WPS6479

Policy Research Working Paper 6479

Partial Consumption Insurance and Financial Openness Across the World

Constantino Hevia
Luis Servén

The World Bank
Development Research Group
Macroeconomics and Growth Team
May 2013

Policy Research Working Paper 6479

Abstract

This paper examines the extent of international consumption risk sharing for a group of 50 industrial and developing countries. The analysis is based on the empirical implementation of a model of partial consumption insurance whose parameters have the natural interpretation of coefficients of partial risk sharing even when the null hypothesis of perfect risk sharing is rejected. Estimation results show that rich countries exhibit higher degrees of risk sharing than developing countries, and that the gap between both country groups appears to have widened over the period of financial globalization. Moreover, the pattern of consumption risk sharing is related to the degree of financial openness: countries with larger stocks of foreign assets or liabilities exhibit larger degrees of risk sharing. Furthermore, countries whose foreign asset stocks are more tilted towards foreign direct investment assets also show higher degrees of consumption risk sharing.

This paper is a product of the Macroeconomics and Growth Team, Development Research Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://econ.worldbank.org. The authors may be contacted at chevia@worldbank.org and Lserven@worldbank.org.

Partial consumption insurance and financial openness across the world*

Constantino Hevia[a,b], Luis Servén[a],

[a]World Bank, Washington DC, USA

[b]Department of Economics, Universidad Torcuato Di Tella, Buenos Aires, Argentina

Keywords: Risk sharing, Financial globalization, Permanent Income.

JEL classification: E21; F36; F41

*We thank Aart Kraay for useful discussions, and Vivian Norambuena for excellent research assistance. This work was supported by the World Bank's Knowledge Change Program. The views expressed here are only ours and do not necessarily reflect those of the World Bank, its Executive Directors, or the countries they represent.

1 Introduction

A considerable analytical and empirical literature has been concerned with the extent of consumption risk sharing across countries. The relatively high sensitivity of aggregate consumption to domestic income shocks – or, equivalently, the low degree of comovement of aggregate consumption observed across countries – has been singled out as one of the major puzzles in international macroeconomics (Obstfeld and Rogoff 2001). In recent years, it has attracted renewed interest due to the growing degree of financial integration of economies across the world. The theoretical argument is straightforward: if financial markets are complete – in the sense that the available assets suffice to span the set of all idiosyncratic risks – the ratio of the marginal utilities of consumption of any pair of agents (or countries in our case) must be constant across dates and states of nature. That is, the economy features perfect risk sharing.[1] In turn, if risk sharing is imperfect and markets are incomplete, financial innovation that expands the set of tradeable assets (or reduces the costs of trading existing assets) should lead to enhanced risk sharing. To the extent that the global increase in cross-border asset holdings over the last quarter century (Lane and Milesi-Ferretti, 2007) reflects such kind of innovation, it should be associated with a rise in consumption risk sharing across countries.

A number of papers have performed empirical tests of the null hypothesis of perfect consumption risk sharing. The standard approach is to adopt a parametric assumption for the utility function – typically, constant relative or absolute risk aversion – and use it to check if the theoretical prediction is satisfied in the data. This is often done indirectly, through least squares regressions testing whether idiosyncratic income shocks have a significant effect on individual consumption after controlling for the average consumption of all agents (or countries). Obstfeld (1994), Canova and Ravn (1996), and Lewis (1996) are some leading

[1] Strictly speaking, market completeness represents a sufficient, but not necessary, condition for perfect risk-sharing. The same outcome might be achieved by, say, agreement between governments to a suitable system of transfers across countries.

examples of this literature applied to cross-country aggregate data.[2]

These conventional tests can be informative about whether perfect risk sharing holds empirically. But once their null hypothesis has been rejected – as is almost invariably the case – in general they cannot say much about the *extent* of imperfect risk sharing present in the data. Put differently, in most cases one cannot draw inferences about the degree of risk sharing just from the magnitude of consumption correlations or estimated regression coefficients. To do this in a rigorous and meaningful way, one needs a model describing the channels through which the risk sharing arrangement is implemented.

One leading example is the model of partial risk sharing developed by Crucini (1999). It assumes that, prior to any income realization, agents contribute a common fraction of their incomes to an income pool in exchange for the right to the same common fraction of the pooled income after shocks are realized. Thus, the fraction of the income that agents contribute to the pool has the natural interpretation of a coefficient of partial risk sharing. After income shocks are realized and transfers made according to the risk sharing agreement, agents act as permanent income consumers borrowing and lending at a fixed interest rate. This basic framework has been used by Crucini (1999), Crucini and Hess (2000), Asdrubali and Kim (2008), and Artis and Hoffman (2008) to obtain estimates of the extent of partial risk sharing. However, this setting suffers from a major limitation, in that all participants in the risk sharing arrangement are assumed to contribute the same fraction of their income to the common pool. This makes the framework unsuitable for a situation in which different agents or countries may engage in different extents of risk sharing.

This is precisely the focus of this paper. It analyzes to what extent different countries are able to diversify their idiosyncratic risks, and relates their respective degree of risk sharing to selected country characteristics, in particular regarding their international asset portfolios. To do this, we develop an expanded version of the model in Crucini (1999) that allows each agent / country to contribute a different fraction of their income to the income pool,

[2]Cochrane (1991) and Townsend (1994) present early implementations of risk sharing tests based on household level data.

and assume that, after income shocks are realized, the transfer that each agent receives is determined by her initial relative contribution to the income pool.

While this modification may seem pretty intuitive, it has some important consequences. On the one hand, it complicates the solution of the model and its empirical implementation. Specifically, Crucini's (1999) model leads to an estimating equation in which the growth rate of an agent's consumption depends linearly on the innovation to that agent's permanent income and on the growth rate of average consumption across agents. Moreover, the parameters multiplying these variables are linear functions of the common contribution to the income pool, which makes the model suitable for OLS estimation. In contrast, empirical implementation of the model with heterogeneous contributions to the pool requires estimating a *system of equations* in which the path of consumption of each agent depends on the innovation to the permanent income of *all* agents. Moreover, the coefficients multiplying these innovations depend in nonlinear fashion on the contributions to the income pool of *all* agents participating in the risk sharing agreement.

On the other hand, the expanded model solves a problem that plagues all conventional risk sharing regressions that include average consumption among the explanatory variables. As noted by Deaton (1990), pooling the observations on consumption of all agents in a linear regression that includes the cross-sectional average of the dependent variable among the regressors leads, mechanically, to a regression coefficient of unity on that variable. Crucini (1999) dealt with this problem by estimating different coefficients of partial risk sharing for each agent (regions or countries in his analysis), even though his model is built on the assumption of a common coefficient for everyone. To obtain a single estimate of partial risk sharing, as implied by the model, Crucini then takes the average of the estimated risk sharing parameters of all agents.[3] In contrast, our model does not suffer from this problem, because our estimating equation is stated only in terms of innovations to the permanent incomes of all agents. Unlike conventional risk sharing regressions, it does not include the cross-sectional

[3]The paper reports a Monte Carlo experiment to check if this procedure leads to major biases in the risk-sharing estimate.

average of the dependent variable among the regressors.[4]

We estimate the model using a time-series cross-section dataset comprising 50 industrialized and developing countries over the period 1970-2007. We find a cross-country average coefficient of partial risk sharing around 0.42. In terms of the paper's theoretical model, this means that the average country contributes less than half its income to the income pool. The estimates, however, vary systematically between industrial and developing countries, with the former exhibiting, on average, significantly higher degrees of risk sharing than the latter (centered around 0.52 and 0.35, respectively). Re-estimating the model over rolling time samples, we find evidence that the degree of risk sharing has been on the rise among industrial countries during most of the sample period, while for developing countries average risk sharing remained flat or even declined somewhat during the globalization period (although there seems to be a modest uptick at the end of the sample). This result is consistent with the view that the benefits from financial globalization do not accrue evenly to all countries, and may prove elusive in particular for economies with relatively low levels of financial development.[5]

In the paper's framework, these partial risk sharing coefficient estimates can be interpreted as measuring agents' chosen contributions to the global income pool. However, the model is silent on the factors behind those choices. To gain further insight, we assume that they reflect cross-country variation in policies and institutions that help or hinder consumption risk sharing. Specifically, we re-estimate the model expressing the partial risk sharing coefficient as a function of selected measures of financial and trade openness, which are commonly viewed as reflecting the mechanisms through which risk sharing may be actually implemented.

While this adds further complication to an already-challenging nonlinear estimation problem, it yields fairly robust empirical results. The main conclusion is that international financial integration is a significant factor behind the observed cross-country patterns of consumption risk sharing. The degree of financial integration, as summarized by total foreign

[4] Moreover, even if we had written the estimating equation in terms of a weighted consumption average, the nonlinearity of the model would break the linear relation between right and left-hand side variables inherent to the linear regression.

[5] On the costs and benefits of financial globalization, see Kose, Prassad, Rogoff, and Wei (2006).

asset and liability positions, is positively correlated with the coefficient of partial risk sharing, consistent with the view that financial integration improves risk sharing across countries. In addition, the FDI share of the foreign asset position – i.e., the proportion of residents' foreign assets that takes the form of direct investment abroad – is also positively correlated with consumption risk sharing even after controlling for total foreign assets. The implication is that FDI assets are more conducive to risk sharing than other kinds of foreign assets. In contrast, the FDI share of foreign liabilities – i.e., the portion that takes the form of foreigners' direct investment in the home country – is negatively associated with consumption risk sharing. This finding seems to be at odds with the conventional wisdom that capital inflows in the form of FDI are preferable, from the point of view of risk, to other forms of capital inflows.[6]

Importantly, once these measures of international financial integration are taken into account, trade integration, domestic financial development, and the level of GDP per capita do not contribute significantly to explain the cross-country pattern of consumption risk sharing.

Our paper is related to a substantial empirical literature assessing international consumption risk sharing. Sorensen et al. (2007) relate trends in risk sharing among OECD countries to their foreign assets and liabilities. They find that the degree of consumption risk sharing appears to be positively related to foreign asset holdings, while the relation with foreign liabilities is not robust. Kose, Prasad, and Terrones (2009) argue that risk sharing has risen among industrial countries (but not among developing countries), and the rise is correlated with the increase in gross foreign assets and liabilities over the globalization period. Holinski et al. (2012) also examine how international consumption risk sharing relates to various features of countries' equity portfolios. These papers base their conclusions on conventional risk sharing regressions, and thus they are subject to the criticism that, strictly speaking, such regressions do not provide a solid basis for inferences about the extent of partial risk sharing. Fratzscher

[6]We conjecture that the result might reflect the action of third factors that both hinder risk sharing and tilt the composition of inflows towards direct investment. In other words, a large share of FDI liabilities could be the constrained optimal mix of foreign financing for a country with poor institutions and weak governance and, therefore, with a riskier environment and less opportunities for risk sharing (Albuquerque, 2003, Daude and Fratzscher, 2008).

and Imbs (2009) develop a model with transaction costs and discuss their effects on conventional tests of risk sharing. They find that larger holdings of foreign capital (especially in the form of equity or bonds) are associated with higher consumption risk sharing. In contrast, larger holdings of FDI or bank loans are not. Flood, Marion, and Matsumoto (2009) use the variance of a country's share of world consumption as a measure of consumption risk sharing. Perfect consumption risk sharing occurs when this share is constant. Using rolling windows, they argue that consumption risk sharing rose during the recent era of financial globalization, particularly when the focus is on low-frequency movements in consumption. Finally,Bai and Zhang (2012) argue that financial frictions can explain why risk sharing failed to improve during the globalization era. They develop a model with incomplete financial markets and enforceability constraints and show that removing capital controls leads to the emergence of default risk that limits the extent of consumption risk sharing.

The rest of the paper is organized as follows. Section 2 develops the model of partial risk sharing and section 3 describes the data and the details of the empirical implementation of the model. Section 4 discusses the main empirical results of the paper and section 5 concludes. The appendices contain some proofs and additional tables.

2 A model of partial risk sharing

In this section we describe a model of partial risk sharing that generalizes the model in Crucini (1999). Time is discrete and denoted by $t = 0, 1, 2, ..., \infty$. There is a group of $N > 1$ families. Family $i = 1, 2, .., N$ is composed of H_i members, each of whom receives a random flow of income Y_{it} in period t. We assume that families enter into a partial risk sharing agreement before income is realized. The agreement requires each family to contribute a possibly different fraction of their income into an income pool in exchange for a claim to a fraction of the pooled income of all families. In particular, if each member of family i contributes a fraction λ_i of her income to the pool, the contribution to the fund at time t is equal to $\lambda_i H_i Y_{it}$, and the aggregate pooled income equals $\sum_{i=1}^{N} H_i \lambda_i Y_{it}$. We assume that the agreement specifies that each family i gets back a fraction of the income pool equal to its share in the total initial contribution, or

$$\theta_i = \frac{\lambda_i H_i Y_{i0}}{\sum_{j=1}^{N} \lambda_j H_j Y_{j0}}. \tag{1}$$

Therefore, total income of family i at time t after the risk sharing agreement is

$$H_i \bar{Y}_{it} = (1 - \lambda_i) H_i Y_{it} + \frac{\lambda_i H_i Y_{i0}}{\sum_{j=1}^{N} \lambda_j H_j Y_{j0}} \sum_{k=1}^{N} \lambda_k H_k Y_{kt}.$$

In per capita terms, each member of family i receives

$$\bar{Y}_{it} = (1 - \lambda_i) Y_{it} + \lambda_i \sum_{k=1}^{N} \left(\frac{\omega_{i0}}{\omega_{k0}} \right) \theta_k Y_{kt}, \tag{2}$$

where $\omega_{i0} = Y_{i0} / \sum_{j=1}^{N} Y_{j0}$ is the share of family i's per capita income in aggregate per-capita income.

We also assume that, on top of the risk sharing agreement, each family member acts as a permanent income consumer borrowing or lending at a fixed interest rate r. Following Campbell and Mankiw (1990), the permanent income hypothesis is represented by the following decision rule for the change in consumption,

$$\Delta C_{it} = \mu_i + \bar{\varepsilon}_{it}, \tag{3}$$

where, for any variable x_t, $\Delta x_t = x_t - x_{t-1}$ is the difference operator and $\bar{\varepsilon}_{it}$ is the innovation between time t and $t-1$ in family i's per-capita permanent disposable income. The latter, in turn, is given by

$$\bar{\varepsilon}_{it} = (1 - \beta) \sum_{t=0}^{\infty} \beta^t \left(E_t \bar{Y}_{it+k} - E_{t-1} \bar{Y}_{it+k} \right) \tag{4}$$

where $\beta = 1/(1+r)$ is the reciprocal of the gross interest rate and E_t is the expectations operator conditional on all the information available to the family at time t. The above expression measures the innovation to permanent income after the risk sharing agreement. A related object is the innovation to permanent income before the risk sharing agreement, denoted by

$$\varepsilon_{it} = (1-\beta) \sum_{t=0}^{\infty} \beta^t \left(E_t Y_{it+k} - E_{t-1} Y_{it+k} \right). \tag{5}$$

Using equations (2), (4), and (5) into equation (3) delivers the following expression for the change in consumption as a function of the innovation to all families' permanent incomes before risk sharing

$$\Delta C_{it} = \mu_i + (1-\lambda_i)\,\varepsilon_{it} + \lambda_i \sum_{j=1}^{N} \frac{\omega_{i0}}{\omega_{j0}} \theta_j \varepsilon_{jt}. \tag{6}$$

The left side of this expression is the change in family i's per-capita consumption. The term on the right side includes the deterministic drift in consumption plus the innovation to the family's permanent disposable income. The latter term, in turn, can be written as

$$\lambda_i \left(\sum_{j=1}^{N} \frac{\omega_{i0}}{\omega_{j0}} \theta_j \varepsilon_{jt} - \varepsilon_{it} \right) + \varepsilon_{it},$$

which illustrates the risk sharing agreement. In particular, if family i's has a bad income realization relative to that of the other families (if the term in parentheses is positive), it receives a transfer proportional to the difference in the income shocks, where the factor of proportionality is the contribution of family i to the income pool. If family i has a good income realization relative to that of the other families (if the term in parentheses is negative), it transfers resources to the other families in proportion to its contribution λ_i. Furthermore, note that if $\lambda_i = 1$, so that country i shares all of its income with the pool, the idiosyncratic shock ε_{it} disappears as an isolated term from expression (6) and only enters through its impact on the innovation to the pool's permanent income $\sum_{j=1}^{N} \frac{\omega_{i0}}{\omega_{j0}} \theta_j \varepsilon_{jt}$.

Before proceeding, we discuss an important specification issue. As noted in Mankiw and Campbell (1989,1990) and Crucini (1999), aggregate consumption and income time series seem to follow log-linear rather than linear processes. We thus follow Mankiw and Campbell (1989,1990), Crucini (1999), and Asdrubali and Kim (2008) and replace all variables by their logs and consider the resulting expression as a log-linear approximation to the permanent income hypothesis equations. (The Appendix shows how to log-linearize the permanent income innovation.) Therefore, (6) is approximated according to

$$\Delta c_{it} = \mu_i + (1 - \lambda_i)\, \varepsilon_{it} + \lambda_i \sum_{j=1}^{N} \frac{\omega_{i0}}{\omega_{j0}} \theta_j \varepsilon_{jt},$$

where $c_{it} = \log C_{it}$, $y_{it} = \log Y_{it}$, and

$$\varepsilon_{it} \cong \sum_{k=0}^{\infty} \beta^k \left[E_t \Delta y_{it+k} - E_{t-1} \Delta y_{it+k} \right].$$

In the econometric analysis below we interpret each family as a country and the number of family members as that country's population. Our objective is then to estimate the degree of risk sharing λ_i for each country $i = 1, 2, ..., N$. To that end, we follow the usual procedure in the rational expectations econometric literature (Hansen and Sargent, 1980) and assume that, when performing revisions to their permanent incomes, agents have more information than the econometrician. Let $\hat{\varepsilon}_{it}$ denote the innovation to the permanent income estimated by the econometrician and rewrite (6) as

$$\Delta c_{it} = \mu_i + (1 - \lambda_i)\, \hat{\varepsilon}_{it} + \lambda_i \sum_{j=1}^{N} \frac{\omega_{i0}}{\omega_{j0}} \theta_j \hat{\varepsilon}_{jt} + u_{it}, \tag{7}$$

where u_{it} is a residual that represents the superior information of agents relative to the econometrician, and is given by

$$u_{it} = (1 - \lambda_i)\, \varepsilon_{it} + \lambda_i \sum_{j=1}^{N} \frac{\omega_{i0}}{\omega_{j0}} \theta_j \varepsilon_{jt} - \left((1 - \lambda_i)\, \hat{\varepsilon}_{it} + \lambda_i \sum_{j=1}^{N} \frac{\omega_{i0}}{\omega_{j0}} \theta_j \hat{\varepsilon}_{jt} \right),$$

It is easy to verify that u_{it} is orthogonal to the econometrician's information set at time

$t-1$.[7] This information set, however, may omit some relevant information used by the agents to forecast their future incomes. This omitted information can lead to serially correlated residuals for a given agent as well as contemporaneously correlated residuals across agents. The former could be due to a persistent variable that is used by an agent, but unobserved by the econometrician, to perform revisions to his or her permanent income. The latter could be due to an aggregate shock observed by all agents, but unobserved by the econometrician, that affects everyone's permanent income.

It is important to stress the key difference between our framework, as summarized by (7), and that in Crucini (1999). In both cases, the starting point is the assumption that agents contribute a fraction of their income to an income pool and behave as permanent-income consumers after transfers are realized. In Crucini (1999), however, all agents are assumed to contribute the same fraction of their income to the pool, which forces their respective partial risk sharing coefficients λ_i to be all equal and allows a drastic simplification of the model.[8] However, instead of estimating a single λ for all agents, Crucini estimates λ_i separately for each agent through OLS regressions, and then computes a single coefficient λ by averaging the individual coefficients. While this procedure might be justified in Crucini's empirical application, whose focus is on relatively homogeneous groups of regions or countries – U.S. states, Canadian provinces, and G-7 countries – it is clearly less defensible in our case given the likely great degree of heterogeneity across the fifty countries in our sample. Moreover, one of our objectives is precisely to assess the extent to which the coefficients of partial risk sharing vary across countries. This requires a more general framework such as (7), in which different countries may exhibit different coefficients of partial risk sharing.

However, the added generality of this framework relative to the conventional model used in the literature poses additional econometric challenges too. For each country, (7) involves a nonlinear function of the risk sharing parameters of *all* countries. Thus, empirical imple-

[7] The proof of this result is an application of the law of iterated expectations (Hansen and Sargent, 1980). In addition, any classical measurement error in consumption can also be included in the error term u_{it}.

[8] This simplifying assumption is maintained in subsequent empirical applications of Crucini's framework; see for example Artis and Hoffman (2008) and Asdrubali and Kim (2008).

mentation of the model requires more complex econometric techniques than those found in earlier literature. This is the topic of the next section.

3 Empirical implementation

We first discuss the econometric issues surrounding the estimation of the parameters of the risk sharing model, as well as those related to the construction of the permanent income innovations required for such task. Then we briefly summarize our data sources.

3.1 Implementing the risk sharing model

The core of our empirical analysis is the estimation of (7) using a cross-country time-series dataset. This poses a challenging problem because each country's consumption path depends in nonlinear fashion on the risk sharing parameters of *all* countries. To make this explicit in what follows, it is useful to rewrite (7) as

$$\Delta c_{it} = \mu_i + \lambda_i \sum_{j=1}^{N} \left(\frac{\omega_{i0}}{\omega_{j0}} \right) \frac{H_j \lambda_j \omega_{j0}}{\sum\limits_{k=1}^{N} H_k \lambda_k \omega_{k0}} \hat{\varepsilon}_{jt} + (1 - \lambda_i)\, \hat{\varepsilon}_{it} + u_{it} \tag{8}$$

where u_{it} is potentially heteroskedastic and possibly correlated across countries and over time.

To express the empirical model in more compact form, it is convenient to introduce some notation. Let $\Delta \mathbf{c}_i = (\Delta c_{i1}, ..., \Delta c_{iT})'$and $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, ..., \varepsilon_{iT})'$ denote the (column) vectors of consumption growth and permanent income innovations of country i, and define the $T \times N$ matrix $\mathbf{Z} = (\boldsymbol{\varepsilon}_1, ..., \boldsymbol{\varepsilon}_N)$. In addition, let $\boldsymbol{\lambda} = (\lambda_1, ..., \lambda_N)'$, the vector of risk sharing parameters for all countries, and define the $N \times N$ diagonal matrices $\mathbf{H}$ and $\mathbf{W}_0$, which have along their main diagonal the population of each country H_{j0} and its per capita GDP share ω_{j0}, respectively, with all off-diagonal elements equal to zero. After some straightforward manipulations the observations of (8) corresponding to the i-th country can be compactly written

$$\Delta \mathbf{c}_i = \boldsymbol{\iota}_T \mu_i + \mathbf{Z}\left(\mathbf{e}_i \mathbf{e}_i' \left(\boldsymbol{\iota}_N - \boldsymbol{\lambda}\right) + \mathbf{H}\boldsymbol{\lambda}(\boldsymbol{\iota}_N' \mathbf{H}\mathbf{W}_0 \boldsymbol{\lambda})^{-\mathbf{1}} \mathbf{e}_i' \mathbf{W}_0^{-1} \boldsymbol{\lambda}\right) + \mathbf{u}_i \tag{9}$$

where $\boldsymbol{\iota}$ denotes a column vector of ones, and $\mathbf{e}_i$ is an $N \times 1$ vector of zeros with an 1 in the i-th entry.

Finally, stacking the observations on consumption growth for all countries into the $NT \times 1$ vector $\Delta \mathbf{c} = (\Delta \mathbf{c}_1', ..., \Delta \mathbf{c}_N')'$ and letting $\boldsymbol{\mu} = (\mu_1, ..., \mu_N)'$, the full system of equations can be written

$$\Delta \mathbf{c} = \boldsymbol{\mu} \otimes \boldsymbol{\iota}_T + (\mathbf{I}_N \otimes \mathbf{Z}) \left[\left(\mathbf{W}_0^{-1} \boldsymbol{\lambda} \otimes \mathbf{H}\boldsymbol{\lambda}(\boldsymbol{\iota}_{\mathbf{N}}' \mathbf{H}\mathbf{W}_0 \boldsymbol{\lambda})^{-\mathbf{1}}\right) + vec\left(\mathbf{I}_N - diag(\boldsymbol{\lambda})\right)\right] + \mathbf{u} \tag{10}$$

where $\otimes$ denotes the Kronecker product, and $diag(\boldsymbol{\lambda})$ is an $N \times N$ diagonal matrix with its main diagonal equal to $\boldsymbol{\lambda}$ and zeros elsewhere.

Thus, the empirical model amounts to a system of N equations with cross-equation parameter restrictions. The restrictions imply that a system estimation procedure is needed, even though the explanatory variables (the innovations to permanent income in all N countries, contained in the $T \times N$ matrix $\mathbf{Z}$) are the same in all equations. In this context, we use system NLS to estimate the parameters of (10). We first partial out $\boldsymbol{\mu}$ by expressing $\Delta \mathbf{c}$ and $\mathbf{Z}$ as deviations from their country-specific means; let $\Delta \tilde{\mathbf{c}}$ and $\tilde{\mathbf{Z}}$ denote the transformed variables. Then we compute the NLS estimator that solves the problem

$$\min_{\lambda} \left(\Delta \tilde{\mathbf{c}} - \mathbf{f}(\tilde{\mathbf{Z}}, \mathbf{H}, \mathbf{W_0}, \boldsymbol{\lambda})\right)' \left(\Delta \tilde{\mathbf{c}} - \mathbf{f}(\tilde{\mathbf{Z}}, \mathbf{H}, \mathbf{W_0}, \boldsymbol{\lambda})\right) \tag{11}$$

where $\mathbf{f}(\tilde{\mathbf{Z}}, \mathbf{H}, \mathbf{W_0}, \boldsymbol{\lambda}) = \left(\mathbf{I}_N \otimes \tilde{\mathbf{Z}}\right) \left[\left(\mathbf{W}_0^{-1} \boldsymbol{\lambda} \otimes \mathbf{H}\boldsymbol{\lambda}(\boldsymbol{\iota}_{\mathbf{N}}' \mathbf{H}\mathbf{W}_0 \boldsymbol{\lambda})^{-\mathbf{1}}\right) + vec\left(\mathbf{I}_N - diag(\boldsymbol{\lambda})\right)\right]$. Since the residuals may exhibit heteroskedasticity, serial correlation, cross-sectional dependence, or any combination of all three, to perform inference on $\boldsymbol{\lambda}$ we use the robust covariance matrix estimator proposed by Driscoll and Kraay (1998) and Vogelsang (2012).[9]

[9] In the formula of the covariance matrix, the usual matrix of regressors is replaced in our case by the

This procedure yields a set of unrestricted estimates of the risk sharing coefficients $\hat{\boldsymbol{\lambda}}$. However, we are also interested in learning about the covariates of risk sharing. To do this, we restrict the risk sharing coefficients so that they satisfy $\boldsymbol{\lambda} = \mathbf{X}\boldsymbol{\delta}$, where $\mathbf{X}$ is an $N \times K_{\lambda}$ matrix whose i-th row contains the (time-invariant) covariates of risk sharing for country i (including a constant). Replacing $\boldsymbol{\lambda}$ in (10), estimation proceeds along the same lines as above, with the parameter vector now given by $\boldsymbol{\delta}$. That is, we solve

$$\min_{\delta} \left(\Delta\tilde{\mathbf{c}} - \mathbf{f}(\tilde{\mathbf{Z}}, \mathbf{H}, \mathbf{W_0}, \mathbf{X}, \boldsymbol{\delta})\right)' \left(\Delta\tilde{\mathbf{c}} - \mathbf{f}(\tilde{\mathbf{Z}}, \mathbf{H}, \mathbf{W_0}, \mathbf{X}, \boldsymbol{\delta})\right) \tag{12}$$

where now $\mathbf{f}(\tilde{\mathbf{Z}}, \mathbf{H}, \mathbf{W_0}, \mathbf{X}, \boldsymbol{\delta}) = \left(\mathbf{I}_N \otimes \tilde{\mathbf{Z}}\right) \left[\left(\mathbf{W}_0^{-1}\mathbf{X}\boldsymbol{\delta} \otimes \mathbf{H}\mathbf{X}\boldsymbol{\delta}(\boldsymbol{\iota}_{\mathbf{N}}'\mathbf{H}\mathbf{W_0}\mathbf{X}\boldsymbol{\delta})^{-\mathbf{1}}\right) + vec\left(\mathbf{I}_N - diag(\mathbf{X}\boldsymbol{\delta})\right)\right]$. Like in the unrestricted case, we use the Driscoll and Kraay (1998) robust covariance matrix estimator to conduct inference on $\boldsymbol{\delta}$.

3.2 Income prediction

Empirical implementation of the risk sharing model requires suitable forecasts of permanent income. To construct them, we use a simple time-series setting allowing for common factors affecting per capita income growth across countries. Specifically, we assume that per capita GDP growth depends on its lagged value as well as the current and lagged values of world per capita income:

$$\Delta y_{it} = \alpha_i + a_i \Delta y_{it-1} + b_i \Delta\bar{y}_t + c_i \Delta\bar{y}_{t-1} + v_{it} \qquad i = 1, ..., N;\ t = 1, ..., T \tag{13}$$

where we have defined the growth rate of world per capita GDP $\Delta\bar{y}_t = \sum_{j=1}^{N} s_j \Delta y_{jt} = \mathbf{s}'\Delta\mathbf{y}_t$; $\Delta\mathbf{y}_t = (\Delta y_{1t}, ..., \Delta y_{Nt})'$, and $\mathbf{s} = (s_1, s_2, ..., s_N)'$ is a set of shares, which will be taken as constant in what follows, with $\boldsymbol{\iota}_N'\mathbf{s} = 1$. In turn, v_{it} is a disturbance assumed independent across t and i. If $b_i = c_i = 0$ for all i, (13) reduces to a set of country-specific first order autoregressive models (AR1), similar to that employed by Crucini (1999). In turn, if $c_i = 0$ for

matrix of partial derivatives $\nabla_{\boldsymbol{\lambda}}\mathbf{f}$.

all i but $b_i \neq 0$, we have an intermediate case in which common factors have contemporaneous effects but not lagged effects.

Stacking the T observations for country i, we can write

$$\Delta \mathbf{y}_i - \left(\mathbf{s}^{'} \otimes \mathbf{I}_T\right) \Delta \mathbf{y} b_i = \boldsymbol{\iota}_T \alpha_i + \Delta \mathbf{y}_{i,-1} a_i + \left(\mathbf{s}^{'} \otimes \mathbf{I}_T\right) \Delta \mathbf{y}_{-1} c_i + \mathbf{v}_i \tag{14}$$

where $\Delta \mathbf{y}_i, \mathbf{v}_i$ and $\Delta \mathbf{y}_{i,-1}$ are $T \times 1$ vectors, and $\Delta \mathbf{y} = (\Delta \mathbf{y}_1', ..., \Delta \mathbf{y}_N')'$. Defining the $N \times 1$ parameter vectors $\boldsymbol{\alpha}, \mathbf{b}$, and $\mathbf{c}$, the $NT \times 1$ vector $\mathbf{v} = (\mathbf{v}_1', ..., \mathbf{v}_N')'$,and the matrix $diag(\mathbf{a})$ with $\mathbf{a} = (a_1, ..., a_N)'$ along the main diagonal and zeros everywhere else, the full system can be compactly expressed

$$[(\mathbf{I}_N - \mathbf{bs}') \otimes \mathbf{I}_T] \Delta \mathbf{y} = (\boldsymbol{\alpha} \otimes \boldsymbol{\iota}_T) + \left[\left(diag(\mathbf{a}) + \mathbf{cs}^{'}\right) \otimes \mathbf{I}_T\right] \Delta \mathbf{y}_{-1} + \mathbf{v}. \tag{15}$$

Like in the case of the risk sharing model described above, estimation of $\boldsymbol{\theta} \equiv (\boldsymbol{\alpha}', \mathbf{a}', \mathbf{b}', \mathbf{c}')'$ can be accomplished by solving the problem

$$\min_{\theta} \mathbf{g}(\Delta \mathbf{y}, \Delta \mathbf{y}_{-\mathbf{1}}, \mathbf{s}, \boldsymbol{\theta})' \mathbf{g}(\Delta \mathbf{y}, \Delta \mathbf{y}_{-\mathbf{1}}, \mathbf{s}, \boldsymbol{\theta})$$

where $\mathbf{g}(\Delta \mathbf{y}, \Delta \mathbf{y}_{-\mathbf{1}}, \mathbf{s}, \boldsymbol{\theta}) = [(\mathbf{I}_N - \mathbf{bs}') \otimes \mathbf{I}_T] \Delta \mathbf{y} - (\boldsymbol{\alpha} \otimes \boldsymbol{\iota}_T) + \left[\left(diag(\mathbf{a}) + \mathbf{cs}^{'}\right) \otimes \mathbf{I}_T\right] \Delta \mathbf{y}_{-1}$. With v_{it} uncorrelated across i and t, a heteroskedasticity-robust covariance matrix for $\hat{\boldsymbol{\theta}}$ can be constructed as

$$Cov(\hat{\boldsymbol{\theta}}) = [\nabla_{\boldsymbol{\theta}} \hat{\mathbf{g}}' \nabla_{\boldsymbol{\theta}} \hat{\mathbf{g}}]^{-1} \left[\nabla_{\boldsymbol{\theta}}' \hat{\mathbf{g}} \ diag(\hat{\mathbf{v}}) \ diag(\hat{\mathbf{v}})' \ \nabla_{\boldsymbol{\theta}} \hat{\mathbf{g}}\right] \nabla_{\boldsymbol{\theta}} \hat{\mathbf{g}}' \nabla_{\boldsymbol{\theta}} \hat{\mathbf{g}}^{-1},$$

where $diag(\hat{\mathbf{v}})$ is an $NT \times NT$ matrix with the estimated residuals along the main diagonal.

Using the parameter estimates, the innovation to permanent income can be constructed using the recursive formula

$$\Delta \mathbf{y}_t = \tilde{\boldsymbol{\alpha}} + \hat{\mathbf{P}} \Delta \mathbf{y}_{t-1} + \hat{\mathbf{w}}_t \tag{16}$$

where $\hat{\mathbf{P}} \equiv \left(\mathbf{I}_N + \hat{\mathbf{b}}\left(1 - \mathbf{s}'\hat{\mathbf{b}}\right)^{-1}\mathbf{s}'\right)\left(\hat{\mathbf{A}} + \hat{\mathbf{c}}\mathbf{s}'\right)$, $\tilde{\boldsymbol{\alpha}} \equiv \left(\mathbf{I_N} + \hat{\mathbf{b}}\left(\mathbf{1} - \mathbf{s}'\hat{\mathbf{b}}\right)^{-\mathbf{1}}\mathbf{s}'\right)\hat{\boldsymbol{\alpha}}$, and $\hat{\mathbf{w}}_t \equiv \left(\mathbf{I}_N + \hat{\mathbf{b}}\left(1 - \mathbf{s}'\hat{\mathbf{b}}\right)^{-1}\mathbf{s}'\right)\hat{\mathbf{v}}_t$. Simple algebra shows that the innovation to permanent income $\mathbf{yp}_t$ then is

$$\begin{aligned} \mathbf{yp}_t - E_{t-1}\mathbf{yp}_t &= \left(\mathbf{I}_N + \beta\left(\mathbf{I}_N - \beta\hat{\mathbf{P}}\right)^{-1}\hat{\mathbf{P}}\right)\hat{\mathbf{w}}_t \\ &= \left(\mathbf{I}_N + \beta\left(\mathbf{I}_N - \beta\hat{\mathbf{P}}\right)^{-1}\hat{\mathbf{P}}\right)\left(\mathbf{I}_N + \hat{\mathbf{b}}\left(1 - \mathbf{s}'\hat{\mathbf{b}}\right)^{-1}\mathbf{s}'\right)\hat{\mathbf{v}}_t \qquad (17) \end{aligned}$$

3.3 Data

The empirical sample is dictated by data availability. We work with the fifty largest countries in the world (in terms of their respective GDP in U.S. dollars in the year 2000) for which complete annual data on consumption, income and foreign asset and liability positions (as described below) could be assembled. Table B1 in the appendix lists the countries included in the analysis. Taken together, these countries account for over 93 percent of world GDP in the year 2000.

Data on real GDP growth, aggregate consumption growth and total population are taken from the Penn World Tables 6.3. Along the time dimension, our sample runs from 1970 to 2007. We use the 1970-71 averages to construct the initial shares $\mathbf{W}_0$ and population $\mathbf{H}$.[10] The regression sample therefore runs from 1972 to 2007, and comprises 1,800 country-year observations. For the calculation of permanent income, we set the discount factor $\beta = 1/1.05$.

For our analysis of the covariates of risk sharing, we use data on international asset and liability positions drawn from the Lane and Milesi-Ferretti dataset released on August 2009,[11] which covers the period 1970-2007. Specifically, we consider the ratios of total foreign assets and liabilities to GDP, as well as the ratios of FDI assets and liabilities to GDP. These variables relate to countries' international financial integration. In addition, we also

[10] However, estaimtion results change very little if we use instead the avearges over 1970-74 or the year 1970 only.

[11] Specifically, the data are drawn from the updated and extended version (as of August 2009) of the External Wealth of Nations Mark II database developed by Lane and Milesi-Ferretti (2007).

experiment with trade openness (as measured by imports plus exports divided by GDP) as a measure of real integration. We include a variable measuring real integration because, in the end, any redistribution of endowments due to (implicit or explicit) risk sharing agreements should be materialized through flows of goods. Lastly, because domestic financial depth can also contribute to risk sharing among domestic agents (perhaps lessening the need for, or the value of, international risk sharing) we also use data on domestic financial depth, as captured by the ratio of domestic credit to GDP. Like the trade openness indicator, the credit depth measure is drawn from the World Development Indicators. To avoid simultaneity concerns, the empirical exercises assessing the covariates of risk sharing employ these variables measured at the beginning of the sample – that is, in 1972 (or the first year in which data is available).

4 Empirical results

We proceed in two stages. First, we estimate the income process given by (15) and construct the innovations to permanent income (17); then we use the latter to estimate the risk sharing model (10).

4.1 Estimation of the income process

To estimate (15), we construct the country shares of global GDP $\mathbf{s}$ using averages over the entire sample period.[12] We estimate three different specifications of (15). The first one ignores common shocks by setting $\mathbf{b} = \mathbf{c} = \mathbf{0}$, and therefore simplifies to a system of N AR1 equations, similar to that employed by Crucini (1999). The second specification allows for contemporaneous effects of common factors by leaving $\mathbf{b}$ unrestricted while still imposing $\mathbf{c} = \mathbf{0}$. The third specification allows for both current and lagged common shocks by leaving both $\mathbf{b}$ and $\mathbf{c}$ unrestricted.

Table 1 summarizes the estimation results. The first column reports the results of the AR1

[12] The correlation between the growth rate of global GDP per capita and the sum of individual-country per capita GDP growth rates weighted by these shares equals 0.95.

specification. The average of the country-specific autoregressive parameters (denoted a_i in (13)) equals 0.244. The Wald test shows that they are jointly highly significant. Inspection of the individual-country estimates shows that they are all smaller than one in absolute value (the largest one equals 0.658), so the estimated growth dynamics are stable. However, Pesaran's (2004, 2012) cross-sectional dependence test overwhelmingly rejects the null that the estimation residuals are uncorrelated across countries, which suggests the presence of common factors – although, strictly speaking, the performance of the test in settings like ours featuring cross-equation parameter restrictions has not been examined in the literature.

Column 2 of table 1 turns to the model with contemporaneous common shocks. Overall, these are highly significant (as shown by the corresponding Wald test) and relatively large: the mean of the country-specific parameters on global GDP growth (denoted b_i in (13)) equals 0.835. In turn, the mean of the country-specific autoregressive parameter estimates remains roughly similar to that in column 1. However, inspection of the characteristic roots of the implied transition matrix $\hat{\mathbf{P}}$ in (16) now reveals the presence of one negative root outside the unit circle. Lastly, the cross-section dependence test statistic is not significant anymore at the 5 percent level, although it still exceeds the 10 percent critical value.

Column 3 reports estimates of the full model including current and lagged global growth. The lagged effect (c_i in (13)) is, on average, negative and small, and not significantly different from zero. However, the Wald test statistic reported in the lower half of the table shows that the country-specific parameters on lagged global growth are jointly highly significant, just like those on current global growth and lagged country-specific growth. Inspection of the individual-country estimates shows that 11 countries exhibit estimates of the c_i parameter significant at the 5 percent level. In turn, the country-specific autoregressive parameter estimates, as well as those on contemporaneous global growth, remain roughly similar to those in column 2. Further, all the roots of the $\hat{\mathbf{P}}$ matrix are smaller than one in absolute value (the largest one equals 0.879). Finally, the test statistic of cross-sectional dependence falls short of 10 percent significance, suggesting that current and lagged global GDP growth suffice to capture common shocks to the countries in the sample.

Using the parameter estimates summarized in Table 1, the innovations to permanent income can be readily constructed employing (17). We construct two sets of innovations, corresponding to the first and third model specifications in the table; we leave aside the second specification because of the unstable root mentioned earlier. The two sets of innovations constructed in this way are not very different from each other: the correlation between them equals 0.92.

4.2 Estimation of the risk sharing model

Table 2 summarizes the NLS estimates of the risk sharing parameters obtained using the two sets of permanent income innovations constructed in the previous step. The results are very consistent across the two specifications. They can be summarized in three main facts. First, 46 of the 50 individual-country estimates lie between zero and one, as predicted by theory. When using the AR1-based innovations, the four remaining point estimates are negative; in contrast, when using the system-based innovations, three are negative and one exceeds 1. However, in no case can we reject the null hypothesis that these parameters lie between zero and one, whether individually or jointly. Moreover, the vast majority of the individual estimates in the admissible region (41 out of 46) are statistically significant at least at the 10 percent level; they are overwhelmingly significant jointly.

Second, the country-specific risk sharing parameters are centered around 0.40-0.42. This is roughly similar to the result obtained by Crucini (1999) using the restricted homogenous model of partial risk sharing; he reports an average partial risk sharing parameter between 0.37 and 0.60 for the group of G7 countries over the period 1970-1987 (his estimates vary depending on the assumed income process). In turn, Asdrubali and Kim (2008) obtain higher coefficients of partial risk sharing, namely 0.77 for a group of 15 European Union countries and 0.82 for all OECD countries, both over the period 1960-2004.

And third, the data overwhelmingly reject the null hypothesis that the risk sharing coefficient is the same for all sample countries – as shown by the Wald test reported in the middle of the table. Furthermore, the estimates differ systematically between industrial and developing economies: on average, the former exhibit higher degrees of risk sharing than the latter. The same conclusion is reached by Kose, Prasad, and Terrones (2009) using conventional risk sharing regressions, and by Flood, Marion, and Matsumoto (2012) using a different measure of imperfect risk sharing. In our case, the mean for industrial countries exceeds 0.5, while for developing countries it equals 0.35. As the table also shows, the difference between both means is highly significant. However, additional tests show that even within these two groups the individual-country coefficients display significant variation. Indeed, Wald tests of equality of the risk sharing coefficients of all the countries in each income group overwhelmingly reject the null, both among developing countries and among industrial countries.

Table B2 in the appendix reports the parameter estimates for the individual countries for both specifications in Table 2. For the majority of countries, the parameters show only minor changes across specifications; indeed, the correlation between the two sets of estimates is 0.91. One notable exception to this finding is the U.S., whose risk sharing parameter estimate equals 0.93 when we allow for common factors across countries, but rises to a whooping 1.5 when we use the independent AR1 processes for income growth. This difference could be due to the large effect of U.S. growth on global economic activity and, thereby, on growth in all the other countries – a feature that is built into the system estimates of the income process but not into the AR1 estimates. Using the latter therefore results in a distorted indication of the degree of U.S. risk-sharing.

The estimates reported in Table 2 are obtained without imposing any restriction on their admissible values. As a result, four of the individual point estimates lie outside the [0,1] interval assumed by theory, although – as already noted – not significantly so. Still, given the fact that in our empirical model the risk sharing parameters of all countries are linked through cross-equation restrictions, one may wonder how the estimates would change if they were explicitly constrained to lie within the unit interval. To investigate this issue, we re-estimated the two specifications in Table 2 constraining all parameters to fall within the theoretically-admissible region. The estimates that result are virtually indistinguishable from those shown in Table 2. In fact, the correlation between matching pairs of unrestricted and restricted estimates is around 0.98 in both the AR1 and the common-factor models. Thus, to save space we do not report the restricted estimates.

4.2.1 Patterns of consumption risk sharing over time

Since the mid-1980s, the world has seen a large and ever-increasing level of financial integration, facilitated by the removal of barriers to international capital movements and reflected in a steady rise in cross-border asset and liability positions (Lane and Milesi-Ferretti, 2007). This observation leads to a natural question, namely: has the rise in global integration been associated with a corresponding rise in consumption risk sharing across countries? Put differently, are countries doing a better job at sharing their idiosyncratic income shocks? To address this question, we use our model of partial risk sharing to examine the evolution of consumption risk sharing over time.

For this purpose, we calculate time-varying estimates of the parameters of the consumption risk sharing model over rolling time samples. Specifically, we use moving windows of 21 years each. With 36 years of data in total, this yields 16 sets of estimates. Importantly, prior to estimation of the risk sharing model we also recalculate the permanent income innovations in a matching way, by estimating the income process over rolling samples of 21 years of data each; this ensures that both construction of the permanent income innovations and estimation of the risk sharing parameters use data from the same time period.[13]

The results from this exercise are summarized in figures 1 and 2. Figure 1 plots the mean of the risk sharing parameter estimates obtained in each estimation window, along with their two standard-error bands. The two graphs correspond to the estimates obtained using the rolling AR1 GDP growth forecasts and the estimates obtained using the system-based growth

[13] Similarly, for each window we recalculate population $\mathbf{H}$ and the initial income shares $\mathbf{W}_0$ using the average of the two annual observations preceding the initial year of the window.

forecasts over rolling windows.

In both graphs, mean risk sharing displays a cycle: a slight initial decline, which is reversed in the windows that start in the late 1970s, and then a steady increase beginning around the windows that start in the mid-1980s. Closer inspection reveals that in the top graph the mean estimate obtained in the final window lies outside the 95-percent confidence region of the initial window, suggesting that over the entire sample period there has been a statistically significant increase in the average degree of consumption risk sharing, at least when the calculation of the latter is based on income forecasts that allow for common factors. Examining the individual-country risk sharing parameter estimates we find that a majority of countries (33 when using the common factor-inclusive forecasts, and 27 when using the AR1 forecasts) exhibit higher coefficients in the final estimation window than in the initial one. Further, for 16 countries (15 when using the AR1 forecasts) the increase in the risk sharing parameter is statistically significant at least at the 10 percent level. At the other end, 10 countries (or eight, when using AR1 income growth forecasts) show statistically significant declines in risk sharing between the initial and final estimation windows.

Figure 2 plots the time path of the average risk sharing coefficients separately for industrial and developing countries, along with their respective two standard-error bands, for the same two sets of permanent income innovations as in the previous figure. Both panels reveal a clear contrast between the mean estimates of the two country groups. The mean for industrial countries displays a rising trend over time, particularly marked in the top panel, which is based on the permanent income forecasts inclusive of common factors. Indeed, this rising trend of average risk sharing among rich countries is the force behind the trend increase in overall mean risk sharing found in Figure 1. More formally, in both panels the 95-percent confidence region for the industrial country mean corresponding to the final window lies fully outside that obtained in the initial window. In other words, over the period of analysis there has been a highly significant rise in mean risk sharing among industrial countries. In contrast, there has been no discernible trend among developing countries. While the figure suggests, if anything, a slight decline in their mean risk sharing, the 95 percent confidence region of the final window is contained almost in full in the initial one, which suggests that no significant change has taken place. The contrasting time path of the two group means reflects the fact that 10 of the 16 countries (11 out of 15, when using the AR1 forecasts) showing a statistically significant (at the 10 percent level) increase in the coefficient of risk sharing belong to the industrial country group. In contrast, all of the countries whose coefficient of risk sharing exhibits a statistically significant decline belong to the developing country group.

4.2.2 Covariates of consumption risk sharing

The empirical results reported earlier unambiguously showed that the coefficients of partial risk sharing display significant heterogeneity across countries. While the country-level risk sharing arrangement that underlies our model is admittedly an abstraction, we can interpret the variation in the estimated coefficients of partial risk sharing as reflecting cross-country variation in policies or institutions that help or hinder consumption risk sharing. The natural next step is to investigate if the cross-country pattern of coefficients of partial risk sharing is related to measures of financial and trade openness commonly viewed as reflecting the mechanisms through which risk sharing may be actually implemented. This approach has been used in previous empirical literature, which has examined the relation between summary measures of the extent and form of international financial integration, and conventional reduced-form estimates of risk sharing coefficients. For example, using this approach, Kose, Prasad, and Terrones (2009) conjecture that emerging economies have failed to attain the levels of consumption risk sharing of the developed countries because their international liabilities have been dominated by foreign debt instead of other, more resilient, liabilities like FDI or portfolio flows.[14]

Formally, to explore the covariates of consumption risk sharing we re-estimate the risk sharing model as specified in (12), that is, with the vector of country-specific risk sharing

[14] See also Sorensen et al (2007) and Kolinsli et al (2012). In turn, Fratzscher and Imbs (2009) correlate their coefficients of bilateral risk-sharing to measures of financial opennes and institutional quality.

coefficients given by $\boldsymbol{\lambda} = \mathbf{X}\boldsymbol{\delta}$, where $\mathbf{X}$ is a matrix of time-invariant risk sharing determinants, and $\boldsymbol{\delta}$ now is the vector of parameters to be estimated. To avoid simultaneity concerns, the variables in $\mathbf{X}$ are measured at the beginning of the sample period, and expressed in logs in order to mitigate the influence of extreme observations. [15]

Table 3 reports the estimates of $\boldsymbol{\delta}$ obtained using different choices of the $\mathbf{X}$ variables. All specifications include also a constant, not reported in the table. The results shown correspond to estimates obtained with permanent income innovations constructed using income forecasts inclusive of common factors (that is, based on the model in the last column of Table 1); results using instead AR1-based forecasts are roughly similar but are not reported to save space.

The first column of Table 3 relates the degree of risk sharing to total foreign assets and total foreign liabilities, both as percent of GDP. Both variables carry positive and significant coefficients, suggesting that higher degrees of international financial integration, as measured by gross foreign asset and liability stocks, come along with higher consumption risk sharing.[16]

Using the parameter estimates, we can compute the implied risk sharing coefficients of the different countries as $\hat{\boldsymbol{\lambda}} = \mathbf{X}\hat{\boldsymbol{\delta}}$. As shown in the bottom part of the table, they are centered around 0.47, with the individual coefficients ranging from .097 to 0.903. Further, on average they are higher for industrial than for developing countries, in line with the estimation results reported in Table 2 above. Finally, the last row of Table 3 shows the correlation between these implied risk sharing coefficients and their unrestricted counterparts (shown in the first column of Table B2). For the specification in column 1, the correlation is just 0.31, which suggests that foreign asset and liability stocks account for a significant but relatively modest proportion of the cross-country variation in consumption risk sharing.

Column 2 of Table 3 adds to the total foreign asset and liability stocks a measure of

[15] For a few countries in the sample, the initial observation (corresponding to the year 1972) on some of the foreign assets and liabilities included among the $\mathbf{X}$ variables is missing from the Lane-Milesi dataset. In those cases, we use instead the earliest available value. Additionally, for some countries the ratios of FDI assets and/or liabilities to total foreign assets and liabilities contain some zeros in the initial sample years. To allow taking logs, we replace them with very small values (1.e-6).

[16] In reduced-form regressions, Kose et al (2009) obtain a similar result for industrial countries, while Sorensen et al (2007) find that gross foreign assets are positively related to risk sharing among OECD countries, while for gross foreign liabilities the relationship is negative instead.

their composition, specifically distinguishing between FDI and other types of foreign assets and liabilities, to test whether direct investment is different regarding its contribution to risk sharing.[17] In contrast with the preceding column, the estimated coefficients on FDI asset and liability positions, relative to overall asset and liability stocks, carry opposing signs, suggesting that residents' direct investment abroad helps risk sharing more than other types of foreign asset holdings, but non-residents' direct investment at home has the opposite effect. While this latter result seems to be at odds with theoretical expectations, we note that there is a body of evidence that documents that the share of FDI in total inflows tends to be larger in countries that are riskier and possess weaker institutions (Hausmann and Fernandez-Arias, 2000; Albuquerque, 2003, and Daude and Fratzscher, 2008). In other words, a large share of FDI liabilities could be the constrained optimal mix of foreign financing for a country with both a riskier environment and fewer opportunities for risk sharing.

The range of the implied risk sharing coefficients narrows somewhat relative to that in column 1, but the opposite happens to the gap between the industrial-country and developing-country averages. Moreover, the correlation with the unrestricted risk sharing coefficient estimates in Table B2 rises substantially, to 0.47.

Going beyond foreign assets and liabilities, columns 3 to 5 assess the contribution of two other potentially important variables, namely trade integration, as measured by imports plus exports over GDP, and domestic financial development, as measured by the stock of credit to the private sector over GDP. However, the estimation results suggest that these variables add little to the determination of consumption risk sharing patterns, since they are both insignificant, whether individually or jointly. Indeed, a Wald test of the null that in column 5 both variables carry a zero coefficient yields a p-value of 0.80, suggesting that they can be safely dropped from the specification. On the other hand, the estimated coefficients on the volume and composition of foreign assets and liabilities show minimal change relative to those in column 2.

[17] The estimation results are virtually identical if instead of FDI we focus on equity plus FDI assets and liabilities vs the rest of assets and liabilities.

Since the unrestricted risk sharing coefficient estimates shown in table 2 above were, on average, higher for industrial than for developing countries, it is natural to wonder if their cross-country pattern is somehow related to countries' level of development – as captured by (log) per capita GDP – even after taking into account their degree of international financial integration. This question is addressed in columns 6 and 7, which add per capita GDP to the list of regressors of columns 2 and 5, respectively. In both cases the per capita GDP parameter estimate falls well short of statistical significance. Further, in column 7 a Wald test fails to reject the null that all the variables other than total foreign assets and liabilities, and their respective FDI shares, carry zero coefficients (with a p-value of 0.89).

Overall, these empirical exercises suggest that international financial integration is a significant factor behind the cross-country patterns of consumption risk sharing. Specifically, both the degree of overall integration, as reflected in total foreign asset and liability positions, and the FDI asset position, are positively related to countries' performance in terms of consumption risk sharing, while FDI liabilities show the opposite pattern. Moreover, once these summary measures of international financial integration are taken into account, trade integration, domestic financial development, and the overall level of development have little to contribute.

5 Concluding remarks

A considerable literature is concerned with assessing the extent to which countries share their consumption risk. Much of it, however, makes use of empirical models designed to test the null hypothesis of perfect consumption risk sharing. Once such hypothesis is rejected – as is almost invariably the case in practice – those models do not offer a rigorous basis for inferences about the degree of imperfect risk sharing present in the data. Drawing such inferences requires an empirical model of partial risk sharing. Furthermore, unless one is willing to assume that the extent of consumption risk diversification is the same across all countries in the world, the model needs to allow for cross-country heterogeneity in the degree of risk sharing.

This paper extends the existing literature by developing an empirical model that fits those two requirements. In the model, countries contribute possibly different fractions of their income to a common pool, in exchange for a claim on the aggregate income contributed to the pool by all countries. The fraction of income contributed to the global pool by each country can be naturally interpreted as its respective degree of risk sharing. Solution of the model yields a system of equations in which each country's consumption path depends on the innovations to the permanent income of all countries. Moreover, the system features nonlinear parameter restrictions across equations.

The model is implemented empirically using panel data for industrial and developing economies. Estimation results show that consumption risk sharing varies significantly across countries. On the whole, rich countries exhibit higher degrees of risk sharing than developing countries, and the gap between both country groups appears to have widened over the period of financial globalization. Moreover, the pattern of consumption risk sharing across countries is significantly related to their degree of financial openness. Countries possessing larger stocks of international assets and/or liabilities exhibit larger degrees of risk sharing. Countries whose foreign asset stocks are more tilted towards FDI assets show higher consumption risk sharing, while the opposite happens with countries whose foreign liability stocks involve larger FDI liabilities. Once these measures of international financial integration are taken into account, trade integration, domestic financial development and the level of GDP per capita do not contribute significantly to explain the cross-country pattern of consumption risk sharing.

References

[1] Artis M.J. and M. Hoffman. (2008). Financial globalization, international business cycles and consumption risk sharing. *Scandinavian Journal of Economics*, Vol. 110, No. 3, 447-471.

[2] Asdrubali P. and S. Kim. (2008). Incomplete intertemporal consumption smoothing and incomplete risk sharing. *Journal of Money, Credit and Banking*, Vol. 40, No. 7, 1521-1531.

[3] Bai, Y. and J. Zhang. (2012). Financial integration and international risk sharing. *Journal of International Economics*, Vol 88, No.1, 17-32.

[4] Baltagai, B., S. Song, B. Jung and W. Koh (2007). Testing for serial correlation, spatial autocorrelation and random effects using panel data. *Journal of Econometrics* 140, 5-51.

[5] Campbell, J. Y. and N. G. Mankiw (1989). Consumption, income, and interest rates: reinterpreting the time series evidence. In O. Blanchard and S. Fisher (eds.), *NBER Macroeconomics Annual*. Cambridge: MIT Press, 185-216.

[6] Campbell, J. Y. and N. G. Mankiw (1990). Permanent income, current income, and consumption. *Journal of Business & Economic Statistics*, Vol. 8, No. 3, 265-279.

[7] Canova, F. and M. Ravn (1996). International consumption risk sharing. *International Economic Review* 37, 573-601.

[8] Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy*, Vol. 99, No. 5, 957-976.

[9] Crucini, M. J. (1999). On international and national dimensions of risk sharing. *The Review of Economics and Statistics*, Vol. 81, No.1, 73-84.

[10] Crucini, M. J. and G.D. Hess. (2000). International and intranational risk sharing. In Hess, G. D. and E. Van Wincoop (eds.), *Intranational Macroeconomics*, Cambridge University Press, 37–59.

[11] Daude, C. and M. Fratzscher. (2008). The pecking order of cross-border investment. *Journal of International Economics* Vol. 74, No. 1, 99-119.

[12] Deaton, A. (1990). On Risk, Insurance and Intra-Village Consumption Smoothing. *Research Program in Development Studies*, Princeton University, N.J

[13] Driscoll, J. and A. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. *The Review of Economics and Statistics,* Vol 80, No. 4, 549-560.

[14] Flood, R.P., N.P. Marion, and A. Matsumoto. (2012). International risk sharing during the globalization era. *Canadian Journal of Economics.* Vol. 45, No. 2, 394-416.

[15] Fratzscher M. and J. Imbs. (2009). Risk sharing, finance, and institutions in international portfolios. *Journal of Financial Economics.* Vol. 94, No. 3, 428-447.

[16] Hansen, L.P. and T. J. Sargent. (1980). Formulating and estimating dynamic linear rational expectations models. *Journal of Economic Dynamics and Control*, Vol. 2, 7-46.

[17] Holinski, N., C. Kool and J. Muyksen (2012). The impact of international portfolio composition on consumption risk sharing. *Journal of International Money and Finance* 31, 1715-1728.

[18] Kose, M. A., E. S. Prasad, K. Rogoff, and S. J. Wei. (2006). Financial globalization: a reappraisal. IMF Working Paper 06/189.

[19] Kose, M. A., E. S. Prasad, and M. E. Terrones. (2009). Does financial globalization promote risk sharing? *Journal of Development Economics* Vol. 89, No. 2, 258–270.

[20] Lane P. R. and G. M. Milesi-Ferretti (2007). The external wealth of nations mark II: Revised and extended estimates of foreign assets and liabilities, 1970–2004. *Journal of International Economics* Vol. 73 No. 2, 223-250.

[21] Lewis, K. (1996). What can explain the apparent lack of consumption risk sharing? *Journal of Political Economy* 104, 267-297.

[22] Obstfeld, M. (1994). Are industrial country risks globally diversified? in L. Leiderman and A. Razin (eds.), *Capital Mobility: the Impact on Consumption, Investment and Growth.* Cambridge University Press.

[23] Obstfeld, M. and K. Rogoff (2001). The Six Major Puzzles in International Macroeconomics: Is There a Common Cause? *NBER Macroeconomics Annual 2000.*

[24] Pesaran, H. (2004). General diagnostic tests for cross section dependence in panels, CESifo Working Paper 1229.

[25] Pesaran, H. (2012). Testing weak cross-sectional dependence in large panels, IZA Discussion Paper 6432.

[26] Sorensen, B., Y. Wu, O. Yosha and Y. Zhu (2007). Home bias and international risk sharing: twin puzzles separated at birth. *Journal of International Money and Finance* 26, 587-605.

[27] Townsend, R. (1994). Risk and Insurance in Village India. *Econometrica*, Vol. 63, No. 3, 539-591.

[28] Vogelsang, T. (2012). Heteroskedasticity, autocorrelation and spatial correlation robust inference in linear panel models with fixed-effects. *Journal of Econometrics* 166, 303-319.

Appendix A. Proofs.

Here we show how to log-linearize the innovation to the permanent income. Consider the present value of a history of future incomes in some country (we omit the subindex i),

$$\begin{aligned} Y_t^p &= (1-\beta)\sum_{k=0}^{\infty}\beta^k Y_{t+k} \\ &\cong Y_{t-1}(1-\beta)\sum_{k=0}^{\infty}\beta^k\left(1+y_{t+k}-y_{t-1}\right), \end{aligned}$$

where $y_t \equiv \log Y_t$ and we use the approximation $y_{t+k}-y_{t-1} \cong \left(Y_{t+k}-Y_{t-1}\right)/Y_{t-1}$. By adding and subtracting y_{t+s} for $s = 1, 2, ...$ it follows that Y_t^p can be written as

$$\begin{aligned} Y_t^p &\cong Y_{t-1}\left[1+(1-\beta)\sum_{k=0}^{\infty}\beta^k\sum_{s=0}^{k}\Delta y_{t+s}\right] \\ &\cong Y_{t-1}\left[1+\sum_{k=0}^{\infty}\beta^k\Delta y_{t+k}\right]. \end{aligned}$$

Taking logs and using the approximation $\log(1+x) \cong x$ then gives

$$y_t^p = y_{t-1} + \sum_{k=0}^{\infty}\beta^k\Delta y_{t+k}.$$

The approximation to the innovation to the permanent income is defined as

$$E_t y_t^p - E_{t-1} y_t^p = \sum_{k=0}^{\infty}\beta^k\left[E_t\Delta y_{t+k} - E_{t-1}\Delta y_{t+k}\right].$$

so that the innovation in permanent income is approximated by the innovation in future growth rates.

Appendix B. Additional Tables

Table B1
List of countries

Industrial	Developing
Australia	Algeria
Austria	Argentina
Belgium	Bangladesh
Canada	Brazil
Denmark	Chile
Finland	China
France	Colombia
Germany	Dominican Republic
Greece	Egypt
Ireland	Hong Kong
Italy	Hungary
Japan	India
Netherlands	Indonesia
New Zealand	Israel
Norway	Korea, Republic of
Portugal	Malaysia
Spain	Mexico
Sweden	Morocco
Switzerland	Pakistan
United Kingdom	Peru
United States	Philippines
	Poland
	Romania
	Singapore
	South Africa
	Thailand
	Turkey
	Venezuela
	Vietnam

Table B2
Individual risk sharing coefficients

Industrial	Income prediction model		Developing	Income prediction model	
	System	AR1		System	AR1
Australia	0.457	0.293	Algeria	0.447	0.390
Austria	-0.024	0.156	Argentina	-0.108	-0.082
Belgium	0.331	0.371	Bangladesh	0.054	0.090
Canada	0.527	0.498	Brazil	0.275	0.265
Denmark	0.250	0.173	Chile	0.277	0.382
Finland	0.696	0.704	China	0.080	0.059
France	0.488	0.524	Colombia	0.221	0.380
Germany	0.681	0.646	Dominican Republic	0.261	0.335
Greece	0.615	0.707	Egypt	0.675	0.658
Ireland	0.908	0.868	Hong Kong	0.518	0.436
Italy	0.158	0.352	Hungary	0.530	0.570
Japan	0.653	0.604	India	0.133	0.069
Netherlands	0.423	0.231	Indonesia	0.454	0.375
New Zealand	0.628	0.455	Israel	0.380	0.394
Norway	0.495	0.330	Korea, Republic of	0.307	0.274
Portugal	0.647	0.712	Malaysia	-0.044	-0.220
Spain	0.654	0.680	Mexico	0.419	0.357
Sweden	0.562	0.373	Morocco	-0.333	-0.290
Switzerland	0.512	0.354	Pakistan	0.498	0.495
United Kingdom	0.486	0.407	Peru	0.411	0.415
United States	0.931	1.525	Philippines	0.825	0.777
			Poland	0.416	0.479
			Romania	0.757	0.779
			Singapore	0.663	0.659
			South Africa	0.650	0.643
			Thailand	0.550	0.514
			Turkey	0.067	0.008
			Venezuela	0.375	0.469
			Vietnam	0.411	0.397

Table 1
NLS estimates of the income prediction model

	(1)	(2)	(3)
Means of the country-specific estimates			
Lagged GDP growth	0.24	0.23	0.23
	(9.39)	(9.63)	(9.30)
Global GDP growth	--	0.84	0.82
		(12.05)	(12.22)
Lagged global GDP growth	--	--	-0.07
			(1.01)
Wald tests of joint significance (p-values)			
Lagged GDP growth	0.00	0.00	0.00
Global GDP growth	--	0.00	0.00
Lagged global GDP growth	--	--	0.00
Test of cross-sectional dependence (p-value)	0.00	0.09	0.12

The top half of the table reports the means of the country-specific parameter estimates and their robust t-statistics in parentheses. The Wald test statistics in the bottom half of the table are distributed as chi-square with 50 degrees of freedom. The last row in the table reports the test of cross sectional dependence of Pesaran (2004, 2012).

Table 2
NLS estimates of the country-specific risk-sharing parameters

	(1)	(2)
Income forecasting model	System	AR1
Statistics of risk sharing estimates		
Number of risk sharing estimates between 0 and 1	46	46
of which significantly positive	41	41
Number of risk sharing estimates greater than 1	0	1
of which significantly greater than 1	0	0
Number of risk sharing estimates smaller than zero	4	3
of which significantly negative	0	0
Median	0.42	0.40
Maximum	0.93	1.53
Minimum	-0.33	-0.29
Wald test of equality of all estimates (p-value)	0.00	0.00
Average estimates		
All countries	0.42	0.42
	(17.1)	(11.4)
Industrial countries	0.53	0.52
	(20.9)	(11.7)
Developing countries	0.35	0.35
	(9.9)	(8.2)
Wald test of difference in means (p-value)	0.00	0.00
Wald test of joint significance of all risk-sharing parameters	0.00	0.00

Robust t-statistics in brackets, using spatial-correlation consistent standard errors of Driscoll and Kraay (1998). The Wald test statistics at the bottom of the table are distributed as chi-square with 50 degrees of freedom.

Table 3
NLS estimates of the covariates of risk-sharing

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Total foreign assets / GDP	0.121	0.085	0.102	0.085	0.105	0.085	0.105
	(6.10)	(3.33)	(2.67)	(3.34)	(2.67)	(3.33)	(2.67)
Total foreign liabilities / GDP	0.105	0.085	0.100	0.085	0.101	0.083	0.100
	(9.92)	(7.14)	(3.48)	(6.50)	(3.56)	(4.76)	(2.93)
FDI assets / Total assets		0.016	0.016	0.016	0.016	0.016	0.016
		(3.06)	(3.00)	(2.88)	(2.92)	(2.39)	(2.45)
FDI liabilities / Total liabilities		-0.028	-0.031	-0.028	-0.032	-0.028	-0.032
		(-4.88)	(-4.11)	(-4.37)	(-3.62)	(-4.79)	(-3.62)
Trade openness (imports + exports / GDP)			-0.033		-0.038		-0.037
			(-0.63)		(-0.69)		(-0.68)
Private sector credit / GDP				0.000	-0.010		-0.011
				(-0.01)	(-0.23)		(-0.24)
GDP per capita						0.003	0.002
						(0.12)	(0.06)
Wald test of joint significance	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Implied risk-sharing coefficients							
Median	0.440	0.477	0.478	0.477	0.475	0.479	0.475
Maximum	0.903	0.825	0.836	0.825	0.832	0.823	0.832
Minimum	0.097	0.102	0.103	0.102	0.104	0.101	0.103
Mean - All countries	0.468	0.483	0.483	0.483	0.483	0.483	0.475
Industrial countries	0.543	0.585	0.589	0.585	0.588	0.586	0.569
Developing countries	0.414	0.410	0.407	0.410	0.406	0.409	0.382
Correlation with unrestricted risk-sharing coefficients	0.306	0.471	0.477	0.471	0.476	0.471	0.476

Robust t-statistics in brackets, using spatial-correlation consistent standard errors of Driscoll and Kraay (1998). All the explanatory variables are expressed in logs. All specifications include also a constant not reported here. The last row in the table reports the correlation coefficient between the individual-country risk-sharing parameters predicted by the regression in each column with those reported in the first column of Table B2.

Figure 1. Average coefficient of partial risk sharing across countries

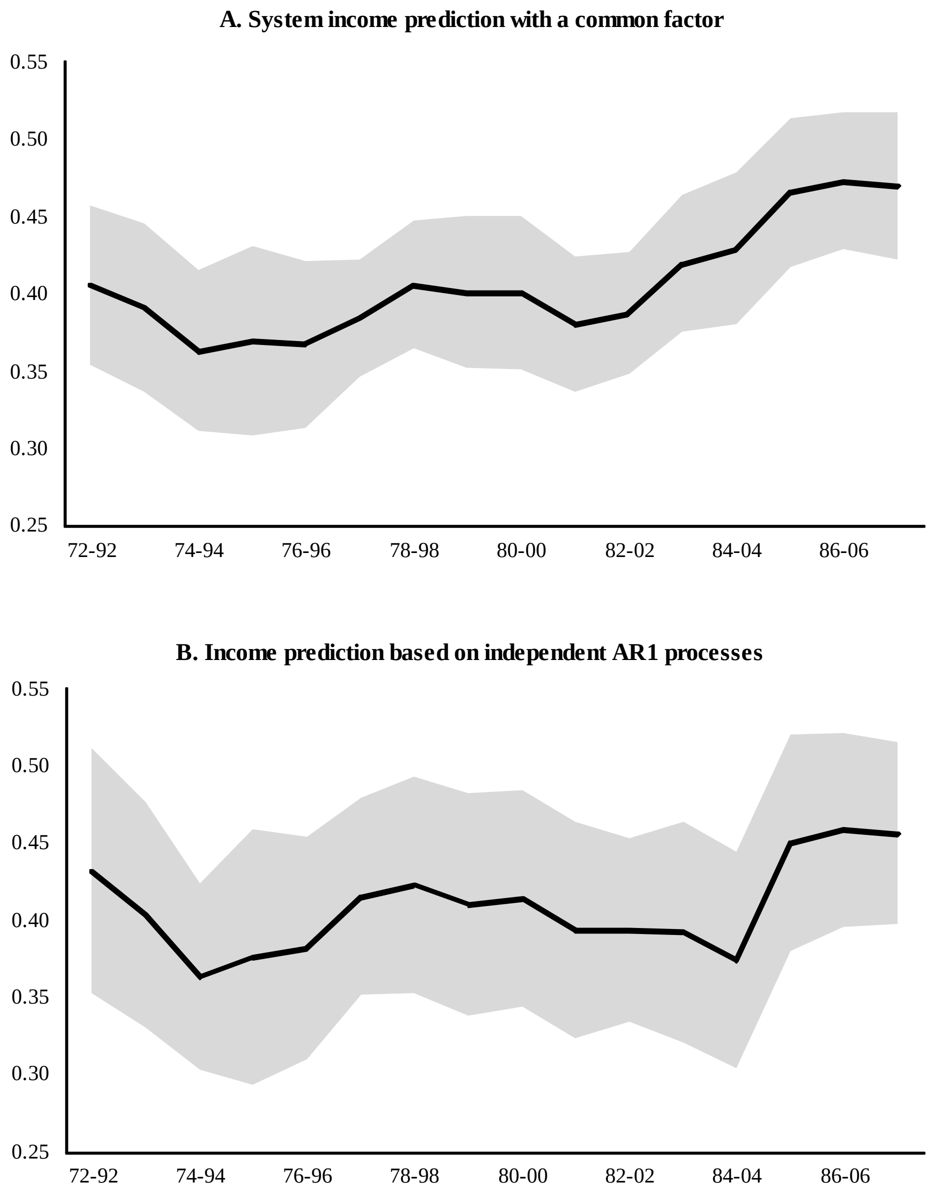

Solid lines represent average point estimate of the coefficients of partial risk sharing. Shaded areas represent two standard errors around the average estimates.

Figure 2. Coefficients of partial risk sharing by income group

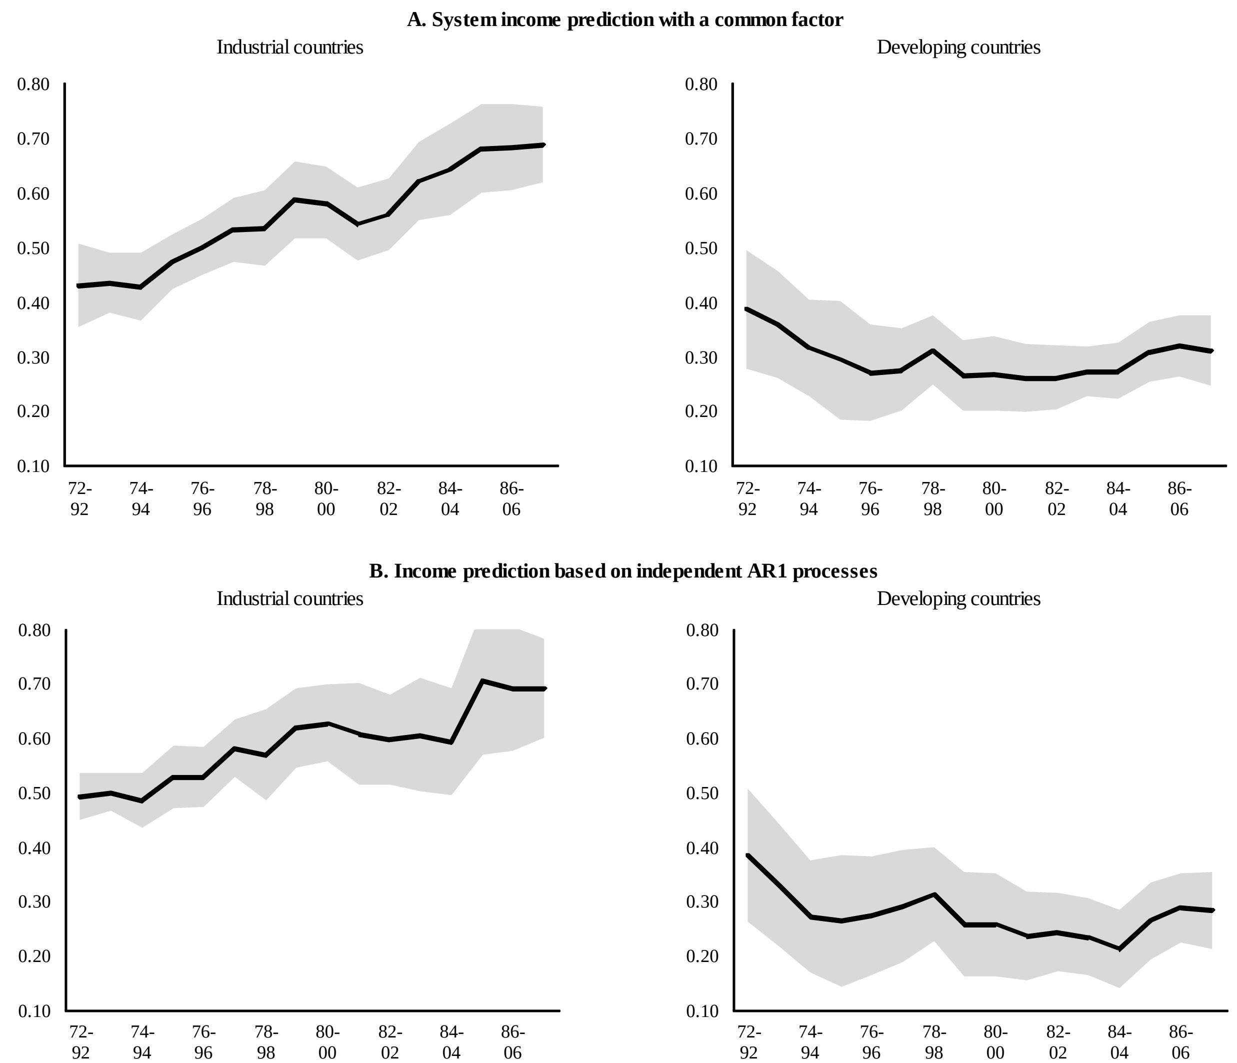

Solid lines represent average point estimate of the coefficients of partial risk sharing. Shaded areas represent two standard errors around the average estimates.