

Household Heterogeneity and Real Exchange Rates

Narayana R. Kocherlakota*

University of Minnesota and Federal Reserve Bank of Minneapolis

Luigi Pistaferri

Stanford University and CEPR

February 19, 2007

Abstract

We assume that individuals can fully insure themselves against cross-country shocks, but not against individual-specific shocks. We consider two particular models of limited risk-sharing: domestically incomplete markets (DI) and private information-Pareto optimal (PIPO) risk-sharing. For each model, we derive a restriction relating the cross-sectional distributions of consumption and real exchange rates. We evaluate these restrictions using household-level consumption data from the US and the UK. We show that the PIPO restriction fits the data well when households have a coefficient of relative risk aversion of around 5. The restrictions implied by the complete risk-sharing model and the DI model fare poorly.

*This paper was prepared for the *Economic Journal* Lecture at the 2006 Royal Economic Society meetings. Kocherlakota thanks Winnie Choi for many stimulating conversations about real exchange rates. Kocherlakota acknowledges the support of NSF 0305833 and 0606695. We thank an associate editor of the journal, George Alessandria and Ivan Werning for helpful comments, Jaromir Nosal and Sonam Sherpa for research assistance, Orazio Attanasio and Andrew Leicester for help with the FES data, and Joan Gieseke and Hakki Yazici for editorial assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

It is increasingly common in international economics to use *incomplete markets* models. In these models, countries are affected by shocks, but asset markets are insufficiently rich to allow agents to insure themselves against these shocks. In the extreme, agents within the model can only trade a risk-free bond.

Virtually all of these models are *representative agent* models, in which agents are assumed to be identical within a country. This treatment creates a peculiar asymmetry. On the one hand, the identicalness of agents within a country implicitly means that they are able to insure themselves fully against person-specific shocks, like unemployment or disability. On the other hand, the incompleteness of markets means that these same agents are unable to trade assets that pay off contingent on country-specific shocks – shocks that are typically fully observable and non-manipulable.

In this paper, we take the opposite approach. We assume that markets are complete with respect to country-specific shocks, so that agents are able to insure themselves fully against these aggregate shocks. However, we assume that they are only able to partially insure themselves against individual-specific skill shocks. This limited insurance creates ex-post heterogeneity across individuals and means that there is no sense in which there is a representative agent within each country.

We use this approach to address a particular anomaly: the Backus and Smith (1993) real exchange rate puzzle. Backus and Smith (1993) consider a world in which there is a representative agent in each country, only some

goods are tradeable, and financial markets are complete. (They impose no restrictions on firm price-setting behavior.) In this setting, they prove that the growth rate in the real exchange rate between any two countries should equal the difference in the growth rates of marginal utilities of the representative consumers. Crucially, this equality is not just true on average, but rather in *all* dates and states. They document that this implication is strongly falsified in the data. In particular, real exchange rates are much more volatile and persistent than the corresponding ratios of marginal utilities.¹

To address the Backus-Smith puzzle, we study an economy in which agents are faced with skill shocks that affect their ability to transform effort into output. The agents are able to insure themselves fully against country-specific shocks, but only partially insure themselves against their individual-specific skill shocks. We consider two forms of partial insurance against these shocks. The first we call *domestically incomplete* markets (*DI*). Under this form of partial insurance, agents sequentially trade assets and goods. The assets consist of claims that pay off contingent on aggregate shocks to their country or any other. Hence, the agents in the model are unable to directly insure themselves against individual-specific shocks.

The second we call *private information, Pareto optimal* (*PIPO*). Under this formulation, agents are able to sign insurance contracts that completely

¹Colacito and Croce (2005) modify the Backus-Smith model in two ways: they allow preferences to be of the Epstein-Zin (1989) form, and they incorporate a small, highly persistent component in consumption growth rates. They do not directly evaluate their model in terms of the Backus-Smith puzzle, but do argue that the model does a better job in explaining other aspects of real exchange rate data.

insure themselves against their skill shocks, subject to the incentive compatibility constraint generated by the restriction that the realizations of the shocks are private information. The insurance companies are then allowed to trade on the agents' behalf in the sequential asset/goods markets described above.

Our theoretical results parallel those of Kocherlakota and Pistaferri (2005). As in that paper, we are able to derive our key findings while imposing only minimal assumptions on the stochastic process generating individual-level or country-level shocks. Let γ be the coefficient of relative risk aversion of the various agents, let e_t^{jk} be the real exchange rate between country j and k , and let $C_{\eta t}^j$ be the η th moment of the cross-sectional distribution of consumption in country j . (Throughout the paper, the term “moment” refers to non-central moments.) We prove that in the DI setup:

$$\Delta \ln e_t^{jk} = \Delta \ln C_{-\gamma t}^k - \Delta \ln C_{-\gamma t}^j, \quad (\text{DI})$$

where Δ is the first-difference operator. This formula says that the growth rate of the real exchange rate between country j and k equals the difference in the growth rates between country j and k of the negative γ -th moments of their cross-sectional consumption distributions. In the PIPO setup, the growth rate of the real exchange rate between country j and k equals the difference in the growth rates between country k and j of the γ -th moments of their cross-sectional consumption distributions

$$\Delta \ln e_t^{jk} = \Delta \ln C_{\gamma t}^j - \Delta \ln C_{\gamma t}^k. \quad (\text{PIPO})$$

Note that if the cross-sectional consumption distribution is degenerate, then both PIPO and DI are equivalent to Backus and Smith's characterization.

We next use household-level data on consumption expenditures from the United States (the Consumer Expenditure Survey (CEX)) and the United Kingdom (the Family Expenditure Survey (FES)) to assess the validity of these restrictions. A key difficulty here is that we only have a sample of households, not the full population. This means that, unlike Backus and Smith (1993), we cannot directly compare the growth rate of the real exchange rates to the difference in the growth rates of the various moments of the consumption distribution. We use the following procedure. We create an error term by subtracting the difference in the growth rates of the appropriate moments of the consumption distributions from the growth rate of the real exchange rate. We regress this error term on the growth rate of the real exchange rate. Under the null hypothesis, the error term is due only to cross-sectional sampling error and the regression coefficient should be zero.

We examine this implication for the representative agent model, the DI model, and the PIPO model for various values of the coefficient of relative risk aversion. The regression coefficient is non-zero, and is in fact close to 1, for realistic parameterizations of the representative agent model and the incomplete markets models. However, the regression coefficient is zero when we use the PIPO model with γ approximately equal to 5.

We conclude that the PIPO model is able to account for movements in real exchange rates, at least if we set γ to be approximately 5. It is

interesting to compare this result with our finding for U.S. asset pricing data in Kocherlakota and Pistaferri (2005). There, we find that the PIPO model does a good job of accounting for the equity premium when γ is also around 5, while the other two models are not able to account for the equity premium. The similarity of the findings strikes us as highly suggestive.

The Backus-Smith puzzle has received a great deal of attention recently. Chari, Kehoe, and McGrattan (2002) explore the ability of an incomplete markets model with a representative agent in each country to explain the puzzle. They show that their model significantly overpredicts the correlation between relative consumptions and the real exchange rate. More recently, Corsetti, Dedola, and Leduc (2005) have augmented the model of Chari et al. (2002) by assuming that there are nontrivial distribution costs associated with non-traded goods. In contrast to Chari et al., Corsetti et al. show that their model is able to rationalize the low correlation between relative consumptions and real exchange rates that is so puzzling from the perspective of the Backus-Smith complete markets setup.² In both of these papers, the authors assume that agents can only trade risk-free bonds across countries.

We view our paper as complementary to Chari et al. (2002) and Corsetti et al. (2005). As we have stressed, our modeling approach is quite different from theirs. However, we also differ in methodology. These papers solve for a full dynamic stochastic general equilibrium, given an exogenous specification of productivity shocks, and compare some aspects of this equilibrium to the

²See also Benigno and Thoenissen (2004) for a similar approach.

data. We undertake a different kind of exercise. We assume that household-level consumption is consistent with the implications of our model(s), and then, conditional on that assumption, evaluate the models' implications for real exchange rates. Hence, we do not evaluate the ability of either model to generate the observed cross-sectional distribution of consumption in the United States or United Kingdom.

1 A Multi-Country Environment

Consider an economy with J countries and $(T + 1)$ periods. Time is indexed $0, 1, 2, \dots, T$; period 0 is a contracting/trading period in which no production or consumption takes place. Each country has a unit measure of agents. A typical agent has expected utility preferences, with cardinal utility function

$$\sum_{t=1}^T \beta^{t-1} \{ [u(c_t^{TR}, c_t^{NT})]^{1-\gamma} / (1-\gamma) - v(l_t^{TR}, l_t^{NT}) \}, 0 < \beta < 1, 0 < \gamma$$

Here, c_t^{TR} is the agent's period t consumption of a tradeable good that is available in all countries and c_t^{NT} is the agent's consumption of a non-tradeable good specific to his country.³ The variables l_t^{TR} and l_t^{NT} represent the agent's labor in the tradeable goods and non-tradeable goods sectors, respectively. We also assume that the partial derivatives of u are positive, that the own second derivatives of u are negative, and that u is homogeneous of degree one.

³We follow Backus and Smith (1993) in considering an economy with one costlessly tradeable good and a non-tradeable good. However, we could readily extend our analysis to allow for the presence of other goods which are tradeable, but only by paying a stochastic good-specific iceberg cost.

There are two kinds of shocks in the economy: public aggregate shocks and private idiosyncratic shocks. The first kind of shocks works as follows. Let Z be a finite set, and let Ψ be a probability density over Z^T that assigns positive probability to all elements of Z^T . At the beginning of period 1, an element z^T of Z^T is drawn according to Ψ . The random vector z^T is the sequence of public aggregate shocks; z_t is the realization of the shock in period t .

The idiosyncratic shocks work as follows in country j . Let Θ be a finite set, and let π_j be a probability density defined over Θ^T . At the beginning of period 1, an element of Θ^T is drawn for each agent according to the density π_j . Conditional on z^T , the draws are independent across agents and π_j is the same for all realizations of z^T ; we require $\pi_j(\theta^T) > 0$ for all θ^T in Θ^T . We assume that a law of large numbers applies across agents: conditional on any z^T , the measure of agents in the population of country j with type θ^T is given by $\pi_j(\theta^T)$.

Any given agent learns the realization of z_t and his own θ_t at the beginning of period t and not before. Thus, at the beginning of period t , the agent knows his own private history $\theta^t = (\theta_1, \dots, \theta_t)$ and the history of public shocks $z^t = (z_1, \dots, z_t)$. This implies that his choices in period t can only be a function of this history.

The individual-specific and aggregate shocks jointly determine skills. In period t , an agent in country j produces output y_t^i in sector $i = TR, NT$

according to the function

$$y_t^i = \phi_t^{ij}(\theta^t, z^t)l_t^i \quad (1)$$

$$\phi_t^{ij} : \Theta^t \times Z^t \rightarrow (0, \infty) \quad (2)$$

We assume that an agent's output in either sector is observable at time t , but his labor input is known only to him. We refer to $\phi_t^{i,j}(\theta^t, z^t)$ as a country j agent's skill in sector i in history (θ^t, z^t) .

We define an allocation in this society to be $(c, y) = (((c_t^{ij}, y_t^{ij})_{i \in \{TR, NT\}})_{j=1}^J)_{t=1}^T$ where

$$c_t^{ij} : \Theta^t \times Z^t \rightarrow R_+ \quad (3)$$

$$y_t^{ij} : \Theta^t \times Z^t \rightarrow [0, \bar{y}] \quad (4)$$

Here, $y_t^{ij}(\theta^t, z^t)$ is the amount of sector i output that an agent in country j with shock history θ^t produces, given that the public shock history is z^t . Similarly, $c_t^{ij}(\theta^t, z^t)$ is the amount of sector i goods consumption for an agent in country j with shock history θ^t , given that the public shock sequence is z^t . We define an allocation (c, y) to be *feasible* if for all t, z^t :

$$\sum_{j=1}^J \sum_{\theta^t \in \Theta^t} c_t^{TR,j}(\theta^t, z^t) \pi_j(\theta^t) \leq \sum_{j=1}^J \sum_{\theta^t \in \Theta^t} y_t^{TR,j}(\theta^t, z^t) \pi_j(\theta^t) \quad (5)$$

$$\sum_{\theta^t \in \Theta^t} c_t^{NT,j}(\theta^t, z^t) \pi_j(\theta^t) \leq \sum_{\theta^t \in \Theta^t} y_t^{NT,j}(\theta^t, z^t) \pi_j(\theta^t) \text{ for all } j \quad (6)$$

Because θ_t is only privately observable, allocations must respect incentive-compatibility conditions. (The following definitions correspond closely to those in Golosov, Kocherlakota and Tsyvinski (2003).) A *reporting strategy*

$\sigma : \Theta^T \times Z^T \rightarrow \Theta^T \times Z^T$, where σ_t is (θ^t, z^t) -measurable and $\sigma(\theta^T, z^T) = (\theta^{T'}, z^{T'})$. An agent can alter his consumption and labor by changing his reporting strategy. Suppose the allocation is (c, y) . If an agent in country j uses reporting strategy σ , his consumption of the two types of goods in history (θ^t, z^t) is

$$(c_t^{TR,j}(\sigma^t(\theta^t, z^t)), c_t^{NT,j}(\sigma^t(\theta^t, z^t))),$$

where $\sigma^t(\theta^t, z^t) = (\sigma_s(\theta^s, z^s))_{s=1}^t$, and his labor is

$$\left(\frac{y_t^{TR,j}(\sigma^t(\theta^t, z^t))}{\phi_t^{TR,j}(\theta^t, z^t)}, \frac{y_t^{NT,j}(\sigma^t(\theta^t, z^t))}{\phi_t^{NT,j}(\theta^t, z^t)} \right).$$

Let $V_j(\sigma; c, y)$ be the ex-ante utility that a country j agent receives from using strategy σ . Let σ_{TT} be the truth-telling strategy $\sigma_{TT}(\theta^T, z^T) = (\theta^T, z^T)$ for all θ^T, z^T . Then, an allocation (c, y) is *incentive-compatible* if

$$V_j(\sigma_{TT}; c, y) \geq V_j(\sigma; c, y) \text{ for all } j \text{ and all } \sigma.$$

An allocation which is incentive-compatible and feasible is said to be incentive-feasible.

2 Real Exchange Rates

Given this definition of the environment, we now examine the properties of real exchange rates in two different trading setups. In the first setup, agents trade a complete set of z^t -contingent securities directly with one another. In the second setup, intermediaries trade a complete set of z^t -contingent securities on behalf of the agents. In this latter formulation, the intermediaries directly insure the agents against the realization of the θ shocks, given

that they are private information. We derive useful necessary conditions of equilibrium in both settings.

In both setups, we use the following notation. For all t, z^t , let $p_t^{TR}(z^t)$ be the price of tradeable goods in public history z^t , and let $p_t^{NT,j}(z^t)$ be the price of non-tradeable goods in public history z^t in country j . Both prices are denominated in terms of a common numeraire (for concreteness, suppose they are quoted in dollars). A crucial element of our analysis is the notion of a *price index*. Given price processes $(p^{TR}, p^{NT,j})$ for the tradeable and non-tradeable goods, define the price index process \bar{p}_t^j to be the minimal amount of expenditure necessary to achieve one utile in history z^t :

$$\begin{aligned} \bar{p}_t^j(z^t) &= \min_{c_t^{TR}, c_t^{NT}} p_t^{TR}(z^t)c_t^{TR} + p_t^{NT,j}(z^t)c_t^{NT} \\ &\text{s.t. } u(c_t^{TR,j}, c_t^{NT,j}) = 1 \end{aligned}$$

We can then define the real exchange rate between country j and country k to be the ratio of the price indices for the two countries:

$$e_t^{jk}(z^t) = \frac{\bar{p}_t^k(z^t)}{\bar{p}_t^j(z^t)}$$

(Note that if there were no non-tradeable goods, then the real exchange rate would always be one.)

The following fact is useful. Suppose the pair $(c_t^{TR,j}(z^t, W), c_t^{NT,j}(z^t, W))$ solves the static budget problem

$$\begin{aligned} \max_{a,b} u(c_t^{TR,j}, c_t^{NT,j}) \\ \text{s.t. } p_t^{TR}(z^t)c^{TR,j} + p^{NT,j}(z^t)c^{NT,j} = W \end{aligned}$$

for some W . Because u is homogeneous of degree one, it is straightforward to show, using Euler's Theorem, that

$$\bar{p}_t^j(z^t)u(c_t^{TR,j}(z^t, W), c_t^{NT,j}(z^t, W)) = W. \quad (7)$$

Thus, an agent's total expenditure equals his momentary utility multiplied by the price index.

2.1 *Domestically Incomplete Markets*

Suppose that at each date $t = 0, 1, 2, \dots, (T-1)$, all agents can trade a complete set of z^{t+1} -contingent securities that pay off in terms of the numeraire good. (The date 0 is added simply to allow agents to trade before the realization of uncertainty in period 1.) Fix agents' labor choices at their individually optimal levels. Then, the choice problem of a typical agent in country j is

$$\begin{aligned} & \max_{c^{TR,j}, c^{NT,j}, W^j} E_0 \sum_{t=1}^T \beta^{t-1} [u(c_t^{TR,j}, c_t^{NT,j})]^{1-\gamma} / (1-\gamma) \\ & \text{s.t. } p_t^{TR}(z^t) c_t^{TR,j}(\theta^t, z^t) + p_t^{NT,j}(z^t) c_t^{NT,j}(\theta^t, z^t) + \sum_{z^{t+1} \geq z^t} q_{t+1}(z^{t+1}|z^t) W_{t+1}^j(\theta^t, z^{t+1}) \\ & \leq p_t^{TR}(z^t) y_t^{TR,j}(\theta^t, z^t) + p_t^{NT,j}(z^t) y_t^{NT,j}(\theta^t, z^t) + W_t^j(\theta^t, z^t) \\ & \sum_{z_1 \in Z} q_1(z_1) W_1^j(z_1) \leq 0 \\ & W_T^j(\theta^T, z^T) \geq 0 \text{ for all } (\theta^T, z^T). \end{aligned}$$

Here, W_{t+1}^j represents the agent's state contingent purchases and $q_{t+1}(z^{t+1}|z^t)$ is the history z^t price of a claim to consumption in history z^{t+1} . The static

first-order necessary conditions for this problem are

$$\beta^{t-1}\pi_j(\theta^t)\Psi(z^t)u_t^j(\theta^t, z^t)^{-\gamma}u_{TR,t}^j(\theta^t, z^t) = \lambda_t^j(\theta^t, z^t)p_t^{TR}(z^t) \quad (8)$$

$$\beta^{t-1}\pi_j(\theta^t)\Psi(z^t)u_t^j(\theta^t, z^t)^{-\gamma}u_{NT,t}^j(\theta^t, z^t) = \lambda_t^j(\theta^t, z^t)p_t^{NT,j}(z^t) \quad (9)$$

Here, $\lambda_t^j(\theta^t, z^t)$ is the multiplier on the history (θ^t, z^t) flow budget constraint and $u_t^j(\theta^t, z^t) = u(c_t^{TR,j}(\theta^t, z^t), c_t^{NT,j}(\theta^t, z^t))$. As well, $u_{i,t}^j(\theta^t, z^t)$ is the marginal utility of an agent in country j with respect to good i ($i = TR, NT$) in history (θ^t, z^t) . It is useful to renormalize the multipliers; define

$$\widehat{\lambda}_t^j(\theta^t, z^t) = \lambda_t^j(\theta^t, z^t)/[\pi_j(\theta^t)\Psi(z^t)]$$

As we shall see, our data contain observations on expenditures for each household and on the price index in each country. For each household with history (θ^t, z^t) , define its real consumption to be its expenditures divided by the price index in its country:

$$c_t^j(\theta^t, z^t) \equiv \frac{p_t^{TR}(z^t)c_t^{TR,j}(\theta^t, z^t) + p_t^{NT,j}(z^t)c_t^{NT,j}(\theta^t, z^t)}{\bar{p}_t^j(z^t)}$$

Note that (7) implies that the household's real consumption $c_t^j(\theta^t, z^t)$ equals its momentary utility $u_t^j(\theta^t, z^t)$.

Using (8) and (9), we find that

$$\beta^{t-1}c_t^j(\theta^t, z^t)^{1-\gamma} = \widehat{\lambda}_t^j(\theta^t, z^t)\bar{p}_t^j(z^t)c_t^j(\theta^t, z^t)$$

Integrating over θ^t , we get

$$\bar{p}_t^j(z^t) = \beta^{t-1} \frac{E\{c_t^{j-\gamma}|z^t\}}{E\{\widehat{\lambda}_t^j|z^t\}}$$

(because $E\{c_t^{j-\gamma}|z^t\}$ is the integral of $c_t^{j-\gamma}$ with respect to θ^t , given z^t). All together, this implies that the real exchange rate equals

$$e_t^{jk}(z^t) = \frac{E\{c_t^{k-\gamma}|z^t\} E\{\widehat{\lambda}_t^j|z^t\}}{E\{c_t^{j-\gamma}|z^t\} E\{\widehat{\lambda}_t^k|z^t\}}$$

This particular restriction on the behavior of the real exchange rate is not econometrically useful, because the multipliers are not observable. As in Backus and Smith (1993), though, the assumption of complete markets (with respect to aggregate shocks) imposes sharp restrictions on the behavior of the multipliers. In particular, agents must satisfy the first-order condition with respect to W_{t+1} :

$$\begin{aligned} q_{t+1}(z^{t+1}|z^t)\widehat{\lambda}_t^j(\theta^t, z^t) &= E\{\widehat{\lambda}_{t+1}^j(\theta^{t+1}, z^{t+1})|\theta^t, z^{t+1}\}, t \geq 1 \\ q_1(z_1)\lambda_0^j &= E\{\widehat{\lambda}_1^j(\theta_1, z_1)|z_1\} \end{aligned}$$

where λ_0^j is the multiplier on the period 0 budget constraint.

Then, integrating over θ^t , we get

$$q_{t+1}(z^{t+1}|z^t)E\{\widehat{\lambda}_t^j|z^t\} = E\{\widehat{\lambda}_{t+1}^j|z^{t+1}\}, t \geq 1.$$

It follows that

$$\frac{E\{\widehat{\lambda}_t^j|z^t\}}{E\{\widehat{\lambda}_t^k|z^t\}} = \frac{\lambda_0^j}{\lambda_0^k}$$

for all dates and states. Hence, for each country pair (j, k) , there exists μ_{jk} such that

$$e_t^{jk}(z^t) = \frac{E\{c_t^{k-\gamma}|z^t\}}{E\{c_t^{j-\gamma}|z^t\}} \mu_{jk}.$$

2.2 *Private Information-Pareto Optimal Equilibrium*

It is well-known from the work of Green (1987) and others that in the above trading structure, equilibrium allocations are not Pareto optimal, even given the restriction that skills are private information. In this subsection, we describe an alternative trading mechanism such that the equilibrium allocation is (constrained) Pareto optimal. This trading protocol is similar to that described in Golosov and Tsyvinski (2006) (and originally in Atkeson and Lucas (1992)).

In period 0, insurance companies compete by offering contracts that specify consumption and output as a function of agents' reports. These contracts feature full commitment on the part of both sides. At the same time, the insurance companies trade a complete set of z^t -contingent assets with one another. Then, at each date, the companies trade tradeable and non-tradeable goods in each country. Let $q^t(z^t)$ be the price of a claim to the numeraire good in public history z^t in terms of the numeraire good in period 0.

In the contract market, the agents receive some equilibrium level of utility. The profit-maximizing insurance companies structure contracts so as to minimize the costs of providing that utility. Let y^* be the output process specified by the equilibrium contract. Then, the equilibrium consumption contract in

country j solves the problem

$$\begin{aligned}
& \min_{c^j} \sum_{j=1}^J \sum_{z^t} q^t(z^t) \left[\sum_{\theta^t \in \Theta^t} \{p_t^{TR}(z^t) c_t^{TR,j}(\theta^t, z^t) + p_t^{NT,j}(z^t) c_t^{NT,j}(\theta^t, z^t)\} \pi_j(\theta^t) \right] \\
& \text{s.t.} \quad \sum_{t=1}^T \beta^{t-1} \sum_{z^t} \Psi(z^t) \sum_{\theta^t} \pi_j(\theta^t) u(c_t^{TR,j}(\theta^t, z^t), c_t^{NT,j}(\theta^t, z^t))^{1-\gamma} / (1-\gamma) = u_*^j \\
& \text{s.t.} \quad V_j(c^j, y^{j*}; \sigma_{TT}) \geq V_j(c^j, y^{j*}; \sigma) \text{ for all } \sigma.
\end{aligned}$$

Here, u_*^j is the equilibrium level of ex-ante utility that agents receive from consumption contracts in country j .

It is straightforward to see that the insurance company equates individual *intratemporal* marginal rates of substitution to relative prices. Hence, in equilibrium:

$$\frac{u_{TR,t}^j(\theta^t, z^t)}{u_{NT,t}^j(\theta^t, z^t)} = \frac{p_t^{TR}(z^t)}{p_t^{NT,j}(z^t)}. \quad (10)$$

The insurance company's intertemporal first-order conditions can be derived as in Proposition 1 of Kocherlakota (2005). They are

$$u_{TR,t}^j u_t^{j-\gamma} E \left[\frac{1}{u_{TR,t+1}^j u_{t+1}^{j-\gamma}} \mid \theta^t, z^{t+1} \right] = \beta \frac{q^t(z^t) p_t^{TR}(z^t)}{q^{t+1}(z^{t+1}) p_{t+1}^{TR}(z^{t+1})}, t \geq 1 \quad (11)$$

$$E \left[\frac{q_1(z_1) p_1^{TR}(z_1)}{u_{TR,1}^j u_1^{j-\gamma}} \mid z_1 \right] = \mu_0^j, \quad (12)$$

where μ_0^j is independent of z_1 . In both of these sets of first-order conditions, we are using the notation of the prior subsection.

As in the previous subsection, the price index in country j satisfies the restriction:

$$\bar{p}_t^j(z^t) u_t^j(\theta^t, z^t) = [p_t^{TR}(z^t) c_t^{TR,j}(\theta^t, z^t) + p_t^{NT,j}(z^t) c_t^{NT,j}(\theta^t, z^t)]$$

By using this restriction together with (10), we can show that

$$\bar{p}_t^j(z^t) = \frac{p_t^{TR}(z^t)}{u_{TR,t}^j(z^t)} \quad (13)$$

so that the price index in country j equals the price of tradeable goods divided by the marginal utility of tradeable goods.

Because u is homogeneous of degree one and the own second derivatives of u are negative, u_{TR}/u_{NT} is a strictly decreasing function of $c^{TR,j}/c^{NT,j}$. Because the relative prices depend only on z^t , relative consumptions depend only on z^t . It follows too that the partial derivatives $u_{TR,t}^j(\theta^t, z^t)$ and $u_{NT,t}^j(\theta^t, z^t)$ are actually independent of θ^t . Hence, we can pull $1/u_{TR,t+1}^j$ outside of the conditional expectation in (11). We can then use (13) to rewrite (11)-(12) as

$$\begin{aligned} c_t^{j-\gamma} E[c_{t+1}^{j\gamma} | \theta^t, z^{t+1}] &= \beta \frac{q^t(z^t) \bar{p}_t^j(z^t)}{q^{t+1}(z^{t+1}) \bar{p}_{t+1}^j(z^{t+1})} \\ E[c_1^{j\gamma} | z_1] &= \frac{\mu_0^j}{q_1(z_1) \bar{p}_1^j(z_1)}, \end{aligned}$$

where we use the fact that in equilibrium, real consumption $c_t^j(\theta^t, z^t)$ equals momentary utility $u_t^j(\theta^t, z^t)$.

By integrating over θ^t , it follows that

$$\begin{aligned} E[c_t^{j\gamma} | z^t] \bar{p}_t^j(z^t) &= \beta \frac{q^{t+1}(z^{t+1})}{q^t(z^t)} E[c_{t+1}^{j\gamma} | z^{t+1}] \bar{p}_{t+1}^j(z^{t+1}) \\ E[c_1^{j\gamma} | z_1] &= \frac{\mu_0^j}{q_1(z_1) \bar{p}_1^j(z_1)} \end{aligned}$$

Taking cross-country ratios, we find that the real exchange rate can be rep-

resented as follows:

$$e_t^{jk}(z^t) \equiv \frac{\bar{p}_t^k(z^t)}{\bar{p}_t^j(z^t)} = \frac{E[c_t^{j\gamma}|z^t]}{E[c_t^{k\gamma}|z^t]} \nu_{jk}$$

for some ν_{jk} .

2.3 *Three Models and Their Implications for Real Exchange Rates*

In this subsection, we use the above results to derive the key testable implication for real exchange rates of three different trading models. In what follows, we use the notation $C_{\eta t}^j(z^t)$ to refer to the cross-sectional η -th (non-central) moment of consumption in country j and history z^t .

First, we proved above that if markets are domestically incomplete, then for some constant μ_{jk}

$$e_t^{jk}(z^t) = \frac{E\{c_t^{k-\gamma}|z^t\}}{E\{c_t^{j-\gamma}|z^t\}} \mu_{jk}.$$

At this point, we take a crucial step. As stated earlier, we assume that a law of large numbers applies across agents: conditional on any z^T , the measure of agents in the population of country j with type θ^T is given by $\pi_j(\theta^T)$. This assumption implies that the conditional expectation $E\{c_t^{k-\gamma}|z^t\}$ is equivalent to the cross-sectional average of the marginal utility of all agents in country k , given public history z^t . Hence:

$$e_t^{jk}(z^t) = \frac{C_{-\gamma t}^k(z^t)}{C_{-\gamma t}^j(z^t)} \mu_{jk}$$

By taking logs and first differences, we get the key implication of domestically incomplete markets for real exchange rates:

$$\Delta \ln e_t^{jk}(z^t) = \Delta \ln C_{-\gamma t}^k(z^t) - \Delta \ln C_{-\gamma t}^j(z^t), \quad (\text{DI})$$

where Δ is the first-difference operator.

We proved that in an equilibrium of the private information-Pareto optimal (PIPO) trading model,

$$e_t^{jk}(z^t) = \frac{E[c_t^{j\gamma} | z^t]}{E[c_t^{k\gamma} | z^t]} \nu_{jk}.$$

We can use logic similar to that above to conclude that the PIPO model implies that

$$\Delta \ln e_t^{jk}(z^t) = \Delta \ln C_{\gamma t}^j(z^t) - \Delta \ln C_{\gamma t}^k(z^t) \quad (14)$$

Finally, we derive the restrictions of the representative agent model. Suppose that Θ is a singleton, so that there is only a single type of agent in each country. Under either the incomplete markets model or the PIPO model, we know that

$$\Delta \ln e_t^{jk}(z^t) = -\gamma \Delta \ln C_{1t}^k(z^t) + \gamma \Delta \ln C_{1t}^j(z^t) \quad (\text{RA})$$

The growth rate of the real exchange rate is equal to the difference in growth rates of per capita real consumption, multiplied by the coefficient of relative risk aversion. This formula is the same as that of Backus-Smith (1993).

The key difference between these three restrictions is how cross-sectional consumption inequality is connected to real exchange rates. In (RA), only

per capita consumption enters the restriction. Hence, consumption inequality does not affect real exchange rates at all.

In contrast, in (DI), $C_{-\gamma t}^k(z^t)$ is driven by the left tail of the cross-sectional distribution of consumption. When the left tail grows thicker, then $C_{-\gamma t}^k(z^t)$ grows. Hence, in the (DI) model, the growth rate of the real exchange rate e_t^{jk} is large when the growth rate of the consumption of *poor* agents in country k is small relative to that of poor agents in country j .

Why does the real exchange rate respond in this way to the cross-sectional distribution of consumption? If the left tail of the cross-sectional distribution grows thicker in country k , agents are facing more uninsurable risks in that country. Because marginal utility is convex, the relative increase in the magnitude of uninsurable risk generates a relative increase in the *precautionary* demand for goods. It follows that goods grow relatively more valuable in country k when the left tail of the cross-sectional distribution of consumption grows relatively thicker.

In (??), $C_{\gamma t}^k(z^t)$ is influenced by the right tail of the cross-sectional distribution of consumption (at least when $\gamma > 1$, which is the empirically relevant range). In particular, when the right tail grows thicker, $C_{\gamma t}^k(z^t)$ grows larger. Hence, in the (PIPO) model, the real exchange rate between country j and k grows more slowly when the consumption of the rich in country k grows faster than the consumption of the rich in country j .

The intuition for this effect comes from incentives. Consider a country with a very thick right tail in the cross-sectional distribution of consumption,

so that one agent has accumulated almost all of the wealth in the society. It is hard to provide sufficient incentives to get this agent to produce. However, because of curvature of utility, it is relatively easy to provide incentives to get the large number of (possibly high-skilled) poor agents to produce. The incentive costs of production are much lower when the right tail of the consumption distribution is thick, and so goods are relatively *less* valuable in countries in which the right tail of consumption distribution is *thick*.

We can summarize this discussion as follows. The RA model predicts that only the mean of consumption matters for real exchange rates. The DI model predicts that goods in country k get *more* valuable when the *left* tail of the consumption distribution gets thicker in that country. The PIPO model predicts that goods in country k get *less* valuable when the *right* tail of the consumption distribution gets thicker. Our empirical work assesses which of these predictions fit the available data best.

3 Measurement Error and Sampling Error

If we had observations on consumption expenditures for all households in country j and k , we could directly test (DI), (PIPO), and (RA). However, there are two problems with this procedure. First, household consumption expenditures are probably measured with error. Second, we do not have access to the full population of all households. We consider each problem in turn. The discussion of measurement error is essentially the same as that of Kocherlakota and Pistaferri (2005).

Suppose that in country j , household j 's consumption expenditure is measured with multiplicative error ξ_t , so that observed consumption equals:

$$\widehat{c}_t^j(\theta^t, z^t, \xi_t^j) = c_t^j(\theta^t, z^t) \exp(\xi_t^j).$$

Assume that $\{\xi_t^j\}_{t=1}^T$ is a stationary stochastic process, where $\{\xi_t^j\}_{t=1}^T$ is i.i.d. over households and independent of z^T . Assume too that ξ_t^j is independent of true consumption $c_t^j(\theta^t, z^t)$, conditional on z^t . Then, it is simple to show that the η th moment of observed consumption equals

$$\begin{aligned} & E\{c_t^j(\theta^t, z^t)^\eta \exp(\eta \xi_t^j) | z^t\} \\ &= C_{\eta t}^j(z^t) E\{\exp(\eta \xi_t^j)\}. \end{aligned}$$

Under these conditions on ξ_t^j , we can conclude that

$$\Delta \ln \widehat{C}_{\eta t}^j(z^t) = \Delta \ln C_{\eta t}^j(z^t).$$

The first-difference of logged cross-sectional moments of mismeasured consumption data, with measurement error ξ , is the same as the first-difference of logged cross-sectional moments of true consumption data. The key to this result is the assumption that the measurement error ξ^j is stationary over time, so that the distribution of errors across households is the same in all dates and states. Note that this does not mean that ξ_t^j is i.i.d. over time.

This analysis provides a justification for us to ignore the presence of measurement error in what follows. However, the situation is different with *sampling* error. Let $\overline{C}_{\eta t}^j(z^t, \varepsilon_t^j)$ be the sample η -th moment of the consumption

distribution in country j in history z^t , given sampling error $\varepsilon_{\eta t}^j$. Then:

$$\begin{aligned}\ln \bar{C}_{\eta t}^j(z^t, \varepsilon_{\eta t}^j) &= \ln(C_{\eta t}^j(z^t) + \varepsilon_{\eta t}^j) \\ &\approx \ln(C_{\eta t}^j(z^t)) + \varepsilon_{\eta t}^j / C_{\eta t}^j(z^t).\end{aligned}$$

If the sample error is identically zero, it is possible to compare the marginal distribution of the left-hand side of equations like (RA) to the analogous properties of the right-hand side. For example, Backus and Smith (1993) compare the variance and autocorrelation of real exchange rates to the variance and autocorrelation of the difference in growth rates of marginal utilities. They conclude that these properties of real exchange rates are quite different from the corresponding properties of the right-hand side of (RA) . It is not possible to proceed in this fashion when sampling error is nontrivial; the properties of the right-hand side will be at least affected, and possibly overwhelmed, by sampling error.⁴

To get around this problem, we use the following approach. Define

$$\chi_{t,DI}^{jk}(z^t, \varepsilon_t) = \Delta \ln e_t^{jk}(z^t) - \Delta \ln \bar{C}_{-\gamma t}^k(z^t, \varepsilon_t) + \Delta \ln \bar{C}_{-\gamma t}^j(z^t, \varepsilon_t) \quad (15)$$

$$\chi_{t,PIPO}^{jk}(z^t, \varepsilon_t) = \Delta \ln e_t^{jk}(z^t) - \Delta \ln \bar{C}_{\gamma t}^j(z^t, \varepsilon_t) + \Delta \ln \bar{C}_{\gamma t}^k(z^t, \varepsilon_t) \quad (16)$$

$$\chi_{t,RA}^{jk}(z^t, \varepsilon_t) = \Delta \ln e_t^{jk}(z^t) - \gamma \Delta \ln \bar{C}_{1t}^j(z^t, \varepsilon_t) + \gamma \Delta \ln \bar{C}_{1t}^k(z^t, \varepsilon_t), \quad (17)$$

⁴Brandt, Cochrane, and Santa-Clara (2005) derive necessary conditions for two country-specific stochastic discount factors to be simultaneously consistent with real exchange rates and with high within-country equity premia. They show that the stochastic discount factors must necessarily be highly correlated with one another. Unfortunately, the sampling error issue means that we cannot evaluate this implication for our SDFs.

where the bars represent sample moments as opposed to population moments.

Under the null hypothesis that the *DI* model is true, then:

$$\chi_{t,DI}^{jk}(z^t) \approx -\Delta\left[\frac{\varepsilon_{\gamma t}^j}{C_{\gamma t}^j(z^t)} + \frac{\varepsilon_{\gamma t}^k}{C_{\gamma t}^k(z^t)}\right].$$

Because $\varepsilon_{\gamma t}^j$ is sampling error, the mean of $\Delta(\varepsilon_{\gamma t}^j/C_{\gamma t}^j(z^t))$, conditional on z^t , is zero. It follows that the first differenced sampling errors $\Delta\varepsilon_{\gamma t}^j$ and $\chi_{t,DI}^{jk}$ are both uncorrelated with Δe_t^{jk} .

This implication suggests the following testing methodology. If we regress $\chi_{t,DI}^{jk}$ on $\Delta \ln e_t^{jk}$, we should get a slope coefficient of zero. Similarly, under the null that the PIPO (RA) model is true, a regression of $\chi_{t,PIPO}^{jk}$ ($\chi_{t,RA}^{jk}$) on $\Delta \ln e_t^{jk}$ should result in a slope coefficient of zero. In contrast, if we obtain a slope coefficient near 1, we can conclude that the models are not faring well: there is little empirical relationship between the left-hand and right-hand sides of the relevant exchange rate restrictions.

4 Data Description

In this section, we describe the data that we used. We begin with the real exchange rates, and then turn to the two household consumption data sets.

4.1 *Real Exchange Rates*

We constructed the monthly real exchange rate between the United States and the United Kingdom as follows. We began with the International Financial Statistics data on the monthly average of daily pound-dollar exchange

rates.⁵ We then divided this exchange rate by the monthly United Kingdom Consumer Price Index (for all items). We then multiplied this exchange rate by the monthly United States Consumer Price Index (for all items). All CPIs are in 2000 basis. The source for both of these price index series was the OECD Main Economic Indicators, vol. 2005, release 07.⁶ In this way, we construct a monthly series for the real exchange rate between the United States and the United Kingdom.

4.2 *The CEX*

The following description closely follows that in Kocherlakota and Pistaferri (2005).

The micro data for the United States are drawn from the 1980-2004 Consumer Expenditure Survey (CEX). The CEX provides a continuous and comprehensive flow of data on the buying habits of American consumers. The data are collected by the Bureau of Labor Statistics and used primarily for revising the Consumer Price Index. Consumer units are defined as members of a household related by blood, marriage, adoption, or other legal arrangement, single person living alone or sharing a household with others, or two or more persons living together who are financially dependent. The definition of the head of the household in the CEX is the person or one of the persons who owns or rents the unit.

⁵The series code is 112..AE.ZF....

⁶The UK price level is reported in series GBR.CPALTT01.IXOB. The US price level is reported in series USA.CPALTT01.IXOB.

The CEX is based on two components, the Diary, or record-keeping survey, and the Interview survey. The Diary sample interviews households for two consecutive weeks, and it is designed to obtain detailed expenditures data on small and frequently purchased items, such as food, personal care, and household supplies. The Interview sample is in the form of a rotating panel, and it follows survey households for a maximum of five quarters, although only inventory and basic sample data are collected in the first quarter (these data are not publicly available). The database covers about 95% of all expenditure, with the exclusion of expenditures for housekeeping supplies, personal care products, and non-prescription drugs. Following most previous research, our analysis below uses only the Interview sample.

The CEX collects information on a variety of socio-demographic variables, including characteristics of members, characteristics of housing unit, geographic information, inventory of household appliances, work experience and earnings of members, unearned income, taxes, and other receipts of consumer unit, credit balances, assets and liabilities, occupational expenses and cash contributions of consumer unit. Expenditure is reported in each interview (after the first) and refers to the months of the previous quarter. Thus, a household interviewed in April 1980 reports expenditure for January, February, and March 1980. Income is reported in the second and fifth interview, and it refers to the previous 12 months. Holdings of financial assets are reported only in the fifth interview.

Our sample selections are aimed at eliminating the most severe report-

ing errors in consumption. We end up discarding about 25% of observations through our selection procedure. The definition of total non-durable consumption is similar to Attanasio and Weber (1995). It includes food (at home and away from home), alcoholic beverages and tobacco, heating fuels and utilities, transports (including gasoline), personal care, clothing and footwear, entertainments, and other services (including domestic services). It excludes expenditure on various durables, housing (furniture, appliances, etc.), education, and health. We refer the reader to the Appendix for step-to-step details on sample selection and consumption definition.

We “deflate” consumption data to account for three phenomena: price differences over time, seasonal differences (i.e., month effects) within a year, and households’ demographic differences at a certain point in time. Thus, nondurable consumption is first expressed in real terms using the CPI (all items) described above. Then, data are de-seasonalized by simple multiplicative regression adjustments. Finally, we convert it into adult-equivalent consumption data.⁷ Given the overlapping panel nature of the CEX, each month a certain number of households enter the panel and an approximately equal number leave it. Monthly consumption data are aggregated to form quarterly consumption data for each household in the sample.⁸ Thus, the

⁷The number of adult equivalents is defined as $(A + \alpha K)^\beta$, where A is the number of adults (aged 18 or more), K the number of kids, and α and β parameters. We set $\alpha = 0.7$ and $\beta = 0.65$ (following recommendations contained in Citro and Michaels, 1995, which in turn draws from Betson, 1990). Similar results are obtained if we use a more sophisticated Engel approach.

⁸Recently, researchers have noted that for many commodities, the aggregation of CEX data matches poorly National Income and Product Accounts (NIPA) Personal Consump-

CEX provides us with a *monthly* sample of household *quarterly* consumption.

4.3 *The FES*

The micro data for the UK come from the 1975–1999 Family Expenditure Survey (FES). The FES, conducted by the Office for National Statistics, surveys a random sample of households in the UK. The FES is primarily a survey of household expenditure on goods and services, and household income. Similarly to the CEX, the main goal of the survey was originally that of providing information on spending patterns for revising the Consumer Price Index. However, with time the survey has become multi-purpose, providing a wealth of information on household economic and social variables, such as household composition, size, social class, occupation, etc.

The FES has been in operation since 1957 and up to and including 1993, data are available by calendar year. From 1994–1995 data are available by financial year (April–March). The basic unit of the survey is the household. Starting with 2000–2001, a household is defined as a group of people living at the same address with common housekeeping, i.e., sharing household

tion Expenditure (PCE) data. Some of the discrepancy is undoubtedly due to differences in covered population and definitional issues. But the amount of underestimation of consumer expenditure is sometimes substantial, and it raises some important warning flags. Furthermore, there is evidence that the detachment between the CEX aggregate and the NIPA PCE has increased over time. At present, it is not clear why this is so, and whether this is necessarily due to a worsening in the quality of the CEX. For example, Bosworth et al. (1991) conclude that most of the discrepancy is explained by the failure of the CEX to sample the super-rich; others have suggested a greater incidence of attrition. According to the BLS, however, the CEX has maintained representativeness of the US population over time, and attrition has not changed much since the redesign of the survey of the early 1980s.

expenses such as food and bills, *or* sharing a living room. Before 2000, the definition of household required both common housekeeping *and* a shared living room. On average, about 7,000 households are interviewed each year. Each individual in the household aged 16 or more is asked to keep diary records of daily expenditure for two weeks. Information about regular expenditure, such as rent and mortgage payments, is obtained from a household interview along with retrospective information on certain large, infrequent expenditures such as those on vehicles. Data are collected throughout the year to cover seasonal variations in expenditures.

The FES consists of three main modules, a Household Schedule, an Income Schedule, and Diary Records. The Household Schedule is taken at the main interview. Information for most of the questions is obtained from the head of household (or housewife). Information is collected about the household, the sex and age of each member, and also details about the type and size of the household accommodation. The main part of the questionnaire relates to expenditure both of a household and an individual nature, but the questions are mainly confined to expenses of a recurring nature.

The Income Schedule data are collected for each household spender. The schedule is concerned with income, national insurance contributions, and income tax. Information collected includes: employment status and recent absences from work, earnings of an employee, self-employed earnings, National Insurance contributions, pensions and other regular allowances, occasional benefits, social security benefits and other types, investment income, mis-

cellaneous earnings of a “once-only” character, tax paid directly to Inland Revenue or refunded, and income of a child.

The Diary Records cover 14 days. As said, each household member aged 16 or more is asked to record all expenditure made during the 14 days. As for the CEX, we refer the reader to the Appendix for step-to-step details on sample selection and consumption definition. The definition of total non-durable consumption is as in Attanasio and Weber (1993), and it is meant to be comparable to the one for the CEX.

Similarly to the CEX, the FES data we construct are “deflated” to account for price differences over time, seasonal differences (i.e., month effects) within a year, and households’ demographic differences at a certain point in time. Thus, nondurable consumption is first expressed in real terms using the CPI (all items) described above (on the same basis as the US CPI). Then, data are de-seasonalized by simple multiplicative regression adjustments. Finally, we convert it into adult-equivalent consumption data. Given that data are provided for only two weeks, we multiply FES expenditure by 6.5 to form quarterly consumption data for each household in the sample.⁹

As stated above, the FES randomly selects households throughout the

⁹As far as reliability is concerned, great care is taken in collecting information from households and comprehensive checks are applied during processing, so that errors in recording and processing are minimized. The main factors that affect the reliability of the survey results are sampling variability, non-response bias and some incorrect reporting of certain items of expenditure and income. Procedures are in place to ensure that users are provided with high quality data. For example, quality control is carried out to ensure that any outliers are genuine, and checks are made on any unusual changes in average spending compared with the previous year.

year. The survey keeps track of those households' consumption expenditures for a two week period. We identify households interviewed in a certain month as those completing their two-week reporting in that month. We treat their (scaled-up) non-durable consumptions as being a random sample of household consumption over the prior quarter.

5 Estimation Results

We have data from the FES from 1975–1999 and data from the CEX from 1980–2004. We use data from the period 1980–1999. From the FES and the CEX, we have *monthly* random samples of *quarterly* household consumption. Given the nature of these data, when we first difference, we subtract the observation from month $(t - 3)$ from the observation from month t . Thus, the structure of our sample is that we have monthly observations on quarterly first-differences; the calculation of standard errors must take into account this overlapping structure.

In Table 1, we report the results of regressing $\chi_{t,j}^{US,UK}$, as defined in (15)–(17) on the first differenced logged real exchange rate between the United States and the United Kingdom. Standard errors are corrected for heteroskedasticity and serial correlation induced by the use of overlapping data using a correction of the form proposed by Hansen and Hodrick (1980). However, we downweight longer lags as in Newey and West (1987) to obtain positive definite covariance matrices.

Look first at the representative agent results. The estimated coefficients

are all statistically insignificantly different from one, but statistically significantly different from zero. This column confirms Backus and Smith's results: the difference in per-capita consumption growth rates between the United States and United Kingdom has little to do with real exchange rates.

Next, look at the PIPO results. When γ is low (in the range 1–4), the estimated regression coefficient is close to 1. For these values of γ , as in the representative agent case, the model seems to have little bite. But for γ near 5, the estimated regression coefficient is zero. The difference in the growth rates of the 5-th moments of the cross-sectional consumption distribution tracks the real exchange rate well.

Finally, look at the DI results. Here, the estimated regression coefficients are all positive, and they grow larger with γ , not smaller. However, as is true in the PIPO case, sampling error is large. Hence, it is not possible to reject the null hypothesis that the true regression coefficient is zero for high values of γ .

In Table 1 we do not formally estimate γ . In contrast, in Table 2, we define α to be the intercept term of the regression of $\chi_{t,j}^{US,UK}$ on the logged first-differenced real exchange rate. We then use the Generalized Method of Moments (GMM) to obtain estimates of α and of the coefficient of relative risk aversion γ . For the PIPO model, the estimate of γ is close to 5, as indicated in Table 1. However, the standard error is high. In the DI case, the estimate that makes the moment condition as close as possible to zero is around 3. The usual GMM estimate of the standard error is undefined be-

cause the (one) moment condition is not set equal to zero in sample. Finally, in the RA case, the estimate of γ is negative (but highly imprecise).

In Tables 1 and 2, we exploited the restriction that the slope coefficient is supposed to be zero in the regression under the null hypothesis. Since we have one parameter and one restriction, the models are exactly identified. However, under the null hypothesis, the intercept term is also supposed to be zero. In Table 3, we exploit this second restriction. We estimate γ using the following moment conditions:

$$\begin{aligned} E\left(\chi_{t,DI}^{jk}(z^t)\right) &= 0 \\ E\left(\chi_{t,DI}^{jk}(z^t)\Delta \ln e_t^{jk}(z^t)\right) &= 0 \end{aligned}$$

The results of this estimation offer more evidence in favor of the PIPO model. In the PIPO model, the estimated value of γ is again around 5. The estimate of γ is around 2 for the DI model and around 0.2 for the RA model. All of these estimates, in particular in the RA case, are imprecise. More interestingly, the J-statistic is very large for the DI and RA models; they would both be rejected at standard levels of significance. In contrast, the J-statistic for the PIPO model is small.

In Figures 1-3, we give a visual representation of the goodness of fit of the various models. We use the estimate of γ from Table 3 to create an estimate of the right hand side of (DI) (we repeat this for the other two models (PIPO) and (RA)). We then regress this onto the growth of the real exchange rate

(the left hand side of the equation). If the model fits well, the projection of this regression should be close to the growth of the real exchange rate. The figures show that the PIPO fits extremely well, while the other two models fit the growth of the real exchange rate rather poorly.

We did two robustness checks on our results. In constructing the real exchange rates and the real consumption data, we use the Consumer Price Index based on all consumption goods. However, we actually restrict attention to non-durable consumption goods. For this reason, it would be desirable to use an index based only on non-durable consumption goods. Unfortunately, we have not been able to find such a price index for the UK. We did re-estimate the models using the non-durable consumer goods price index from the BLS in the United States. Doing so made little change in the point estimates. However, the test statistics in Table 3 for the DI and RA models were considerably larger. We conclude that our use of the CPI for all goods did not affect our results in important ways.

In our formulation of the DI model, we assume that all agents participate in the asset markets, and that no agent ever faces a binding short-sales constraint in those markets. In reality, it is likely that at least some agents do face such binding constraints. To allow for this possibility, we need information about household asset holdings. There is no information about asset holdings in the FES. In the CEX, households are asked in their fifth interview about their current asset holdings and about the change in their asset holdings relative to a year before. We first eliminate all households in

the sample that do not complete all interviews. We then discard all those who did not report positive holdings of stocks and bonds in their second interview.

We re-estimated the DI model using this sample of households from the CEX (without changing the selected sample from the FES). The results in Table 1 did not change, the point estimate of γ in Table 2 became negative, and the point estimate of γ in Table 3 became small and insignificantly different from zero. We concluded that accounting for binding short-sales constraints and/or non-participation in this fashion did not help the performance of the DI model.¹⁰

¹⁰We reach similar conclusions in our earlier paper (Kocherlakota and Pistaferri (2005)).

Table 1

Slope Estimates in Regressions of Model Errors on Exchange Rates

γ	<i>PIPO</i>	<i>Domestically incomplete markets</i>	<i>Representative agent</i>
1	1.06 (0.08)	1.04 (0.06)	1.06 (0.08)
2	1.36 (0.47)	1.05 (0.21)	1.11 (0.16)
3	1.62 (1.61)	1.17 (0.78)	1.17 (0.24)
4	1.11 (2.97)	1.63 (1.63)	1.23 (0.32)
5	0.16 (4.32)	2.17 (2.57)	1.28 (0.40)
6	-0.91 (5.62)	2.65 (3.53)	1.34 (0.48)
7	-1.96 (6.87)	3.06 (4.49)	1.39 (0.56)
8	-2.94 (8.10)	3.40 (5.43)	1.45 (0.64)
9	-3.86 (9.30)	3.68 (6.36)	1.51 (0.72)
10	-4.74 (10.49)	3.94 (7.27)	1.56 (0.80)

Note: Under the null hypothesis of each model, the true slope coefficient is zero. Standard errors are of the Hansen-Hodrick (1980) form but downweight higher order correlations as in Newey and West (1987) to ensure positive definiteness. The number of lags used is three.

Table 2GMM Estimates of γ and α

	<i>PIPO</i>	<i>Domestically Incomplete Markets</i>	<i>Representative Agent</i>
γ	5.1541 (4.2212)	2.1544 (<i>n.a.</i>)	-17.7682 (25.8401)
α	0.0041 (0.1990)	0.0007 (<i>n.a.</i>)	-0.0920 (0.1464)

Note: Under the null hypothesis, the intercept term α is equal to zero. In the incomplete markets case, the usual asymptotic formulae for standard errors cannot be used because there are no parameters that set the minimized objective equal to zero, even though the model is exactly identified. Standard errors are of the Hansen-Hodrick (1980) form but downweight higher order correlations as in Newey and West (1987) to ensure positive definiteness. The number of lags used is three.

Table 3GMM Estimate of γ

	<i>PIPO</i>	<i>Domestically Incomplete Markets</i>	<i>Representative Agent</i>
γ	5.2306 (3.8724)	2.2270 (1.5674)	0.1880 (1.4481)
<i>J</i> -stat.	0.0004	3.9927	10.2259

Note: Standard errors are of the Hansen-Hodrick (1980) form but downweight higher order correlations as in Newey and West (1987) to ensure positive definiteness. The number of lags used is three. The *J*-statistic has a $\chi^2(1)$ asymptotic distribution under the null.

6 Conclusion

Backus and Smith (1993) consider a model in which all countries have a representative agent and markets are complete. They prove that in this setting, the growth rate in the real exchange rate should equal the difference in marginal utility growth rates in every date and state. They show that this restriction is not satisfied by available exchange rate data.

In this paper, we relax the assumption that there is a representative agent in each country, and instead assume that agents can only partially share person-specific risks within each country. We consider two forms of limited risk-sharing: first, Pareto optimal risk-sharing conditional on asymmetric information about skills, and second, competitive exchange of a limited set of securities. We derive analogs of the Backus-Smith restriction for these two forms of risk-sharing. We show that this analog restriction fits the data in the constrained Pareto optimal case, at least for coefficients of relative risk aversion around 5. In Kocherlakota and Pistaferri (2005), we show that the PIPO model rationalizes the equity premium in the United States when the coefficient of relative risk aversion is also around 5.

Our analysis is limited in a number of respects. Nonetheless, we believe that our results are quite suggestive. There are a large number of asset pricing and trade anomalies in international economics. Our results indicate that limited within-country risk-sharing might help explain at least some of these.

There are at least a couple of useful directions for future research. We treat the real exchange rate in isolation. It would be informative to explore the ability of the model(s) to rationalize financial asset returns in multiple countries, together with the real exchange rate. The work of Brandt et al. (2005) suggests that this test would be a stringent one.

We do not evaluate the ability of either model to generate the observed cross-sectional distribution of consumption in the United States or United Kingdom. We believe that taking this step is an essential one to build further confidence in the PIPO model. Doing so requires us to solve two difficult problems. First, we have to find a way to measure the process generating productivity shocks and aggregate shocks. Then, given that process, we have to solve for the data generation process for the cross-sectional distribution of consumption. We hope to make progress on these questions in future research.

References

- [1] Atkeson, A., and Lucas, R. E., Jr. (1992). ‘On efficient distribution with private information’, *Review of Economic Studies* vol. 59(3) (July), pp. 427–53.
- [2] Attanasio, O., and Weber, G. (1993). ‘Consumption growth, the real interest rate, and aggregation’, *Review of Economic Studies* 60(3) (July), pp. 631–49.
- [3] Attanasio, O., and Weber, G. (1995). ‘Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey’, *Journal of Political Economy* 103(6) (December), pp. 1121–57.
- [4] Backus, D., and Smith, G. (1993). ‘Consumption and real exchange rates in dynamic economies with non-traded goods’, *Journal of International Economics* 35(3-4) (November), pp. 297–316.
- [5] Benigno, G., and Thoenissen, C. (2004). ‘Consumption and real exchange rates with incomplete markets and non-traded goods’, London School of Economics working paper.
- [6] Betson, D. (1990). ‘Alternative estimates of the cost of children from the 1980–1986 Consumer Expenditure Survey’, US Department of Health and Human Services, Washington, D.C..

- [7] Bosworth, B., Burtless, B., and Sabelhaus, J. (1991). ‘The decline in saving: evidence from household surveys’. *Brooking Papers on Economic Activity*, pp. 183–241.
- [8] Brandt, M. W., Cochrane, J. H., and Santa-Clara, P. (2006). ‘International risk sharing is better than you think or exchange rates are much too smooth’, *Journal of Monetary Economics*.53(4) (May), pp. 671-698.
- [9] Chari, V., Kehoe, P., and McGrattan, E. (2002). ‘Can sticky price models generate volatile and persistent real exchange rates?’ *Review of Economic Studies* 69(3) (July), pp. 533–63.
- [10] Citro, C., and Michaels, R., eds. (1995). *Measuring Poverty: A New Approach (1995)*, Washington, D.C.: National Academy Press.
- [11] Colacito, R., and Croce, M. (2005). ‘Risks for the long run and the real exchange rate’, New York University working paper.
- [12] Corsetti, G., Dedola, L., and Leduc, S. (2005). ‘International risk-sharing and the transmission of productivity shocks’, International Finance Discussion Paper 826, Board of Governors of the Federal Reserve System.
- [13] Epstein, L., and Zin, S. (1989). ‘Substitution, risk aversion, and the temporal behavior of consumption and asset returns’, *Econometrica* 57(4) (July), pp. 937–69.

- [14] Golosov, M., Kocherlakota, N., and Tsyvinski, A. (2003). ‘Optimal indirect and capital taxation’, *Review of Economic Studies* 70(3) (July), pp. 569–87.
- [15] Golosov, M., and Tsyvinski, A. (forthcoming). ‘Optimal taxation with endogenous insurance markets’, *Quarterly Journal of Economics*.
- [16] Green, E. (1987). ‘Lending and the smoothing of uninsurable income shocks’, in *Contractual Agreements for Intertemporal Trade*, ed. Prescott, E., and Wallace, N., Minneapolis: University of Minnesota Press, pp. 3–25.
- [17] Hansen, L., and Hodrick, R. (1980). ‘Forward exchange rates as optimal predictors of future spot rates: An econometric analysis’, *Journal of Political Economy* 88(5) (October), pp. 829–53.
- [18] Kocherlakota, N. (2005). ‘Zero expected wealth taxes: A Mirrlees approach to dynamic optimal taxation’, *Econometrica* 73(5) (September), pp. 1587–1621.
- [19] Kocherlakota, N., and Pistaferri, L. (2005). ‘Asset pricing implications of Pareto optimality with private information’, University of Minnesota working paper.
- [20] Newey, W., and West, K. (1987). ‘A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix’, *Econometrica* 55(3) (May), pp. 703–8.

7 Appendix: The CEX

Each annual tape of the CEX contains four groups of core files, MTAB, ITAB, FMLY, and MEMB. There are also many auxiliary files, which are not used here. The FMLY, MEMB, MTAB, and ITAB files are organized by the calendar quarter of the year in which the data are collected. The FMLY files contain household characteristics, income, and summary level expenditures; the MEMB files contain member characteristics and income data; the MTAB files contain expenditures organized on a monthly basis at the Universal Classification Code (UCC) level; and the ITAB files contain income data converted to a monthly time frame and assigned to UCCs.

There are five quarterly data sets for each of these files. For example, in the 1980 tape, there are files running from the first quarter of 1980 through the first quarter of 1981. This is for the purpose of allowing computation of calendar year statistics. In fact, the MTAB 1980:Q1 file, say, contains expenditure information as reported by households interviewed in the first quarter of 1980. Since households report data for the previous three months, the 1980:Q1 file, say, has expenditure data from October 1979 to February 1980. The 1981:Q1 file is included in the 1980 tape because it contains information that spans the last three months of 1980. With just two exceptions (discussed below), the Y:Q1 file contained in the Y-1 tape is identical to the Y:Q1 file contained in the Y tape. The FMLY file for a given quarter has one record per household. Similarly, the MEMB file for a given quarter has

one record per household member.

The CEX has three “household tracking” problems, detailed as follows. In 1980–1981 households that re-entered the survey after missing an interview were assigned a new ID. While this is probably a minor proportion of the whole sample, it is for all purposes not a problem in our context, given that we do not focus on the longitudinal aspect of the survey. In 1986 the CEX changed its sample design. The consequence of this is that the core 1986:Q1 files contained in the 1985 tape survey different households than the core 1986:Q1 files contained in the 1986 tape. Indeed, issuing of household IDs starts from scratch beginning with the 1986 tape. Again, this is not a problem for us. We use both “samples,” and so end up with a size that is larger than usual. Another sample design change occurs in 1996, but in this case some of the households that are surveyed in the core 1996:Q1 files contained in the 1995 tape appear also in the core 1996:Q1 files of the 1996 tape, although with the same ID. Of course, we eliminate the duplicates.

We use the MTAB files from 1980:Q1 throughout 2004:Q1 to create monthly expenditure records for each household ever surveyed in the CEX. Since households report data for at most four quarters, there are between 3 and 12 observations per household. We merge this information with household characteristics from the FMLY file. The file so compiled contains 1,848,352 observations (where, to reiterate, each observation is a household/month data point).

Our measure of nondurable consumption is as in Attanasio and Weber

(1995), and it is the sum of the following items (the UCC codes are in parentheses):¹¹ *Food at home* (790220, 790230, 190904), *Food away from home* (190901-190903, 790410, 790430, 800700), *Alcohol* (200900, 790310, 790320, 790420), *Tobacco* (630110, 630210), *Clothing and footwear* (360110-420120), *clothing services* (440110-440140, 440210, 440900), *Fuel, light and power* (the sum of Heating (250111-250904), Gas (260211-260214), Electricity (260111-260114)), *Telephone communications* (270000-270104, 270310), *Housing* (the sum of Water and sewerage (270211-270214, 270411-270414, 270901-270904), Rent (210110-210902, 800710, 350110), Home insurance (220111-220122), Home maintenance and repairs (230111-230902, 330511, 340914, 790600), and Other home services (340510-340530, 340906, 340911-340912, 340915)), *Domestic services* (the sum of Babysitting (340210-340212) and Other domestic services (340310-340420)), *Public transport* (530110-530902), *Private transport* (the sum of Vehicle expenses (520110-520907), Gasoline and oil (470111-470212), Vehicle maintenance and repairs (470220-490900), and Parking fees (220901-220902)), *Entertainment* (the sum of Club membership fees (620110-620115), Ticket admissions (620121-620310), and Miscellaneous entertainment expenses (610900, 620330-620926)), and *Miscellaneous expenditures* (the sum of Newspapers and magazines (590110-590212), Books (590220-590230), Personal care (650110-650900), and Rentals (340610-340905, 340907-340908, 440150).

¹¹Erich Battistin at IFS kindly provided assistance in replicating Attanasio and Weber's aggregation procedures.

Here is a description of our sample selection. As said, we start with 1,848,352 monthly observations. We drop 164,137 observations for which our measure of total nondurable consumption is missing or zero. We further drop 4,259 observations for households that report zero food spending (at home and away from home) during an entire interview (three-month period). Households interviewed in a certain month are supposed to report consumption data only for the previous three months. We drop 4,116 observations reporting data for the same month in which they are interviewed. We next eliminate 398,047 observations corresponding to households that are classified as “incomplete income respondents” or report less than three months of data for a given interview. At this point, we aggregate monthly data in (overlapping) quarters, indexed by the last month in the quarter. The resulting sample has 425,931 quarterly observations, corresponding to 147,412 households. We drop 9,635 observations corresponding to households that jump interviews (i.e., exit and re-enter the survey), and 5,498 observations corresponding to households living in college dorms. We end up with a final sample of 410,798 quarterly observations, or 140,364 households.

8 Appendix: The FES

The FES is composed of various modules: The Household schedule A (which asks questions about demographics, housing, durables), the Income schedule B (which asks questions about employment, wages, transfers, assets), the Diary of expenditure schedule D (which asks questions about expenditure on

goods and services), the Checking schedule K and the Checking and outcome schedule K+L (which ask a number of consistency check questions), the England, Wales, and Northern Ireland-only schedule M (which asks questions – especially about publicly provided housing – pertinent to all countries but Scotland), and the Northern Ireland schedule N.I. (which asks questions – especially about publicly provided housing – only pertinent to Northern Ireland). Household expenditure on nondurable goods and services, described below, uses information from schedules A, B, and D.

The expenditure codes in the FES have changed several times between 1968 and 2004. In what follows, we report the expenditure codes and their allocation to the various consumption categories for a single representative year, 1990. For other years, the allocation is similar, although it is based on different expenditure codes.¹² The definition of non-durable consumption and services includes the following expenditure categories:¹³ *Food at home* (101-137, 139-195, 198-199), *Food away from home* (840-857, 138, 196-197), *Alcohol* (260-289), *Tobacco* (211-213), *Fuel, light and power* (the sum of Coal and solid fuel (240,242,A321), Electricity (225,255,175,B222,B178⁻), Gas (226,254,B170,B221,B173⁻), and Other fuels (258,B017,B027)), *Telephone communications* (227,752,B166), *Housing* (the sum of Rent, Rates and com-

¹²We thank Orazio Attanasio and Andrew Leicester for providing information on the allocation procedure.

¹³The acronym “A” before an expenditure code indicates that the data come from the *Household* schedule A. The acronym “B” before an expenditure code indicates that the data come from the *Income* schedule B. A superscript ⁻ on an expenditure code means that the value is subtracted from the total. Expenditure codes without acronyms come from the *Diary of Expenditures* schedule D.

munity charges, Repair and maintenance (B102,B104,B107,B108), Own repair (232-238), Mortgage payments (B130,B198-B200), and Home insurance (B60,B110)), *Domestic services* (the sum of Domestic help (780), Repairs (782,788), and Laundry services (790,791)), *Clothing and footwear* (301-349), *Private transports* (the sum of Maintenance (510,513,514,545,546,548,549), Gasoline and oil (538,539,542), Vehicle tax and insurance (512,B187,B188,B179), and others (555,556)), *Public transports* (the sum of Rail fares (550,551,B216,B218,B220), Bus and coach fares (552,B217), Air travel (553), and School travel (B158)), *Entertainment* (the sum of TV license (229,760,768,B181), TV and cable rental (B195,B253,B254), Fees and admissions to events (753,755,761,763,764,765,B162), Hotels (B441-B452), and Holidays (754-759)), and *Miscellaneous items* (the sum of Household consumables (437,623,648,742,746), Petcare (731,732), Postage (751), Fees and subscriptions (219,220,228,770,772,796,797,799,805,806,807,B168,B180,B273,B283), Gardening (733,734), and Books, newspapers, and magazines (721-723)).

We use the FES files to create bi-weekly expenditure records for each household ever surveyed in the FES. There is a single record for each household. In other words, there is no panel component in the FES. The file so compiled (after excluding households from Northern Ireland, less than 2% of the sample, and those for which our measure of total nondurable consumption is missing or zero) contains 172,163 observations (approximately 7,000 per year). We create quarterly expenditures by multiplying monthly expenditure by 6.5.

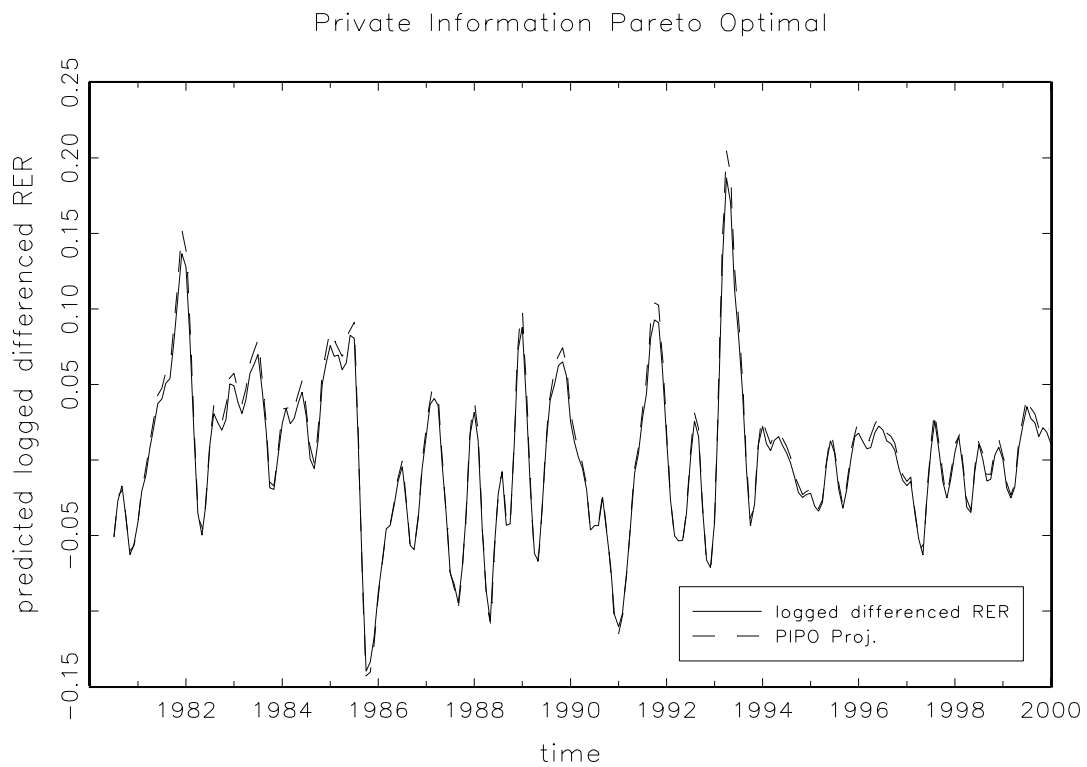


Figure 1: To construct this graph, we regressed the first difference of the log of the PIPO SDF ratio on the first difference of the log of the real exchange rate. We then plot the resultant projection and first difference of the log of the real exchange rate against time.

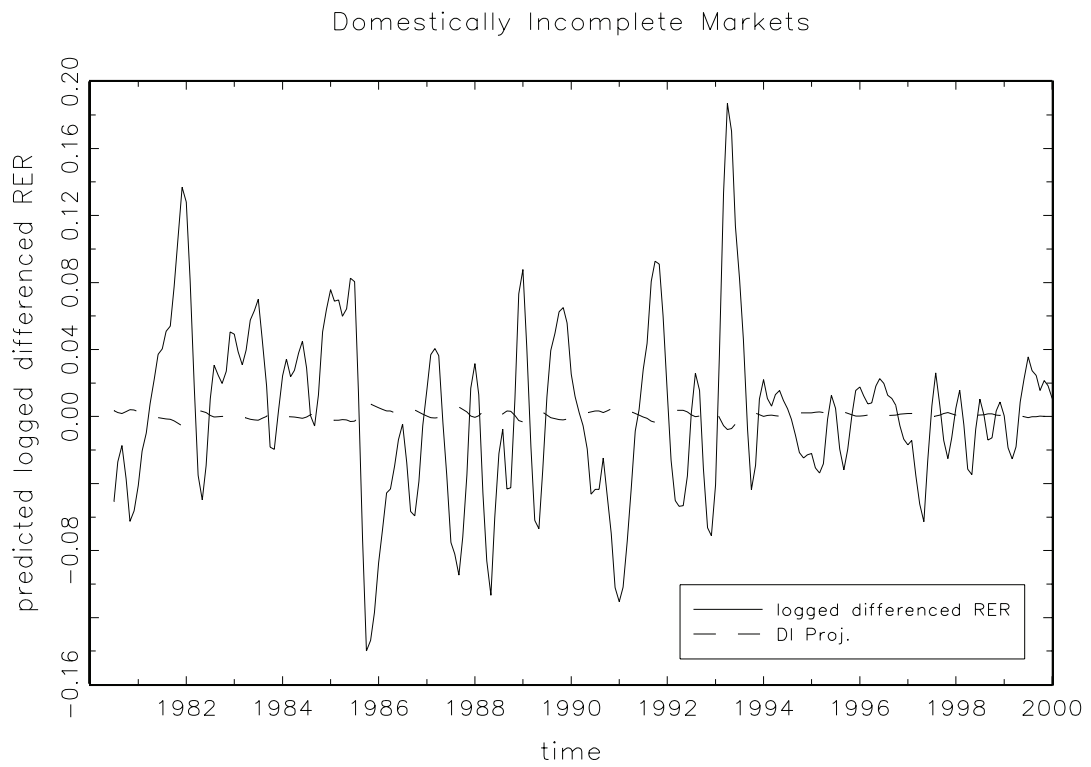


Figure 2: To construct this graph, we regressed the first difference of the log of the DI SDF ratio on the first difference of the log of the real exchange rate. We then plot the resultant projection and first difference of the log of the real exchange rate against time

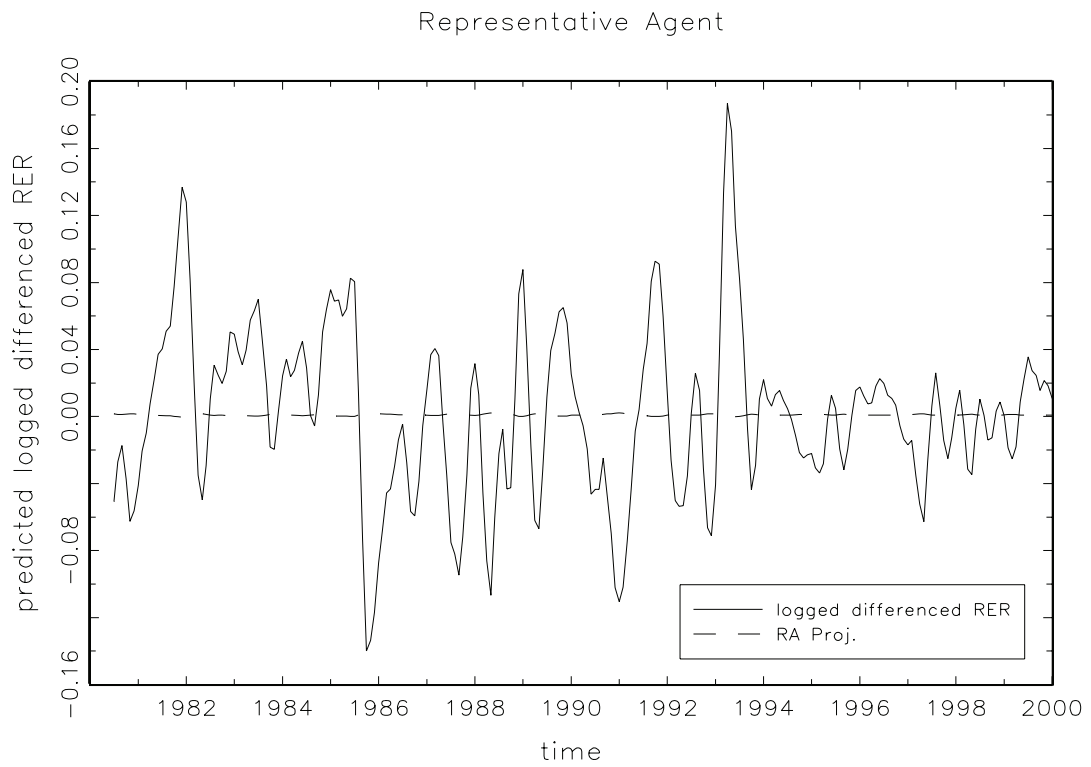


Figure 3: To construct this graph, we regressed the first difference of the log of the RA SDF ratio on the first difference of the log of the real exchange rate. We then plot the resultant projection and first difference of the log of the real exchange rate against time