EUROPEAN UNIVERSITY INSTITUTE
DEPARTMENT OF ECONOMICS

EUI Working Paper **ECO** No. 2003/15

Home Bias and the Structure of International and Regional Business Cycles

MICHAEL J. ARTIS
and
MATHIAS HOFFMANN

BADIA FIESOLANA, SAN DOMENICO (FI)

Published in Italy in June 2003
European University Institute
Badia Fiesolana
I-50016 San Domenico (FI)
Italy

Home Bias and the Structure of international and regional Business Cycles.[1]

Michael J. Artis[2] Mathias Hoffmann[3]

First version: June 2002. This version: May 2003

[1]We are grateful to seminar participants at Harvard, the University of Victoria, the European University Institute, University of Kent at Canterbury, Queen's University Belfast and the CEPR-EPRU workshop on International Capital Markets in Copenhagen and the MEA workshop on International Capital Flows in Mannheim. In particular we thank Philip Lane and Oved Yosha for comments and discussions.

[2]Dept. of Economics, European University Institute, Villa San Paolo, Via della Piazzuola 43, 50133 Florence, Italy E-Mail: artis@iue.it

[3]Corresponding author. Full address: Dept. of Economics, University of Dortmund, D-44221 Dortmund, Germany. E-Mail: M.Hoffmann@wiso.uni-dortmund.de
Hoffmann gratefully acknowledges financial support from the Economic and Social Research Council. This paper is part of the ESRC project 'International Capital Flows and Macroeconomic Fluctuations' in the framework of the 'Evolving Macroeconomy' programme (Grant # L138251037).

Abstract

We estimate Shiller portfolio weights for OECD countries and US states. We find that the income of US federal states is derived to about 50 percent from own output, that of OECD countries to about 60 percent.

This suggests that US states display considerable 'home bias at home' and that the international portfolio home bias may be relatively less severe than evidence based on models of optimal portfolio allocation would suggest. We relate the estimated portfolio weights to the structure of regional and international business cycles. In particular, we can reproduce the empirical evidence on inter-state and international income flows.

Keywords: Consumption Risk Sharing, International and regional business cycles, Capital flows, Home Bias, Non-stationary panel data

JEL classification: C23, E21, F36

1 Introduction

This paper links the evidence on international and interregional portfolio holdings to the structure of business cycles.

The empirical literature documents that financial markets are a lot less integrated between countries than within them.[1] Most studies consider regressions of idiosyncratic consumption growth on other idiosyncratic variables, notably relative output growth, to find the associated coefficients are a lot higher in international data than in national (regional) data sets.

This lack of international consumption risk sharing seems to find its natural correspondence in the home bias puzzle of international finance: French and Poterba (1991) but also studies based on more recent data such as Lane (2000) and Kraay, Loayza, Serven and Ventura (2000) find that country portfolios are biased towards domestic assets. In contrast to the relative lack of international consumption risk sharing, the extent of portfolio home bias is typically measured against the benchmark of models of optimal portfolio allocation. Such calculations may tend to overestimate the extent of portfolio diversification that international financial integration may bring along since it is not clear *a priori* that regional portfolios themselves are completeley diversified. That there could be a 'home bias at home'[2] seems plausible once one recognizes that many of the potential explanations of the home bias puzzle that have been suggested are not necessarily specific to countries: labour income is likely to be non-insurable; local equity may provide a better hedge against (local) inflation risk (since its returns are more strongly inversely related to domestic inflation than foreign returns) and international or interregional diversification may be implicit since domestic companies have operations in other regions or countries that generate returns that are a hedge for local consumption.[3]

While this list of plausible causes of a regional home bias could probably be extended, one message from the consumption risk sharing literature is, in fact, that even (US) regions do not share all their idiosyncratic risk. (see e.g. the results in Asdrubali, Sørensen and Yosha (1996)). It is therefore possible and likely that even regional portfolios in well integrated national financial

[1] Asdrubali, Sørensen and Yosha (1996), Sørensen and Yosha (1998), Hess and Shin (1998), Crucini (1999), Lane (2001), Becker and Hoffmann (2001)

[2] Coval and Moskowitz (1999) show that there is a home bias even in portfolios of domestic stocks. U.S. Fund managers seem to prefer equity of locally headquartered companies.

[3] These potential explanations are explored in Lewis (1999) who concludes that non of these factors in itself is sufficient to explain the home bias but that they are likely to interact in a complex way.

markets may display considerable home bias.

Unfortunately, the extent of regional home bias at the regional level cannot be directly assessed since good data on regional portfolios do not exist. Furthermore, even to the extent that these data exist - as they do for countries - the international income flows that these portofolios generate cannot be observed.

The first contribution of this paper is to develop empirical measures of international regional portfolio diversification based on the concept of a Shiller (1993) portfolio of perpetual claims to a country's income. Our approach offers the advantage that regional portfolios can be estimated and compared to those of entire countries. In this way, we overcome the lack of data on portfolio cross-holdings at the regional level. Furthermore, the simplicity of our framework allows us to generate artificial income flow data which can then be compared with actually observed cross-country and cross-regional income flows.

Our evidence is based on US state level data from 1960-90 and data for OECD countries from 1960-2000. We find that US states still have a home bias of around 50 percent, whereas countries have a home bias of 90 percent in the period 1960-90 but only of 60 percent in the period 1980-2000. This provides a new benchmark as to the degree of portfolio diversification that we should expect to see at the international level as financial markets continue to become more integrated. We then ask to what extent the estimated portfolio structure can replicate income flows that are in line with actual GNP and state disposable income data. Our results suggest that the estimated portfolio structure supports income flows that are similar to those observed in the real world. In particular, we can replicate the evidence on the relative roles of *ex-ante* and *ex-post* risk sharing in national and international data.

The observation that both US federal states as well as industrialized countries seem incompletely diversified leads us to the second contribution of this paper: we recognize that if diversification is incomplete, the amount of risk sharing that is eventually achieved will also depend on the structure of idiosyncratic risk, i.e. the structure of business cycles. It is therefore important to ask, to what extent differences in business cycle structure (e.g. w.r.t to synchronization, relative persistence and volatility) between industrialized countries on the one hand and between US states on the other can account for the apparent lack of international consumption risk sharing as it is detected by standard risk sharing regressions.

Here we find that industrialized countries would indeed be more insured if international business cycles had the same features as US regional cycles. However, this lever is not sufficiently important to account for the lack of international consumption risk sharing. Rather, we document that US federal states are a lot more successful in obtaining consumption insurance because they are better insured against permanent variation in their relative outputs. Because insurance against permanent variation cannot be achieved ex-post, i.e. through borrowing or lending, our findings are in line with Asdrubali, Sørensen and Yosha (1996) who find a lot more ex-ante risk sharing in US state level data than in international data and with Kraay, Loayza, Serven and Ventura (2000) who also document that countries' net foreign asset positions are biased towards loans and bonds.

A by-product of our paper is that it can reconcile apparently contradictory findings in the literature on consumption risk sharing between regions and countries. So, Asdrubali, Sørensen and Yosha (1996)[4] find that roughly a quarter of all idiosyncratic output risk remains uninsured in US data. Using a different data set, Crucini (1999) finds that the null that all idiosyncratic risk is insured in US data cannot be rejected. Using the same data set for the two different regression specifications run by ASY and Crucini, we can reproduce their respective results. As our discussion is going to show, the interpretation of the estimated coefficients is different: Crucini estimates how much of the permanent variation in the data gets insured, given a country's portfolio whereas ASY obtain an estimate that measures how much risk is shared, given the country portfolio *and* the structure of business cycles.

The remainder of this paper is structured as follows: in section two we discuss the measurement of risk sharing using consumption data. In section three we introduce the model and its implications for the measurement of risk sharing and diversification. In section four we present our data and estimate the Shiller portfolio weights. In section five we relate the estimated portfolio weights to international income flows and discuss to what extent differences in business cycle structure can account for the apparent lack of risk sharing at the international level. Section six summarizes and concludes.

2 Measuring risk sharing and portfolio diversification

In a world with complete capital markets, countries and regions will insure completely against any idiosyncratic risk. Therefore, *ex post*, there should not be any correlation between a country's or region's relative output and consumption. This fundamental insight was first applied to household level data by Mace (1991) and Cochrane (1991) and Townsend (1994). These authors suggested to run regressions of idiosyncratic consumption growth on

[4] We will henceforth abbreviate the reference to this paper as 'ASY'.

idiosyncratic income growth.

In country-level data, household income is typically replaced by per capita GDP. The basic risk sharing regression to which we will refer throughout this paper will therefore have the form:

$$\Delta c_t^k - \Delta c_t^* = b\left[\Delta y_t^k - \Delta y^*\right] + \varepsilon_t \tag{1}$$

where c_t^k and y_t^k are the logarithms of consumption and GDP in country k respectively and ε_t is a disturbance term. 'Rest of the World'-variables are denoted by an asterisk. Under the null of complete markets this regression should yield a coefficient of zero.

The acknowledgment that real world financial markets are likely to be incomplete in many ways has subsequently led to a more pragmatic approach in applied work. Rather than testing the null of complete markets, i.e. $b = 0$, Asdrubali, Sørensen and Yosha (1996)as well as Sørensen and Yosha (1998) have argued very convincingly that the coefficient in the risk sharing regression may be of interest in itself and that it should be interpreted as a *measure* of the extent of risk sharing. Applying this insight to US state level data, ASY find that roughly a quarter of idiosyncratic output fluctuations remain uninsured. Conversely, Sørensen and Yosha (1998) show that among OECD countries, more than 70 percent of idiosyncratic fluctuations remain uninsured.

The basic risk sharing regression is motivated by a benchmark model with complete markets. The coefficient b is then a natural metric of the deviation from the complete markets outcome. However, one has to be wary not to directly interpret b as a measure of international diversification. In this paper, we introduce a distinction between the degree of diversification and the degree of insurance that a given degree of diversification may achieve. By diversification we mean the share of international assets in a country's total wealth. By insurance we mean to what extent variability in idiosyncratic consumption is shielded from variability in output.

While these two concepts are identical in the long-run, they may mean different things in the short-run: b measures the extent of insurance, but we will see that it does not necessarily measure the degree of international portfolio diversification if diversification is incomplete. If we want to use the difference of country-level and regional measures of risk sharing to learn about the underlying home bias in portfolios, we will therefore need to gauge the extent to which different business cycle structure may account for the differential success in obtaining insurance.

Why do we not base our notion of the degree of diversification on a measure of net foreign assets? Firstly, data on cross-regional portfolio holdings do not exist, at least not on a broader basis and over sufficiently long time horizons. It is therefore impossible to gauge the extent of the national home bias relative to a (potentially present) regional home bias.

Secondly, for our purposes, it may be problematic to classify a given asset as domestic or foreign. While this may be quite simple with respect to the equity of a local barber shop (unless it is run by an international or interregional chain), it is clearly problematic with respect to multinational companies that may be listed on one country's stock exchange but own subsidiaries or plants in many countries.

More importantly, even if we knew the composition of both country-level as well as regional portfolios, we would still need to know the associated interest rates and dividend payments in order to reconstruct international or interregional income flows. This would be a fabulously complex task. We avoid it by using a simple model that allows us to distinguish between portfolio diversification and insurance.

We therefore prefer to use the construct of a portfolio of Shiller-securities to measure a country' s or regions diversification. While such a portfolio may not be directly observable, Shiller (1993) has argued very convincingly that it is a useful abstract concept that can serve as a benchmark. Contrasting this 'portfolio' with standard consumption based measures of risk sharing can shed light on the sources of the home bias and the structure of international and interregional risk sharing arrangements. We turn to a formal exposition of our model in the next section.

3 Country portfolios and business cycle structure

In this paper we adapt the simple theoretical framework suggested by Crucini (1999) to estimate country portfolio shares. The asset market structure is given by a set of Shiller (1993) securities. Each country chooses to allocate its wealth across two assets: one perpetual claim to domestic output and a mutual fund of perpetual claims to world average output. The degree of portfolio diversification is then given by the share of holdings of the world mutual fund in the country's total stock of assets.

In presenting our approach, it will be useful to emphasize the distinction between *ex-ante* and *ex-post* risk sharing.[5] International and interregional diversification pertains to an *ex-ante* sharing of risk through the cross-border holdings of assets.

[5] ASY (1996) were the first to empirically explore this distinction.

Ex post, that is after the realization of this period's uncertainty, consumption should be determined by permanent income. Under full risk sharing, country k's per capita consumption should correspond to world per capita permanent income,which equals per capita world permanent output

$$C_t^k = Inc_t^{kP} = \frac{r}{1+r}\sum_{l=0}^{\infty}\left[\frac{1}{1+r}\right]^l E\left[Y_t^*\right] = Y_t^{*P}$$

where r is the world real interest rate.[6] Once we recognize that *ex-ante* risk sharing may be incomplete country k's permanent income should be a weighted average of world and home permanent output:

$$C_t^k = Inc_t^{kP} = \lambda Y_t^{*P} + (1-\lambda)Y_t^{kP} \tag{2}$$

Here λ can be interpreted as the degree of diversification into a perpetual claim on world average per capita output. Our setup so far is identical to that of Crucini (1999). But Crucini, in keeping with most of the earlier literature uses (2) to motivate a risk sharing regressions in relative growth rates. We will use (2) to derive a panel regression of relative (logarithmic) consumption and output *levels*.

We therefore now proceed to derive an estimable levels regression from (2). Using the fact that world consumption equals world permanent income, we can write

$$\frac{C_t^k}{C_t^*} = \lambda + (1-\lambda)\frac{Y_t^{kP}}{Y_t^{*P}} + \pi_k \tag{3}$$

where we have introduced π_k as a country-specific fixed effect to account for differences in the mean of relative consumption. These may arise due to differences in preferences or due to a country-specific consumption risk premium.

We rewrite

$$\frac{C_t^k}{C_t^*} = 1 + \frac{C^K - C^*}{C^*}$$

and use the approximation

$$\log\left(\frac{C_t^k}{C_t^*}\right) \approx \frac{C^K - C^*}{C^*}$$

[6] Hereafter, we will not repeat the qualification 'per capita', as it will generally be obvious where it should appear.

Doing the same for relative permanent outputs, we get

$$\begin{aligned} 1+\log\left(\frac{C_t^k}{C_t^*}\right) &= \lambda+(1-\lambda)\left[1+\log\left(\frac{Y_t^{kP}}{Y_t^{*P}}\right)\right]+\pi_k \\ \log\left(\frac{C_t^k}{C_t^*}\right) &= (1-\lambda_k)\log\left(\frac{Y_t^{kP}}{Y_t^{*P}}\right)+\pi_k \end{aligned} \tag{4}$$

where π_k can now be removed as a country-specific fixed effect. Having lower-case letters denote logarithms, we can write

$$c_t^k - c_t^* = (1-\lambda)\left[y_t^{kP} - y_t^{*P}\right] + \pi_k \tag{5}$$

where we have used the approximation that $\log\left(Y^P\right) = \left(\log Y\right)^P = y^P$ for which we provide a justification in the appendix.

This equation is reminiscent of Mace (1991), Asdrubali, Sørensen and Yosha (1996), Cochrane (1991) or Crucini (1999). It relates relative consumption to relative output. Under full risk sharing, the coefficient on relative (permanent) output should be zero. The decisive difference vis-à-vis the earlier literature is that equation (5) relates relative log-levels whereas the basic risk sharing regression is formulated in differences.

We will show that the coefficients of the levels regressions are not only numerically different from those of the differenced regressions, but also have a different interpretation: the levels regression identifies the degree of portfolio diversification, whereas the differenced regression tells us how much insurance is achieved, given the degree of portfolio diversification and the structure of business cycles. We discuss this important distinction in the next subsection

Once we have estimated λ, we can then use (2) to generate artificial income data. We then use the decomposition suggested by ASY (1996), to gauge the plausibility of these estimated country and regional portfolio shares by asking whether the artificial income data reproduce the structure of interregional and international income flows observed in actual data. Here, the simplicity of our framework pays off because it allows us to link estimated stocks to observed flows

3.1 Risk sharing regressions: levels vs. differences

Why can we not just difference equation (5) to obtain the well-known basic risk sharing regression?

The reason for this is that in log-linearizing (3) we have made the assumption that the cross-sectional dispersion of output and consumption is not too high, so that

$$\log\left(\frac{C_k}{C^*}\right) \approx \frac{C_k - C^*}{C^*} \text{ and } \log\left(\frac{Y_k}{Y^*}\right) \approx \frac{Y_k - Y^*}{Y^*}$$

Since we believe this approximation only to hold for the average country in the cross-section, it is an approximation of the long-run behaviour of the data. This cross-sectional perspective on risk sharing has not been used in macro studies. It is similar in spirit to the test for market completeness suggested by Cochrane (1991). Cochrane argues that measuring risk sharing based on cross-sectional data is likely to be more robust to changes in relative shadow prices, or non-separabilities in utility. A similar argument applies in our setting: the level (panel) regression is much more likely to be robust to short-term fluctuations in the economic environment, i.e. to the structure of idiosyncratic business cycles.

The logarithmic levels equation (5) will describe a long-run relation in the data.[7] To see that the coefficient estimate in the differenced regression can deviate substantially from that of the levels regression, let us again consider equation (3). Neglecting the country-specific fixed effect we have

$$\frac{C_t^k}{C_t^*} = \lambda + (1-\lambda)\frac{Y_{kt}}{Y_t^*}$$

Dividing by lagged relative levels we can derive the following log-linear approximation (see appendix I):

$$\Delta c_{kt} - \Delta c_t^* = \frac{(1-\lambda)Z_{kt-1}}{(1-\lambda)Z_{kt-1} + \lambda}\left[\Delta y_{kt} - \Delta y_{kt}^*\right] \tag{6}$$

where $Z_{kt-1} = Y_{kt-1}/Y_{t-1}^*$.

Clearly, whenever $Z_{kt} = 1$ or whenever the degree of diversification is zero or unity, the differenced regression will carry the same message as the level-regression. However, for intermediate values of λ, Z_{kt} will be correlated with $[\Delta y_{kt} - \Delta y_{kt}^*]$ and the sample (time-dimension) mean of Z_{kt} may not equal unity. Therefore, the regression of relative consumption growth on relative output growth does not necessarily identify the degree of diversification. To this end, one should rather use the levels equation (5).

[7] The distinction between a long-term levels regression and a short-term differenced regression is similar to the notion of long-term relations in the literature on cointegrated systems. Indeed, interpreting the level relation as a cointegrating relation is perfectly in line with our reasoning. However, we prefer a more general interpretation as a long-run relation in a non-stationary but potentially non-cointegrated panel data set. For a discussion of the econometric issues involved, see Phillips and Moon (2000).

This is not tantamount to saying that the basic risk sharing regression in differences is misspecified. As becomes apparent from (6), the estimate of b will also depend on the covariance structure of $\frac{(1-\lambda)Z_{kt-1}}{(1-\lambda)Z_{kt-1}+\lambda}$ and $\Delta y - \Delta y^*$, i.e. on the structure of business cycles. One of the main arguments we wish to put forward in this paper is that differenced regressions reveal how much risk is shared, given the degree of portfolio diversification, λ, *and* the structure of business cycles, whereas the level equation can help us to identify the underlying degree of portfolio diversification itself.

In a separate subsection at the end of the paper, we provide an application of this insight: we explore if and how the relative roles of permanent and transitory variation in macroeconomic fluctuations both at the international and the regional level matter for consumption insurance. To the extent that it is harder to insure against permanent than against transitory variation in relative outputs, countries could be less insured simply because they experience more persistent shocks to output than do US states. While we find that this channel cannot explain the magnitude of the lack of consumption risk sharing, our findings emphasize the importance of distinguishing between the amount of portfolio diversification and the amount of insurance that this diversification achieves.

We now turn to our empirical implementation and discuss data and econometric issues.

4 Econometric implementation

4.1 Data

We apply our approach two data sets: one for U.S. states and one for a group of 23 OECD countries. All data are annual.

The US-data set is the one also used by Asdrubali, Sørensen and Yosha[8] and is based on gross-state product and income data from the Bureau of Economic Analysis (BEA). Since consumption data at the state level is not available, it is common practice[9] to use retail sales data by state. These retail sales data are rescaled by the share of retail sales in aggregate (US-wide) consumption too obtain measures of state level consumption data. All data are deflated by the US-wide consumpion price index. The US-data range from 1960 to 1990.

[8] The data base is available at Oved Yosha's web page `http://econ.tau.ac.il/research/riskshare/channels/channels.htm`

[9] Asdrubali, Sørensen and Yosha (1996), Hess and Shin (1998) and DelNegro(2002) all follow this approach.

Country-level data are from the Penn World Table, release 6.1 (PWT 6.1.) by Heston, Summers and Aten (2002) and range from 1960 to 2000. All data are in constant (1996) international prices. The countries included in our estimation are:

1. Canada, 2. the United States,3. Japan, 4. Austria, 5. Belgium, 6. Denmark , 7. Finland, 8. France, 9. Germany (West), 10. Greece, 11. Iceland, 12. Ireland, 13. Italy, 14. Luxemburg, 15. Netherlands, 16. Norway, 17. Portugal, 18. Spain, 19. Sweden, 20. Switzerland, 21. United Kingdom, 22. Australia, 23. New Zealand.

Most of these countries are OECD countries and we will refer to them under this label. As regards the US, we follow the general practice in the US regional business cycle literature and include all states except Washington D.C.

We express all data in per capita terms. Rest of the World (RoW) aggregates are the US- or OECD-wide average per capita values. Population data are from the BEA and PWT respectively.

Over the sample period covered by our international data set, international financial markets have become increasingly liberalized. To take account of this change, we will report results obtained for two subperiods: the first covers the period 1960-1990, the second covers 1980-2000. The results we obtain from the first sub period can be compared directly to others in the literature (the studies by Sørensen and Yosha (1998) and Crucini (1999) cover the same period), while the results from the second sub period shoulld provide insights into the effects of the daramatic increase in net international asset positions that took place in the 1990s and 1990s (compare e.g. the data in Lane (2000) and Kraay, Loayza, Serven and Ventura (2001)).

4.2 First results - basic risk sharing regressions

For good econometric reasons, the empirical literature on consumption risk sharing has focused on regressions of growth rates rather than levels. Both idiosyncratic output and consumption are typically found to be integrated processes. The risk sharing regression in levels is likely to be spurious in time series unless there is a cointegrating restriction between the two variables. In panel data differencing also removes any individual specific trend or mean in the variables.

The first line in table 1 provides the results of basic risk sharing regressions for both U.S. and international data. Roughly three quarters of idiosyncratic output variability remains uninsured in country-level data, in the later period (1980-2000), more than 80 percent. Only 15 percent of idiosyncratic variability spills over into consumption in U.S. state level data, results.

The state-level results are somewhat below the estimate obtained by Asdrubali, Sørensen and Yosha (1996) (who use a more elaborate estimation procedure than we do here) but slightly higher than the estimates obtained by Crucini (1999). At the country level, the results are very close to the estimates in Sørensen and Yosha (1998). Hence, our basic risk sharing regressions clearly reproduce the general pattern that is documented in the literature: there is a lot more risk sharing in regional data than there is in country-level data, but even at the regional or state level, risk sharing is not complete.

We will not discuss these results any further but present them here as a point of reference. We will refer back to table 1 as the results of the 'baseline specification', or equivalently, of the 'basic risk sharing regression'.

4.3 Estimating portfolio shares

We now turn to estimating the degree of home bias in country and regional portfolios based on the risk sharing regression in relative log-levels.

Because relative consumption and output levels are likely to be integrated variables, country-by-country regressions could potentially yield spurious results. The spurious regression problem, does not arise, however, once we exploit the panel dimension of the data set, even if the data are not cointegrated (see Phillips and Moon (2000)). In all our panel regressions, we remove country-specific means.

Equation (5) tells us that the identification of λ requires to regress relative consumption levels on relative *permanent levels* of output.

We construct two different measures of y^P and y^{*P} : one is based on a univariate AR(1) process for (log) home and foreign output respectively, whereas the other one models the joint dynamics of (log) foreign and home output by means of a VAR. We then use the following log-linear approximation:

$$\begin{aligned} \log\left[Y_{kt}^P/Y_t^{*P}\right] &= y_{kt} - y_t^* + \mathbf{E}_t\left\{\sum_{l=1}^{\infty}\frac{\Delta y_{t+l}^k - \Delta y_{t+l}^*}{(1+r)^k}\right\} && (7) \\ &= y_t^P - y_t^{*P} && (8) \end{aligned}$$

In the appendix, we describe the construction of y^P and y^{*P} in detail.

To check that our results are robust with respect to the particular way in which we construct the permanent levels, we also run the level risk sharing regression with actual relative log-levels instead of permanent levels values. Since business cycle deviations from trend are small as a percentage of the level variables, we should expect that $y - y^P$ is relatively small. Irrespective of what measure of y^P we use, we should therefore obtain broadly similar results from the regressions with plain and with constructed permanent levels.

We present our results in table 2, both for the 'plain levels' as well as for the permanent levels specifications. The results are very consistent across the different specifications and carry a clear message:

For US federal states, we find a home bias of around 40-50 percent. In international data, we estimate a home bias of 96 percent in the 1960-90 period. For the later period, we estimate $(1-\lambda)$ to be around 0.66, considerably lower than in the 1960-90 period. This latter finding is in line with the evidence reported by researchers who have examined international portfolio holdings directly (Lane (2000), Kraay et al. (2000)) and who report a considerably increase in the net foreign asset positions of OECD countries over the last two decades.

These results suggest that, at least for the 1980-2000 period, in the metric of Shiller portfolio, the home bias for the average industrialized country is not much higher than it is for the average US state.

Figures 1 and 2 further illustrate the evolution of the home bias in OECD data. Figures 1 a) and b) display the link between relative consumption and output in the international data set for the two subperiods. It becomes apparent that the cloud of points is a lot more dispersed in the later period. In figure 2, we superimpose the US-data set (dots) on the OECD data set (crosses), again in the 1960-90 (panel a)and in the 1980-90 (panel b)periods. At least in the later period, the OECD cloud displays virtually the same degree of dispersion as does the US data set.

The estimates in table 2 as well as figure 2 suggest that diversification in the sense of a portfolio of country- or region- Shiller securities is by far less complete in U.S. state level data than would appear from the basic risk sharing regression in table 1. There is certainly a huge home bias in international portfolios, but even at the regional level we find that U.S. citizens own a disproportionate share of the claims to output of the federal state in which they live.

Our results provide a fresh perspective on the home bias by taking account of those components of a nation's or region's output risk that are not traded in financial markets: the equity of small firms or companies is most likely not traded across countries or regions nor are claims to the labour share of national or regional outputs. Our estimates seem to reflect this.

5 Diversification and business cycles

The central question of this paper is how a given degree of portfolio diversification interacts with the structure of macroeconomic disturbances to generate the stylized facts from the basic risk sharing regression in differences. To understand this interaction, it is important to know whether the portfolio shares that we have estimated are in line with cross-border income flows. This helps us to gauge the plausibility of our estimated portfolio shares (that are unobservable, after all) and it should give us further insights into the nature of the home bias in both national and international data.

In the next subsection we therefore compare the flow evidence on *ex-post* and *ex-ante* risk sharing with the evidence on 'stocks' of cross-holdings of assets as it is captured by our estimates of λ.

As we will see, the estimated portfolio shares do indeed support the evidence on cross-regional and international income flows. Since the estimated international portfolio shares for OECD countries and US federal states estimated from data of the last two decades are relatively similar, this raises the question, why the basic risk sharing regression detects so much less insurance in international than in US state level data.

Clearly, according to equation (6), differences in the structure of business cycles could be responsible for this finding. Conversely, it is possible that the existing portfolio structure allows US states to insure against certain types of risk that countries cannot insure against. One important dimension along which this may be the case is the relative contribution of permanent and transitory components in business cycle fluctuations. Insurance against permanent idiosyncratic shocks requires trade in state contingent assets. In this respect, (intra-) national financial markets may indeed be more complete than international ones. We explore this issue in the second part of this section.

5.1 Ex-ante and ex-post risk sharing

ASY (1996) were the first authors to explore empirically to what extent consumption smoothing is achieved through either *ex-ante* insurance or *ex-post* smoothing. While ASY report that more than 40 percent of relative income variability gets smoothed *ex-ante*, Sørensen and Yosha (1998) report that this channel is virtually inactive in international data. The contribution of *ex-post* smoothing, is, however, comparable in both international and national data. Can this evidence be reconciled with estimates of international or interregional portfolio shares of around 50 percent in regional and 20 percent in international data, as we have obtained them from our analysis?

In their work, ASY and Sørensen and Yosha associate the *ex-ante* channel of risk sharing with cross-border capital income flows, as measured by the difference between GDP and GNP (income). *Ex ante* risk sharing can therefore be thought of as income smoothing. *Ex-post* risk sharing is then the smoothing of consumption through borrowing and lending in credit markets.

Our framework formalizes these notions of *ex ante* and *ex post* risk sharing. Income smoothing is achieved through the *ex-ante* diversification into world mutual fund of Shiller securities. Smoothed (or 'pooled') income is given by

$$INC_t^k = \lambda Y_t^* + (1-\lambda)Y_t^k \tag{9}$$

Our estimates of λ appear quite plausible in the sense that US states are by no means fully diversified but clearly much more so than industrialized countries. In order to gauge the quantitative plausibility of our estimates of λ, we now generate artificial income data according to (9), using the estimated portfolio structure. We can then compare these simulated data to actual GNP or US state income data. More specifically, we can ask whether the artificial data give us quantitatively similar relative contributions of *ex-ante* and *ex-post* risk sharing as actual data do.

To assess this issue, we run regressions of the type suggested by ASY (1996):

$$\Delta\left[y_t^k - y_t^*\right] - \Delta\left[inc_t^k - inc_t^*\right] = \beta_a \Delta\left[y_t^k - y_t^*\right] + u_t \tag{10}$$

Here, β_a measures the extent of ex-ante risk sharing and '*inc*' denotes the logarithm of income constructed according to (9).

In analogy, we also run

$$\Delta\left[inc_t^k - inc_t^*\right] - \Delta\left[c_t^k - c_t^*\right] = \beta_p \Delta\left[y_t^k - y_t^*\right] + v_t \tag{11}$$

and β_p now measures *ex-post* risk sharing.

ASY also consider a fiscal transfer channel. We cannot identify such a channel based on our simple theoretical setup. However, we note that most fiscal transfers are not discretionary but based on rules or laws that have been set *ex-ante*. We would therefore argue that fiscal transfers provide mainly *ex-ante* insurance.

Table (3) contains our results and juxtaposes them to the findings obtained in ASY(1996) and Sørensen and Yosha (1998). In U.S. data, we do indeed find that the amount of *ex-ante* risk sharing in actual state (disposable) income data and in the artificial data is almost the same. Our point estimate for the artificial data is also virtually the same as the one obtained in ASY (1996), if one includes fiscal transfers in the *ex-ante* channel. This result is strong evidence that a constant home bias of around 40 to 50 percent in U.S. state level data is also quantitatively very plausible because it supports the (flow) evidence from national accounting data.

Our model is also able to approximate the relative roles of ex-ante and ex-post risk sharing in international data, but only in the early period, i.e. 1960-90. Hence, a home bias of 0.95 is broadly consistent with cross-border income flows that help to smooth between 0 and ten percent of volatility in idiosyncratic output growth.

In the period 1980-2000 however, an estimated degree of diversification of around 0.6 does not support GDP-GNP differentials that can approximate the stylized fact identified by Sørensen and Yosha and replicated here: the virtual shut-off of the ex-ante channel in international data. Our estimates of ex-ante risk sharing turn out to be far too high (0.43).

In table 4 we report the correlations between observed relative GNP growth and the growth rates of relative incomes generated from the estimated model. Relative income and actual relative GNP growth are highly correlated in both periods, but particularly so in the later period. This suggests that our model does indeed replicate the direction of relative income fluctuations very precisely. However, it generates income series that are much too smooth relative to output.

We explore two possible solutions to this problem:

The first solution is to allow for different degrees of diversification acrosss countries. After all, assuming λ to be constant across regions may be an appropriate assumption for a well integrated financial market such as the United States and the fact that we are able to reproduce cross-state income flows with considerable precision supports this view. OECD countries on the other hand, differ considerably in size, population structure and quality of their financial institutions. It appears plausible that λ differs across countries. We experimented with this possibility. Even though our conclusion is negative on this issue, we briefly report on our exercise here:

Letting λ differ across countries raises a technical issue: as long as λ is constant across countries the payoff of the world mutual fund, Y^*, is just the population-weighted average of per capita outputs, hence independent of λ. As soon as we let λ vary, the aggregate feasibility constraint requires that Y^* is weighted with $\{\lambda_k\}_{k=1..K}$. In estimating the portfolio vector $\boldsymbol{\lambda} = \begin{bmatrix} \lambda_1 & \lambda_2 & ... & \lambda_K \end{bmatrix}$ from the levels risk sharing regression we are therefore faced with a non-linear least squares problem:

$$\min_{\{\lambda_k\}_{k=1..K}} \sum_{k,t} [c_{kt} - c_t^*(\boldsymbol{\lambda}) - \lambda_k[y_{kt} - y_t^*(\boldsymbol{\lambda})] \tag{12}$$

In solving this problem we set all $K = 23$ starting values equal to the estimate of λ obtained in table 2. With these starting values, the algorithm converged very quickly. To the extent that the algorithm detected portfolio weights outside the unit interval, we set the weights to zero or unity respectively. Then we reiterated with these new starting values. This procedure typically converged to a solution within three iterations. Changing the starting value did not alter the estimate of $\boldsymbol{\lambda}$, which suggests that the solution to the non-linear problem is unique.

The estimated portfolio weights differed considerably across countries, but it is interesting to note that the cross-sectional mean of the entries of $\boldsymbol{\lambda}$, was around 0.6 in the later period and around 0.9 in the earlier period. While the estimated weights should not be taken too literally (they are available from the authors upon request), they appeared plausible in that smaller, more open economies also appeared more diversified.

We then generated artificial income data to run the same regressions as those reported in table 3. Our findings were virtually identical to those obtained in tables 3 and 4. This suggests that letting λ vary in the cross-section may be broadly consistent with the data (after all we replicate the relative roles of *ex-ante* and *ex-post* risk sharing in the 1960-90 period even with different λ_k), but it cannot help to explain the relative smoothness of fitted relative income vis a vis the data.

While-cross-sectionally different but time invariant portfolio weights cannot explain the lack of *ex-ante* risk sharing in the data, our results are consistent with considerable short-term volatility in national portfolio holdings: Lane (2001) has emphasized that ex-ante risk sharing can take place through capital gains: if the price of foreign assets is appropriately correlated with income uncertainty at home, capital gains could contribute to insure a country's consumption against adverse output shocks. Capital gains made from sales of such assets would not be recorded in national income (GNP). Still they would lead to international capital flows through the current account. If all ex-ante risk sharing takes place through the purchase of (state-contingent) assets with a view to reap capital gains, the wedge between GDP and GNP that measures only capital income (i.e. profit or dividend) flows can be small or even zero.

Our findings are therefore consistent with high turnover in international equity portfolio positions as documented by Tesar and Werner (1995). They are also in line with Kraay and Ventura (2002) who find country portfolio positions to be remarkably stable in the long-run but very volatile in the short run. Our results may help to integrate the risk sharing literature with this recent evidence on international capital flows.

We summarize the results from this sub-section as indicating that the estimated portfolio shares for both the U.S. and the OECD are quantitatively supported by interregional and international income flows. This provides an important link between the evidence on home bias in stocks and the lack of international consumption risk sharing, as measured by correlations between relative output and consumption growth. Our results suggest that the home bias in portfolios at the international level is substantially lower than the risk sharing regression evidence would suggest: In our level regressions, our estimate of OECD home bias $(1 - \lambda)$ is only 20-40 percent higher than that for US states (0.6-0.7 relative to around 0.5). At the same time, basic risk sharing regression suggest that four times more risk is shared among US states than among OECD countries (estimates of b of 0.15 (US) and 0.80 (OECD) respectively. In the last substantive section of the paper we therefore assess to what extent differences in the structure of business cycles can explain this result.

5.2 Insurance of permanent and transitory shocks

One important respect in which business cycle structure may influence the outcome of risk sharing regressions is the relative contribution of permanent and transitory idiosyncratic shocks to output variability. In existing financial markets, it is harder to insure against permanent fluctuations than against transitory shocks. The reason for this is that insurance against permanent variability in consumption requires countries or regions to insure *ex-ante*, whereas transitory fluctuations could be smoothed *ex-ante* or *ex-post*, through borrowing and lending. *Ex-ante* insurance can only be achieved through state-contingent assets, e.g. equity. However, such assets will only exist to the extent that the state of the world, on which they are contingent, is observable. Otherwise, problems of moral hazard or enforceability may arise which may render markets endogenously incomplete.

This suggests two possible explanations for the lack of international consumption risk sharing: firstly, countries could appear less insured because idiosyncratic output fluctuations are more persistent between countries than within them. Against the backdrop of our results from the previous section, this appears as a distinct possibility. After all, portfolio home bias among US states is not that much smaller than in the OECD, but US states appear a lot more insured.

Secondly, the existing portfolio structure may provide more insurance against permanent shocks within countries than between them. This may have to do with different legal frameworks or limited enforceability, as recently argued by Kehoe and Perri (2002), but also with different compositions of national portfolios.

Our framework allows us to distinguish between these two explanations.

We run separate regressions of idiosyncratic consumption growth rates on relative growth rates of the permanent and transitory components of output respectively. The two regressions

$$\begin{aligned}\Delta c - \Delta c^* &= b_P \left[\Delta y^P - \Delta y^{*P}\right] + \xi_t \\ \Delta c - \Delta c^* &= b_T \left[\Delta y^T - \Delta y^{*T}\right] + v_t\end{aligned}$$

then give us two separate measures of how consumption is insured against permanent (b_P) and transitory (b_T) fluctuations. The permanent components are constructed in the way described in the previous subsections and the change in the transitory part is then just $\Delta y_t^T = \Delta y - \Delta y^P$.

In table 5 we present the results from this exercise. In U.S. data we find that 85-95 percent of permanent variability is insured whilst a similar number obtains for transitory fluctuations. In U.S. data we cannot find evidence that there is a qualitative difference between permanent and transitory shocks to output in as far as their degree of insurability is concerned. This result is in line with earlier findings by ASY (1996) who document that idiosyncratic persistence does not seem to have a big effect on the overall extent of insurance in U.S. data but that regions with more persistent idiosyncratic fluctuations rather tend to insure *ex-ante*.

The picture changes quite substantially once we turn to the regression with international data. OECD consumption is less insured against both permanent and transitory shocks than is U.S. consumption. While OECD countries seem to insure against virtually all transitory variation (the respective coefficients are insignificant), the coefficient on permanent output variation tells us that only 50 percent of permanent idiosyncratic output variability is insured at the international level. The difference between our estimate of b_P at the national and the international level is around 0.35 to 0.4. This is almost the difference between the coefficient estimates in the baseline regression. It seems that the home-bias detected in the baseline regression largely reflects a failure of countries to insure against permanent idiosyncratic risk.

Still, while there is a lot less insurance at the international than at the regional level, it is important to note that the baseline regression reports estimates of the overall amount of risk sharing that are below both the degree of risk sharing that we find for either permanent or transitory shocks. This finding is true in both regional and international data. Again, the difference can be ascribed to the structure of business cycles: to the extent that permanent and transitory shocks are not completely insured, the overall extent of insurance that is achieved will also depend on the covariance structure of permanent and transitory fluctuations in output. To understand the anatomy of this result, we will write the coefficient of the baseline regression as a function of the coefficients on the transitory and permanent parts of output respectively:

To keep the following equations tractable, let the tilde denote relative growth rates of a variable, i.e. $\widehat{c} = \Delta c - \Delta c^*$ and $\widehat{y} = \Delta y - \Delta y^*$. Then the regression coefficient b of the baseline risk sharing regression can be written as

$$
\begin{aligned}
b &= \frac{cov(\widehat{c},\widehat{y})}{var(\widehat{y})} = \frac{cov(\widehat{c},\widehat{y}^P) + cov(\widehat{c},\widehat{y}^T)}{var(\widehat{y}^P) + 2cov(\widehat{y}^P,\widehat{y}^T) + var(\widehat{y}^T)} && \text{(13a)}\\
&= \left[b_P + b_T \frac{var(\widehat{y}_T)}{var(\widehat{y}^P)}\right]\left[1 + \frac{2cov(\widehat{y}^P,\widehat{y}^T)}{var(\widehat{y}^P)} + \frac{var(\widehat{y}_T)}{var(\widehat{y}^P)}\right]^{-1} && \text{(13b)}\\
&= [\alpha b_P + (1-\alpha) b_T]\left[\frac{var(\widehat{y}^P) + var(\widehat{y}^T)}{var(\widehat{y})}\right] && \text{(13c)}
\end{aligned}
$$

where

$$
\begin{aligned}
b_P &= \frac{cov(\widehat{c},\widehat{y}^P)}{var(\widehat{y}^P)}\\
b_T &= \frac{cov(\widehat{c}^T,\widehat{y}^T)}{var(\widehat{y}^T)}
\end{aligned}
$$

are the regression coefficients of idiosyncratic consumption on idiosyncratic changes in the permanent and transitory component of output and the weight α is given by

$$
\alpha = \left[1 + \frac{var(\widehat{y}_T)}{var(\widehat{y}^P)}\right]^{-1}
$$

If the covariance between changes in the permanent and transitory component of output is zero, then the coefficient of the baseline regression is a weighted average of the extent of insurance achieved for either permanent or transitory variation in relative outputs. There is, however, no reason to believe that the covariance term will generally vanish. Therefore, the overall extent of insurance that is measured by the baseline Mace-Cochrane-ASY regression is again a function of the structure of macroeconomic fluctuations, in this case the covariance structure of permanent and transitory components (or, in other words: the covariance structure of innovations in trend and cycle). If the term $cov(\widehat{y}^P, \widehat{y}^T)$ is positive, the baseline regression will detect a lower b, hence more insurance; if the covariance is negative, we see less insurance.

While equation (13) has been derived without any restrictions from theory, the intuition behind it can be understood in the framework of the permanent income hypothesis. Suppose a country enjoys a permanent positive idiosyncratic shock against which it was not or only partially insured ex-ante. Assume for simplicity that the country is however completely insured against transitory shocks. According to the relative variant of the PIH we have considered, relative consumption should instantaneously adjust to the new permanent level of relative output. If the adjustment of output to its new permanent level is gradual, then current consumption should overshoot current relative output fluctuations, making consumption appear more volatile than output. But a gradual adjustment in output just means that an increase (decrease) in the permanent level decreases (increases) the transitory component. Hence, the change in the transitory and the change in the permanent component are negatively correlated - and this case, equation (13) would indeed predict that, *ceteris paribus* we find less insurance.

Conversely, if the permanent positive shock is also associated with a positive change in the transitory component, then this implies that current output changes will be larger than permanent changes. Because consumption will mainly adjust to the permanent part, it will react less strongly than current output changes, making consumption appear more insured in the basic risk sharing regression.

Against this backdrop, equation (13) provides a decomposition of the original risk sharing coefficient into a 'business-cycle adjusted' risk sharing coefficient and into a component that accounts for different business cycle structures. The adjusted risk sharing coefficient, b_{adj} is the first term on the RHS of equation (13):

$$b_{adj} = [\alpha b_P + (1-\alpha) b_T] \tag{14}$$

The impact of business cycle structure, i.e. in particular of the covariance between Δy^P and Δy^T, is given by the second term

$$struc = \left[\frac{var(\widehat{y}^P) + var(\widehat{y}^T)}{var(\widehat{y})} \right]$$

Using this decomposition, we can now answer to what extent the apparent lack of consumption risk sharing at the national and the international levels is due to *a)* a relatively low degree of insurance against permanent shocks, *b)* a relatively large component of permanent shocks (as measured by α) or *c)* differences in the covariance of permanent and transitory components of the idiosyncratic business cycle.

Table 6 gives estimates of the share of permanent shocks in the business cycle, α, and of $struc$, the measure of business cycle structure. Our estimates were obtained from the panel regressions based on the VAR-permanent component. As is apparent from the table, US federal states do not have a systemically lower share of permanent variation in their idiosyncratic business cycles. Nor is there a marked difference in business cycle structure, $struc$, between OECD countries and US states. In terms of our measures, US regional and world business cycles seem surprisingly similar.

What is the contribution of business cycle structure to the lack of consumption risk sharing? The third column of table 6 provides the fitted or adjusted b-coefficient, according to (14). In calculating b_{adj}, we have used the estimates from table 5. It turns out that both OECD countries and US federal states would display twice as much risk sharing if there was no correlation between the transitory and the permanent parts of relative output (i.e. if $struc = 1$). However, while business cycle structure seems important in explaining how much risk is shared, US and world business cycles are far too similar for this to make a difference in explaining the lack of international consumption risk sharing. According to the results in table 4, it seems that countries fail to insure against permanent variation in relative outputs, whereas US states are able to obtain insurance against virtually all idiosyncratic fluctuations in permanent income.

6 Conclusion

In this paper we have examined the link between the structure of business cycles and consumption insurance. Consumption based tests suggest that there is much less risk sharing at the international than at the regional level. However, risk sharing, both in international as well as in regional data sets seems incomplete.

To the extent that risk sharing is incomplete, the structure of business cycles may matter for the degree of insurance that is eventually achieved, given the degree of international diversification. We have used the concept of a portfolio of Shiller securities to distinguish between the extent of international diversification (as measured by the value weighted share of a world mutual fund in a country's income) and the degree of insurance. We find that U.S. federal states are much less diversified than simple risk sharing regressions would suggest: their home bias is about 50 percentage points. Conversely, in data from OECD countries, we find a home bias of about 90 percent in the period 1960-90, but of only 60 percent in the period 1980-2000. This is much lower than evidence based on international equity portfolio holdings (French and Poterba (1991)) would suggest. While standard risk sharing regressions find a lot less consumption insurance in international data, our findings also suggest that a sizable component may be due to differences in the structure of international business cycles.

Our estimates of national and regional degrees of diversification are also consistent with flow evidence on international and interregional income flows. In U.S. data, our model exactly matches the evidence on *ex-ante* and *ex-post* consumption insurance as examined by Asdrubali, Sørensen and Yosha. In international data, fitted relative income growth rates derived from our model are highly correlated with relative GNP growth rates but generally less volatile. We find that country heterogenity cannot account for this result. Rather the high volatility of international asset positions, as documented in Tesar and Werner (1995) and Kraay and Ventura (2002), seems to play an important role in explaining this finding.

Our method has also allowed us to investigate whether international macroeconomic fluctuations represent an altogether different type of risk than national business cycles. This could be the case to the extent that idiosyncratic output fluctuations have a bigger permanent component in international data than in regional data. There are good theoretical reasons to believe that permanent shocks are less insurable in existing financial markets than transitory fluctuations. Therefore, different degrees of persistence of asymmetric business cycles may be responsible for the home bias, without markets *per se* being less complete between countries than within. However, we find that such structural differences cannot account for the home bias. Rather, we document that U.S. federal states are much more successful than OECD countries in achieving insurance in particular against permanent idiosyncratic output shocks. The lack of international consumption risk sharing (as measured by basic risk sharing regressions) can almost exclusively be ascribed to this factor.

There is a wealth of studies that have looked at risk sharing at either the U.S. state or at the international level. With only a few exceptions (Hess and Shin (1998), DelNegro (2002)) most of them reach the conclusion that there is much more risk sharing within countries than among them. Most of them also find that consumption risk-sharing is generally incomplete, be it at the international or national level. However, the exact numbers vary widely across studies. While this is partly due to different data sets, the econometric specifications employed are also often quite different (compare e.g. the setup in Asdrubali, Sørensen and Yosha (1996) to Crucini (1999)). This paper has shown that slight changes in specification can have a huge impact on the outcome of risk sharing regressions. In this paper, we have suggested an economic reinterpretation of the typical outcomes from different specifications of the risk sharing regression. In so doing, we have offered a new way to understand how consumption risk is allocated, both at the national as well as the international level.

References

[1] Asdrubali, P., Sørensen, P. E. and Yosha, O. (1996). 'Channels of interstate risk sharing: United States 1963-90.' *Quarterly Journal of Economics*, vol. 111, pp.1081-1110.

[2] Backus, D., Kehoe, P. and Kydland, F. (1992), 'International Real Business Cycles', *Journal of Political Economy* 100(4): 745-775.

[3] Becker, S. O. and Hoffmann, M. (2001) 'International Risk Sharing in the Short-Run and the Long-Run.' European University Institute Discussion Paper, ECO 2001/03.(http://www.iue.it/PUB/ECO2001-3.pdf).

[4] Canova, F. and Ravn, M. (1996). 'International consumption risk sharing.' *International Economic Review*, vol 37, no. 3, pp. 573-601.

[5] Cochrane, J. H. (1991) 'A simple test of consumption insurance', *Journal of Political Economy,* vol. 99(5), pp. 957-976.

[6] Coval, J. D. and T.J. Moskowitz (1999). 'Home bias at home : Local equity preference in domestic portfolios', *Journal of Finance* 54, 2045-74.

[7] Crucini, M. (1999). 'On international and national dimensions of risk sharing', *The Review of Economics and Statistics*, February 1999, 81(1): 73-84

[8] DelNegro, M. (2002). 'Asymmetric shocks among U.S. states',*Journal of International Economics* 56 (2002), pp. 273-297.

[9] Feldstein, M. and Horioka, C., 1980. Domestic Saving and International Capital Flows. *Economic Journal* 90, 314-29.

[10] French, K. and Poterba, J. (1991). 'Investor diversification and international equity markets.' *American Economic Review*, vol. 81(2), pp. 222-26.

[11] Hess, G.D. and Shin, K. (1998). 'Intranational Business Cycles in the United States', *Journal of International Economics*, 44, pp. 289-313.

[12] Heston, A., Summers, R. and Aten, B. (2002) Penn World Table Version 6.1, *Center for International Comparisons at the University of Pennsylvania (CICUP)*, October 2002.

[13] Kehoe, P. and Perri, F. (2000). 'International Business Cycles with Endogenous Incomplete Markets.' *Federal Reserve Bank of Minneapolis, Research Department Staff Report #265* (March), forthcoming in *Econometrica*.

[14] Kraay, A. and Ventura, J. (2000). 'Country Portfolios', NBER Working Paper #7795.

[15] Kraay, A. and Ventura, J. (2002). 'Current Accounts in the Short Run and the Long Run', NBER Macreconomics Annual 2002.

[16] Lane, P. D. (2000). 'International Investment Positions: A Cross-Sectional Analysis' *Journal of International Money and Finance*, 19(4), 513-534.

[17] Lane, P. D. (2001). 'Do International Investment Income Flows Smooth Income?', Weltwirtschaftliches Archiv 137(4), 714-736.

[18] Lewis, K. (1999). 'Trying to Explain Home Bias in Equities and Consumption.' *Journal of Economic Literature*, 37(2), 571-608.

[19] Mace, B. (1991) 'Full Insurance in the Presence of Aggregate Uncertainty', *Journal of Political Economy,* vol. 99(5), pp. 928-956.

[20] Phillips, P. C. B. and Moon, H. R. (1999) 'Linear Regression Limit Theory for Nonstationary Panel Data', *Econometrica* 67 (5), 1057-1111.

[21] Sørensen, B. E. and Yosha, O. (1998). 'International risk sharing and European Monetary Unification.' *Journal of International Economics*, vol. 45, pp. 211-38.

[22] Shiller, R. J. (1993). *Macro Markets: creating institutions for managing society's largest economic risk.* Clarendon Lectures in Economics Series, Oxford University Press.

Appendix

I. Approximating the risk sharing coefficient (equation (6)):

From

$$\frac{C_t^k}{C_t^*} = \lambda + (1-\lambda)\frac{Y_{kt}}{Y_t^*}$$

we get

$$\begin{aligned}
1 + \Delta c_t - \Delta c_t^* &= \frac{\lambda + (1-\lambda)\frac{Y_{kt}}{Y_t^*}}{\lambda + (1-\lambda)\frac{Y_{kt-1}}{Y_{t-1}^*}} \\
&= 1 + \frac{(1-\lambda)}{\lambda + (1-\lambda)\frac{Y_{kt-1}}{Y_{t-1}^*}}\left[\frac{Y_{kt}}{Y_t^*} - \frac{Y_{kt-1}}{Y_{t-1}^*}\right] \\
\Delta c_t - \Delta c_t^* &= \frac{(1-\lambda)Z_{t-1}}{\lambda + (1-\lambda)Z_{t-1}}\left[\frac{Y_{kt}}{Y_t^*} / \frac{Y_{kt-1}}{Y_{t-1}^*} - 1\right] \\
&= \frac{(1-\lambda)Z_{t-1}}{\lambda + (1-\lambda)Z_{t-1}}\left[\Delta y_t - \Delta y_t^*\right]
\end{aligned}$$

q.e.d.

II. Log linearization of permanent values (Equation(7)):

By the definition of Y_t^P we have

$$Y_t^P = (1-R)\sum_{k=0}^{\infty} R^k \mathbf{E}\left(Y_{t+k}\right)$$

where

$$R = (1+r)^{-1}$$

Then we write

$$\begin{aligned}
Y_t^P &= (1-R)Y_t\left[\sum_{k=0}^{\infty} R^k \mathbf{E}\left(Y_{t+k}\right)/Y_t\right] \\
&= (1-R)Y_t\left[1 + \sum_{k=1}^{\infty} R^k \mathbf{E}\left(Y_{t+k}\right)/Y_t\right] \\
&= (1-R)Y_t\left[1 + \sum_{k=1}^{\infty} R^k \mathbf{E}\left(Y_{t+k}\right)/Y_t\right] \\
&= (1-R)Y_t\left[1 + \sum_{k=1}^{\infty} R^k \mathbf{E}\left(1 + \frac{Y_{t+k} - Y_t}{Y_t}\right)\right]
\end{aligned}$$

Now use a log-linear approximation to write

$$\begin{aligned}
Y_t^P &\approx (1-R)Y_t\left[1+\sum_{k=1}^{\infty}R^k\mathbf{E}\left(1+\log Y_{t+k}-\log Y_t\right)\right] \\
&= (1-R)Y_t\left[1+\frac{1}{1-R}-1+\sum_{k=1}^{\infty}R^k\mathbf{E}\left(\log Y_{t+k}-\log Y_t\right)\right] \\
&= (1-R)Y_t\left[\frac{1}{1-R}+\sum_{k=1}^{\infty}R^k\sum_{l=1}^{k}\mathbf{E}\left(\Delta y_{t+k+l}\right)\right]
\end{aligned}$$

where $y_t = \log Y_t$. Then

$$\begin{aligned}
\sum_{k=1}^{\infty}R^k\sum_{l=1}^{k}\Delta y_{t+k+l} &= R(\Delta y_{t+1}) \\
&\quad +R^2(\Delta y_{t+1}+\Delta y_{t+2}) \\
&\quad +R^3(\Delta y_{t+1}+\Delta y_{t+2}+\Delta y_{t+3...} \\
&\quad ...etc. \\
&= R\left(\frac{1}{1-R}\Delta y_{t+1}+R\left(\frac{1}{1-R}\Delta y_{t+2}+...\right)\right) \\
&= \frac{1}{1-R}\sum_{k=1}^{\infty}R^k\Delta y_{t+k}
\end{aligned}$$

So that

$$\begin{aligned}
Y_t^P &\approx (1-R)Y_t\left[\frac{1}{1-R}+\frac{1}{1-R}\sum_{k=1}^{\infty}R^k\mathbf{E}\left(\Delta y_{t+k}\right)\right] \\
&= Y_t\left[1+\sum_{k=1}^{\infty}R^k\mathbf{E}\left(\Delta y_{t+k}\right)\right]
\end{aligned}$$

Appendix III: construction of permanent (log)output levels

We follow Crucini (1999) in comparing two alternative specifications for the permanent components of home and foreign output. First, we consider a univariate AR(1) process in growth rates of home and foreign output.

$$\begin{aligned}
\Delta y_t^k &= \rho_k\Delta y_{kt-1}+v_{kt} \\
\Delta y_t^* &= \rho^*\Delta y_{t-1}^*+v_t^*
\end{aligned} \tag{15}$$

This specification implicitly assumes that there are no spillovers between home and RoW output. We therefore also consider a VAR specification in output growth rates. In this specification, we also take into consideration that the maintained hypothesis in this paper is that aggregate consumption equals permanent income. If this is the case, then aggregate consumption should be a sufficient statistic for expected future levels of output. We therefore use the methodology first suggested by Campbell and Shiller (1989) and include consumption as an endogenous state in the VAR.

Now let

$$\mathbf{x}_t = \begin{bmatrix} y_{kt} & y_t^* & c_k & c_t^* \end{bmatrix}'$$

denote the vector of endogenous variables. Then we estimate the VAR-model

$$\boldsymbol{\Delta}\mathbf{x}_t = \mathbf{A}\boldsymbol{\Delta}\mathbf{x}_{t-1} + \boldsymbol{\varepsilon}_t$$

Using the approximation from appendix II

$$Y_t^P \approx Y_t \left[1 + \sum_{k=1}^{\infty} R^k \mathbf{E} \left(\Delta y_{t+k} \right) \right]$$

we write

$$\begin{aligned} \log \left[Y_{kt}^P / Y_t^{*P} \right] &\approx y_{kt} - y_t^* + \mathbf{E}_t \left\{ \sum_{l=1}^{\infty} \frac{\Delta y_{t+l}^k - \Delta y_{t+l}^*}{(1+r)^k} \right\} \\ &= y_t^P - y_t^{*P} \end{aligned}$$

To construct the relative permanent values $y^p - y^{*P}$ from the VAR-process, we use the Hansen-Sargent prediction formula to get

$$\mathbf{E}_t \left\{ \sum_{l=1}^{\infty} \frac{\Delta y_{t+l}^k - \Delta y_{t+l}^*}{(1+r)^k} \right\} = \mathbf{h}' \left[\frac{\mathbf{A}}{1+r} \right] \left[\mathbf{I} - \frac{1}{1+r} \mathbf{A} \right]^{-1} \Delta \mathbf{x}_t$$

where $\mathbf{h}' = \begin{bmatrix} 1 & -1 & 0 & 0 \end{bmatrix}$. In the case of the AR(1)-process we obtain a similar expression in which $\mathbf{A}$ gets replaced by the relative degrees of persistence of the two processes.

In constructing $y^P - y^{*P}$, we set the real interest rate, r, to 0.02. throughout.

Table 1: Basic Risk Sharing Regressions

Regression	United States 1960-90		OECD 1960-1990		OECD 1980-2000	
$\Delta c - \Delta c^* = b(\Delta y - \Delta y^*)$	0.15	(0.03)	0.79	(0.04)	0.84	(0.09)

Table 2: Estimates of the home bias $(1 - \lambda)$

Regression	United States (1960-90)		OECD 1960-90		OECD 1980-2000	
$c - c^* = (1 - \lambda)(y^P - y^{*P})$						
AR(1)	0.42	(0.02)	0.90	(0.02)	0.61	(0.04)
VAR(1)	0.41	(0.02)	0.94	(0.03)	0.60	(0.04)
$c - c^* = (1 - \lambda)(y - y^*)$	0.50	(0.02)	0.96	(0.02)	0.66	(0.03)

Table 3: Ex-ante and ex-post risk sharing in fitted and actual GNP data

Regression	United States		OECD			
			1960-90		1980-2000	
Data	*ex-ante*	*ex-post*	*ex-ante*	*ex-post*	*ex-ante*	*ex-post*
fitted	0.53	0.31	0.08	0.16	0.43	-0.28
	(0.02)	(0.21)	(0.01)	(0.06)	(0.01)	(0.09)
actual	0.57	0.27	-0.01	0.24	0.01	0.14
	(0.10)	(0.21)	(0.04)	(0.06)	(0.02)	(0.09)
	ASY (1996)		SY98 (1960-90)			
Literature	0.52	0.30	0.03	0.25	—	—

artificial data generated with the portfolio shares estimated from the panel regression with perm. levels based on the VAR

Table 4
OECD: correlations between fitted and actual income growth

Country	1960-90	1980-2000
Canada	—	0.98
USA	0.69	0.93
Japan	0.72	0.99
Austria	0.60	0.97
Belgium	0.56	0.85
Denmark	—	0.95
Finland	0.73	0.98
France	0.61	0.97
Germany	—	0.96
Greece	0.79	0.97
Iceland	0.90	0.98
Ireland	0.84	0.91
Italy	0.72	0.93
Luxemburg	0.65	0.44
Netherlands	0.64	0.89
Norway	0.83	0.98
Portugal	0.76	0.95
Spain	0.77	0.99
Sweden	0.72	0.95
Switzerland	0.74	0.88
United Kingdom	—	0.93
Australia	0.78	0.94
New Zealand	—	0.92

artificial data generated with the portfolio shares estimated from the panel regression with perm. levels based onthe VAR

Table 5: Insurance against permanent and transitory shocks

Regression	United States (1960-90)	OECD 1960-90	OECD 1980-2000
$\Delta c - \Delta c^* = b_P(\Delta y^P - \Delta y*^P)$			
b_P	0.04	0.53	0.41
	(0.17)	(0.04)	(0.08)
$\Delta c - \Delta c^* = b_T(\Delta y^T - \Delta y*^T)$			
b_T	0.12	0.10	0.20
	(0.24)	(0.08)	(0.10)

Table 6: **Business cycle structure, US states and OECD countries**

	share of perm. shocks α	Cov. structure $struc$	BC-'Adjusted' coefficient b_{adj}
United States	0.64	2.2	0.07
OECD			
1960-90	0.73	2.01	0.41
1980-2000	0.61	3.01	0.32

In the calculation of the adjusted b, the estimates from table 5 have been used

Figure 1: The OECD data set

a) OECD countries 1960-90

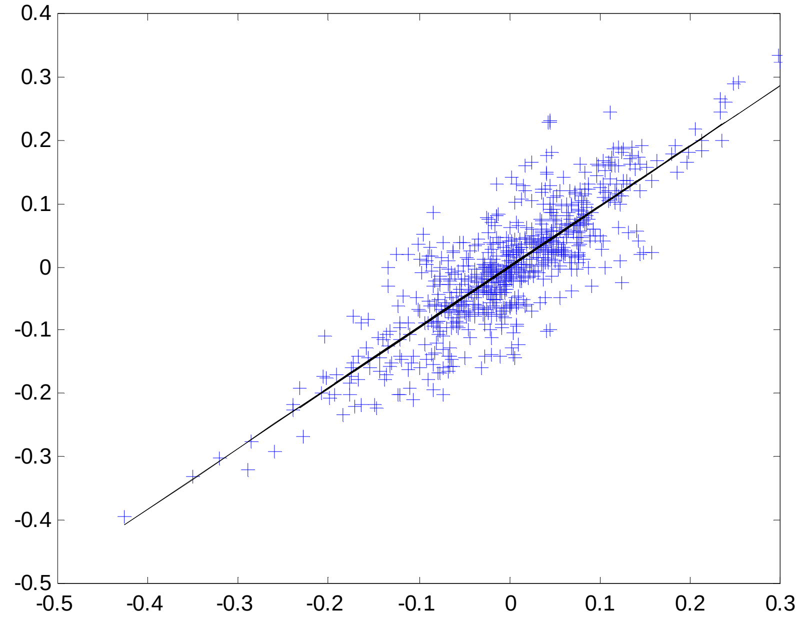

b) OECD countries 1980-2000

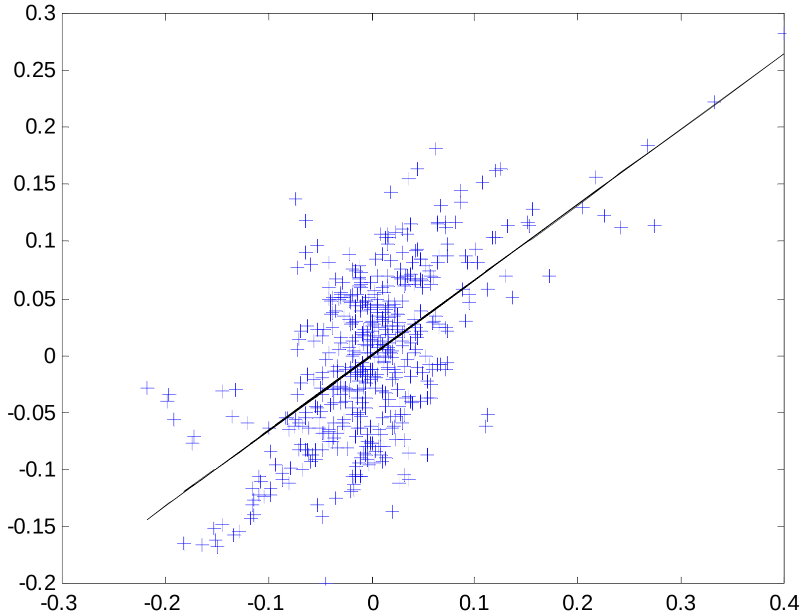

Figure 2: US and OECD data superimposed
a) US and OECD(60-90)data

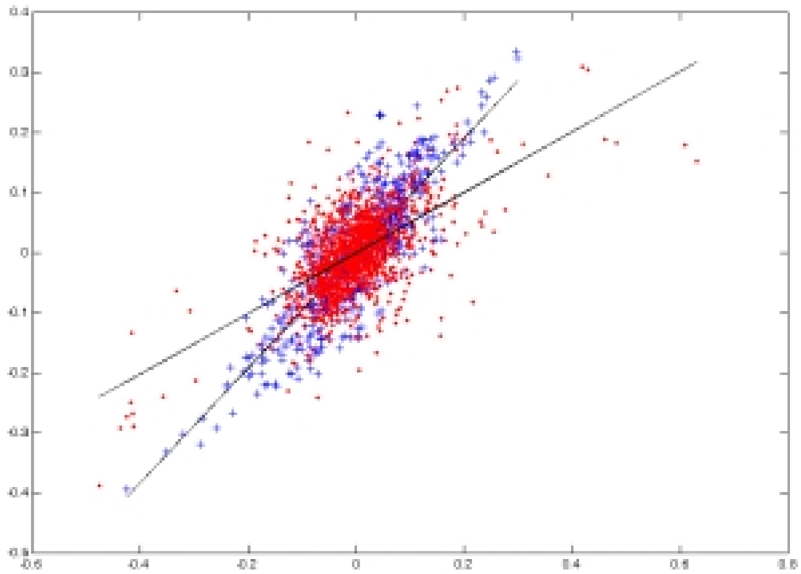

b)US and OECD(1980-2000)data