The Forward Premium Puzzle in a Two-Country World

Ian Martin*

March, 2013

Abstract

I explore the behavior of asset prices and the exchange rate in a two-country world. When the large country has bad news, the relative price of the small country's output declines. As a result, the small country's bonds are risky, and uncovered interest parity fails, with positive excess returns available to investors who borrow at the large country's interest rate and lend at the small country's interest rate. I use a diagrammatic approach to derive these and other results in a calibration-free way.

JEL code: G15.

*ian.martin@gsb.stanford.edu; http://www.stanford.edu/∼iwrm/. First draft: April, 2010. I am grateful to Fernando Alvarez, Dave Backus, Thomas Baranga, Jeremy Bulow, John Cochrane, Darrell Duffie, Bernard Dumas, Lars Hansen, Tarek Hassan, Bruce Lehmann, Robert Lucas, Monika Piazzesi, Roberto Rigobon, Martin Schneider, Andy Skrzypacz, Nancy Stokey, and Harald Uhlig, and to seminar participants at Central European University, the University of Chicago Money & Banking Workshop, the Stanford Institute for Theoretical Economics conference, Stanford GSB, and UCSD for their comments.

An extensive literature has documented the fact that interest-rate differentials across countries are not on average counteracted by offsetting currency movements. As a result, it is possible to earn excess returns by investing in high-interest-rate currencies and borrowing in low-interest-rate currencies. More specifically, Hassan (2012) documents an empirical relationship between interest rates and country size: smaller economies tend to have persistently higher interest rates, and the returns on smaller economies' bonds have higher covariances with long-run consumption growth of US stockholders.

This paper presents a two-country model in which this failure of uncovered interest parity (UIP)—also known as the forward premium puzzle—emerges in equilibrium, and with this size-related pattern. The outputs of the two countries are imperfect substitutes for one another, so units matter, and there are two term structures of interest rates, one for each good. Global financial markets are assumed to be perfectly integrated; assets are real rather than nominal, and are priced by marginal global investors with power utility; there are no non-tradable goods, no liquidity issues, no portfolio constraints.

Even so, UIP fails to hold. Suppose, for example, that both countries have the same distribution of output growth, and that one country is much smaller than the other. In equilibrium, the smaller country has a higher short-term real interest rate. But the small country's exchange rate does not depreciate enough, on average, to offset the interest-rate differential. In fact, on the contrary, the exchange rate is expected to appreciate (in an example of Siegel's (1972) "paradox"), thereby *increasing* the expected excess return on the carry trade.

Why does the smaller country have a higher interest rate? Any risk-based explanation must provide a story for why the small country's bond underperforms in bad states of the world. In the model considered here, the marginal investor is particularly concerned with states in which the large country experiences low output growth, since its output contributes the majority of the marginal investor's consumption. In such states, there is an increase in the relative supply of the small country's output, so its exchange rate depreciates, causing the small-country bond to underperform; hence the risk premium.

These effects do not depend on finely tuned parameter values: UIP fails in *any* calibration of the model, and I provide economically interpretable nonparametric conditions under which I am able to sign the direction in which UIP fails—that is, to show when it is the small country whose bonds are risky, as in the example above. An unusual feature of the paper is that it develops a method of demonstrating that this relationship holds via diagrammatic proofs. (More conventional, but less comprehensible, algebraic proofs are provided in the supplementary appendix.) I use this approach to show that various other relationships hold within the model, for example between risk premia on the large and small countries' output claims, in own and foreign units.

Although the ingredients of the model are entirely standard, it has not previously been solved even in the lognormal case: the interactions between the model's additive features (a CES aggregator over multiple goods) and multiplicative features (power utility and i.i.d. log consumption growth) that make it interesting also make it hard to solve. These features capture the familiar tension in finance that portfolio formation is additive while returns compound multiplicatively. The loglinearization approach of Campbell and Shiller (1987) provides an approximate solution to this problem. In the context of international finance, an alternative approach (adopted, for example, by Cole and Obstfeld (1991), Zapatero (1995), Pavlova and Rigobon (2007), and Stathopoulos (2009)) has been to focus on the knife-edge, but tractable because multiplicative, Cobb-Douglas case of unit elasticity of substitution between goods.

This paper avoids the need to loglinearize or to assume unit elasticity of substitution and provides an exact solution via Fourier transform methods, extending Martin (2013b). Prices, interest rates and so on are expressed in terms of integral formulas that allow for a range of specifications of output dynamics and can be evaluated numerically almost instantly. They can also be solved analytically in the limit in which one country becomes very small relative to the other, and this limit turns out to capture many important features of the model in a highly tractable setting. A further benefit is that it is possible to provide explicit conditions that ensure that expected utility is finite; perhaps surprisingly, it does not appear to be conventional to do so in the literature (Hansen and Scheinkman (2012) being a notable exception).

Last but not least, the approach allows for jumps. Motivation for doing so is provided by an increasingly large body of work that emphasizes the importance of nonlognormality in financial markets (Rietz (1988), Barro (2006), Backus, Chernov, and Martin (2011)); in particular, Martin (2011) shows that *no* conditionally lognormal model is consistent with index option prices in the US. Jumps also enable the model to replicate the finding that carry trade returns are negatively skewed (Brunnermeier, Nagel and Pedersen (2008), Jurek (2009)) and to generate sensible risk premia without implausibly high risk aversion. The key to retaining tractability in the presence of jumps is the cumulant-generating function (CGF), which was also exploited in Martin (2013b). This paper introduces some new properties that CGFs may or may not possess, and which provide nonparametric ways of restricting possible calibrations. For example, in a lognormal world one might want to assume that countries have *linked fundamentals* in the sense of having (weakly) positively correlated log output growth. A contribution of the paper is to identify—for this example and others—the right definition that generalizes beyond the lognormal case.

A final benefit of allowing for jumps is that the results hold independently of Merton's (1973) ICAPM and Breeden's (1979) consumption-CAPM, which do not hold once asset prices can jump; this is reassuring, given the mixed empirical performance of the ICAPM and consumption-CAPM.

When not included in the body of the paper, proofs are in the appendix.

Related literature. Frankel (1980), Hansen and Hodrick (1980), and Fama (1984) are amongst the early contributions to the literature on the forward premium puzzle. More generally, Dumas, Harvey and Ruiz (2003) present evidence that international markets are well integrated; Lustig, Roussanov and Verdelhan (2009) provide support for the idea that the behavior of international asset prices is amenable to a risk-based explanation; and Hollifield and Yaron (2003) argue that models of currency risk premia should focus on real, as opposed to nominal, risk. Alvarez, Atkeson and Kehoe (2007) emphasize that the empirical fact that exchange rates are approximately random walks implies that fluctuations in interest rates are associated (in their lognormal setting) with changes in conditional variances of the log SDF. This is a feature of the model presented below.[1]

Hassan (2012) motivates his empirical results with a two-period, lognormal and loglinearized model with one tradable good and multiple non-tradable goods, under an assumption that a transfer before the start of period 1 equates all agents'

[1] Of the theoretical papers mentioned in this introduction, Alvarez, Atkeson and Kehoe (2007), Cochrane, Longstaff and Santa-Clara (2008), Colacito and Croce (2010), Hassan (2012), Lustig, Roussanov and Verdelhan (2009), Pavlova and Rigobon (2007), Stathopoulos (2009), Verdelhan (2009), and Zapatero (1995) impose conditional lognormality.

wealth. Verdelhan (2010) considers a model with habit formation. Colacito and Croce (2010) and Bansal and Shaliastovich (2010) argue that long-run risk, together with Epstein-Zin preferences, can explain the forward premium puzzle. Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2009) construct a model featuring rare disasters, and present supporting evidence from FX option data.

Piazzesi, Schneider and Tuzel (2007) present a model with housing and non-housing goods that are imperfect substitutes for one another. Their stochastic discount factor takes the same form as that of the present paper, but the driving stochastic processes are modelled differently so that their analog of my state variable u mean-reverts rather than following a Lévy process.

To my knowledge, no other paper in this literature derives calibration-free results of the type presented here.

1 UIP and the forward premium puzzle

UIP is a conjectured relationship between next year's spot exchange rate between two countries, e_{t+1}, today's spot exchange rate, e_t, and 1-year interest rates in each country, $i_{1,t}$ and $i_{2,t}$:

$$\mathbb{E}_t \log e_{t+1} \stackrel{?}{=} \log e_t + i_{1,t} - i_{2,t} \,. \tag{1}$$

The thought behind (1) is this: if country 2 has a lower interest rate than country 1, surely this should be compensated by the expected appreciation of its currency? Unfortunately this plausible idea is decisively rejected by the data: in the relationship

$$\log e_{t+1} - \log e_t = a_0 + a_1(i_{1,t} - i_{2,t}) + \varepsilon_{t+1} \,, \tag{2}$$

UIP holds if $a_0 = 0$ and $a_1 = 1$. But typical estimates place a_1 close to (and often less than) zero.

An equivalent formulation of UIP exploits *covered* interest parity, which is the no-arbitrage—and therefore robust—relationship between today's 1-year forward exchange rate, f_t, today's spot exchange rate, and the two countries' 1-year rates: $\log f_t = \log e_t + i_{1,t} - i_{2,t}$. Using this relationship, equation (1) can be expressed as $\mathbb{E}_t \log e_{t+1} = \log f_t$, so the failure of UIP is equivalent to the fact that log forward rates are biased predictors of future log exchange rates.

If we write $M_{i,t+1}$ for the stochastic discount factor (SDF) that prices assets whose payoffs are in units of country i's good, then (as we will see below) the relationship $e_{t+1}/e_t = M_{2,t+1}/M_{1,t+1}$ holds. Taking logs then expectations we have the identity

$$\mathbb{E}_t \, \Delta \log e_{t+1} = \mathbb{E}_t \log M_{2,t+1} - \mathbb{E}_t \log M_{1,t+1} \,. \tag{3}$$

Now, if $M_{1,t+1}$ and $M_{2,t+1}$ were roughly constant—as would be the case if either there were little risk in the economy, or if investors were roughly risk-neutral—then we could approximate (3) by

$$\mathbb{E}_t \, \Delta \log e_{t+1} \approx \log \mathbb{E}_t \, M_{2,t+1} - \log \mathbb{E}_t \, M_{1,t+1} = i_{1,t} - i_{2,t} \,. \tag{4}$$

That is, UIP would approximately hold in economies in which either the price or quantity of risk was very low. Empirically, however, the fact that high Sharpe ratios are attainable tells us, thanks to the Hansen-Jagannathan (1991) bound, that the SDFs $M_{i,t+1}$ are volatile. Thus the move from (3) to (4) was not justified, and we must take the effects of Jensen's inequality into account, arriving at the identity

$$\mathbb{E}_t \, \Delta \log e_{t+1} = i_{1,t} - i_{2,t} + \underbrace{\log \mathbb{E}_t \, M_{1,t+1} - \mathbb{E}_t \log M_{1,t+1}}_{L_t(M_{1,t+1})} - \underbrace{\left(\log \mathbb{E}_t \, M_{2,t+1} - \mathbb{E}_t \log M_{2,t+1}\right)}_{L_t(M_{2,t+1})} \,. \tag{5}$$

The terms $L_t(M_{i,t+1})$ measure SDF variability; I call $L_t(M_{i,t+1})$ the *entropy* of $M_{i,t+1}$, following Backus, Chernov and Martin (2011). High attainable expected log returns translate into high SDF entropy (Bansal and Lehmann (1997), Alvarez and Jermann (2005)), much as high attainable Sharpe ratios translate into high SDF volatility.

The identity (5) reveals some necessary ingredients of any model in which UIP fails to hold (see also Fama (1984), and Backus, Foresi and Telmer (2001)). First, as discussed, the entropies of $M_{1,t+1}$ and $M_{2,t+1}$ must be non-zero and economically significant: risk must matter. Second, there must be an asymmetry: if the entropies were equal they would cancel out, returning us to a world in which UIP held. Third, to generate the patterns found when estimating (2), $L_t(M_{1,t+1}) - L_t(M_{2,t+1})$ should be small at times when $i_{1,t} - i_{2,t}$ is high: assets denominated in the high-interest-rate currency should have relatively low risk-adjusted returns. The model that follows has these properties.

2 The model

There are two countries with output streams $\{D_{1t}\}$ and $\{D_{2t}\}$ respectively, at least one of which is nondeterministic. There are three classes of agents. The first two classes are made up of *locals* in each country; these are hand-to-mouth consumers who, in aggregate, consume a fraction $1-\phi$ of their own country's output and do not eat the foreign country's output. The third class comprises the *jetsetters*: these individuals eat goods from both countries, consuming the remaining fraction ϕ of each country's output. (One could embed this structure in a model in which output is produced as a Cobb-Douglas function of capital, owned by jetsetters, and labor, owned by locals.) The goods are viewed as imperfect substitutes by jetsetters, and $\eta \in [1,\infty)$ is the elasticity of intratemporal substitution between goods 1 and 2. Locals do not participate in financial markets, so assets are priced by the jetsetters, each of whose expected utility takes the form (up to the irrelevant constant multiple $\phi^{1-\gamma}$)

$$\mathbb{E}\int_0^\infty e^{-\rho t}\frac{C_t^{1-\gamma}}{1-\gamma}\,dt\,, \quad \text{where} \quad C_t \equiv \left[w^{1/\eta}D_{1t}^{\frac{\eta-1}{\eta}} + (1-w)^{1/\eta}D_{2t}^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}.$$

The constant w controls the relative importance of goods 1 and 2 for jetsetters.

These assumptions implicitly embed the view that financial markets are well integrated, since the jetsetters are the marginal agents pricing both countries' output claims and riskless bonds; as a result, asset prices will be more highly correlated across countries than outputs, and shocks will be transmitted from one market to another. On the other hand, the model is also consistent with the view that aggregate consumption risk is imperfectly shared across countries (Backus, Kehoe and Kydland (1992)) due to the presence of the locals who do not participate in asset markets. Nonparticipation is a feature of the data: Calvet, Campbell and Sodini (2007) analyze a comprehensive dataset that covers all Swedish households and report that 38 per cent of households were nonparticipants.

Write $v(D_{1t},D_{2t}) \equiv C_t^{1-\gamma}/(1-\gamma)$ for the instantaneous felicity function, $v_i(D_{1t},D_{2t})$ for the marginal utility of good i, $e_t \equiv v_2(D_{1t},D_{2t})/v_1(D_{1t},D_{2t})$ for the intratemporal price of a unit of good 2 in units of good 1, and $M_{i\tau} \equiv e^{-\rho(\tau-t)}v_i(D_{1\tau},D_{2\tau})/v_i(D_{1t},D_{2t})$ for the SDF that prices time-τ claims to good i. The SDFs $M_{1\tau}$ and $M_{2\tau}$ and the relative price e_t are related, as noted above, by the equation

$$\frac{e_\tau}{e_t} = \frac{M_{2\tau}}{M_{1\tau}}\,,$$

which appears as Proposition 1 of Backus, Foresi and Telmer (2001) and as equation 1 of Brandt, Cochrane and Santa-Clara (2006). I refer to the relative price e_t as the exchange rate, because it is the relative price of locals' consumption baskets in the two countries. When the relative supply of country 2's good declines, the exchange rate e_t increases.[2]

The price at time t of asset i is given by Lucas's (1978) Euler equation

$$P_{it} = \mathbb{E} \int_t^\infty M_{i\tau} D_{i\tau} \, d\tau \, . \tag{6}$$

Similarly, the price (in own units) of a perpetuity that pays out a constant stream of good i dividends at a rate of one per unit time, is given by (6) with $D_{i\tau}$ replaced by 1.

It is important to emphasize, first, that we need two different SDFs because there are two different types of assets, with payouts denominated in two different goods and, second, that the perturbation logic underlying (6) implies that the price P_{it} is *denominated in units of good i* ("in i-units"). I use stars to indicate a price *not* expressed in its own units: $P_{2t}^* = P_{2t} e_t$ and $P_{1t}^* = P_{1t}/e_t$.

The outputs of the two countries, D_{1t} and D_{2t}, are taken as exogenous—though they could also be thought of as being determined by a production side of the economy that is left unmodelled here—and are assumed to have dividend growth that is i.i.d. over time, though not necessarily across countries. Formally, $y_{it} - y_{i0} \equiv \log D_{it} - \log D_{i0}$ is a Lévy process for $i = 1, 2$.

Definition 1. *The* cumulant-generating function (CGF) $\boldsymbol{c}(\theta_1, \theta_2)$ *is defined by*

$$\boldsymbol{c}(\theta_1, \theta_2) \equiv \log \mathbb{E}\, e^{\theta_1 (y_{1,t+1} - y_{1,t}) + \theta_2 (y_{2,t+1} - y_{2,t})} = \log \mathbb{E} \left[\left(\frac{D_{1,t+1}}{D_{1,t}} \right)^{\theta_1} \left(\frac{D_{2,t+1}}{D_{2,t}} \right)^{\theta_2} \right] .$$

The CGF captures all relevant information about the technological side of the model. Mean log output growth of country j is $\frac{\partial \boldsymbol{c}}{\partial \theta_j}(0,0)$, the variance of log output

[2]The model does not address the finding of Backus and Smith (1993) that there is considerable variation in exchange rates that is not explained by variation in relative consumption. Stockman and Tesar (1995) suggest that this shortcoming can be addressed by introducing taste shocks. In the present setup we can do so by allowing the weights w and $1 - w$ in the consumption aggregator to follow exponential Lévy processes. This extension comes at the cost of an increase in complexity, however, and the goal of this paper is to demonstrate that the failure of UIP emerges robustly in a simple model.

growth of country j is $\frac{\partial^2 \boldsymbol{c}}{\partial \theta_j^2}(0,0)$, and the covariance between the log output growth of countries j and k is $\frac{\partial^2 \boldsymbol{c}}{\partial \theta_j \partial \theta_k}(0,0)$. Third-order and higher partial derivatives at the origin capture higher cumulants and co-cumulants of output growth: skewness, excess kurtosis, and so on.

By exploiting general properties of CGFs, it is possible to establish features of asset prices in this economy that hold not just for a particular calibration but for a whole family of driving stochastic processes. The most important such property is that CGFs are always *convex*. But I also find it helpful to introduce three nonparametric properties, each of which the CGF may or may not possess. The first is the *exchangeability* property,[3] which holds if $\boldsymbol{c}(\theta_1,\theta_2) = \boldsymbol{c}(\theta_2,\theta_1)$. This is a *cet. par.* assumption: it ensures the two countries have the same means and volatilities of output growth, the same arrival rates of jumps, and so on. It therefore focusses attention on the underlying economic mechanism and on the consequences of asymmetry in country size alone. The second is the *convex difference* property, which is a restriction on the higher cumulants of log output growth. It holds in the lognormal case, and more generally it ensures, roughly speaking, that output growth in each country is not positively skewed. The third is the *linked fundamentals* property. In the lognormal case, it is natural to consider imposing the economically plausible assumption that the correlation between the two countries' output growth is nonnegative. The linked fundamentals property, which will turn out to be the appropriate generalization for arbitrary Lévy processes, requires that the CGF is supermodular.

It may be helpful to consider some parametric examples of CGFs. Write $\boldsymbol{\theta} \equiv (\theta_1,\theta_2)'$ and $\boldsymbol{y}_t \equiv (y_{1t}, y_{2t})'$.

(i) If output growth is lognormal, i.e. $\boldsymbol{y}_t = \boldsymbol{y}_0 + \boldsymbol{\mu} t + \boldsymbol{A}\boldsymbol{Z}_t$, where $\boldsymbol{\mu}$ is a 2-dimensional vector of drifts, $\boldsymbol{A}$ a 2×2 matrix of factor loadings, and $\boldsymbol{Z}_t$ a 2-dimensional Brownian motion, then the CGF is

$$\boldsymbol{c}(\theta_1,\theta_2) = \boldsymbol{\mu}'\boldsymbol{\theta} + \boldsymbol{\theta}'\boldsymbol{\Sigma}\boldsymbol{\theta}/2, \tag{7}$$

where $\boldsymbol{\Sigma} \equiv \boldsymbol{A}\boldsymbol{A}'$ is the covariance matrix of log output growth, whose elements I write as σ_{jk}. If the countries have independent output growth then $\boldsymbol{\Sigma}$ is a

[3] Two random variables X_1, X_2 are *exchangeable* if the distribution of (X_1, X_2) is the same as that of (X_2, X_1).

diagonal matrix, so the CGF decomposes as $\boldsymbol{c}(\theta_1,\theta_2)=\boldsymbol{c}_1(\theta_1)+\boldsymbol{c}_2(\theta_2)$ where $\boldsymbol{c}_k(\theta_k)\equiv\mu_k\theta_k+\sigma_{kk}\theta_k^2/2$.

(ii) If log dividends follow a jump-diffusion, we can write $\boldsymbol{y}_t=\boldsymbol{y}_0+\boldsymbol{\mu}t+\boldsymbol{A}\boldsymbol{Z}_t+\sum_{k=1}^{K(t)}\boldsymbol{J}^k$, where $K(t)$ is a Poisson process with arrival rate $\widetilde{\omega}$ that represents the number of jumps that have taken place by time t, and $\boldsymbol{J}^k$ are 2-dimensional random variables that are i.i.d. across k. There may be arbitrary correlations between the elements of $\boldsymbol{J}$. (I will write $\boldsymbol{J}\equiv\boldsymbol{J}^1$ when I discuss the distribution of these random variables.) The CGF then acquires an extra term, whose precise nature depends on the assumptions made about the jump size distribution: $\boldsymbol{c}(\theta_1,\theta_2)=\boldsymbol{\mu}'\boldsymbol{\theta}+\boldsymbol{\theta}'\boldsymbol{\Sigma}\boldsymbol{\theta}/2+\widetilde{\omega}\left(\mathbb{E}\,e^{\boldsymbol{\theta}'\boldsymbol{J}}-1\right)$.

(iii) By specifying the jump distribution $\boldsymbol{J}$ appropriately, this framework allows for multiple types of jump. Letting p_1, p_2, and p_3 be probabilities summing to one, suppose that with probability p_1, $\boldsymbol{J}$ shocks the log output of country 1 by a Normal random variable with mean $\mu_1^{(1)}$ and variance $\sigma_{11}^{(1)}$; with probability p_2, $\boldsymbol{J}$ shocks the log output of country 2 by a Normal random variable with mean $\mu_2^{(2)}$ and variance $\sigma_{22}^{(2)}$; and with probability p_3, $\boldsymbol{J}$ shocks both log outputs simultaneously by a bivariate Normal random variable with mean $(\mu_1^{(3)},\mu_2^{(3)})'$ and covariance matrix $\begin{pmatrix}\sigma_{11}^{(3)} & \sigma_{12}^{(3)}\\ \sigma_{12}^{(3)} & \sigma_{22}^{(3)}\end{pmatrix}$. This is a special case of Example (ii), so the CGF is

$$\begin{aligned}\boldsymbol{c}(\theta_1,\theta_2)&=\mu_1\theta_1+\mu_2\theta_2+\frac{1}{2}\sigma_{11}\theta_1^2+\sigma_{12}\theta_1\theta_2+\frac{1}{2}\sigma_{22}\theta_2^2+\omega_1\left(e^{\mu_1^{(1)}\theta_1+\frac{1}{2}\sigma_{11}^{(1)}\theta_1^2}-1\right)+\\&+\omega_2\left(e^{\mu_2^{(2)}\theta_2+\frac{1}{2}\sigma_{22}^{(2)}\theta_2^2}-1\right)+\omega_3\left(e^{\mu_1^{(3)}\theta_1+\mu_2^{(3)}\theta_2+\frac{1}{2}\sigma_{11}^{(3)}\theta_1^2+\sigma_{12}^{(3)}\theta_1\theta_2+\frac{1}{2}\sigma_{22}^{(3)}\theta_2^2}-1\right),\end{aligned}$$

where $\omega_k=\widetilde{\omega}p_k$ for $k=1,2,3$.

(iv) Specializing still further, suppose that output is i.i.d. across countries. Independence requires that $\sigma_{12}=\omega_3=0$, and identical distribution means that we can simplify the notation by writing $\mu_k=\mu$, $\sigma_{kk}=\sigma^2$, $\mu_k^{(k)}=\mu_J$, $\sigma_{kk}^{(k)}=\sigma_J^2$, and finally $\omega_1=\omega_2=\omega$ for the arrival rate of each country's jumps. The CGF is then

$$\boldsymbol{c}(\theta_1,\theta_2)=\mu\theta_1+\mu\theta_2+\frac{1}{2}\sigma^2\theta_1^2+\frac{1}{2}\sigma^2\theta_2^2+\omega\left(e^{\mu_J\theta_1+\frac{1}{2}\sigma_J^2\theta_1^2}-1\right)+\omega\left(e^{\mu_J\theta_2+\frac{1}{2}\sigma_J^2\theta_2^2}-1\right). \tag{8}$$

The symmetry of the two countries is reflected in the fact that exchangeability holds in (8), i.e. $\boldsymbol{c}(\theta_1,\theta_2)=\boldsymbol{c}(\theta_2,\theta_1)$. Their independence is reflected in the fact that the CGF can be decomposed as a function of θ_1 plus a function of θ_2 (so for example there are no cross-terms $\theta_1\theta_2$).

For the purposes of illustration (only), I consider a calibration in which output growth is i.i.d. across countries so that any correlations or asymmetries that emerge do so endogenously; the associated CGF is given by (8). I assume that log output in each country has a Brownian motion component with drift $\mu = 0.02$ and volatility $\sigma = 0.1$. Each country is also afflicted by jumps, which arrive at times dictated by a Poisson process with arrival rate $\omega = 0.02$. When a country experiences a jump, the shock to its log output is Normal with mean $\mu_J = -0.2$ and volatility $\sigma_J = 0.1$. I also assume that the marginal investor has time preference rate $\rho = 0.04$ and risk aversion $\gamma = 4$; that the weight of country 1 in the consumption aggregator is $w = 0.2$; and that the elasticity of substitution between goods is $\eta = 2$.

For the general results, I assume that $\gamma\eta \geq 2$; typical estimates in the literature put both γ and η somewhere in the range 2–10, so this is a mild assumption. It will be convenient to write $\chi = (\eta - 1)/\eta$ and $\widehat{\gamma} = (\gamma + \chi - 1)/\chi$. The variable χ ranges between 0 (the Cobb-Douglas case) and 1 (the perfect substitutes case). Result 6 requires the further assumption that $\widehat{\gamma}$ is an integer.

Given the specification of tastes and technology—specifically, given that preferences are power utility over a CES aggregator and log outputs follow Lévy processes—price-dividend ratios, interest rates and expected returns are invariant to scalings $(D_{1t}, D_{2t}) \mapsto (kD_{1t}, kD_{2t})$. These quantities do however depend on a single state variable that indexes the *relative* sizes of the two countries. In principle this state variable could be any monotonic function of D_{2t}/D_{1t}, but it will be most convenient to work with

$$u_t \equiv (1-\chi)\log\frac{1-w}{w} + \chi\log\frac{D_{2t}}{D_{1t}} = (1-\chi)\log\frac{1-w}{w} + \chi(y_{2t} - y_{1t}).$$

The state variable u_t is a Lévy process because y_{1t} and y_{2t} are Lévy processes. The constant term $(1-\chi)\log[(1-w)/w]$ ensures that $u_t = 0$ if countries 1 and 2 have equal shares of world output. I plot graphs against the alternative state variable $s_t \equiv 1/(1+e^{u_t})$. This is less mathematically convenient than u_t because it is not a Lévy process, but it has a nice interpretation as country 1's share of world output,

$$s_t = \frac{D_{1t}}{D_{1t} + e_t \cdot D_{2t}}.$$

Before jumping into the solution of the model, it may be helpful to explore one of its most important components. If we define the marginal investor's *risk aversion to a good-i shock* as $\gamma_i \equiv -D_i v_{ii}/v_i$, some algebra shows that

$$\gamma_1 = \frac{1}{\eta} + \left(\gamma - \frac{1}{\eta}\right) s_t \tag{9}$$

and, symmetrically, that risk aversion to a good-2 shock is $\gamma_2 = 1/\eta + (\gamma - 1/\eta)(1 - s_t)$. Thus the marginal investor's attitude to the risk associated with a country's output depends on the size of that country. As country 1's share of world output, s_t, declines from 1 to 0, risk aversion to good-1 shocks declines linearly from γ to $1/\eta$. Equation (9) also provides a way to understand the models of Cochrane, Longstaff and Santa-Clara (2008) and of Martin (2013b)—but in those models $\eta = \infty$, so that when country 1 is negligibly small the marginal investor is risk-neutral to good-1 shocks. Here, with $\eta < \infty$, the agent remains averse to the risk of good 1 even if country 1 is very small because good 2 is only an imperfect substitute for it. This effect will drive up risk premia on a small country's output claim relative to the $\eta = \infty$ case.

2.1 Prices

The first result characterizes the behavior of the exchange rate in terms of the driving output processes. The proof exploits the following useful, and completely general, property of CGFs.

Property 0 (Convexity). *Any cumulant-generating function $\mathbf{c}(\theta_1, \theta_2)$ is convex.*

The convexity property makes it possible to prove results that are valid for dividends driven by *any* Lévy process with finite moments. This property is used in the result below, in a rather simple form, to show that Siegel's (1972) "paradox" holds, and will be used in more complicated situations in subsequent results. Since the convexity logic will be used throughout the paper, I illustrate its application in Figure 1.

Result 1 (Exchange rate). *The exchange rate e_t is*

$$\begin{aligned} e_t &= [(1-w)/w]^{1/\eta} (D_{2t}/D_{1t})^{-1/\eta} \\ &= [(1-w)/w]^{(1-\chi)/\chi} \cdot \exp\left\{-[(1-\chi)/\chi]\, u_t\right\}. \end{aligned} \tag{10}$$

This means that the log exchange rate is a Lévy process.

The expected appreciation in country 1's exchange rate, FX_1^*, is

$$FX_1^* \, dt \equiv \frac{\mathbb{E}\, d(1/e_t)}{1/e_t} = \boldsymbol{c}(\chi - 1, 1 - \chi)\, dt, \tag{11}$$

and the expected appreciation in country 2's exchange rate, FX_2^*, is

$$FX_2^* \, dt \equiv \frac{\mathbb{E}\, de_t}{e_t} = \boldsymbol{c}(1 - \chi, \chi - 1)\, dt. \tag{12}$$

If $\eta < \infty$, so that exchange rates are nonconstant, then the average expected appreciation, $(FX_1^* + FX_2^*)/2$, is positive—an instance of Siegel's paradox.

Proof. Equation (10) follows from the fact that $e_t = v_2(D_{1t}, D_{2t})/v_1(D_{1t}, D_{2t})$ and the definition of the state variable u_t. For arbitrary constants w_1 and w_2,

$$\frac{\mathbb{E}\left[d\left(e^{w_1 y_{1t} + w_2 y_{2t}}\right)\right]}{e^{w_1 y_{1t} + w_2 y_{2t}}} = \boldsymbol{c}(w_1, w_2)\, dt\,. \tag{13}$$

Equations (11) and (12) follow by applying (13) to equation (10).

It remains to show that Siegel's paradox holds, i.e. that $\boldsymbol{c}(\chi - 1, 1 - \chi) + \boldsymbol{c}(1 - \chi, \chi - 1) > 0$. If dividend growth is lognormal, so that the CGF is given by (7), then we must show that $(1 - \chi)^2(\sigma_{11} - 2\sigma_{12} + \sigma_{22}) > 0$. The first factor, $(1 - \chi)^2$, is positive if $\eta < \infty$ since then $\chi < 1$, and the second factor is positive because it equals the variance of $(y_{1,t+1} - y_{1,t}) - (y_{2,t+1} - y_{2,t})$. This establishes the result for the lognormal case.

Property 0 provides exactly the right line of attack to prove the result in the general case. Since $\boldsymbol{c}(0, 0) = 0$ for any CGF, the goal is to show that $\boldsymbol{c}(\chi-1, 1-\chi) - 2\boldsymbol{c}(0, 0) + \boldsymbol{c}(1-\chi, \chi-1)$ is positive. Figure 1a plots the convex surface parametrized by $(\theta_1, \theta_2, \boldsymbol{c}(\theta_1, \theta_2))$. The points $(\chi - 1, 1 - \chi, \boldsymbol{c}(\chi - 1, 1 - \chi))$ and $(1 - \chi, \chi - 1, \boldsymbol{c}(1 - \chi, \chi - 1))$ lie on this surface and are indicated with plus signs; the point $(0, 0, \boldsymbol{c}(0, 0))$—i.e., the origin, since $\boldsymbol{c}(0, 0) = 0$—is correspondingly indicated with (two, overlapping) minus signs. It is now clear that, by convexity, $\boldsymbol{c}(\chi - 1, 1 - \chi) - 2\boldsymbol{c}(0, 0) + \boldsymbol{c}(1 - \chi, \chi - 1)$ is positive, as shown in Figure 1a. For comparison, Figure 1c illustrates the fact that $\boldsymbol{c}(0, -\gamma) + \boldsymbol{c}(-\gamma, 0) - \boldsymbol{c}(\chi-1, 1-\chi-\gamma) - \boldsymbol{c}(1-\chi-\gamma, \chi-1)$ is positive, which will be useful in Result 8, below. □

Convexity arguments, as employed in the above proof, will be useful throughout the paper. Figures 1b and 1d therefore represent the information in Figures 1a and

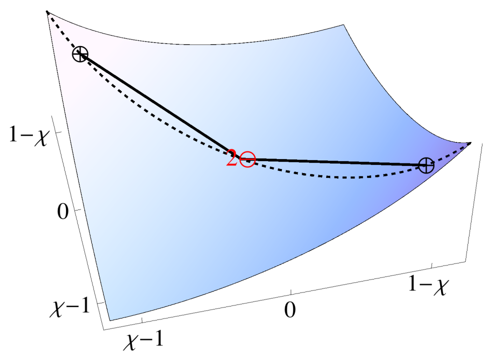

(a) $\boldsymbol{c}(\chi-1,1-\chi)-2\boldsymbol{c}(0,0)+\boldsymbol{c}(1-\chi,\chi-1)>0$

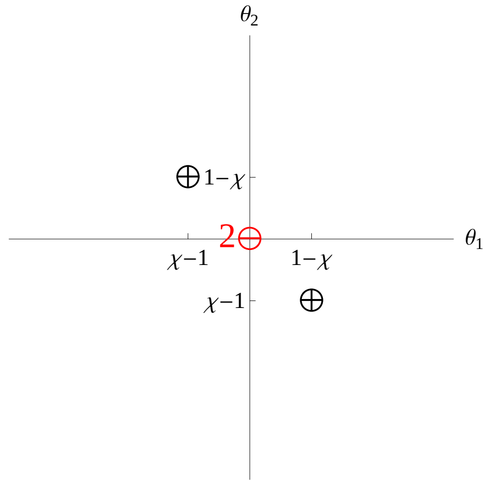

(b) A simpler representation of panel (a)

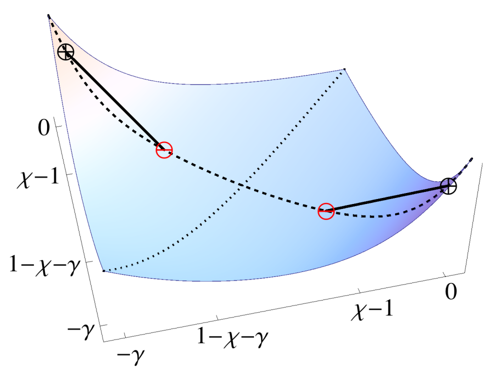

(c) $\boldsymbol{c}(0,-\gamma)+\boldsymbol{c}(-\gamma,0)-\boldsymbol{c}(\chi-1,1-\chi-\gamma)-\boldsymbol{c}(1-\chi-\gamma,\chi-1)>0$

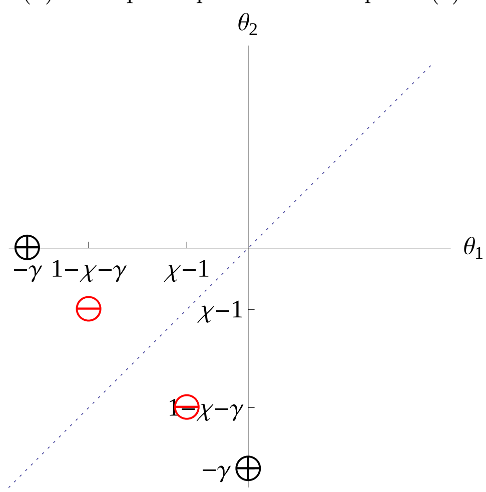

(d) A simpler representation of panel (c)

Figure 1: Exploiting CGF convexity. Panels (a) and (b) illustrate Siegel's paradox.

1c in a more convenient way, so that it is immediately obvious if an expression is positive by virtue of convexity: we require the four points—counted with multiplicity, so that the origin in Figures 1a and 1b counts as two points—to be arranged on any straight line in (θ_1,θ_2)-space, symmetrically about their midpoint, and with sign pattern as in Figures 1b and 1d.

Figure 2a shows how the exchange rate varies as a function of country 1's output share in the calibration described in the previous section. When country 1 is small,

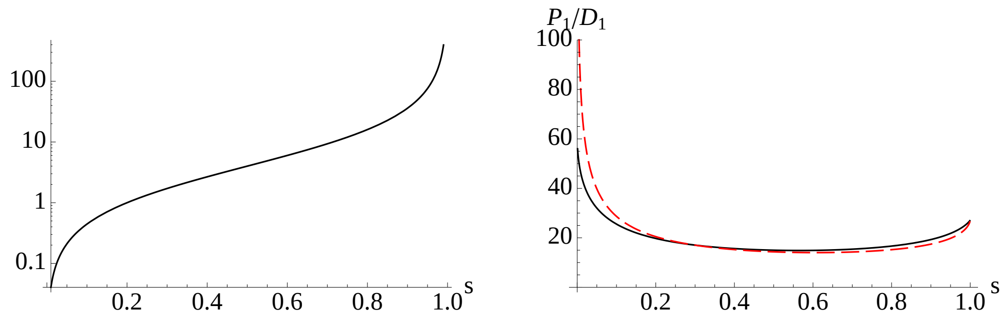

(a) e_t (b) P_1/D_1

Figure 2: Left: The exchange rate—the relative price of good 2 in units of good 1—plotted against s, the output share of country 1. Right: Price-dividend ratio on asset 1, plotted against s, in the imperfect substitution case $\eta = 2$ (black) and the perfect substitution case (dashed red).

its goods are in short supply so command a high price. (In the perfect-substitutes case the relative price would always equal 1, independent of country 1's output share.) Since the weight of country 1 in the consumption aggregator is $w = 0.2$, the exchange rate e_t equals 1 when $s = 0.2$.

The next result provides integral formulas for the price-dividend ratios of the two countries' output claims and for the prices of each country's perpetuity—where, within the model, country i's perpetuity is a security that delivers one unit of good i per unit of time. In the case with no jumps these integrals can be expressed in closed form in terms of hypergeometric functions, as in Martin (2013b). I do not pursue this approach here because the integral formulas apply more generally and can be numerically evaluated almost instantaneously.

Result 2 (Valuation ratios)**.** *The price-dividend ratios, P_{it}/D_{it}, of each country's output claim, and the prices of each country's perpetuity, B_{it}, $i = 1, 2$, are*

$$\begin{aligned} P_{1t}/D_{1t} &= V_{1,0}(u_t) \\ P_{2t}/D_{2t} &= V_{0,1}(u_t) \\ B_{1t} &= V_{1-1/\chi,0}(u_t) \qquad (14) \\ B_{2t} &= V_{0,1-1/\chi}(u_t), \qquad (15) \end{aligned}$$

where

$$V_{\alpha_1,\alpha_2}(u) \equiv \left(e^{u/2}+e^{-u/2}\right)^{\widehat{\gamma}} \int_{-\infty}^{\infty} \frac{e^{iuz}\mathscr{F}(z)}{\rho - \boldsymbol{c}[\chi(\alpha_1 - \widehat{\gamma}/2 - iz), \chi(\alpha_2 - \widehat{\gamma}/2 + iz)]}\, dz$$

and $\mathscr{F}(z) \equiv \frac{1}{2\pi} \cdot \Gamma(\widehat{\gamma}/2 - iz)\Gamma(\widehat{\gamma}/2 + iz)/\Gamma(\widehat{\gamma})$.

The proof of Result 2 shows that we need some assumptions to guarantee that asset prices and expected utility are well defined. Specifically, the discount rate, ρ, must be high enough that $\rho - \boldsymbol{c}\,[\chi(\alpha_1 - \widehat{\gamma}/2), \chi(\alpha_2 - \widehat{\gamma}/2)] > 0$ for a range of values of α_1 and α_2: namely, $\alpha_1 = 1, \alpha_2 = 0$, so that the output claim of country 1 has a finite price; $\alpha_1 = 0, \alpha_2 = 1$, so that the output claim of country 2 has a finite price; $\alpha_1 = 1 - 1/\chi, \alpha_2 = 0$, so that country 1's perpetuity has a finite price; and $\alpha_1 = 0, \alpha_2 = 1 - 1/\chi$, so that country 2's perpetuity has a finite price. The first two of these assumptions—all of which are summarized in Table 1, after rewriting $\widehat{\gamma}$ and χ in terms of γ and η—ensure that expected utility is finite.

Restriction	Reason
$\rho - \boldsymbol{c}(1 - \gamma/2 - 1/(2\eta), -\gamma/2 + 1/(2\eta)) > 0$	finite price of asset 1
$\rho - \boldsymbol{c}(-\gamma/2 + 1/(2\eta), 1 - \gamma/2 - 1/(2\eta)) > 0$	finite price of asset 2
$\rho - \boldsymbol{c}(-\gamma/2 - 1/(2\eta), -\gamma/2 + 1/(2\eta)) > 0$	finite price of perpetuity 1
$\rho - \boldsymbol{c}(-\gamma/2 + 1/(2\eta), -\gamma/2 - 1/(2\eta)) > 0$	finite price of perpetuity 2

Table 1: The restrictions imposed on the model.

Valuation ratios move around over time as dividends, and hence u_t and s_t, move around. Figure 2b plots the price-dividend ratio of the claim to country 1's output stream, P_1/D_1, against country 1's output share, s. The solid line is the price-dividend ratio in the imperfect substitution case, using the same calibration as above, and the dashed line shows the price-dividend ratio in the perfect substitutes case. The price-dividend ratio increases sharply as country 1's share of output declines in both the perfect and imperfect substitution cases, though the effect is muted in the latter case.

Figure 3 shows an illustrative sample realization over a two-year period. Panel 3a plots the paths of exogenous fundamentals: the outputs, or dividends, produced by the two countries. The larger country (black line) initially contributes 80% of global output. After about 0.6 years, it experiences a disaster that causes its output to drop; the smaller country experiences a similar disaster after about 1.5 years.

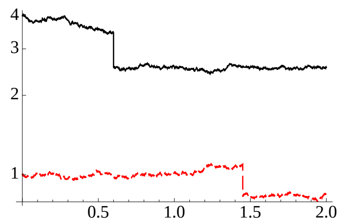

(a) Dividends

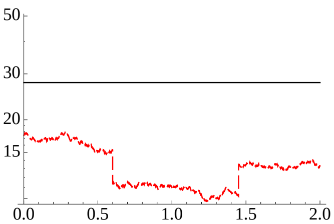

(b) Bond prices in large country's currency, $\eta = 1$

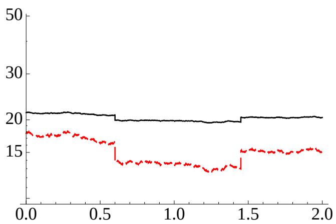

(c) Bond prices in large country's currency, $\eta = 2$

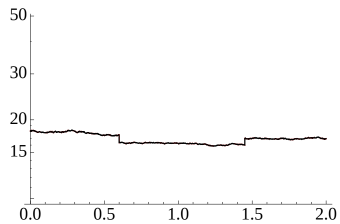

(d) Bond prices in large country's currency, $\eta = \infty$

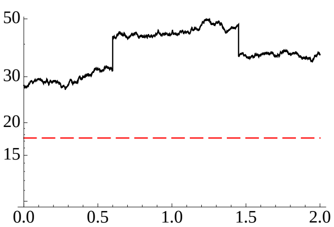

(e) Bond prices in small country's currency, $\eta = 1$

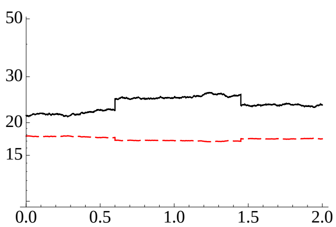

(f) Bond prices in small country's currency, $\eta = 2$

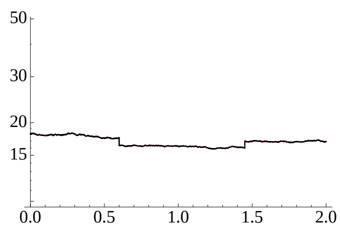

(g) Bond prices in small country's currency, $\eta = \infty$

Figure 3: Prices, in each currency, of the large country's perpetuity (black) and small country's perpetuity (red) along a particular sample path, for different values of the elasticity of substitution between goods, η. Each figure is plotted on a log scale, and is based on the same underlying path of fundamentals, shown in panel (a).

The panels below show how the behavior of perpetuity prices, given by (14) and (15), depends on the elasticity of substitution between goods, η, with the large country's bond in black and the small country's bond in red. The left-hand column sets $\eta = 1$, so that consumption is a Cobb-Douglas aggregator of the two goods. Panels 3b and 3e illustrate the distinctive feature of the Cobb-Douglas setup that bond prices are constant in their own currency, despite the large shocks each country experiences. On the other hand, the exchange rate is extremely volatile, so for example the price of the large country's bond, denominated in the small country's units, jumps up when the large country experiences its output disaster. The right-hand column shows the other extreme, in which the two goods are perfect substitutes. The exchange rate effect disappears: panels 3d and 3g show that bond prices in the perfect substitutes case are the same in each set of units. Put crudely, in the Cobb-Douglas case all the action is in exchange rates and none in valuation ratios, in conflict with the overwhelming empirical evidence that movements in valuation ratios are a major driver of asset returns;[4] and in the perfect substitutes case all the action is in valuation ratios and none in exchange rates.

In between, both effects are present. Panel 3c shows that in large country units, exchange rate movements exacerbate the poor performance of the small country's bond when the large country suffers its disaster. As a result, the small bond is riskier than the large bond, and so the overall level of the small bond's price is lower, reflecting higher interest rates in the small country and hence the emergence of a carry trade. Notice, also, that the carry trade experiences severe underperformance at times of large-country disaster. Panel 3f shows the corresponding plots viewed in small country units. The large country's disaster reduces the relative supply of its good, so its currency appreciates. In small-country units, the large country's bond therefore outperforms at the time of disaster and hence is a hedge, so earns a negative excess return.

The price-dividend ratio of each country's output claim also depends on the relative size of the two economies, as in Cochrane, Longstaff and Santa-Clara (2008) and Martin (2013b). If, say, the large country experiences bad output news then the small country's output share increases; its output claim is now riskier—more correlated with consumption growth—so requires a higher excess return and a lower price-dividend ratio. Thus shocks to one country affect asset valuations in the other

[4]See Cochrane (2008) for a recent survey.

country. Furthermore, if goods are imperfect substitutes, then the small country's currency depreciates when the large country experiences bad output news, as before. This amplifies the underperformance of the small-country output claim in large-country units. I illustrate this in the supplementary appendix, using the same output sample paths as in Figure 3.

2.2 Interest rates

To understand why valuation ratios behave as they do, we must turn to the behavior of interest rates and risk premia. Each good has its own set of zero-coupon bond prices, and attached to these bond prices are zero-coupon yields that move around over time as shocks to the outputs of the two countries induce changes in u. Three measures of interest rates are particularly natural: the riskless rates for each good, calculated from zero-coupon yields in the limit as $T \to 0$; coupon yields on perpetuities, $1/B_{iT}$, provided by equations (14) and (15), which can be viewed as a weighted average of yields of all maturities; and long rates, which are calculated from zero-coupon yields as $T \to \infty$.

Result 3 (Interest rates)**.** *Writing $\mathscr{Y}_{T,i}(u)$ for the continuously compounded T-period zero-coupon yield in i-units when the current state is u, we have*

$$\begin{aligned}
\mathscr{Y}_{T,1}(u) &= \frac{-1}{T}\log\left\{\left(e^{u/2}+e^{-u/2}\right)^{\widehat{\gamma}}\int_{-\infty}^{\infty} e^{iuz}\mathscr{F}(z)e^{-\{\rho-\boldsymbol{c}[\chi(1-1/\chi-\widehat{\gamma}/2-iz),\chi(-\widehat{\gamma}/2+iz)]\}T}\,dz\right\}\\
\mathscr{Y}_{T,2}(u) &= \frac{-1}{T}\log\left\{\left(e^{u/2}+e^{-u/2}\right)^{\widehat{\gamma}}\int_{-\infty}^{\infty} e^{iuz}\mathscr{F}(z)e^{-\{\rho-\boldsymbol{c}[\chi(-\widehat{\gamma}/2-iz),\chi(1-1/\chi-\widehat{\gamma}/2+iz)]\}T}\,dz\right\}.
\end{aligned}$$

The riskless rates, $R_{f,i}(u)=\lim_{T\to 0}\mathscr{Y}_{T,i}(u)$, $i=1,2$, are

$$\begin{aligned}
R_{f,1}(u) &= \left(e^{u/2}+e^{-u/2}\right)^{\widehat{\gamma}}\int_{-\infty}^{\infty} e^{iuz}\mathscr{F}(z)\left\{\rho-\boldsymbol{c}[\chi(1-1/\chi-\widehat{\gamma}/2-iz),\chi(-\widehat{\gamma}/2+iz)]\right\}\,dz\\
R_{f,2}(u) &= \left(e^{u/2}+e^{-u/2}\right)^{\widehat{\gamma}}\int_{-\infty}^{\infty} e^{iuz}\mathscr{F}(z)\left\{\rho-\boldsymbol{c}[\chi(-\widehat{\gamma}/2-iz),\chi(1-1/\chi-\widehat{\gamma}/2+iz)]\right\}\,dz.
\end{aligned}$$

In this notation, the currently prevailing one-year rate in country j, as discussed in Section 1, is $i_{j,t}=\mathscr{Y}_{1,j}(u)$. In contrast to short rates, we will now see that long rates are always constant over time, and often constant across countries too.

Result 4 (The behavior of long interest rates)**.** *Long rates,* $\mathscr{Y}_{\infty,i} \equiv \lim_{T\to\infty} \mathscr{Y}_{T,i}(u)$*, are independent of* u*, hence constant over time:*

$$\mathscr{Y}_{\infty,1} = \max_{\theta\in[0,\gamma+\chi-1]} \rho - \boldsymbol{c}(\theta-\gamma,-\theta) \tag{16}$$

$$\mathscr{Y}_{\infty,2} = \max_{\theta\in[0,\gamma+\chi-1]} \rho - \boldsymbol{c}(-\theta,\theta-\gamma)\,. \tag{17}$$

When the exchangeability property holds, $\mathscr{Y}_{\infty,1} = \mathscr{Y}_{\infty,2} = \rho - \boldsymbol{c}(-\gamma/2,-\gamma/2)$*, so long rates are independent of the elasticity of substitution between goods. Even if exchangeability does not hold, long rates are equated across countries whenever the maximum in* (16) *is attained for* $\theta \in (1-\chi, \gamma+\chi-1)$*.*

Given the result of Dybvig, Ingersoll and Ross (1996), it is not particularly surprising that long rates are constant across time. But it is interesting, and non-obvious, that under mild conditions long rates should be constant across countries, and should be independent of the substitutability of their outputs. For a given CGF, it is easy to specify these "mild conditions" precisely by checking the condition provided in the last sentence of Result 4 explicitly. More generally, long rates will be equated if the countries' output growth processes have sufficiently similar moments and their goods are sufficiently good substitutes.

Figure 4a shows how the riskless rate (black solid line), perpetuity yield (red dashed line), and long rate (blue dotted line) depend on s. Riskless rates are low when s is close to 0 or to 1 due to precautionary savings demand in the face of an unbalanced—because poorly technologically diversified—economy, and high when s is close to 0.5. Country 1's yield curve can be upward-sloping, downward-sloping, or hump-shaped.

The figure is not symmetric: country 1's interest rate is higher when it is small than when it is large. In this example, as country 1's share of global output declines to zero, its interest rate approaches 3.25% while the large country's interest rate drops to 1.18%. From the perspective of an investor (or economist) thinking in large country units, this might suggest the following carry trade: borrow at the large-country interest rate of 1.18%, and invest in the small-country interest rate of 3.25%. We have not yet taken into account the effects of exchange-rate movements, however. Before doing so, recall that in this example the *exchangeability property* holds:

Property 1 (Exchangeability)**.** $\boldsymbol{c}(\theta_1,\theta_2) = \boldsymbol{c}(\theta_2,\theta_1)$ *for all* θ_1, θ_2*.*

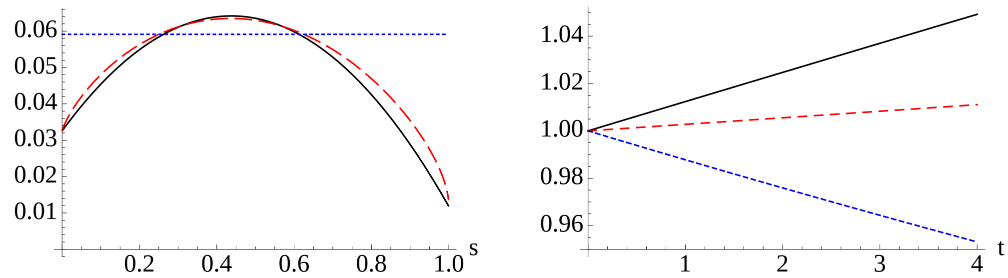

(a) Interest rates plotted against output share

(b) Forward and expected spot exchange rates

Figure 4: Left: Country 1's riskless rate (black solid), perpetuity yield (red dashed), and long rate (blue dotted). Right: Forward price to time t of good 2 in 1-units ($F_{0\to t}$, black solid), expected future spot prices ($\mathbb{E}\, e_t = \mathbb{E}\, 1/e_t$, red dashed), and forward price of good 1 in 2-units ($1/F_{0\to t}$, blue dotted), plotted against t, assuming starting share $s = w = 20\%$.

From Result 1, we know that $(FX_1^* + FX_2^*)/2$, is positive. But the exchangeability property and equations (11) and (12) imply that $FX_1^* = FX_2^*$, so it follows that both are strictly positive. Thus exchange rate movements actually work in favor of the carry trade. The reason for the carry trade's excess return is (of course) that it is risky: if the large country has bad news, the small country's exchange rate deteriorates, and the carry trade has a low return.

Returning to the general case, we can now see how the model generates the failure of uncovered interest parity in the regression (2). We saw in Result 1 that $\log e_t$ is a Lévy process, so $\log e_{t+1} - \log e_t$ is independent of information known at time t, and in particular of $i_{1,t}$ and $i_{2,t}$. Thus $\mathrm{cov}(\log e_{t+1} - \log e_t, i_{1,t} - i_{2,t}) = 0$; combining this with the fact that $\mathrm{var}(i_{1,t} - i_{2,t}) \neq 0$, we have

Result 5 (Failure of UIP)**.** *Interest-rate differentials are totally uninformative about future movements of the exchange rate: plim*$(a_1) = 0$*.*

Given that the random walk nature of $\log e_{t+1} - \log e_t$ was hard-wired in, the interesting aspect of the model is not that it generates $\mathrm{cov}(\log e_{t+1} - \log e_t, i_{1,t} - i_{2,t}) = 0$ but that interest rates vary across countries, $\mathrm{var}(i_{1,t} - i_{2,t}) \neq 0$, despite the random walk character of exchange rates.

The failure of UIP can also be recast in terms of forward rates. Let $F_{0\to t}$ be the time-0 forward price of good 2 in 1-units, for settlement at t. By a no-arbitrage argument, this forward exchange rate is determined by the spot exchange rate and t-period interest rates in the two countries: $F_{0\to t} = e_0 \cdot \exp\{(\mathscr{Y}_{t,1} - \mathscr{Y}_{t,2})t\}$. Figure 4b shows how the forward exchange rates $F_{0\to t}$ (black solid line) and $1/F_{0\to t}$ (blue dotted line) compare to expected future spot exchange rates $\mathbb{E}\, e_t$ and $\mathbb{E}\, 1/e_t$ (red dashed line) in the numerical example. The starting share of country 1 is $s = w = 0.2$, so the current spot exchange rate is $e_0 = 1$. Since the example features symmetric output growth processes, expected future spot exchange rates, shown as a dashed red line, are the same from the perspective of both countries—$\mathbb{E}\, e_t = \mathbb{E}\, 1/e_t$—and they lie above the spot price (Siegel's paradox again—but note its limited *quantitative* importance). The forward price of good 2 is higher than its expected future spot price, while the forward price of good 1 moves in the opposite direction to its expected future spot price, because interest rates are higher in country 1 than in country 2. This is another manifestation of the violation of UIP.

2.3 Expected returns

The expected return, ER, on an asset with price P and instantaneous dividend D is $ER\, dt \equiv \mathbb{E}\, dP/P + (D/P)dt$. This expected return is calculated in the asset's own units. The expected return on asset 2 *in units of country 1* is

$$ER_2^*\, dt \equiv \frac{\mathbb{E}\, d(eP)}{eP} + \frac{eD\, dt}{eP} = \frac{\mathbb{E}\, d(eP)}{eP} + \frac{D\, dt}{P}.$$

The dividend yield component of expected returns is unit-free, but the expected capital gains component depends on exchange rate movements. If these are correlated with asset prices, there will be an associated risk premium.

Result 6 (Expected returns)**.** *Expected returns, $ER_{\alpha_1,\alpha_2,\lambda_1,\lambda_2}(u)$, are given by*

$$ER_{\alpha_1,\alpha_2,\lambda_1,\lambda_2}(u) = \frac{1 + G_{\alpha_1,\alpha_2,\lambda_1,\lambda_2}(u)}{V_{\alpha_1,\alpha_2}(u)}$$

where $G_{\alpha_1,\alpha_2,\lambda_1,\lambda_2}(u)$, and the values of α_i and λ_i (which index the asset and reference units of interest), are provided in the appendix.

Suppose for example that country 1 is small, so that s is close to zero. Figure 5 plots risk premia on the various assets, denominated in (large) country 2's units, against s. Country 1's output claim earns a significantly lower risk premium than

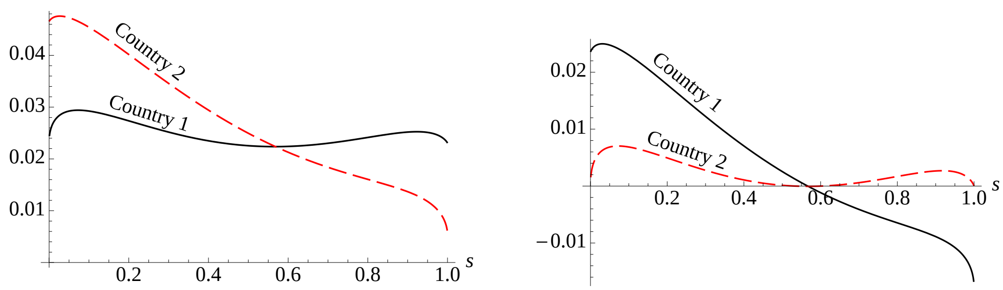

(a) Output claim risk premia

(b) Bond risk premia

Figure 5: Left: Risk premia, in 2-units, on output claims of country 1 (black solid) and country 2 (red dashed). Right: Risk premia, in 2-units, on good-1 bond (black solid) and good-2 bond (red dashed).

country 2's output claim because its output makes up a smaller share of, and is therefore less correlated with, the marginal investor's consumption aggregator. But country 1's bond earns a *higher* risk premium than country 2's bond. This is due to currency risk: with s small, marginal investors particularly dislike states in which (large) country 2 experiences a negative shock to fundamentals. In such states, the *relative* supply of good 1 increases, and hence country 1's currency depreciates. Thus country 1's bond underperforms in bad times, and so requires a higher risk premium. Note, also, that s increases in these bad states, so interest rates rise in both countries and hence bonds underperform even in their own currencies; this explains why country 2's bond earns a positive risk premium in its own units.

This exchange rate effect also operates in the case of output claim risk premia; this explains why country 1's output claim earns a positive risk premium in the limit $s \to 0$ even though its output stream is independent of the larger country's output stream, and hence, in the limit, of consumption growth. In contrast if $\eta = \infty$, as in Cochrane, Longstaff and Santa-Clara (2008) and Martin (2013b), there is no currency risk, so no risk premium, in the limit.

3 The small-country limit

These complicated characterizations of riskless rates, price-dividend ratios, and expected returns simplify considerably in the *small-country limit* in which country

1 is very small and country 2 very large; the ability to take limits of the integral formulas is a major advantage of my analytical approach over the loglinearization approach. Several features of the model emerge more clearly in the limit, and those emphasized in the example above turn out to be characteristic of a whole family of possible calibrations. By continuity, the strict inequalities presented in Results 8, 9 and 10 also hold away from the limit point, so long as country 1 is sufficiently small relative to country 2.

From now on I assume that $\rho - \boldsymbol{c}(\chi, 1-\chi-\gamma) > 0$ so that the price-dividend ratios of the two output claims are finite in the limit, rather than merely for $s \in (0,1)$ as is ensured by previous assumptions. This corresponds to the *subcritical case* considered in Martin (2013b). In the *supercritical case* with $\rho - \boldsymbol{c}(\chi, 1-\chi-\gamma) < 0$, which tends to be associated with a low interest rate environment, the small country's output claim displays in some respects more interesting and unexpected behavior; I rule it out here in the interest of brevity, to focus on the behavior of bonds rather than output claims.

Result 7 (Asset pricing in the small-country limit)**.** *Interest rates are*

$$R_{f,1} = \rho - \boldsymbol{c}(\chi - 1, 1 - \chi - \gamma) \tag{18}$$

$$R_{f,2} = \rho - \boldsymbol{c}(0, -\gamma). \tag{19}$$

Country i's perpetuity earns a risk premium in foreign units, written $XS^*_{B,i}$, *where*

$$XS^*_{B,1} = \boldsymbol{c}(\chi - 1, 1 - \chi) + \boldsymbol{c}(0, -\gamma) - \boldsymbol{c}(\chi - 1, 1 - \chi - \gamma) \tag{20}$$

$$XS^*_{B,2} = \boldsymbol{c}(1 - \chi, \chi - 1) + \boldsymbol{c}(\chi - 1, 1 - \chi - \gamma) - \boldsymbol{c}(0, -\gamma). \tag{21}$$

The dividend yields on the output claims are

$$D_1/P_1 = \rho - \boldsymbol{c}(\chi, 1 - \chi - \gamma) \tag{22}$$

$$D_2/P_2 = \rho - \boldsymbol{c}(0, 1 - \gamma). \tag{23}$$

Excess returns on output claims denominated in own units, XS_i*, are*

$$XS_1 = \boldsymbol{c}(1, 0) + \boldsymbol{c}(\chi - 1, 1 - \chi - \gamma) - \boldsymbol{c}(\chi, 1 - \chi - \gamma) \tag{24}$$

$$XS_2 = \boldsymbol{c}(0, 1) + \boldsymbol{c}(0, -\gamma) - \boldsymbol{c}(0, 1 - \gamma). \tag{25}$$

Excess returns on output claims denominated in foreign units, XS_i^*, are

$$XS_1^* = \boldsymbol{c}(\chi, 1-\chi) + \boldsymbol{c}(0, -\gamma) - \boldsymbol{c}(\chi, 1-\chi-\gamma) \tag{26}$$

$$XS_2^* = \boldsymbol{c}(1-\chi, \chi) + \boldsymbol{c}(\chi-1, 1-\chi-\gamma) - \boldsymbol{c}(0, 1-\gamma). \tag{27}$$

The Gordon growth model holds: $D_i/P_i = XS_i + R_{f,i} - G_i$, where $G_1 \equiv \boldsymbol{c}(1,0)$ and $G_2 \equiv \boldsymbol{c}(0,1)$ are the log mean growth rates of output in each country.

In one sense, asset pricing in the large country is just closed-economy asset pricing: equations (19), (23) and (25) correspond directly to those derived in the one-tree economy of Martin (2013a). But, for example, the risk premium on the large stock market in small-country units, given by equation (27), is a natural object of interest in a multi-country world that has no counterpart in a single closed economy.

The excess return on investment in a foreign country's bond can be decomposed as the sum of an interest-rate differential and an expected currency return, since we can rewrite equations (20) and (21) as $XS_{B,1}^* = FX_1^* + R_{f,1} - R_{f,2}$ and $XS_{B,2}^* = FX_2^* + R_{f,2} - R_{f,1}$.

To put these expressions in more familiar form, suppose that output growth is lognormal, and make the *cet. par.* assumption that the exchangeability property holds. Then $\boldsymbol{c}(\theta_1, \theta_2) = \mu\theta_1 + \mu\theta_2 + \sigma^2\theta_1^2/2 + \kappa\sigma^2\theta_1\theta_2 + \sigma^2\theta_2^2/2$, where μ is the mean, σ the volatility, and κ the cross-country correlation of log output growth in the two countries, and we have

$$\begin{aligned}
R_{f,1} &= \rho + \mu\gamma - \gamma^2\sigma^2/2 + \sigma^2(1-\kappa)(1-\chi)(\gamma+\chi-1) \\
R_{f,2} &= \rho + \mu\gamma - \gamma^2\sigma^2/2 \\
XS_{B,1}^* &= \gamma\sigma^2(1-\kappa)(1-\chi) \\
XS_{B,2}^* &= -(\gamma+2\chi-2)\sigma^2(1-\kappa)(1-\chi) \\
D_1/P_1 &= \rho + \mu(\gamma-1) - \sigma^2(\gamma-1)^2/2 - \sigma^2\chi(1-\kappa)(\gamma+\chi-1) \\
D_2/P_2 &= \rho + \mu(\gamma-1) - \sigma^2(\gamma-1)^2/2 \\
XS_1 &= \gamma\kappa\sigma^2 + \sigma^2(1-\kappa)(1-\chi) \\
XS_2 &= \gamma\sigma^2 \\
XS_1^* &= \gamma\kappa\sigma^2 + \gamma\sigma^2(1-\kappa)(1-\chi) \\
XS_2^* &= \gamma\sigma^2 - (\gamma+2\chi-1)\sigma^2(1-\kappa)(1-\chi).
\end{aligned}$$

In the perfect substitutes case, $\chi = 1$, there is no exchange-rate risk, so interest rates are equal in each country, bonds are riskless, and the small country's equity claim is risky only to the extent that its fundamentals are correlated with the large country's fundamentals. All this changes if the goods are imperfect substitutes ($\chi < 1$). The interest rate is higher in the small than in the large country; the excess return on the small country's bond in large units is positive, while that on the large country's bond in small units is negative; and the small country's dividend yield increases, and its equity risk premium increases—particularly when denominated in foreign units—as exchange-rate risk becomes important.

Without the exchangeability assumption, the signs of most of these risk premia can be set arbitrarily even in the lognormal case (for example, by making the small country's output extremely volatile, and adjusting its correlation with the large country). The only risk premium for which this is not true is that of the large country's equity claim denominated in large units, which is always positive. In the lognormal case this is obvious. Although it is no surprise that $XS_2 > 0$ in general, the proof provides another illustration of how the convexity property of CGFs can be exploited. For $XS_2 = \boldsymbol{c}(0,1) + \boldsymbol{c}(0,-\gamma) - \boldsymbol{c}(0,1-\gamma)$; and note that $(0,1)$ and $(0,-\gamma)$, considered as points in $\mathbb{R}^2$, are a midpoint-preserving-spread of $(0,1-\gamma)$ and $(0,0)$. Convexity therefore implies that $\boldsymbol{c}(0,1)+\boldsymbol{c}(0,-\gamma) > \boldsymbol{c}(0,1-\gamma)+\boldsymbol{c}(0,0)$, and the result follows because $\boldsymbol{c}(0,0) = 0$.

To put some discipline on the model without making strong parametric assumptions, it is helpful to focus attention away from the details of the countries' output processes by assuming that the exchangeability property holds. We then have the following result.

Result 8 (Strong failure of UIP for the small country)**.** *Suppose Property 1 holds. Then the interest rate in the small country is higher than the interest rate in the large country: $R_{f,1} > R_{f,2}$. But this higher interest rate is not offset by expected exchange rate movements. On the contrary, the expected appreciations in the relative price of each country's good are equal, and positive—Siegel's paradox once again. Thus uncovered interest parity (UIP) fails in a strong sense: not only do expected exchange rate movements not fully offset the small country's higher interest rate, they actually* increase *the expected return on the carry trade. That is, $XS^*_{B,1} > FX^*_1 > 0$.*

Proof. We already know that $(FX^*_1 + FX^*_2)/2 > 0$. If Property 1 holds, we must have $FX^*_1 = FX^*_2$. It follows that $FX^*_1 > 0$. In light of the fact that $R_{f,1} - R_{f,2} = XS^*_{B,1} - FX^*_1$, it only remains to prove that $XS^*_{B,1} > FX^*_1$.

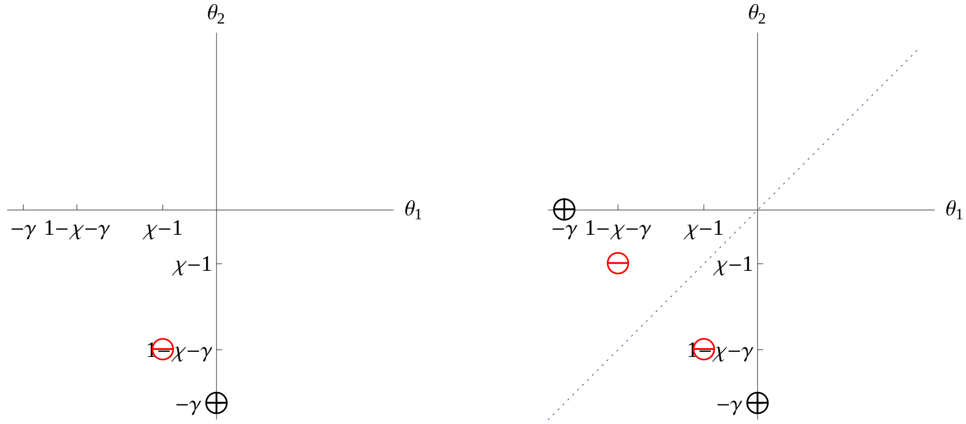

(a) This is positive...   (b) ...by exchangeability & convexity

Figure 6: A visual proof of the strong failure of UIP: $XS^*_{B,1} > FX^*_1$.

From equations (11) and (20), this is equivalent to showing that $\boldsymbol{c}(0,-\gamma)-\boldsymbol{c}(\chi-1,1-\chi-\gamma) > 0$. Panel (a) in Figure 6 represents the left-hand side of this inequality, in the sense that the signed sum of values taken by the CGF at the indicated points (with signs indicated by black plus and red minus signs) is the expression on the left-hand side of the inequality. The aim is to show that the panel is positive, i.e. that the signed sum that it represents is positive. By exchangeability, the desired inequality is equivalent to $\boldsymbol{c}(0,-\gamma)+\boldsymbol{c}(-\gamma,0)-\boldsymbol{c}(\chi-1,1-\chi-\gamma)-\boldsymbol{c}(1-\chi-\gamma,\chi-1) > 0$. Graphically, this corresponds to reflecting in the 45 degree line, as shown in panel (b). Panel (b) is positive by convexity, so the result follows. $\square$

The corresponding result for the large country relies on a property that restricts the behavior of the higher cumulants of output growth.

Property 2 (Convex difference property)**.** *The CGF* $\boldsymbol{c}(\cdot,\cdot)$ *has the* convex difference property (CDP) *if* $\boldsymbol{c}(\theta_1,\theta_2)-\boldsymbol{c}(\theta_1+t,\theta_2+t)$ *is convex in* (θ_1,θ_2) *for all* $t>0$*,* θ_1*, and* θ_2 *such that* (θ_1,θ_2) *and* (θ_1+t,θ_2+t) *lie in the triangle* $\Delta\subset\mathbb{R}^2$ *whose corners are at* $(1,1)$*,* $(1,-\gamma-1)$ *and* $(-\gamma-1,1)$*.*

This property imposes a restriction that neither country has positively skewed log output growth. If, for example, output growth is independent in the two countries, so that $\boldsymbol{c}(\theta_1,\theta_2)$ can be expressed as $\boldsymbol{c}_1(\theta_1)+\boldsymbol{c}_2(\theta_2)$, then it is equivalent to $\boldsymbol{c}_i'''(\theta_i)\le 0$, $i=1,2$, in the triangle Δ. In particular, it implies that $\boldsymbol{c}_i'''(0)\le 0$, i.e. that the third cumulant—skewness—is nonpositive. It is satisfied if output growth is lognormal,

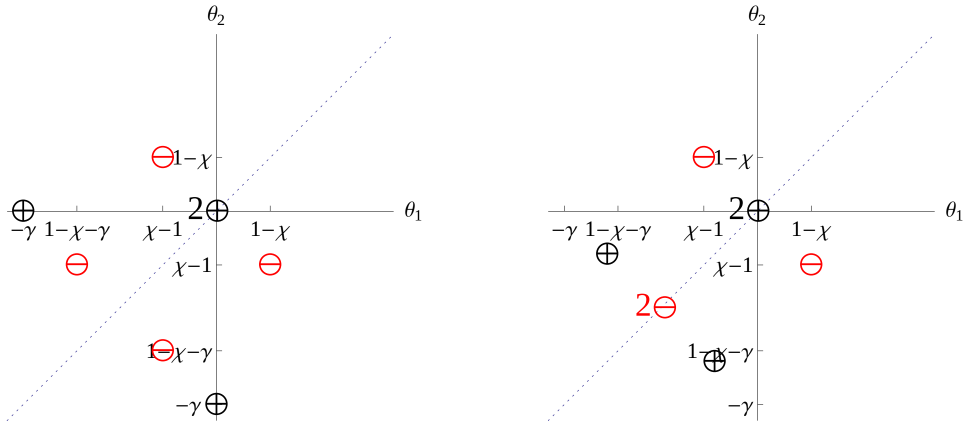

(a) To prove this is positive...   (b) ...use convexity then CDP

Figure 7: The exorbitant privilege: a visual proof that $XS^*_{B,2} < 0$.

in disaster calibrations of the type suggested in Barro (2006), and in the numerical example used in the illustrations. The proofs of the next two results, and specifically Figures 7b and 9b, will reveal why the CDP is formulated as it is.

When exchangeability and the CDP hold, the large country's bond earns a negative risk premium in foreign units—a type of "exorbitant privilege" (Gourinchas and Rey (2007), quoting Valéry Giscard d'Estaing).

Result 9 (An exorbitant privilege)**.** *Suppose Properties 1 and 2 hold. Then UIP also fails for the large country, whose bonds pay a negative risk premium in small-country units:* $XS^*_{B,2} < 0$*.*

Proof. From equation (21), $XS^*_{B,2} < 0$ if and only if $\boldsymbol{c}(0,-\gamma) - \boldsymbol{c}(1-\chi, \chi-1) - \boldsymbol{c}(\chi-1, 1-\chi-\gamma) > 0$. By exchangeability, this is equivalent to showing that $\boldsymbol{c}(0,-\gamma) + \boldsymbol{c}(-\gamma, 0) - \boldsymbol{c}(\chi-1, 1-\chi-\gamma) - \boldsymbol{c}(1-\chi-\gamma, \chi-1) - \boldsymbol{c}(1-\chi, \chi-1) - \boldsymbol{c}(\chi-1, 1-\chi) > 0$.

Panel (a) of Figure 7 represents this inequality graphically. In panel (b), the four southwest-most points have been compressed towards their mid-point. By convexity, this makes the sum smaller. But it remains positive by the CDP, and thus the panel (a) was also positive. [Note: The CDP was formulated as it was precisely so that sign patterns like the one in panel (b) would be positive.] $\square$

Result 8 showed only that the riskless rate is higher in the small country than in the large country. Result 9 is stronger: it can be rephrased as saying that the unfavorable riskless rate differential faced by an investor who borrows at the small country's interest rate and invests at the large country's interest rate is sufficiently large that it overcomes the favorable expected exchange rate movement, i.e. $R_{f,1} - R_{f,2} > FX_2^*$.

3.1 The failure of uncovered equity parity

To characterize the risk premia on the two countries' output claims—and to see the failure of "uncovered *equity* parity"—I make a final assumption that the countries have *linked fundamentals*. In the lognormal case, for example, we want to rule out the possibility that the correlation between the two countries' output growth is negative so that the small country's output claim is a hedge. The following definition is the appropriate one for the general, nonlognormal, case.

Property 3 (Linked fundamentals)**.** *The two countries have linked fundamentals if the CGF is* supermodular,[5] *meaning that for all* $\theta_1, \theta_2, \phi_1, \phi_2$ *in* Δ*,*

$$\boldsymbol{c}(\theta_1, \theta_2) + \boldsymbol{c}(\phi_1, \phi_2) \leq \boldsymbol{c}\left(\max\left\{\theta_1, \phi_1\right\}, \max\left\{\theta_2, \phi_2\right\}\right) + \boldsymbol{c}\left(\min\left\{\theta_1, \phi_1\right\}, \min\left\{\theta_2, \phi_2\right\}\right).$$

By Topkis's (1978) Characterization Theorem, Property 3 holds if

$$\frac{\partial^2 \boldsymbol{c}(\theta_1, \theta_2)}{\partial \theta_1 \partial \theta_2} \geq 0 \tag{28}$$

for all θ_1 and θ_2 in some open set containing Δ. It immediately follows that the linked fundamentals property holds (with equality) if output growth is independent across countries. In any given parametric example, it is easy to check whether (28) holds. In the lognormal case the cumulant-generating function is quadratic in θ_1 and θ_2, so (28) implies that the linked fundamentals property is equivalent to the correlation between the two countries' log output growth being nonnegative. More generally, the easiest and best way to understand what supermodularity imposes is through a diagram: see Figure 8c.

Result 10. *Suppose Properties 1, 2 and 3 hold. Then* $XS_{B,1}^* \leq XS_1^*$*, and there is a critical value* $\eta^* \in (1, \infty)$*—where* $\eta^* = 2$ *in the lognormal case—such that*

$$\begin{aligned} 0 < XS_1 < XS_1^* < XS_2^* < XS_2 \qquad & \textit{if } \eta > \eta^* \\ 0 < XS_1 < XS_2^* < XS_1^* < XS_2 \qquad & \textit{if } \eta < \eta^*. \end{aligned}$$

[5] Vives (1990), Milgrom and Roberts (1990), and Athey (2002) present other applications of supermodularity, notably to games with strategic complementarities.

If the countries have strictly linked fundamentals (i.e. if the inequality in the definition of Property 3 is strict) and η is sufficiently large then we have a total ordering of risk premia: $XS^*_{B,2} < 0 < XS^*_{B,1} < XS_1 < XS^*_1 < XS^*_2 < XS_2$.

Result 10 extends the model's predictions regarding bond risk premia to risky assets. The risk premium on the small country's output claim is greater in foreign units than in own units, $XS^*_1 > XS_1$, while the risk premium on the large country's claim is smaller in foreign units than in own units, $XS^*_2 < XS_2$. The size of η indexes the amount of currency risk. If the goods of the two countries are sufficiently poor substitutes ($\eta < \eta^*$), then currency risk is so great that the risk premium on the small country's output claim in foreign units exceeds the risk premium on the large country's claim in foreign units, $XS^*_1 > XS^*_2$, even though the small country contributes a negligible proportion of the marginal investor's consumption.

4 Conclusion

This paper's basic prediction, for which Hassan (2012) provides empirical support, is that small countries should have higher interest rates than large countries, all else equal. The fundamental asymmetry is that the marginal investor cares more about the large country than the small country, since it provides a larger share of consumption. This means that bad (high-marginal-utility) states are those in which the large country has bad news. When such states occur, the small country's output is in greater relative supply, so its exchange rate depreciates. The small country's higher interest rate is compensation for this risk. Since, by construction, the exchange rate follows a random walk, UIP fails in *any* calibration.

On the methodological side, the paper makes two contributions. It extends the results of Martin (2013b) to allow for imperfect substitution between goods; this is a simple and natural way to generate multiple distinct yield curves in an equilibrium model with perfectly integrated capital markets. It also introduces the exchangeability, convex difference, and linked fundamentals properties and shows how to apply them in diagrammatic proofs. The motivation for doing so is that the quantitative implications of models that feature jumps are sensitively dependent on assumptions made about the extreme tails of the driving stochastic processes. The nonparametric approach taken here offers a novel way to understand the economic mechanism underpinning the model without tying oneself to a particular calibration.

5 References

Alvarez, F., A. Atkeson, and P. J. Kehoe (2007), "If Exchange Rates are Random Walks, Then Almost Everything We Say About Monetary Policy is Wrong," *American Economic Review*, 97:2:339–345.

Alvarez, F., and U. J. Jermann (2005), "Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth," *Econometrica*, 73:6:1977–2016.

Athey, S. (2002), "Monotone Comparative Statics under Uncertainty," *Quarterly Journal of Economics*, 117:1:187–223.

Backus, D. K., M. Chernov, and I. W. R. Martin (2011), "Disasters Implied by Equity Index Options," *Journal of Finance*, 66:6:1969–2012.

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

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1992), "International Real Business Cycles," *Journal of Political Economy*, 100:4:745–775.

Backus, D. K., and G. W. Smith (1993), "Consumption and Real Exchange Rates in Dynamic Economies with Non-Traded Goods," *Journal of International Economics*, 35:297–316.

Bansal, R., and B. N. Lehmann (1997), "Growth-Optimal Portfolio Restrictions on Asset Pricing Models", *Macroeconomic Dynamics*, 1:333–354.

Bansal, R., and I. Shaliastovich (2010), "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," working paper.

Barro, R. J. (2006), "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121:3:823–866.

Brandt, M. W., J. H. Cochrane, and P. Santa-Clara (2006), "International Risk Sharing is Better Than You Think, or Exchange Rates are Too Smooth," *Journal of Monetary Economics*, 53:671–698.

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

Brunnermeier, M. K., S. Nagel, and L. Pedersen (2008), "Carry Trades and Currency Crashes," in Acemoglu, D., K. Rogoff, and M. Woodford (eds.), *NBER Macroeconomics Annual*, 23:1:313–347, University of Chicago Press.

Calvet, L. E., J. Y. Campbell, and P. Sodini (2007), "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes," *Journal of Political Economy* 115:5:707–747.

Campbell, J. Y. and R. J. Shiller (1988), "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1:3:195–228.

Cochrane, J. H. (2008), "Financial Markets and the Real Economy," chapter 7 of Rajnish Mehra, ed., *The Handbook of the Equity Premium*, Elsevier.

Cochrane, J. H., Longstaff, F. A., and P. Santa-Clara (2008), "Two Trees," *Review of Financial Studies*, 21:1:347–385.

Colacito, R., and M. M. Croce (2010), "Risks for the Long Run and the Real Exchange Rate," working paper.

Cole, H. L., and M. Obstfeld (1991), "Commodity Trade and International Risk Sharing: How Much Do Financial Markets Matter?" *Journal of Monetary Economics*, 28:3–24.

Dumas, B., C. R. Harvey, and P. Ruiz (2003), "Are Correlations of Stock Returns Justified by Subsequent Changes in National Outputs?" *Journal of International Money and Finance*, 22:777–811.

Dybvig, P. H., J. E. Ingersoll, and S. A. Ross (1996), "Long Forward and Zero-Coupon Rates Can Never Fall," *Journal of Business*, 69:1:1–25.

Farhi, E., S. P. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan (2009), "Crash Risk in Currency Markets," working paper.

Fama, E. F. (1984), "Forward and Spot Exchange Rates," *Journal of Monetary Economics*, 14:3:319–338.

Frankel, J. A. (1980), "Tests of Rational Expectations in the Forward Exchange Market," *Southern Economic Journal*, 46:4.

Gourinchas, P. O., and H. Rey (2007), "From World Banker to World Venture Capitalist: U.S. External Adjustment and the Exorbitant Privilege," in Richard Clarida, ed., *G7 Current Account Imbalances: Sustainability and Adjustment*, University of Chicago Press, IL.

Hansen, L. P., and R. J. Hodrick (1980), "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, 88:5:829–853.

Hansen, L. P. and R. Jagannathan (1991), "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99:2:225–262.

Hansen, L. P. and J. A. Scheinkman (2012), "Recursive Utility in a Markov Environment with Stochastic Growth," *Proceedings of the National Academy of Sciences*, forthcoming.

Hassan, T. A. (2012), "Country Size, Currency Unions, and International Asset Returns," working paper, Chicago Booth.

Hollifield, B., and A. Yaron (2003), "The Foreign Exchange Risk Premium: Real and Nominal Factors," working paper.

Jurek, J. W. (2009), "Crash-Neutral Currency Carry Trades," working paper, Princeton University.

Lucas, R. E. (1978), "Asset Prices in an Exchange Economy," *Econometrica*, 46:6:1429–1445.

Lustig, H., N. Roussanov, and A. Verdelhan (2009), "Common Risk Factors in Currency Markets," working paper.

Martin, I. W. R. (2011), "Simple Variance Swaps," NBER WP 16884.

Martin, I. W. R. (2013a), "Consumption-Based Asset Pricing with Higher Cumulants," *Review of Economic Studies*, 80:2:745–773.

Martin, I. W. R. (2013b), "The Lucas Orchard," *Econometrica*, 81:1:55–111.

Merton, R. C. (1973), "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41:5:867–887.

Milgrom, P., and J. Roberts (1990), "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58:6:1255–1277.

Piazzesi, M., M. Schneider, and S. Tuzel (2007), "Housing, consumption and asset pricing," *Journal of Financial Economics*, 83:531–569.

Siegel, J. J. (1972), "Risk, interest rates, and the forward exchange," *Quarterly Journal of Economics*, 86:303–309.

Stathopoulos, A. (2009), "Asset Prices and Risk Sharing in Open Economies," working paper, University of Southern California.

Stockman, A. C., and L. L. Tesar (1995), "Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements," *American Economic Review*, 85:1:168–185.

Topkis, D. M. (1978), "Minimizing a Submodular Function on a Lattice," *Operations Research*, 26:305–321.

Verdelhan, A. (2010), “A Habit-Based Explanation of the Exchange Rate Risk Premium,” *Journal of Finance*, 65:1:123–145.

Vives, X. (1990), “Nash Equilibrium with Strategic Complementarities,” *Journal of Mathematical Economics*, 19:305–321.

Zapatero, F. (1995), “Equilibrium Asset Prices and Exchange Rates,” *Journal of Economic Dynamics and Control*, 19:787–811.

A Appendix

Recall that $\chi = (\eta - 1)/\eta$ and $\widehat{\gamma} = (\gamma + \chi - 1)/\chi$, and let $\widehat{y}_{1t} \equiv y_{1t} + [(1-\chi)/\chi]\log w$ and $\widehat{y}_{2t} \equiv y_{2t} + [(1-\chi)/\chi]\log(1-w)$, so that $u_t = \chi(\widehat{y}_{2t} - \widehat{y}_{1t})$. Finally, write $\widetilde{y}_{it} \equiv y_{it} - y_{i0}$. The consumption aggregator can be expressed as

$$C_t = \left[e^{\chi\widehat{y}_{10} + \chi\widetilde{y}_{1t}} + e^{\chi\widehat{y}_{20} + \chi\widetilde{y}_{2t}}\right]^{\frac{1}{\chi}},$$

and the price of good 2 in 1-units is

$$e_t = \left(\frac{1-w}{w}\right)^{1-\chi}\left(\frac{D_{1t}}{D_{2t}}\right)^{1-\chi} = \left(\frac{1-w}{w}\right)^{(1-\chi)/\chi} \cdot e^{-[(1-\chi)/\chi]u_t}.$$

If $\eta = \infty$—the perfect substitutes case—then $\chi = 1$ so e_t is constant. If, on the other hand, $\eta = 1$, then C_t is a Cobb-Douglas aggregator of the two goods so that $C_t \propto D_{1t}^w D_{2t}^{1-w}$. It is easy to check that each asset's price-dividend ratio is then constant. More generally, from (6), good 1's price-dividend ratio is

$$\begin{aligned}
\frac{P_1}{D_{10}} &= \mathbb{E}\int_0^\infty e^{-\rho t}\left(\frac{C_t}{C_0}\right)^{-\chi\widehat{\gamma}}\left(\frac{D_{1t}}{D_{10}}\right)^{\chi} dt \\
&= C_0^{\chi\widehat{\gamma}} \cdot \int_0^\infty e^{-\rho t}\,\mathbb{E}\left(\frac{e^{\chi\widetilde{y}_{1t}}}{\left[e^{\chi\widehat{y}_{10}+\chi\widetilde{y}_{1t}} + e^{\chi\widehat{y}_{20}+\chi\widetilde{y}_{2t}}\right]^{\widehat{\gamma}}}\right) dt,
\end{aligned}$$

and the price of a zero-coupon bond that pays a unit of good 1 at t is

$$\begin{aligned}
Z_{t,1} &= \mathbb{E}\, e^{-\rho t}\left(\frac{C_t}{C_0}\right)^{-\chi\widehat{\gamma}}\left(\frac{D_{1t}}{D_{10}}\right)^{\chi-1} \\
&= C_0^{\chi\widehat{\gamma}} e^{-\rho t} \cdot \mathbb{E}\left(\frac{e^{\frac{\chi-1}{\chi}\cdot\chi\widetilde{y}_{1t}}}{\left[e^{\chi\widehat{y}_{10}+\chi\widetilde{y}_{1t}} + e^{\chi\widehat{y}_{20}+\chi\widetilde{y}_{2t}}\right]^{\widehat{\gamma}}}\right). \qquad (29)
\end{aligned}$$

Integrating over t, we find the perpetuity price

$$B_1 = C_0^{\chi\widehat{\gamma}} \cdot \int_0^\infty e^{-\rho t}\,\mathbb{E}\left(\frac{e^{\frac{\chi-1}{\chi}\cdot\chi\widetilde{y}_{1t}}}{\left[e^{\chi\widehat{y}_{10}+\chi\widetilde{y}_{1t}} + e^{\chi\widehat{y}_{20}+\chi\widetilde{y}_{2t}}\right]^{\widehat{\gamma}}}\right) dt\,.$$

Correspondingly, the price-dividend ratio of asset 2 is

$$\frac{P_2}{D_{20}} = C_0^{\chi\widehat{\gamma}} \cdot \int_0^\infty e^{-\rho t}\,\mathbb{E}\left(\frac{e^{\chi\widetilde{y}_{2t}}}{\left[e^{\chi\widehat{y}_{10}+\chi\widetilde{y}_{1t}}+e^{\chi\widehat{y}_{20}+\chi\widetilde{y}_{2t}}\right]^{\widehat{\gamma}}}\right)dt\,,$$

and the price, in 2-units, of the t-period zero-coupon good-2 bond is

$$Z_{t,2} = C_0^{\chi\widehat{\gamma}} e^{-\rho t}\cdot\mathbb{E}\left(\frac{e^{\frac{\chi-1}{\chi}\cdot\chi\widetilde{y}_{2t}}}{\left[e^{\chi\widehat{y}_{10}+\chi\widetilde{y}_{1t}}+e^{\chi\widehat{y}_{20}+\chi\widetilde{y}_{2t}}\right]^{\widehat{\gamma}}}\right), \tag{30}$$

so the price-dividend ratio of the good-2 perpetuity is

$$B_2 = C_0^{\chi\widehat{\gamma}} \cdot \int_0^\infty e^{-\rho t}\,\mathbb{E}\left(\frac{e^{\frac{\chi-1}{\chi}\cdot\chi\widetilde{y}_{2t}}}{\left[e^{\chi\widehat{y}_{10}+\chi\widetilde{y}_{1t}}+e^{\chi\widehat{y}_{20}+\chi\widetilde{y}_{2t}}\right]^{\widehat{\gamma}}}\right)dt\,.$$

The above six equations each feature an expectation of the form

$$E(\alpha_1,\alpha_2) \equiv \mathbb{E}\left(\frac{e^{\alpha_1\chi\widetilde{y}_{1t}+\alpha_2\chi\widetilde{y}_{2t}}}{\left[e^{\chi\widehat{y}_{10}+\chi\widetilde{y}_{1t}}+e^{\chi\widehat{y}_{20}+\chi\widetilde{y}_{2t}}\right]^{\widehat{\gamma}}}\right), \tag{31}$$

for appropriate values of α_1 and α_2.

Proof of Result 2: Expression (31) can be rewritten

$$E(\alpha_1,\alpha_2) = e^{-\chi\widehat{\gamma}(\widehat{y}_{10}+\widehat{y}_{20})/2}\cdot\mathbb{E}\left(\frac{e^{\chi(\alpha_1-\widehat{\gamma}/2)\widetilde{y}_{1t}+\chi(\alpha_2-\widehat{\gamma}/2)\widetilde{y}_{2t}}}{\left[e^{\chi(\widehat{y}_{20}+\widetilde{y}_{2t}-\widehat{y}_{10}-\widetilde{y}_{1t})/2}+e^{-\chi(\widehat{y}_{20}+\widetilde{y}_{2t}-\widehat{y}_{10}-\widetilde{y}_{1t})/2}\right]^{\widehat{\gamma}}}\right).$$

Now, for $\omega\in\mathbb{R}$ and $\widehat{\gamma}>0$,

$$\frac{1}{\left(e^{\omega/2}+e^{-\omega/2}\right)^{\widehat{\gamma}}} = \int_{-\infty}^{\infty} e^{i\omega z}\mathscr{F}(z)\,dz\,,$$

where i is the complex number $\sqrt{-1}$ and $\mathscr{F}(z)\equiv\frac{1}{2\pi}\cdot\Gamma(\widehat{\gamma}/2-iz)\Gamma(\widehat{\gamma}/2+iz)/\Gamma(\widehat{\gamma})$ is defined in terms of the gamma function (Martin (2013b)). It follows that

$$\begin{aligned}
E(\alpha_1,\alpha_2) &= e^{-\chi\widehat{\gamma}(\widehat{y}_{10}+\widehat{y}_{20})/2}\cdot\mathbb{E}\left(e^{\chi(\alpha_1-\widehat{\gamma}/2)\widetilde{y}_{1t}+\chi(\alpha_2-\widehat{\gamma}/2)\widetilde{y}_{2t}}\int_{-\infty}^{\infty}e^{i\chi(\widehat{y}_{20}+\widetilde{y}_{2t}-\widehat{y}_{10}-\widetilde{y}_{1t})z}\mathscr{F}(z)\,dz\right)\\
&= e^{-\chi\widehat{\gamma}(\widehat{y}_{10}+\widehat{y}_{20})/2}\int_{-\infty}^{\infty}e^{iuz}\mathscr{F}(z)\cdot e^{\mathbf{c}(\chi(\alpha_1-\widehat{\gamma}/2-iz),\chi(\alpha_2-\widehat{\gamma}/2+iz))t}\,dz. \qquad (32)
\end{aligned}$$

The generic expression we want to evaluate is

$$\begin{aligned}
V_{\alpha_1,\alpha_2}(u) &= C_0^{\chi\widehat{\gamma}} \int_0^\infty e^{-\rho t}\,\mathbb{E}\,\frac{e^{\alpha_1\chi\widetilde{y}_{1t}+\alpha_2\chi\widetilde{y}_{2t}}}{\left[e^{\chi(\widehat{y}_{10}+\widetilde{y}_{1t})}+e^{\chi(\widehat{y}_{20}+\widetilde{y}_{2t})}\right]^{\widehat{\gamma}}}\,dt \\
&= \frac{\left[e^{\chi\widehat{y}_{10}}+e^{\chi\widehat{y}_{20}}\right]^{\widehat{\gamma}}}{e^{\chi\widehat{\gamma}(\widehat{y}_{10}+\widehat{y}_{20})/2}} \int_0^\infty e^{-\rho t}\cdot e^{\chi\widehat{\gamma}(\widehat{y}_{10}+\widehat{y}_{20})/2}\cdot E(\alpha_1,\alpha_2)\,dt \\
&= \left[e^{u/2}+e^{-u/2}\right]^{\widehat{\gamma}} \int_{t=0}^\infty \int_{z=-\infty}^\infty e^{-\{\rho-\boldsymbol{c}[\chi(\alpha_1-\widehat{\gamma}/2-iz),\chi(\alpha_2-\widehat{\gamma}/2+iz)]\}t} e^{iuz}\mathscr{F}(z)\,dz\,dt \\
&= \left[e^{u/2}+e^{-u/2}\right]^{\widehat{\gamma}} \int_{-\infty}^\infty \frac{e^{iuz}\mathscr{F}(z)}{\rho-\boldsymbol{c}\left[\chi(\alpha_1-\widehat{\gamma}/2-iz),\chi(\alpha_2-\widehat{\gamma}/2+iz)\right]}\,dz
\end{aligned}$$

using (32). The last equality only holds if Re $\rho-\boldsymbol{c}\left[\chi(\alpha_1-\widehat{\gamma}/2-iz),\chi(\alpha_2-\widehat{\gamma}/2+iz)\right] > 0$ for all $z\in\mathbb{R}$. By Lemma 2 of Martin (2013b), this holds if the (superficially weaker) restriction that $\rho-\boldsymbol{c}\left[\chi(\alpha_1-\widehat{\gamma}/2),\chi(\alpha_2-\widehat{\gamma}/2)\right]>0$ holds; I impose this assumption for the relevant values of α_1 and α_2. □

Proof of Result 3: Using (32) in equations (29) and (30), we have

$$Z_{t,1} = \left(e^{u/2}+e^{-u/2}\right)^{\widehat{\gamma}}\cdot\int_{-\infty}^\infty e^{iuz}\mathscr{F}(z)e^{-\{\rho-\boldsymbol{c}[\chi(1-1/\chi-\widehat{\gamma}/2-iz),\chi(-\widehat{\gamma}/2+iz)]\}t}\,dz$$

and

$$Z_{t,2} = \left(e^{u/2}+e^{-u/2}\right)^{\widehat{\gamma}}\cdot\int_{-\infty}^\infty e^{iuz}\mathscr{F}(z)e^{-\{\rho-\boldsymbol{c}[\chi(-\widehat{\gamma}/2-iz),\chi(1-1/\chi-\widehat{\gamma}/2+iz)]\}t}\,dz.$$

These give the zero-coupon yields immediately; and the riskless rates follow, using l'Hôpital's rule to take the limit as $t\to 0$. □

Proof of Result 4: To calculate long rates, we use the method of steepest descent. The long rate in 1-units is

$$\mathscr{Y}_{\infty,1}(u) = \lim_{T\to\infty} -\frac{1}{T}\log\left\{\int_{-\infty}^\infty e^{iuz}\mathscr{F}(z)e^{-\{\rho-\boldsymbol{c}[\chi(1-1/\chi-\widehat{\gamma}/2-iz),\chi(-\widehat{\gamma}/2+iz)]\}T}\,dz\right\},$$

so we are interested in a stationary point of $\rho-\boldsymbol{c}[\chi(1-1/\chi-\widehat{\gamma}/2-iz),\chi(-\widehat{\gamma}/2+iz)]$, considered as a function of $z\in\mathbb{C}$. If $z=ix$ is pure imaginary, this function is concave when considered as a function of $x\in\mathbb{R}$ (Martin (2013b)), and has a maximum at some ix^*, $x^*\in\mathbb{R}$. If $|x^*|<\widehat{\gamma}/2$ then the contour of integration can be continuously deformed to pass through the stationary point without crossing the pole of $\mathscr{F}(z)$ at $i\widehat{\gamma}/2$. In this case, by the method of steepest descent, $\mathscr{Y}_{\infty,1}(u) = \rho-\boldsymbol{c}[\chi(1-1/\chi-\widehat{\gamma}/2+x^*),\chi(-\widehat{\gamma}/2-x^*)]$. If on the other hand the stationary point occurs for $x^* > \widehat{\gamma}/2$, then we must take the residue at $z = i\widehat{\gamma}/2$ into account: in this case, $\mathscr{Y}_{\infty,1}(u) = \rho - \boldsymbol{c}[\chi - 1, 1 - \chi - \gamma]$. Similarly, if the stationary point occurs at $x^* < -\widehat{\gamma}/2$ then $\mathscr{Y}_{\infty,1}(u) = \rho - \boldsymbol{c}[-\gamma, 0]$. These cases can be summarized by writing $\mathscr{Y}_{\infty,1} = \max_{\theta \in [-\widehat{\gamma}/2, \widehat{\gamma}/2]} \rho - \boldsymbol{c}[\chi(1 - 1/\chi - \widehat{\gamma}/2 + \theta), \chi(-\widehat{\gamma}/2 - \theta)]$, or equivalently, $\mathscr{Y}_{\infty,1} = \max_{\theta \in [0, \gamma+\chi-1]} \rho - \boldsymbol{c}(\theta - \gamma, -\theta)$.

The long rate in 2-units follows by interchanging the roles of countries 1 and 2, which corresponds to switching the arguments of the CGF.

When the exchangeability property holds, it is immediate that long rates are equal. To get the closed-form solution, use exchangeability to rewrite equation (16) as $\mathscr{Y}_{\infty,1} = \max_{\theta \in [0, \gamma+\chi-1]} \rho - \frac{1}{2}\boldsymbol{c}(\theta - \gamma, -\theta) - \frac{1}{2}\boldsymbol{c}(-\theta, \theta - \gamma)$. Differentiating, the maximum is attained at $\theta = \gamma/2$ (which lies in the admissible region, i.e. $\gamma/2 \leq \gamma + \chi - 1$, because $\gamma\eta \geq 2$).

The final statement holds because (17) can be rewritten as

$$\mathscr{Y}_{\infty,2} = \max_{\theta \in [1-\chi, \gamma]} \rho - \boldsymbol{c}(\theta - \gamma, -\theta). \tag{33}$$

Thus if the maximum in (16) is attained for θ in the *open* set $(1 - \chi, \gamma + \chi - 1)$, then it is attained at the same point in (33). $\square$

	Expected returns in 1-units				Expected returns in 2-units			
	α_1	α_2	λ_1	λ_2	α_1	α_2	λ_1	λ_2
tree 1	1	0	1	0	1	0	χ	$1-\chi$
tree 2	0	1	$1-\chi$	χ	0	1	0	1
bond 1	$1-1/\chi$	0	0	0	$1-1/\chi$	0	$\chi-1$	$1-\chi$
bond 2	0	$1-1/\chi$	$1-\chi$	$\chi-1$	0	$1-1/\chi$	0	0

Table 2: Values of $\alpha_1, \alpha_2, \lambda_1, \lambda_2$ for Result 6.

Proof of Result 6: The dividend yield is provided by Result 2. Since

$$\begin{aligned} P &= e^{\lambda_1 \widehat{y}_{1t} + \lambda_2 \widehat{y}_{2t}} \left(e^{\chi(\widehat{y}_{1t} - \widehat{y}_{2t})/2} + e^{\chi(\widehat{y}_{2t} - \widehat{y}_{1t})/2} \right)^{\widehat{\gamma}} \int_{-\infty}^{\infty} \frac{e^{i\chi(\widehat{y}_{2t} - \widehat{y}_{1t})z} \mathscr{F}(z)}{\rho - \boldsymbol{c}[\chi(\alpha_1 - \widehat{\gamma}/2 - iz), \chi(\alpha_2 - \widehat{\gamma}/2 + iz)]} \, dz \\ &= \sum_{m=0}^{\widehat{\gamma}} \binom{\widehat{\gamma}}{m} \int_{-\infty}^{\infty} \frac{\mathscr{F}(z) e^{\widehat{y}_{1t}(\lambda_1 + m\chi - \widehat{\gamma}\chi/2 - i\chi z) + \widehat{y}_{2t}(\lambda_2 - m\chi + \widehat{\gamma}\chi/2 + i\chi z)}}{\rho - \boldsymbol{c}[\chi(\alpha_1 - \widehat{\gamma}/2 - iz), \chi(\alpha_2 - \widehat{\gamma}/2 + iz)]} \, dz \,, \end{aligned}$$

where $\alpha_1, \alpha_2, \lambda_1, \lambda_2$, which vary from asset to asset, are supplied in Table 2, we find that $\mathbb{E}\, dP$ equals

$$\sum_{m=0}^{\widehat{\gamma}} \binom{\widehat{\gamma}}{m} \int \frac{\mathscr{F}(z)\boldsymbol{c}(\lambda_1 + m\chi - \widehat{\gamma}\chi/2 - i\chi z, \lambda_2 - m\chi + \widehat{\gamma}\chi/2 + i\chi z) e^{\widehat{y}_{1t}(\lambda_1+m\chi-\widehat{\gamma}\chi/2-i\chi z)+\widehat{y}_{2t}(\lambda_2-m\chi+\widehat{\gamma}\chi/2+i\chi z)}}{\rho - \boldsymbol{c}[\chi(\alpha_1 - \widehat{\gamma}/2 - iz), \chi(\alpha_2 - \widehat{\gamma}/2 + iz)]} \, dz$$

by using (13). Dividing by P and rearranging, the result follows, after defining $G_{\alpha_1,\alpha_2,\lambda_1,\lambda_2}(u)$ to equal

$$\sum_{m=0}^{\widehat{\gamma}} \binom{\widehat{\gamma}}{m} e^{(\widehat{\gamma}/2-m)u} \int_{-\infty}^{\infty} \frac{e^{iuz}\mathscr{F}(z)\boldsymbol{c}(\lambda_1 + m\chi - \widehat{\gamma}\chi/2 - i\chi z, \lambda_2 - m\chi + \widehat{\gamma}\chi/2 + i\chi z)}{\rho - \boldsymbol{c}[\chi(\alpha_1 - \widehat{\gamma}/2 - iz), \chi(\alpha_2 - \widehat{\gamma}/2 + iz)]} \, dz. \quad \square$$

Proof of Result 7: Follows the proof of Proposition 6 in Martin (2013b). $\square$

Proof of Result 10: Figure 8 establishes that $XS_1 > 0$. Figure 9 shows that $XS_1 < XS_1^*$; the proof that $XS_1 < XS_2^*$ is almost identical. Figure 10, panel (a) shows that $XS_1^* < XS_2$; the proof that $XS_2^* < XS_2$ is almost identical. The final inequality that holds independent of conditions on η is that $XS_{B,1}^* \leq XS_1^*$. This is established in Figure 10, panel (b).

The remaining inequalities are easier to derive algebraically. By exchangeability, the sign of $XS_2^* - XS_1^*$ is the same as the sign of $Q(\chi) \equiv \boldsymbol{c}(\chi-1, 1-\chi-\gamma) + \boldsymbol{c}(1-\chi-\gamma, \chi-1) - \boldsymbol{c}(0, 1-\gamma) - \boldsymbol{c}(1-\gamma, 0) - \boldsymbol{c}(0, -\gamma) - \boldsymbol{c}(-\gamma, 0) + \boldsymbol{c}(\chi, 1-\chi-\gamma) + \boldsymbol{c}(1-\chi-\gamma, \chi)$. Now, $Q(0) = \boldsymbol{c}(-1, 1-\gamma) + \boldsymbol{c}(1-\gamma, -1) - \boldsymbol{c}(0, -\gamma) - \boldsymbol{c}(-\gamma, 0) < 0$ and $Q(1) = -\boldsymbol{c}(0, 1-\gamma) - \boldsymbol{c}(1-\gamma, 0) + \boldsymbol{c}(1, -\gamma) + \boldsymbol{c}(-\gamma, 1) > 0$ by strict convexity of $\boldsymbol{c}(\cdot, \cdot)$. Note that $Q(\chi)$ is continuous and strictly convex, so $Q(\chi^*) = 0$ for some unique $\chi^* \in (0, 1)$. Defining η^* via $\chi^* = 1 - 1/\eta^*$, we have $XS_2^* = XS_1^*$ if $\eta = \eta^*$; $XS_2^* < XS_1^*$ for $\eta < \eta^*$; and $XS_2^* > XS_1^*$ for $\eta > \eta^*$.

It remains to be shown that $XS_{B,1}^* < XS_1$ if the CGF is strictly supermodular and η is sufficiently large; equivalently, $\boldsymbol{c}(\chi - 1, 1 - \chi) + \boldsymbol{c}(0, -\gamma) - 2\boldsymbol{c}(\chi - 1, 1 - \chi - \gamma) - \boldsymbol{c}(1, 0) + \boldsymbol{c}(\chi, 1 - \chi - \gamma) < 0$. In the case $\chi = 1$, we must show that $\boldsymbol{c}(0, -\gamma) + \boldsymbol{c}(1, 0) - \boldsymbol{c}(1, -\gamma) > 0 = \boldsymbol{c}(0, 0)$. This follows by strict supermodularity. Therefore by continuity, $XS_{B,1}^* < XS_1$ for χ in some neighborhood of 1; equivalently, $XS_{B,1}^* < XS_1$ for sufficiently large η. $\square$

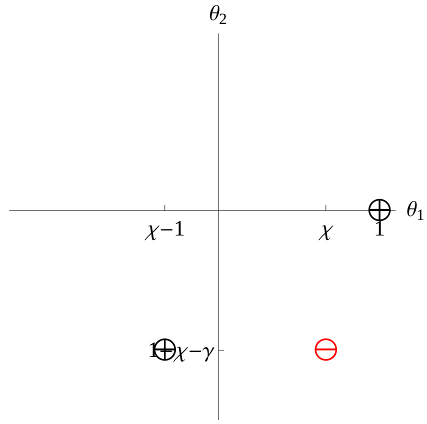

(a) This is positive because. . .

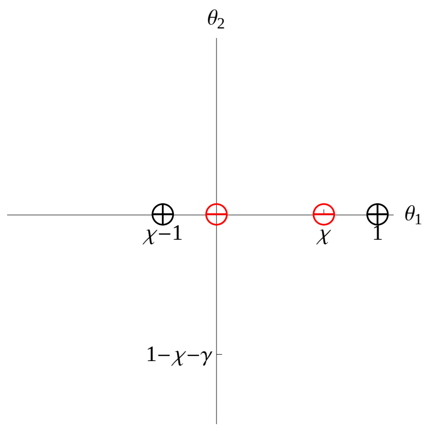

(b) . . . this is, by convexity. . .

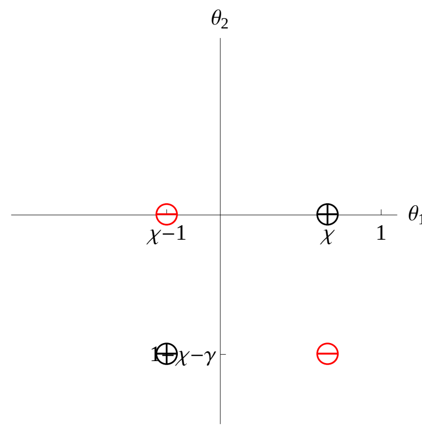

(c) . . . and this is, by supermodularity

Figure 8: The risk premium on the small country's output claim is positive in own units, $XS_1 > 0$. The top panel, which represents XS_1, is the sum of the bottom two panels. Now, panel (b) is strictly positive by convexity of the CGF and panel (c) is nonnegative by supermodularity. (The CGF is zero at the origin, so the top panel really is the sum of the bottom two panels.)
Note: panel (c) illustrates what supermodularity "is", namely the property that every rectangle whose edges are parallel to the axes and that has this sign pattern is positive.

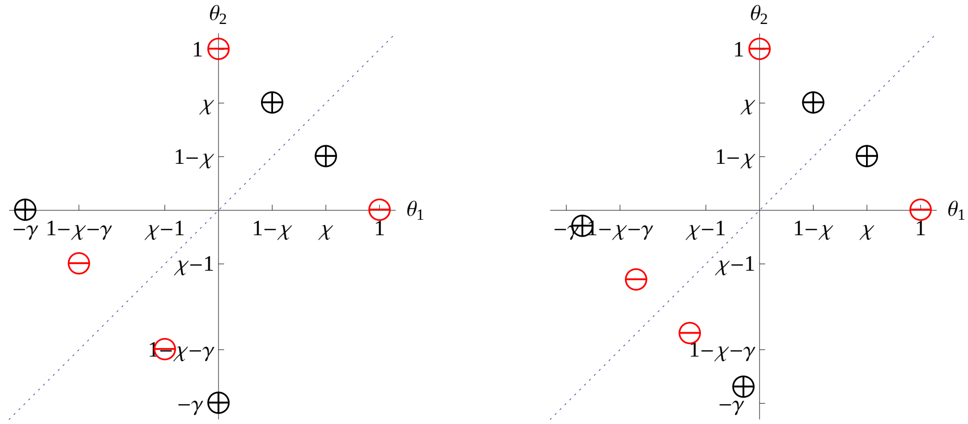

(a) To prove this is positive. . .

(b) . . . use convexity then CDP

Figure 9: The risk premium on the small country's output claim is higher in foreign units than it is in own units, $XS_1 < XS_1^*$. The logic is as in Figure 7. (An almost identical argument shows that $XS_1 < XS_2^*$.)

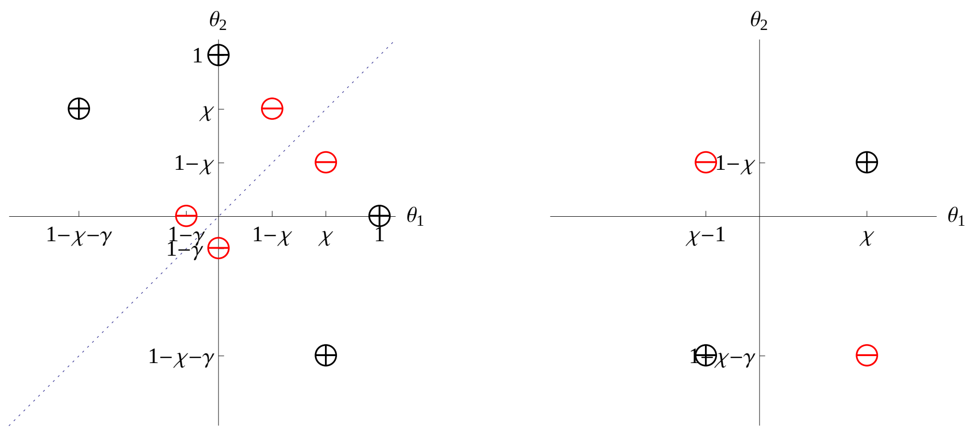

(a) Positive by convexity

(b) Positive by supermodularity

Figure 10: Left: The risk premium on the large country's output claim, in large-country units, is higher than the risk premium on the small country's output claim, in large-country units, $XS_1^* < XS_2$. The two groups of points are both positive by convexity of the CGF. An almost identical argument shows that $XS_2^* < XS_2$.
Right: In large-country units, the small country's output claim requires a (weakly) higher risk premium than its bond, $XS_{B,1}^* \leq XS_1^*$. The result follows by supermodularity.