

Financial Globalization, International Business Cycles and Consumption Risk Sharing¹

Michael J. Artis² Mathias Hoffmann³

September 2004

¹Our work on this paper has benefitted greatly from discussions with Oved Yosha. We also thank seminar participants at the University of Essen, the 2004 ‘Summer at CEU’ Workshop in Budapest and the 2004 ESRC programme conference ‘Understanding the Evolving Macroeconomy’ in Oxford. Special thanks go to our discussant Sebnem Kalemli-Ozcan and to Bent Sørensen.

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

³Corresponding author. Full address: Dept. of Economics, University of Dortmund, D-44221 Dortmund, Germany. E-Mail: M.Hoffmann@wiso.uni-dortmund.de

This paper is part of the project *The International Allocation of Risk* funded by Deutsche Forschungsgemeinschaft in the framework of SFB 475. Early stages of the research on which the paper is based were also funded by an ESRC project grant (Grant # L138251037).

Abstract

Consumption based measures of international risk sharing seem to defy the effects of more than two decades of ongoing financial globalization. We put forward an explanation of this puzzle: under incomplete risk sharing and if there are several sources of risk, consumption based measures of risk sharing will also be a function of the structure of business cycles, i.e. their degree of synchronization and persistence. We argue that permanent and transitory shocks to output constitute such qualitatively different sources of risk. Using OECD data, we then illustrate that countries have indeed become more insured against permanent shocks, in line with the ever better integration of financial markets. Basic measures of risk sharing have however not picked up this change because globalization has also affected the structure of business cycles. In particular, our results are consistent with the observation recently made by several authors that the globalization period has seen the emergence of less volatile and internationally more synchronized business cycles among industrialized countries.

Keywords: Consumption Risk Sharing, International and regional business cycles, Capital flows, Home Bias

JEL classification: C23, E21, F36

1 Introduction

Since the beginning of the 1980s, international financial markets have become increasingly integrated. It would seem that this would unambiguously lead to better international consumption risk sharing. But conventional measures of international consumption risk sharing seem to defy the effects of more than two decades of financial globalization and continue to display low levels of financial market integration. In this paper, we document that consumption risk sharing has actually improved during the globalization period but that the impact of this process on consumption may have been blurred by concurrent changes in the structure of the underlying risks: international business cycles have changed as well.

The literature on consumption risk sharing falls into two broad strands. The first strand emphasizes that in complete financial markets, marginal utility growth should be equated across countries so that consumption growth rates should be highly correlated. This correlation-based approach has encountered the now famous consumption correlation or quantity puzzle: in the data the correlation in consumption is not only low (consistent with low degrees of international risk sharing) but also generally lower than the correlation in the underlying risks, i.e. output growth rates. (Backus, Kehoe and Kydland (1992).

A second strand of the empirical consumption risk sharing literature has focused on regression-based measures. These studies (notably Asdrubali, Sørensen and Yosha (1996), Crucini (1999) and Kalemli-Ozcan, Sørensen and Yosha (2001)) emphasize an alternative prediction of the complete markets model: fluctuations in relative (i.e. idiosyncratic) marginal utility growth should be independent of idiosyncratic risk (as measured by relative output growth rates). Therefore, the coefficient of a regression of relative consumption growth on relative output growth should be zero. Similar regressions were first proposed by Mace (1991), Cochrane (1991) and Townsend (1991) as tests of the null of market completeness, but as argued very convincingly by Asdrubali, Sørensen and Yosha (1996) they are more generally useful: the estimated coefficient, typically between zero and one, can be directly interpreted as a measure of the deviation from the complete markets outcome. Throughout the paper we will refer to such regressions as ‘risk sharing’ (RS) regressions.

Both risk sharing regressions as well as consumption correlations document a lack of risk sharing in international relative to regional data. ¹

¹Generally, the coefficient in risk sharing regressions is lower in international than in regional data (compare e.g. Asdrubali, Sørensen and Yosha (1996) and Sørensen and Yosha (1998)). Consumption correlations are also generally lower in international data.

But they speak much less clearly about how the extent of risk sharing has evolved over time. None of the papers cited above detects a major increase in risk sharing in international data in what we refer to as the globalization period, i.e. after 1980, even though international capital flows as well as cross-holdings of equity and foreign direct investment have seen spectacular growth (see Lane and Milesi-Ferretti (2001, 2004)). Recently, Moser et al. (2003) have investigated risk sharing regressions based on European data and conclude that consumption risk sharing has not increased since capital markets have been liberalized nor since the move to a common currency. Heathcote and Perri (2003) even document a decrease in international consumption correlations for U.S. data. Labhard and Sawicki (2004), using UK regional and international data, equally find a slight decrease in risk sharing based on a factor-analytical approach.

We argue that in order to uncover the impact of financial globalization on consumption risk sharing over time it is important to distinguish between permanent and transitory variation in the country-specific component of output growth. Once risk sharing is incomplete, we are back in a permanent income model, so that consumption should react in qualitatively different ways to permanent and transitory changes in output. We therefore account separately for permanent (trend growth) and transitory (purely cyclical) variation in country-specific business cycles. Based on a data set from 22 OECD countries spanning the period 1960-2000, we detect a considerable increase in risk sharing during the globalization period – i.e. during the 1980s, but particularly during the 1990s. For the period 1960-90, we also compare our results to those obtained from US state level data.

So why does the basic risk sharing regression (or consumption correlations) not pick up the increase in international consumption risk sharing while our method does?

The answer that we offer to this question emphasizes that the response of output to permanent and transitory shocks, along with the structure of the underlying shocks themselves, has changed over time.

First, we document a decrease in the variability of the country-specific component of output growth. This drop in the variance of country-specific output growth is mainly driven by a drop in the covariance between growth in the trend and in the cyclical component of output: it is negative and has fallen considerably. As we discuss, this suggests that the response of relative output to the underlying permanent shocks must have become more gradual. In a permanent income model, a more gradual response of output to

But Hess and Shin (1998) document a quantity puzzle in U.S. state level data, whereas Pakko (1998) casts doubt on the robustness of the quantity puzzle in international data.

a permanent shock does, however, imply that consumption becomes statistically more sensitive to current output changes: in response to permanent shocks, people will want to de-smooth consumption and this effect will bias the coefficient in risk sharing regressions upwards.

Secondly, we also find a decrease in the variability of the global component of trend output growth. This effect tends to lower the international correlations between consumption growth rates, offsetting the increase in correlations that would have occurred if financial integration was the only force at work.

Our results tie in with a recent theoretical and empirical literature that emphasizes that financial globalization may in itself have important effects on international business cycles. In particular, they are in line with the observation that international business cycle have become more synchronized over the last two decades and that business cycle volatility has decreased globally (see e.g. Kose et al. (2003, 2004), Bordo and Helbling (2004)). Our paper is also related to the work of Imbs (2004) who provides cross sectional evidence that financial integration seems to lead to more business cycle symmetry and that it may increase risk sharing. But Imbs also acknowledges that the consumption correlation puzzle does not seem to vanish and that changes in consumption based measures of international risk sharing over time are notoriously hard to detect. We get at this issue in this paper.

The remainder of this paper is structured as follows: in section two we discuss both the regression- and correlation-based approaches to the measurement of consumption risk sharing. Section three offers the theoretical and empirical rationale that leads us to treat permanent and transitory fluctuations in output as qualitatively different risks. We then propose a decomposition of regression-based measures of risk sharing into a business cycle component and into a (business cycle-) adjusted measure of risk sharing. In section four we present our data and the details of the empirical implementation. Section five offers a discussion of our main results. While we find it useful to discuss most results in this paper based on a regression-based approach, this section also applies our framework to the analysis of consumption correlations. Section six concludes.

2 Consumption-based measures of risk sharing

All consumption-based measures of risk sharing that we consider in this paper are motivated by a benchmark model with complete markets. In the

simplest complete markets model, marginal utility growth in country or region k equals the growth in the shadow price of consumption and is therefore equalized across countries:

$$\frac{u'_k(C_{t+1}^k)}{u'_k(C_t^k)} = \frac{\mu_{t+1}}{\mu_t} \quad (1)$$

where $u'(\cdot)$ is the period utility function and C_t^k measures consumption in country k . The shadow price of consumption is μ_t . There are two related readings of this fundamental equation that have both found their reflections in the empirical literature. The first is that marginal utility growth should be perfectly correlated across countries. One branch of the literature therefore looks at consumption correlations. This line of research has encountered the now famous consumption correlation puzzle (Backus, Kehoe, Kydland (1992)) that consists in the stylized fact that international consumption correlations are lower than the corresponding output correlations. Stockman and Tesar (1995) have argued that consumption is likely to be driven by preference shocks and subject to considerable measurement error so that low empirical consumption correlations could in principle arise even in complete markets. While the quantity puzzle is an important stylized fact that calls for a theoretical explanation, the argument by Stockman and Tesar demonstrates that it is difficult to directly interpret consumption correlations as *measures* of risk sharing: the very fact that consumption correlations (as measure of relative marginal utility) are less correlated than the underlying risks (i.e. output) would suggest that people use financial markets to destabilize their relative marginal utilities.

Our main focus in this paper will therefore be on an alternative reading of equation (1) that has equally made a profound impact on the risk sharing literature: marginal utility growth in country k should be independent of country-specific risk-variables. Since growth in the shadow price is common to all countries, the difference between marginal utility growth in two countries should be statistically independent of the relative endowment variables.

In order to obtain an estimable equation, specific assumptions on the form of utility are typically made. Under log-utility, the optimality condition can therefore be written

$$\mathbf{E} [\Delta c_t^k - \Delta c_t^* | \mathbf{X}_t^k] = 0$$

where \mathbf{X} is a vector of idiosyncratic risk factors, such as relative output growth. Therefore, under full insurance, the regression

$$\Delta c_t^k - \Delta c_t^* = \mathbf{b}' \mathbf{X}_t^k + \varepsilon_t$$

should yield a coefficient of zero. Mace (1991) and Cochrane (1991) were the first authors to investigate regression of this type in household-level data.

Clearly, this test of the full risk sharing proposition assumes that \mathbf{X}_t^k is not coincidentally correlated with relative consumption growth for reasons that are unrelated to market incompleteness. For example, in household level data, income and consumption are likely to be highly correlated since a large share of household income is derived from labour. So if there are non-separabilities in utility between consumption and leisure, then income and consumption will both be endogenous and correlated. A similar reasoning could of course be applied at the aggregate level, where a country's output will also depend on the labour supply decision of the representative agent and therefore, there may be a correlation between the relative marginal utility of consumption and relative output movements. However, one of the central messages of the international real business cycle literature (notably Backus, Kehoe and Kydland (1992)) is that the substitution effect between consumption and leisure can only account for a very limited degree of comovement in relative consumption and output at the aggregate level. Therefore, many researchers, including Asdrubali, Sørensen and Yosha, as well as Hess and Shin (1998) and Crucini (1999) have formulated regressions of the form

$$\Delta c_t^k - \Delta c_t^* = b [\Delta y_t^k - \Delta y_t^*] + \varepsilon_t \quad (2)$$

where y^k is the logarithm of output in country k and the asterisk denotes the world average. We call this equation the basic risk sharing equation (RSR). Clearly, in model with complete markets (and against the backdrop of the qualifications given above) the coefficient estimate of b should be zero or close to zero. The acknowledgement that real world financial markets are likely to be incomplete in many ways has led researchers to adopt 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 b may be of interest in itself and that it should be interpreted as a *measure* of the deviation from the complete markets outcome. Applying this insight to US state level data, Asdrubali, Sørensen and Yosha 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. Hence, there is a lack of international consumption risk sharing when risk sharing within countries is taken as the benchmark.

In the next section, we will discuss how the results from the basic risk sharing regression are affected if there are several sources of country-specific

consumption risk. First we will argue why this distinction is particularly important with respect to permanent and transitory variation in $\Delta y - \Delta y^*$. In a separate sub-section we will then discuss what this may imply for the results obtained from the basic RSR.

3 Permanent and transitory shocks

We illustrate our point by employing a simple theoretical framework inspired by Crucini (1999) that acknowledges that incomplete diversification resuscitates a version of the permanent income model. In an economy with risk sharing, a country's income and output do not have to coincide. Rather, by diversifying *ex ante*, a country can buy into a mutual fund that pays world average output as a dividend, so that income growth is a weighted sum of domestic and world average output growth:

$$\Delta inc_t^k = (1 - \omega)\Delta y_t^k + \omega\Delta y_t^*$$

where ω measures the country's degree of diversification into the world mutual fund. Following the standard permanent income framework, consumption growth in the home country should only respond to the growth in permanent income

$$\Delta c_t^k = \Delta inc_t^{kP} = (1 - \omega)\Delta y_t^{kP} + \omega\Delta y_t^{*P} \quad (3)$$

In this model, permanent and transitory shocks are two different kinds of risk: equation (3) implies that the country should (and assumes that it can) actually smooth away any transitory fluctuation in income *ex post* (i.e. after output and income have been observed). Variation in consumption is driven by permanent shocks alone, as is always the case in permanent income models.

Equation (3) also implies a version of the risk sharing regression in which only the permanent component of country-specific output growth will figure. To see this note that for the world as a whole it must be that $\Delta c^* = \Delta inc^{*P} = \Delta y^{*P}$ so that

$$\Delta c_t^k - \Delta c_t^* = (1 - \omega) [\Delta y_t^k - \Delta y_t^*] \quad (4)$$

The model formalizes the idea that insurance against permanent (idiosyncratic) shocks would require countries or regions to insure *ex-ante* (by choosing ω), whereas transitory fluctuations can also be smoothed *ex-post*, e.g. through borrowing and lending. The assumption that transitory fluctuations can actually be smoothed away completely whereas insurance against

permanent shocks is generally incomplete, i.e. ω is between zero and one, is just a metaphor for saying that existing financial markets make it harder for countries to insure against permanent shocks than against transitory fluctuations.

The idea that it may be harder for countries and regions to insure against permanent shocks has by now some tradition in the literature. Insurance against permanent shocks requires *ex-ante* insurance which is possible only through state-contingent assets, e.g. equity. But markets for state contingent assets are more likely to be subject to frictions and market incompleteness than are markets for non-state contingent assets such as bonds or loans. State-contingent assets will only exist to the extent that the state of the world, on which they are contingent, is not too costly to verify. As pointed out in e.g. Kocherlakota (1996), problems of moral hazard or enforceability are particularly likely to arise in such markets which may render them endogenously incomplete.

This line of reasoning has recently also been brought to bear on the lack of international risk sharing: Kehoe and Perri (2001) investigate a quantitative theoretical model in which financial contracts have to be self-enforceable which makes more persistent shocks harder to insure. They posit this model as a solution to Backus, Kehoe and Kydland's (1992) consumption correlation puzzle. Canova and Ravn (1996) find less risk sharing in the lower frequency components of international data than in the business cycle components. Becker and Hoffmann (2003) find fluctuations at the business cycle frequency to be very well insured in both regional and international data. Insurance against permanent shocks, however, is found to be a lot less complete, but in regional data it is still a lot better than it is in international data. They use a cointegrated VAR model to show that insurance *ex post* declines over longer horizons, reflecting the fact that ultimately insurance against permanent shocks is possible only *ex ante*. Their results also suggest that the lack of international consumption risk sharing worsens at longer horizons.

Based on the above considerations, we start our empirical analysis by running separate regressions of idiosyncratic consumption growth rates on relative growth rates of the permanent and transitory components of output respectively. The two regressions

$$\Delta c - \Delta c^* = b_P [\Delta y^P - \Delta y^{*P}] + \xi_t \quad (5a)$$

$$\Delta c - \Delta c^* = b_T [\Delta y^T - \Delta y^{*T}] + v_t \quad (5b)$$

then give us two separate measures of how consumption is insured against permanent (b_P) and transitory (b_T) fluctuations. Under the null, $b_P = (1 - \omega)$

so that the first regressions reveals the extent of diversification whereas b_T should actually equal zero - transitory variation in relative outputs should not have an impact on relative consumption. Requiring that b_T is zero may be too strict a condition to impose. After all, real world financial markets might not allow countries to smooth all of their transitory fluctuations. But as we have argued, insurance against permanent shocks should be harder to obtain than against transitory shocks. Our expectation would therefore be that $b_p = (1 - \omega) > b_T$.

We consider b_P and b_T as the 'primitive' risk sharing coefficients that are associated with the true sources of output risk. We note, however, that the the empirical implementation of (5) does not require us to identify permanent and transitory shocks to output directly. For our argument it is sufficient to identify the *response* of output growth to a potentially varied set of underlying shocks. Clearly, these responses, $\Delta y^P - \Delta y^{*P} \Delta y^T - \Delta y^{*T}$, should not be purely statistical but based on economic theory. In a separate subsection in the empirical part of the paper, we will discuss how we construct the permanent and transitory components involved in our regressions as the change in the present value of lifetime income. Before doing so, we express the risk sharing coefficient b as a function of the 'primitive' risk sharing coefficients b_P and b_T .

3.1 Link to the basic risk sharing regression

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

$$b = \frac{cov(\hat{c}, \hat{y})}{var(\hat{y})} = \frac{cov(\hat{c}, \hat{y}^P) + cov(\hat{c}, \hat{y}^T)}{var(\hat{y}^P) + 2cov(\hat{y}^P, \hat{y}^T) + var(\hat{y}^T)} \quad (6a)$$

$$= \left[b_P + b_T \frac{var(\hat{y}^T)}{var(\hat{y}^P)} \right] \left[1 + \frac{2cov(\hat{y}^P, \hat{y}^T)}{var(\hat{y}^P)} + \frac{var(\hat{y}^T)}{var(\hat{y}^P)} \right]^{-1} \quad (6b)$$

$$= [\alpha b_P + (1 - \alpha) b_T] \left[\frac{var(\hat{y}^P) + var(\hat{y}^T)}{var(\hat{y})} \right] \quad (6c)$$

$$= [\alpha b_P + (1 - \alpha) b_T] \left[1 - 2 \frac{cov(\hat{y}^P, \hat{y}^T)}{var(\hat{y})} \right] \quad (6d)$$

where

$$b_P = \frac{\text{cov}(\widehat{c}, \widehat{y}^P)}{\text{var}(\widehat{y}^P)}$$

$$b_T = \frac{\text{cov}(\widehat{c}, \widehat{y}^T)}{\text{var}(\widehat{y}^T)}$$

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

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

Equation (6) 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 right hand side of equation (6):

$$b_{adj} = [\alpha b_P + (1 - \alpha)b_T] \quad (7)$$

As we will argue the covariance between the permanent and the transitory component can be interpreted as a summary statistics for the dynamic adjustment of output to permanent and transitory shocks. We therefore measure the effect of business cycle structure by the second term, i.e.

$$\phi = \left[\frac{\text{var}(\widehat{y}^P) + \text{var}(\widehat{y}^T)}{\text{var}(\widehat{y})} \right] = \left[1 - 2 \frac{\text{cov}(\widehat{y}^P, \widehat{y}^T)}{\text{var}(\widehat{y})} \right]$$

If the covariance between changes in the permanent and transitory component of output is zero, then the adjusted coefficient just equals the coefficient of the baseline regression, i.e. it is a weighted average of the extent of insurance achieved for either permanent or transitory variation in relative outputs.

But in general, the overall extent of insurance that is measured by the baseline RS regression is a function of the covariance structure of permanent and transitory components (or, in other words: the covariance structure of innovations in trend and cycle). If the term $\text{cov}(\widehat{y}^P, \widehat{y}^T)$ is positive, the baseline regression will detect a lower b , hence more insurance; if the covariance is negative, we see less insurance.

While our decomposition (6) has been derived without any restrictions from theory, the intuition behind it can be understood in the framework of the permanent income hypothesis:

As the model in the previous subsection demonstrates, under incomplete diversification, domestic income will at least in part be determined by domestic output. As a result, permanent fluctuations in relative output lead to permanent shifts in relative income and therefore in relative consumption.

If adjustment to the new permanent level of (relative) output is gradual, (relative) consumption will already adjust today - an instance of Deaton's (1992) paradox. As a consequence, consumption growth may be very volatile in relation to current output growth which will then lead to a high coefficient estimates in the RS regression. 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 in this case, equation (6) would indeed predict that, *ceteris paribus* we find a higher risk sharing coefficient. The conclusion from such a high coefficient estimate would then appear to be that little risk is shared. However, a part of the un-smoothing of consumption is an optimal response of the consumer to the permanent shock, not a failure to share consumption risk.

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 *ceteris paribus* react less strongly than current output changes, making consumption appear more insured in the basic risk sharing regression.

Finally, if the adjustment of output to a its long-run level is immediate, then future realizations of \hat{y}^P are not predictable and therefore, \hat{y}^P cannot be correlated with \hat{y}^T , which by its very nature predicts changes in output. Only in this case will the coefficient of the basic risk sharing regression and the adjusted coefficient coincide.

We therefore interpret the correlation between changes in the permanent and the transitory components of relative output as a summary statistics for the dynamic adjustment of output to permanent shocks. The term $\frac{cov(\hat{y}^P, \hat{y}^T)}{var(\hat{y})}$ is a measure of the optimal degree of consumption smoothing or de-smoothing in response to permanent shocks.

Using this decomposition, we can now attempt to identify the sources of variation in b across space as well as across time: Is the apparent lack of consumption risk sharing at the national and the international levels or the apparent failure of risk sharing regressions to pick up twenty-five years of financial globalization 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 transi-

tory components, i.e. the structure of business cycles?

4 Econometric implementation

4.1 Constructing permanent components

Estimating our set of risk sharing regressions (5) involves the identification of a permanent component of domestic and foreign output growth. We construct these as the changes in the annuity-value of domestic or foreign output according to

$$Y_t^P = (1 - R) \sum_{k=0}^{\infty} R^k \mathbf{E}(Y_{t+k}) \quad (8)$$

where $R = (1 + r)^{-1}$ and r is the world real interest rate.

We do not proxy the expectations involved in constructing Y_t^P directly. Rather we estimate a process for the growth rates of Y (and Y^*). The reason for doing so is twofold: first, our estimating equations (5) are formulated in relative growth rates. Furthermore, macroeconomic data are generally well-described by standard linear processes such as ARs or VAR only after logarithmic transformation. In the appendix we show that the annuity value of permanent log output can still be used as an approximation for the logarithm of the annuity value of output, so that

$$\log [Y_t^P] \approx y_t + \sum_{k=1}^{\infty} R^k \mathbf{E}(\Delta y_{t+k}) = (1 - R) \sum_{k=1}^{\infty} R^k \mathbf{E}(y_{t+k}) = y_t^P \quad (9)$$

and hence

$$\Delta y^P - \Delta y^{*P} = \Delta \left[\log \left(\frac{Y^k}{Y^*} \right) \right]^P \approx \Delta \log \left(\frac{Y^{kP}}{Y^{*P}} \right)$$

In this way we can construct changes in trend component of home and foreign output that are firmly grounded in economic theory but at the same time consistent with the requirements of our econometric specification.

To check our results for robustness, we then follow Crucini (1999) and consider two different specifications for the stochastic processes driving Δy and Δy^* : in the first, home and foreign output follow separate AR(1) processes whereas in the second, home and foreign output and consumption follow a VAR. We then use the estimated processes to compute the expectations involved in the construction of Δy^P and Δy^{*P} in (9).

We note that, by construction, Y^P follows a random walk, almost irrespective of what the driving process of Y_t is. This feature is essential for our purposes here, since a random-walk permanent component is minimal in the sense that in any other trend-cycle decomposition the trend will contain a predictable component. But the presence of a predictable component implies that there is scope for further consumption smoothing through borrowing and lending. The hypothesis explored in this paper, however, is that permanent and transitory fluctuations in output and therefore income are qualitatively different with respect to their degree of insurability. This qualitative difference stems from the very fact that ultimately, permanent fluctuations in output *cannot* be smoothed through borrowing and lending and therefore have to be insured through more sophisticated assets such as equity.² This is why it is appropriate to restrict our attention on a permanent component that follows a random walk.³

4.2 Data and estimation

Our empirical findings are based on two data sets: one for U.S. states and one for a group of 22 OECD countries. All data are annual.

The US-data set is the one also used by Asdrubali, Sørensen and Yosha⁴ 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 practice⁵ to use retail sales data by state. These retail sales data are re-scaled by the share of retail sales in aggregate (US-wide) consumption to obtain measures of state level consumption data. All data are deflated by the US-wide consumption price index. The US-data range from 1960 to 1990.

²See e.g. Baxter and Crucini (1995) who show that the dynamic properties of bonds-only and complete markets economies are different only for very persistent shock processes.

³One apparent alternative to our procedure would be to check the robustness of our results to various, purely statistical, filtering methods. However, being eclectic in the choice of filter would entail a serious risk of mismeasurement in our application. This is particularly easy to illustrate in the case of the HP filter. The permanent component obtained from an HP-filter is smooth for standard parameter values. It does contain a random walk but it does not *equal* a random walk. Rather, its smoothness implies a high degree of predictability. Therefore changes in the HP permanent component will not generally correspond closely to changes in the theoretically relevant permanent component which is the annuity value of output.

⁴The data base is available at Oved Yosha's web page <http://econ.tau.ac.il/research/riskshare/channels/channels.htm>

⁵Asdrubali, 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. Netherlands,
15. Norway,
16. Portugal,
17. Spain,
18. Sweden,
19. Switzerland,
20. United Kingdom,
21. Australia,
22. 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.

Our main objective is to ascertain to what extent changes in the structure of business cycles over the globalization period could have blurred the effect that financial market integration may have had on consumption risk sharing.

We therefore report results obtained from the international data set for three subperiods: the first covers the period 1960-1990, the second covers 1980-2000. Finally, we also look at the 1990s alone. 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 and third subperiods should provide insights into the effects of the dramatic increase in net international asset positions that took place in the second half of the 1980s and in particular throughout the 1990s (compare e.g. the data in Lane (2000) and Kraay, Loayza, Serven and Ventura (2001)). We refer jointly to the second and third periods as the globalization period.

We estimated all risk sharing regressions with a panel two-staged least squares procedure. First, we removed country-specific fixed effects. Then we estimate the panel by OLS. To control for heteroskedasticity, all variables are then weighted with the country-specific variance of the first stage residuals and the model is re-estimated.

5 Empirical Results

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 according to the results obtained from U.S. state level data.

The state-level results are somewhat below the estimate obtained by Asdrubali, Sørensen and Yosha (1996) 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. Note further that in international data the estimate for the globalization period is actually higher than that obtained for the earlier period. Even though the difference between the coefficients may not be significant, the effects of financial globalization do not seem to work through to the risk sharing regression. This result is also robust to changes in the sample period. If we restrict the sample range to the 1990s alone, the picture does not change either - again we estimate a coefficient of 0.85.

These results constitute our 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’. We now turn to dissecting these results: can differences in the structure of business cycles account for the failure of financial globalization to materialize in risk sharing regressions?

5.1 Insurance of permanent and transitory shocks

In table 2 we present the results from the risk sharing regressions on permanent and transitory output variation. In U.S. data we find that only 5-15 percent of permanent variability remain uninsured. A similar number obtains for transitory fluctuations. Hence, 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 Asdrubali, Sørensen and Yosha (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 in particular against permanent shocks than is U.S. consumption. While OECD countries seem to insure against virtually all transitory variation (the respective coefficients are insignificant in all three subperiods, at least for the VAR), the

coefficient on permanent output variation tells us that only 50 percent of permanent idiosyncratic output variability is insured at the international level. But this coefficient has also come down quite markedly over time. In the period 1980-2000 it is 0.43 (0.42 for the AR) already considerably lower than for the 1960-2000 period even though the difference is barely significant. In the 1990s however, there is a dramatic drop that is also significant vis-a-vis the 1960-90 period.

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. The difference can only be ascribed to the structure of business cycles as given by α , the share of permanent shocks in the total output variance, and ϕ . We now turn to this part of our results.

5.2 Effects of business cycle structure

Table 3 reports the decomposition of the panel risk sharing coefficient into the share of permanent shocks in the business cycle, α , the measure of business cycle structure, ϕ , and the business-cycle adjusted risk sharing coefficient b_{adj} . In calculating the covariances involved in the calculation of α , ϕ and b_{adj} , we used the permanent and transitory components of output growth obtained from the VAR along with the corresponding coefficients b_p and b_T from Table 2.⁶

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, ϕ , between OECD countries and US states, at least for the 1960-90 periods. In terms of our

⁶To obtain an exact decomposition of the panel risk sharing coefficient according to (6), we calculate the variances and covariances involved in α , ϕ and b_{adj} based on the stacked vector of panel observations. Hence our measure of e.g. the variance of idiosyncratic output – the term $var(\hat{y})$ – is obtained as the variance of $\hat{\mathbf{y}} = vec[\hat{\mathbf{y}}^1, \hat{\mathbf{y}}^2, \dots, \hat{\mathbf{y}}^K]$ where $\hat{\mathbf{y}}^k = [\hat{y}_1^k, \dots, \hat{y}_T^k]'$ is the vector of observations for country k and where T is the time dimension of the panel and K the number of countries or regions. We interpret the measures of $var(\hat{y})$, $var(\hat{y}^P)$, $cov(\hat{y}^P, \hat{y}^T)$ etc. obtained in this way as the variances / covariances for the representative country in the panel. We note that these are not identical to the cross-sectional average. Since we remove country-specific means in our panel estimation, it is easy to show that the link between e.g. the variance of the panel regressor and the average of the country-specific variances is given by the identity $var(\hat{\mathbf{y}}) = \frac{KT-1}{KT-K} \left[\frac{1}{K} \sum_{k=1}^K var(\hat{y}^k) \right]$.

measures, US regional and world business cycles seem very similar. There is no way to explain away the lack of international risk sharing by differences in the structure of business cycles.

In as far as changes over time are concerned, however, it becomes immediately apparent that the structure of international business cycles has changed dramatically over the last two decades and that this change has been particularly marked after 1990.

The share of permanent shocks to idiosyncratic fluctuations has decreased markedly, from 0.77 to 0.55. This finding that the random walk component in idiosyncratic fluctuations has decreased over time is in line with the results obtained by many authors who document that international business cycles may have become less volatile over time (see e.g. Kose et al. (2003)). In line with this development as well as with the better insurance against permanent shocks, the business-cycle adjusted risk sharing coefficient drops dramatically, from 0.45 to around 0.19 during the 1990s - financial globalization can, after all be detected in consumption data.

But why then does the basic risk sharing regression not reveal this change? Turning to Table 3 again, we see that our indicator ϕ has increased from around 2 to more than 4. Our findings imply that the last term on the right hand side, $cov(\hat{y}^P, \hat{y}^T)/var(\hat{y})$, must have fallen from around -0.5 in 1960-90 to around -1.5 during the 1990s. The first thing that is noteworthy here is that the correlation between permanent and transitory components of idiosyncratic output growth is negative. This is in line with one of the scenarios we discussed earlier: if permanent shocks to output trigger a monotonic but gradual response, this will generate a negative correlation between the permanent and transitory components. As a result, the risk sharing coefficient is biased upwards, implying that too little risk sharing is detected.

In Table 4 we decompose the variance of idiosyncratic output growth into the variance of the permanent and transitory components as well as into the correlation component according to

$$\begin{aligned} var(\hat{y}) &= var(\hat{y}^P) + var(\hat{y}^T) + 2cov(\hat{y}^P, \hat{y}^T) \\ &= var(\hat{y}^P) + var(\hat{y}^T) + 2\rho\sqrt{var(\hat{y}^P)var(\hat{y}^T)} \end{aligned}$$

where ρ is the correlation between \hat{y}^P and \hat{y}^T .

It is apparent that the variance of idiosyncratic output growth has decreased quite dramatically during the globalization period, implying that world business cycles have overall become more synchronized. This decrease in overall idiosyncratic volatility has two main sources: on the one hand, the volatility of permanent fluctuations has decreased (second column), on

the other hand, the correlation between the permanent and the transitory component, ρ , which was negative throughout, has decreased further (fourth column). Third, note that the variance of the transitory component has increased as well.

The first and the third observations are in line with the decreasing variance share of permanent fluctuations reported above and is tantamount to saying that the random walk component of idiosyncratic business cycles has become less important over time, so that to the least the idiosyncratic part of business cycles has become less volatile. The second, i.e. the increase in the correlation between the permanent and transitory components indicates that at least some of the increase in the variance of the transitory component reflects a more sluggish adjustment of output to permanent shocks, implying bigger transitory components and a correlation between permanent and transitory components that is closer to minus unity. Note also that not only the correlation has decreased but that the variance of the transitory component has increased by more than the permanent has fallen so that the covariance term $cov(\hat{y}^P, \hat{y}^T)$ must also have decreased. Recall that we can interpret the term $\frac{cov(\hat{y}^P, \hat{y}^T)}{var(\hat{y})}$ as measuring the optimal amount of smoothing or unsmoothing of permanent shocks.

Taking stock of these findings, it seems that the decrease in the volatility of country-specific business cycle components has two partly offsetting effects: first, the share of permanent idiosyncratic variability has decreased considerably. We would therefore expect to see more risk sharing because countries are better insured against transitory fluctuations. But this effect is outweighed by the fact that the impulse response to permanent shocks seems to have become more protracted: permanent shocks covary more with the business cycle, suggesting that adjustment to permanent shocks has become a lot more gradual, again in line with the view that world business cycles have become more stable and synchronized. But the more gradual response to permanent fluctuations offers a lot more scope for optimal de-smoothing, and induces the risk sharing coefficient to signal less diversification.

5.3 Implications for consumption correlations

So far, our analysis has concentrated on the impact that changes in the structure of business cycles may have on risk sharing regressions. It is now straightforward to use our framework to explore some implications for consumption correlations.

The first part of table 5 presents the average consumption correlation for the U.S. and for the OECD countries for the three subperiods. The

second row of Table 5 also presents the average correlation of output. As is apparent, consumption correlations for the U.S. are not that much higher than international correlations. Both international *and* regional correlations indicate a significant lack of risk sharing. Note further that, very much as risk sharing regressions, consumption correlations do not indicate an increase in international risk sharing over time. There is also a manifest quantity anomaly, both in U.S. state level and the international data sets and the anomaly does not seem to disappear over time. As we have argued, the quantity puzzle limits the interpretability of consumption correlations as a measure of risk sharing. Still, one should expect that financial integration would have led to an increase in consumption correlations. But, very much as in the case of the risk sharing regression, changes in the structure of world business cycles could have offset this effect.

We build on Rigobon and Forbes (2001) to calculate a consumption correlation measure that corrects for changes in the structure of business cycles. Note that our simple theoretical model implies a factor structure for domestic and international consumption:

$$\begin{bmatrix} \Delta c_t^* \\ \Delta c_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \omega & 1 - \omega \end{bmatrix} \begin{bmatrix} \Delta y_t^{*P} \\ \Delta y_t^P \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 1 - \omega \end{bmatrix} \begin{bmatrix} \Delta y_t^{*P} \\ \Delta y_t^P - \Delta y_t^{*P} \end{bmatrix} \quad (10)$$

where Δy_t^{*P} can be interpreted as the 'common factor' in domestic and foreign consumption growth.

As (10) shows, changes in consumption correlations can come about either due to changes in the relative variance of common or idiosyncratic shocks to trend output growth or as a change in ω .

Now assume that between the baseline (pre-globalization) period 1960-90 and the globalization period, the variance of the common factor, $var(\Delta y_t^{*P})$ changes so that the new variance is

$$var(\Delta y_t^{*P})_{t>1990} = (1 + \delta)var(\Delta y_t^{*P})_{t<1990}$$

Then, as Rigobon and Forbes show, the correlation coefficient implied for the globalization period can be written as

$$\rho_{imp} = \rho \sqrt{\frac{1 + \delta}{1 + \delta \rho^2}} \quad (11)$$

where ρ is the correlation coefficient from the baseline period. If the imputed and the actual correlation coefficient in the globalization period are

significantly different, then ω must have changed - provided that the variance of the idiosyncratic component has not changed as well.

In the third and fourth rows of table (5) we present the standard deviation of the common factor, i.e. of trend output growth along with the implied correlation coefficient from (11). The common output (trend-) growth component has become considerably less variable over time, whereas from table 4 we see that the idiosyncratic part of trend output growth displays approximately constant volatility between the baseline and the globalization periods. Accordingly, keeping ω constant, international consumption correlations should actually have been falling in the globalization period. This is what we see: The implied consumption correlations are lower than in the baseline period, both in the 1980-2000 as well as in the 1990-2000 periods. The effect is particularly dramatic in the nineties, where the implied correlation is 0.14 is also considerably lower than the correlation of 0.34 observed in the data. Given that the variability of the idiosyncratic factor has remained relatively constant, the fact that observed correlations have remained constant can therefore only be explained by an offsetting increase in ω .

How can these findings be interpreted? Recently a number of authors have documented that the variability of global output shocks has decreased during the globalization period. Otrok, Kose and Whiteman (2004) identify the period between the demise of the Bretton-Woods system as the common shocks period, characterized by the two major oil shocks as well as by a common disinflationary policy response to them. The globalization period has not seen comparably large common shocks. Our correlation results are also in line with those reported by Heathcote and Perri (2003) who find that the U.S. economy has become considerably less symmetric with respect to the rest of the world since the mid-eighties.

We therefore conclude that, very much as for the case of risk sharing regressions, the effect of financial globalization on consumption correlations has been virtually offset by concurrent changes in the structure of world output cycles. But note that the drivers of these offsetting effects seem to be different for risk sharing regressions and correlations: consumption correlations have not increased because *common* movements in trend output growth have become less volatile. The risk sharing regression coefficient, on the other hand, does not pick up improvements in risk sharing because the response to permanent but *idiosyncratic* shocks has changed.

6 Conclusion

Consumption based measures of international risk sharing seem to defy the effects of more than two decades of ongoing financial globalization. In this paper, we put forward an explanation of this puzzle: under incomplete risk sharing and if there are several sources of risk, standard consumption based measures of risk sharing, are not only a function of the degree of diversification, but also of the structure of business cycles.

We have argued that permanent and transitory shocks to output constitute such qualitatively different sources of risk and that they should be treated separately in risk sharing regressions. Therefore, differences in the structure of business cycles (defined as the relative contribution of permanent and transitory shocks, their degree of comovement etc.) rather than differences in the structure of financial markets can be responsible for differences in risk sharing coefficients (or consumption correlations) in different samples.

In principle, this insight can be applied to different samples over space and time. While we find that differences in the structure of business among OECD countries on the one hand and U.S. federal states on the other hand cannot explain away the lack of consumption risk sharing at the international level, we *do* find that changes in international business cycles over time have indeed blurred the effect of financial globalization on consumption.

Using OECD data, we illustrate that countries have become more insured in particular against permanent idiosyncratic shocks, in line with the ever better integration of financial markets. Basic measures of risk sharing have however not picked up this change because the structure of international business cycles has changed at the same time:

Firstly, country-specific factors have become less important for the business cycles of most economies over the globalization period. Secondly, the global component of output fluctuations has become less volatile. Both of these stylized facts are in line with the findings of many other researchers.

The first stylized fact affects regression-based measures of consumption risk sharing. We have shown that, for the average OECD economy, the decrease in the importance of country-specific fluctuations is due to a more gradual response of output to permanent shocks. But a more gradual response to permanent shocks also opens the possibility for optimal consumption de-smoothing, an instance of Deaton's paradox. This effect will induce an upward bias in risk sharing regressions and can explain why the regression-based approach to measuring consumption risk sharing has not picked up the effects of financial globalization: world business cycles have changed in a way that makes less smoothing optimal.

The second stylized fact affects consumption correlations. Since world

trend output growth is the common component in domestic and rest-of-the-world consumption, international consumption correlations will fall *ceteris paribus* if trend output growth becomes less volatile. This effect offsets the increase in consumption correlations that is due to financial integration.

Virtually all measures of financial integration tell us that financial globalization is happening. So far, consumption-based measures were an exception to this rule. This is troubling, because the very *raison d'être* of financial markets is the allocation of consumption risk and any assessment of the welfare benefits of financial integration will ultimately amount to a statement about consumption. It is therefore crucial to identify the factors that may have led international consumption comovements to defy the predictions of a wide class of theoretical models. To our knowledge, the results in this paper provide the first account of these factors: once we control for changes in international business cycles, consumption growth rates have become more correlated and relative consumption growth has become more independent from relative output growth – very much as theory would predict.

We emphasize that this finding has nothing to say about whether the changes in international business cycles that we and others have identified over the globalization period could not by themselves be the outcome of financial or trade integration. Clearly, there are good theoretical reasons to believe that this may indeed be the case.

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] Baxter, M. and Crucini, M. J. (1995), 'Business Cycles and the Asset Structure of Foreign Trade', *International Economic Review* 36(4):821-853.
- [4] Becker, S. O. and Hoffmann, M. (2003) 'Intra- and International Risk Sharing in the Short-Run and the Long-Run.' CESifo Discussion Paper no. 1111.

- [5] Bordo, Michael D. and Thomas F. Helbling (2004). 'Have National Business Cycles Become More Synchronized?', forthcoming in Siebert, H. and R. Langhammer *Macroeconomic Policies in the World Economy*.
- [6] Canova, F. and Ravn, M. (1996). 'International consumption risk sharing.' *International Economic Review*, vol 37, no. 3, pp. 573-601.
- [7] Cochrane, J. H. (1991) 'A simple test of consumption insurance', *Journal of Political Economy*, vol. 99(5), pp. 957-976.
- [8] Crucini, M. (1999). 'On international and national dimensions of risk sharing', *The Review of Economics and Statistics*, February 1999, 81(1): 73-84
- [9] Deaton, A. (1992) *Understanding Consumption*, Oxford. Clarendon Lectures in Economics Series. Clarendon Press.
- [10] DelNegro, M. (2002). 'Asymmetric shocks among U.S. states', *Journal of International Economics* 56 (2002), pp. 273-297.
- [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] Forbes, Kristin and Roberto Rigobon (2001). 'Measuring Contagion: Conceptual and Empirical Issues' in (eds). Stijn Claessens and Kristin Forbes. *International Financial Contagion*, Kluwer Academic Publishers. 2001.
- [14] Heathcote Jonathan and Fabrizio Perri (2004) 'Financial Globalization and Real Regionalization', forthcoming, *Journal of Economic Theory*
- [15] Imbs, Jean (2004). 'The Real Effects of Financial Integration', mimeo, Princeton University and London Business School.
- [16] Kalemli-Ozcan, S., B. Sørensen, and O. Yosha, (2001). Economic Integration, Industrial Specialization, and the Asymmetry of Macroeconomic Fluctuations. *Journal of International Economics* 55 (1): 107-137.

- [17] Kalemli-Ozcan, S., B. Sørensen, and O. Yosha (2003). Industrial Specialization and the Asymmetry of Shocks Across Regions. *American Economic Review* 93 (3): 903–918.
- [18] 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*.
- [19] Kocherlakota, N. R. (1996). 'Implications of Efficient Risk Sharing without Commitment', *Review of Economic Studies*, vol 63(4), 595-609.
- [20] Kose, M. Ayhan. , Cristopher Otrok, and Charles. H. Whiteman (2003). 'International Business Cycles: World, Region, and Country-Specific Factors.' *American Economic Review* 93 (4): 1216–1239.
- [21] Kose, M. Ayhan. , Cristopher Otrok, and Charles. H. Whiteman (2004). 'Understanding the Evolution of World Business Cycles', paper presented at the Summer at CEU workshop, Budapest June 2004.
- [22] Kraay, A. and Ventura, J. (2000). 'Country Portfolios', NBER Working Paper #7795.
- [23] Kraay, A. and Ventura, J. (2002). 'Current Accounts in the Short Run and the Long Run', NBER Macroeconomics Annual 2002.
- [24] Lane, Philip D. (2000). 'International Investment Positions: A Cross-Sectional Analysis' *Journal of International Money and Finance*, 19(4), 513-534.
- [25] Labhard, Vincent and Michael Sawicki (2004). 'International and Intranational Consumption Risk Sharing: The evidence for the UK and OECD', forthcoming as *Bank of England Working Paper*.
- [26] Lewis, K. (1999). 'Trying to Explain Home Bias in Equities and Consumption.' *Journal of Economic Literature*, 37(2), 571-608.
- [27] Lane, P. D. (2001). 'Do International Investment Income Flows Smooth Income?', *Weltwirtschaftliches Archiv* 137(4), 714-736.
- [28] Lane, P. D. and G. M. Milesi-Ferretti (2001). 'The External Wealth of Nations: Measures of Foreign Assets and Liabilities for Industrial and Developing Nations', *Journal of International Economics* 55, 263-294.

- [29] Lane, P. D. and G. M. Milesi-Ferretti (2004). 'International Investment Patterns', mimeo Trinity College Dublin and International Monetary Fund.
- [30] Lewis, K. (1999). 'Trying to Explain Home Bias in Equities and Consumption.' *Journal of Economic Literature*, 37(2), 571-608.
- [31] Mace, B. (1991) 'Full Insurance in the Presence of Aggregate Uncertainty', *Journal of Political Economy*, vol. 99(5), pp. 928-956.
- [32] Moser, Gabriel , Wolfgang Pointner and Johann Scharler (2003). 'International Risk Sharing in Europe: Has Anything Changed?' mimeo, Foreign Research Department, Austrian National Bank.
- [33] Pakko, Michael R. (1998). 'Characterizing cross-country consumption correlations', *Review of Economics and Statistics*, 80(1), pp. 169-174.
- [34] Sørensen, B. E. and Yosha, O. (1998). 'International risk sharing and European Monetary Unification.' *Journal of International Economics*, vol. 45, pp. 211-38.
- [35] van Wincoop, E. (1999) 'How big are potential welfare gains from international risk sharing?' *Journal of International Economics*, vol. 47, pp. 109-135.

Appendix: Construction of permanent GDP-values

By the definition of Y_t^P we have

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

where

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

Then we write

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

We can think of $\sum_{k=1}^{\infty} R^k \mathbf{E} \left(1 + \frac{Y_{t+k} - Y_t}{Y_t} \right)$ as the percentage of total output that is transitory and can therefore be associated with the business cycle component. Since business cycles account for a small percentage of the level of GDP, we can use the log-linear approximation:

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

where $y_t = \log Y_t$.

Then

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

So that

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

and therefore with $y_t = \log Y_t$

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

We now 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.

$$\begin{aligned}
\Delta y_t^k &= \rho_k \Delta y_{kt-1} + v_{kt} \\
\Delta y_t^* &= \rho^* \Delta y_{t-1}^* + v_t^*
\end{aligned} \tag{12}$$

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 (relative) consumption as an endogenous state in the VAR.

Now let

$$\mathbf{x}_t = [y_{kt} \quad y_t^* \quad c_k - c_t^*]'$$

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

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

Using the approximation from above

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

we write

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

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 $\mathbf{h}' = [1 \quad -1 \quad 0]$. 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		OECD	
	1960-90	1960-1990	1980-2000	1990-2000
$\Delta c - \Delta c^* = b(\Delta y - \Delta y^*)$		Estimate of b		
	0.16 (4.6)	0.84 (19.56)	0.89 (10.22)	0.85 (6.27)

NOTES: Numbers in parentheses are t -statistics based on the weighted least squares procedure described in the main text.

Table 2: Insurance against permanent and transitory shocks

Regression	United States		OECD	
	(1960-90)	1960-90	1980-2000	1990-2000
$\Delta c - \Delta c^* = b_P(\Delta y^P - \Delta y^{*P})$		b_P		
AR(1)	0.08 (3.18)	0.52 (15.39)	0.42 (7.42)	0.30 (2.71)
VAR(1)	0.04 (1.52)	0.55 (13.46)	0.43 (5.84)	0.19 (1.66)
$\Delta c - \Delta c^* = b_T(\Delta y^T - \Delta y^{*T})$		b_T		
AR(1)	0.01 (0.08)	-0.47 (-5.73)	-0.28 (-2.93)	0.07 (0.41)
VAR(1)	0.16 (3.23)	0.10 (1.13)	0.14 (1.33)	0.20 (1.54)

NOTES: Numbers in parentheses are t -statistics based on the weighted least squares procedure described in the main text.

Table 3: Business cycle structure, US states and OECD countries

	share of perm. shocks α	Cov. structure ϕ	BC-‘Adjusted’ coefficient b_{adj}
United States			
1960-90	0.76 (0.002)	2.15 (0.04)	0.07 (0.02)
OECD			
1960-90	0.77 (0.001)	1.91 (0.06)	0.45 (0.02)
1980-2000	0.66 (0.002)	2.89 (0.11)	0.33 (0.02)
1990-2000	0.55 (0.001)	4.23 (0.40)	0.19 (0.04)

Notes: Decomposition of the panel risk sharing coefficient in Table (1). Calculations of permanent components based on the VAR. See section 5.2 for details. Standard errors(in parentheses) reported were obtained as follows: we repeatedly estimated the panel, dropping the observations for one country or region at a time. The respective α , ϕ and b_{adj} were then re-calculated and stored. This procedure gives us 22 (or 50 for the U.S. without Washington D.C.) different realizations of each the three parameters. The standard errors reported are the standard deviations over these 22 (50) realizations.

Table 4: Decomposition of country-specific output variance

OECD	Output $var(\hat{y})^{1/2}$	Permanent $var(\hat{y}^P)^{1/2}$	Transitory $var(\hat{y}^T)^{1/2}$	Correlation ρ
1960-90	0.026	0.031	0.017	-0.52
1980-2000	0.020	0.027	0.020	-0.65
1990-2000	0.017	0.025	0.023	-0.80

Notes: Calculations of \hat{y}^P and \hat{y}^T based on the permanent components based on the VAR. Numbers reported are averages across the panel of 22 industrialised countries, calculated as described in section

5.2.

Table 5: Intl' Consumption Correlations and the quantity puzzle

	United States 1960-90	1960-1990	OECD 1980-2000	1990-2000
	Data			
$corr(\Delta c, \Delta c^*)$	0.54	0.37	0.34	0.31
$corr(\Delta y, \Delta y^*)$	0.78	0.51	0.52	0.60
	Variability of global component and implied international consumption correlations			
$var(\Delta y^{*P})^{1/2}$		0.023	0.019	0.008
ρ_{imp}		—	0.33	0.15

NOTES: Reported correlations are the cross-sectional (across countries or US states) averages of the correlations of consumption/output growth rates with the respective world aggregate. $\rho_{imp} = \rho \left((1 + \delta) (1 + \delta \rho^2)^{-1} \right)^{1/2}$, where ρ is the correlation between consumption growth rates from the preceding subperiod. δ is obtained is the percentage increase in the variance of the global component, $var(\Delta y^{*P})$.