Working Paper no. 302

International and intranational consumption risk sharing: the evidence for the United Kingdom and OECD

Vincent Labhard and Michael Sawicki

July 2006

Bank of England

International and intranational consumption risk sharing: the evidence for the United Kingdom and OECD

Vincent Labhard[*]

and

Michael Sawicki[**]

Working Paper no. 302

* Conjunctural Assessment and Projections Division, Bank of England.
Email: vincent.labhard@bankofengland.co.uk

** External Monetary Policy Committee Unit, Bank of England.
Email: michael.sawicki@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. The authors wish to thank the referees and seminar participants at the Bank of England for their comments and suggestions. This paper was finalised on 21 April 2006.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/publications/workingpapers/index.htm.

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone +44 (0)20 7601 4030, fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.

ISSN 1749-9135 (on-line)

Contents

Abstract

'Consumption risk sharing' refers to the ability of agents to insure or protect their consumption against shocks to their income, for example, by borrowing and lending or holding claims on foreign equity. So measuring the extent of risk sharing informs us about how consumption is likely to respond to country or region-specific shocks to income. This paper presents the evidence for consumption risk sharing by UK consumers, both across regions of the United Kingdom (intranationally) and between the United Kingdom and other countries (internationally). The main motivation for collecting this evidence is to establish to what extent UK consumers insure against income risks, and whether they do so to the same extent intranationally and internationally. Such evidence can tell us whether risk sharing functions effectively as an absorber of country or region-specific shocks in the United Kingdom. We find that there is more risk sharing across the UK regions than between the United Kingdom and other OECD countries. To test the robustness of our conclusions, we document the evidence for risk sharing using recent econometric techniques, which allow explicitly for country or region-specific factors impacting on consumption and output, including measurement errors in the data. We find that our results remain robust when we account for the possible impact of measurement error and preference shocks, and are consistent with results reported in the existing literature. Additionally, our paper also makes a separate contribution to the literature by illustrating the role of the choice of deflators in estimating the true extent of risk sharing for the United Kingdom and OECD. In terms of the channels through which risk sharing occurs, we find that the main mechanism of regional risk sharing operates via cross-regional asset holdings. Internationally, the main source of income smoothing comes from international borrowing and lending.

Key words: Risk sharing, consumption smoothing, asymmetric shocks, UK regions, OECD countries.

JEL classification: E21, E32, E44.

Summary

'Consumption risk sharing' refers to the ability of households to protect their consumption against shocks to their income. This could take the form of holding equity claims on output that is unrelated to their income, by receiving transfers from other agents or by borrowing and lending. Successful risk sharing should imply a smooth pattern of consumption, not greatly affected by fluctuations in households' incomes. However, empirical work has shown that this does not appear to be the case. For example, households hold a much smaller proportion of foreign equity than would be expected if they decided on their asset holdings on the basis of risk and return, a phenomenon known in the literature as the home bias in equities. A related puzzle is the home bias in consumption. Under full risk sharing, domestic consumption growth should be more highly correlated with foreign consumption growth than with domestic output growth, but the empirical evidence again suggests the contrary: cross-country consumption growth correlations are relatively low, and often lower than correlations with domestic output growth. This would suggest that 'idiosyncratic' – ie country or region-specific – output shocks are not effectively smoothed away and hence materially affect consumption.

Measuring the overall extent of risk sharing in consumption is interesting for monetary policy makers because of its impact on the transmission of shocks. Gauging the extent of risk sharing – as well as the channels through which it occurs – can help us understand business cycle developments and imbalances in an economy by informing us about how consumption is likely to respond to country or region-specific output shocks, and can shed light on how much policymakers might need to react to such shocks. As another potential adjustment mechanism against country-specific shocks, a large enough degree of international risk sharing may mitigate the effects of such shocks if other adjustment mechanisms, for example exchange rate flexibility or labour mobility, are absent or limited. And finally, understanding the channels by which risk sharing is achieved would also provide an insight into the effects of capital and credit market integration, both domestically and internationally.

In this paper, we present the empirical evidence for consumption risk sharing by UK consumers, both between the United Kingdom and other countries (internationally) and across regions of the United Kingdom (domestically). Such evidence can tell us whether risk sharing acts as an absorber of country or region-specific shocks in the United Kingdom. The key questions we seek

to answer are whether there is more or less risk sharing domestically than internationally; through which channels it occurs; whether risk sharing has increased or decreased over time; and finally, whether these estimation results are robust. We address these questions by employing two specific methodologies. First, we use an established panel regression analysis – using a data set spanning a set of countries or regions across time – to illustrate both the extent of risk sharing and the channels through which it is carried out, updating existing UK results with more recent data. The incremental information on the channels by which risk sharing occurs means that we prefer this methodology to simple correlations of consumption and output data. However, as with correlations, the panel analysis may be distorted by factors affecting output or consumption but not related to risk sharing, such as changes in household preferences and measurement error in the data. This prompts us to utilise a second, more recent, methodology, which takes these influences explicitly into account using a factor model – a technique which aims to separate out the key drivers or factors (here at the regional and national level, for example) from a potentially large set of data – on consumption and output. Applying this factor model to the United Kingdom and OECD data is the main contribution of this paper.

We find that there is more risk sharing across the UK regions than between the United Kingdom and OECD countries. This baseline result is robust to accounting for the possible impact of measurement error and changes in household preferences, and consistent with results reported in the previous work. We find that the main mechanism of regional risk sharing operates via cross-regional asset holdings. Internationally, the main source of income smoothing comes from international borrowing and lending. Consistent with previous work in the field, we find tentative evidence that risk sharing has declined over time, although the importance of capital market smoothing has gradually increased, consistent with recent increases in capital market integration. However, these trends may require caution in interpretation, because the methodologies we use may not fully detect changes in the nature of the risks to output occurring during the course of our sample.

Finally, our paper also makes a separate contribution to the literature by illustrating the role of the choice of price measures (deflators) in estimating the true extent of risk sharing for the United Kingdom and OECD. While estimates of the extent of risk sharing within the United Kingdom are relatively invariant to the choice of deflator, using our preferred choice of deflator for the OECD data set yields higher estimates of risk sharing than typically reported in previous work.

Therefore, it appears as if in the United Kingdom regional consumption fluctuations may be largely unaffected by regional output fluctuations, and that UK consumption – while affected by global output fluctuations – may also be more robust than suggested in previous work.

1 Introduction

'Consumption risk sharing' refers to the ability of agents to protect their consumption against shocks to their income, for example by holding assets that offer high returns when their income is low. Insuring against income shocks is often seen as involving diversifying portfolios with foreign assets, as their returns are likely to be less correlated with domestic income. However, empirical work has shown that agents in fact hold a much smaller proportion of foreign assets – particularly for equities – than would be optimal under standard assumptions about investors' portfolio preferences. This phenomenon, known as home bias in equities, was first illustrated by French and Poterba (1991) and is one of the major puzzles in international macroeconomics discussed in Obstfeld and Rogoff (2000).[(1)] A related puzzle is the home bias in consumption. In a simple model, under certain assumptions about the structure of financial markets and the shape of agents' utility functions,[(2)] full consumption risk sharing would imply that domestic consumption growth should be more highly correlated with consumption growth in the world as a whole than with domestic output growth. In essence, risk sharing can act as an absorber of country-specific shocks and can therefore decouple movements in domestic consumption from those in domestic output. However, the empirical evidence suggests the contrary: cross-country consumption growth correlations are relatively low and in fact often lower than correlations with domestic output growth.[(3)] This would suggest that 'idiosyncratic' (ie country or region-specific) output shocks are in fact not effectively smoothed away and hence materially affect consumption, a phenomenon also referred to as the 'quantity anomaly', and originally documented by Backus, Kehoe and Kydland (1992).

Agents can smooth their consumption in several ways. Internationally, *ex-ante* risk sharing in consumption would involve agents holding equity claims on foreign output. So home biases in consumption and equity holdings are consistent in the sense that if individuals hold too few claims on foreign assets, then they will not optimally share risk with foreigners. This implies that their marginal rates of substitution in consumption – and with further assumptions about utility

---

(1) For a detailed discussion, see Lewis (1999). Obstfeld and Rogoff (2000) suggest trade transaction costs as one of the possible explanations.

(2) Specifically, this holds if markets are complete in the Arrow-Debreu sense, asset prices are actuarily fair, ie reflect the true probability of events, and utility functions are such that agents have constant relative risk aversion. But the home bias result observed abstracts from any explicit bias in preferences towards home-produced goods.

(3) For example, in our international data set described in Section 4.1, the average pairwise consumption correlation between OECD countries is 0.27, compared with an average consumption-output correlation across OECD countries of 0.59. The corresponding figures for our regional data set (described in Section 3.1) are 0.63 and 0.59.

functions, consumption growth rates – will not be equalised internationally. But the link between the two biases need not necessarily hold: *ex-post* risk sharing in consumption can be duplicated through *ex-ante* consumption smoothing behaviour without international trade in equities. For example, in response to temporary shocks, individuals can borrow and lend internationally through another asset. Or output risks could be smoothed by non-market-based means altogether, for example by intergovernment transfers or payments to relatives abroad.

Measuring the overall extent of risk sharing in consumption is of interest to monetary policy makers because of its impact on the transmission of shocks. Gauging the extent of risk sharing – as well as the channels through which it occurs – can help us understand business cycle developments and imbalances in an economy by informing us about how consumption is likely to respond to country or region-specific output shocks, and can shed light on how much policymakers might need to react to such shocks. As another potential adjustment mechanism against country-specific shocks, sufficient international risk sharing may also compensate for the absence of other mechanisms, such as exchange rate flexibility or labour mobility. And finally, understanding the channels by which risk sharing is achieved would also provide an insight into the effects of capital and credit market integration, both intranationally and internationally.

In this paper, we present the empirical evidence for consumption risk sharing by UK consumers, both between the United Kingdom and other countries (internationally) and across regions of the United Kingdom (intranationally), based on national accounts data for the period 1970-2001 and regional accounts data for the period 1974-99. The main motivation for collecting this evidence is to establish the extent to which UK consumers share risks, whether they do so to the same extent intranationally and internationally, as well as the channels through which it occurs. Such evidence can tell us whether risk sharing functions effectively as an absorber of country or region-specific shocks in the United Kingdom.

Some empirical evidence on the extent of risk sharing for the United Kingdom exists, but it is rather mixed. At a regional level for the United Kingdom, Dedola, Usai and Vannini (1999) estimate that 75% of long-run risks and 100% of short-run risks have been shared during 1979-94. These results use a methodology due to Asdrubali, Sorensen and Yosha (1996) based on a panel analysis of the cross-sectional variance of GDP, with the duration of the shared risks identified by the method of detrending data. Using a modified approach, however, Melitz and Zumer (1999) report a percentage of around 30% for a similar period (1972-96). The evidence is equally mixed on risk sharing between the United Kingdom and other countries. Sorensen and Yosha (1998), for example, find that risk sharing with other European countries has fallen from 42% to 21% after 1980, but has increased slightly relative to the OECD, from 34% to 36%.[4] These results raise a number of questions. First, is there more or less risk sharing across the regions of the United Kingdom than between the United Kingdom and other countries? Second, has risk sharing increased or decreased over time? And finally, are these results robust to the methodology used? In particular, the two issues we address here are the choice of the appropriate deflator, and the potential existence of measurement error in the data. These have recently become the subject of some debate in the literature, and not accounting for them properly can lead to the mismeasurement of risk sharing. This paper's main contribution is to shed light on these issues by investigating the robustness of UK results.

We address these questions by employing two specific methodologies. The first, the panel analysis by Asdrubali *et al* (1996) and Sorensen and Yosha (1998), cited above, can help to illustrate both the extent of risk sharing and, importantly, the channels through which risk sharing is carried out. We use it to update existing UK results in the literature with more recent data. The incremental information on the proportions of risk sharing carried out via specific channels means that we prefer this methodology to simple correlations, one naive way of looking at the pattern of risk sharing in consumption and output data. However, similarly to correlations, the risk-sharing regressions we use are susceptible to the distortions caused by idiosyncratic factors affecting output or consumption, such as preference shocks and measurement error in the data.[5] This prompts us to utilise a second methodology, developed originally by Del Negro (2002), which takes these influences explicitly into account using a factor model of consumption and output. Applying this factor model to the UK and OECD data is the main innovation of this paper.[6] We describe both methodologies in detail in Section 2. In Section 3 we focus on risk sharing across the regions of the United Kingdom, in Section 4 we turn to risk sharing between the United Kingdom and the OECD, looking at the results from each methodology in turn. We then provide a synthesis and suggestions for further research in Section 5. Section 6 concludes.

---

(4) A summary of other estimates, alongside estimates from a dynamic approach to measuring risk sharing and how it occurs is presented by Asdrubali and Kim (2004).

(5) Throughout the paper we are interested how idiosyncratic shocks to output are smoothed away. But in doing so we have to see through other idiosyncratic factors potentially affecting both output and consumption, which do not actually reflect output risks.

(6) Del Negro (2002) applies the factor model to US regional data.

2 Methodology

2.1 Estimating the extent and channels of risk sharing

In the first instance, we aim to provide evidence on the extent of risk sharing and the mechanisms through which it is achieved. For this purpose, we use the methodology suggested by Asdrubali *et al* (1996) and Sorensen and Yosha (1998), based on a panel analysis of the cross-sectional variance of GDP. Empirically, Becker and Hoffmann (2003) find this approach is consistent with the idea that shocks originate in output fluctuations and get smoothed via the channels of risk sharing. The original formulation, as used by Asdrubali *et al* (1996), was expressed in real terms, but the decompositions hold for both nominal and real quantities; for reasons explained in Section 2.3, we use several alternative price deflators.

We begin by examining risk sharing with the regions of the United Kingdom, using the methodology of Asdrubali *et al* (1996). We start from an identity of regional product for region i (GRP_i), using the ratios of regional product to regional income (GRP_i/RI_i), regional income to disposable regional income (RI_i/RDI_i) and disposable regional income to regional consumption (RDI_i/C_i), as well as the level of regional consumption (C_i), as below

$$GRP_i = \frac{GRP_i}{RI_i}\frac{RI_i}{RDI_i}\frac{RDI_i}{C_i}C_i \quad \textbf{(1)}$$

At a conceptual level, regional income (RI) comprises regional output (ie, GRP) and net interest, dividend and rental payments across regions (ie factor income flows), while disposable regional income (RDI) includes net national taxes and transfers. So movements in regional output not reflected in regional income suggest risk sharing via income flows resulting from cross-regional asset holdings. For example, dividend payments from a company based in one region can help support consumption when output in another region falls. Second, movements in gross regional income not reflected in disposable regional income suggest smoothing via the tax-transfer system: so redistribution of income via the national tax system can help support consumption when income falls in a particular region. And finally, movements in disposable income not reflected in consumption suggest risk sharing via simple intertemporal consumption smoothing: individuals can smooth their consumption by varying levels of borrowing and lending.

Formally, to recover the relative importance of these channels we take log differences,(7) multiply by the log difference of GRP_i, and then take expected values to obtain a decomposition of the period-by-period cross-sectional variance of GRP_i.(8) Subsequently by scaling by the variance of the log difference of GRP_i ($var(\Delta grp_i)$), we obtain an adding-up constraint for the covariances such that $1 = \beta_K + \beta_F + \beta_C + \beta_U$, where

$$\beta_K = cov(\Delta grp_i - \Delta ri_i, \Delta grp_i)/var(\Delta grp_i) \quad \textbf{(2)}$$

$$\beta_F = cov(\Delta ri_i - \Delta rdi_i, \Delta grp_i)/var(\Delta grp_i) \quad \textbf{(3)}$$

$$\beta_C = cov(\Delta rdi_i - \Delta c_i, \Delta grp_i)/var(\Delta grp_i) \quad \textbf{(4)}$$

$$\beta_U = cov(\Delta c_i, \Delta grp_i)/var(\Delta grp_i) \quad \textbf{(5)}$$

The terms β_K, β_F and β_C can be interpreted respectively as estimates of the incremental percentage of smoothing through capital markets, fiscal transfers and intertemporal smoothing, averaged across the regions. The intuition is that smoothing through capital markets would be reflected in a positive covariance between $\Delta grp_i - \Delta ri_i$ and Δgrp_i, smoothing through fiscal mechanisms in a positive covariance between $\Delta ri_i - \Delta rdi_i$ and Δgrp_i, and intertemporal smoothing in a positive covariance of $\Delta rdi_i - \Delta c_i$ and Δgrp_i.(9) The remaining covariance between Δc_i and Δgrp_i, captured by the coefficient β_U, would therefore reflect unsmoothed consumption, ie consumption risk which is not shared across regions. Note that as the expressions **(2)** to **(5)** relate to regressions (respectively of $\Delta grp_i - \Delta ri_i$ on Δgrp_i, $\Delta ri_i - \Delta rdi_i$ on Δgrp_i, $\Delta rdi_i - \Delta c_i$ on Δgrp_i, and Δc_i on Δgrp_i) in which the β coefficients add up to unity, they are not independent. The solution used by Asdrubali *et al* (1996) which avoids estimation problems as a result of this interdependence is to jointly estimate a panel of equations as follows:(10)

$$\Delta ri_i = \alpha_K + \lambda_K \Delta grp_i + \varepsilon_{K,i} \quad \textbf{(6)}$$

$$\Delta rdi_i = \alpha_F + \lambda_F \Delta grp_i + \varepsilon_{F,i} \quad \textbf{(7)}$$

$$\Delta c_i = \alpha_C + \lambda_C \Delta grp_i + \varepsilon_{C,i} \quad \textbf{(8)}$$

(7) Logs of variables in this section are denoted by lower-case letters. We focus throughout on one-year differences, although we could also consider risk sharing at different horizons by looking at differences across several periods. However, Becker and Hoffmann (2003) note several caveats to doing so in the context of a panel analysis. Instead of taking differences, some papers have applied a Hodrick-Prescott or linear filter, in order to obtain an estimate of the degree of risk sharing at lower cycle frequencies (see for example Dedola *et al* (1999)). Asdrubali and Kim (2004) use an approach based on a panel VAR, which allows them to trace the role of the smoothing channels over time while capturing the feedback from output to the various smoothing channels.

(8) This makes use of the fact that $E(XY) = Cov(X,Y) + E(X)E(Y)$ and $E(X^2) = Var(X) + [E(X)]^2$.

(9) For example, a positive covariance between the two terms would indicate that Δgrp_i would move by more than Δri_i, ie that some smoothing occurs.

(10) See also Melitz and Zumer (1999).

In the above the α coefficients are time fixed effects designed to eliminate any national or global factors common to all the regions or countries, allowing the regressions to focus specifically on idiosyncratic movements in consumption and output.(11) In equations **(6)** to **(8)** the coefficient λ_K reflects smoothing through channels other than capital markets, the coefficient λ_F smoothing except that through capital market or fiscal mechanisms, and the coefficient λ_C reflects what is not smoothed through any of these mechanisms.(12) The smoothing coefficients in equations **(2)** to **(5)** then can be recovered by calculating

$$\beta_K = 1 - \lambda_K \tag{9}$$

$$\beta_F = \lambda_K - \lambda_F \tag{10}$$

$$\beta_C = \lambda_F - \lambda_C \tag{11}$$

$$\beta_U = \lambda_C \tag{12}$$

Sorensen and Yosha (1998) modify this methodology to measure the extent of risk sharing internationally, ie across countries. Because national accounts data are typically more disaggregated, the appropriate identity for domestic product (GDP_i) breaks down into the ratio of domestic to national product (GDP_i/GNP_i), national product to national income (GNP_i/NI_i), national income to national disposable income (NI_i/DNI_i), the ratio of national disposable income to consumption (DNI_i/C_i) and the level of consumption (C_i), as below

$$GDP_i = \frac{GDP_i}{GNP_i}\frac{GNP_i}{NI_i}\frac{NI_i}{DNI_i}\frac{DNI_i}{C_i}C_i \tag{13}$$

As at the regional level, factor income flows can provide a channel of risk sharing between countries, which would be reflected in the difference between domestic product (GDP) and national product (GNP). The second term in the decomposition above captures the difference between national product and national income (NI), which reflects capital depreciation.(13) Sorensen and Yosha (1998) argue that patterns of capital depreciation in national accounts data might in fact be expected to reduce risk sharing if depreciation is expected to be relatively invariant to changes in output.(14) At the national level, the difference between national income (NI) and disposable national income (DNI) reflects international transfers – akin to a tax-transfer system

(11) The fixed effects are conceptually separate from the variance decomposition. In the estimation we leave them unconstrained, their use being solely to add flexibility.

(12) In each case the intuition is that a coefficient closer to 1 implies more comovement with output, ie that shocks are not smoothed away.

(13) Conceptually at least, capital depreciation could also play a role at the regional level. However, in our estimation this channel is effectively subsumed in smoothing via regional capital markets captured by β_K, as regional accounts data are not sufficiently disaggregated to identify it separately.

(14) This would be sensitive to our assumptions of how the capital stock responds to output innovations, both in terms of direction and composition. Sorensen and Yosha (1998) assume that capital-output ratio is countercyclical.

within a country – which can also contribute to income smoothing. And as before, movements in disposable income not reflected in consumption suggest risk sharing via simple intertemporal smoothing: individuals can smooth their consumption by varying levels of borrowing and lending. Formally, the presence of an additional 'channel' for risk sharing via capital depreciation modifies our adding-up constraint, such that it becomes $1 = \beta_F + \beta_D + \beta_T + \beta_S + \beta_U$, where

$$\beta_F = cov(\Delta gdp_i - \Delta gnp_i, \Delta gdp_i)/var(\Delta gdp_i) \quad \textbf{(14)}$$

$$\beta_D = cov(\Delta gnp_i - \Delta ni_i, \Delta gdp_i)/var(\Delta gdp_i) \quad \textbf{(15)}$$

$$\beta_T = cov(\Delta ni_i - \Delta dni_i, \Delta gdp_i)/var(\Delta gdp_i) \quad \textbf{(16)}$$

$$\beta_S = cov(\Delta dni_i - \Delta c_i, \Delta gdp_i)/var(\Delta gdp_i) \quad \textbf{(17)}$$

$$\beta_U = cov(\Delta c_i, \Delta gdp_i)/var(\Delta gdp_i) \quad \textbf{(18)}$$

Our interpretation of the coefficients is broadly as before, with the individual coefficients reflecting the various smoothing mechanisms, now in an international context: factor income flows (β_F), capital depreciation (β_D), international transfers (β_T) and saving (β_S). As before, the final coefficient is effectively a residual (β_U), taken to reflect any lack of international risk sharing. The jointly estimated panel equations using international data become

$$\Delta gnp_i = \alpha_K + \lambda_F \Delta gdp_i + \varepsilon_{F,i} \quad \textbf{(19)}$$

$$\Delta ni_i = \alpha_K + \lambda_D \Delta gdp_i + \varepsilon_{D,i} \quad \textbf{(20)}$$

$$\Delta dni_i = \alpha_C + \lambda_T \Delta gdp_i + \varepsilon_{T,i} \quad \textbf{(21)}$$

$$\Delta c_i = \alpha_C + \lambda_S \Delta gdp_i + \varepsilon_{S,i} \quad \textbf{(22)}$$

so that the smoothing coefficients can be recovered as

$$\beta_F = 1 - \lambda_F \quad \textbf{(23)}$$

$$\beta_D = \lambda_F - \lambda_D \quad \textbf{(24)}$$

$$\beta_T = \lambda_D - \lambda_T \quad \textbf{(25)}$$

$$\beta_S = \lambda_T - \lambda_S \quad \textbf{(26)}$$

$$\beta_U = \lambda_S \quad \textbf{(27)}$$

As discussed in Melitz and Zumer (1999) however, the interpretation of the coefficients in Asdrubali *et al* (1996) and Sorensen and Yosha (1998) may not be entirely straightforward, meaning it may be necessary to modify the approach to obtain more reliable estimates. The first point is that the methodology may not be capable of capturing correctly the relative importance of the market-based risk sharing channels (β_K and β_C). In the case of regional risk sharing, for example, if the corporate sector increases savings, hence pushing up the value for β_K, households may simply respond by lowering their savings, thus forcing a commensurate reduction in β_C. As a result, we may be less confident about the individual importance of either channel. Second, the coefficients β_C and β_K may simply respond to a change in the nature and duration of income shocks, which in turn may be related to industrial composition. For example, long-term risks – which might reflect permanent shocks to specific industries – are unlikely to be shared by intertemporal smoothing, because individuals are unlikely to be able to dissave indefinitely. This is an important caveat, consistent with research showing that risk sharing tends to be higher for short-term risks relative to long-term risks.[(15)]

The second point raised by Melitz and Zumer (1999) is that the methodology of Asdrubali *et al* (1996) and Sorensen and Yosha (1998) may not capture the full extent of risk sharing, as the coefficient β_U derived from the regressions above may in fact not be the best measure of unsmoothed consumption. The reason for this is the existence of preference shocks to consumption – for example, reflecting changes in regional or national tastes – which may lead to a lower covariance of output and consumption but have little to do with risk sharing. As a remedy, they suggest imposing the average degree of risk sharing from the start, ie fix β_U, and to estimate the other coefficients subject to this constraint.[(16)] In their paper, they derive the estimate of β_U from the relative variance of the (levels of) consumption and output, averaged across time. They do not provide a value of β_U for the United Kingdom, but for the United States impose a value of 39%, compared with an estimate of 25% quoted by Asdrubali *et al* (1996) for the same data sample. However, this adjustment is sensitive to the statistical properties of the data on output and consumption: for example, looking at the relative variance of the data in log-levels may give a slightly different estimate of β_U. And furthermore, it is in any case a second-best solution: as suggested in the introduction, we can gauge the robustness of the estimate of β_U directly using a factor model (described in detail in Section 2.2), which can account more explicitly for preference shocks to consumption.

For our estimates of the channels of risk sharing, which we describe in Sections 3 and 4, we therefore retain the original formulation of the methodology of Asdrubali *et al* (1996) and Sorensen and Yosha (1998). However, to gauge the effects of idiosyncratic factors potentially

(15) See for example Artis and Hoffmann (2003), Becker and Hoffmann (2003) and Dedola *et al* (1999).
(16) This tends to affect the coefficient on smoothing via the credit markets, β_C.

affecting output or consumption, we also obtain measures of the extent of risk sharing using the factor model of Del Negro (2002), described in the next section.

2.2 Accounting for preference shocks and measurement error

The existence of preference shocks and measurement error in the data is one of the key issues in the empirical literature on consumption risk sharing. The intuition for their importance is that the presence of either can affect the measured relative variance of consumption and output. So in empirical work which aims to find evidence of a lack of comovement between consumption and output innovations – the result implied by full risk sharing – this can make a substantial difference to our interpretation of the results. In this section we explain how by using a factor model we can account for idiosyncratic factors such as preference shocks and measurement error such as to get another read of the true pattern of the comovement of consumption and output.[17]

The logic behind the importance of measurement error is best illustrated formally (see Crucini and Hess (2000)). If we assume that consumption is subject to preference shocks (μ) and consumption data are measured with error (η_c), while there is also a measurement error in output data (η_{gdp}), then the measured growth rates are related to the true growth rates (marked by asterisks in the expressions below) via $\Delta \log(gdp) = \Delta \log(gdp)^* + \eta_{gdp}$ and $\Delta \log(c) = \Delta \log(c)^* + \eta_c + \mu$.[18] This implies that the true variance of (underlying) consumption and output growth are overstated, as

$$var(\Delta \log(gdp_{i,t})^*) = \sigma^2_{gdp} - \sigma^2_{\eta_{gdp}} < \sigma^2_{gdp} = var(\Delta \log(gdp_{i,t})) \quad \textbf{(28)}$$

and

$$var(\Delta \log(c_{i,t})^*) = \sigma^2_c - (\sigma^2_{\eta_c} + \sigma^2_\mu) < \sigma^2_c = var(\Delta \log(c_{i,t})) \quad \textbf{(29)}$$

As a result, measures of risk sharing which are derived from correlations of observable consumption and output (or regression-based methodologies such as the panel approach outlined above), may be distorted – potentially in either direction – and the true extent of risk sharing can only be inferred if consumption and output are measured 'cleanly enough'.[19] For example,

(17) The original argument for benchmarking measurement error in both consumption and output is from Hess and Shin (1998).
(18) For simplicity, in what follows it is assumed that preference shocks and measurement error are independently distributed, so that there are no covariance terms to consider. It would be possible in principle to also consider situations where measurement error in consumption and output are correlated, perhaps as a result of data collection methodology, for example in early releases of national accounts data. But this is less likely to be the case with the vintages of data we use, which typically would have already been revised several times.
(19) See Crucini and Hess (2000).

consider the simplest test of the risk-sharing hypothesis based on the relative size of cross-country consumption and output correlations. If data are measured without error, evidence consistent with risk sharing would be provided by correlations that are larger for (true) consumption than output growth, ie

$$cov(\Delta \log(c_{i,t})^*, \Delta \log(c_{j,t})^*) > cov(\Delta \log(gdp_{i,t})^*, \Delta \log(gdp_{j,t})^*) \quad \textbf{(30)}$$

for $j \neq i$. However, more generally, this would require that the relationship holds when (observed) consumption and output growth correlations are normalised by the ratio of measured and observed variance

$$\frac{cov(\Delta \log(c_{i,t}), \Delta \log(c_{j,t}))}{\frac{\sigma_c^2 - (\sigma_{\eta_c}^2 + \sigma_\mu^2)}{\sigma_c^2}} > \frac{cov(\Delta \log(gdp_{i,t}), \Delta \log(gdp_{j,i}))}{\frac{\sigma_{gdp}^2 - \sigma_{\eta_{gdp}}^2}{\sigma_{gdp}^2}} \quad \textbf{(31)}$$

which can be rewritten as

$$\frac{cov(\Delta \log(c_{i,t}), \Delta \log(c_{j,t}))}{cov(\Delta \log(gdp_{i,t}), \Delta \log(gdp_{j,t}))} > \frac{1 - (\sigma_{\eta_c}^2 + \sigma_\mu^2)/\sigma_c^2}{1 - \sigma_{\eta_{gdp}}^2/\sigma_{gdp}^2} \quad \textbf{(32)}$$

In other words, we would need to adjust our observed pattern of consumption and output correlations to recover the appropriate underlying measure of correlations to test the risk-sharing hypothesis.

The second methodology we employ in this paper provides us with a means to take preference shocks and measurement error into account explicitly. Originally suggested by Del Negro (2002), it consists of a factor model of output and consumption. This model is based on the idea there is a small set of unobserved 'factors' which explain the dynamics of output and consumption data for an individual region or country by summarising the information in the output and consumption data for the entire cross-section. In Del Negro (2002), the assumption is that the key factors are linked to three geographical layers (corresponding to the nation, the region, and the state in US data), while measurement error in the data and preference shocks are possible additional factors. In such a framework, the state factor only affects individual states, a regional factor affects all states in the region (potentially with varying impact), while the national factor affects all states (again, potentially with differential impact).

Using GDP_i and C_i to denote real per capita output and consumption in region or country i, and $\overline{GDP}$ and $\overline{C}$ to denote real per capita output and consumption for the regional / country aggregate, the factor model is written for $gdp_i = \log(GDP_i) - \log(\overline{GDP})$ and $c_i = \log(C_i) - \log(\overline{C})$, detrended and demeaned.[20] The terms gdp_i and c_i are referred to as asymmetric shocks to output

(20) In the regional case, GDP should be replaced by the equivalent measure of regional product, GRP.

and consumption, respectively, as they capture the lack of comovement between $\log(GDP_i)$ and $\log(\overline{GDP})$ as well as $\log(C_i)$ and $\log(\overline{C})$. In the factor model, therefore, the focus is on the movements of output and consumption relative to the average, while in regression-based model the focus is on output and consumption – this is one reason why the results may not be entirely comparable (see also the discussion in Section 5).

In its simplest form as used by Del Negro (2002), the factor model reads

$$gdp_i = \beta^1_{gdp_i} f^1 + \sum_{r=1}^{\bar{r}} \beta^r_{gdp_i} f^r + \sum_{j=1}^{n} \beta^j_{gdp_i} f^j \tag{33}$$

$$c_i = \beta^1_{c_i} f^1 + \sum_{r=1}^{\bar{r}} \beta^r_{c_i} f^r + \sum_{j=1}^{n} \beta^j_{c_i} f^j + \varepsilon_{c_i} \tag{34}$$

where f^k are the factors at layer k, and ε_{c_i} is a disturbance term which can be either a preference shock or measurement error in consumption data.(21) In this paper, we consider three layers of factors, so $k = 3$. When looking at risk sharing across the United Kingdom we use regions (denoted j), groups of regions (denoted r) and the UK aggregate; for risk sharing between the United Kingdom and OECD, we use countries (denoted j), groups of countries (denoted r) and the OECD aggregate.(22) The β's in equation **(33)** and **(34)** are the factor loadings, and are estimated by maximum likelihood methods. We are interested in these because they can be used to calculate the standard deviations of the asymmetric shocks to output and consumption as $\sigma(gdp_i) = \sqrt{(\beta^1_{gdp_i})^2 + (\beta^r_{gdp_i})^2 + (\beta^j_{gdp_i})^2}$ and $\sigma(c_i) = \sqrt{(\beta^1_{c_i})^2 + (\beta^r_{c_i})^2 + (\beta^j_{c_i})^2}$. Intuitively, these standard deviations can tell us whether output or consumption are more correlated (ie less asymmetric) across regions or countries, and hence whether the quantity anomaly holds or not. We would expect risk sharing to be reflected in higher standard deviations of asymmetric shocks to output than to consumption, suggesting that output is less correlated (more asymmetric) than consumption. Whether this is due to risk sharing can then be gauged by looking at the standard deviations corrected for preference shock or measurement error to consumption, ie excluding the term ε_{c_i}.(23) If measurement error or preference shocks in consumption are important, then this second standard deviation should be much smaller, and hence a quantity anomaly be less obvious. This would suggest that not accounting for measurement error understates the extent of risk

(21) Of course, measurement error could also enter our output equation. We discuss this possibility in subsequent paragraphs.
(22) We describe how we group the UK regions and OECD countries when presenting the data in Section 3.1 (for the United Kingdom) and Section 4.1 (for the OECD).
(23) Stripping off measurement error and preference shocks in this way implies that the asymmetric shocks in consumption are by construction related to the asymmetric shocks in output. So the relative standard deviations then tell us to what extent the output shocks are smoothed.

sharing.

Assigning the disturbance term ε_{c_i} to the consumption equation is motivated by a prior that consumption is more likely to be measured with error than output.[24] But we can cross-check this question explicitly by examining what would happen if we moved the disturbance term to the output equation, writing the factor model as

$$gdp_i = \beta^1_{gdp_i} f^1 + \sum_{r=1}^{\bar{r}} \beta^r_{gdp_i} f^r + \sum_{j=1}^{n} \beta^j_{gdp_i} f^j + \varepsilon_{gdp_i} \quad \textbf{(35)}$$

$$c_i = \beta^1_{c_i} f^1 + \sum_{r=1}^{\bar{r}} \beta^r_{c_i} f^r + \sum_{j=1}^{n} \beta^j_{c_i} f^j \quad \textbf{(36)}$$

If there is a measurement error in output, then the standard deviation of asymmetric shocks to output should become smaller when the term ε_{gdp_i} is excluded, suggesting that output would be more correlated across regions or countries than consumption. However, in this case the implication would be that any quantity anomaly would be accentuated, and the evidence turned against the risk-sharing result.

Given these two sets of results, it is also possible to ascertain the relative size of measurement error in output and consumption, both at the regional and international level, by comparing by how much $\sigma(gdp_i)$ and $\sigma(c_i)$ change, in itself an interesting issue we can comment on in this paper. But it is also quite possible that measurement error would in fact enter both consumption and output. It is possible to address this question by extending these models to include a disturbance term for both consumption and output. Indeed, one potential advantage this would have over the simple factor models presented above is that preference shocks and measurement error could be identified separately, as illustrated by equation **(32)**. Del Negro (2002) does so by considering alternative series for output and consumption, which in part rely on using alternative deflators.[25] However, using alternative deflators in this way in fact may not be extracting the measurement error, as argued by Sorensen and Yosha (2002), but instead distorting the relevant measure of output or consumption, and hence our interpretation of the extent of risk sharing. The next section explains this in more detail. Further, being able to distinguish between preference shocks and

(24) Of course, consumption is a large component of GDP measured in expenditure terms, so measurement error in consumption can affect measured GDP. But to the extent that GDP can also be measured in output and income terms, the use of 'coherence adjustments' means that output may be relatively less affected by movements in any individual component of expenditure.

(25) Using this approach, Del Negro (2002) shows that contrary to the widely held view that output is less likely to be measured with error, the evidence for the states in the United States suggests that at least at the regional level there is a measurement error also in output.

measurement error is not crucial to correctly interpreting the observed pattern of comovement in output and consumption in the presence of idiosyncratic factors. This is because both have similar effects on the comovement of output and consumption, such that making the distinction explicit will not affect the conclusions on the risk-sharing hypothesis, and so we do not pursue this option here.

Two aspects of the factor model though are worth highlighting. First, the model is formulated in terms of the movements of output and consumption relative to the regional / country average. This is in contrast to the panel approach followed by Asdrubali *et al* (1996) and Sorensen and Yosha (1998), which focuses on output and consumption for each region or country. To the extent that there is a difference in risk identified when looking at patterns relative to (i) all other regions or countries and (ii) the average of all regions or countries, the measured degree of risk sharing may differ in this approach. Related, this means that the factor model may be less useful as an independent metric of the overall degree of risk sharing, but most useful as a complement to (and robustness check on) the panel approach, by giving an indication of the direction of the distortion caused by not gauging the impact of preference shocks to consumption and measurement error in consumption and output.

2.3 Accounting for price differentials

As with preference shocks and measurement error in the data, the choice of deflator is another key aspect which can affect our interpretation of observed consumption and output growth correlations, and hence is a key issue in the empirical literature on consumption risk sharing. And this issue is potentially relevant for both methodologies employed in this paper. The reason is that relative price changes provide another risk-sharing mechanism. Intuitively, the risk faced by a region or country is associated with how the changes in the total value of its output, relative to the rest of the country or world, affect its consumption possibilities. But changes in consumption possibilities can arise from both changes in output (quantities) and local prices. For example, a negative output shock specific to a region could be partly offset if the relative price of consumption in that region were to also fall. To appropriately measure risk sharing we therefore need to consider the relevant measure of local prices, as our results – for example, the decompositions of the channels of risk sharing, based on identities **(1)** and **(13) –** may vary with the choice of deflator. This may be particularly relevant for international data where the relative prices of consumption between two countries can also change due to movements in exchange rates. Such deviations from purchasing power parity at an international level may in fact be one possible reason why the existing empirical literature has found more risk sharing across regions of the same country than across countries, as argued in Hoffmann (2005).

In terms of the key variables, it is clear that consumption must be deflated with a consumption deflator in order to obtain a series in real terms which represents consumption possibilities. However, the question is which deflator to use for output. Existing empirical work has used, alternatively, a GDP deflator or a deflator based on consumer prices.[26] Sorensen and Yosha (2002) discuss these alternatives, and illustrate why the appropriate deflator for output data is also one based on consumer prices. The key reason is that risk sharing is about consumption possibilities: so only using a deflator based on consumption prices would capture the true value of output in terms of consumption possibilities. Using the implicit GDP deflator instead would have the effect of eliminating changes in consumption possibilities which arise from changes in the relative price of output in terms of consumption. For example, a country or region that is a major producer of a particular commodity will have the value of its nominal output related to the variation in the price of that commodity. So deflating that nominal output by the GDP deflator – which would reflect the movements in the price of the commodity – would lead us to mismeasure the impact on consumption possibilities. Using a measure of a consumption deflator would in this case more accurately reflect changes in the purchasing power of output.

In order to capture the role of relative price changes in risk sharing, we therefore use three sets of deflators. First we use a set of implied deflators: country-specific for the OECD, and region-specific for the United Kingdom.[27] These ignore relative price effects, and hence provide a benchmark for the analysis below. Second, we use the region (for the United Kingdom) or country-specific (for the OECD) consumption deflators. By comparing the two we highlight the importance of relative price effects between output and consumption emphasised in Sorensen and Yosha (2002). We can argue that this set of deflators would be our preferred measure of local prices for gauging the extent of risk sharing. Finally, we use the aggregate (UK or OECD) consumption deflators to capture the effect of relative price changes across regions or countries. This will allow us to shed light on whether relative price effects play a different role across the UK

(26) See Del Negro (2002) for an overview of the deflators used in data sets for the United States.
(27) In the case of the United Kingdom, we deflate GRP by the national GDP deflator, and the other components by a regional consumption deflator.

regions and between the United Kingdom and OECD.

In the subsequent sections of the paper, we use these alternative deflators both for our estimates of the channels and extent of risk sharing. We apply them first to the panel analysis of the cross-sectional variance of GDP described in Section 2.1. We also use them in the factor model, which accounts for idiosyncratic factors such as preference shocks and measurement error in the data, as described in Section 2.2. This should help us highlight the true extent of risk sharing, both in our UK and OECD data sets.

3 Risk sharing across UK regions

3.1 Data

The data set for the analysis of risk sharing between the UK regions is based on data from the 2002 edition of *Regional Trends* compiled by the United Kingdom's Office for National Statistics (ONS). The series available are: individual consumption expenditure ('consumption', from Table 12.11), total household income ('regional income', from Table 12.7), disposable household income ('regional disposable income', Table 12.7) and GDP at basic prices ('output', excluding 'Extra-Regio', Table 12.1).[28] These series are in nominal terms, and are available at an annual frequency for the period 1974 to 1999. Our effective sample extends from 1975-99, so that relative to previous work on the UK regions by Dedola *et al* (1999) we therefore have an additional ten years of data. Our data set has no outliers, and all variables have a unit root (see the appendix).

Table A lists the regions covered in our data set. These regions correspond to the twelve Government Office Regions (GOR), as covered by the ONS.[29] The table also shows how we assign these regions to regional aggregates, based on geography: the South (comprising London, the South East, the South West and the East), the North (consisting of the West Midlands, the East Midlands, the North West, the North East as well as Yorkshire and the Humber) and the other regions (Northern Ireland, Scotland and Wales). We will employ these aggregates in the factor model, and use geography as a proxy for factors that can be a source of shocks above the regional

(28) The output data are on a residence basis, such that the income of commuters is allocated to where they live, rather than to their place of work. Output data on a workplace basis could potentially distort upwards the measured amount of risk sharing.
(29) For the period prior to 1989, when this classification came into effect, the series have been converted accordingly. We are grateful to Peter Hayes of the University of Sheffield for advice and help in compiling this data.

and below the UK level, such as industrial specialisation, transport costs and regional tastes.[30]

Table A: The UK data set

Region	Aggregate region	C	GRP	RI	RDI
East	South		1974-99 for all		
East Midlands	North		1974-99 for all		
London	South		1974-99 for all		
North East	North		1974-99 for all		
North West	North		1974-99 for all		
Northern Ireland	Other		1974-99 for all		
Scotland	Other		1974-99 for all		
South East	South		1974-99 for all		
South West	South		1974-99 for all		
Wales	Other		1974-99 for all		
West Midlands	North		1974-99 for all		
Yorkshire and the Humber	North		1974-99 for all		

The table shows the regions and samples for the UK data set. Abbreviations: C consumption expenditure, GRP gross regional product, RI regional income, RDI regional disposable income.

3.2 Evidence

3.2.1 Extent and channels of risk sharing

In this section, we consider the extent of risk sharing across the UK regions, and the channels through which it occurs using the model described in Section 2.1. Table B shows the generalised least squares estimates of the smoothing coefficients **(9)** to **(12)** and associated one standard error bands for the full sample, as well as the sample split in half, 1975-87 and 1988-99, using different deflators.[31] Going from left to right, the estimates of these coefficients are shown under each channel of risk sharing, combined with a total extent of smoothing and the proportion of output risks left unsmoothed. Several results are worth noting. First, it appears that there is a high degree of risk sharing across the regions of the United Kingdom. Using the results based on the regional consumption deflators (our preferred measure), the percentage of risks shared is 81%. In terms of the channels through which this degree of risk sharing is achieved, the results point to risk sharing via factor income flows resulting from cross-regional asset holdings as the most important

(30) Del Negro (2002) argues that geographic proximity can play a role in defining the productive structures of a given area, and hence proxy regional groupings based on other criteria.
(31) We have also obtained estimates for the sample analysed by Dedola *et al* (1999), 1979-94. Using our preferred price measure, the regional consumption deflators, we find the following estimates and confidence intervals: 53 [46, 60], 1 [-1, 3], 31 [19, 43], 14 [5, 23]. For comparison, their estimates and confidence intervals are: 48 [40, 56], -6 [-7, -4], 56 [51, 61], -1 [-1, -1]. The differences may in part reflect the choice of deflators: it is not clear what deflator Dedola *et al* (1999) use.

explanation, accounting for 46%. Risk sharing through capital markets might suggest that, from a regional perspective, agents hold a broad portfolio of corporate equity, perhaps facilitated by stock exchange listings. Risk sharing through saving and dissaving is also an important channel: intertemporal substitution accounts for 36%.(32) Fiscal transfers, by contrast, do not significantly contribute to the sharing of risk. These results are comparable to those recently reported in Kalemli-Ozcan, Sorensen and Yosha (2003).

Table B: Risk sharing across the UK regions (%)

	Capital markets		Fiscal transfers		Intertemporal		Smoothed		Unsmoothed	
Panel 1: Data deflated by implied deflators										
1975-99	49	[43, 55]	-4	[-6, -2]	38	[28, 48]	84	[77, 91]	16	[9, 23]
1975-87	49	[41, 57]	-8	[-10, -6]	44	[32, 56]	85	[77, 93]	15	[7, 23]
1988-99	55	[44, 66]	0	[-3, 3]	28	[12, 44]	83	[70, 96]	17	[4, 30]
Panel 2: Effect of relative price changes (consumption to output)										
1975-99	3	..	-3	..	2	..	3	..	-3	..
1975-87	2	..	-1	..	0	..	1	..	-1	..
1988-99	4	..	-6	..	10	..	8	..	-8	..
Panel 3: Data deflated by regional consumption deflators (preferred measure)										
1975-99	46	[40, 52]	-1	[-3, 1]	36	[26, 46]	81	[74, 88]	19	[12, 26]
1975-87	47	[39, 55]	-7	[-9, -5]	44	[31, 57]	84	[76, 92]	16	[8, 24]
1988-99	51	[41, 61]	6	[2, 10]	18	[3, 33]	75	[73, 87]	25	[13, 37]
Panel 4: Effect of relative price changes (regional to UK consumption)										
1975-99	1	..	-3	..	1	..	-2	..	2	..
1975-87	0	..	-1	..	-1	..	-2	..	2	..
1988-99	0	..	-6	..	7	..	1	..	-1	..
Panel 5: Data deflated by UK consumption deflator										
1975-99	47	[41, 53]	-4	[-6, -2]	37	[27, 47]	79	[72, 86]	21	[14, 28]
1975-87	47	[39, 55]	-8	[-10, -6]	43	[31, 55]	82	[74, 90]	18	[10, 26]
1988-99	51	[40, 62]	0	[-4, 4]	25	[9, 41]	76	[64, 88]	24	[12, 36]

Notes: The table shows estimates and associated one standard error bands for the percentage of consumption risks smoothed through the channels of intranational risk sharing.

The table also shows the role of relative price adjustments for regional risk sharing, by comparing estimates of risk sharing based on different deflators. Relative price changes between regions appear to have a limited impact on the amount of risk sharing found. This can be seen by comparing the results based on the UK consumption deflator with those based on regional consumption deflators (panels 3 and 5 of the table).(33) For the whole sample, the difference is no

(32) However, these coefficients may be imprecisely estimated, and it may therefore be more sensible to consider these estimates jointly than individually, as explained in Melitz and Zumer (1999).

(33) We are grateful for the use of regional consumption deflators provided to us by Peter Hayes of the University of Sheffield. The indices are constructed on the basis of the bi-annual Croner-Reward Regional Cost of Living Surveys, and adjusted to an annual frequency. See Hayes (2004).

more than 2 percentage points. This is a plausible result for regional data, given that such relative price changes are likely to be small across regions. Likewise, there is only a limited impact from relative price changes between output and consumption of no more than 3 percentage points, as can be seen from a comparison of the results based on regional consumption deflators and implied deflators (panels 1 and 3 of the table).[34] This is again a plausible result, and consistent with region-specific shocks not being very important in the United Kingdom.

Chart 1: Risk sharing across the UK regions

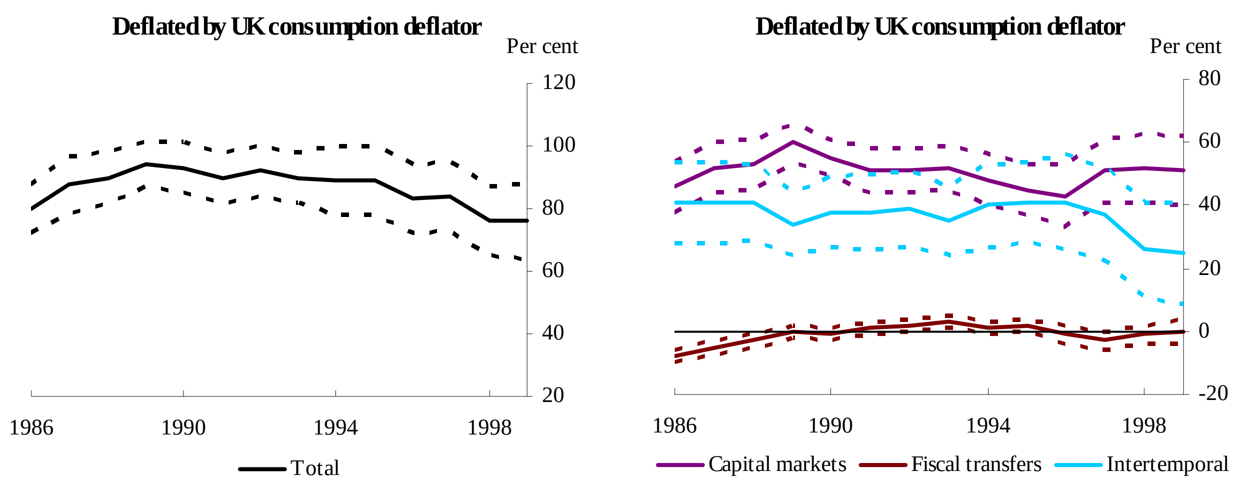

In terms of variation over time, using our preferred results based on regional consumption deflators, we find that the extent of risk sharing may have fallen between the first and second part of the sample (see again Table B). Mostly, this seems to be accounted for by a reduction in smoothing via intertemporal substitution, which has fallen from 44 to 18 percentage points. This would suggest that individuals have been less likely in recent years to smooth their consumption by borrowing or lending. Although this is hard to explain given increased financial liberalisation providing easier access to credit, for example, it could potentially reflect the changing nature of risks to be insured. For example, longer-duration risks might be harder to smooth by intertemporal substitution, and these may have become more prevalent over the sample. The 1980s and 1990s have arguably been a period of substantial changes to the structure of the UK economy. By contrast, capital market smoothing seems to have increased slightly from the early to the later period, although not by enough to compensate for the fall in borrowing and lending. In addition, there is no dis-smoothing due to fiscal transfers in the latter subsample. In order to assess these

(34) As regional output deflators are not available for the United Kingdom, due to the fact that *Regional Trends* data are only available in nominal terms, we use the UK GDP deflator to deflate all income series.

changes over time in more detail, we provide rolling estimates of the smoothing coefficients using twelve-year windows running from 1975-86 to 1988-99, the latter being the second subsample for which results are reported in Table B. These are plotted in Chart 1, with the period on the horizontal axis denoting the end of the rolling window. The chart reveals that in fact, the degree of risk sharing seems to have increased up to 1989, but after that has fallen back. It also illustrates that this fall in regional risk sharing has been driven by the decline in intertemporal smoothing from 1995 onwards, which has been offset only partly by the subsequent increase in capital market smoothing.

We discuss these results further in the context of the international picture and findings from the literature in Section 3, but broadly speaking, the results in this section point to a substantial degree of risk sharing across the United Kingdom. In addition to these findings – consistent with the results reported by Dedola *et al* (1999) and other research on regional risk sharing – our results also seem to be robust to the choice of deflator, implying that relative price movements are not that important in the UK data. Therefore, it appears as if regional consumption fluctuations may be largely unaffected by regional output fluctuations. Nevertheless, the results outlined in this section may be affected by preference shocks and possible measurement error in output or consumption, and we address this in the subsequent section.

3.2.2 Preference shocks and measurement error

As explained in the earlier sections of the paper, idiosyncratic factors entering consumption and output can affect our interpretation of the actual extent of risk sharing in measured output and consumption data. At the regional level, it is unlikely that preference shocks to consumption are very large, but it is quite possible that output and consumption are affected by measurement error. It is therefore important to establish how important these factors are, so as to be able to take a view on whether the results from the previous section are likely to be of value.

In this section, we therefore adopt the methodology developed by Del Negro (2002). We first consider the estimates and confidence intervals for the standard deviations of asymmetric shocks to output and consumption,[35] and then how these change once preference shocks and measurement error are taken into account explicitly (see Table C: compare the rows for the

(35) We recall that 'asymmetric shocks' is the term used for the lack of comovement between consumption and output at the regional and UK level or, when looking at international risk sharing, the country and OECD level.

uncorrected data, labelled either ‘Growth rates’ or ‘HP filter’, with that corrected for idiosyncratic factors immediately below). Panel 1 shows the results assuming that these factors affect consumption, as in equations **(33)** and **(34)**, panel 2 when they are assigned to output, as in equations **(35)** and **(36)**. As explained earlier, we can compare these two sets of results to assess the relative size of the idiosyncratic factors in output and consumption. By using two versions of the simple factor model described in Section 2.2 – as opposed to the more complicated factor model used in Del Negro (2002) to model measurement error in output and consumption jointly – we avoid the need for two separate data series for both output and consumption. This avoids the potential pitfalls of mismeasuring the extent of risk sharing due to using series deflated by alternative deflators, as argued by Sorensen and Yosha (2002).

Table C: Asymmetric shocks and risk sharing (UK regions)

	Consumption (a)		Output (b)		Ratio (a)/(b)	Smoothed (c)	Unsmoothed (d)
Panel 1: Standard deviations of asymmetric shocks: idiosyncratic factors assumed to affect consumption							
Growth rates							
not corrected	1.98	[1.33, 2.63]	1.83	[1.17, 2.49]	1.08	81%	19%
corrected	0.95	[0.42, 1.48]	1.83	[1.17, 2.49]	0.52	90%	10%
HP filter							
not corrected	1.47	[0.85, 2.09]	1.26	[0.73, 1.79]	1.17	79%	21%
corrected	0.80	[0.12, 1.48]	1.26	[0.73, 1.79]	0.63	88%	12%
Panel 2: Standard deviations of asymmetric shocks: idiosyncratic factors assumed to affect output							
Growth rates							
not corrected	2.15	[1.48, 2.72]	1.85	[1.19, 2.51]	1.16	79%	21%
corrected	2.15	[1.48, 2.72]	1.03	[0.44, 1.62]	2.09	66%	34%
HP filter							
not corrected	1.41	[0.96, 1.88]	1.28	[0.78, 1.78]	1.10	80%	20%
corrected	1.41	[0.96, 1.88]	0.75	[0.28, 1.22]	1.88	69%	31%

The table shows estimates and confidence intervals for the standard deviation of asymmetric shocks as defined in Section 2.2. It also shows the ratio of standard deviations for consumption and output, and a measure of risk sharing derived from this ratio via an exponential transformation, which produces a value between 0 and 100%. All estimates are for data deflated by the region-specific consumption deflator, sample 1974-99.

Table C shows that for consumption the standard deviation of asymmetric shocks is roughly halved when idiosyncratic factors are taken into account, falling from 1.98 to 0.95 (for data in growth rates) and 1.47 to 0.80 (for HP-filtered data). This points to preference shocks or measurement error in this variable as being important in UK data. When we account for idiosyncratic factors in output, the standard deviation of asymmetric shocks also falls, but by somewhat less than for consumption: from 1.85 to 1.03 (for data in growth rates), and 1.28 to 0.75 (for HP-filtered data). This suggests that while measurement error may be an issue for both variables, this error (which includes preference shocks) is larger for consumption than output in the regional data set.[36] As consumption becomes relatively more smooth than output when idiosyncratic factors are taken into account – reinforcing the risk-sharing result that there is a high degree of risk sharing across the UK regions – not taking idiosyncratic factors into account implies that the extent of risk sharing across the UK regions is likely to be underestimated.

To illustrate this, we derive a measure of risk sharing from the ratio of the standard deviations of asymmetric shocks, shown in the final column of Table C. This measure is computed via an exponential transformation, such that the measure becomes zero when the ratio goes to infinity (ie there is no risk sharing) and the measure becomes one when the ratio goes to zero (ie there is full risk sharing). For the proportion of risks smoothed, η_S, we use the formula $\eta_S = \mu\ (e^r)^{-2}$ where r is the ratio of asymmetric shocks to consumption and output (the ratio of columns (a) and (b) in Table C) and μ is a scaling factor;[37] the proportion of risks not smoothed is then simply $\eta_N = 100 - \mu\ (e^r)^{-2}$. Note that this is a highly stylised summary measure of risk sharing based on a different methodology and therefore not directly comparable to the proportion of risk sharing derived from the panel estimates in Sections 3.2 and 4.2.[38] Nevertheless, the measure illustrates that smoothing is higher when idiosyncratic factors in consumption are taken into account (90% compared with 81% for growth rates, and 88% compared with 79% for HP-filtered data), due to a fall in the standard deviation of asymmetric shocks to consumption. Conversely, when we correct for idiosyncratic factors in output, hence reducing the standard deviation of asymmetric shocks to output, the measure shows a smaller degree of risk sharing (66% compared with 79% for growth rates, and 69% compared with 80% for HP-filtered data).

(36) Not surprisingly, and consistent with the results presented in the previous section, we find very similar results when using the data sets based on the alternative deflators (regional consumption deflators or regional implied deflators).
(37) We set $\mu = 0.2$ to match the 81% estimate for proportion of risks smoothed we obtain using the panel approach with our preferred deflator (we can set the value of μ to calibrate the measure for any benchmark).
(38) This measure will be affected by the relative volatilities of the underlying consumption and output series, which makes it difficult to achieve a high degree of risk sharing. Consequently measures presented in this table also need not be directly comparable with similar measures quoted for the OECD data later in the paper.

Chart 2: Asymmetric shocks to output and consumption (UK regions)

2a: Growth rates

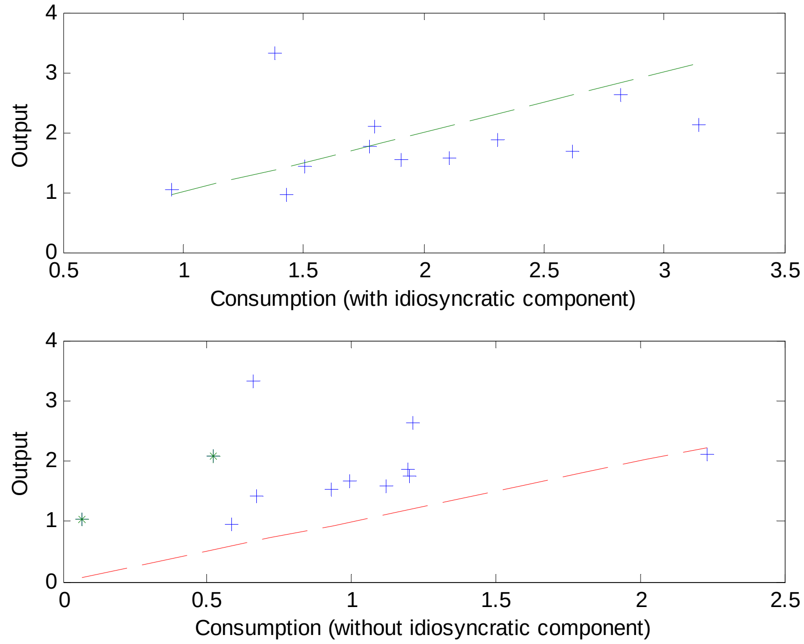

2b: HP filter

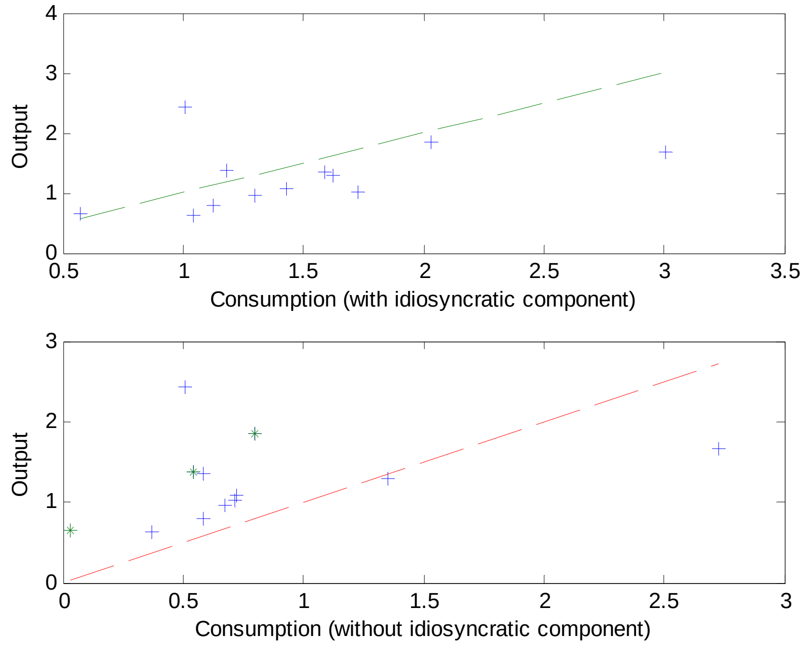

Note: The charts in the top panel show asymmetric shocks based on first-order differencing, the charts at the bottom show asymmetric shocks based on a Hodrick-Prescott filter with a smoothing parameter of 10.

Another way to look at the findings presented in Table C is to consider the scattergram of the standard deviation of the asymmetric shocks to consumption and output for all the regions. This is provided in Chart 2 for the case in which the idiosyncratic factors are associated with consumption, which appears to be the relevant case given the results reported in Table C. As a robustness check, we provide the scattergram for two detrending methods: first-order differencing (on the left) and a Hodrick-Prescott filter with a smoothing parameter of $\lambda = 10$ (on the right).[39] The effect of making the adjustment for idiosyncratic factors affecting consumption becomes apparent by comparing the top and bottom panels. When the adjustment is not made (top panel), most (albeit not all) observations are below the 45 degree line, thus indicating that asymmetric shocks are larger and cross-region correlations lower for consumption than output, effectively revealing a possible quantity anomaly. However, when the adjustment is made (bottom panel), all observations are above the 45 degree line, suggesting that if it were not for idiosyncratic factors, consumption correlations would be higher than output correlations and no evidence for a quantity anomaly would be detectable. The points marked with stars suggest observations for which the adjustment for idiosyncratic factors makes a difference at the 5% significance level.

4 Risk sharing between the United Kingdom and OECD

4.1 Data

For the analysis of risk sharing between the United Kingdom and other countries we use national accounts data available from the OECD. We extract the series on private final consumption expenditure ('consumption', code EXPPC), government final consumption expenditure ('government', code EXPGC), gross domestic product ('GDP', code EXPGDP), gross national income ('GNI', code RELGNI), net national income ('NNI', code RELNNI) and net national disposable income ('NNDI', code RELNNDI).[40] Because of data limitations, we exclude four countries from our data set: the Czech Republic, Hungary, Poland and the Slovak Republic. We also exclude Turkey, due to its hyperinflation, leaving us a data set of 25 countries covering 1970-2001. As for the UK regions, we find no outliers and a unit root in all the variables (see the appendix).

(39) As Dedola *et al* (1999), we have also experimented with alternative values for λ, without finding any major changes in results.

(40) The use of quarterly data would be desirable, but turns out not to be feasible, as national income measures at this frequency are only available for a small number of countries in the data set and, more generally, samples are much shorter.

Our data set is summarised in Table D. Table D also shows the four country aggregates we use for this data set, again based on geography: the Americas (which comprises Canada, the United States and Mexico), the euro area (with a total of twelve countries), other Europe (which includes Denmark, Iceland, Norway, Sweden, Switzerland and the UK) and finally the Asia Pacific region (comprising Australia, Japan, Korea and New Zealand). These aggregates as well as the OECD aggregate are computed using GDP weights also provided in the OECD National Accounts.

Table D: The OECD data set

Country	Aggregate	C	G	GDP	GNI	NNI	$NNDI$
Austria	Euro area	1970-2000 for all					
Australia	Asia Pacific	1970-2001 for all					
Belgium	Euro area	1970-2001 for all					
Canada	Americas	1970-2001 for all					
Denmark	Other Europe	1970-2001 for all					
Finland	Euro area	1970-2001 for all					
France	Euro area	1970-2001 for all					
Germany	Euro area	1970-2001 for all					
Greece	Euro area	1970-2000 for all					
Iceland	Other Europe	1970-2001 for all					
Ireland	Euro area	1970-2001	1970-2000	1970-2001	1970-2000	1970-2000	1970-2000
Italy	Euro area	1970-2001 for all					
Japan	Asia Pacific	1970-2001	1970-2001	1970-2001	1970-2001	1970-2000	1970-2000
Korea	Asia Pacific	1970-2001 for all					
Luxembourg	Euro area	1970-2001 for all					
Mexico	Americas	1970-2000 for all					
Netherlands	Euro area	1970-2001	1970-2001	1970-2001	1970-2000	1970-2000	1970-2000
New Zealand	Asia Pacific	1970-2000	1970-2000	1970-2000	1970-99	1970-99	1970-99
Norway	Other Europe	1970-2000 for all					
Portugal	Euro area	1995-2002 for all					
Spain	Euro area	1970-2001	1970-2001	1970-2001	1970-2001	1970-2000	1970-2000
Sweden	Other Europe	1970-2001 for all					
Switzerland	Other Europe	1970-2001	1970-2001	1970-2001	1970-2000	1970-99	1970-99
UK	Other Europe	1970-2001 for all					
USA	Americas	1970-2001 for all					

The table shows the countries and samples for the OECD data set. Abbreviations: C private final consumption expenditure, G final government consumption expenditure, GDP gross domestic product, GNI gross national income, NNI net national income, $NNDI$ net national disposable income. For Luxembourg, not available $NNDI$ data have been proxied by GDP.

4.2 Evidence

4.2.1 Extent and channels of risk sharing

In this section, we report results based on the panel analysis of cross-sectional variance of GDP for the full panel of 25 OECD countries. We also obtain estimates for two other panels. The first of these also excludes Mexico, another country with much higher inflation rates than the OECD average. The second panel extends the time dimension of the original panel to 2000, but excludes three countries from the cross-section for which the data stop in 1999: Luxembourg, Switzerland and New Zealand. Both these panels produce very similar estimates and lead us to the same conclusions, and so are not reported in full in the paper.(41)

Table E: Risk sharing between the United Kingdom and OECD (%)

	Factor income		Depreciation		Transfers		Savings		Smoothed		Unsmoothed	
Panel 1: Deflated by country-implied deflators												
1971-99	-2	[-3, -1]	-10	[-11, -9]	-4	[-5, -3]	45	[42, 48]	34	[31, 37]	66	[63, 69]
1971-87	3	[1, 5]	-8	[-9, -7]	-6	[-8, -4]	48	[44, 52]	38	[35, 41]	62	[59, 65]
1988-99	-2	[-4, 0]	-11	[-12, -10]	-2	[-3, -1]	45	[41, 50]	31	[27, 35]	69	[65, 73]
Panel 2: Effect of relative price changes (consumption to output)												
1971-99	-2	..	-4	..	-1	..	-12	..	-14	..	14	..
1971-87	3	..	-4	..	-2	..	-14	..	-16	..	16	..
1988-99	-1	..	-2	..	-1	..	-8	..	-11	..	11	..
Panel 3: Deflated by country consumption deflators (preferred measure)												
1971-99	0	[-1, 1]	-6	[-7, -5]	-3	[-4, -2]	57	[54, 60]	48	[46, 50]	52	[50, 54]
1971-87	0	[-1, 1]	-4	[-5, -3]	-4	[-6, -2]	62	[59, 65]	54	[51, 57]	46	[43, 49]
1988-99	-1	[-3, 1]	-9	[-10, -8]	-1	[-2, 0]	53	[49, 57]	42	[38, 46]	58	[54, 62]
Panel 4: Effect of relative price changes (country to OECD consumption)												
1971-99	0	..	5	..	2	..	-51	..	-43	..	43	..
1971-87	1	..	4	..	3	..	-54	..	-49	...	49	...
1988-99	1	..	6	..	1	..	-47	..	-39	..	39	..
Panel 5: Deflated by OECD consumption deflator												
1971-99	0	[0, 0]	-1	[-1, -1]	-1	[-1, -1]	6	[4, 8]	5	[4, 6]	95	[94, 96]
1971-87	1	[1, 1]	0	[0, 0]	-1	[-1, -1]	6	[4, 8]	5	[3, 7]	95	[93, 97]
1988-99	0	[-1, 1]	-3	[-4, -2]	0	[0, 0]	6	[3, 9]	3	[0, 6]	97	[94, 100]

The table shows estimates and associated one standard error bands for the percentage of consumption risks smoothed through the channels of international risk sharing.

Table E reports the results for three samples comparable to those used in Section 3: 1971-99, 1971-87 and 1988-99. Again, we show results for three different sets of deflators. As emphasised in Section 2.3, the choice of deflator may have an important impact on the extent of risk sharing

(41) The full results for these panels are available from the authors on request.

estimated from the data if there are large movements in relative prices, which we would expect to be more pronounced in international data. Table E suggests this may indeed be the case. Using country-specific consumption deflators – our preferred measure highlighting the importance of local purchasing power – we find the percentage of risks shared is 48%. By contrast, when using implied deflators, we find that the percentage of risk sharing amounts to 34%, so that about one third of risks are smoothed between the panel of OECD countries. This difference of 14 percentage points reflects the importance of relative price changes between output and consumption as a mechanism for international risk sharing. However, when we use a bespoke OECD-wide consumption deflator to deflate all series, the extent of risk sharing found in the data drops to 5%, suggesting that on this basis risks are not shared across the OECD. But it is likely that the aggregate OECD deflator constructed in fact may be a weak proxy for price dynamics in any individual country, obscuring the importance of relative price changes as a mechanism for risk sharing.

These differences obviously affect the answer to the question whether there is more risk sharing between the United Kingdom and OECD or across the UK regions, so that the use of deflator indeed becomes a crucial issue, as suggested by Hoffmann (2002). But although our analysis shows that our preferred estimate of the extent of international risk sharing may be higher than would be estimated with the deflator most commonly used in the literature (48% compared with 34%), we are not able to find the same extent of risk sharing as estimated for the UK regions (81%). From these results, it therefore appears as if differences in risk sharing across countries and regions remain a stylised fact of the data, consistent with the findings of most of the existing literature.

In terms of the main risk-sharing channels, our results indicate that credit markets are the main mechanism for smoothing consumption internationally, which we would expect to show up in international flows of borrowing and lending. In fact, smoothing through credit markets is more extensive than overall smoothing, and factor income flows, depreciation and transfers actually involve some dis-smoothing, as indicated by the negative numbers in Table E.[42] The table also provides a rough illustration of the changes that have occurred over time. Most importantly, it

(42) A negative contribution from depreciation is consistent with the aforementioned assumptions of Sorensen and Yosha (1998). We might also expect dis-smoothing as a result of factor income flows if equity home bias was prevalent. Finally, dis-smoothing from international transfers might be expected if the sizes of the flows do not mostly reflect purely economic motives.

Chart 3: Risk sharing between the United Kingdom and OECD

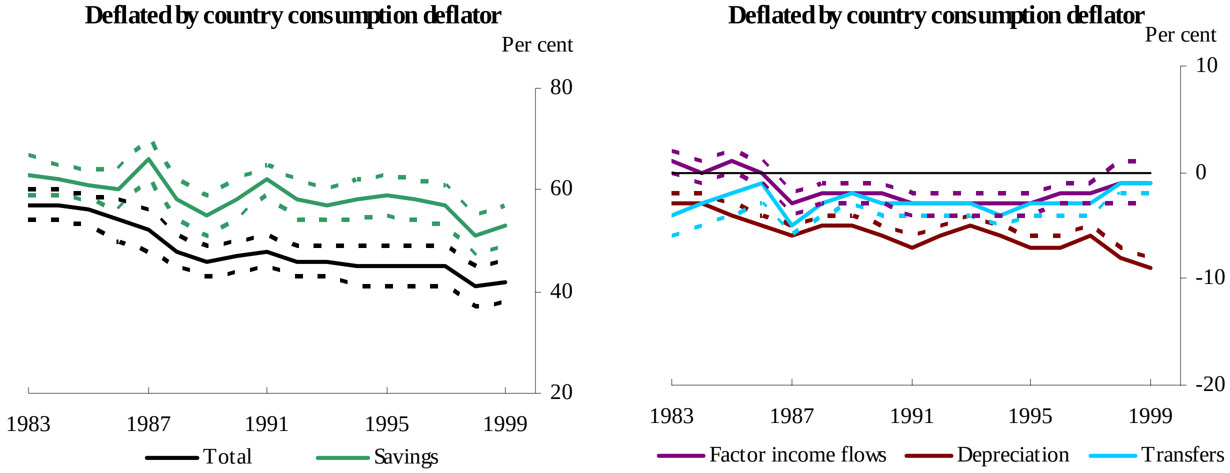

shows a decline in overall risk sharing from 54% to 42%, as well as more dis-smoothing through factor income flows and capital markets. However, it also shows that the relative importance of the different channels has remained broadly unchanged. As in Section 3.2, we also estimate rolling panels to provide more detail on the question of time variation (see Chart 3). The results using twelve-year windows from 1971-83 to 1988-99 illustrate further that the decline in savings and more dis-smoothing through depreciation have contributed to the overall fall in risk sharing. But it can also be seen that factor income flows, the analogue to the capital market smoothing in national data, have been increasing slightly, consistent with increased capital market integration among OECD countries, although in terms of magnitude this mechanism is not very important.

The surprising result of a fall in international consumption risk sharing despite a liberalisation of financial markets and an increase in international capital flows is one that has also been noted by other authors.[43] We draw on two possible explanations from the literature. The first is that despite a fall in consumption risk sharing estimated by the risk-sharing regressions as specified in Section 2.1, output risk may in fact have become much more diversified, but this is not being picked up by the regressions. Artis and Hoffmann (2003) propose a framework with the regressions specified in levels (as opposed to first differences), which suggests an increase in risk sharing in the period 1980-2000 relative to 1960-90.[44] They suggest that this risk sharing may reflect insurance against longer-duration risks and occurs as a result of larger international asset

(43) See Artis and Hoffmann (2003), for example.
(44) The specification of the regression in levels terms results from an assumed framework of 'Shiller securities'. See Artis and Hoffmann (2003).

holdings. An alternative second hypothesis is that fluctuations in output could become relatively more rather than less synchronised between financially integrated regions or countries, if the availability of finance induces specialisation in certain activities, and as a result of trade linkages. A recent paper by Imbs (2004) provides evidence on this second hypothesis. Even though consumption correlations might also be expected to increase, if output risks would appear to become more internationally correlated, and hence harder to diversify away, the relative increase in output correlations might appear as reduced risk sharing in the kinds of regressions investigated here.[45] While it is therefore difficult to make clear-cut claims about trends in risk sharing, the stylised fact that there appears to be more output risk diversification at the UK rather than international level is more likely to be robust.

4.2.2 Preference shocks and measurement errors

As for the UK regions, we also need to consider the possibility of preference shocks and measurement error in international data, and can usefully apply the factor model methodology of Del Negro (2002) to this effect. Apart from helping us to establish how much risk sharing takes place internationally when we account for these idiosyncratic factors, this analysis will also provide useful information on the importance of measurement error in OECD relative to UK data. The results are shown in Table F.

Again we consider how the standard deviations of asymmetric shocks to output and consumption change once preference shocks and measurement error are taken into account explicitly (comparing the rows for the data not corrected with those for the data corrected for idiosyncratic factors immediately below). They suggest that also for the OECD, taking into account preference shocks and measurement error in consumption matters is potentially important, though the impact of these factors is less marked than for the UK regions. Specifically, for the OECD data, the standard deviations of asymmetric shocks either to consumption or output fall by around a quarter when we correct for idiosyncratic factors (for data in growth rates, from 2.12 to 1.54 and 2.88 to 2.13 for consumption and output respectively), relative to a fall of around a half for the UK regions. Moreover, in the case of the OECD data, these factors impact to a very similar degree on consumption and output. This implies that correcting the data might not be as important as in the

(45) An alternative assumption might be that the response of consumption volatility to financial integration is in fact non-linear, and financial integration might be expected to lead to increases in consumption correlations beyond a certain threshold. See Kose, Prasad and Terrones (2003).

Table F: Asymmetric shocks and risk sharing (UK and OECD)

	Consumption (a)		Output (b)		Ratio (a)/(b)	Smoothed (c)	Unsmoothed (d)
Panel 1: Standard deviations of asymmetric shocks: idiosyncratic factors assumed to affect consumption							
Growth rates							
not corrected	2.12	[1.14, 3.10]	3.06	[1.58, 4.54]	0.69	87%	13%
corrected	1.54	[0.65, 2.43]	3.06	[1.58, 4.54]	0.50	90%	10%
HP filter							
not corrected	1.47	[0.92, 2.12]	2.08	[1.03, 3.13]	0.71	87%	13%
corrected	1.08	[0.46, 1.70]	2.08	[1.03, 3.13]	0.52	90%	10%
Panel 2: Standard deviations of asymmetric shocks: idiosyncratic factors assumed to affect output							
Growth rates							
not corrected	2.34	[1.16, 3.52]	2.88	[1.46, 4.30]	0.81	85%	15%
corrected	2.34	[1.16, 3.52]	2.13	[0.94, 3.32]	1.10	80%	20%
HP filter							
not corrected	1.63	[0.76, 2.50]	2.04	[1.09, 2.99]	0.80	85%	15%
corrected	1.63	[0.76, 2.50]	1.59	[0.72, 2.46]	1.03	81%	19%

The table shows estimates and confidence intervals for the standard deviation of asymmetric shocks as defined in Section 2.2. It also shows the ratio of standard deviations for consumption and output, and a measure of risk sharing derived from this ratio via an exponential transformation, which produces a value between 0 and 100%. All estimates are for data deflated by the country-specific consumption deflator, sample 1971-2000.

regional case: so we are probably not underestimating the extent of risk sharing by as much as in the regional case.[46] From this we also conclude that it appears that output is measured with a similar degree of precision in the UK and OECD data sets, while consumption is measured more accurately at a country rather than regional level, especially if one assumes that consumption must also be affected by preference shocks.

(46) This is borne out by the measure of risk sharing reported in column (c) of Table F, which is not affected by the correction as much in the case of the OECD data as in the case of the data for the United Kingdom.

Chart 4: Asymmetric shocks to output and consumption (UK and OECD)

4a: Growth rates

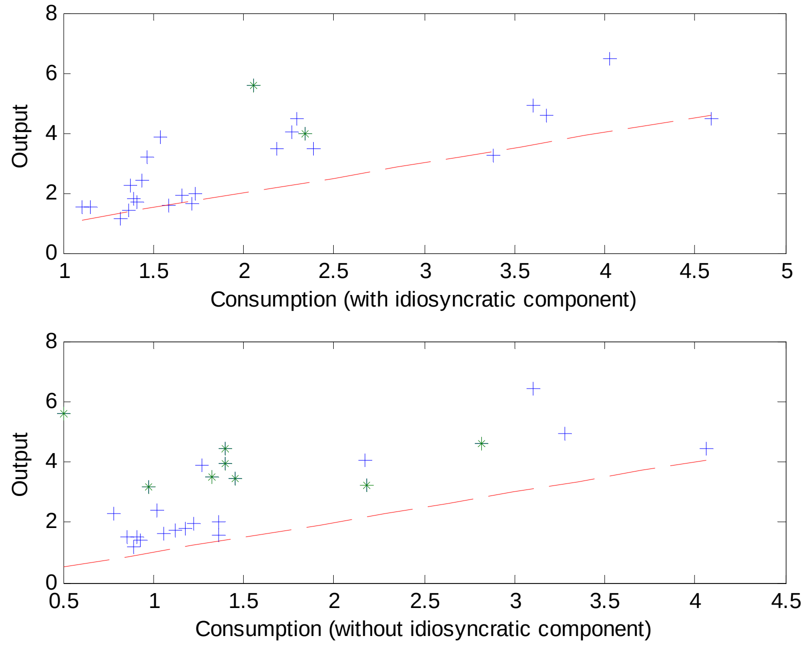

4b: HP filter

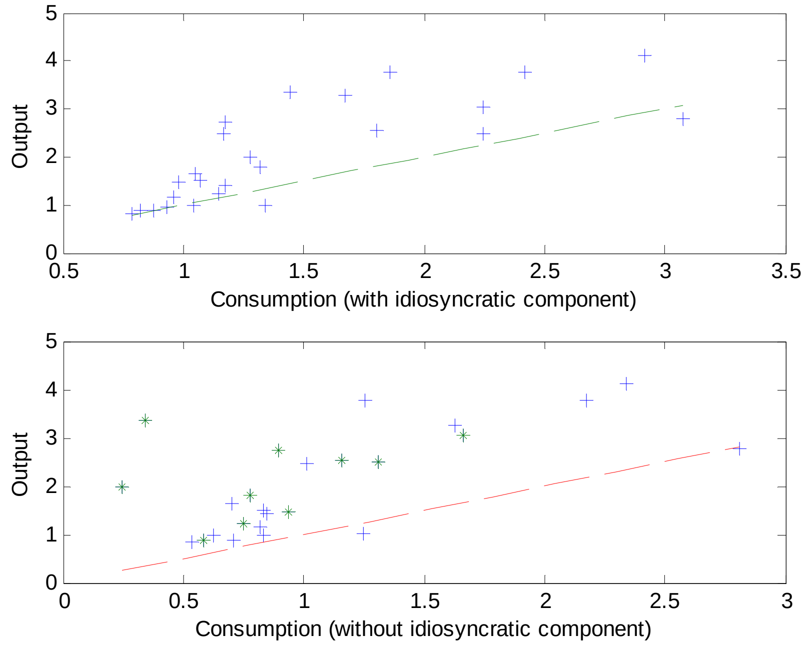

Note: The charts in the top panel show asymmetric shocks based on first-order differencing, the charts at the bottom show asymmetric shocks based on a Hodrick-Prescott filter with a smoothing parameter of 10.

In Chart 4, we provide the graphical illustration, based on the data deflated by the country consumption deflators and detrended alternatively with first differences (left panel) and the Hodrick-Prescott filter (right panel). We show the chart for the case in which the idiosyncratic factors are assumed to affect consumption, as in Section 3. As we have just shown, it is not obvious for the international data that the idiosyncratic factors are more important for consumption than output, but given they are smaller than in the regional data, the chart based on the alternative assumption of these factors affecting output would give a very similar picture. Indeed, the chart shows that correcting for the idiosyncratic factors makes little difference. In both the top and bottom panels, the observations tend to cluster on or above the diagonal, indicating as above that the evidence against the quantity anomaly would be little affected. We contrast this with the case of the UK regions, where accounting for idiosyncratic factors in consumption meant that the observations moved from below the diagonal to above it.

5 Synthesis and suggestions for further research

From the summary measure reported in column (c) of Table F, we can see that the estimated extent of risk sharing rises when we correct for idiosyncratic factors in consumption (for example, from 87% to 90% for data in growth rates), to a level similar to that reported for the UK regions in Table C. Although this would tentatively support the view that there is a similar degree of risk sharing internationally and intranationally, these two sets of results may only be spuriously similar. A natural question to ask is whether these results can be reconciled with those from the panel approach. In order to address this, it is helpful to separate two issues. First, what is the key factor that explains the differences between the factor model and panel approach? Is it measurement error, preference shocks, the choice of deflators (ie how we account for relative prices), or is there something more fundamental that tells them apart? And second, more specifically, why are the factor model results for the OECD countries and the UK regions so similar – in other words, why does the factor model results suggest a higher degree of international risk sharing than the panel regressions?(47)

In relation to the first point, note that it is not necessary to distinguish between measurement error and preference shocks to say in which direction our read of risk sharing might be distorted. So even if we believe that the factor model goes some way towards measuring the extent of risk

(47) Remember that for the UK regions, the measure of risk sharing shown for the factor model has been calibrated to match the degree of risk sharing obtained with the panel approach.

sharing more accurately – by providing an illustration of the effect of accounting correctly for measurement error – splitting up the sources of difference between our two sets of results does not necessarily help us further. Similarly, as regards the deflators, using the same (preferred) measure of relative prices means that we get as close as possible to a comparable measure of risk sharing relative to the regression-based approach. So do the differences between the two sets of results in the case of the OECD regions therefore reflect something 'fundamental' related to differences in methodology?

First, it is important to recognise that it is difficult to talk with precision about the proportion of risks being smoothed in the factor model framework, which instead permits us to make statements about consumption being more or less volatile than output. As a result, the use of the factor model lies most in emphasising the direction in which the presence of idiosyncratic shocks can affect the results, as opposed to pinning down the implied degree of risk sharing. Second, it is worth noting that the factor model results could potentially differ quite significantly from those from the risk-sharing regressions, which could help explain the observed differences between the two sets of results in the case of the OECD regions mentioned above.

One reason we've alluded to in Section 2.2 is that the regression-based approach focuses on the individual response of consumption in a country or region to an output shock in that country or region, while the factor model looks at how a country or region behaves relative to a set of countries of regions as a whole. Another reason is that the factor model, when applied at the cross-country level, may be focusing on a different set of shocks than when applied to the regional level. For example, in the OECD case, some region-specific shocks, which contribute to the country-level factor, could be offsetting, so that it reduces the standard deviation of asymmetric shocks to consumption by more than those to output. Conversely, in the regional case, the country-level factor might pick up global shocks, which would be relatively harder to smooth away, impacting similarly on asymmetric shocks to both consumption and output. Consistent with this, in our data set consumption for individual UK regions appears to be more volatile than consumption for individual OECD countries. In combination, this could also spuriously bring closer together the level of risk sharing from our summary measure in the two data sets, and thus explain why the factor model results for the case of the OECD reveal such a high degree of risk sharing internationally (the second point made above). Unfortunately, on the basis of our data set, we are not able to demonstrate to what extent this conjecture is correct.

Additionally, when talking about asymmetric shocks in the factor model, the volatility of an individual region is likely to have a bigger impact on the UK average than the volatility of an individual country on the OECD average. Nevertheless, the relative impact of this on either asymmetric shocks to consumption or output – and hence on the perceived smoothness of consumption against output – will depend on the consumption and output behaviour in individual large regions or countries.

Another observation we can make is that the asymmetric shocks to output and consumption in the factor model are more akin to consumption or output innovations when we detrend the data using an HP or linear filter, as opposed to first-order differencing. But, as Dedola *et al* (1999) find, there is not much difference for the overall degree of risk sharing between the three ways of detrending. However, insofar as the proportion of unsmoothed risks increases somewhat in regressions on HP or linear-filtered data, this might reflect the impact of longer-duration shocks being harder to smooth away, if we believe the structural interpretation in Dedola *et al* (1999). So, one interpretation would be that the factor model focuses on different shocks, which may bias the implied amount of risk sharing.

In light of the above, we can legitimately ask which approach we would feel more comfortable with recommending to applied researchers producing new estimates of the degree of risk sharing, and how they fit with the results reported in the existing literature. A related question is whether our results help to select between the two camps in the risk-sharing literature as far as the degree of risk sharing across regions is concerned – the works by Kalemli-Ozcan *et al* (2003) and others, which suggest that regional risk sharing is larger than risk sharing across national borders, and the research of others including Del Negro (2002) which suggests there may not be much more regional than cross-border risk sharing. Instead of either deciding between the two alternatives, or developing an approach nesting the two analyses used in this paper – one avenue we leave for future research – we emphasise that the two approaches naturally complement each other, and that each of the approaches has advantages and shortcomings. As discussed in Section 2.2, the factor model approach has the advantage of potentially providing an indication of the direction of the effect of idiosyncratic factors, which can potentially impact significantly on estimates of the proportion of risk sharing. For providing a self-standing metric of risk sharing, the factor model on its own is not as useful. It also is not suited to looking directly at the channels of risk sharing. Therefore, it is most usefully combined with a method which indicates the degree of risk sharing and the channels through which it occurs, which serves as a baseline for such an analysis. Overall, the key conclusion we can draw from our work is that given these baseline results from the panel analysis described in Sections 3.2 and 4.2 – which suggest a higher degree of risk sharing between the UK regions than countries of the OECD – these are not overturned by the evidence we obtain from the factor model.

6 Conclusion

In this paper, we have examined the extent of risk sharing across the UK regions and between the United Kingdom and OECD. We find that there is more risk sharing across the UK regions than between the United Kingdom and OECD. We find that this baseline result is robust to accounting for the possible impact of measurement error and preference shocks, and consistent with the previous results in the literature where such results exist. We find that the main mechanism of regional risk sharing operates via cross-regional asset holdings. Internationally, the main source of income smoothing comes from international borrowing and lending. Consistent with previous work in the field, we find tentative evidence that risk sharing has declined over time, although the importance of capital market smoothing has gradually increased, consistent with recent increases in capital market integration. However, these trends may be subject to caveats about the changing nature of output risks which our methodologies may not fully detect, and hence require caution in interpretation. Finally, our paper also makes a separate contribution to the literature by illustrating the role of the choice of deflators in estimating the true extent of risk sharing for the United Kingdom and OECD. While estimates of the extent of risk sharing within the United Kingdom are relatively invariant to the choice of deflator, using our preferred choice of deflator for the OECD data set yields higher estimates of risk sharing than typically reported in the literature.

Appendix: Econometric issues

Unit root tests

Both of the methodologies we use in the paper – the panel regression analysis and the factor model – employ some way of detrending the data, and hence assume that the underlying series are I(1). This section documents tests carried out to confirm that this is the case for the series in our data set: both for the absolute data used in our risk-sharing regressions described in Section 2.1 (real per capita consumption or output), and the relative data used in our factor models described in Section 2.2 (real per capita consumption or output relative to the regional or country aggregate).

The results are shown in Tables G and H. For the absolute data, the tests suggested mixed conclusions as to the presence of a unit root in the underlying series: for country-level consumption data in particular, the overall conclusion is sensitive to the inclusion of time trends in the test regressions; however, tests on differenced series suggested that a unit root was not present, in both data sets. For the relative data, our tests more uniformly suggest the presence of a unit root, indicating the data are I(1) and making their detrending valid in subsequent analysis. This applies for both our data sets for the UK regions and OECD countries.

Outliers

We also conducted checks for outliers. For the UK regions (Section 3.1), there are no persistent outliers in either the absolute or relative data which may visibly distort our results, although some regions, such as Northern Ireland, on the whole appear somewhat more volatile in both consumption and output data. In our OECD data set (Section 4.1), individual countries have also exhibited differing patterns of volatility, with our sample of 25 countries including very different economies: for example, small open economies such as Iceland or New Zealand, or countries which have industrialised over our sample period, such as Mexico or South Korea. The persistently divergent experience of some of these economies relative to the rest of the OECD makes the differences in the relative data more acute. Nevertheless, excluding these countries would have neglected a large part of the OECD and potentially biased our results.

Table G: Unit root tests – UK regions

	IPS Wald		ADF-Fischer		PP-Fischer	
	statistic	p-value	statistic	p-value	statistic	p-value
Exogenous regressors: individual fixed effects only						
C absolute	2.71	1.00	10.05	0.99	7.50	1.00
Y absolute	3.12	1.00	5.33	1.00	4.03	1.00
C relative	-1.43	0.08	36.25	0.05	27.00	0.30
Y relative	-1.28	0.10	32.90	0.11	26.56	0.33
Exogenous regressors: individual fixed effects only						
C absolute	-2.06	0.02	43.88	0.01	18.46	0.78
Y absolute	-0.80	0.21	25.00	0.41	12.16	0.98
C relative	0.58	0.72	21.57	0.61	15.17	0.92
Y relative	-1.29	0.10	29.97	0.19	23.50	0.49

	Levin, Lin, Chu		Breitung		Hadri	
	statistic	p-value	statistic	p-value	statistic	p-value
Exogenous regressors: individual fixed effects only						
C absolute	-1.38	0.08	0.77	0.78	12.99	0.00
Y absolute	-0.27	0.39	1.92	0.97	12.58	0.00
C relative	-0.36	0.36	-1.68	0.05	4.81	0.00
Y relative	-0.74	0.23	-0.95	0.17	10.08	0.00
Exogenous regressors: individual fixed effects only						
C absolute	-1.13	0.13	1.61	0.95	5.36	0.00
Y absolute	-0.22	0.41	2.99	1.00	2.24	0.01
C relative	0.71	0.76	1.73	0.96	7.14	0.00
Y relative	-0.46	0.32	1.13	0.87	5.84	0.00

The tables presents test statistics and p-values for five panel unit root tests. The Im, Pesaran and Shin test and the Fisher-ADF and PP tests have the null of a unit root, so ‘acceptance’ of the null (ie a p-value larger than the significance level) suggests that the data should be differenced. These tests allow for different degrees of persistence across the panel. The Levin, Leu and Cheung, Breitung and Hadri tests assess whether there is common persistence in the panel. Of these, the Hadri test has the null of no unit root, so rejection of the null (ie a p-value smaller than the significance level) suggests that the data should be differenced.

Table H: Unit root tests – OECD countries

	IPS Wald		ADF-Fischer		PP-Fischer	
	statistic	p-value	statistic	p-value	statistic	p-value
Exogenous regressors: individual fixed effects only						
C absolute	1.64	0.95	52.57	0.37	67.11	0.05
Y absolute	4.53	1.00	27.91	1.00	26.84	1.00
C relative	0.09	0.54	58.80	0.18	52.32	0.38
Y relative	-0.11	0.46	64.75	0.08	50.16	0.47
Exogenous regressors: individual fixed effects only						
C absolute	-3.42	0.00	85.70	0.00	60.94	0.14
Y absolute	-3.32	0.00	94.35	0.00	46.95	0.60
C relative	-0.16	0.44	50.70	0.45	36.85	0.92
Y relative	0.49	0.69	55.19	0.29	40.07	0.84

	Levin, Lin, Chu		Breitung		Hadri	
	statistic	p-value	statistic	p-value	statistic	p-value
Exogenous regressors: individual fixed effects only						
C absolute	-4.97	0.00	-1.51	0.07	18.60	0.00
Y absolute	-0.52	0.30	-3.03	0.00	18.32	0.00
C relative	-1.18	0.12	0.60	0.73	15.29	0.00
Y relative	-0.27	0.39	-0.47	0.32	15.34	0.00
Exogenous regressors: individual fixed effects only						
C absolute	-2.97	0.00	-1.15	0.12	7.65	0.00
Y absolute	-0.39	0.35	-2.33	0.01	8.59	0.00
C relative	-0.87	0.19	-0.81	0.21	7.81	0.00
Y relative	0.20	0.58	2.17	0.99	9.11	0.00

The tables presents test statistics and p-values for five panel unit root tests. The Im, Pesaran and Shin test and the Fisher-ADF and PP tests have the null of a unit root, so ‘acceptance’ of the null (ie a p-value larger than the significance level) suggests that the data should be differenced. These tests allow for different degrees of persistence across the panel. The Levin, Leu and Cheung, Breitung and Hadri tests assess whether there is common persistence in the panel. Of these, the Hadri test has the null of no unit root, so rejection of the null (ie a p-value smaller than the significance level) suggests that the data should be differenced.

References

Artis, M and Hoffmann, M (2003), 'Home bias and the structure of international and regional business cycles', *European University Institute Working Paper*, no. 15.

Asdrubali, P and Kim, S (2004), 'Dynamic risk sharing in the United States and Europe', *Journal of Monetary Economics*, Vol. 51, pages 809–36.

Asdrubali, P, Sorensen, B E and Yosha, O (1996), 'Channels of interstate risk sharing: United States 1963-1990', *The Quarterly Journal of Economics*, Vol. 102, pages 293–328.

Backus, D K, Kehoe, P J and Kydland, F E (1992), 'International real business cycles', *Journal of Political Economy*, Vol. 100, pages 745–75.

Becker, S and Hoffmann, M (2003), 'Intra- and international risk-sharing in the short run and the long run', *CESifo Working Paper*, no. 1111.

Crucini, M J and Hess, G D (2000), 'International and intranational risk sharing', in Hess, G D and Van Wincoop, E (eds), *Intranational macroeconomics*, Cambridge University Press, pages 37–59.

Dedola, L, Usai, S and Vannini, M (1999), 'An assessment of regional risk sharing in Italy and the United Kingdom', in Adams, J and Pigliaru, F (eds), *Economic growth and change: national and regional patterns of convergence and divergence*, Elgar.

Del Negro, M (2002), 'Asymmetric shocks among US states', *Journal of International Economics*, Vol. 56, pages 273–97.

French, K and Poterba, J (1991), 'Investor diversification and international equity markets', *American Economic Review*, Vol. 81, pages 222–26.

Hayes, P (2004), 'Estimating regional cost of living indices 1974-96', *Regional Studies*, Vol. 38, forthcoming.

Hess, G D and Shin, K (1998), 'Intranational business cycles in the United States', *Journal of International Economics*, Vol. 44, pages 289–313.

Hoffmann, M (2002), 'Financial integration, specialisation and trade', unpublished.

Hoffmann, M (2005), 'International and regional evidence on consumption risk sharing: can trading costs explain the difference?', unpublished.

Imbs, J (2004), 'The real effects of financial integration', *CEPR Discussion Paper*, no. 4335.

Kalemli-Ozcan, S, Sorensen, B E and Yosha, O (2003), 'Risk sharing and industrial specialization: regional and international evidence', *American Economic Review*, Vol. 93, pages 903–18.

Kose, M A, Prasad, E S and Terrones, M E (2003), ‘Financial integration and macroeconomic volatility’, *IMF Working Paper*, no. 50.

Lewis, K (1999), ‘Trying to explain the home bias in equities and consumption’, *Journal of Economic Literature*, Vol. 37, pages 571–608.

Melitz, J and Zumer, F (1999), ‘Interregional and international risk-sharing and lessons for EMU’, *Carnegie-Rochester Series for Public Policy*, Vol. 51, pages 149–88.

Obstfeld, M and Rogoff, K (2000), ‘The six major puzzles in international macroeconomics: is there a common cause?’, *NBER Macroeconomics Annual*, Vol. 15, pages 339–90.

Sorensen, B and Yosha, O (1998), ‘International risk sharing and European monetary unification’, *Journal of International Economics*, Vol. 45, pages 211–38.

Sorensen, B and Yosha, O (2002), ‘Producer prices versus consumer prices in the measurement of risk sharing’, unpublished.