ELSEVIER

International risk sharing is better than you think, or exchange rates are too smooth ☆

Michael W. Brandt[a], John H. Cochrane[b,*], Pedro Santa-Clara[c]

[a]*Fuqua School of Business, Duke University, Durham, NC 27708, USA*
[b]*Graduate School of Business, University of Chicago, Chicago, IL 60637, USA*
[c]*Anderson Graduate School of Management, UCLA, Los Angeles, CA 90095, USA*

Abstract

Exchange rates depreciate by the difference between domestic and foreign marginal utility growth or discount factors. Exchange rates vary a lot, as much as 15% per year. However, equity premia imply that marginal utility growth varies much more, by at least 50% per year. Therefore, marginal utility growth must be highly correlated across countries: international risk sharing is better than you think. Conversely, if risks really are not shared internationally, exchange rates should vary more than they do: exchange rates are too smooth. We calculate an index of international risk sharing that formalizes this intuition. We treat carefully the realistic case of incomplete capital markets. We contrast our estimates with the poor risk sharing suggested by consumption data and home-bias portfolio calculations.
© 2006 Elsevier B.V. All rights reserved.

*Corresponding author. Tel.: +1 773 702 3059; fax: +1 773 834 2031.

E-mail address: john.cochrane@gsb.uchicago.edu (J.H. Cochrane).

1. Introduction

Exchange rates e_{t+1} are linked to domestic and foreign marginal utility growth or discount factors m_{t+1}^d and m_{t+1}^f by the equation[1]

$$\ln \frac{e_{t+1}}{e_t} = \ln m_{t+1}^f - \ln m_{t+1}^d. \tag{1}$$

Take variances of both sides of the equation. Exchange rates vary a lot, as much as 15% per year. However, equity premia imply that marginal utility growth varies much more, by at least 50% per year.[2] For Eq. (1) to hold, therefore, marginal utility growth must be highly correlated across countries: international risk sharing is better than you think. Conversely, if risks are not shared internationally—if marginal utility growth is uncorrelated across countries—exchange rates should vary by $\sqrt{2} \times 50\% = 71\%$ or more: exchange rates are too *smooth*.

To quantify this idea, we compute the following index of international risk sharing:

$$1 - \frac{\sigma^2(\ln m_{t+1}^f - \ln m_{t+1}^d)}{\sigma^2(\ln m_{t+1}^f) + \sigma^2(\ln m_{t+1}^d)} = 1 - \frac{\sigma^2\left(\ln \frac{e_{t+1}}{e_t}\right)}{\sigma^2(\ln m_{t+1}^f) + \sigma^2(\ln m_{t+1}^d)}. \tag{2}$$

The numerator measures how different marginal utility growth is across the two countries—how much risk is *not* shared. The denominator measures the volatility of marginal utility growth in the two countries—how much risk there is to share. A 10% exchange rate volatility and a 50% marginal utility growth volatility in each country imply a risk sharing index of $1 - 0.1^2/(2 \times 0.5^2) = 0.98$. Our detailed calculations in Table 2 give similar results. It seems that a lot of risk is shared.

Our index is not quite the same as a correlation. Like a correlation, our index is equal to one if $\ln m^f = \ln m^d$, it is equal to zero if $\ln m^f$ and $\ln m^d$ are uncorrelated, and it is equal to minus one if (pathologically) $\ln m^f = -\ln m^d$. However, risk sharing requires that domestic and foreign marginal utility growth are *equal*, not just perfectly correlated, and our index detects violations of scale as well as of correlation. For example, if $\ln m^f = 2 \times \ln m^d$, risks are not perfectly shared despite perfect correlation. In this case our index is 0.8. Like the correlation of discount factors or of consumption growth, our index is a statistical description of how far we are from perfect risk sharing. Higher risk sharing indices do not necessarily imply higher welfare.

The risk sharing index is essentially an extension of Hansen–Jagannathan (1991) bounds to include exchange rates. Hansen and Jagannathan showed that marginal utility growths must be highly volatile in order to explain the equity premium. We show that marginal

utility growths must also be highly correlated across countries in order to explain the relative smoothness of exchange rates.

Like the equity premium, we regard this result as a puzzle. Common intuition and calculations based on consumption data or observed portfolios suggest much less risk sharing. For example, if agents in each country have the same power utility function over a consumption composite, $\sum \beta^t c_t^{(1-\gamma)}/(1-\gamma)$, then the discount factors are

$$\ln m_{t+1}^d = \ln \beta - \gamma \ln(\Delta c_{t+1}^d) \quad \text{and} \quad \ln m_{t+1}^f = \ln \beta - \gamma \ln(\Delta c_{t+1}^f). \tag{3}$$

Consumption growth is correlated across countries—we find correlations as large as 0.41 below—but both the correlation and risk sharing indices as in Eq. (2) come nowhere near those implied by asset-market data.

Yet the conclusion is hard to escape. Our calculation uses only price data, and no quantity data or economic modeling (utility functions, income or productivity shock processes, and so forth). A large degree of international risk sharing is an inescapable logical conclusion of Eq. (1), a reasonably high equity premium (over 1%, as we show below), and the basic economic proposition that price ratios measure marginal rates of substitution.

If the puzzle is eventually resolved following the path of the equity premium literature, with utility functions and environments that deliver high equity premia and relatively smooth exchange rates (volatile and highly correlated marginal rates of substitution), then the puzzle will be resolved in favor of the asset market view that international risk sharing is, in fact, better than we now think. The only way it can be resolved in favor of poor international risk sharing is if the vast majority of consumers are far from their first order conditions, further off than can be explained by observed financial market imperfections, so that measurements of marginal rates of substitution from price ratios are wrong by orders of magnitude. However, like the equity premium, our contribution is to articulate the puzzle, not to opine on which underlying view will have to be modified in order to resolve the puzzle.

What exactly does the calculation mean? To start a more careful interpretation of the calculation, we discuss two major potential impediments to risk sharing: transport costs and incomplete asset markets. In the remainder of the paper we also consider the poor correlation of consumption across countries, the home bias puzzle, frictions that separate marginal utility from prices, and the possibility that the equity premium is a lot lower than past average returns.

1.1. Transport costs

Risk sharing requires frictionless *goods* markets. The container ship is a risk sharing innovation as important as 24 hour trading.[3] Suppose that Earth trades assets with Mars by radio, in complete and frictionless capital markets. If Mars enjoys a positive shock, Earth-based owners of Martian assets rejoice in anticipation of their payoffs. But trade with Mars is still impossible, so the real exchange rate between Mars and Earth must

---

[3]Obstfeld and Rogoff (2000) emphasize the importance of transport costs for understanding a wide range of puzzles in international economics. Alesina et al. (2003) find that real exchange rate volatility increases with distance between two countries. Ravn (2001) emphasizes the importance of real exchange rate movements in risk sharing measurements.

adjust exactly to offset any net payoff. In the end, Earth marginal utility growth must reflect Earth resources, and the same for Mars. Risk sharing is impossible. If the underlying shocks are uncorrelated, the exchange rate variance is the sum of the variances of Earth and Mars marginal utility growth, and we measure a zero risk sharing index despite perfect capital markets.

At the other extreme, if there is costless *trade* between the two planets (teletransportation), and the real exchange rate is therefore constant, marginal utilities can move in lockstep. With constant exchange rates, we measure a perfect risk sharing index of one.

Actual economies produce a mix of tradeable and non-tradeable goods, and transport costs vary by good and country-pair. Thus, actual economies lie somewhere between the two extremes. If there is a positive shock in one country, asset holdings by the other countries should lead to an outflow of goods. But it is costly to ship goods, and those costs increase with the volume being shipped. Therefore, real exchange rates move and their fluctuations blunt risk sharing. Our index lies below one.

Thus, our index answers the question: How much do exchange rate movements blunt the risk sharing opportunities offered by capital markets? Since exchange rates vary so much less than marginal utility growth (as revealed by asset markets), our answer is "not much".

1.2. Incomplete asset markets

Incomplete asset markets are the second reason that risk sharing may be imperfect. If no state-contingent payments are promised, marginal utilities can diverge across countries, even without transport costs and with constant exchange rates.

With incomplete asset markets, the discount factors m_{t+1}^d and m_{t+1}^f that we can recover from asset market data are not unique. Given a discount factor m_{t+1}, that generates the prices p_t of payoffs x_{t+1} by $p_t = E_t(m_{t+1}x_{t+1})$, any discount factor of the form $m_{t+1} + \varepsilon_{t+1}$ also prices the payoffs, where ε_{t+1} is any random variable with $E_t(\varepsilon_{t+1}x_{t+1}) = 0$. Therefore, Eq. (1) does not hold for arbitrary pairs of domestic and foreign discount factors. If Eq. (1) holds for one choice of m^d and m^f, then it will not hold for another, e.g., $m^d + \varepsilon$. In particular Eq. (1) need not hold between domestic and foreign marginal utility growth in an incomplete market. Of course, there always *exist* pairs of discount factors for which (1) holds; Eq. (1) can be used to construct a valid discount factor m^f from a discount factor m^d and the exchange rate.

To think about incomplete markets and multiple discount factors, start by considering a single economy with many agents. In an incomplete market, individuals' marginal utility growths may not be equal. However, the *projection* of each individual's marginal utility growth on the set of available asset payoffs should still be the same.[4] In this sense, individuals should use the available assets to share risks as much as possible. For example, we should not see one consumer heavily invested in tech stocks and another in blue chip stocks, so that one is more affected than the other when tech stocks fall (unless, of course,

---

[4]The projection is the fitted value of a regression of marginal utility growth on the asset payoffs, and the "mimicking portfolio" for marginal utility growth. Each individual's first order conditions state that his marginal utility growth m^i satisfies $p_t = E_t(m_{t+1}^i x_{t+1})$, and thus (since residuals are orthogonal to fitted values) $p_t = E_t[\text{proj}(m_{t+1}^i|X)x_{t+1}]$ where X denotes the payoff space of all traded assets. There is a unique discount factor $x^* \in X$ such that $p_t = E_t(x_{t+1}^* x_{t+1})$ for all $x \in X$. Since this discount factor is unique, $\text{proj}(m_{t+1}^i|X) = x^*$ is the same for all individuals. See Cochrane (2004, p. 68) for elaboration.

the consumers differ significantly in preferences or exposures to non-marketed risks, i.e. if the first consumer works in the old economy and the second in tech). We could imagine running regressions of individual marginal utility growth on asset market returns. The correlation between the fitted values of those regressions would tell us how well individuals use *existing* markets to share risk.

With this example in mind, we evaluate Eqs. (1) and (2) using the unique discount factors m_{t+1}^d and m_{t+1}^f that are in the space of domestic and foreign asset payoffs (evaluated in units of domestic and foreign goods, respectively). These discount factors are the projections of any possible domestic and foreign discount factors onto the relevant spaces of asset payoffs, and they are also the minimum-variance discount factors. We show that Eq. (1) continues to hold with this particular choice of discount factors. If we start with $m_{t+1}^d \in X$, form $m_{t+1}^f = m_{t+1}^d \times e_{t+1}/e_t$ to satisfy (1), we can quickly see that m_{t+1}^f is in the payoff space available to the foreign investor.

In an incomplete markets context, then, our international risk sharing index answers the questions: How well do countries share risks spanned by *existing* asset markets? How much do exchange rate changes blunt risk sharing using *existing* asset markets?

Marginal utility growth is equal to the minimum-variance discount factors m_{t+1}^d and m_{t+1}^f that we examine, plus additional risks that are orthogonal to, and hence uninsurable by, asset markets. Overall risk sharing—the risk sharing index computed using true marginal utility growth—can be better or worse than our asset-based measure. It can be better if the additional risks, not spanned by asset markets, generate a component of marginal utility growth that is positively correlated across countries. It can be worse if the unspanned risks generate a component of marginal utility growth that is negatively correlated across countries. However, to have a *quantitatively* important effect on the risk sharing measure, additional non-spanned risks would have to be a good deal larger than the already puzzling 50% volatility of marginal utility growth implied by the equity premium. Ruling out such huge numbers, our asset-based measure implies that overall risks are also well shared. We make detailed calculations to address this point below.

2. Calculation

2.1. Discount factors and the risk sharing index

We adopt a continuous time formulation, which allows us to translate more easily between logs and levels. We first describe how to recover the minimum-variance discount factor from asset markets in general, and then we specialize the discussion to our international setting.

2.1.1. Discount factors and asset returns
Suppose that a vector of assets has the following instantaneous excess return process:

$$\mathrm{d}R = \mu \mathrm{d}t + \mathrm{d}z \quad \text{with} \quad \mathrm{E}(\mathrm{d}z\,\mathrm{d}z') = \Sigma\,\mathrm{d}t. \tag{4}$$

(An excess return is the difference between any two value processes, e.g., $\mathrm{d}R = \mathrm{d}S/S - \mathrm{d}V/V$. We often use a risk-free asset $\mathrm{d}V/V = r\,\mathrm{d}t$, but this is not necessary.) We will keep the expressions simple by specifying that a risk-free asset

is traded,

$$\frac{\mathrm{d}B}{B} = r\,\mathrm{d}t. \tag{5}$$

With this notation, the discount factor

$$\frac{\mathrm{d}\Lambda}{\Lambda} = -r\,\mathrm{d}t - \mu'\Sigma^{-1}\,\mathrm{d}z \tag{6}$$

prices the assets, i.e., it satisfies the basic pricing conditions

$$r\,\mathrm{d}t = -\mathrm{E}\left(\frac{\mathrm{d}\Lambda}{\Lambda}\right) \quad \text{and} \quad \mu\,\mathrm{d}t = -\mathrm{E}\left(\frac{\mathrm{d}\Lambda}{\Lambda}\,\mathrm{d}R\right). \tag{7}$$

In this formulation, $\mathrm{d}\Lambda/\Lambda$ plays the role of m in the introduction.

We can construct the log discount factor[5] required in Eqs. (1) and (2) via Ito's lemma:

$$\mathrm{d}\ln\Lambda = \frac{\mathrm{d}\Lambda}{\Lambda} - \frac{1}{2}\frac{\mathrm{d}\Lambda^2}{\Lambda^2} = -\left(r + \frac{1}{2}\mu'\Sigma^{-1}\mu\right)\mathrm{d}t - \mu'\Sigma^{-1}\,\mathrm{d}z. \tag{8}$$

We can then evaluate the variance of the discount factor as

$$\frac{1}{\mathrm{d}t}\sigma^2(\mathrm{d}\ln\Lambda) = \mu'\Sigma^{-1}\mu. \tag{9}$$

Eqs. (6) and (9) show why *mean* asset returns determine the *volatility* of marginal utility growth. Holding constant Σ, the higher μ is, the more $\mathrm{d}\Lambda/\Lambda$ loads on the shocks $\mathrm{d}z$ in Eq. (6), and the more volatile is $\mathrm{d}\Lambda/\Lambda$ in Eq. (9).

2.1.2. International discount factors and the risk sharing index

We now specialize these formulas to our international context. We write the real returns on the domestic risk-free asset B^d, domestic stocks S^d, exchange rate e (in units of foreign goods/domestic goods), foreign risk-free asset B^f, and foreign stocks S^f as

$$\frac{\mathrm{d}B^d}{B^d} = r^d\,\mathrm{d}t, \quad \frac{\mathrm{d}S^d}{S^d} = \theta^d\,\mathrm{d}t + \mathrm{d}z^d,$$

$$\frac{\mathrm{d}e}{e} = \theta^e\,\mathrm{d}t + \mathrm{d}z^e,$$

---

[5]The discrete-time versions of these calculations condition down. Starting with $1 = \mathrm{E}_t(m_{t+1}R_{t+1})$ we can condition down to $1 = \mathrm{E}(mR)$ and hence $\mathrm{E}(R) = 1/\mathrm{E}(m) - \mathrm{cov}(m, R)$ where E and cov are unconditional moments. We can represent these unconditional moments analogously to (6) with the discount factor $m_{t+1} = 1/R^f - (\mathrm{E}(R) - R^f)'\mathrm{cov}(R, R')(R_{t+1} - E(R))/R^f$ where $R^f \equiv 1/\mathrm{E}(m)$. These constructions go through even if there is arbitrary variation in the conditional moments of returns. This conditioning down property does not go through in the continuous time formulation. While $0 = \mathrm{E}_t[d(\Lambda P)]$ conditions down to $0 = \mathrm{E}[d(\Lambda P)]$, we cannot apply Ito's lemma to the latter expression. Thus, the covariance on the right-hand side of Eq. (7) must be the conditional covariance. To estimate using sample averages as we do, we must assume constant conditional covariances. With constant conditional covariances, the means in (7) do condition down, so our calculation based on unconditional moments is still correct if there is variation in conditional means. One can derive the same formulas for log discount factors from the discrete time formulas that do condition down, but approximations are required to translate from levels to logs. The approximations are essentially that conditional variances do not change too much.

$$\frac{\mathrm{d}B^f}{B^f} = r^f \, \mathrm{d}t, \quad \frac{\mathrm{d}S^f}{S^f} = \theta^f \, \mathrm{d}t + \mathrm{d}z^f. \tag{10}$$

We collect the shocks in a vector $\mathrm{d}z$ and write their covariance matrix as

$$\mathrm{d}z = \begin{bmatrix} \mathrm{d}z^d \\ \mathrm{d}z^e \\ \mathrm{d}z^f \end{bmatrix} \quad \text{and} \quad \Sigma = \frac{1}{\mathrm{d}t}\mathrm{E}(\mathrm{d}z\,\mathrm{d}z') = \begin{bmatrix} \Sigma^{dd'} & \Sigma^{de} & \Sigma^{df'} \\ \Sigma^{ed'} & \Sigma^{ee} & \Sigma^{ef'} \\ \Sigma^{fd'} & \Sigma^{fe} & \Sigma^{ff'} \end{bmatrix}. \tag{11}$$

(We use the notation $\Sigma^{dd'}$ etc. because there may be several risky assets in each economy.)

We show in the Appendix that the vector of expected excess returns on domestic stock, exchange rate, and foreign stock are, for the domestic investor,

$$\mu^d = \begin{bmatrix} \theta^d - r^d \\ \theta^e + r^f - r^d \\ \theta^f - r^f + \Sigma^{ef} \end{bmatrix} \tag{12}$$

and for the foreign investor,

$$\mu^f = \begin{bmatrix} \theta^d - r^d - \Sigma^{ed} \\ \theta^e + r^f - r^d - \Sigma^{ee} \\ \theta^f - r^f \end{bmatrix}. \tag{13}$$

The covariance matrix of these returns is Σ for both investors.

The discount factors are, as in Eq. (6),

$$\frac{\mathrm{d}\Lambda^i}{\Lambda^i} = -r^i \, \mathrm{d}t - \mu^{i\prime}\Sigma^{-1}\,\mathrm{d}z, \quad i = d,f. \tag{14}$$

By Ito's lemma we then have

$$d \ln \Lambda^i = -\left(r^i + \tfrac{1}{2}\mu^{i\prime}\Sigma^{-1}\mu^i\right)\mathrm{d}t - \mu^{i\prime}\Sigma^{-1}\,\mathrm{d}z, \quad i = d,f \tag{15}$$

and

$$\frac{1}{\mathrm{d}t}\sigma^2(\mathrm{d}\ln\Lambda^i) = \mu^{i\prime}\Sigma^{-1}\mu^i, \quad i = d,f. \tag{16}$$

Our risk sharing index from Eq. (2) is therefore

$$1 - \frac{\sigma^2(\mathrm{d}\ln\Lambda^d - \mathrm{d}\ln\Lambda^f)}{\sigma^2(\mathrm{d}\ln\Lambda^d) + \sigma^2(\mathrm{d}\ln\Lambda^f)} = 1 - \frac{\Sigma^{ee}}{\mu^{d\prime}\Sigma^{-1}\mu^d + \mu^{f\prime}\Sigma^{-1}\mu^f}. \tag{17}$$

As one might expect, the domestic and foreign discount factors load equally on the domestic and foreign stock return shocks, and their loadings on the exchange rate shock differ by exactly one. To see this result, note from the definitions (12) and (13) that we can write

$$\mu^f = \mu^d - \begin{bmatrix} \Sigma^{ed} \\ \Sigma^{ee} \\ \Sigma^{ef} \end{bmatrix} = \mu^d - \Sigma^e, \tag{18}$$

where Σ^e denotes the middle column of Σ. Thus, we can relate the domestic and foreign minimum-variance discount factors by

$$
\begin{aligned}
\frac{\mathrm{d}\Lambda^f}{\Lambda^f} &= \frac{\mathrm{d}\Lambda^d}{\Lambda^d} + (r^d - r^f)\,\mathrm{d}t + \Sigma^{e\prime}\Sigma^{-1}\,\mathrm{d}z \\
&= \frac{\mathrm{d}\Lambda^d}{\Lambda^d} + (r^d - r^f)\,\mathrm{d}t + \mathrm{d}z^e.
\end{aligned}
\tag{19}
$$

2.2. Incomplete markets

When markets are incomplete, the discount factors recovered from asset markets as above are not the only possible ones. If we add noise to the discount factor that is mean zero (to price the interest rate) and uncorrelated with the asset return shocks, the pricing implications are not affected. So long as $E_t(\mathrm{d}w^i) = 0$ and $E_t(\mathrm{d}w^i\,\mathrm{d}z) = 0$,

$$
\frac{\mathrm{d}\Lambda^{i*}}{\Lambda^{i*}} = \frac{\mathrm{d}\Lambda^i}{\Lambda^i} + \mathrm{d}w^i
\tag{20}
$$

satisfies the pricing conditions (7) and is hence a valid discount factor.

Since $\mathrm{d}w^i$ is orthogonal to $\mathrm{d}z$ and $\mathrm{d}\Lambda^i/\Lambda^i$ is driven only by $\mathrm{d}z$, we have three properties of the discount factors $\mathrm{d}\Lambda^i/\Lambda^i$ that we recover from asset markets by Eq. (14): (i) they are the minimum variance discount factors; (ii) they are the unique domestic and foreign discount factors formed as a linear combination of shocks to assets $\mathrm{d}z$; (iii) they are the discount factors formed by the projection of any valid discount factor $\mathrm{d}\Lambda^{i*}/\Lambda^{i*}$ on the space of asset returns.

In an incomplete market, there *exist* pairs of discount factors that satisfy the relation

$$
\mathrm{d}\ln e = \mathrm{d}\ln\Lambda^f - \mathrm{d}\ln\Lambda^d,
\tag{21}
$$

but *not all* pairs of discount factors do so. Existence follows by construction. A domestic discount factor satisfies $0 = E[\mathrm{d}(\Lambda^d S)]$, for any asset S, and a foreign discount factor satisfies $0 = E[\mathrm{d}(\Lambda^f (S/e))]$. Given any domestic discount factor Λ^d that satisfies the first relation, the choice $\Lambda^f = ke\Lambda^d$ (with any constant k) obviously satisfies the second relation and the identity (21). But another foreign discount factor formed by adding $\mathrm{d}w^f$ no longer satisfies the identity (21).

The basic identity (21) does hold for the minimum-variance discount factors. To see this result, start with the minimum-variance discount factor Λ^d, and form a candidate $\Lambda^f = ke\Lambda^d$. As above, this is a valid foreign discount factor. To check that this is the minimum-variance foreign discount factor, we can check any of the three properties above. It is easiest to check the second. $\mathrm{d}\ln\Lambda^d$ is driven only by the shocks $\mathrm{d}z$, and $\mathrm{d}z^e$ is one of the shocks $\mathrm{d}z$, so $\mathrm{d}\ln\Lambda^f$ and hence $\mathrm{d}\Lambda^f$ is also driven only by the shocks $\mathrm{d}z$. (The Appendix also includes a more explicit but algebraically more intensive proof.)

3. Results

3.1. Data and summary statistics

We implement the continuous time formulas in Section 2 with straightforward discrete time approximations and monthly data. We start with domestic excess stock returns $R^d_{t+\Delta}$,

foreign excess stock returns $R^f_{t+\Delta}$, the exchange rate e_t (in units of foreign currency per dollar), and domestic and foreign interest rates $r^d_{t+\Delta}$ and $r^f_{t+\Delta}$, where $\Delta = \frac{1}{12}$ years. Denoting sample mean by E_T (where T is the sample size), we estimate the instantaneous risk premia and variances required in the risk sharing measure (17) by the obvious sample counterparts to the continuous time moments:

$$\theta^d - r^d = \frac{1}{\Delta}E_T R^d_{t+\Delta}, \quad \theta^f - r^f = \frac{1}{\Delta}E_T R^f_{t+\Delta},$$

$$\theta^e + r^f - r^d = \frac{1}{\Delta}E_T\left(\frac{e_{t+\Delta} - e_t}{e_t} + r^f_{t+\Delta} - r^d_{t+\Delta}\right),$$

$$dz^d = R^d_{t+\Delta} - E_T R^d_{t+\Delta}, \quad dz^f = R^f - E_T R^f_{t+\Delta},$$

$$dz^e = \left(\frac{e_{t+\Delta} - e_t}{e_t}\right) - E_T\left(\frac{e_{t+\Delta} - e_t}{e_t}\right),$$

$$\Sigma = \frac{1}{\Delta}E_T(dz\,dz'). \tag{22}$$

We use real returns, real interest rates, and a real exchange rate, each adjusted ex post by realized inflation. Since the risk sharing index is based entirely on excess returns, the calculation is fairly insensitive to how one handles interest rates and inflation. However, excess nominal returns are not quite the same as excess real returns, so we start with real returns to keep the calculation as pure as possible.

We use the US as the domestic country, and the UK, Germany, and Japan as the foreign countries. Table 1 presents summary statistics important to our calculations. The top panel shows the mean and standard deviation of the excess returns on the domestic and foreign

Table 1
Summary statistics

	US Stock	UK		Germany		Japan	
		Stock	X-rate	Stock	X-rate	Stock	X-rate
Returns (%)							
Mean	9.21	10.31	0.16	7.21	1.75	4.78	−0.84
	(2.81)	(1.69)	(2.69)	(1.59)	(3.01)	(1.46)	(2.85)
Std Dev.	14.70	18.12	11.51	16.65	11.56	18.74	12.89
	(0.09)	(0.17)	(0.03)	(0.12)	(0.03)	(0.17)	(0.04)
Return correlations							
US stock		0.57	0.01	0.45	0.01	0.34	0.03
Foreign stock			0.04		0.13		−0.08

This table shows annualized summary statistics for real excess returns on stock indices and exchange rates for the US, UK, Germany, and Japan. The stock indices are total market returns from Datastream, the interest rates are for one-month Eurocurrency deposits from Datastream, and the CPI is from the International Monetary Fund's IFS database. The stock returns (Stock) are excess returns over the same country's one-month interest rates. The exchange rate returns (X-rate) are excess returns for borrowing in dollars, converting to the foreign currency, lending at the foreign interest rate, and converting the proceeds back to dollars. Monthly data from January 1975 through June 1998. Serial-correlation adjusted standard errors in parenthesis.

stock indices and the exchange rate (annualized and reported in percent). The bottom panel shows correlations between the returns. Autocorrelation-corrected standard errors are in parentheses.[6]

The table reminds us of the high equity risk premium, not only in the US but also abroad. The mean excess stock index returns range from 4.78% in Japan to 10.31% in the UK and are all statistically significant. The standard deviation ranges from 14.70% in the US to 18.74% in Japan, resulting in annualized Sharpe ratios of 0.63 in the US, 0.57 in the UK, 0.43 in Germany, and 0.26 in Japan. The estimates of the exchange rate risk premium $\theta^e + r^f - r^e$ are positive for the UK and Germany and negative for Japan, but are all small and statistically indistinguishable from zero. The volatility of the exchange rates is about $\frac{2}{3}$ that of stocks, ranging from 11.51% to 12.89%. Suggestively, this is about the same as the volatility of long-term nominal bond returns. The stock returns are positively correlated across countries, with correlations ranging from 0.34 to 0.57. The exchange rates are poorly correlated with stock returns both in the US and abroad, with correlations less than 0.15.

Table 1 also reminds us of how difficult it is to estimate equity risk premia. For example, the US risk premium of 9.21% is measured with a 2.81% standard error. The reason is, of course, the high volatility of stock returns relative to the size of the mean return. The standard error of the mean return $\sigma(R_{t+\Delta})/\sqrt{T}$, and its more precise incarnation in our GMM standard errors, dooms precise measurement of the risk premium. This is the central source of uncertainty in the risk sharing index.

3.2. Risk sharing index

Table 2 presents our central result, estimates of the risk sharing index (17). The numbers are all at or above 0.98, implying a very high level of international risk sharing.

The high risk sharing indices are driven by the relatively *low* volatilities of the exchange rates compared to the volatilities of the discount factors. To see this fact, Table 2 calculates the volatilities of the minimum-variance domestic and foreign discount factors. We see that discount factor volatilities are between 0.63 and 0.69, much higher than the 0.11–0.12 exchange rate volatilities of Table 1. The discount factor volatilities are very similar for investors in each pair of countries, despite the quite different equity premia in each country. Remember that the discount factor volatility has the interpretation of the maximum Sharpe ratio obtainable by trading in all of the international assets, not just the assets in each investor's home country. Thus, all investors essentially face the same set of assets, distorted only by exchange rate volatility. For the observed exchange rate volatilities, this distortion is small. If exchange rates were more volatile, the discount factor volatilities would differ more across countries.

The standard errors on the discount factor volatilities are quite high, ranging from 0.13 to 0.21. This is because the *variance* of the discount factor depends on *mean* excess returns μ^d and μ^f (see Eq. (16)), which, as we saw in Table 1, are difficult to estimate precisely. However, the standard errors on the risk sharing index are very small. The risk sharing index of 0.98 is measured with a standard error of about 0.01. The index is so close to one

---

[6]To obtain standard errors, we stack Eqs. (22) into a vector of moment conditions and treat that vector in the framework of GMM estimation (Hansen, 1982). We use Newey and West's (1987) estimator of the moment covariance matrix with a six-month correction for serial correlation.

Table 2
Risk sharing index

	US vs. UK	US vs. Germany	US vs. Japan
Risk sharing index			
	0.986	0.985	0.980
	(0.005)	(0.007)	(0.016)
Standard Deviation of marginal utility growth			
Domestic	0.69	0.67	0.64
	(0.18)	(0.18)	(0.21)
Foreign	0.69	0.66	0.67
	(0.13)	(0.19)	(0.16)

This table shows the risk sharing index and annualized standard deviations of the discount factors recovered from asset markets. The domestic country is the US and the foreign country is the UK, Germany, or Japan. The investable assets are the domestic interest rate, the domestic stock market, the foreign interest rate, and the foreign stock market. The risk sharing index is defined in Eq. (17),

$$1 - \frac{\Sigma^{ee}}{\mu^{d'}\Sigma^{-1}\mu^d + \mu^{f'}\Sigma^{-1}\mu^f}.$$

Discount factor volatility is calculated by $\mu^{i'}\Sigma^{-1}\mu^i$. Serial-correlation adjusted standard errors in parenthesis.

Table 3
Discount factor loadings

	US vs. UK		US vs. Germany		US vs. Japan	
	Domestic	Foreign	Domestic	Foreign	Domestic	Foreign
dz^d	3.13	3.13	3.77	3.77	4.22	4.22
dz^e	0.31	−0.69	1.05	0.05	−0.76	−1.76
dz^f	1.73	1.73	1.13	1.13	0.14	0.14

This table shows the loading of the discount factors $\mu^{d'}\Sigma^{-1}$ and $\mu^{f'}\Sigma^{-1}$ on the domestic stock return shocks dz^d, the exchange rate shocks dz^e and the foreign stock return shocks dz^f. The risk premium vectors μ^d and μ^f are given in Eqs. (12) and (13).

that even substantial uncertainty about the discount factor volatility does not much affect its value. Even if the variance of discount factors were 20% lower, the risk sharing index would still be 0.96.

3.3. Discount factor loadings and plots

To get a better sense of the discount factors recovered from asset markets, we show in Table 3 the loadings $\mu^{d'}\Sigma^{-1}$ and $\mu^{f'}\Sigma^{-1}$ of the domestic and foreign discount factors on the excess return shocks dz. As Eq. (19) shows, the discount factor loadings on the stock return shocks are exactly the same, and the foreign discount factor loads by minus one more on the exchange rate shock, so that the difference between the two discount factors is exactly equal to the exchange rate. Also sensibly, the loadings on the stock market shocks are all positive, which, together with the minus sign in Eq. (14), means that marginal utility

declines when either stock market rises. Finally, since the weights are larger on the stocks than on the exchange rate, the volatility in marginal utility growth comes mostly from stock market shocks.

Fig. 1 plots the discount factors for the three country pairs (i.e., US vs. UK, US vs. Germany, and US vs. Japan) to obtain a visual idea of the discount factors' correlation and relative magnitudes. (We use the average ex post real one-month interest rate to proxy for the real risk-free rate in each country. These interest rates are required to display the discount factors in Eq. (14), but they do not appear in discount factor variances or the risk sharing measure.) On the left, we plot the logarithms of US and foreign discount factors $\log \Lambda_t^i$. On the right, we show a scatter plot of the US versus foreign discount factor growth $d\Lambda^i/\Lambda^i$. As the pictures show, the US and foreign discount factor growth are nearly the same. They do differ, by the exchange rate, but discount factors are so volatile relative to the exchange rate that they do not differ by much. Therefore, domestic and foreign discount factor growth are very close. This is our basic point.
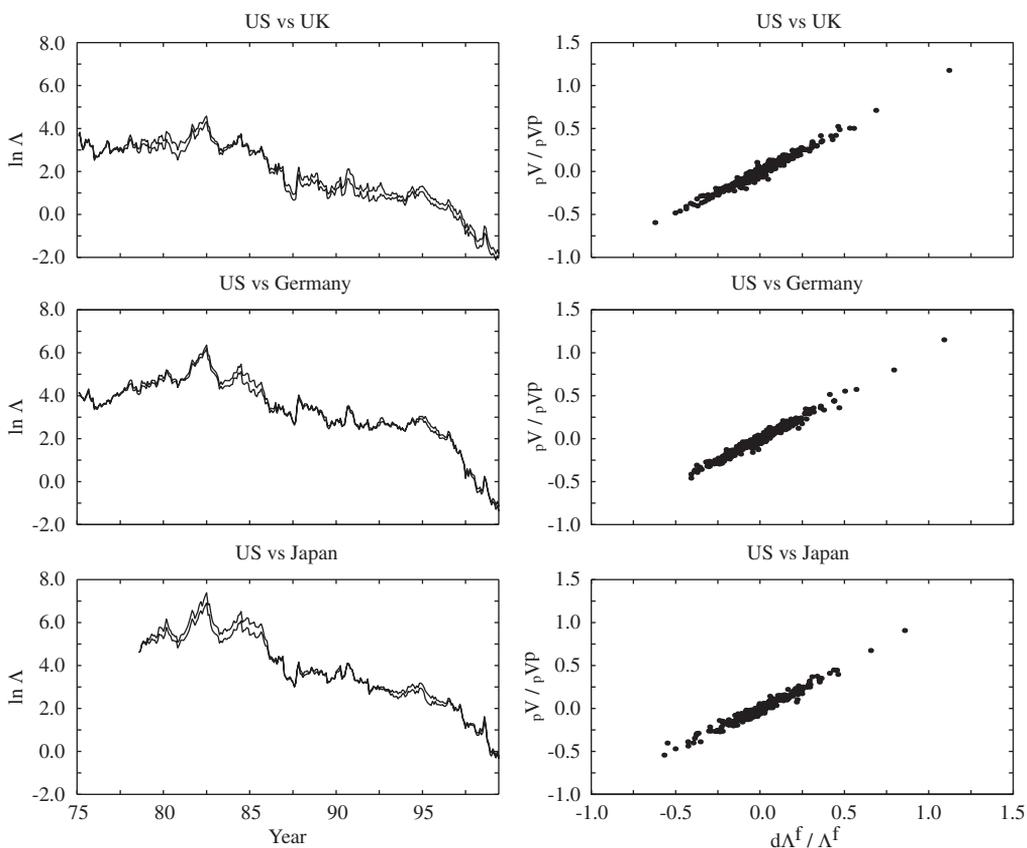
Fig. 1. Discount factors. This figure presents discount factors implied by asset markets. The left-hand column plots log levels of the discount factors $\ln \Lambda_t$ over time. The right-hand column is a scatter plot of growth in the domestic (US) discount factor vs. growth in the foreign discount factors, $d\Lambda/\Lambda$. The investable assets are the domestic risk-free rate and stock market as well as the foreign risk-free rate and stock market. The discount factors are given in (14), $d\Lambda^i/\Lambda^i = -r^i dt - \mu^{i\prime}\Sigma^{-1} dz$, $i = d, f$.

Table 4
Risk sharing index from consumption data

	US vs. UK	US vs. Germany	US vs. Japan
Standard Deviation of log consumption growth (%)			
Quarterly			
Domestic	1.44	1.44	1.44
Foreign	3.11	1.66	2.65
Annual			
Domestic	1.80	1.80	1.80
Foreign	3.14	2.09	1.58
Log consumption growth correlations			
Quarterly	0.31	0.17	0.27
Annual	0.42	0.24	0.35
Risk sharing index			
Quarterly	0.245	0.166	0.265
Annual	0.361	0.233	0.350

This table shows annualized summary statistics for domestic and foreign log consumption growth and the corresponding risk sharing index. Consumption is real per-capita consumption of non-durables and services from the International Monetary Fund's IFS database. Quarterly and annual data from Q1 1975 through Q2 1998. Quarterly standard deviations are annualized. The domestic country is the US and the foreign country is the UK, Germany, or Japan. The risk sharing index is calculated according to Eq. (23),

$$1 - \frac{\sigma^2(\mathrm{d} \ln c^d - \mathrm{d} \ln c^f)}{\sigma^2(\mathrm{d} \ln c^d) + \sigma^2(\mathrm{d} \ln c^f)}.$$

3.4. Consumption data

Consumption growth is poorly correlated across countries, so studies based on consumption data typically find that international risks are poorly shared. This finding is a major puzzle of international economics.[7] Consumption growth is also not very volatile, so when it is multiplied by reasonable risk aversion to produce marginal utility growth, consumption-based models do not produce the observed volatility of the exchange rate via Eq. (1). By this standard, exchange rates seem too volatile.[8] The low volatility of consumption growth, and hence marginal utility growth, is also the heart of the equity premium puzzle.

To examine quantitatively the discrepancy between the asset market view of risk sharing and the view based on consumption data, we calculate marginal utility growth and risk sharing measures implied by consumption data and standard power utility. Table 4 presents the results.

Table 4 starts with annualized standard deviations and correlations of log consumption growth for the three country pairs. Since high-frequency consumption data is notoriously noisy, we consider both quarterly and annual data. (Temporally aggregating consumption

---

[7]Important contributions include Backus et al. (1992), Backus and Smith (1993), Brennan and Solnik (1989), Davis et al. (2001), Lewis (1996, 2000), Obstfeld (1992, 1994) and Tesar (1993).

[8]See Flood and Rose (1995), Mark (1995), Meese and Rogoff (1983), Obstfeld and Rogoff (2000), and Rogoff (1999).

data is likely to reduce the effect of measurement errors. See Bell and Wilcox, 1993, and Wilcox, 1992.) The table confirms the usual results. Consumption growth is relatively smooth, with standard deviations of 1.4% to 3.1%. Using CRRA utility with $u'(c) = c^{-\gamma}$, the variance of log marginal utility growth is $\sigma^2(\mathrm{d}\ln\Lambda) = \sigma^2(\gamma\,\mathrm{d}\ln c)$. As is well known from the equity risk premium literature, we therefore need a very large γ, at least 20, to generate the 60% volatility of marginal utility growth implied by asset markets from such smooth consumption growth.

Table 4 also shows that consumption growth is imperfectly correlated across countries, with correlations ranging from 0.17 to 0.42. The correlations are highest between the US and UK and lowest between the US and Germany. Although nicely above zero, these are well below one, and surprisingly lower than the correlations of output growth across countries (Backus et al., 1992). This is the heart of the usual observation that international risks do not seem to be well shared.

Assuming for simplicity that the domestic and foreign representative investors have the same level of relative risk aversion, we can compute the risk sharing index based on consumption data as

$$1 - \frac{\sigma^2(\mathrm{d}\ln\Lambda^d - \mathrm{d}\ln\Lambda^f)}{\sigma^2(\mathrm{d}\ln\Lambda^d) + \sigma^2(\mathrm{d}\ln\Lambda^f)} = 1 - \frac{\sigma^2(\mathrm{d}\ln c^d - \mathrm{d}\ln c^f)}{\sigma^2(\mathrm{d}\ln c^d) + \sigma^2(\mathrm{d}\ln c^f)}. \tag{23}$$

(In incomplete markets, marginal utility growth does not necessarily obey Eq. (21), so we cannot use exchange rates in the numerator of the risk sharing index. The risk aversion coefficients conveniently drop out of the index when both countries have the same level of risk aversion.) As shown in Table 4, this calculation results in risk sharing indices of 0.17–0.36. These numbers are much lower than the risk sharing indices implied by asset markets in Table 2, and they capture in our index the usual conclusion that risks are poorly shared internationally. (The index is lower than the consumption growth correlation because the index penalizes size as well as correlation. The greater standard deviation of foreign consumption growth would count against risk sharing even if consumption growths were perfectly correlated.)

These consumption-based results do not contradict our calculations, in the sense of pointing out an error. We measure marginal utility growth directly from asset markets, ignoring quantity data. The consumption-based calculations measure marginal utility growth from quantity data via a utility function, ignoring asset prices. Since the results are so different, the quantity-based and price-based calculations need to be reconciled of course, and we discuss this reconciliation below.

Consumption-based calculations also measure *overall* risk sharing, while we only measure that component of risk sharing achieved through incomplete asset markets. If there are substantial additional risks, one could see good risk sharing in asset markets despite poor risk sharing in total consumption. As we quantify in the next section, though, this view falls apart on the required magnitude of the additional risks.

3.5. Non-market risks and bounds on overall risk sharing

What does the asset market-based risk sharing index tell us about overall risk sharing—the risk sharing index computed using marginal utility growth—in incomplete markets?

Table 5
Risk sharing index with incomplete markets

$\rho(\mathrm{d}w^d, \mathrm{d}w^f)$	$\sigma(\mathrm{d}w^d) = \sigma(\mathrm{d}w^f)$					
	0%	10%	20%	30%	50%	100%
−1.0	0.975	0.909	0.735	0.507	0.060	−0.557
−0.8	0.975	0.916	0.760	0.554	0.152	−0.402
−0.4	0.975	0.929	0.808	0.649	0.338	−0.091
0.0	0.975	0.943	0.857	0.744	0.523	0.219
0.4	0.975	0.956	0.905	0.839	0.710	0.529
0.8	0.975	0.969	0.954	0.934	0.894	0.839
1.0	0.975	0.976	0.978	0.981	0.987	0.994

This table shows the risk sharing index in the case of incomplete markets for different levels of the volatility and correlation of additional risks. The domestic and foreign mean excess stock returns $\theta^d - r^d$ and $\theta^f - r^f$ are 8%. The volatility of both stock returns is 18%. The volatility of the exchange rate is 12%. The correlation between the stock returns is 0.4 and the stock returns are both uncorrelated with the exchange rate. The domestic foreign exchange risk premium $\theta^e + r^f - r^d$ is zero. The implied standard deviation of the minimum-variance discount factors is 56.7%. The risk sharing index is calculated as in Eq. (27),

$$1 - \frac{\Sigma^{ee} + \sigma^2(\mathrm{d}w^d - \mathrm{d}w^f)}{\mu^{d\prime}\Sigma^{-1}\mu^d + \mu^{f\prime}\Sigma^{-1}\mu^f + \sigma^2(\mathrm{d}w^d) + \sigma^2(\mathrm{d}w^f)}.$$

How much and what kind of non-market risk would it take to drive the asset market-based risk sharing index down to the values suggested by consumption data and power utility?

Overall risk sharing can be better or worse than the asset market-based risk sharing index. Marginal utility growth equals the minimum-variance discount factor, plus the effects of additional non-spanned risks.

The additional risks[9] are orthogonal to asset payoffs, but they may be arbitrarily correlated across countries. If the additional risks are highly correlated across countries, marginal utility growths can move together more closely than the asset market-based discount factors, and the risk sharing index for marginal utility growth can be *higher* than we calculate. If additional risks are uncorrelated across countries, marginal utility growth is less correlated than the asset-based discount factors, and overall risk sharing is lower than our calculation. Finally, if the unspanned risks are somehow negatively correlated across countries, they will drive down the risk sharing index more quickly, and ultimately drive it to negative values. In each of these cases, there must be *more risk* as well. The non-traded risks are orthogonal to asset market-based discount factors, so to be less correlated across countries, marginal utility growth must be even more volatile than the already highly volatile asset-based discount factors.

To make this discussion quantitative and explicit, Table 5 presents some calculations. Consider a pair of candidates for marginal utility growth $\mathrm{d}\tilde{\Lambda}^i/\tilde{\Lambda}^i$ formed from the minimum-variance discount factors $\mathrm{d}\Lambda^i/\Lambda^i$ by the addition of noises $\mathrm{d}w^i$ with

---

[9]Precisely, we mean here and below "the effects on marginal utility growth of the additional risks". If utility is power and the additional risks take the form of endowment shocks, the consumption risk translates directly to marginal utility growth risk. Under other circumstances, additional risks affect marginal utility growth in more complex ways.

$E(dw^i dz) = 0$. The risk sharing index for these marginal utility growths is

$$1 - \frac{\sigma^2(d \ln \tilde{\Lambda}^d - d \ln \tilde{\Lambda}^f)}{\sigma^2(d \ln \tilde{\Lambda}^d) + \sigma^2(d \ln \tilde{\Lambda}^f)}. \tag{24}$$

To calculate this index and to relate this expression to the risk sharing index using the minimum-variance discount factors, note that

$$\sigma^2(d \ln \tilde{\Lambda}^i) = \sigma^2(d \ln \Lambda^i) + \sigma^2(dw^i), \quad i = d, f, \tag{25}$$

and

$$\sigma^2(d \ln \tilde{\Lambda}^d - d \ln \tilde{\Lambda}^f) = \sigma^2(d \ln \Lambda^d - d \ln \Lambda^f) + \sigma^2(dw^d - dw^f). \tag{26}$$

Substituting the last two expressions into Eq. (24), the risk sharing index using candidates for marginal utility growth is

$$
\begin{aligned}
1 &- \frac{\sigma^2(d \ln \Lambda^d - d \ln \Lambda^f) + \sigma^2(dw^d - dw^f)}{\sigma^2(d \ln \Lambda^d) + \sigma^2(d \ln \Lambda^f) + \sigma^2(dw^d) + \sigma^2(dw^f)} \\
&= 1 - \frac{\Sigma^{ee} + \sigma^2(dw^d - dw^f)}{\mu^{d\prime}\Sigma^{-1}\mu^d + \mu^{f\prime}\Sigma^{-1}\mu^f + \sigma^2(dw^d) + \sigma^2(dw^f)}.
\end{aligned} \tag{27}
$$

There is no point in redoing the calculations for all three countries, so we use a simple set of representative numbers for Table 5 and following experiments. We use a 8% mean stock excess return and 18% standard deviation of stock excess returns in both countries. We set the exchange rate volatility to 12%. We assume a correlation of 0.4 between the stock index returns and zero correlation between the stock index returns and the exchange rate. We set the domestic exchange rate risk premium to zero.[10] In the notation of the formulas, our assumptions for the sensitivity analysis are

$$\theta^e + r^f - r^e = 0, \quad \theta^d - r^d = \theta^f - r^f = 0.08,$$

$$\sigma(dz^d) = \sigma(dz^f) = 0.18, \quad \sigma(dz^e) = 0.12,$$

$$\rho(dz^d, dz^f) = 0.4, \quad \rho(dz^d \, dz^e) = \rho(dz^f \, dz^e) = 0. \tag{28}$$

The risk sharing index for these parameter values is 0.975, typical of Table 2. We then vary the volatility of the additional risks dw^d and dw^f from zero to 100% and their correlation between -1 and 1. Table 5 reports the corresponding values of the risk sharing index, calculated from Eq. (27).

Table 5 and Eq. (27) verify the above intuition about additional risk. The risk sharing index rises if the additional risks are more highly correlated than the asset market-based discount factors, as shown in the $\rho = 1$ row. In this case, the numerator of (27) is unchanged and the denominator increases. If additional risks are uncorrelated, the risk sharing index declines towards zero as the size of the additional risks increases. Here, $\sigma^2(dw^d - dw^f) = \sigma^2(dw^d) + \sigma^2(dw^f)$. The extra terms in the numerator and denominator of (27) are the same, which increases the ratio (since it starts below one) and pulls the risk sharing index towards zero. For intermediate positive correlations, the risk sharing index

---

[10]This assumption gives foreigners a currency risk premium of $-\Sigma^{ee} \approx 0.1^2 = 1\%$. However, the results are insensitive to the assumed exchange rate risk premium.

declines as one adds additional risks, but more slowly than for the $\rho = 0$ case. Finally, if the additional risks are somehow *negatively* correlated, the risk sharing index declines quickly, and eventually becomes negative, reflecting the now negative correlation of marginal utility growth across countries.

Most importantly, Table 5 shows *how much* risk must be added to the discount factor to have a quantitatively important impact on the risk sharing index. We find some positive correlation of additional risks to be the most plausible case. One reason is that we omit a wide variety of assets from the analysis. Additional assets make discount factors more volatile, but typically do so in the same way for domestic and foreign investors. They provide additional risk-sharing opportunities. Non-traded insurance mechanisms such as international remittances, government aid, insurance and reinsurance, etc. have the same effect. Finally, many shocks, such as international business cycles, commodity prices, etc. have common effects across countries, making uninsured marginal utility growths move together. Table 5 shows that if the additional risks have only a $\rho = 0.4$ correlation, similar to that of consumption or GDP growth across countries, then additional risks with as much as 50% volatility leave the risk sharing index above 0.70. Even additional risks with 100% volatility do not bring the risk sharing index down to the maximum 0.35 risk sharing index obtained from consumption growth. Keep in mind that 50% discount factor volatility launched the entire equity premium puzzle literature.

If we assume instead that the extra risks are completely uncorrelated across countries, the $\rho = 0$ line of Table 5 shows that the volatility of extra risks still has to be more than 30% to reduce the risk sharing index to below 0.75, and we need additional risks with volatility on the order of 100% to reach values of 0.2–0.3 suggested by consumption-based calculations in Table 4.

It is hard to think of quantitatively important negatively correlated shocks. Beggar-thy-neighbor policies are a possibility, though changes in such policies pale relative to 10% exchange rate volatility, let alone 50% marginal utility growth volatility. Commodity (oil) price changes might affect certain country pairs, but not US vs. Germany, for example. Still, even with a correlation of negative one, shocks with 30% volatility only bring down the risk sharing index to 0.5.

In sum, additional risks of plausible magnitude and correlation do not have a large impact on the asset market based risk sharing measures. Additional risks have to be larger than the 50% volatility of the minimum-variance discount factor to seriously affect the risk sharing index. It is difficult enough to understand the 50% volatility of the minimum-variance discount factor; do we really believe that there are other country risks, orthogonal to asset market returns, that add up to an *additional* 50–100% volatility of marginal utility growth? Without *huge* additional risks, we can conclude that overall risk sharing, measured from marginal utility growth, is large as well.[11]

For this reason, additional risks also do not offer a reconciliation between the asset market-based risk sharing measure and that derived from consumption data. The consumption-based calculations predict marginal utility growth that is poorly correlated, but they also predict that marginal utility growth has a low volatility, typically below 20% for risk aversion γ less than 10. If we add enough additional risks to our calculation to bring the correlation of

---

[11]If one imposes an a priori upper limit on the volatility of marginal utility growth as in Cochrane and Saá-Requejo (2000), one can obtain lower and upper bounds on overall risk sharing from the asset market-based index to formalize this observation.

marginal utility growth down to the consumption-based levels, we increase the volatility of marginal utility growth from 50% to well over 100%, doubling the equity premium puzzle. This is a sideways move, not one that brings the two calculations closer together.

3.6. A lower equity premium?

Can we avoid the puzzles altogether? Perhaps the ex ante equity premium is a lot lower than sample averages suggest. It is attractive to view the exchange rate from Eq. (1) as a *direct* measure of marginal utility growth, one that gives much lower volatility than the Hansen–Jagannathan (1991) calculation coming from the mean equity premium. If so, this paper adds to the view that a lot of the historical equity premium was luck.

To explore in quantitative detail the effect of uncertainty about the equity premium, we report in Table 6 the risk sharing index for different equity premia. Again, we start with the baseline parameters (28). Then, we try several lower values for the domestic and foreign equity premium. Table 3 shows that even a 5% equity premium, rather than the 9.2% point estimate for the US, still produces a risk sharing index of 0.94. Even a very low 3% equity risk premium in both countries implies a risk sharing index of 0.85. To get to an index below 0.35, the largest value we found from consumption data in Table 4, we need both equity premia to be below 1%!

The basic point follows from Eq. (1). To produce exchange rate volatility of only 10% with uncorrelated discount factors, the discount factors must have a standard deviation no larger than $10\%/\sqrt{2} = 7.1\%$. In turn, a 7.1% standard deviation of marginal utility growth, together with 15% standard deviation of stock returns, implies an equity premium of no more than $0.15 \times 7.1 = 1.1\%$ per year. (And this calculation assumes that marginal utility growth is perfectly correlated with stock returns. Any additional risk lowers the equity premium for given volatility of marginal utility.)

There are good reasons to think that the equity premium may be lower than the sample mean, and that our samples reflect a certain amount of good luck. Over the postwar period, for example, the price–dividend ratio and price–earnings ratio have increased substantially. Careful corrections for this luck, however, only reduce the equity premium to about 6% (Constantinides, 2002). Furthermore, even in our short sample, the US, UK, German and Japanese equity premia are 2.9, 5.5, 3.9, and 2.6 standard errors above 1%,

Table 6
How the risk sharing index depends on equity risk premia

$\theta^d - r^d$	$\theta^f - r^f$				
	0.5%	1%	3%	5%	8%
0.5%	0.133	0.303	0.804	0.922	0.969
1%		0.380	0.800	0.918	0.968
3%			0.846	0.918	0.965
5%				0.939	0.967
8%					0.975

This table calculates the risk sharing index for different values of the mean excess returns of the domestic stock market $\theta^d - r^d$ and foreign stock market $\theta^f - r^f$. The volatility of the exchange rate is 12%. The volatility of both stock indices is 18%. The correlation between the stock returns is 0.4 and the stock returns are both uncorrelated with the exchange rate. The domestic exchange rate risk premium $\theta^e + r^f - r^d$ is zero.

and the significance grows in longer samples. If the true equity premium is below 1%, the last century was extraordinarily lucky.

Table 6 also shows that the risk sharing index is driven by the *larger* of the domestic and foreign equity risk premia. As long as one of the risk premia is high, the risk sharing index is also high. Only when *both* premia drop does the risk sharing index decrease substantially. The risk sharing index is driven by the maximum Sharpe ratio in portfolios of the domestic and international assets. This fact contributes to the small standard errors of our risk sharing index in Table 2, relative to the standard errors of the individual country equity premia. Since stock returns are not perfectly correlated, cross-country data adds to the statistical evidence against very small equity premia.

3.7. ...Or exchange rates are too smooth

Suppose that risks are in fact not well shared internationally. Then, the exchange rates we observe are surprisingly *smooth*. To make this point quantitatively, in Table 7 we solve Eq. (27) for the exchange rate volatility Σ^{ee} that is consistent with a risk sharing index of 0.35, the highest value implied by the consumption data in Table 4. As in Table 5 and Table 6, we start with baseline parameters (28). The result, in the 8% row and 0% column of Table 7, is that exchange rates should have a standard deviation of 102.4%!

As above, perhaps the equity premium is lower than 8%. Table 7 considers a reasonable range of equity risk premia by varying $\theta^d - r^d$ and $\theta^f - r^f$ from 4% to 8%. (Since our index is driven by the maximum risk premium, we set $\theta^d - r^d = \theta^f - r^f$.) The implied exchange

Table 7
Exchange rate volatility implied by a risk sharing index of 0.35

$\mu^d = \mu^f$	$\sigma(\mathrm{d}w^d) = \sigma(\mathrm{d}w^f)$				
	0%	10%	20%	30%	40%
$\rho = 0.4$					
4%	51.2	51.5	52.3	53.6	55.5
6%	76.8	77.0	77.5	78.4	79.7
8%	102.4	102.5	102.9	103.6	104.6
$\rho = 0.0$					
4%	51.2	49.2	42.7	28.6	0.1
6%	76.8	75.5	71.4	70.0	51.9
8%	102.4	101.4	98.4	93.2	85.3

This table shows the exchange rate volatility implied by a risk sharing index of 0.35 in the case of incomplete markets for different levels of the domestic and foreign mean excess stock returns $\theta^d - r^d$ and $\theta^f - r^f$ and different levels of the volatility and correlation of additional risks. The volatility of both stock returns is 18%. The correlation between the stock returns is 0.4 and the stock returns are both uncorrelated with the exchange rate. The domestic foreign exchange risk premium $\theta^e + r^f - r^d$ is zero. The implied standard deviations of the minimum-variance discount factors for the three different equity risk premia are 28.3%, 42.5%, and 56.7%. The implied exchange rate volatility is computed by setting Eq. (27) equal to 0.35, i.e. we solve

$$0.35 = 1 - \frac{\Sigma^{ee} + \sigma^2(\mathrm{d}w^d - \mathrm{d}w^f)}{\mu^{d\prime}\Sigma^{-1}\mu^d + \mu^{f\prime}\Sigma^{-1}\mu^f + \sigma^2(\mathrm{d}w^d) + \sigma^2(\mathrm{d}w^f)}$$

for Σ^{ee}. (Note Σ^{ee} is an element of Σ.)

rate volatility declines to 76.8% and then 51.2%. Cutting the equity premium in half still leaves a prediction that exchange rates should be four times as volatile as they are in fact.

As unspanned risks dw^d and dw^f lower risk sharing for a given exchange rate volatility, they can also lower the predicted exchange rate volatility for a given risk sharing index. Thus, Table 7 also considers additional unspanned risks with volatilities ranging from zero to 40% and a correlation of either zero or 0.4.

When the additional risks are positively correlated, they *raise* the predicted volatility of exchange rates; numbers increase across the columns of Table 7 in the $\rho = 0.4$ panel. Marginal utility growth equals the asset-market discount factor plus the extra risks. If the extra risks are more correlated than is marginal utility growth, then the asset-market discount factors must be even less correlated than marginal utility growth. The ratio of asset-market discount factors generates the exchange rate, and the less correlated the asset-market discount factors, the more volatile the exchange rate. We can also digest this result by looking at Eq. (27). We can solve (27) approximately for exchange rate volatility Σ^{ee} given a risk sharing index $\bar{\rho}$ as

$$\Sigma^{ee} \approx 2[(1 - \bar{\rho})\mu'\Sigma^{-1}\mu + (\rho - \bar{\rho})\sigma^2(dw)]. \tag{29}$$

(This is only an approximate solution since Σ^{ee} also appears in Σ^{-1}, though its effect there is minor.) Extra risks that are more correlated than the presumed risk sharing index $\rho > \bar{\rho}$ increase the predicted volatility of exchange rates.

With uncorrelated extra risks, the addition of extra risks does help to predict smoother exchange rates given a view that overall risk sharing is only 0.35. Numbers decline across columns of Table 7 in the $\rho = 0$ rows. But again, this specification is not quantitatively plausible. Only by *simultaneously* assuming an equity premium of 4% or less (less than half the sample value), additional risks with a huge 40% volatility, and extra risks completely uncorrelated with each other, do we produce exchange rates with volatility less than or equal to sample values of about 12%, as seen in the top right corner of the $\rho = 0$ panel.

4. Interpreting the calculation

4.1. Clarifications and misconceptions

Domestic and foreign discount factors each price the full set of assets: domestic and foreign stocks and bonds. One might derive a "domestic discount factor" that only prices domestic stocks and bonds. Such a discount factor is a different object, and an uninteresting one in our context. In the continuous-time limit, in fact, the domestic and foreign discount factors lie in the same space, diffusions driven by the three shocks dz^d, dz^e, dz^f.

For this reason, the correlation of international stock markets has little bearing on the risk sharing index. The risk sharing index is driven by the mean returns, through the maximum Sharpe ratio. It also is not directly related to the question, how much do mean variance frontiers expand by the inclusion of international stocks?[12] We always allow investors to trade in both domestic and foreign stocks.

Do not confuse the *discount factor* with the *optimal portfolio*. The optimal portfolio, together with a utility function and a driving process for income, supports a marginal utility growth process, and the discount factor is the projection of that marginal utility

---

[12]For example, Chen and Knez (1995) and De Santis and Gerard (1997).

growth process on the space of asset returns. The beginning and the end of this trail are totally different objects. The discount factor is not the optimal portfolio, the "market" portfolio, or the portfolio held by any investor. The minimum-variance discount factor is proportional to the minimum second-moment portfolio, which is on the *bottom* of the mean—variance frontier, and is therefore a portfolio nobody wants to hold. In a mean—variance setting, the optimal, tangency, and market portfolios are all on the *top* of the mean—variance frontier. With more general utility functions and non-marketed income processes, asset portfolios are not even on the mean—variance frontier. Discount factors can be highly correlated, even when portfolios are not correlated.

Although the index has the above interpretation as the ratio of risk not shared to overall risk, so country pairs with lower risk sharing indices have greater scope to reduce overall risk (variance of discount factor) if capital markets can be opened or transport costs lowered, we do not connect the risk sharing index to a measure of welfare. Also, a high risk sharing index does not mean that further risk sharing through the introduction of new assets or institutions, or the lowering of transport costs, is not possible or desirable. The risk sharing index is high in large part because the denominator is high. Improvements in the numerator can have substantial welfare benefits, even when the index is already high.

One might ask, "If two countries fix their exchange rates, will this not lead to a perfect risk sharing index?" The answer is no. Two countries can fix their *nominal* exchange rates, but they cannot fix their *real* exchange rates, except by allowing and enjoying perfect shipment of goods, which of course is the perfect risk-sharing case.

The notion of risk-sharing measured by our index does not respect transport costs. For example, we would measure poor risk sharing between Earth and Mars, or between countries with integrated capital markets, free trade, but a preponderance of locally provided non-tradeable goods. Yet the allocations between countries in these examples may be Pareto optimal *if* the planning problem respects the costs of transport. Our measure does not say anything about this concept of risk sharing. Our measure asks to what extent real exchange rate variability, which is the shadow price of transport costs, blunts risk sharing, defined relative to costless trade in all goods. Whether removing the transport cost is a physical impossibility (haircuts) or a result of government policy (tariffs) is not relevant to our question. Also for this reason we do not complicate the analysis with multiple (tradeable and non-tradeable) goods.

4.2. Home bias

US investors largely hold US securities, Japanese investors largely hold Japanese securities, and so forth. This observation is the "home bias puzzle".[13]

Exchange rate volatility blunts optimal risk sharing, so it implies that optimal portfolios should display some home bias. Once Earth-based investors understand that exchange rate

---

[13]The home bias puzzle goes back a long way, for example Solnik (1974). French and Poterba (1991) observe that Americans held 94% of their equity wealth in the US stock market. The analogous figure for Japan was 98%. Lewis (1999) surveys the home bias puzzle. Baxter and Jermann (1997) argue that human capital is a domestic asset, so adding labor income risk worsens the puzzle. Tesar and Werner (1995) show that foreign portfolios turned over much faster than domestic portfolios, and argue that this suggests that transactions costs do not drive the home bias. "Home bias" puzzles include the propensity of people and managers to hold stocks in companies near them or in their industries. For example, Coval and Moskowitz (1999) find that fund managers hold stocks of companies that are on average somewhat closer to them than the market as a whole.

movements will offset any gains from holding Martian securities, their optimal portfolios will ignore those securities. In mean—variance portfolio calculations, exchange rates are volatile, poorly correlated with stock returns, and offer virtually no risk premium. As a result, foreign stocks are less desirable to domestic investors than similar domestic stocks. However, neither observation has yet generated the amount of home bias that we observe.[14]

In one sense, the home bias calculations complement ours: the home bias literature agrees that exchange rates do not move enough to blunt the risk-sharing possibilities of asset markets. On the other hand, it suggests that risks are not in fact that well shared since people do not hold as many foreign stocks as the portfolio calculations recommend.

Like the consumption-based evidence, the home bias puzzle does not contradict our calculation. The home bias puzzle refers to an optimal portfolio calculation for a specific utility function, and underlying risky income and asset return process or production function and technology shock specification. *Portfolios* do not need to be similar across countries for *marginal utilities* to be similar. (Again, do not confuse the *discount factor* with the *optimal portfolio*.) For example, we can see volatile and highly correlated marginal utility growth despite very little asset cross-holding if the underlying income shocks are highly correlated, if risk sharing is achieved by other means (e.g., government transfers, direct payments to relatives abroad, assets not included in the analysis such as international reinsurance), or if preferences are not the simple one-period quadratic utility over wealth envisioned by mean–variance calculations, or the simple forms used in general-equilibrium models. It is possible to see large amounts of risk sharing in marginal utility growth, but still to be puzzled by the underlying economic model and the portfolio positions that support it. It is an open question whether the home bias puzzle represents a failure of portfolio models to capture the right utility function and the right income process, or whether it represents a systematic and dramatic failure of people to follow their first order conditions for portfolio optimization.

Still, like the consumption-based evidence, one is uncomfortable that two independent ways of calculating risk sharing lead to such different results, so some thoughts on reconciling the calculations are in order.

4.3. Reconciliation of asset-market, consumption, and home-bias calculations

Asset prices give a quite different picture of marginal utility growth from that painted by consumption or portfolio data, using standard utility functions, environments, and parameters. Based on asset prices, it appears that there is a lot of risk ($\sigma(\ln m^d)$ and $\sigma(\ln m^f)$ are large) and that risk is shared well across countries (m^d and m^f move together), since exchange rates are surprisingly smooth ($\sigma(\Delta \ln e)$ is a lot smaller than $\sigma(\ln m)$). Based on consumption data and power utility with low risk aversion, it appears that there is little risk, that risk is poorly shared across countries, and that exchange rates are surprisingly volatile. Based on portfolio calculations with standard utility functions and simple income processes, it appears that people do not share risks through asset markets as much as they could.

---

[14]For example, see Baxter et al. (1995) for a dynamic general equilibrium model with non-tradeable goods, and Black and Litterman (1992) for mean-variance calculations.

As a result, we present our calculations as a contribution to a puzzle, rather than a definitive answer to the risk sharing question. Our contribution is to tie together the asset-market evidence on the equity premium, exchange rate volatility, and risk-sharing aspects of the puzzle, not to reconcile all three phenomena with consumption and portfolio data. We knew that asset markets present more volatile exchange rates than are implied by consumption data and standard utility functions; we knew that asset markets present higher equity premia than are implied by consumption data and standard utility functions. Now we know that asset markets imply substantially higher risk sharing than are implied by consumption data and standard utility functions as well. Furthermore, we see that risk sharing, the equity premium, and exchange rate volatility are all tied together.

Which view of risk sharing will prevail in the end? We cannot be sure until the asset-market and quantity-based views are reconciled. One cannot take the quantity-based measures of risk-sharing at face value while the quantity-based models remain completely at odds with basic facts such as the observed equity premium and exchange rate volatility. While the asset-price measures do not suffer the same internal inconsistency because they are not based on models that make counterfactual predictions, one is justifiably uncomfortable in believing their results with no economic model in hand that can generate the measured pattern of volatility and high correlation of marginal utility growth.

How will the asset-market and quantity-based calculations eventually be resolved? Our asset-market calculations are like the observation that the relative price of apples and oranges is pretty much the same across supermarkets in Los Angeles. We conclude that shoppers' marginal rates of substitution for apples and oranges must be pretty much the same. We do not look inside shopping baskets to make this observation. The quantity-based calculations look into shopping baskets, ignoring prices. Shoppers bring home grocery bags with vastly different apple/orange ratios. Translated through standard utility functions, this quantity data implies that marginal rates of substitution vary a lot across people.

There are two ways to resolve such contrasting calculations: either the specification of the utility function or environment needs modification, or some friction causes marginal rates of substitution not to equal price ratios. In the former category, we can consider different functional forms for preferences, including non-separabilities (habits, durability, non-state-separable preferences), preference heterogeneity (some like oranges better than apples; preference shocks in our one-good intertemporal context), heterogeneity in endowments (people who have apple trees at home buy more oranges; different income or production processes in our context).

The first path is familiar from the equity premium literature. Reconciling the equity premium with consumption data seems to require dramatic departures from the conventional power utility setup.[15] Since these departures generate large discount factor volatility, they will generate volatile exchange rates as well. If they are not to generate *too* volatile exchange rates, they will have to imply good international risk sharing as well, by the inescapable logic of Eq. (1). Thus, a reconciliation along these lines will bring the quantity-based risk sharing measures up to what we measure from asset markets, not the other way around.

---

[15] For example, habits as in Campbell and Cochrane (1999), or countercyclical idiosyncratic non-traded risk as in Constantinides and Duffie (1996).

The second path is also familiar. If it costs me $10 to get to the grocery store, my marginal rate of substitution for apples vs. oranges may deviate substantially from the price ratio. A large literature has tried quantitatively to explain the equity premium, exchange rate volatility, poor risk sharing, and home bias puzzles by explicit frictions such as transactions costs, taxes, legal portfolio constraints, asymmetric information, borrowing constraints, participation constraints and so forth,[16] and the implication that many consumers do not participate in asset markets.[17] Alas, this literature has not yet been entirely successful—which is equivalent to saying the puzzles are still puzzles rather than ex-puzzles. The frictions are just not large enough to keep people far from frictionless first order conditions. The equity premium implies that contingent claims prices (discount factors) vary by 50% or more. Marginal rates of substitution inferred from consumption growth, standard utility functions and low risk aversion vary by 10% or less. That means the wedge between them must vary by 40% or more. Observed transactions costs of a few percent or any reasonable information and adverse selection costs, or taxes just do not vary this much. Furthermore, if some consumers are in fact disconnected from asset markets by transactions costs, the puzzles all remain intact for the investors who do obey their first order conditions. If limited participation becomes the resolution, we will conclude that risks are surprisingly well shared among people who do hold any stocks.

5. Conclusion

We present a calculation of international risk sharing based on asset market data. The high Sharpe ratios obtainable in international asset markets imply that marginal utility growth is very volatile—there is a lot of risk to share. Compared to this volatility, the volatility of exchange rates is very small. We conclude that there is a surprisingly high level of international risk sharing. Alternatively, if risks really are poorly shared, exchange rates are much too smooth. Both conclusions are in stark contrast to the standard findings from consumption data, general equilibrium models, or portfolio calculations.

Since markets are incomplete, our risk sharing index answers the question, "to what extent do exchange rate movements blunt the risk-sharing possibilities of *existing* asset markets?" The answer is "not much". Additional risks uncorrelated with asset movements can lower overall risk sharing, but they simultaneously raise the total amount of risk. Therefore, we argue that reasonably sized and correlated risks cannot lower overall risk sharing by much, and we conclude that the asset-based results indicate a surprisingly large amount of overall risk sharing as well.

The large equity premium drives the large volatility of marginal utility growth, so perhaps a lower equity premium can resolve the puzzle. Like additional risks, this is not a quantitatively plausible resolution however. The equity premium needs to be below 1% in both countries for our risk sharing measure to go down to conventional levels, or, equivalently, for the *smoothness* of exchange rates to be explained.

We conduct our analysis only for country pairs and two assets in each country—the stock market and a risk-free asset. Taking into account more countries or more assets will make our results even stronger, since the maximum Sharpe ratio obtainable in international asset markets can only increase. Adding conditioning information is likely

---

[16]Some examples include Heaton and Lucas (1997), Lewis (1996, 1999), Luttmer (1999), and Davis et al. (2003).
[17]For example Mankiw and Zeldes (1991) and Vissing-Jorgenson (2002).

to increase the implied level of risk sharing as well, since the presence of conditioning information is just like the presence of additional assets (managed portfolios) that can only increase unconditional Sharpe ratios. Conditioning information may however give time-variation in the amount of risk that is shared. For example, Brandt and Santa-Clara (2002) find that the volatility of exchange rates drops at times very significantly, which suggests even higher values of the risk sharing index at those times.

We interpret our results as a puzzle. Asset markets allow the *potential* of a great deal of risk sharing. To conclude that risks are *in fact* well shared, all we need to believe is that investors obey first order conditions for portfolio allocation. Even approximately optimal portfolios will do, as the magnitudes are so large. This conclusion needs *no* statement about preferences, technology, endowments, shocks, etc. It is just the conclusion that consumer's marginal rates of substitution are equated once we observe a common price ratio. To conclude that risks are well shared *overall*, all we need believe is that additional risks, orthogonal to asset returns, are of reasonable magnitude, i.e. they contribute less than 100% per year to the volatility of marginal utility growth. These conclusions are counterintuitive, and they contradict calculations made with quantity data and simple utility functions, as do a wide range of asset price phenomena. Still, they are a simple and inescapable logical result of Eq. (1) and an equity premium above one percent.

Appendix A

Here, we go through the algebra to find the excess returns in Eqs. (12), (13) given the processes for stock and bond returns and the exchange rate in (10). The excess domestic stock return is

$$\frac{\mathrm{d}S^d}{S^d} - \frac{\mathrm{d}B^d}{B^d} = (\theta^d - r^d)\,\mathrm{d}t + \mathrm{d}z^d. \tag{A.1}$$

To invest in the foreign risk-free bond, the domestic investor also faces currency risk. Hence, the excess return is

$$\frac{\mathrm{d}(eB^f)}{eB^f} - \frac{\mathrm{d}B^d}{B^d} = \frac{\mathrm{d}e}{e} + r^f\,\mathrm{d}t - r^d\,\mathrm{d}t = (\theta^e + r^f - r^d)\,\mathrm{d}t + \mathrm{d}z^e. \tag{A.2}$$

The foreign excess return to foreign stocks is

$$\frac{\mathrm{d}S^f}{S^f} - \frac{\mathrm{d}B^f}{B^f} = (\theta^f - r^f)\,\mathrm{d}t + \mathrm{d}z^f. \tag{A.3}$$

Accounting for exchange risk, the domestic returns to holding a foreign stock and risk-free asset are $\mathrm{d}(eS^f)/eS^f$ and $\mathrm{d}(eB^f)/eB^f$, respectively. Thus, the corresponding domestic excess return is

$$
\begin{aligned}
\frac{\mathrm{d}(eS^f)}{eS^f} - \frac{\mathrm{d}(eB^f)}{eB^f} &= \frac{\mathrm{d}S^f}{S^f} - \frac{\mathrm{d}B^f}{B^f} + \frac{\mathrm{d}e}{e}\left(\frac{\mathrm{d}S^f}{S^f} - \frac{\mathrm{d}B^f}{B^f}\right) \\
&= \left(1 + \frac{\mathrm{d}e}{e}\right)\left(\frac{\mathrm{d}S^f}{S^f} - \frac{\mathrm{d}B^f}{B^f}\right) \\
&= (1 + \theta^e\,\mathrm{d}t + \mathrm{d}z^e)[(\theta^f - r^f)\,\mathrm{d}t + \mathrm{d}z^f]
\end{aligned}
$$

$$= (\theta^f - r^f)\,dt + dz^e\,dz^f + dz^f$$

$$= (\theta^f - r^f + \Sigma^{ef})\,dt + dz^f. \tag{A.4}$$

(This excess return differences two risky assets, foreign stock less foreign bond, rather than foreign stock over domestic bond. This choice gives slightly simpler formulas with no change in results.)

Stacking the three excess returns (A.1), (A.2), and (A.4), the vector of mean excess returns to a domestic investor is (12) and the covariance matrix of excess returns is just Σ.

The foreign excess return from holding foreign stocks is already given by Eq. (A.3). The foreign excess return from borrowing at the domestic interest rate and investing at the foreign interest rate is[18]

$$-\frac{d(1/e)}{(1/e)} - (r^d - r^f)\,dt = (r^f - r^d)\,dt + \frac{de}{e} - \frac{de^2}{e^2}$$

$$= (\theta^e + r^f - r^d - \Sigma^{ee})\,dt + dz^e. \tag{A.5}$$

The foreign returns to holding a domestic stock and risky bond are $d(S^d/e)/(S^d/e)$ and $d(B^d/e)/(B^d/e)$, respectively. Therefore, the excess return to a foreign investor is

$$\frac{d(S^d/e)}{S^d/e} - \frac{d(B^d/e)}{B^d/e} = \left(\frac{dS^d}{S^d} - \frac{de}{e} + \frac{de^2}{e^2} - \frac{de}{e}\frac{dS^d}{S^d}\right) - \left(\frac{dB^d}{B^d} - \frac{de}{e} + \frac{de^2}{e^2} - \frac{de}{e}\frac{dB^d}{B^d}\right)$$

$$= \frac{dS^d}{S^d} - \frac{dB^d}{B^d} - \frac{de}{e}\left(\frac{dS^d}{S^d} - \frac{dB^d}{B^d}\right)$$

$$= (1 - \theta^e\,dt - dz^e)[(\theta^d - r^d)\,dt + dz^d]$$

$$= (\theta^d - r^d)\,dt - dz^e\,dz^d + dz^d$$

$$= (\theta^d - r^d - \Sigma^{ed})\,dt + dz^d. \tag{A.6}$$

Stacking the excess returns (A.3), (A.5), and (A.6), the vector of mean excess returns faced by foreign investors is given in (13) and the covariance matrix is again Σ.

Next, we prove directly that the minimum-variance discount factors satisfy the identity (21). Starting from the definitions (12) and (13) and using (19),

$$d\ln\Lambda^f = -(r^f + \tfrac{1}{2}\mu^{f\prime}\Sigma^{-1}\mu^f)\,dt - \mu^{f\prime}\Sigma^{-1}dz$$

$$= -(r^f + \tfrac{1}{2}(\mu^d - \Sigma^e)'\Sigma^{-1}(\mu^d - \Sigma^e))\,dt - (\mu^d - \Sigma^e)\Sigma^{-1}\,dz$$

$$= -(r^f + \tfrac{1}{2}\mu^{d\prime}\Sigma^{-1}\mu^d - \mu^{d\prime}\Sigma^{-1}\Sigma^e + \tfrac{1}{2}\Sigma^{e\prime}\Sigma^{-1}\Sigma^e)\,dt - \mu^{d\prime}\Sigma^{-1}\,dz + dz^e$$

$$= -(r^f + \tfrac{1}{2}\mu^{d\prime}\Sigma^{-1}\mu^d - (\theta^e + r^f - r^d) + \tfrac{1}{2}\Sigma^{ee})\,dt - \mu^{d\prime}\Sigma^{-1}\,dz + dz^e$$

$$= -(r^d + \tfrac{1}{2}\mu^{d\prime}\Sigma^{-1}\mu^d - \theta^e + \tfrac{1}{2}\Sigma^{ee})\,dt - \mu^{d\prime}\Sigma^{-1}\,dz + dz^e$$

$$= d\ln\Lambda^d - (-\theta^e + \tfrac{1}{2}\Sigma^{ee})\,dt + dz^e$$

$$= d\ln\Lambda^d + d\ln e. \tag{A.7}$$

---

[18]Note that $d(1/e) = -(1/e^2)\,de + (1/e^3)\,de^2$ and $\frac{d(1/e)}{1/e} = -\frac{de}{e} + \frac{de^2}{e^2}$. We write the foreign excess return slightly unconventionally as the return to borrowing domestically to invest in foreign assets. This specification highlights the symmetry between domestic and foreign investors. Of course, the sign of an excess return makes no difference to the volatility of the discount factor (see Eq. (9)).

References

Alesina, A., Barro, R., Tenreyro, S., 2003. Optimal currency areas. NBER Macroeconomics Annual, 2002, 2003, Vol. 17, pp. 301–345. MIT Press, Cambridge and London.

Backus, D.K., Smith, G.W., 1993. Consumption and real exchange rates in dynamic economies with non-traded goods. Journal of International Economics 35, 297–316.

Backus, D.K., Kehoe, P.J., Kydland, F.E., 1992. International real business cycles. Journal of Political Economy 100, 745–755.

Backus, D.K., Foresi, S., Telmer, C.I., 1996. Affine models of currency pricing. Working paper, NBER #5623.

Backus, D.K., Foresi, S., Telmer, C.I., 2001. Affine term structure models and the forward premium anomaly. Journal of Finance 56, 279–304.

Baxter, M., Jermann, U.J., 1997. The international diversification puzzle is worse than you think. American Economic Review 87, 170–191.

Baxter, M., Jermann, U.J., King, R., 1995. Nontraded goods, nontraded factors, and international nondiversification. Journal of International Economics 46, 211–229.

Bell, W.R., Wilcox, D.W., 1993. The effect of sampling error on the time series behavior of consumption data. Journal of Econometrics 55, 235–265.

Black, F., Litterman, R., 1992. Global portfolio optimization. Financial Analysts Journal 48, 28–43.

Brandt, M.W., Santa-Clara, P., 2002. Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets. Journal of Financial Economics 63, 161–210.

Brennan, M.J., Solnik, B., 1989. International risk sharing and capital mobility. Journal of International Money and Finance 8, 359–373.

Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. Journal of Political Economy 107, 205–251.

Chen, Z., Knez, P.J., 1995. Measurement of market integration and arbitrage. Review of Financial Studies 8, 287–325.

Cochrane, J.H., 2004. Asset Pricing, revised edition. Princeton University Press, Princeton, NJ.

Cochrane, J.H., Saá-Requejo, J., 2000. Beyond arbitrage: good-deal asset price bounds in incomplete markets. Journal of Political Economy 108, 79–119.

Constantinides, G.M., 2002. Rational asset prices. Journal of Finance 57, 1567–1591.

Constantinides, G.M., Duffie, D., 1996. Asset pricing with heterogeneous consumers. Journal of Political Economy 104, 219–240.

Coval, J.D., Moskowitz, T.J., 1999. Home bias at home: local equity preference in domestic portfolios. Journal of Finance 54, 2045–2074.

Davis, S., Kubler, F., Willen, P., 2003. Borrowing costs and the demand for equity over the life cycle. Working paper, University of Chicago.

Davis, S.J., Nalewaik, J., Willen, P., 2001. On the gains to international trade in risky financial assets. Working paper, University of Chicago.

De Santis, G., Gerard, B., 1997. International asset pricing and portfolio diversification with time-varying risk. Journal of Finance 52, 1881–1912.

Flood, R.P., Rose, A.K., 1995. Fixing the exchange rate regime: a virtual quest for fundamentals. Journal of Monetary Economics 36, 3–37.

French, K., Poterba, J., 1991. Investor diversification and international equity markets. American Economic Review 81, 222–226.

Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. Econometrica 50, 1029–1054.

Hansen, L.P., Jagannathan, R., 1991. Implications of security market data for models of dynamic economies. Journal of Political Economy 99, 225–262.

Heaton, J., Lucas, D., 1997. Market frictions, savings behavior, and portfolio choice. Macroeconomic Dynamics 1, 76–101.

Lewis, K.K., 1996. What can explain the lack of international consumption risk sharing? Journal of Political Economy 104, 267–297.

Lewis, K.K., 1999. Trying to explain home bias in equities and consumption. Journal of Economic Literature 37, 571–608.

Lewis, K.K., 2000. Why do stocks and consumption imply such different gains from international risk sharing? Journal of International Economics 52, 1–35.

Luttmer, E.G.J., 1999. What level of fixed costs can reconcile consumption and stock returns. Journal of Political Economy 107, 969–997.

Mankiw, N.G., Zeldes, S.P., 1991. The consumption of stockholders and non-stockholders. Journal of Financial Economics 29, 97–112.

Mark, N., 1995. Exchange rates and fundamentals: evidence on long horizon predictability. American Economic Review 85, 201–218.

Meese, R., Rogoff, K., 1983. Empirical exchange rate models of the seventies: do they fit out of sample. Journal of International Economics 14, 3–24.

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Obstfeld, M., 1992. International risk sharing and capital mobility: another look. Journal of International Money and Finance 11, 115–121.

Obstfeld, M., 1994. Are industrial-country consumption risks globally diversified? In: Leiderrman, L., Razin, A. (Eds.), Capital Mobility: The Impact on Consumption, Investment, and Growth. Cambridge University Press, Cambridge, MA.

Obstfeld, M., Rogoff, K., 2000. The six major puzzles in international macroeconomics: is there a common cause. In: Bernanke, B., Rogoff, K. (Eds.), NBER Macroeconomics Annual. MIT Press, Cambridge, MA.

Ravn, M., 2001. Consumption dynamics and real exchange rates. Working paper, London Business School.

Rogoff, K., 1999. Perspectives on exchange rate volatility. In: Feldstein, M. (Ed.), International Capital Flows. University of Chicago Press, Chicago.

Solnik, B.H., 1974. Why not diversify internationally rather than domestically. Financial Analysts Journal 30, 48–54.

Tesar, L.L., 1993. International risk sharing and non-traded goods. Journal of International Economics 35, 69–89.

Tesar, L.L., Werner, I.M., 1995. Home bias and high turnover. Journal of International Money and Finance 14, 467–492.

Vissing-Jorgenson, A., 2002. Limited asset market participation and the elasticity of intertemporal substitution. Journal of Political Economy 110, 825–853.

Wilcox, D.W., 1992. The construction of US consumption data: some facts and their implications for empirical work. American Economic Review 82, 922–941.