

DISCUSSION PAPER SERIES

No. 3463

FOREIGN CURRENCY FOR LONG-TERM INVESTORS

John Y Campbell, Luis M Viceira
and Josh S White

FINANCIAL ECONOMICS



Centre for **E**conomic **P**olicy **R**esearch

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP3463.asp

FOREIGN CURRENCY FOR LONG-TERM INVESTORS

John Y Campbell, Harvard University
Luis M Viceira, Harvard University and CEPR
Josh S White, Harvard University

Discussion Paper No. 3463
July 2002

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **FINANCIAL ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: John Y Campbell, Luis M Viceira and Josh S White

July 2002

ABSTRACT

Foreign Currency for Long-Term Investors*

Conventional wisdom holds that conservative investors should avoid exposure to foreign currency risk. Even if they hold foreign equities, they should hedge the currency exposure of these positions and should hold only domestic Treasury bills. This Paper argues that the conventional wisdom may be wrong for long-term investors. Domestic bills are risky for long-term investors because real interest rates vary over time, and bills must be rolled over at uncertain future interest rates. This risk can be hedged by holding foreign currency if the domestic currency tends to depreciate when the domestic real interest rate falls, as implied by the theory of uncovered interest parity. Empirically this effect is important and can lead conservative long-term investors to hold more than half their wealth in foreign currency.

JEL Classification: G12

Keywords: foreign exchange rates, home bias, intertemporal hedging demand, portfolio choice and uncovered interest parity

John Y Campbell
Department of Economics
Harvard University
Littauer Centre 213
Cambridge MA 02138
USA
Tel: (1 617) 496 6448
Fax: (1 617) 495 7730
Email: john_campbell@harvard.edu

Luis M Viceira
Graduate School of Business
Administration
Morgan Hall 367
Harvard University
Boston MA 02163
USA
Tel: (1 617) 495 6331
Fax: (1 617) 495 6592
Email: lviceira@hbs.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=108043

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=139086

Josh S White
Graduate School of Business
Administration
Harvard University
Boston MA 02163
USA
Tel: (1 617) 495 6995
Fax: (1 617) 496 7357
Email: jswwhite@fas.harvard.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=157912

*Campbell acknowledges the financial support of the National Science Foundation and Viceira the financial support of the Division of Research of the Harvard Business School. A preliminary version of this Paper was delivered by Campbell as the Frank Hahn lecture to the annual meeting of the Royal Economic Society at the University of Warwick on 27 March 2002. We are grateful for comments by participants at the meeting, and at seminars at HEC and London Business School.

Submitted 26 June 2002

Home bias is one of the most striking phenomena in international finance. Investors have a strong tendency to concentrate their portfolios in domestic assets, assigning relatively little weight to foreign assets. Although home bias may have been diminishing gradually over time, it remains important as documented by Lewis (1999) and others.

It is helpful to distinguish two different aspects of home bias. First, there is home bias in long-term real assets such as equities. Second, there is home bias in short-term debt instruments such as Treasury bills. In the international context, these instruments are sometimes called “currency”, and foreign bills are called “foreign currency”, but this should not be confused with literal holdings of cash for transactions purposes.

Equity home bias and currency home bias are conceptually separate. A portfolio that is concentrated in domestic equities can still include foreign currency, and a diversified portfolio of foreign equities can include only domestic currency. Even if the domestic-currency value of foreign equities is correlated with the value of foreign currency (as is generally the case), the exposure of a diversified international equity portfolio to foreign exchange rates can be hedged by shorting foreign currency. Such foreign exchange hedging is widely practiced by institutional investors.

Financial economists have asked whether modern portfolio theory offers any possible justification for home bias. Home bias in equities appears to be extremely hard to justify. To the extent that different national stock markets are imperfectly correlated, a diversified portfolio of international equities offers reduced risk with no loss of average return, and a domestic equity portfolio is likely to be mean-variance inefficient. Although cross-country equity correlations have risen in recent decades (Goetzmann, Li, and Rouwenhorst 2002), the argument for international equity diversification remains compelling.²

The case for currency diversification is much less clear-cut. One argument for

²Nontraded assets such as private businesses and human capital can affect the demand for equities. Baxter and Jermann (1997) argue that these assets are positively correlated with domestic equities and strengthen the case for international equity diversification, while Bottazzi, Pesenti, and van Wincoop (1996) argue that they may be negatively correlated with domestic equities and may justify some degree of equity home bias. Nontraded assets can also generate home bias through external habit formation, as agents without nontraded assets mimic the portfolios of agents with nontraded assets (Chue 2002, DeMarzo, Kaniel and Kremer 2002, Shore and White 2001, Wheatley 2001).

diversifying across currencies emphasizes the importance of inflation risk. Suppose that purchasing power parity holds, so that real exchange rates are constant and nominal exchange rates move only with inflation shocks. A stable currency will be a safer investment than a currency that is subject to large changes in its real value. This fact explains why dollars and euros are popular in some developing countries with a history of volatile monetary policy. More generally, a diversified portfolio of currencies might be a safer real investment than any one currency alone.

This argument applies well to some emerging markets, but it is much less satisfactory for developed markets. Since the end of the Bretton Woods exchange rate system in 1973, real exchange rates of developed countries have been highly volatile and correlated with nominal exchange rates. In the short term at least, nominal exchange rate movements are almost equal to real exchange rate movements.

Real exchange rate risk has asymmetric effects that can justify a strong home bias in currency holdings. A domestic investor is concerned with the domestic purchasing power of her portfolio, so real exchange rate movements make foreign currency risky from her perspective. A foreign investor, meanwhile, is concerned with the purchasing power of her portfolio in the foreign country where she is located; real exchange rate movements make domestic currency risky from her perspective. Each investor should hold her own country's currency.

The classic model of this is due to Solnik (1974). Solnik studies the portfolio choice problem of an investor who can hold domestic and foreign currency and equities. Solnik assumes that local-currency equity values are uncorrelated with exchange rate movements, and that real exchange rates are variable. He finds that the optimal portfolio is internationally diversified in equities, but home-biased in currency; the investor should hedge the currency exposure of foreign equities.

Solnik's analysis underlies the conventional wisdom that foreign equity investments should be hedged. On this view, foreign currency is a speculative asset that should be held only for tactical reasons to enhance returns, for example by hedge funds that seek to exploit short-term deviations from uncovered interest parity.

What factors could overturn Solnik's justification for currency home bias? One possibility is that local-currency equity returns are correlated with exchange rate movements. This would be the case, for example, if equities are international assets that are implicitly valued in dollars. In this case a depreciation of a foreign currency would cause an offsetting increase in the foreign-currency value of foreign

equities, leaving the dollar value of those equities unchanged. Some emerging equity markets appear to behave in this way, at least in response to currency crises. Froot (1993) argues that even in developed markets, equity returns are correlated with exchange rate movements over long horizons, although not over short horizons. Wilson (2002) shows that correlations between local-currency equity returns and exchange rate movements have tended to increase over the past few decades.

In this paper we explore a second argument for international currency diversification, which has to do with the distinction between short-term and long-term risk. In the short term, domestic currency is almost riskless in real terms; the only risk arises from shocks to the price level, which are modest over short periods. It is conventional in finance to ignore this short-term real currency risk, and to plot domestic bills on the riskless vertical axis of a mean-standard deviation diagram. In the long term, however, domestic currency is not riskless because the real interest rate varies over time. A long-term investor must roll over short-term debt at uncertain future real interest rates.

The long-term risk created by this effect can be quite substantial. Campbell and Viceira (2002) estimate a vector autoregressive (VAR) model for quarterly postwar US asset returns in which the annualized standard deviation of US Treasury bill returns grows from 1.5% over a quarter to 3% at a 20-year horizon. The effect is even more dramatic in long-term annual US data, which includes volatile inflation and real interest rates from the early 20th Century. In data since 1890, Campbell and Viceira estimate that the annualized standard deviation of US bill returns grows from 8% over a year to 13% over 20 years, almost as large as the annualized standard deviation of stock returns. Siegel (1998) also emphasizes the increased riskiness of bills for long-term investors. These authors have studied the US, which has had comparatively stable real interest rates; the effect is likely to be even more important in other countries.

The long-term risk of domestic currency implies that conservative long-term investors should be interested in alternative assets that hedge real interest rate fluctuations. One possibility, emphasized by Campbell and Shiller (1996) and Campbell and Viceira (2001, 2002) is a long-term inflation-indexed bond. Such a bond offers a long-term riskless real return. In the short term, if the real interest rate falls, an inflation-indexed bond will increase in value, offering the investor increased wealth to compensate for the deterioration in investment opportunities. Inflation-indexed bonds have been issued by the governments of a number of developed countries, in-

cluding Canada, the UK, and the US. But they do not exist in all countries and are relatively new financial instruments.

A second possibility is a long-term nominal bond. This hedges investors against real interest rate fluctuations, but exposes them to movements in expected inflation. Empirically, Campbell and Viceira (2001) find that nominal bonds are good substitutes for inflation-indexed bonds in the period of anti-inflationary monetary policy since 1983, but are poor substitutes in the earlier postwar period when US inflation was volatile and persistent.

A third possibility is foreign currency. To see how foreign short-term debt can hedge an investor against domestic real interest rate fluctuations, consider what happens when the domestic real interest rate falls. The theory of uncovered interest parity holds that expected returns are equated across currencies. This theory, as embodied in exchange rate models such as the classic Dornbusch (1976) model, implies that a decline in the domestic real interest rate should depreciate the domestic currency in the short run, with a subsequent gradual domestic currency appreciation that offsets the low real interest rate at home. Equivalently, a decline in the domestic real interest rate should immediately increase the domestic value of foreign currency. Thus a domestic investor who holds foreign currency is compensated for the deterioration in domestic investment opportunities by an immediate increase in wealth.

Of course, there are also movements in the exchange rate that are unrelated to the domestic real interest rate. These might be driven by changes in the long-run equilibrium real exchange rate (Campbell and Clarida 1987), by movements in the foreign real interest rate, or by short-term deviations from uncovered interest parity. These exchange rate movements create short-term risk in foreign currency. The implications for a long-term investor depend on the extent to which foreign exchange rate movements persist in the long term. If purchasing power parity holds, then movements in the real exchange rate are temporary and will be of little concern to a long-term investor.

Recent empirical research in international economics suggests the relevance of these arguments. The hypothesis of uncovered interest parity (UIP) can be rejected (Fama 1984, Hodrick 1987, Froot and Thaler 1990, Engel 1996). But recent data are more favorable to the hypothesis, especially at the longer horizons which are of greater concern here (Baillie and Bollerslev 2000, Chinn and Meredith 2001, Bekaert and Hodrick 2001, Bekaert, Wei, and Xing 2002). There is evidence for slow reversion of

real exchange rates to a stable long-run mean, as implied by the theory of purchasing power parity (Lothian 1990, Froot and Rogoff 1995, M. Taylor 1995, Frankel and Rose 1996, Lothian and M. Taylor 1996, A. Taylor 2002).

In this paper we document the empirical importance of foreign currency as a hedge against real interest rate risk for long-term investors. We argue that foreign currency is not necessarily a purely speculative asset; it can play an important role in the portfolios of conservative long-term investors. Thus we argue against the presumption that short-term debt portfolios should always be fully domestic.

To keep the analysis simple and focused, we assume that the only available assets are domestic currency and foreign currency. We use the long-term portfolio choice theory of Campbell, Chan, and Viceira (henceforth CCV, 2002), an empirical implementation of Merton (1969, 1971, 1973), as a framework for the analysis. Campbell and Viceira (2002) provides a general overview of recent research on long-term portfolio choice.

The organization of the paper is as follows. Section 1 lays out the CCV framework. Section 2 estimates VAR models for real exchange rates and real interest rates. Section 3 reports implications for long-term portfolio choice. Section 4 concludes.

1. Long-Term Portfolio Choice in a VAR Model

The CCV model is set in discrete time. It assumes an infinitely-lived investor with Epstein-Zin (1989, 1991) recursive preferences defined over a stream of consumption. It allows a general specification for the number of securities and the dynamics of their expected returns; in particular it does not require that markets are complete.

1.1. Budget constraint

The model assumes that the investor lives only off financial wealth. Thus, in common with much of the finance literature, it ignores the existence of labor income. The intertemporal budget constraint is

$$W_{t+1} = (W_t - C_t) R_{p,t+1}, \tag{1}$$

where C_t is consumption and W_t is wealth at time t . $R_{p,t+1}$ denotes the gross portfolio return from t to $t + 1$.

There are n assets available for investment, so the real portfolio return $R_{p,t+1}$ is given by

$$R_{p,t+1} = \sum_{i=2}^n \alpha_{i,t} (R_{i,t+1} - R_{1,t+1}) + R_{1,t+1}, \quad (2)$$

where $\alpha_{i,t}$ is the portfolio weight on asset i . The first asset is a short-term instrument whose real return is $R_{1,t+1}$. Although we use the short-term return as a benchmark and measure other returns relative to it, we do not assume that this return is riskless. In practice we use the domestic nominal bill as the short-term asset; the nominal return on a nominal bill is riskless, but the real return is not because it is subject to short-term inflation risk.

1.2. Dynamics of state variables

CCV postulate that the dynamics of the relevant state variables are well captured by a first-order vector autoregressive process or VAR(1). In principle the use of a VAR(1) is not restrictive since any vector autoregression can be rewritten as a VAR(1) through an expansion of the vector of state variables.

We define a vector of log excess returns

$$\mathbf{x}_{t+1} \equiv \begin{bmatrix} r_{2,t+1} - r_{1,t+1} \\ r_{3,t+1} - r_{1,t+1} \\ \vdots \\ r_{n,t+1} - r_{1,t+1} \end{bmatrix}, \quad (3)$$

where $r_{i,t+1} \equiv \log(R_{i,t+1})$ for all i . In our empirical application, $r_{1,t+1}$ is the domestic real short rate and $r_{2,t+1}$ refers to the real domestic return on foreign currency.

We allow the system to include other state variables \mathbf{s}_{t+1} , such as the real exchange rate. Stacking $r_{1,t+1}$, \mathbf{x}_{t+1} , \mathbf{s}_{t+1} into an $m \times 1$ vector \mathbf{z}_{t+1} , we have

$$\mathbf{z}_{t+1} \equiv \begin{bmatrix} r_{1,t+1} \\ \mathbf{x}_{t+1} \\ \mathbf{s}_{t+1} \end{bmatrix}. \quad (4)$$

We will call \mathbf{z}_{t+1} the state vector and we assume a first order vector autoregression for \mathbf{z}_{t+1} :

$$\mathbf{z}_{t+1} = \Phi_0 + \Phi_1 \mathbf{z}_t + \mathbf{v}_{t+1}, \quad (5)$$

where Φ_0 is the $m \times 1$ vector of intercepts, Φ_1 is the $m \times m$ matrix of slope coefficients, and \mathbf{v}_{t+1} are the shocks to the state variables, assumed to be homoskedastic and normally distributed with arbitrary covariances.

The assumption of homoskedasticity is of course restrictive. It rules out the possibility that the state variables predict changes in risk; they can affect portfolio choice only by predicting changes in expected returns. Chacko and Viceira (1999) and Ait-Sahalia and Brandt (2001), among others, show how to solve long-term portfolio choice problems with changing risk.

1.3. Preferences and optimality conditions

We assume that the investor has Epstein-Zin (1989, 1991) recursive preferences. This preference specification has the desirable property that the notion of risk aversion is separated from that of the elasticity of intertemporal substitution. Following Epstein and Zin, we let

$$U(C_t, E_t(U_{t+1})) = \left[(1 - \delta) C_t^{\frac{1-\gamma}{\theta}} + \delta (E_t(U_{t+1}^{1-\gamma}))^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \quad (6)$$

where C_t is consumption at time t , $\gamma > 0$ is the relative risk aversion coefficient, $\psi > 0$ is the elasticity of intertemporal substitution, $0 < \delta < 1$ is the time discount factor, $\theta \equiv (1 - \gamma)/(1 - \psi^{-1})$, and $E_t(\cdot)$ is the conditional expectation operator. Epstein-Zin recursive utility nests as a special case the standard, time-separable power utility specification, in which $\gamma = \psi^{-1}$. Log utility obtains when we impose the additional restriction $\gamma = \psi^{-1} = 1$.

Epstein and Zin (1989, 1991) have shown that given the budget constraint (1), the Euler equation for consumption is

$$E_t \left[\left\{ \delta \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{1}{\psi}} \right\}^{\theta} R_{p,t+1}^{-(1-\theta)} R_{i,t+1} \right] = 1, \quad (7)$$

for any asset i , including the portfolio p itself. This first-order condition reduces to the standard one in the power utility case where $\gamma = \psi^{-1}$ and $\theta = 1$.

The investor's optimal consumption and portfolio policies must satisfy the Euler equation (7). When investment opportunities are constant, the optimal policies imply

a constant consumption-wealth ratio and a myopic portfolio rule—that is, the investor chooses her portfolio as if her investment horizon was only one period. However, when investment opportunities are time-varying, there are no known exact analytical solutions to this equation except for some specific values of γ and ψ . Giovannini and Weil (1989) have shown that with $\gamma = 1$, it is optimal for the investor to follow a myopic portfolio rule. They also show that with $\psi = 1$, the investor optimally chooses a constant consumption-wealth ratio equal to $(1 - \delta)$. However, with $\gamma = 1$, the optimal consumption-wealth ratio is not constant unless $\psi = 1$ and, conversely, with $\psi = 1$ the optimal portfolio rule is not myopic unless $\gamma = 1$. Thus the solution is fully myopic only when $\gamma = \psi = 1$, that is, with log utility. To solve for the optimal rules in all other cases, we extend the approximate analytical solution method in Campbell and Viceira (1999, 2001) to a multivariate framework.

1.4. Approximate solution

CCV show that the optimal portfolio and consumption rules take the form

$$\boldsymbol{\alpha}_t = \mathbf{a}_0 + \mathbf{A}_1 \mathbf{z}_t, \quad (8)$$

$$c_t - w_t = b_0 + \mathbf{b}'_1 \mathbf{z}_t + \mathbf{z}'_t \mathbf{B}_2 \mathbf{z}_t \quad (9)$$

That is, the optimal portfolio rule is linear in the VAR state vector but the optimal consumption rule is quadratic. \mathbf{a}_0 , \mathbf{A}_1 , b_0 , \mathbf{b}_1 , and \mathbf{B}_2 are constant coefficient matrices to be determined, with dimensions $(n - 1) \times 1$, $(n - 1) \times m$, 1×1 , $m \times 1$, and $m \times m$, respectively.³

The solution (8) is explained in detail by Campbell and Viceira (1999, 2001, 2002) and CCV, following Merton (1969, 1971, 1973) and Kim and Omberg (1996). The basic intuition is that conservative long-term investors, with risk aversion greater than one, will hold assets that covary negatively with shocks to expected returns on the optimal portfolio. Such assets hedge investors against the risk that investment opportunities will deteriorate.

³Only $m + (m^2 - m)/2$ elements of \mathbf{B}_2 are determined. The diagonal elements of \mathbf{B}_2 are unique, but the consumption-wealth ratio is determined by the sums of off-diagonal elements $b_{2,ij} + b_{2,ji}$ because $z_{i,t}z_{j,t} = z_{j,t}z_{i,t}$. Thus we can impose arbitrary normalizations on \mathbf{B}_2 provided that we leave each sum $b_{2,ij} + b_{2,ji}$ unrestricted. For example, we could restrict \mathbf{B}_2 to be symmetric, upper triangular, or lower triangular.

This intuition is easiest to understand in the special case where risk premia are constant, and only real interest rates vary. In this case the portfolio rule becomes a constant, and the log consumption-wealth ratio becomes linear. Campbell and Viceira (2002, Chapter 3) show that the portfolio rule can then be written as

$$\boldsymbol{\alpha} = \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} (\mathbf{E} \mathbf{x}_{t+1} + \boldsymbol{\sigma}^2 / 2) + \left(1 - \frac{1}{\gamma} \right) \boldsymbol{\Sigma}^{-1} (\boldsymbol{\sigma}_h - \boldsymbol{\sigma}_1), \quad (10)$$

where $\boldsymbol{\Sigma}$ is the variance-covariance matrix of excess returns relative to the benchmark return and $\boldsymbol{\sigma}^2$ is the vector of excess-return variances, the main diagonal of $\boldsymbol{\Sigma}$. The vector $\boldsymbol{\sigma}_h$ contains the covariances of each excess return with declines in expected future real interest rates:

$$\boldsymbol{\sigma}_h \equiv \text{Cov}(\mathbf{x}_{t+1}, -(\mathbf{E}_{t+1} - \mathbf{E}_t) \sum_{j=1}^{\infty} \rho^j r_{1,t+1+j}). \quad (11)$$

The use of the letter h here is intended to evoke Merton's concept of intertemporal hedging demand. $\boldsymbol{\sigma}_1$ is the vector of covariances of each risky asset's excess return with the benchmark return itself.

The first term in equation (10) is the traditional myopic demand for risky assets, which is based on their expected excess simple return (or equivalently, their expected excess log return plus one-half their variance), relative to their risk. This myopic demand is weighted by risk tolerance $1/\gamma$. The second term in (10) is the intertemporal hedging demand, which is weighted by $1 - 1/\gamma$. Conservative investors with $\gamma > 1$ hold assets that covary positively with declines in real interest rates. There is also an adjustment for the covariances of excess returns with the benchmark return, but this is small if one chooses a relatively stable benchmark asset.

Campbell and Viceira (2001) use these formulas to study the demand for long-term bonds in a model of the term structure of interest rates. In this paper, we can use these formulas if we assume that uncovered interest parity holds, so that expected excess returns on foreign assets over domestic assets are constant. If uncovered interest parity does not hold, we must use the more general solution (8).

The solution also simplifies if we assume that the elasticity of intertemporal substitution $\psi = 1$. In this case the log consumption-wealth ratio is constant and equal to $(1 - \delta)$, while the approximate solution (8) is exact in the limit of continuous time. In the empirical work of this paper, we assume $\psi = 1$ and set $\delta = 0.92$ at an annual rate; however our portfolio solutions are not highly sensitive to these assumptions.

2. Application to Foreign Exchange

2.1. Data description

In our empirical analysis we use quarterly data from the first quarter of 1973 through the fourth quarter of 2001. We consider four countries, the US, the UK, Germany, and Japan. We include the US as one member of each country pair; thus we consider three country pairs, the US and UK, the US and Germany, and the US and Japan.

For each country, we form the log ex-post real interest rate and for each pair of countries, we construct the log real exchange rate level. The real interest rate is the log three-month nominal short rate, less log realized inflation over the period, measured by the log change in the CPI. The log real exchange rate level is the log real foreign currency price of domestic currency, the sum of the log nominal exchange rate and the log domestic CPI, less the log foreign CPI. US data are taken from the CRSP database, and UK, German and Japanese data come from the IMF's International Financial Statistics (IFS). The German currency is the deutschemark through 1998, and then the euro; for simplicity we refer to the euro in our discussion.⁴

Table 1 reports summary statistics. The first two rows of the table show the mean and standard deviation of the ex post real interest rate, at an annual rate. Average real interest rates have been highest in Germany, and lowest in Japan. The unconditional standard deviation of real interest rates has been highest in the UK and lowest in Germany. This suggests that German short-term debt has been safest for investors, but of course investors should be concerned with conditional volatility rather than unconditional volatility, and long-term investors should be concerned with volatility over longer holding periods.

⁴Sources for the data are as follows. The nominal short rate in the US is the 90-day bill rate, from the CRSP Fama Treasury and Inflation file. In the UK it is the treasury bill rate, IFS 11160.ZF. In Japan and Germany it is the money market rate, IFS 15860B.ZF and 13460B.ZF. Nominal exchange rates are collected as the nominal US dollar price of foreign currency. All exchange rates are market exchange rates as of the end of the quarter. Exchange rates for the UK and Japan are IFS 112.AG.ZF and IFS 158.AE.ZF. Germany exchange rate data is the Dollar/Deutschemark exchange rate through 1/1/99, IFS 134.AE.ZF. After 1/1/99, we use the Dollar/Euro rate, IFS 163.AE.ZF. The conversion Deutschemark/Euro conversion ratio is 1.95583. The consumer price index for the US comes from the U.S. Treasury and Inflation Data Series file in CRSP. The UK, German and Japan CPI are IFS 11264.ZF, IFS 13464.ZF, and IFS 15864.ZF, respectively.

The next two rows report the mean change in the real exchange rate and the standard deviation of this change, both at annual rates. Since exchange rates are measured relative to the dollar, there are no entries for the US column. The UK and Japanese currencies have appreciated on average against the dollar over the 1973-2001 period, while the German currency has depreciated. The volatility of exchange rates is so great, however, that these average changes are not statistically significant if one assumes that exchange rate changes are serially uncorrelated. The standard deviation of the change in the dollar-pound real exchange rate is over 10% per quarter at an annual rate (5% in natural units), while for Germany and Japan these standard deviations are even larger at around 13% per quarter at an annual rate.

We have tested all our data series for unit roots. Standard Dickey-Fuller tests strongly reject a unit root in the ex post real interest rate for each country, but do not reject a unit root in the real exchange rate for any country pair. Nonetheless, recognizing the low power of these tests and the long-run evidence in favor of mean-reversion in the real exchange rate, we assume that the real exchange rate is stationary. The real exchange rate is an index number whose mean has no interpretation, but in the next row of the table we report its standard deviation. This is largest for Japan and smallest for the UK, reflecting both the volatility of quarterly changes and the persistence of these changes.

Finally, in the last three rows of Table 1 we report the mean inflation rate, the standard deviation of the inflation rate, and the standard deviation of inflation after removing deterministic seasonal fluctuations by regression on quarterly dummy variables. Since nominal interest rates and nominal exchange rates have no seasonals, we assume that inflation seasonals correspond to measurement error rather than to seasonal effects on the prices that are of ultimate concern to investors. Thus we use deseasonalized inflation to calculate real interest rates and exchange rates in Table 1 and in our subsequent empirical work. None of our results are very sensitive to seasonal adjustment. Deseasonalizing improves the fit of the VARs we estimate in the next section, but all coefficients that are significant using deseasonalized data are also significant in the raw data, and vice versa. The table shows that the UK has had the most inflation volatility, while Germany has had the least. This helps to explain why the UK real interest rate has been volatile and the German real interest rate comparatively stable.

2.2. VAR specification

To characterize the dynamics of real interest rates and inflation, we estimate a simple VAR that treats the home country and the foreign country symmetrically, and that includes each country's ex post real interest rate and the real exchange rate. We assume that the real exchange rate is stationary, so we include its level in the VAR. The VAR allows each country's real interest rate to be predicted by its own lagged real rate, the other country's real rate, and the lagged real exchange rate. Other variables could of course be included in the system, but this is the simplest VAR that allows us to impose and test the hypothesis of uncovered interest parity.

Our notation is as follows. We write r_{t+1} for the log domestic ex post real interest rate at time $t + 1$, r_{t+1}^* for the log foreign ex post real interest rate at time $t + 1$, and e_{t+1} for the log real domestic price of foreign currency. Thus a rise in e_{t+1} corresponds to an appreciation of the foreign currency and a depreciation of the domestic currency.

The VAR system can now be written as

$$z_{t+1} = \begin{bmatrix} r_{t+1} \\ r_{t+1}^* \\ e_{t+1} \end{bmatrix} = \begin{bmatrix} a_{10} \\ a_{20} \\ a_{30} \end{bmatrix} + \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} r_t \\ r_t^* \\ e_t \end{bmatrix} + \begin{bmatrix} u_{1,t+1} \\ u_{2,t+1} \\ u_{3,t+1} \end{bmatrix} \quad (12)$$

We estimate this system with free constant terms, since we have no prior beliefs about the average levels of real interest rates, and the average real exchange rate has no inherent meaning because the exchange rate is an index number. In our portfolio choice exercise, however, we will assume that average real interest rates are equal at home and abroad in order to focus on the risk aspects of foreign currency investment.

It is straightforward to impose uncovered interest parity on this system. Working in logs and neglecting constant terms, UIP implies that the expected domestic return on a foreign currency investment, $E_t[r_{t+1} + \Delta e_{t+1}]$, equals the expected domestic real interest rate, $E_t[r_{t+1}]$. Equivalently,

$$E_t \Delta e_{t+1} = E_t [r_{t+1} - r_{t+1}^*]. \quad (13)$$

In the VAR, we have

$$\begin{aligned}
E_t \Delta e_{t+1} &= a_{30} + a_{31}r_t + a_{32}r_t^* + (a_{33} - 1) e_t & (14) \\
E_t [r_{t+1} - r_{t+1}^*] &= (a_{10} + a_{11}r_t + a_{12}r_t^* + a_{13}e_t) - (a_{20} + a_{21}r_t + a_{22}r_t^* + a_{23}e_t)
\end{aligned}$$

Thus uncovered interest parity implies that

$$a_{30} + a_{31}r_t + a_{32}r_t^* + (a_{33} - 1) e_t = (a_{10} + a_{11}r_t + a_{12}r_t^* + a_{13}e_t) - (a_{20} + a_{21}r_t + a_{22}r_t^* + a_{23}e_t) \quad (15)$$

Since this restriction must hold for all values of the state variables, it determines three sets of restrictions on the VAR coefficients:

$$\begin{aligned}
\text{R1} &: a_{31} = a_{11} - a_{21} \\
\text{R2} &: a_{32} = a_{12} - a_{22} \\
\text{R3} &: a_{33} - 1 = a_{13} - a_{23}
\end{aligned}$$

We leave the constants unrestricted.

Although the VAR system (12) provides a succinct description of the data, it is not in the right form for application of the CCV machinery. The CCV model requires a VAR in which the first state variable is the domestic real interest rate r_{t+1} , the second state variable is the excess foreign-currency return $r_{t+1}^* + \Delta e_{t+1} - r_{t+1}$, and the other state variables capture the dynamics of (12). It turns out that this requires the addition of two state variables that are deterministically related, the change in the exchange rate and the level of the exchange rate. Thus we define a new state vector

$$Z_{t+1} = \begin{bmatrix} r_{t+1} \\ r_{t+1}^* + \Delta e_{t+1} - r_{t+1} \\ \Delta e_{t+1} \\ e_{t+1} \end{bmatrix}. \quad (16)$$

In the Appendix we show how (12) determines the elements of Φ_0 , Φ_1 in the VAR

$$Z_{t+1} = \Phi_0 + \Phi_1 Z_t + v_{t+1} \quad (17)$$

as well as the elements of the covariance matrix of the error vector v_{t+1} .

2.3. VAR estimation

Table 2 reports the estimates of the VAR system (12) for the US and the UK. The US is taken to be the home country, with real interest rate r_{t+1} , and the UK is the foreign country, with real interest rate r_{t+1}^* . However the system could be written the other way round and nothing would change except the notation.

The top part of the table reports the estimated VAR coefficients, with t statistics in parentheses. The ex post real interest rate in the US depends on its own lag, with a coefficient of 0.51 and a t statistic greater than 5. In addition, the lagged real interest rate in the UK has a statistically significant effect, but the coefficient on this variable is much smaller at 0.13. The lagged real exchange rate seems to have no effect on the US real interest rate. Overall, about 40% of the variation in the ex post real interest rate is predictable in the US.

The UK real interest rate is similar to the US real interest rate in several aspects of its dynamics. The UK rate is significantly predicted by its own lag, with a coefficient of 0.55, and by the US interest rate, with a coefficient of 0.44. In contrast with the US, however, the lagged exchange rate predicts the UK real interest rate with a small but statistically significant coefficient of 0.013. Recalling that the exchange rate is the dollar price of sterling, this says that a strong pound predicts high UK real interest rates, consistent with standard exchange rate theory. The overall predictability of the UK real rate is higher than that of the US rate, with an R^2 statistic of 57%.

Finally, the real exchange rate between the US and the UK follows a highly persistent process with a coefficient of 0.90 on its own lagged value. Lagged US interest rates negatively predict the real exchange rate, with a coefficient of -1.99 and a t statistic above 3. The coefficient on the UK rate is positive, but not significant.

These results say that when the US ex post real interest rate has been high (perhaps because of negative inflation shocks), the dollar tends to strengthen subsequently in real terms. This is inconsistent with uncovered interest parity (UIP) because a high US ex post real interest rate also predicts a higher real rate in the US relative to the UK, implying a positive excess return on US dollar investments relative to UK investments. A test that the VAR coefficients obey the restrictions of UIP rejects the null at the 5% significance level. More generally, however, the low-frequency behavior of the exchange rate seems consistent with the standard view that a country's currency tends to be strong when its real interest rate is high. When the pound is

strong the UK real interest rate is predicted to be high, and positive shocks to US real interest rates predict both high US real interest rates and a strong dollar.

The bottom panel of Table 2 reports the variances and correlations of the innovations to the three equations of the VAR system. Unsurprisingly, innovations in real interest rates in the US and the UK have a weak positive correlation of 0.164. The real exchange rate shock is almost uncorrelated with shocks to the real interest rate in both the UK and the US. The residual variance of the UK real rate is substantially higher than that of the US real rate. Although R^2 statistics show that a greater proportion of the variation in UK real interest rates is explained in the VAR, the unconditional standard deviation (reported in Table 1) is much higher in the UK (2.57%) than in the US (1.72%).

Table 3 reports a VAR system for the US and Germany. Many of the results are similar to those in Table 2, with some interesting differences. The US real interest rate is positively predicted by its own lag, as is the German real interest rate. The US real rate positively predicts the German real rate. Also, the lagged real dollar value of the euro predicts low real interest rates in the US. The dollar-euro real exchange rate, like the dollar-sterling real exchange rate, is strongly negatively predicted by the US real rate, while the coefficient on the German real rate is again positive, but not significant. Shocks to the exchange rate are weakly positively correlated with shocks to the US real rate and negatively correlated with shocks to the German real rate.

Table 4 reports a VAR system for the US and Japan. The Japanese real interest rate follows an autoregressive process, like the US and German real rates. The Japanese real rate is positively predicted by its own lag and by the US real rate. The lagged real exchange rate has no predictive power for either the US or the Japanese real rate. The real exchange rate is highly persistent. It is positively predicted by the Japanese real rate, while the coefficient on the US real rate is negative, but not significant. Shocks to the yen are negatively correlated with shocks to the Japanese real rate, and positively correlated with the shocks to the US real rate.

All these VAR systems are estimated freely, without imposing uncovered interest parity. The restrictions of UIP are rejected at the 10% level for all countries, and at the 5% level for the UK. In all country pairs, the dollar value of foreign currency is either positively predicted by the foreign real rate or negatively predicted by the US real rate. While there is some evidence against UIP on a quarterly basis, our VARs are consistent with a longer run version of UIP, in which exchange rates react to interest rate differentials with a short lag. The dollar-pound and dollar-euro real

exchange rates respond more to US interest rates, while the dollar-yen real exchange rate appears to respond more strongly to the Japanese interest rate.

3. Foreign Currency for Long-Term Investors

We now ask what our VAR systems imply for the optimal portfolio choices of long-term investors. For each country pair, we calculate the optimal portfolio of a long-term US investor and compare it with the optimal portfolio of a long-term investor based in the other country of the pair. We do this both for the unrestricted VARs reported in Tables 1, 2, and 3, and for restricted VARs that impose uncovered interest parity. In all cases we assume that average log real interest rates are equal in the two countries and that the real exchange rate has no time trend; that is, we assume equality of expected log returns on domestic and foreign currency.

We report results for relative risk aversion coefficients of 1, 5, and 2000 (effectively infinite). As risk aversion increases, the investor seeks to minimize risk without regard for any effects on the expected return of the portfolio. For comparison, we also report the portfolio that minimizes the variance of the real one-period portfolio return. This portfolio would be held by an extremely conservative short-term investor.

Table 5 reports results for the US and UK country pair. The top panel gives optimal portfolios for a US investor, while the bottom panel gives results for a UK investor. When risk aversion equals one, the optimal portfolio is 50% domestic, and 50% foreign. This result is an artefact of our assumptions that the average log real interest rate is the same in the US and the UK, and that the real exchange rate is stationary with no time trend. These assumptions imply that domestic currency and foreign currency have the same average log return. It follows from the relation between log portfolio return and log returns on individual assets that the portfolio that maximizes average log return is equally weighted in the two currencies.⁵ An investor with unit risk aversion seeks to maximize average log return, so he places equal weight on domestic and foreign currency.

⁵See Campbell and Viceira (2002), equation (2.21). Up to a loglinear approximation, which is exact in continuous time, the log return on a portfolio of two assets is α_t times the return on the first asset, plus $(1 - \alpha_t)$ times the return on the second asset, plus $\alpha_t(1 - \alpha_t)\sigma_t^2/2$, where α_t is the portfolio weight on the first asset and σ_t^2 is the variance of the excess return on one asset over the other. Thus, if expected log returns on the two assets are equal, $\alpha_t = 1/2$ maximizes the expected log portfolio return for any value of σ_t^2 .

Another way to understand this phenomenon is to note that with random real exchange rates and equal log returns, Jensen's Inequality increases the average simple real return on foreign currency above the average simple real return on domestic currency. This is true no matter which country is the home country; US investors perceive higher average returns on UK currency, while UK investors perceive higher average returns on US currency. Thus investors in each country have a speculative motive for holding the other's currency. A similar effect on nominal returns has been called "Siegel's paradox" in the international finance literature (Siegel 1972, Obstfeld and Rogoff 1996, Chapter 8).

As risk aversion increases, investors pay less attention to average returns and more attention to the risks of domestic and foreign currency. A US investor reduces his holding of UK currency, and moves towards a portfolio that consists almost entirely of US currency. With risk aversion of 2000, the optimal portfolio for the US investor, based on the unrestricted VAR, is long 101% domestic currency, and short 1% foreign currency. This is virtually identical to the short-term minimum-variance portfolio, long 100% domestic currency. The optimal portfolio weights are similar if we impose UIP on the VAR system.

Results are very different for a UK investor. As risk aversion increases, the optimal portfolio for a UK investor places greater weight on US currency. With risk aversion of 2000, the optimal portfolio given an unrestricted VAR is 66% foreign currency and only 34% domestic currency. This is true even though the short-term minimum variance portfolio for the UK investor, like that for the US investor, is dominated by domestic currency (97% domestic and only 3% foreign). Clearly the risks perceived by a long-term UK investor are not the same as those perceived by a short-term UK investor.

One way to understand these results is to look at a plot of risk against investment horizon. The top panel of Fig. 1 shows the annualized standard deviation of the return over k quarters, for k from 1 to 100, as perceived by a US investor who believes the data are generated by the unrestricted VAR model of Table 1. The standard deviation of a US (domestic) Treasury bill rises from about 1.5% at a short horizon to about 3% after three or four years, consistent with the findings of Campbell and Viceira (2002, Chapter 4). The annualized standard deviation of a UK (foreign) Treasury bill is over 10% over a quarter, reflecting the short-term variability of the exchange rate. This risk diminishes gradually because of the estimated mean-reversion in the real exchange rate, but even at long horizons a UK bill has an

annualized standard deviation of about 7%. Thus at all horizons UK currency is far riskier than US currency for a US investor.

The bottom panel of Fig. 1 is the equivalent diagram for a UK investor. At short horizons, UK currency is comparatively safe with a standard deviation of about 2%, while US currency has the 10% standard deviation caused by exchange rate variability. For the UK investor, however, the long-term uncertainty of the real interest rate is a much more serious matter. The annualized standard deviation of a UK bill portfolio rises above 5% at long horizons. The standard deviation of US currency declines rapidly with the horizon, and falls below the standard deviation of UK currency at a horizon of about 10 years.

Fig. 1 is derived from the VAR system estimated in Table 2, and it reflects the dynamics of that system. UK currency is risky at short horizons because of the historical volatility of UK inflation; it is risky at long horizons because shocks propagate over time through the US real interest rate and the real exchange rate. US currency is appealing to UK investors because the US real interest rate has been more stable than the UK real interest rate in both the short and the long term, and because real exchange rate fluctuations are transitory, so real exchange rate risk dies away over time.

It is important to note that long-horizon investors do not choose a portfolio using mean-variance analysis applied to the long-horizon risks plotted in Fig. 1. That would be appropriate only for investors who must make a one-time portfolio allocation without any ability to rebalance over time as in Barberis (2000). Nonetheless, the same features of the VAR that give US currency a low long-term volatility also tilt the portfolio of a conservative UK investor towards US currency. In particular, the fact that a strong pound predicts high UK real interest rates, shown in Table 2, means that UK investors can use US dollar investments to hedge the risk that UK real interest rates will decline. If the pound weakens, UK investors holding dollar assets will be wealthier in sterling terms, and this will compensate them for the predicted decline in UK investment opportunities.

Table 6 reports the optimal portfolio holdings of US and German investors who can hold dollars and euros. Here we find that the euro has great appeal to conservative long-term US investors. A conservative long-term US investor holds 65% of wealth in euros and 35% in dollars. Long-term German investors are 10% short dollars. The long-term portfolio for the US investor contrasts dramatically with the minimum-variance short-term portfolio, which is conventionally home-biased with a weight on

domestic currency of approximately 100%.

Once again these results can be understood by considering the risks of US and German currency at different horizons. The top panel of Fig. 2 shows that euros become safer than dollars for US investors at horizons longer than 11 or 12 years, while dollars are much riskier than euros for German investors at all horizons. These patterns reflect the structure of the VAR estimated in Table 3, and especially the fact that a strong euro predicts low real interest rates in the US. This means that US investors can use German investments to hedge the risk that US real interest rates will decline. If the dollar weakens, US investors holding euros will be wealthier in dollar terms, and this will compensate them for the predicted decline in US investment opportunities.

Table 7 reports results for US and Japanese investors. The results are less dramatic than the UK and German cases, but still reflect the importance of long-term considerations for both US and Japanese investors. The optimal portfolio of a conservative long-term US investor has a 12% weight on the Japanese currency, while conservative long-term Japanese investors hold 28% of their wealth in US dollars. Fig. 3 plots risks against horizon for the US-Japanese case. Foreign currency remains riskier than domestic currency at all horizons in Fig. 3, but the difference in risk narrows considerably as the horizon lengthens.

So far we have found that long-term UK investors prefer dollars to pounds, and long-term US investors prefer euros to dollars. These results suggest that long-term UK investors might have a strong desire to hold euros in a direct comparison of the pound against the euro. When we estimate a VAR system for the UK and Germany, we find that the euro is safer than the pound for UK investors with a horizon longer than about 7 years, as illustrated in Fig. 4. A long-term UK investor with a risk aversion coefficient of 2000 holds 84% of wealth in euros; even if risk aversion is only 5, the optimal weight on euros is still 76%. This contrasts with the short-term minimum-variance portfolio for the UK investor, which has less than 3% weight in euros.

These findings might be used to support a decision by the UK to adopt the euro as its currency. To the extent that one function of monetary policy is to provide a stable long-term store of value, the euro (and its predecessor the deutschmark) has been a more successful currency than the pound. Euros are more attractive than pounds to UK-based long-term investors, but would be still more attractive if real exchange rate risk between the pound and the euro were removed altogether, as

illustrated by a comparison of the dashed line in the top panel of Fig. 4, showing euro risk for pound-based investors, with the solid line in the bottom panel of Fig. 4, showing euro risk for euro-based investors.

This argument should not be pushed too hard, however. Monetary policy in both the UK and the euro area has evolved over time, and the independent Bank of England has a good track record in controlling inflation since 1997. The UK government already offers investors a stable long-term store of value by issuing long-term inflation-indexed bonds (index-linked gilts). Finally, monetary policy has other functions that should also be taken into account in any decision on the future of the pound.

4. Conclusion

In this paper we have argued that foreign currency can be an attractive asset class, not just for exchange rate speculators, but also for long-term investors who wish to hedge the risk that domestic real interest rates will decline. This provides a new way to understand the function of international bond funds that hold portfolios of bonds denominated in different currencies. From a short-term perspective, such funds merely add exchange rate risk to domestic bond funds without any offsetting advantages; they only make sense if one adopts a longer-term perspective and applies Merton's theory of intertemporal hedging demand.

Looking at data on the US dollar, the pound sterling, the deutschmark/euro, and the yen since 1973, we have found that the intertemporal hedging demand for foreign currency can be surprisingly large. For example, we have estimated that a highly conservative long-term UK investor allowed to hold domestic currency and a single foreign currency should hold 66% of wealth overseas if the foreign currency is the dollar, and 84% of wealth overseas if the foreign currency is the euro. In general, the currencies that should be attractive to foreign investors are those with stable real interest rates that are not correlated with their exchange rates. One might expect that the currencies of large economies would have these characteristics, and indeed we find that the US dollar and euro are relatively attractive.

The empirical work in this paper has several limitations that should be kept in mind when evaluating the results. First, we have considered only two countries at a time. Second, for each country pair we have assumed that the only two assets available are domestic bills and foreign bills. We have ruled out other intertemporal

hedging assets, such as long-term nominal and inflation-indexed bonds; we have also excluded equities from the analysis, so we have nothing to say about the optimal currency hedge ratio for a foreign equity position. Third, we have used an extremely parsimonious model that includes only three forecasting variables: each country's lagged real interest rates and the lagged real exchange rate. In particular, we have excluded the nominal interest differential which might be a good predictor of exchange rate and real interest rate movements.

Fourth, we have assumed that the real exchange rate follows a stationary process. To the extent that permanent shocks to the real exchange rate occur—perhaps because of productivity shocks that shift the long-run equilibrium price of nontraded goods relative to traded goods—our model is misspecified. Fifth, we have looked at a relatively short time period covering the last three decades. Mean-reversion in the real exchange rate is notoriously slow and hard to estimate over short sample periods, so our estimates of exchange rate dynamics are quite imprecise. We have ignored this by using the point estimates as if they were known with certainty. Finally, we have assumed that interest rate and exchange rate dynamics are constant throughout our sample period. This ignores the fact that major changes in monetary policy have taken place in all four countries during the last part of the 20th Century. The historical volatility of UK inflation, for example, probably overstates the volatility of UK inflation under the current policy regime with an independent Bank of England.

We plan to address all these limitations in future work. The model of Campbell, Chan, and Viceira (2002), which we use to solve for optimal long-term portfolios, is flexible enough to handle multiple countries, multiple assets within each country, and additional forecasting variables. We can modify the assumption of real exchange rate stationarity if we can identify a stationary component of the real exchange rate that is linked with real interest rate fluctuations. And we can explore the effects of alternative parameter estimates from longer or shorter sample periods or Bayesian estimation procedures. We expect that the basic intuition developed in this paper will survive these extensions to richer and more realistic environments.

References

- Aït-Sahalia, Y. and Brandt, M. (2001). ‘Variable selection for portfolio choice.’ *Journal of Finance* vol. 56, pp. 1297–1351.
- Baillie, R.T. and Bollerslev, T. (2000). ‘The forward premium anomaly is not as bad as you think.’ *Journal of International Money and Finance* vol. 19, pp. 471–488.
- Barberis, N.C. (2000). ‘Investing for the long run when returns are predictable.’ *Journal of Finance* vol. 55, pp. 225–264.
- Baxter, M. and Jermann, U. (1997). ‘The international diversification puzzle is worse than you think.’ *American Economic Review* vol. 87, pp. 170–180.
- Bekaert, G. and Hodrick, R.J. (2001). ‘Expectations hypotheses tests.’ *Journal of Finance* vol. 56, pp. 115–138.
- Bekaert, G., Wei, M. and Xing, Y. (2002). ‘Uncovered interest parity and the term structure.’ NBER Working Paper No. 8795.
- Bottazzi, L., Pesenti, P., and van Wincoop, E. (1996). ‘Wages, profits, and the international portfolio puzzle.’ *European Economic Review* vol. 40, pp. 219–254.
- Campbell, J.Y., Chan, Y.L., and Viceira, L.M. (2002). ‘A multivariate model of strategic asset allocation.’ Forthcoming *Journal of Financial Economics*.
- Campbell, J.Y. and Clarida, R.H. (1987). ‘The dollar and real interest rates.’ *Carnegie-Rochester Conference Series on Public Policy* vol. 27, pp. 103–140.
- Campbell, J.Y. and Shiller, R.J. (1996). ‘A scorecard for indexed government debt.’ In Ben S. Bernanke and Julio Rotemberg eds., *National Bureau of Economic Research Macroeconomics Annual* vol. 11, pp. 155–197.
- Campbell, J.Y. and Viceira, L.M. (1999). ‘Consumption and portfolio decisions when expected returns are time varying.’ *Quarterly Journal of Economics* vol. 114, pp. 433–495.
- Campbell, J.Y. and Viceira, L.M. (2001). ‘Who should buy long-term bonds?’ *American Economic Review* vol. 91, pp. 99–127.

- Campbell, J.Y. and Viceira, L.M. (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford: Oxford University Press.
- Chacko, G. and Viceira, L.M. (1999). ‘Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets.’ NBER Working Paper No. 7377.
- Chinn, M. and Meredith, G. (2001). ‘Testing uncovered interest parity at short and long horizons.’ Unpublished paper, UC Santa Cruz and IMF.
- Chue, T.K. (2002). ‘Explaining the lack of global diversification: Can local envy help?’ Unpublished paper, Hong Kong University of Science and Technology.
- DeMarzo, P.M., Kaniel, R., and Kremer, I. (2002). ‘Community effects and externalities in portfolio choice.’ Unpublished paper, Stanford University.
- Dornbusch, R. (1976). ‘Expectations and exchange rate dynamics.’ *Journal of Political Economy* vol. 84, pp. 1161–1176.
- Engel, C. (1996). ‘The forward discount anomaly and the risk premium: A survey of recent evidence.’ *Journal of Empirical Finance* vol. 3, pp. 123–192.
- Epstein, L. and Zin, S. (1989). ‘Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework.’ *Econometrica* vol. 57, pp. 937–69.
- Epstein, L. and Zin, S. (1991). ‘Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical investigation.’ *Journal of Political Economy* vol. 99, pp. 263–286.
- Fama, E.F. (1984). ‘Forward and spot exchange rates.’ *Journal of Monetary Economics* vol. 14, pp. 319–338.
- Frankel, J.A. and Rose, A.K. (1996). ‘A panel project on purchasing power parity: Mean reversion within and between countries.’ *Journal of International Economics* vol. 40, pp. 209–224.
- Froot, K.A. (1993). ‘Currency hedging over long horizons.’ NBER Working Paper No. 4355.

- Froot, K.A. and Rogoff, K. (1995). ‘Perspectives on PPP and long-run real exchange rates.’ In G. Grossman and K. Rogoff eds. *Handbook of International Economics Vol. 3*. Amsterdam: North-Holland.
- Froot, K.A. and Thaler, R.H. (1990). ‘Foreign exchange.’ *Journal of Economic Perspectives* vol. 4, pp. 179–192.
- Giovannini, A. and Weil, P. (1989). ‘Risk aversion and intertemporal substitution in the capital asset pricing model.’ NBER Working Paper No. 2824.
- Goetzmann, W.N., Li, L., and Rouwenhorst, K.G. (2002). ‘Long-term global market correlations.’ Working paper, International Center for Finance, Yale University.
- Hodrick, R.J. (1987). *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*. Chur, Switzerland: Harwood Academic Publishers.
- Kim, T. S. and Omberg, E. (1996). ‘Dynamic nonmyopic portfolio behavior.’ *Review of Financial Studies* vol. 9, pp. 141–161.
- Lewis, K. (1999). ‘Trying to explain home bias in equities and consumption.’ *Journal of Economic Literature* vol. 37, pp. 571–608.
- Lothian, J.R. and Taylor, M.P. (1996). ‘Real exchange rate behavior: The recent float from the perspective of the past two centuries.’ *Journal of Political Economy* vol. 104, pp. 488–509.
- Merton, R.C. (1969). ‘Lifetime portfolio selection under uncertainty: The continuous time case.’ *Review of Economics and Statistics* vol. 51, pp. 247–257.
- Merton, R.C. (1971). ‘Optimum consumption and portfolio rules in a continuous-time model.’ *Journal of Economic Theory* vol. 3, pp. 373–413.
- Merton, R.C. (1973). ‘An intertemporal capital asset pricing model.’ *Econometrica* vol. 41, pp. 867–87.
- Obstfeld, M. and Rogoff, K. (1996). *Foundations of International Macroeconomics*. Cambridge, MA: MIT Press.
- Shore, S. and White, J. (2001). ‘External habit formation and the home bias puzzle.’ Unpublished paper, Harvard University.

- Siegel, J. (1972). ‘Risk, interest rates, and the forward exchange.’ *Quarterly Journal of Economics* vol. 86, pp. 303–309.
- Siegel, J. (1998). *Stocks for the Long Run*, 2nd ed. New York, NY: McGraw-Hill.
- Solnik, B.H. (1974). ‘An equilibrium model of the international capital market.’ *Journal of Economic Theory* vol. 8, pp. 500–524.
- Taylor, A.M. (2002). ‘A century of purchasing-power parity.’ *Review of Economics and Statistics* vol. 84, pp. 139–150.
- Taylor, M.P. (1995). ‘The economics of exchange rates.’ *Journal of Economic Literature* vol. 33, pp. 13–47.
- Wheatley, S.M. (2001). ‘Keeping up with the Joneses, human capital, and the home equity bias.’ Unpublished paper, University of Melbourne.
- Wilson, M. (2002). ‘The relationship of stock market returns and exchange rate movements over time.’ Unpublished paper, Harvard University.

Appendix

In this appendix we show how to transform (12) into the form (16) for which the portfolio choice results of CCV apply.

Notice the relationship between Z_{t+1} and z_t and z_{t+1} :

$$Z_{t+1} = Hz_t + \Xi z_{t+1},$$

where H and Ξ are defined by

$$\begin{bmatrix} r_{t+1} \\ r_{t+1}^* + \Delta e_{t+1} - r_{t+1} \\ \Delta e_{t+1} \\ e_{t+1} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 0 & -1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} r_t \\ r_t^* \\ e_t \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_{t+1} \\ r_{t+1}^* \\ e_{t+1} \end{bmatrix}.$$

Now note that Hz_t can be recovered linearly from Z_{t+1} , so that $Hz_t = \Psi Z_{t+1}$:

$$Hz_t = \begin{bmatrix} 0 \\ -e_t \\ -e_t \\ 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} r_{t+1} \\ r_{t+1}^* + \Delta e_{t+1} - r_{t+1} \\ \Delta e_{t+1} \\ e_{t+1} \end{bmatrix} = \Psi Z_{t+1}$$

Therefore, we have $Z_{t+1} = \Psi Z_{t+1} + \Xi z_{t+1}$, or $(I - \Psi) Z_{t+1} = \Xi z_{t+1}$. This relationship holds at all time periods, so $(I - \Psi) Z_t = \Xi z_t$, and

$$\begin{aligned} \Xi' (I - \Psi) Z_t &= \Xi' \Xi z_t \\ (\Xi' \Xi)^{-1} \Xi' (I - \Psi) Z_t &= z_t \end{aligned}$$

From the original VAR system (12),

$$\begin{aligned} z_{t+1} &= A_0 + A_1 z_t + u_{t+1}, \\ Hz_t + \Xi z_{t+1} &= \Xi A_0 + (H + \Xi A_1) z_t + \Xi u_{t+1}, \\ Z_{t+1} &= \Xi A_0 + (H + \Xi A_1) (\Xi' \Xi)^{-1} \Xi' (I - \Psi) Z_t + \Xi u_{t+1}. \end{aligned}$$

This is the VAR for Z_{t+1} , (16), with

$$\begin{aligned} v_{t+1} &= \Xi u_{t+1}, \\ \Phi_0 &= \Xi A_0, \\ \Phi_1 &= (H + \Xi A_1) (\Xi' \Xi)^{-1} \Xi' (I - \Psi). \end{aligned}$$

The covariance matrix of the residuals is

$$\Sigma_v = \text{var} [v_{t+1}] = \Xi \Sigma_u \Xi'$$

It is singular because the innovation in e_{t+1} must equal the innovation in Δe_{t+1} .

TABLE 1
Summary Statistics

	US	UK	Germany	Japan
$E[r_t]$	2.238	1.798	2.657	1.289
$\sigma[r_t]$	1.635	2.572	1.192	2.170
$E[\Delta e_t]$	0.000	0.487	-0.640	1.409
$\sigma[\Delta e_t]$	0.000	10.589	12.855	13.011
$\sigma[e_t]$	0.000	26.074	34.379	45.489
$E[\pi_t]$	4.914	7.062	3.010	3.463
$\sigma[\pi_t]$	1.839	3.314	1.370	2.816
$\sigma[\pi_{DS,t}]$	1.765	3.035	1.258	2.670

Note: r_t = ex post real short rate, e_t = real exchange rate, π_t = CPI inflation, $\pi_{DS,t}$ = CPI inflation, deseasonalized using seasonal dummy variables. Real short rates and exchange rates calculated using deseasonalized inflation. All quantities expressed in annualized percentage terms.

TABLE 2
VAR Estimation Results:
US and UK

Dependent Variable	r_t (t)	r_t^* (t)	e_t (t)	R^2 (p)
VAR Estimation Results				
r_{t+1}	0.506 (5.471)	0.129 (2.618)	-0.003 (-0.526)	0.404 (0.000)
r_{t+1}^*	0.444 (2.760)	0.547 (3.868)	0.013 (2.095)	0.569 (0.000)
e_{t+1}	-1.993 (-3.137)	0.422 (1.099)	0.901 (21.752)	0.855 (0.000)
Cross-Correlation of Residuals				
	r	r^*	e	
r	0.627	0.164	0.072	
r^*	-	5.007	0.010	
e	-	-	4.950	

Note: r_t = domestic (US) ex post real short rate, r_t^* = foreign (UK) ex post real short rate, e_t = real exchange rate (dollars per pound). Data are quarterly, 1973Q1 through 2001Q4.

TABLE 3
VAR Estimation Results:
US and Germany

Dependent Variable	r_t (t)	r_t^* (t)	e_t (t)	R^2 (p)
VAR Estimation Results				
r_{t+1}	0.376 (3.699)	0.104 (1.253)	-0.018 (-4.526)	0.469 (0.000)
r_{t+1}^*	0.177 (2.475)	0.325 (4.282)	0.002 (0.612)	0.207 (0.000)
e_{t+1}	-2.148 (-2.605)	0.427 (0.403)	0.888 (22.606)	0.874 (0.000)
Cross-Correlation of Residuals				
	r	r^*	e	
r	0.592	0.248	0.047	
r^*	-	6.048	-0.104	
e	-	-	6.092	

Note: r_t = domestic (US) ex post real short rate, r_t^* = foreign (German) ex post real short rate, e_t = real exchange rate (dollars per euro). Data are quarterly, 1973Q1 through 2001Q4.

TABLE 4
VAR Estimation Results:
US and Japan

Dependent Variable	r_t (t)	r_t^* (t)	e_t (t)	R^2 (p)
VAR Estimation Results				
r_{t+1}	0.550 (5.895)	0.089 (0.980)	-0.002 (-0.742)	0.386 (0.000)
r_{t+1}^*	0.311 (3.519)	0.449 (2.424)	0.005 (1.360)	0.400 (0.000)
e_{t+1}	-1.502 (-1.487)	0.983 (2.108)	0.941 (37.084)	0.925 (0.000)
Cross-Correlation of Residuals				
	r	r^*	e	
r	0.637	0.272	0.079	
r^*	-	6.121	-0.082	
e	-	-	6.174	

Note: r_t = domestic (US) ex post real short rate, r_t^* = foreign (Japanese) ex post real short rate, e_t = real exchange rate (dollars per yen). Data are quarterly, 1973Q1 through 2001Q4.

TABLE 5
Portfolio Allocation Results: US and UK

Domestic US, Foreign UK					
		Relative Risk Aversion			Min. Short-Term
		1	5	2000	Variance
Unrestricted VAR	Domestic	50.0	83.5	100.5	99.7
	Foreign	50.0	16.5	-0.5	0.3
UIP	Domestic	50.0	95.7	107.1	99.7
	Foreign	50.0	4.3	-7.1	0.3
Domestic UK, Foreign US					
		Relative Risk Aversion			Min. Short-Term
		1	5	2000	Variance
Unrestricted VAR	Domestic	50.0	39.3	33.8	97.3
	Foreign	50.0	60.7	66.2	2.7
UIP	Domestic	50.0	34.2	30.3	97.3
	Foreign	50.0	65.8	69.7	2.7

TABLE 6
Portfolio Allocation Results: US and Germany

Domestic US, Foreign Germany					
		Relative Risk Aversion			Min. Short-Term
		1	5	2000	Variance
Unrestricted VAR	Domestic	50.0	39.9	35.1	99.7
	Foreign	50.0	60.1	64.9	0.3
UIP	Domestic	50.0	39.3	36.6	99.7
	Foreign	50.0	60.7	63.4	0.3
Domestic Germany, Foreign US					
		Relative Risk Aversion			Min. Short-Term
		1	5	2000	Variance
Unrestricted VAR	Domestic	50.0	92.0	110.1	100.4
	Foreign	50.0	8.0	-10.1	-0.4
UIP	Domestic	50.0	97.7	109.6	100.4
	Foreign	50.0	2.3	-9.6	-0.4

TABLE 7
Portfolio Allocation Results: US and Japan

Domestic US, Foreign Japan					
		Relative Risk Aversion			Min. Short-Term
		1	5	2000	Variance
Unrestricted VAR	Domestic	50.0	74.0	88.1	100.1
	Foreign	50.0	26.0	11.9	-0.1
UIP	Domestic	50.0	82.6	90.8	100.1
	Foreign	50.0	17.4	9.2	-0.1
Domestic Japan, Foreign US					
		Relative Risk Aversion			Min. Short-Term
		1	5	2000	Variance
Unrestricted VAR	Domestic	50.0	64.6	72.4	99.7
	Foreign	50.0	35.4	27.6	0.3
UIP	Domestic	50.0	70.6	75.7	99.7
	Foreign	50.0	29.4	24.3	0.3

Figure 1: Risk vs. Horizon, US and UK





Figure 2: Risk vs. Horizon, US and Germany





Figure 3: Risk vs. Horizon, US and Japan





Figure 4: Risk vs. Horizon, UK and Germany

