

International and Regional Evidence on Consumption Risk Sharing: can Trading Costs help explain the Difference?¹

Mathias Hoffmann, University of Dortmund & CESifo

This version: June 2005

¹I would like to thank George von Fuerstenberg and seminar participants at the UEM 2003 conference at University College, Oxford, at the Deutsche Bundesbank and at the University of St. Andrews.. This paper is part of the project B6: *The International Allocation of Risk* in the framework of SFB 475 funded by the Deutsche Forschungsgemeinschaft. The UK Economic and Social Research Council funded early stages of the research that led to this paper under its 'Evolving Macroeconomy' scheme (grant # L138251037).

Address for correspondence: Prof. Mathias Hoffmann, University of Dortmund, Dept. of Economics, D-44221 Dortmund, Germany. E-Mail: M.Hoffmann@wiso.uni-dortmund.de

Abstract

The bulk of evidence on the lack of international risk sharing is based on regressions of idiosyncratic consumption growth on idiosyncratic output growth. This paper argues that the results from such regressions obtained from international data are, however, not directly comparable to those based on regional data: the standard practice of running such regressions on international data fails to account for persistent international differentials in consumer prices, whereas – implicitly – most of the literature based on regional data has accounted for these differences. When risk sharing regressions are set up in conceptually the same way in international and regional data sets, the estimated coefficients are also very similar. To explore this result further, we adapt the variance decomposition of Asdrubali Sørensen and Yosha (QJE 1996) to allow for deviations from purchasing power parity across countries. While quantity (income and credit) flows are the dominant channel of risk sharing among regions, relative consumption and output price (internal terms of trade) fluctuations account for the bulk of the deviation from the complete markets outcome in international data. To the extent that persistent differences in consumer prices are an indication of goods market segmentation, our findings provide empirical evidence for the proposition by Obstfeld and Rogoff (2000) that segmented international goods markets rather than asset market incompleteness may account for the (apparent) lack of risk sharing between countries.

1 Introduction

Risk sharing between regions and nations has been the focus of much empirical research over the last decade.¹ The main conclusion that emerges from this literature is that regions within a country share a lot more consumption risk than do countries.

Most of the evidence on international and interregional risk sharing is based on panel regressions of real idiosyncratic consumption growth on other idiosyncratic variables, mostly national or regional output. The motivation behind such risk sharing regressions is that in a world with complete capital markets, countries and regions will insure completely against any idiosyncratic risk. If furthermore, trade in goods markets is frictionless so that prices equalize across countries and regions, then, *ex post*, there should not be any correlation between a country's or region's relative output and consumption. The size of the regression coefficient of idiosyncratic consumption on idiosyncratic output can therefore be interpreted as a measure of the deviation from the complete markets outcome.

Persistent deviations from purchasing power parity (PPP) are, however, one of the salient features of international price data (Rogoff (1996), Engel (1999)). While a number of authors – in particular Backus and Smith (1993), but also Kollmann (1995) and Ravn (2000) — have emphasized that theory predicts a high correlation between real exchange rates and consumption if PPP is violated, the role of international price dynamics for the type of risk sharing regressions typically used in most of the literature has to date not been spelled out. This paper aims to contribute to the closing of this gap.

We put forward a version of the risk sharing regression that accounts explicitly for the role of persistent differences in consumption price levels across countries: when goods markets are segmented and consumer prices do not equalize, then an optimal allocation of risk does not entail that growth rates in the relative quantities of consumption should be independent of measures of idiosyncratic risk. Rather, it is the relative *value* of consumption that should not be affected by idiosyncratic shocks. We therefore suggest to deflate nominal consumption and output with a common price deflator for all countries, so that relative price fluctuations are preserved.

At a conceptual level, this price adjusted risk sharing regression for international data bears a surprising resemblance to the one that most researchers would implicitly use when running risk sharing regressions on regional data. The reason for this is that – for most countries – consumer price indexes

¹Some prominent papers are Asdrubali, Sørensen and Yosha (1996), Sørensen and Yosha (1998), Hess and Shin (1998), Crucini (1999) and Méhitz and Zumer (1999).

are not available at the regional level. Therefore, the commonly applied procedure is to transform nominal into real quantities by deflating with the country-wide CPI. This practice preserves fluctuations in the relative *value* of consumption across regions. In this paper, we advocate this practice also for international data sets. Earlier studies that have examined risk sharing in international data have typically deflated the data with national (i.e. country-specific) CPIs. In this way, only fluctuations in the relative quantities and not in the relative value of consumption are preserved. The procedure commonly used on international data therefore eliminates exactly those relative price effects that are an integral part of the first order condition for the optimal allocation of risk when consumption prices do not equalize internationally – a feature of the data that is particularly important at the international level, where goods markets are highly segmented.

For our empirical investigation, we use data from the Penn World tables for 22 industrialized countries from 1973-2000. Our results show that once prices are treated in the same way, the coefficients estimated from risk sharing regressions are similar in regional and international data sets. Hence, conceptual differences in the preparation of the data used in estimation seem to explain why most studies find very little risk sharing in international data and a lot in regional data.

Does this finding suggest that there is no lack of international risk sharing? To explore the anatomy of this result further, we build on Asdrubali, Sørensen and Yosha (1996) and suggest a decomposition of the variance of country-specific risk that allows to examine to what extent prices and quantity flows respectively contribute to consumption risk sharing. We compare our international results to evidence obtained from regional data sets from Australia, Canada, Germany and Italy – countries for which consumer price data can be obtained at a regional level. We find that regions within countries achieve most of their risk sharing through quantity (income and credit flows), very much as the earlier literature has documented. Also, quantity flows between countries are small, again in line with virtually all of the extant literature. In this sense there is a clear lack of international consumption risk sharing. The reason why we still find a small coefficient when our version of the risk sharing regression is performed on international data, is that the relative internal terms of trade – the ratio between a country’s consumer and output price levels relative to the rest of the world – covary strongly with the relative value of a country’s output. This channel, on the other hand, is virtually absent in regional data, presumably because the cross-regional dispersion of consumer prices is low.

We interpret the degree of international variation in relative prices as an indication of goods market segmentation or, loosely speaking, of trading costs.

This allows us to read the lack of international quantity – i.e. income and credit – flows in the light of a recent literature that emphasizes the role of goods market segmentation in rationalizing some of the major anomalies in international finance. In particular Obstfeld and Rogoff (2000) have argued that relatively small trading costs in goods markets can lead to huge equity portfolio home biases and may therefore also explain the apparent lack of capital flows between countries. In the Obstfeld-Rogoff model this occurs even though financial markets are complete. While our results do not imply that either regional or international financial markets are complete, they provide further empirical support for the view that goods market frictions rather than financial market frictions can explain the lack of risk sharing at the country relative to the regional level: optimal risk sharing under goods market segmentation implies that in response to idiosyncratic output shocks, relative marginal utilities are equated across countries only to the extent that prices are equated. *Ceteris paribus*, larger deviations from the law of one price mean that smaller income and credit flows are required to implement an allocation in which risk is shared optimally.² The comparison between our results obtained from international data on the one hand and regional data on the other suggest that larger variation in relative prices does indeed go in hand with smaller cross-border income and credit flows.

A number of studies have emphasized that market completeness implies a perfect inverse relation between real exchange rates and relative consumption if PPP is violated. The tenor of the studies by Backus and Smith (1993), Kollmann (1995) and Ravn (2000) is that the link between real exchange rates and consumption growth is tenuous at best.

It would therefore seem surprising that our approach reveals such an important role for relative prices in international risk sharing. However, our results are perfectly in line with the observation that real exchange rates and relative consumption are weakly correlated and indeed we corroborate this finding in our data set. As we argue, correlations between consumption and real exchange rates may be low for a variety of reasons that could be unrelated to market incompleteness.³

This is why in this paper we prefer to build on the literature on risk

²We have nothing to say about welfare implications. Clearly, it will be welfare enhancing if there are no transport costs and if prices equalize, even though this may entail more flows of capital (and ultimately shipment of goods). Our interest here is in the optimality (or otherwise) of risk sharing *given* the structure of goods markets, not in assessing the welfare implications of the respective structure.

³Indeed, we find that the correlation between consumption and real exchange rates (i.e. inflation differentials) is very low even in regional data, even though there is wide agreement in the literature that there is quite a lot of risk sharing at the regional level.

sharing *regressions* in the spirit of Cochrane (1991), Mace (1991), Townsend (1994), Asdrubali, Sørensen and Yosha (1996), Crucini (1999) and others. Rather than to examine *correlations* between real exchange rates and consumption, we argue, that a somewhat more robust, reading of the conditions for optimal risk sharing is that the relative value of marginal utility should not be systematically related to a country's idiosyncratic risk.

The remainder of this paper is structured as follows: in section two we review how deviations from purchasing power parity change the optimality conditions for the allocation of consumption risk. We then take stock of the current practice of testing these optimality conditions on regional and international data. We then propose how to adapt the basic specification of risk sharing regressions to take account of deviations from purchasing power parity. In section three, we build on the seminal work by Asdrubali, Sørensen and Yosha (1996) and Sørensen and Yosha (1998) to explore the channels of risk sharing in international data. We suggest a decomposition of business cycle risk into a price channel and a quantity channel. In section four, we present our data set along with our empirical results. Section five discusses and concludes.

2 Risk Sharing and relative prices

We start by considering the general form of the first order condition for utility maximization in a complete markets model with frictionless trade in goods. We write

$$\frac{u_c(C_{t+1}^k, \mathbf{X}_{t+1})}{u_c(C_t^k, \mathbf{X}_t)} = \frac{\mu_{t+1}}{\mu_t} \quad (1)$$

where $u(\cdot)$ is the period utility function, u_c its partial derivative with respect to consumption, C_t^k denotes country k 's consumption level at time t and μ_t is the shadow price of consumption, the Lagrange multiplier associated with the aggregate resource constraint of the economy. The vector \mathbf{X} contains (possibly unobserved) components of utility. In the baseline complete markets model without frictions to trade in goods markets, all countries will face the same shadow price of consumption. If marginal utility growth only depends on consumption, i.e. if \mathbf{X}_t is constant over time, or if utility is additively separable in C and \mathbf{X} , the growth rates of consumption should be highly correlated across countries.

Such assumptions on \mathbf{X}_t or $u(\cdot)$ can be too restrictive for empirical work. It is well known that international consumption correlations are low (see Backus, Kehoe and Kydland (1991)). Even under complete markets, countries or regions may experience idiosyncratic demand or preference shocks

(see e.g. Stockman and Tesar (1995)). An alternative interpretation of equation (1) that is more robust in this respect is that consumption growth may be imperfectly correlated across regions but that it should be uncorrelated with any idiosyncratic country risk characteristics, notably relative output. As a consequence, regressions of consumption growth on indicators of business cycle risk, such as e.g. relative GDP fluctuations should yield a zero coefficient – provided markets are complete. This fundamental insight underlies the regressions first suggested by Mace (1991) and Cochrane (1991) and popularized in the macroeconomic literature by Asdrubali, Sørensen and Yosha (1996) and Sørensen and Yosha (1998). These risk sharing regressions have the form

$$\Delta \log \left[\frac{C_t^k}{C_t^*} \right] = \beta_u \Delta \log \left[\frac{Y_t^k}{Y_t^*} \right] + \varepsilon_t^k \quad (2)$$

where Y is real per capita output and the asterisk denotes the population-weighted rest of the world average of the respective variable. Asdrubali, Sørensen and Yosha (1996) and Sørensen and Yosha (1998) and others have argued very convincingly that the coefficient the coefficient β_u – typically between zero and one in the data – is a measure of the deviation from the complete markets outcome; it indicates *how much* of the idiosyncratic risk represented by fluctuations in $\Delta \log \left[\frac{Y_t^k}{Y_t^*} \right]$ is *not* shared but spills over into fluctuations in relative consumption. We refer to (2) as the quantity-based risk sharing regression that constitutes the standard specification used in many empirical studies of international risk sharing.

Our key argument in this paper is that once the price level of consumption varies across regions, this purely quantity-based regression can be misleading. Rather, the risk sharing regression should take into account that only nominal marginal utilities will be independent of idiosyncratic risk, even under full risk sharing. To see this note that once prices do not equalize across markets, the first order condition for optimal risk sharing can be written as

$$\frac{\mu_t^k u_c(C_{t+1}^k, \mathbf{X}_{t+1})}{\mu_{t+1}^k u_c(C_t^k, \mathbf{X}_t)} = 1 \quad (3)$$

where the shadow price of consumption now varies across countries. A natural measure of μ^k is a country's consumer price index and we will henceforth denote it with CPI_t^k .

There are two empirical implications of (3) of which only one has found its reflection in the literature:

The first is that if \mathbf{X} is constant or if utility is separable in C and \mathbf{X} , then the ‘nominal’ first order condition (3) predicts a perfect inverse relation between relative (i.e. idiosyncratic) consumption growth and the real exchange rate as measured by μ^k/μ^* (where μ^* is the shadow price of consumption in the foreign country). Backus and Smith (1993) as well as Kollmann (1995) and Ravn (2000) have studied this link. All of these papers find the predicted correlation to be highly elusive in the data.⁴

The second implication of (3) is the one we propose to study here: the relative *value* of consumption growth should not be systematically correlated with the realizations of idiosyncratic risk. This reading of (3) is analogous to the reading of the purely quantity-based first-order condition (1) that has motivated the quantity-based risk sharing regressions of the form (2). The major modification we propose is to take account of deviations from the law of one price or purchasing power parity, by regressing the relative growth in the *value* instead of the quantity of consumption on the realizations of idiosyncratic risk.

As we will discuss next, this modification amounts to deflating the data for all regions or countries with a common consumption deflator. This is also the practice typically pursued in regional data – for the simple reason that regional consumer price data are not available for many countries. But virtually all studies using international data would deflate the data for each country with that country’s consumption deflator. An alternative way to read this paper is therefore that we propose to use the same approach when running risk sharing regressions on regional and on international data. We will argue that the use of a common deflator brings the risk sharing regression conceptually closer to a test of the first order condition (3). In discussing our modification to the basic risk sharing regression, we therefore find it useful to start by contrasting the current practice of running risk sharing regressions on regional data with the practice of running them on international data.

⁴Correlations between consumption and real exchange rates may be low for other reasons than market incompleteness. Indeed, market incompleteness alone may not be sufficient to rationalize the correlations between consumption and real exchange rates that is typically found in the data. Baxter and Crucini (1995) have demonstrated that the equilibria in complete market economies are almost identical to those of a bonds-only economy, unless shocks get very persistent. As argued by Corsetti, Dedola and Leduc (2004), it may therefore be rather difficult to generate realistic correlations between real exchange rates and consumption through market incompleteness alone.

2.1 Common vs. region or country-specific consumption deflators

For most countries, notably the U.S., for which we have the most evidence in relation to interregional risk sharing, consumer prices are not available at the regional (say, federal state) level. Most researchers use the country-wide CPI to deflate both GDP and consumption. Since risk sharing regressions are formulated in terms of *relative* growth rates, this amounts to running regressions of region-specific nominal consumption growth on the relative nominal growth rates of output. The risk sharing regression effectively run in regional data is therefore not strictly (2) above, but

$$\Delta \log \left[\frac{CPI_t^k C_t^k}{CPI_t^* C_t^*} \right] = \beta_u^{reg} \Delta \log \left[\frac{P_t^k Y_t^k}{P_t^* Y_t^*} \right] + \varepsilon_{kt}^{reg} \quad (4)$$

Here, P^k/P^* denotes the price level of GDP in region k relative to the national GDP price level, whereas CPI_t^k/CPI_t^* is the relative level of consumer prices in region k , so that $CPI_t^k C_t^k$ is nominal consumption and $P_t^k Y_t^k$ is nominal output.

Conversely, most studies based on international data deflate both consumption and output with the country-specific CPI , so that the risk sharing regression effectively estimated on international data is

$$\Delta \log \left[\frac{C_t^k}{C_t^*} \right] = \beta_u^{int} \Delta \log \left[\frac{P_t^k Y_t^k / CPI_t^k}{P_t^* Y_t^* / CPI_t^*} \right] + \varepsilon_{kt}^{int} \quad (5)$$

The coefficient of the regional regression β_u^{reg} will correctly measure the comovement of the quantities of consumption and output, i.e. β_u^{reg} will equal the coefficient β_u in (2) above, to the extent that $var(\log(CPI_t^k/CPI_t^*))$ and $var(\log(P_t^k/P_t^*))$ are both sufficiently small. In regional data, at least the first condition is relatively likely to be satisfied since consumption price differentials across regions tend to be relatively small, at least a lot smaller than across countries.⁵ Conversely, the coefficient of the regional regression, β_u^{int} , would approximate β_u to the extent that the variance of P_t^k/CPI_t^k is sufficiently small. In most industrialised countries the price levels of GDP and consumption move very closely together, so that this condition is, again, likely to be satisfied in the data.

So, in practice, both the regional and the international risk sharing regressions will therefore generally constitute an appropriate test of the first

⁵Sørensen and Yosha (2002) have argued that P_t^k/P_t^* actually varies quite a lot in regional data, in which case the right hand side of regression (4) will not coincide with that of the quantity-based regression (2) above. We return to this issue below.

order condition (1) pertaining to the model without deviations from purchasing power. Our point here is that this condition does, however, not provide the relevant benchmark in measuring the deviation from the complete market allocation in international data because international price differentials are extremely volatile and persistent. Rather, the modification to the international risk sharing regression that we propose looks very much like the regional risk sharing regression (4).

There is one main reason to believe that the coefficient in (4) is a good parametrization of the deviation of the actual allocation if the full risk sharing allocation characterized by (3): according to (3), the relative growth in the value of marginal utilities should be independent of the realizations of idiosyncratic risk under complete markets. To the extent that we can think of $\Delta \log \left[\frac{C_t^k}{C_t^*} \right]$ as a measure of relative marginal utility growth, $\Delta \log \left[\frac{CPI_t^k C_t^k}{CPI_t^* C_t^*} \right]$ can be interpreted as this relative growth rate in the value of marginal utilities and therefore as the empirical counterpart of $\log \left(\frac{\mu_t^k u_c(C_{t+1}^k, \mathbf{X}_{t+1})}{\mu_{t+1}^k u_c(C_t^k, \mathbf{X}_t)} \right) - \log \left(\frac{\mu_t^* u_c(C_{t+1}^*, \mathbf{X}_{t+1})}{\mu_{t+1}^* u_c(C_t^*, \mathbf{X}_t)} \right)$.⁶ Hence, also the international risk sharing regression should be of the form

$$\Delta \log \left[\frac{CPI_t^k C_t^k}{CPI_t^* C_t^*} \right] = \beta^{int} z_t + \varepsilon_{kt}^{int}$$

where z_t is a measure of the country's idiosyncratic business cycle risk. In choosing z_t , we follow Sørensen and Yosha (2002) in arguing that the regressor $\Delta \log \left[\frac{P_t^k Y_t^k}{P_t^* Y_t^*} \right]$ is a good measure of this idiosyncratic risk because it captures the relative growth rate in the *value* of a country's output rather than the fluctuations in quantity alone.

The maintained hypothesis in all of the literature that has implemented risk sharing regressions such as the ones discussed here is that all of a region's or country's output is perfectly tradeable. This assumption provides a useful benchmark against which to gauge the degree of market completeness. The assumption that output is tradeable also suggests that the same set of prices should be used to value output quantities in different regions and countries. In risk sharing regressions on regional data this issue does not pose a problem: for example, the U.S. state level output price indexes are the output-weighted sums of nationwide sectoral price indexes. Hence, the price used to value

⁶Clearly, $\Delta \log \left[\frac{C_t^k}{C_t^*} \right]$ will only be strictly equal to marginal utility growth under logarithmic utility. But even in more general cases, we can think of $\Delta \log \left[\frac{C_t^k}{C_t^*} \right]$ as a local approximation of marginal utility growth. This is the assumption underlying most of the literature in this area and we do not wish to depart from it here.

e.g. oil production is the same across all states. Since oil accounts for a much larger share of output in some states, such as e.g. Alaska and Texas, fluctuations in the price of oil will still lead to idiosyncratic fluctuations in the value of these states' outputs. Therefore, fluctuations in P^k/P^* by construction cannot reflect deviations from the law of one price but only differences in the sectoral composition of output.

We argue that in order to make the results from risk sharing regressions comparable at the intra-national (regional) and international levels, a common set of international prices should therefore also be used to compare the value of output across countries. At the international level, such a set of prices for the components of GDP is implicit in the data compiled in the Penn World tables. The scope of the PWT is exactly to facilitate such international comparisons of national account data like the one we are conducting here. To this end, the PWT uses a common set of dollar prices to value the components of a country's output basket. This approach is akin to the construction of GDP price indexes at the state level described earlier and we therefore use this data set in our empirical analysis. Hence, the risk sharing regression we suggest in international data has the form

$$\Delta \log \left[\frac{CPI_t^k}{CPI_t^*} \right] + \Delta \log \left[\frac{C_t^k}{C_t^*} \right] = \beta_u \left[\Delta \log \left[\frac{P_{\$}}{P_{\$}^*} \right] + \Delta \log \left[\frac{Y_t^k}{Y_t^*} \right] \right] + \varepsilon_{kt} \quad (6)$$

Here, $P_{\$}/P_{\* denotes the (relative) price level of GDP, in *international* (i.e. PPP) prices. We note again that even though home and foreign GDP are evaluated with the same set of prices, the aggregate GDP price levels do not have to equalize since domestic and rest-of-the-world GDP will generally be composed of very different outputs. It is therefore straightforward to interpret $P_{\$}/P_{\* as a (PPP adjusted) measure of the terms of trade.⁷

One way to interpret this price adjusted risk sharing regression (6) is that we extend the work of Sørensen and Yosha (2002) to those cases in which consumption price levels can diverge between countries. Sørensen and Yosha argue that the proper way to deflate GDP in risk sharing regressions is by the CPI. Fluctuations in the internal price ratio, i.e. of the GDP price level relative to the CPI, reflect changes in the value of a country's output in

⁷It would seem that a similar approach should be adopted for consumption. But we note again that consumption price indexes are not generally constructed at the regional level, therefore the input into risk sharing regressions is just the relative value of consumption deflated with a country-wide deflator, which just amounts to relative nominal consumption expenditure at local prices as shown in (4). Therefore, to get what is conceptually the same regression as on regional data, we also use relative nominal consumption growth in the modified international regression.

terms of its consumption bundle and are therefore important in risk sharing regressions. The specification (6) not only preserves changes in the internal price ratio (i.e. the ratio between the prices of GDP and consumption) but also *relative* fluctuations in consumer and output prices across countries.

All of the above specifications abstract from the role of nominal exchange rate fluctuations. It might seem that our alternative specification should also take account of nominal exchange rates by using nominal exchange rates on both side of the equation so that the regression to be estimated would become

$$\Delta \log \left[\frac{E_t C P I_t^k}{C P I_t^*} \right] + \Delta \log \left[\frac{C_t^k}{C_t^*} \right] = \tilde{\beta}_u \left[\Delta \log \left[\frac{E_t P_t^k}{P_t^*} \right] + \Delta \log \left[\frac{Y_t^k}{Y_t^*} \right] \right] + \varepsilon_{kt}$$

where E_t is the nominal exchange rate expressed as units of foreign currency relative to domestic currency.

There are two main reasons why we do not follow this approach. First, it is important to note that all of the literature on international risk sharing has conditioned on fixed nominal exchange rates and we find it useful to keep with this approach. In particular, we note that the lack of international consumption risk sharing that we seek to understand here has been documented without ever recurring to the role of nominal exchange rates. Our point is that in order to make the outcome of risk sharing regressions as comparable as possible between regions and countries, we should treat the data in the conceptually the same way. Our proposed regression (6) accounts for relative price movements and in this sense captures the logic behind the first order condition (3). But it does so only to the extent that these price movements can in principle play a role in both regional and international data. Clearly, nominal exchange rate movements only matter in international data. So in comparing results from risk sharing regressions in regional and international data, we feel one should abstract from nominal exchange rate risk.

Secondly, nominal exchange rates are known to be disconnected from and much more volatile than macroeconomic fundamentals. In a regression such as (6), the fact that E is likely to be virtually uncorrelated with all other variables in the regression, just implies that $\tilde{\beta}_u$ is close to unity simply because E features on both sides of the equation. While one might interpret this as evidence of low levels of international risk sharing, we would rather read such a result as evidence for the exchange rate disconnect puzzle that will tell us little about consumption risk sharing.

Before moving on to the next section, we introduce some notational simplification: all risk sharing regressions are formulated in idiosyncratic terms, i.e. in relation to a 'rest of the world' aggregate. This just reflects the fact that only idiosyncratic risk can be insured. We will therefore abbreviate

the logarithm of relative levels with the lower case letter, so that $y = \log \frac{Y^k}{Y^*}$, $c = \log \frac{C^k}{C^*}$. It will also prove convenient to abbreviate with gdp the logarithm of the relative value of output, i.e. $gdp = \log \frac{P_{\$}Y^k}{P_{\$}^*Y^*}$.

3 Prices vs. quantities: channels of risk sharing

In a sequence of seminal papers, Asdrubali, Sørensen and Yosha (1996)⁸ and Sørensen and Yosha (1998) also suggested a variance decomposition of relative GDP growth that allows to examine, to which extent capital income and credit flows contribute to consumption risk sharing. They find that at an international level, capital income flows virtually do not contribute at all to consumption insurance whereas borrowing and lending *ex-post* smoothes about one quarter of the variability induced by idiosyncratic output growth. Since their decomposition is based on what we have called the quantity-based risk sharing regression, their setup does not explicitly consider relative price adjustment as a mechanism of risk sharing. We now propose a version of the ASY-decomposition that allows to examine to what extent international risk sharing is achieved through relative price adjustments or through quantity i.e. income or credit flows.

As we have argued, in a world in which prices do not fully equalize, the endowment risk faced by a country is not the sheer quantity of output produced but its relative value, i.e. relative GDP multiplied by the terms of trade. We now write

$$P_{\$}Y = \frac{Y}{INC} \times \frac{INC}{C} \times \frac{P_{\$}}{CPI} \times CPI \times C$$

where the new symbol INC denotes national income and will usually be measured by GNP. We take logarithms and apply the variance operator on both sides. Using the notational convention introduced at the end of the last section, we can write

$$\begin{aligned} var(\Delta \log [P_{\$}Y]) &= cov(\Delta y - \Delta inc, \Delta gdp) \\ &\quad + cov(\Delta inc - \Delta c, \Delta gdp) \\ &\quad + cov(\Delta p_{\$} - \Delta cpi, \Delta gdp) \\ &\quad + cov(\Delta cpi + \Delta c, \Delta gdp) \end{aligned}$$

Dividing through by $var(\Delta gdp)$ we get

$$1 = \beta_{inc} + \beta_{cons} + \beta_{price} + \beta_u$$

⁸For convenience, we will often refer to this paper as ‘ASY’.

where

$$\begin{aligned}\beta_{inc} &= \frac{\text{cov}(\Delta y - \Delta inc, \Delta gdp)}{\text{var}(\Delta gdp)} \\ \beta_{cons} &= \frac{\text{cov}(\Delta inc - \Delta c, \Delta gdp)}{\text{var}(\Delta gdp)} \\ \beta_{price} &= \frac{\text{cov}(\Delta p_s - \Delta cpi, \Delta gdp)}{\text{var}(\Delta gdp)} \\ \beta_u &= \frac{\text{cov}(\Delta cpi + \Delta c, \Delta gdp)}{\text{var}(\Delta gdp)}\end{aligned}$$

Since the wedge between output and income reflects cross-border flows of (capital or dividend) income, β_{inc} measures risk sharing through capital (i.e. equity) markets. In the same mould, β_c measures consumption smoothing through saving or dissaving in credit markets whereas β_u reflects the unsmoothed component of risk. The interpretation of the coefficients β_{inc} , β_{cons} and also β_u is therefore quite analogous to that suggested by ASY (1996) in their version of the variance decomposition of output.⁹ We refer to the sum of β_{inc} and β_{cons} as the quantity channel since it measures how quantity flows in the form of credit or income streams help to stabilize relative marginal utility. We abbreviate the contribution of the quantity channel with $\beta_q = \text{cov}(\Delta y - \Delta c, \Delta gdp) / \text{var}(\Delta gdp)$.

The new channel we introduce is the price channel and its contribution is given by β_{price} . This coefficient measures to what extent international goods market segmentation and hence the possibility of consumer prices to differ across regions can make quantity flows unnecessary for the optimal allocation of risk. The main mechanism we mean to capture with β_{price} was first highlighted by Obstfeld and Rogoff (2000): if consumer prices fall in response to a positive output shock, consumers will take advantage of low prices to increase their consumption. This will induce a positive comovement between (relative) consumption and relative (output). The quantity-based risk sharing regression will register this comovement as a failure of financial markets to provide consumption insurance, even though it may just reflect an optimal response to idiosyncratic price fluctuations.

⁹We note, however, that they are not analytically the same: in ASY and Sørensen and Yosha (1998), the variable with respect to which income and consumption are smoothed is $\Delta \log Y$, in our setup it is $\Delta gdp := \Delta \log [P_s Y]$. We empirically explore the importance of this difference in section 4.4 below.

4 Empirical implementation

4.1 Results from international data

The source of our international data are the Penn World Tables of which we use the most recent release (PWT 6.1.). Besides data from national accounts, the PWT also contain a set of deflators that have been constructed using a set of common international prices. This allows us to obtain a measure of the *international* value of countries' output, i.e. $gdp = p_s + y$.

The PWT expresses all data in per capita terms. We generate the Rest-of-the-world (RoW) aggregate as the population-weighted mean. We construct measures of world-wide (RoW) GDP components using population weighted averages, where the population data is also from the PWT. Our analysis covers a panel of 22 industrialized countries over the period 1973-2000. Virtually all of the countries in the panel are OECD members and we sometimes refer to them under this label. Specifically, the countries in our cross-section are:

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

We start our empirical analysis by estimating our adapted version of the risk sharing regression (6). We then compare the outcomes to that obtained from the pure quantity based regression (2). We take account of country specific fixed effects by removing the mean from each country-time series. By expressing all variables in growth rates relative to the rest of the world, we also account for time-specific fixed effects. We then estimate the risk sharing regressions by means of panel OLS. Throughout the paper, we report heteroskedasticity consistent standard errors based on Newey and West (1987).

In Table (1), the price-adjusted version of the risk sharing regression reveals a lot more risk sharing than the purely quantity based specification: the coefficient on the price adjusted equation is 0.20, whereas the coefficient of quantity-based regression is 0.68. The latter is completely in line with what is typically found in risk sharing regressions based on international data and suggests that only about a quarter to a third of all idiosyncratic country risk is smoothed or insured (e.g. Sørensen and Yosha (1998), Crucini (1999)).

The price-adjusted risk sharing regression, on the other hand, reveals that taking account of differences in consumer prices and correcting for different price levels of GDP matters substantially for the amount of risk that is found

to be shared. Our estimate of 0.20 is rather in the order of magnitude of the coefficient estimated from risk sharing regressions based on US state-level data (Asdrubali, Sørensen and Yosha (1996), Crucini (1999), Méhitz and Zumer (1999)). We address the relation between risk sharing regressions in US state level data and the price adjusted risk sharing regression in a separate subsection below. For now we emphasize that adjusting risk sharing regressions for relative price dynamics seems to reveal a considerable amount of risk sharing in international data.

In the following subsections we first attempt to identify what the channels of risk sharing are in the price adjusted setup. We then briefly relate to the link between real exchange rates and consumption growth rates as it is predicted by the first-order conditions (3) in a baseline complete markets model in which financial markets are complete but in which goods prices do not equalize across countries. Finally, we discuss the relation between our results and those obtained by other authors, notably in US state level data.

4.2 Prices vs. quantities?

Through which channels is the allocation of risk achieved once we take account of both relative price fluctuations and quantity flows? The variant of the ASY decomposition that we suggested in the previous section can shed light on this issue. In table 2 we report the estimates of the β -coefficients. Half of all idiosyncratic risk is buffered by relative movements in the terms of trade and real exchange rates, our point estimate of β_{price} is 0.48. The two quantity-flow channels taken together account for only 30 percent. This suggests that terms of trade fluctuations account for most of the allocation of idiosyncratic risk, much more than do international capital flows.

For comparison, we also estimate a purely quantity-based decomposition of relative GDP growth, i.e. a version of the risk sharing regression where $\Delta y - \Delta inc$ and $\Delta inc - \Delta c$ are regressed on Δy instead of Δgdp . This is the regressions that is typically run on international data, when the data are deflated with the country-specific CPI. The second column of table 2 reports the outcome of this exercise. Even though the regressor is quite different in the price-adjusted regressions, the point estimates for the β s associated with income and consumption smoothing respectively are very similar in both the price adjusted and the pure quantity-based decompositions.

Hence, accounting for deviations from PPP and for terms of trade fluctuations does not alter our conclusions as to what extent international quantity flows contribute to international risk sharing. But it does change the conclusion as to how efficient the international allocation of risk actually is: consumer and output prices covary systematically with idiosyncratic risk in

OECD countries. This effect goes a long way towards explaining why most studies would find more risk sharing in regional than in international data.

4.3 International vs. regional evidence from risk sharing regressions.

The regression specification that we advocate for international data is conceptually equivalent to the one typically used on regional data. It is therefore interesting to find that this approach reveals similar levels of risk sharing in both international and regional data sets.¹⁰ A comparison of our decomposition of risk sharing among countries into a price and the two quantity channels to the results obtained in Sørensen and Yosha (1998) and Becker and Hoffmann (2006) does, however, suggest that quantity flows between countries tend to be a lot smaller than between regions.

We further investigate this issue by looking at regional data from a small group of countries, for which both real and nominal consumption and GDP data are available at the regional level. The countries and the sample of years for which we have data are: Australia (1990-2002), Canada (1980-2000), Italy (1960-96) and Germany (1996-2002). We provide details on the regional data in a separate appendix.

For these countries, we can assess to what extent both the price and the quantity channels contribute to risk sharing. Only for Canada and Germany we have income measures at the regional level. Since the scope of our analysis is not to assess the relative importance of the two quantity channels we just identify the sum of $\beta_{inc} + \beta_c = \beta_q$ by regressing $\Delta y - \Delta c$ on Δgdp . Table (3) provides the estimates of β_q and β_{price} as well as of the unsmoothed part, β_u , for the four countries.

In as far as the size of the unsmoothed component β_u is concerned, our estimates provide a wholesale confirmation of those obtained by ASY and others for U.S. data: roundabout three quarters of all idiosyncratic risk is shared among the regions of a country. While small relative to the standard international quantity regression, the non-insured component is generally significant.

As becomes apparent, the price channel contributes a lot less to risk sharing than it does in international data. This may not appear too surprising since persistent deviations from PPP are a feature rather of international than of regional data. But it is noteworthy that the price effect, though

¹⁰Most of the evidence pertains to risk sharing among U.S. federal states. Compare for example the high levels of risk sharing found by ASY (1996), Crucini (1999), Becker and Hoffmann (2006).

small, is significant in all countries except Germany.¹¹

These results confirm our previous conjecture: controlling for relative price effects in international data reveals a much smaller deviation from the complete markets outcome than is commonly found using what we call a quantity based regression. Eventually, the allocation of risk is comparably efficient to that in regional data. What is different are the *channels* of risk sharing at the regional and the international levels.

Our interpretation of these findings is that goods markets are more integrated among the regions of a country than among countries and therefore price differences are small. Equating the value of marginal utilities across regions therefore virtually amounts to equating real marginal utilities which in turn requires big quantity flows. On the other hand, optimal risk sharing contracts between countries will take account of the fact that goods markets are very segmented internationally. Lane and Milesi-Ferreti (2004) provide evidence that countries that trade a lot with each other also have larger cross-holdings of financial assets. One reason for this may be that trade eliminates price differentials and therefore, consumption insurance can only be achieved through a diversified portfolio of financial assets.

Interestingly, some authors have also documented a quantity anomaly – a high comovement between relative consumption and output – in regional data. Hess and Shin (1998) find that regressions for US state-level income and consumption yield coefficients near unity, not unlike the quantity-based risk sharing regressions obtained from international data. Del Negro (2002) confirms the results obtained by Hess and Shin and claims that the high levels of risk sharing identified by Asdrubali, Sørensen and Yosha (1996) can be explained by measurement error in the ASY data set.

Again, it seems that the principal difference between those studies that find a quantity anomaly and those that find high levels of risk sharing in regional data lies in the way they deflate the data. Asdrubali, Sørensen and Yosha (1996) and Crucini (1999) deflate gross state product (the state level equivalent of GDP) with the consumption price index whereas Hess and Shin deflate GSP with the respective state GSP-deflator. Sørensen and Yosha (2002) argue that the right way to deflate quantities in risk sharing regressions is with the CPI: the endowment risk of an economy is the value of its GDP

¹¹For Australia, the data reveal relative roles of quantity and price channels that are comparable to what we have obtained from international data. While this is an interesting result, we note two things: first, our sample for Australia is rather short. Second, to obtain measure of the regional GDP deflator, we had to use an experimental volume chain index for real state-level GDP. The Australian Bureau of Statistics issues a note of caution regarding the use of this series. We would therefore not overemphasize this particular result.

in terms of the country's consumption basket. This implies that nominal output should be deflated with the CPI. The present paper has extended this logic in two important respects: first, we have argued that the data should be deflated with the area-wide ('common') CPI so that fluctuations in relative consumption price levels are preserved. Secondly, in international data it may also be important to use international prices to price output. We turn to a further discussion of this second point in the next subsection.

4.4 The role of international prices

We have argued that using nominal GDP in national prices rather than international prices may actually lead to a spuriously high correlation between consumption and output fluctuations simply because consumption and output price levels are generally highly correlated. We further illustrate this point in table 4. It presents regression results, based on the national accounts data that are also provided in the PWT. It turns out that the estimated coefficient is even higher than that obtained from a pure quantity regression – 0.86. This result should not be surprising. Panel b) in table 4 shows the regression of CPI inflation on GDP-price inflation based on our panel. The coefficient is virtually unity. Changes in national GDP and consumption price levels are highly positively correlated and this will drive up the estimate of the risk sharing coefficient if national prices are used to evaluate relative fluctuations in the value of GDP. In other words: simply regressing nominal growth rates of consumption and output on each other will also preserve changes in deviations from PPP that are not relevant in determining the risk associated with the relative value of GDP in *international* markets.

In regional data, deviations from PPP in output price levels are ruled out by construction. As we have discussed earlier, it is common practice to construct GDP deflators using a nation-wide set of sectoral price indexes.¹² Therefore, there are no deviations from the law of one price for individual sectoral components of GDP. Relative GDP-price fluctuations therefore only reflect differences in sectoral composition of output across regions and can therefore be thought of as changes in the relative price of output baskets.

In order to make the risk sharing regression in international data as comparable as possible to that generally run on regional data, a data set like the PWT is therefore needed. The PWT imposes an international set of prices on all components of a country's GDP and in so doing just replicates the procedure commonly followed in computing regional GDP price levels.

¹²For example, this is the procedure followed by the U.S. Bureau of Economic Analysis. See their windows help file at <http://www.bea.gov/bean/regional/gsp/OnlineHelp.chm>

One issue that may arise in this context is that the PWT's set of international GDP prices could be subject to measurement error, plausibly on a larger scale than are regional output price levels. This could lead to attenuation bias since relative growth in the value of GDP is the regressor in all our regressions.

We address this problem as follows. We run all regressions based on international data with real rather than nominal relative GDP growth as regressor, i.e. we take out any effect that derives from valuing output at international prices. This has virtually no effect on the coefficients of the quantity channels, a strong indication that measurement error cannot be a problem since it should lead to attenuation bias in both the regressions for the quantity and the price channels. The coefficient on the price channel decreases somewhat. This just indicates that fluctuations in the terms of trade help shield nominal marginal utility from fluctuations in the value of GDP, so that the term $cov(\Delta p_{\$} - \Delta cpi, \Delta p_{\$})$ is positive. While this just suggests that fluctuations in the relative internal terms of trade, $\Delta p_{\$} - \Delta cpi$ shield consumption from fluctuations in the international terms of trade, $\Delta p_{\$}$, the magnitude of this effect seems to be small. Hence, all our findings for international data would remain qualitatively unchanged if we used output Δy instead of $\Delta p_{\$} + \Delta y$ as a regressor in our international risk sharing regression.

4.5 The relation between real exchange rates and consumption

As we have argued, the correlation between relative consumption and relative consumer price levels that is predicted by the optimality conditions (3) is not likely to constitute a good measure of risk sharing. This is why this paper has built on the literature on risk sharing regressions. However, there is a significant literature on the Backus-Smith puzzle and we therefore wrap up the presentation of our results by exploring what our data sets have to say about this issue.

In table (5) we report the results of regressions of inflation differentials and real exchange rate changes on real idiosyncratic consumption growth, both for the international but also for the four regional data sets described above. In international data, the coefficient estimate for the inflation differential regression is -0.20 , whereas the coefficient for the real exchange rate regression is 0.16 . Whereas the first coefficient is highly significant, the second is not. The average (across countries) correlation of relative consumption growth and inflation differentials is -0.19 , the average correlation

between real exchange rates and consumption virtually zero. Regressing inflation differentials on relative consumption growth in regional data, we obtain on average coefficients that are even lower (in absolute value) than in international data. The average correlation ranges from -0.14 to -0.23 . The R^2 statistics are low in all regressions.

Hence, there is virtually no relation between real exchange rates and consumption at the international level, quite in keeping with most results in the literature. If we abstract from nominal exchange rate variability and focus on the link between relative inflation differentials and consumption, we find a noisy but significant inverse relationship in both the regional and international data sets.

We read these findings as follows: it is well known that nominal exchange rates do not move to offset inflation differentials: relative PPP is not generally found to hold in the data. We should therefore not expect to see a close relationship between the real exchange rate and relative consumption movements. However, our results from regional data show that even the elimination of nominal exchange rate variability in financially well integrated areas does not help to establish a particularly pronounced link between the two variables. There is a nascent theoretical literature that may help to explain why the correlation between real exchange rates and real relative consumption is so weak (see e.g. the paper by Corsetti and Dedola (2002)). But the very fact that relative inflation - consumption correlations are low even in regional data seems to suggest that this correlation has relatively little to say about the extent of risk sharing itself.

5 Discussion and Conclusion

A lot more risk is shared among the regions within a country than among countries. This is the evidence from a well-established literature that has looked at risk sharing regressions, i.e. regressions of idiosyncratic consumption growth on idiosyncratic output growth. This paper has argued that the way in which the data have been deflated in risk sharing regressions when applied to international data is conceptually different from the way in which regional data are deflated in such cases: whereas international data are deflated with country-specific CPIs, regional data have mostly been deflated with the country-wide (i.e. a common) CPI.

Our results suggest that accounting for differences in the way that data are deflated can explain why there is an apparent lack of risk sharing between countries. It seems that consumption allocations observed between countries are not as far away from an optimal allocation (of risk) as is often

thought. Movements in the relative price levels of consumption and output account for a lot of the departure from the full risk sharing allocation at the international level. Still, our results corroborate the findings by Sørensen and Yosha (1998), Crucini (1999), Becker and Hoffmann (2006) and others that quantity (income and credit) flows between countries are small relative to quantity flows between regions. Certainly, in this respect there *is* a lack of international risk sharing.

But our interpretation of these findings is that *goods* markets – rather than *financial* markets – are a lot more segmented between countries than between regions. The more segmented goods markets are, the higher the dispersion of prices across regions or countries will generally be. And the more the price of consumption can differ across countries, the less quantity flows are needed to equate changes in the relative *value* of marginal utilities of consumption, which is the condition for an optimal allocation of risk. If it is costly to ship goods, then other things equal, optimal financial contracts will minimize the shipment of capital (and therefore ultimately: goods) between countries. The optimal risk sharing condition (3) reflects this: if prices do not equalize between countries or regions, then consumption should be relatively high when prices are relatively low.

These findings are consistent with a recent strand of the theoretical literature that emphasizes that frictions in international goods markets may be the main ‘culprits’ behind the major quantity anomalies in international macroeconomics (see notably Obstfeld and Rogoff (2000)). They can also help understand the empirical regularity – recently highlighted by Lane and Milesi-Ferretti (2004) – that countries’ international asset portfolio weights are highly correlated with their trade weights: among countries that have highly integrated goods markets and therefore quite similar consumption prices, capital income and credit flows will be needed to achieve an efficient allocation of consumption risk. Therefore, one may expect these countries to have more substantial cross-holdings of financial assets than pairs of countries for which relative price dynamics plays an important role in allocating idiosyncratic risk.

References

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

- [2] Backus, D. and Smith, G. (1993). 'Consumption and Real Exchange Rates in Dynamic Exchange Economies with Nontraded Goods.' *Journal of International Economics* 35, 297-316.
- [3] Becker, S. O. and Hoffmann, M. (2006). 'Intra- and International Risk Sharing in the Short Run and the Long Run.' , *European Economic Review*, forthcoming.
- [4] Cochrane, J. H. (1991) 'A simple test of consumption insurance', *Journal of Political Economy*, vol. 99(5), pp. 957-976.
- [5] Cole, J. and Obstfeld, M. (1991) 'Commodity Trade and International Risk Sharing: How much do financial markets matter?' *Journal of Monetary Economics* 28: 3-24.
- [6] Corsetti, Giancarlo and Luca Dedola (2002). 'The macroeconomics of international price discrimination. ' ECB Working Paper No. 176.
- [7] Corsetti, Giancarlo, Luca Dedola and Sylvain Leduc (2004). 'International risk-sharing and the transmission of productivity shocks.'
- [8] DelNegro, M. (2002), "Asymmetric shocks among U.S. states", *Journal of International Economics*, 56, 273-297.
- [9] Engel, Charles. (1999). 'Accounting for U.S. real exchange rate changes' *Journal of Political Economy*, 107, 507-38.
- [10] Hess, G.D. and Shin, K. (1998). 'Intranational Business Cycles in the United States', *Journal of International Economics*, 44, pp. 289-313.
- [11] Kollmann, R. (1995) 'Consumption, Real Exchange Rates and the Structure of International Capital Markets', *Journal of International Money and Finance*, 14, 191-211.
- [12] Kraay, A. and Ventura, J. (2002). 'Trade integration and risk sharing', *European Economic Review* (June).
- [13] Lane, P. D. and G. M. Milesi-Ferretti (2004). 'International Investment Patterns', mimeo Trinity College Dublin and International Monetary Fund.
- [14] Mace, B. (1991) 'Full Insurance in the Presence of Aggregate Uncertainty', *Journal of Political Economy*, vol. 99(5), pp. 928-956.

- [15] Mélitz, J. and Frédéric Zumer (1999) ‘Smoothing of Regional Shocks within Countries and Lessons for EMU’ *Carnegie-Rochester Conference series on Public Policy*, Volume 51, pp 149-188.
- [16] Newey, W. and K. D. West (1987). "A simple, positive definite, heteroskedasticity and autocorrelation consistent covariance matrix." *Econometrica* 55:3;703-708.
- [17] Obstfeld, M., and K. Rogoff (2000) “The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?,” in *NBER Macroeconomics Annual 2000*, ed. by K. Bernanke, Ben S. Rogoff. MIT Press, Cambridge, Massachusetts.
- [18] Townsend, R. M. (1994), "Risk and insurance in Village India", *Econometrica*, 62, 539-591.
- [19] Ravn, M. (2000), "Consumption dynamics and real exchange rates", mimeo, London Business School.
- [20] Rogoff, K. R. (1996). ‘The Purchasing Power Parity Puzzle’, *Journal of Economic Literature*.
- [21] Sørensen, B. E. and Yosha, O. (1998). ‘International risk sharing and European Monetary Unification.’ *Journal of International Economics*, vol. 45, pp. 211-38.
- [22] Sørensen, B. E. and Yosha, O. (2002). ‘Producer Prices versus Consumer Prices in the Measurement of Risk Sharing’, mimeo, Tel Aviv University.
- [23] Stockman, A C. and Tesar, L. (1995). ‘Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements’, *American Economic Review*, Vol. 85, No. 1. (Mar., 1995), pp. 168-185.

Appendix: Regional Data Sources

Australia: All data are from the Australian Bureau of Statistics and are available at the state level. The CPI data are the CPIs of the respective eight capital cities. Consumption and output are obtained from the breakdown of state level GDP by expenditure and are mid-year estimates (June), ranging from 1990-2002. Income is real gross state domestic income, 1992-2002.

Canada: The data are from Statistics Canada. The data series are personal income, retail sales, population, GDP and CPI by province and range from 1981-2002.

Germany: All data are from the Statistisches Bundesamt, at the federal state level for all 16 federal states. The data range is 1990-2002.

Italy: We used the REGIO-IT data set from the Centro di Ricerche Economiche Nord Sud (CRENoS) at University of Cagliari. The data range from 1960-1996.

Table 1
Price adjusted and quantity based risk sharing regressions

	Price Adjusted		Quantity based	
Coefficient estimate	0.20	(2.48)	0.68	(8.37)

Notes: The price adjusted regression is $\Delta cpi + \Delta c = \beta_u \Delta gdp + u$, the quantity-based regression is $\Delta c = \beta_u \Delta y + v$. Panel OLS estimates with country and time-specific fixed effects. Robust t -statistics based on Newey and West (1987) in parentheses.

Table 2
Channels of Risk Sharing

Channel		Coefficient estimate			
		Price Adjusted		Quantity based	
Quantity	β_{inc}	0.05	(1.52)	0.04	(1.19)
	β_{cons}	0.27	(4.73)	0.29	(3.94)
Price	β_{price}	0.48	(7.89)	—	

Notes: The price adjusted regressions are regressions of $\Delta y - \Delta inc$ (β_{inc}), $\Delta inc - \Delta c$ (β_{cons}) and $\Delta p_{\$} - \Delta cpi$ (β_{price}) on Δgdp . The quantity adjusted regressions are the same variables (except $\Delta p_{\$} - \Delta cpi$) regressed on Δy . On estimation details see table 1.

Table 3

Relative role of price and quantity channels in regional data

Country	Quantity flows	Price	Unsmoothed
	β_q	β_{price}	β_u
Australia	0.43 (4.93)	0.50 (5.34)	0.07 (0.78)
Canada	0.64 (5.80)	0.16 (2.02)	0.18 (2.93)
Germany	0.58 (6.14)	0.01 (0.19)	0.42 (3.55)
Italy	0.77 (16.83)	0.07 (3.67)	0.11 (4.86)

Notes: see table 1.

Table 4

Nominal consumption and GDP growth regression

a)	$[\Delta cpi + \Delta c] = \gamma [\Delta p + \Delta y]$
estimate of γ	0.86 (13.77)
b)	$\Delta cpi = \gamma \Delta p$
estimate of γ	0.98 (22.22)

Notes: $\Delta p + \Delta y$ is nominal GDP growth in local prices. On estimation see table 1.

Table 5
Inflation differentials, real exchange rate changes and relative consumption

		Regression on Δc			average correlation with Δc
		coeff.	t-stat	R^2	
Intl' Data	Δcpi	-0.20	(-2.84)	0.04	-0.19
	$\Delta cpi - \Delta e$	0.16	(1.14)	0.004	0.04
Australia	Δcpi	-0.05	(-2.11)	0.04	-0.19
Canada	Δcpi	-0.03	(-2.41)	0.02	-0.14
Germany	Δcpi	-0.05	(-1.15)	0.03	-0.18
Italy	Δcpi	-0.142	(-3.07)	0.04	-0.23

Notes: The average correlation is $\frac{1}{K} \sum_{k=1}^K \text{corr}(\Delta c_k, \Delta p_k)$ where k is the number of regions or countries.
For notes on regression results see again table 1.