustig, Roussanov, and Verdelhan (forthcoming). They correspond to the returns on their last portfolio (i.e., high interest rate currencies) minus the return on their first portfolio (i.e., low interest rate currencies). The sample period is July 2, 2007 through June 30, 2009.

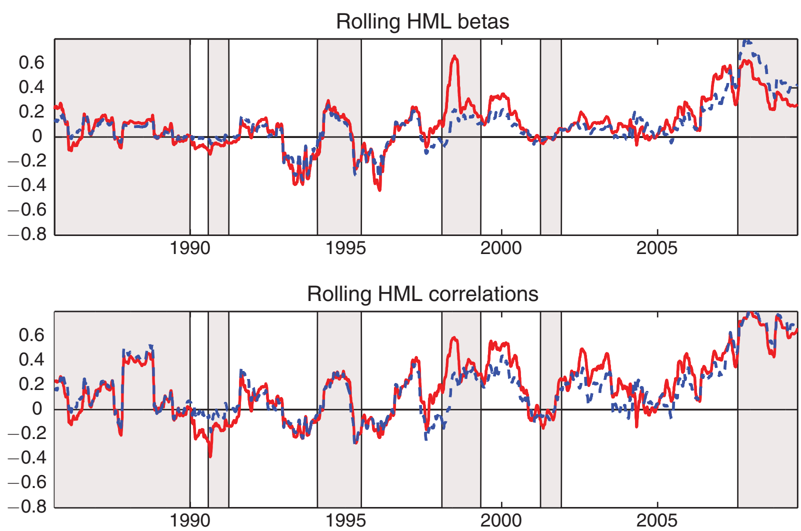

Figure 5. HML Market Correlations and Betas

Notes: This figure first presents $Corr_{\tau}[R_t^m, HML_{FX,t}]$, where $Corr_{\tau}$ is the sample correlation over the previous 12 months (253 days; i.e, over the sample $[\tau - 253, \tau]$), R^m is the return on the US stock market, and HML_{FX} denotes currency carry trade returns. We use monthly returns at daily frequency. The figure also presents β_{HML}, the stock market beta of HML_{FX}. The stock market return is the return on US MSCI index. The solid line uses carry returns on a large sample of developed and emerging countries; the dotted line corresponds to a sample of developed countries. The sample period is October 31, 1983 through June 30, 2009.

correlation of the sixth and the first portfolio with the US stock market excess return. We also plot the stock market beta of HML_{FX}, defined again as the difference in returns between high and low interest rate portfolios. The shaded areas indicate NBER recessions and financial crises. These market correlations exhibit enormous variation. In times of crisis and during US recessions, the difference in market correlation between high and low currencies increases significantly. During the Mexican, Asian, Russian, and Argentinean crises, the correlation difference jumps up by 50–90 basis points.

Three Case Studies.—We explore time variation in the stock market betas over the 1983–2008 sample in more detail for three crisis episodes. To estimate the market β^m, we use daily observations on monthly currency and stock market returns. During these three crisis episodes, stock market betas of carry trades increase dramatically.[9] Table 4 reports the market betas of currency portfolios for six-month windows before the end of May 1998, August 2007, and June 2009.

We observe that β^m_{HML} increases to 1.08 in the run-up to the Russian default in 1998, implying that high interest rate currencies depreciate on average by 1.08 percent relative to low interest rate currencies when the stock market goes down by 1 percent. In times of crisis, low interest rate currencies provide a hedge against market risk, while high interest rate currencies expose US investors to more market risk. And as expected the estimated market betas increase monotonically as we move from low to high interest rate currency portfolios. Similarly, β^m_{HML} increases to 0.64 at the start of the mortgage crisis in August of 2007 and drops to 0.13 toward the end of the crisis in June 2009.

There is substantial high frequency variation in the consumption betas as well. These consumption betas tend to increase during recessions and financial crises. In Table 5, we revisit the same three episodes and report the consumption exposure in these long-short currency strategies. We use 18-month windows to estimate the factor betas. The nondurable betas of the $6 - 1$ (HML_{FX}) zero-cost portfolio are 0.99 in May 1998, 1.53 in August 2007, and 1.73 in June 2009 (as estimated using the most comprehensive sample of currencies). For the sample of developed currencies, the respective consumption betas are 1.98, 3.76, and 2.07. The panel on the right reports the corresponding industry return betas. These produce the same pattern.

IV. US and World Consumption Growth

As in our original paper, we have so far focused on US investors. We show now that a consumption-based explanation of carry trade excess returns holds from the perspective of many foreign investors because currency carry trades comove significantly with the average consumption growth in developed countries. High (low) interest rate currencies tend to depreciate (appreciate) when the *world* component of US consumption growth is low.

To illustrate this point, let us disentangle total US consumption growth into two components—a world and a US-specific component—by regressing US consumption

[9] Charlotte Christiansen, Angelo Ranaldo, and Paul Soderlind (2011) find similar results on a shorter sample (1995–2008) using a logistic smooth transition regression model.

TABLE 4—ESTIMATION OF MARKET LOADINGS ON DAILY DATA—THREE CASE STUDIES

Portfolio	α_m^i	β_m^i	R^2
Panel A. May 26, 1998			
1	−0.83	−0.06	0.63
	[0.73]	[0.15]	
2	−0.52	−0.04	0.39
	[1.03]	[0.17]	
3	−1.29	0.19	11.35
	[0.42]	[0.10]	
4	−1.67	0.33	7.13
	[0.59]	[0.14]	
5	−3.13	0.58	15.44
	[1.50]	[0.27]	
6	−3.68	1.03	40.47
	[1.08]	[0.23]	
6 minus 1	−2.85	1.08	53.36
	[0.63]	[0.16]	
Panel B. Aug 31, 2007			
1	0.18	−0.14	13.52
	[0.37]	[0.05]	
2	0.15	0.22	28.11
	[0.36]	[0.06]	
3	0.72	0.18	29.72
	[0.26]	[0.05]	
4	0.29	0.21	39.76
	[0.26]	[0.04]	
5	0.46	0.15	43.24
	[0.15]	[0.04]	
6	0.60	0.51	54.47
	[0.43]	[0.11]	
6 minus 1	0.42	0.64	67.39
	[0.37]	[0.08]	
Panel C. Jun 30, 2009			
1	0.37	0.31	60.69
	[0.45]	[0.03]	
2	0.47	0.22	54.97
	[0.33]	[0.03]	
3	0.14	0.33	49.16
	[0.78]	[0.06]	
4	0.62	0.29	52.48
	[0.54]	[0.05]	
5	0.87	0.43	59.35
	[0.76]	[0.07]	
6	0.74	0.45	71.19
	[0.77]	[0.09]	
6 minus 1	0.37	0.13	20.04
	[0.67]	[0.08]	

Notes: This table reports estimates of the CAPM betas during crises. The sample period is 129 days (six months) before and including the mentioned date in each panel. The table reports the intercept α_m^i, slope coefficient β_m^i, and R^2 in a regression of each portfolio i's currency excess returns on a constant and the US stock market return. The intercept α_m^i and the R^2 are reported in percentage points. The Newey-West standard error correction is computed with 20 lags. We use the daily currency portfolios from Lustig, Roussanov, and Verdelhan (forthcoming), updated through June 2009, and the MSCI return on the US stock market.

Table 5—Estimation of Factor Loadings on Monthly Data—Three Case Studies

	Betas		Industry return betas	
	6 minus 1 All currencies	5 minus 1 Developed currencies	6 minus 1 All currencies	5 minus 1 Developed currencies
Panel A. May 1998				
Durables	0.38	0.21	0.15	−0.10
	[0.51]	[0.27]	[0.24]	[0.11]
Nondurables	0.99	1.98**	0.43**	0.03
	[2.14]	[0.95]	[0.17]	[0.11]
Market	0.34*	−0.07	0.34*	−0.07
	[0.19]	[0.13]	[0.19]	[0.13]
Panel B. August 2007				
Durables	−0.01	−1.15	0.29*	0.30*
	[0.49]	[0.77]	[0.18]	[0.22]
Nondurables	1.53*	3.76***	0.55**	0.47*
	[0.95]	[1.42]	[0.14]	[0.26]
Market	0.46***	0.27	0.46***	0.27*
	[0.16]	[0.18]	[0.16]	[0.18]
Panel C. June 2009				
Durables	0.91*	1.79**	0.11**	0.22***
	[0.56]	[0.81]	[0.04]	[0.08]
Nondurables	1.73**	2.07**	0.29***	0.65***
	[0.68]	[1.05]	[0.12]	[0.19]
Market	0.26***	0.54***	0.26***	0.54***
	[0.10]	[0.12]	[0.10]	[0.12]

Notes: Table entries are regressions on the 18 months preceding the event, including the month itself. Following Andrews (1991), we use Newey-West heteroskedasticity-consistent standard errors with an optimal number of lags to estimate the spectral density matrix. The left panel reports actual betas for consumption and the market returns. The right panel reports nondurable and durable goods industry betas (ten industry portfolios downloadable from French's website).

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

growth on the average consumption growth in OECD countries (a measure of world consumption growth):

$$\Delta c_{t+1}^{US} = \delta_0 + \delta^w \Delta c_{t+1}^w + \Delta c_{t+1}^{US,res}.$$

To save space, we consider only developed countries post–Bretton Woods. The sample focuses on OECD countries at a quarterly frequency over the 1971:I–2009:II period. The average excess return or the average exchange rate in each portfolio is regressed on a constant and the two orthogonal components of US-consumption growth, i.e., $\Delta c_{t+1}^{US,res}$ and $\Delta c_{t+1}^{US,w}$ (where $\Delta c_{t+1}^{US,w} = \delta^w \Delta c_{t+1}^w$ denotes the projection of US consumption growth on world consumption growth). Results are reported in Table 6. The upper panel corresponds to currency excess returns; the bottom panel corresponds to exchange rates.

Let us start with exchange rates. Slope coefficients (denoted betas) on the *global* component of US consumption growth range from −0.77 for the first portfolio to −3.33 for the last one. The *US-specific* betas are 1.31 and 1.22 for the same portfolios.

Table 6—US-Specific and World Consumption Growth Betas

Portfolio	Mean excess return	α_0^j	$\beta^j_{\Delta c^{US,res}}$	$\beta^j_{\Delta c^{US,w}}$	R^2
Panel A. Excess return betas					
1	−0.08	−1.81	−1.21	0.56	2.15
	[1.76]	[4.09]	[0.73]	[1.15]	
2	2.04	−1.98	−1.06	1.29	2.66
	[1.65]	[5.26]	[0.65]	[1.45]	
3	2.69	−2.00	−0.81	1.51	2.34
	[1.62]	[5.19]	[0.65]	[1.45]	
4	3.01	−6.93	−1.22	3.19	7.20
	[1.75]	[6.44]	[0.67]	[1.87]	
5	6.37	−2.50	−0.76	2.85	5.13
	[1.74]	[6.15]	[0.64]	[1.78]	
5 − 1	6.45	−0.69	0.45	2.29	2.86
	[1.58]	[5.47]	[0.67]	[1.61]	
Panel B. Exchange rate betas					
1	−0.08	0.48	1.31	−0.77	2.70
	[1.76]	[3.88]	[0.73]	[1.12]	
2	2.04	3.52	1.23	−1.64	3.89
	[1.65]	[5.23]	[0.67]	[1.43]	
3	2.69	5.16	1.00	−1.90	3.69
	[1.62]	[5.06]	[0.68]	[1.41]	
4	3.01	12.32	1.39	−3.72	9.63
	[1.75]	[6.27]	[0.67]	[1.82]	
5	6.37	19.10	1.22	−3.33	7.71
	[1.74]	[5.59]	[0.63]	[1.66]	
5 − 1	6.45	18.63	−0.09	−2.56	3.06
	[1.58]	[4.97]	[0.67]	[1.48]	

Notes: The table reports OLS estimates of the factor betas. The first column reports the mean excess return for each portfolio. The upper panel focuses on currency excess return betas, while the bottom panel reports exchange rate betas. R^2s are reported in percentage points. The quarterly excess returns and changes in exchange rates are annualized (i.e., multiplied by four). The alphas are reported in percentage points. The standard errors in brackets are Whitney K. Newey and Kenneth D. West (1987) standard errors computed with the optimal number of lags according to Andrews (1991). Data are quarterly, from the IMF/IFS and OECD databases. The sample period is 1971:I–2009:II.

As a result, the US-specific component cannot account for the cross-section of carry trade excess returns; however, world consumption growth does. Looking at a long-short strategy (long the last portfolio and short the first one), the US-specific consumption beta is small at −0.09 (with a standard error of 0.67), while the world consumption beta is −2.56 (with a standard error of 1.48). The world consumption growth beta is almost 1.8 standard errors away from zero, whereas the US-specific consumption beta is clearly insignificant. Currency excess returns deliver similar results. Overall, exchange rates appear to satisfy a simple macrofinance logic: assets that pay badly in bad times are risky, and risk-averse investors expect excess returns to compensate for that risk. For carry trades, bad times correspond to low world consumption growth rates.

V. Conclusion

Our article on "The Cross-Section of Currency Risk Premia and Consumption Growth" demonstrates that consumption growth risk is priced in currency markets.

To make this point, we use currency portfolios sorted by interest rates. These portfolios average out the idiosyncratic risk in exchange rate changes, and this produces a sharper picture of the risk-return trade-off in currency markets. In our sample, low interest rate currency portfolios have low consumption growth betas, and high interest rate currency portfolios have high consumption growth betas. This implies that the forward premium puzzle has a risk-based explanation. Burnside argues that the data are not informative about the relation between consumption growth and foreign currency returns. We disagree and have pointed out the parts of our article that Burnside overlooked. We have also provided additional evidence in favor of a risk-based explanation based on factor betas that are measured precisely. Our portfolios do not allow us to identify the price of dollar risk, which is Burnside's second point. We agree, but this is not what either our article or the forward premium puzzle is about. Our article is concerned with the spread between high and low interest rate currency returns, and we have shown that the model explains about 80 percent of the variation in these returns.

What Have We Learned?—Our approach of building portfolios of currencies has helped to establish that low and high interest rate currencies have undeniably different risk characteristics. The current crisis provides a painful lesson to anyone who doubts this. Had researchers started by looking at these currency portfolios 25 years ago, the forward premium would probably not have been assigned the "puzzle" label. In fact, these researchers would have been puzzled to find that uncovered interest rate parity actually holds.

Our approach to studying currencies has been adopted by several authors recently (including Burnside). It enabled, for example, Roberto A. DeSantis and Fabio Fornari (2008), Jakub W. Jurek (2008), Lukas Menkhoff et al. (forthcoming), Farhi et al. (2009), Andrew Ang and Joseph S. Chen (2010), Christiansen, Ranaldo, and Soderlind (2011), and Serhiy Kozak (2011) to make further progress on the road to a better understanding of exchange rates.

REFERENCES

Andrews, Donald W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica*, 59(3): 817–58.

Ang, Andrew and Joseph S. Chen. 2010. "Yield Curve Predictors of Foreign Exchange Returns." Unpublished.

Backus, David K., Silverio Foresi, and Chris I. Telmer. 2001. "Affine Term Structure Models and the Forward Premium Anomaly." *Journal of Finance*, 56(1): 279–304.

Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, 59(4): 1481–1509.

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–66.

Bekaert, Geert, and Robert J. Hodrick. 1993. "On Biases in the Measurement of Foreign Exchange Risk Premiums." *Journal of International Money and Finance*, 12(2): 115–38.

Brennan, Michael J., and Yihong Xia. 2006. "International Capital Markets and Foreign Exchange Risk." *Review of Financial Studies*, 19(3): 753–95.

Burnside, Craig. 2007. "The Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk: A Comment." National Bureau of Economic Research Working Paper 13129.

Burnside, Craig, Bing Han, David A. Hirshleifer, and Tracy Yue Wang. 2010. "Investor Overconfidence and the Forward Premium Puzzle." Duke University, Department of Economics Working Paper 10–46.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo. 2006. "The Returns to Currency Speculation." National Bureau of Economic Research Working Paper 12489.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo. 2008. "Do Peso Problems Explain the Returns to the Carry Trade?" National Bureau of Economic Research Working Paper 14054.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2007. "The Returns to Currency Speculation in Emerging Markets." *American Economic Review*, 97(2): 333–38.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2008. "Carry Trade: The Gains of Diversification." *Journal of the European Economic Association*, 6(2-3): 581–88.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2009. "Understanding the Forward Premium Puzzle: A Microstructure Approach." *American Economic Journal: Macroeconomics*, 1(2): 127–54.

Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy*, 107(2): 205–51.

Christiansen, Charlotte, Angelo Ranaldo, and Paul Soderlind. 2011. "The Time-Varying Systematic Risk of Carry Trade Strategies." *Journal of Financial and Quantitative Analysis*, 46: 1107–25. Cochrane, John H. 2001. *Asset Pricing*. Princeton, N.J.: Princeton University Press.

DeSantis, Roberto A. and Fabio Fornari. 2008. "Does Business Cycle Risk Account for Systematic Returns from Currency Positioning? The International Perspective." Unpublished.

Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57(4): 937–69.

Fama, Eugene F. 1984. "Forward and Spot Exchange Rates." *Journal of Monetary Economics*, 14(3): 319–38.

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

Farhi, Emmanuel, Samuel P. Fraiberger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan. 2009. "Crash Risk in Currency Markets." Unpublished.

Froot, Kenneth A., and Richard H. Thaler. 1990. "Foreign Exchange." *Journal of Economic Perspectives*, 4(3): 179–92.

Gabaix, Xavier. 2009. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." Unpublished.

Gomes, Joao F., Leonid Kogan, and Motohiro Yogo. 2009. "Durability of Output and Expected Stock Returns." *Journal of Political Economy*, 117(5): 941–86.

Hansen, Lars Peter, and Robert J. Hodrick. 1980. "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis." *Journal of Political Economy*, 88(5): 829–53.

Harvey, Campbell R., Bruno Solnik, and Guofu Zhou. 2002. "What Determines Expected International Asset Returns?" *Annals of Economics and Finance*, 3(2): 249–98.

Hodrick, Robert. 1987. *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*. Chur, Switzerland: Harwood Academic Publishers.

Hollifield, Burton and Amir Yaron. 2001. "The Foreign Exchange Risk Premium: Real and Nominal Factors." Unpublished.

Jagannathan, Ravi, and Zhenyu Wang. 1996. "The Conditional CAPM and the Cross-Section of Expected Returns." *Journal of Finance*, 51(1): 3–53.

Jagannathan, Ravi, and Zhenyu Wang. 1998. "An Asymptotic Theory for Estimating Beta-Pricing Models Using Cross-Sectional Regression." *Journal of Finance*, 53(4): 1285–1309.

Jurek, Jakub W. 2008. "Crash-Neutral Currency Carry Trades." Unpublished.

Kho, Bong-Chan, Dong Lee, and Rene M. Stulz. 2000. "U.S. Banks, Crises, and Bailouts: From Mexico to LTCM." *American Economic Review*, 90(2): 28–31.

Kozak, Serhiy. 2011. "Bond and Foreign Exchange Risk Premia." Unpublished.

Lettau, Martin, and Sydney Ludvigson. 2001. "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying." *Journal of Political Economy*, 109(6): 1238–87.

Lewis, Karen K. 1995. "Puzzles in International Financial Markets." In *Handbook of International Economics Volume 3*, ed. G. M. Grossman and K. Rogoff, 1913–71. Vol.3, *Handbooks in Economics.* Amsterdam: Elsevier, North-Holland.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan. Forthcoming. "Common Risk Factors in Currency Markets." *Review of Financial Studies.*

Lustig, Hanno, and Adrien Verdelhan. 2005. "The Cross-Section of Currency Risk Premia and US Consumption Growth Risk." National Bureau of Economic Research Working Paper 11104.

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

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. Forthcoming. "Carry Trades and Global Foreign Exchange Rate Volatility." *Journal of Finance*.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–08.

Parker, Jonathan A., and Christian Julliard. 2005. "Consumption Risk and the Cross Section of Expected Returns." *Journal of Political Economy*, 113(1): 185–222.

Ross, Stephen A. 1976. "The Arbitrage Theory of Capital Asset Pricing." *Journal of Economic Theory*, 13(3): 341–60.

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

Shanken, Jay, and Guofu Zhou. 2007. "Estimating and Testing Beta Pricing Models: Alternative Methods and Their Performance in Simulations." *Journal of Financial Economics*, 84(1): 40–86.

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