WPS7867

Policy Research Working Paper 7867

Assessing the Degree of International Consumption Risk Sharing

Constantino Hevia
Luis Servén

World Bank Group
Development Research Group
Macroeconomics and Growth Team
October 2016

Policy Research Working Paper 7867

Abstract

This paper examines the extent of consumption risk sharing for a group of 50 high-income 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. The estimation results show that high-income countries exhibit higher degrees of risk sharing than developing countries, and that the gap between the two 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 more open capital accounts, and larger stocks of foreign assets and liabilities exhibit larger degrees of risk sharing. Yet, larger countries in terms of gross domestic product show lower 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 lserven@worldbank.org.

Assessing the Degree of International Consumption Risk Sharing *

Constantino Hevia
Universidad Torcuato di Tella

Luis Servén
World Bank

*We thank Aart Kraay, Alex Werner and seminar participants at IMF and UAM for useful discussions, and Vivian Norambuena for excellent research assistance. This work was supported by the World Bank's Knowledge Change Program.

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 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 argument is straightforward: if financial markets are complete – in the sense that the available assets span the set of all idiosyncratic risks – the ratio of the marginal utilities of consumption of any pair of agents (or countries, in this paper's representative-agent framework) must be constant across dates and states of nature. That is, the economy features perfect consumption 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 allow enhanced risk diversification, although it might also raise the overall exposure to risk.[2] To the extent that the global increase in international financial openness over the last quarter century (Lane and Milesi-Ferretti, 2007) reflects such kind of innovation, it should be associated with a rise in 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, imposing constant absolute or relative 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. Obstfeld (1994), Canova and Ravn (1996), and Lewis (1996) are some leading examples of this literature applied to cross-country aggregate data.[3]

These conventional tests can be informative about whether perfect risk sharing holds empirically. But if the null hypothesis is rejected – as is almost invariably the case – in

[1] 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.

[2] The reason is that financial integration may also open the door to the propagation of financial turbulence from abroad. For example, Devereux and Yu (2016) find that financial integration reduces the severity of crises as a result of better diversification, but makes them more frequent due to increased opportunities for contagion. As a consequence, the overall degree of risk faced by consumers may rise or fall relative to a situation of financial autarky.

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

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 more precisely how 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 their 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 Hoffmann (2008) to obtain estimates of the extent of partial risk sharing. This setting, however, 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 may engage in different extents of risk sharing.

This is precisely the focus of this paper. It analyzes to what extent different countries diversify their idiosyncratic risks, and relates their respective degree of risk sharing to selected country characteristics, in particular regarding their degree of international financial integration. 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, 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 intuitive, it has some important consequences. On the one hand, it complicates the solution of the model and its empirical implementation. Specifically, Crucini's 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 expected evolution of the future income of all agents. Moreover, the coefficients multiplying these variables 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 the 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.[4] In contrast, our model does not suffer from this problem because our estimating equation is stated only in terms of the (quasi-)innovations to the present value of income growth of all agents. Unlike conventional risk-sharing regressions, it does not include the cross-sectional average of the dependent variable among the regressors.[5]

Another difference between our framework and that employed by earlier literature concerns the decision problem that agents solve after the cross-country transfers are made. While the traditional literature assumes that agents act as permanent-income consumers with quadratic preferences, we assume that agents maximize a utility function of the constant relative risk aversion (CRRA) form allowing for country-specific coefficients of risk aversion and interest rates, along the lines of Blundell et al. (2008). We follow Campbell and Mankiw (1989) and derive an estimating equation log-linearizing the Euler equation and intertemporal budget constraint of the representative household's problem.

We estimate the model using a time-series cross-section dataset comprising 50 industrialized and developing countries over the period 1970-2010. We find an average coefficient of partial risk sharing of around 0.5. In terms of the model, this means that the average country contributes half its income to the income pool. The estimates, however, vary systematically between industrial and developing countries, with the former exhibiting, on average, higher degrees of risk sharing than the latter (centered around 0.59 and 0.43, respectively). Re-estimating the model over rolling time samples, we find evidence that the degree of risk sharing has been on the rise for the average country during most of the sample period. When we split the sample between industrial and developing countries,

[4]Crucini's paper reports a Monte Carlo experiment to check if this procedure leads to major biases in the risk-sharing estimate.

[5]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.

however, we find a widening gap between industrial and developing countries: while the degree of risk sharing increased over time for industrial countries, average risk sharing for developing countries remained flat or even declined somewhat during the globalization period (although there is an 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 countries with relatively low levels of financial and/or institutional development, as is the case of many developing economies.[6]

In the paper's framework, these partial risk-sharing coefficient estimates can be interpreted as measuring agents' chosen contributions to the global income pool. The model, however, is silent on the factors behind those choices. To gain further insight, we assume that they reflect cross-country variation in policies, institutions, and other characteristics that help or hinder consumption risk sharing. Specifically, we re-estimate the model expressing the partial risk-sharing coefficients 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. We also include a measure of country size as determinant of the coefficient of risk sharing, because larger countries should expect smaller benefits from participating in the risk pooling agreement. In particular, for a given contribution to the income pool, the larger is the country the more correlated will be its output with the pooled income.

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 a measure of de jure capital account openness, and/or by a de facto measure of total foreign asset and liability positions, is positively correlated with the coefficient of partial risk sharing, consistent with the view that financial integration improves international risk sharing. After controlling for financial integration, however, trade openness appears to affect negatively risk sharing. While we do not pursue this idea further, one potential explanation is that the impact of global terms of trade shocks on income volatility is amplified in more open economies, which – for a given degree of financial integration – weakens consumption risk sharing. In addition, the degree of domestic financial development, measured as the ratio of domestic private credit to GDP, is positively correlated with the degree of risk sharing. Finally, the size of the economy, as measured by real GDP, is robustly negatively associated with the degree of risk sharing, consistent with the observation that larger countries reap smaller benefits from participating in the risk-sharing agreement.

[6] On the costs and benefits of financial globalization, see Kose et al. (2009a).

Our paper is related to a substantial empirical literature assessing international consumption risk sharing. Sørensen 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 et al. (2009b) 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 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 et al. (2012) 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.[7]

The paper closest to ours is Ho and Ho (2015) which, in independent work, developed a model with heterogeneous contributions to the income pool. There are, however, major differences between their framework and conclusions and ours. First, we are interested in comparing the differential risk sharing in the group of industrial countries relative to that in developing countries. Ho and Ho focus only on rich countries. Second, we use a different risk-sharing rule that implies that larger countries have less to gain from participating in the income pool, a natural assumption to make. Third, Ho and Ho use the standard permanent-income model with quadratic preferences, while we consider CRRA preferences and allow for heterogeneity across countries. And fourth, while Ho and Ho estimate the model imposing that output growth is iid across countries and over time, we allow for more flexible income processes with temporal dependence within countries as

[7]On the theory side, 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. Finally, Bengui et al. (2013) analyze whether the international globalization of financial markets is associated with more risk sharing using a calibrated two-country model with different frictions.

well as cross-sectional dependence through common factors that may affect income growth of all countries simultaneously. This specification for the income process is supported by the data, and it matters for the paper's results. We discuss other differences in the body of the paper.

The rest of the paper is organized as follows. Section 2 develops the model of partial risk sharing. 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 appendix contains some proofs and additional tables.

2 A model of partial risk sharing

In this section we develop a model of partial risk sharing that generalizes the risk-sharing agreement in Crucini (1999). Time is discrete and denoted by $t = 0, 1, 2, ..., \infty$. There are N different locations which we label countries. Country $i = 1, 2, ..., N$ is composed of H_i identical consumers each of whom owns a tree which yields $Y_{i,t}$ fruits at time t. Thus, aggregate income in country i is $H_i Y_{i,t}$.

Consumers enter into a partial risk-sharing agreement with the consumers of other countries at the beginning of time 0, before income is realized. The agreement requires that all agents contribute a fraction of their income, possibly different across countries, into an income pool in exchange for a claim to a fraction of the aggregate pooled income. If consumers of country i contribute a fraction λ_i of their income to the pool, the contribution of country i to the fund at time t is $\lambda_i H_i Y_{i,t}$ and the aggregate pooled income equals $\sum_{j=1}^{N} \lambda_j H_j Y_{j,t}$.

We assume that the agreement specifies that consumers in country i get back a fraction θ_i of the income pool equal to their share in the initial contribution,

$$\theta_i = \frac{\lambda_i H_i Y_{i,0}}{\sum_{j=1}^{N} \lambda_j H_j Y_{j,0}} \tag{1}$$

Therefore, the aggregate income of country i after the risk-sharing agreement is

$$H_i \bar{Y}_{i,t} = (1 - \lambda_i) H_i Y_{i,t} + \theta_i \sum_{j=1}^{N} \lambda_j H_j Y_{j,t}.$$

In per capita terms, each consumer of country i receives

$$\bar{Y}_{i,t} = (1-\lambda_i)Y_{i,t} + \lambda_i \sum_{j=1}^{N} \frac{\omega_{i,0}}{\omega_{j,0}} \theta_j Y_{j,t}, \tag{2}$$

where $\omega_{i,0} = Y_{i,0} / \sum_{k=1}^{N} Y_{k,0}$ is the share of country i's per capita income in aggregate per capita income across countries at time 0.[8]

Suppose now that, on top of the risk-sharing agreement, consumers in country i can invest their wealth in an asset with a gross real return (or interest rate) $R_{i,t+1}$. Following Campbell and Mankiw (1989), we assume that the income flows after the risk-sharing agreement are capitalized into tradable wealth $W_{i,t}$. Thus, the flow budget constraint is

$$W_{i,t+1} = R_{i,t+1}(W_{i,t} - C_{i,t}). \tag{3}$$

The preferences of consumers in country i are represented by the utility function

$$E_t \left[\sum_{h=0}^{\infty} e^{-\delta_i h} \frac{C_{i,t+h}^{1-1/\gamma_i}}{1-1/\gamma_i} \right], \tag{4}$$

where γ_i is the reciprocal of the coefficient of relative risk aversion, δ_i is the subjective discount rate, and E_t is the expectation operator conditional on information at time t.

The Euler equation associated with the problem of maximizing (4) subject to (3) is

$$C_{i,t}^{-1/\gamma_i} = E_t \left[e^{-\delta_i} R_{i,t+1} C_{i,t+1}^{-1/\gamma_i} \right].$$

Following Campbell and Mankiw (1989), we log-linearize the Euler equation and the budget constraint to obtain a function relating log-consumption to the expected path of future returns and log-income after the risk-sharing agreement,

$$c_{i,t} - \bar{y}_{i,t} = \sum_{s=1}^{\infty} \rho_i^s E_t[\Delta \bar{y}_{i,t+s}] - \gamma_i \sum_{s=1}^{\infty} \rho_i^s E_t[r_{i,t+s}] + \frac{\rho_i \gamma_i \delta_i}{1-\rho_i}, \tag{5}$$

[8] Ho and Ho (2015) use the sharing rule $\theta_i = \lambda_i / \sum_{j=1}^{N} \lambda_j$. In terms of our notation, this rule implicitly assumes that average per-capita income and total population is the same across countries. It seems natural to assume that, to contribute a given fraction of their income to the pool, wealthier agents will demand a higher fraction of the pooled income than poorer agents. For otherwise poorer consumers will receive on average a transfer from the wealthier ones due to the differences in their average incomes. Unless we use a sharing rule similar to (1), poorer and less populated countries will receive net transfers from wealthier and larger countries for reasons that have nothing to do with the uncertainty of their income process.

where lowercase letters denote the log of uppercase letters and $\rho_i = 1 - \bar{C}_i/\bar{W}_i$ is one minus the average consumption-wealth ratio.[9]

The consumption function (5) imposes a cointegration relation between log-consumption $c_{i,t}$ and log-income after risk sharing $\bar{y}_{i,t}$. Although in principle we could use this expression to estimate the model, in practice saving rates display considerable persistence, and appear close to being nonstationary.[10] Since our model is not designed to account for persistent movements in the saving rate, we follow earlier literature and estimate the model in differences.[11]

Differentiating expression (5) between times t and $t-1$ gives

$$\Delta c_{i,t} = \gamma_i r_{i,t} + \sum_{s=0}^{\infty} \rho_i^s \left(E_t - \rho_i E_{t-1}\right)\left(\Delta \bar{y}_{i,t+s} - \gamma_i r_{i,t+s}\right), \tag{6}$$

where, for any variable z_{t+s}, the expression

$$(E_t - \rho_i E_{t-1})(z_{t+s}) = E_t(z_{t+s}) - \rho_i E_{t-1}(z_{t+s})$$

represents the quasi-innovation in z_{t+s} as of time t.[12]

In equation (6), the growth rate of consumption is expressed as a function of the growth rates of current and future incomes after risk sharing, and the real interest rate. We next derive an expression relating consumption to the growth rate of the income of all countries before risk sharing. In the appendix we show that a log-linear approximation of equation (2) around $Y_{i,t}/Y_{k,t} = Y_{i,0}/Y_{k,0}$ implies the following relation between income growth after and before risk sharing,

$$\Delta \bar{y}_{i,t} = (1 - \lambda_i)\Delta y_{i,t} + \lambda_i \sum_{j=1}^{N} \theta_j \Delta y_{j,t}, \tag{7}$$

[9] The online appendix contains a detailed derivation of all the results discussed in this section.

[10] For example, China's gross saving rate has kept a persistent upward trend, from 35 percentage points of GDP at the beginning of the 1980s to over 50 percent by 2010. At the other end, the gross saving rate in Greece declined steadily from about 40 percent in 1973 to less than 9 percentage points in 2010. More broadly, formal unit root tests reported by Byrne et al. (2009), for industrial countries, and Coakley et al. (1999), for developing countries, fail to reject nonstationarity of the saving rate for the majority of the countries in their respective samples.

[11] Estimation in differences is commonly used in the permanent-income literature and in models of partial consumption insurance along the lines of Crucini (1999). See Campbell and Mankiw (1990), Artis and Hoffmann (2008), Asdrubali and Kim (2008), and Ho and Ho (2015), among others.

[12] The innovation $(E_t - E_{t-1})z_{t+s}$ is the new information obtained about variable z_{t+s} between times t and $t-1$. We label $(E_t - \rho_i E_{t-1})z_{t+s}$ the "quasi-innovation" in z_{t+s} because the constant ρ_i is typically close to but smaller than 1.

If $\lambda_i = 0$ (when country i does not contribute at all to the income pool) income growth after risk sharing trivially equals income growth before risk sharing. In contrast, if $\lambda_i = 1$ (when country i contributes all its income to the pool) income growth after risk sharing equals a weighted average of the income growth before risk sharing across countries, where the weights are given by the relative contribution of each country to the income pool.

Replacing equation (7) into (6) delivers an expression relating consumption growth in country i to the quasi-innovation to the discounted value of income growth of all countries and the quasi-innovation to the discounted value of returns,

$$\begin{aligned}\Delta c_{i,t} =& \gamma_i r_{i,t} - \gamma_i (E_t - \rho_i E_{t-1}) \sum_{s=0}^{\infty} \rho_i^s r_{i,t+s} + (1-\lambda_i)(E_t - \rho_i E_{t-1}) \sum_{s=0}^{\infty} \rho_i^s \Delta y_{i,t+s} \\ &+ \lambda_i \sum_{j=1}^{N} \theta_j (E_t - \rho_i E_{t-1}) \sum_{s=0}^{\infty} \rho_i^s \Delta y_{i,t+s}. \end{aligned} \tag{8}$$

Unfortunately we do not have data on returns for the 50 countries and for the period that we consider. We thus assume that the current interest rate and quasi-innovations to the present value of returns are orthogonal to the quasi-innovation to the present value of income growth. Although restrictive, this assumption is much weaker than that used in most of the literature on the permanent income model and the literature on partial risk sharing that followed the seminal paper by Crucini (1999). That literature simply assumes that interest rates are constant and equal across countries.[13] In contrast, our approach allows for interest rates that are random and different across countries, although we need the orthogonality assumption mentioned above to rule out problems of endogeneity.

To estimate the model we also impose that ρ_i is constant across countries. To calibrate ρ, rewrite the budget constraint (3) as

$$\left(\frac{W_{i,t+1}/C_{i,t+1}}{W_{i,t}/C_{i,t}}\right)\frac{C_{i,t+1}}{C_{i,t}} = R_{i,t+1}\left(1 - \frac{C_{i,t}}{W_{i,t}}\right)$$

In a balanced growth path, with $W_{i,t+1}/C_{i,t+1} = W_{i,t}/C_{i,t} = W/C$, $R_{i,t+1} = R$, and $C_{i,t+1}/C_{i,t} = 1 + g$, the previous expression collapses to

$$1 + g = R\left(1 - \frac{C}{W}\right)$$

[13] Crucini (1999), Crucini and Hess (2000), Artis and Hoffmann (2008), Asdrubali and Kim (2008), and Ho and Ho (2015) all assume that interest rates are constant and equal across countries.

But noting that $\rho = 1 - C/W$ gives

$$\rho = \frac{1+g}{R}.$$

Using $g = 0.023$, which is the average pooled growth rate of consumption per capita in our sample period, and an average interest rate of 5% , $R = 1.05$, gives $\rho \approx 0.97$. Our results, however, are robust to reasonable variations in the calibrated value of ρ.

It is convenient to rewrite equation (8) compactly as

$$\Delta c_{i,t} = \mu_i + (1-\lambda_i)\epsilon_{i,t} + \lambda_i \sum_{j=1}^{N} \theta_j \epsilon_{j,t} + \nu_{i,t}, \tag{9}$$

where

$$\epsilon_{j,t} = (E_t - \rho E_{t-1}) \sum_{s=0}^{\infty} \rho^s \Delta y_{j,t+s}$$

is the quasi-innovation to the present value of country j's income growth rate, μ_i is the unconditional mean of $\gamma_i r_{i,t} - \gamma_i (E_t - \rho E_{t-1}) \sum_{s=0}^{\infty} \rho^s r_{i,t+s}$, and $\nu_{i,t}$ is defined as

$$\nu_{i,t} = \gamma_i r_{i,t} - \gamma_i (E_t - \rho E_{t-1}) \sum_{s=0}^{\infty} \rho^s r_{i,t+s} - \mu_i.$$

Our objective is to estimate the degree of risk sharing λ_i for each country $i = 1, 2, ..., N$. This requires estimating the quasi-innovations to the present value of country j's income growth $\epsilon_{j,t}$. Let $\hat{\epsilon}_{j,t}$ denote the estimated quasi-innovation and rewrite (9) as

$$\Delta c_{i,t} = \mu_i + (1-\lambda_i)\hat{\epsilon}_{i,t} + \lambda_i \sum_{j=1}^{N} \theta_j \hat{\epsilon}_{j,t} + u_{i,t}, \tag{10}$$

where the residual $u_{i,t}$ is the sum of $\nu_{i,t}$ and the rational expectation's error that represents the superior information of agents relative to the econometrician,

$$u_{i,t} = \nu_{i,t} + (1-\lambda_i)\epsilon_{i,t} + \lambda_i \sum_{j=1}^{N} \theta_j \epsilon_{j,t} - \left((1-\lambda_i)\hat{\epsilon}_{i,t} + \lambda_i \sum_{j=1}^{N} \theta_j \hat{\epsilon}_{j,t} \right).$$

It is easy to verify that, under our assumptions, $u_{i,t}$ is orthogonal to the econometrician's information set at time $t-1$.[14] This information set, however, may omit some relevant

[14]We follow the usual procedure in the rational expectations econometric literature and assume that, when performing revisions to their anticipated output paths, agents have more information than the econometrician. The proof that $u_{i,t}$ is orthogonal to the econometrician's information set at $t-1$ uses the law of iterated

information used by the agents to forecast their future incomes. The 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 the consumer, but unobserved by the econometrician, to perform revisions to his present value of output growth. The latter could be due to an aggregate shock observed by all agents, but unobserved by the econometrician, that affects everyone's output growth.

To sum up, the key differences between our framework and that in Crucini (1999) can be summarized as follows. In both cases, the starting point is the assumption that agents contribute a fraction of their income to an income pool and have access to an asset to smooth consumption over time 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.[15] 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 because we will work with a large country sample comprising both advanced and developing countries, and hence exhibiting a greater degree of heterogeneity. Indeed, 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 (10), in which different countries may exhibit different coefficients of partial risk sharing. A second difference between the two frameworks is that, rather than assuming quadratic preferences (as in the permanent income literature), we obtain our estimating equation by log-linearizing the Euler equation of a consumer with preferences displaying constant relative risk aversion.

The added generality of our risk-sharing agreement relative to the conventional model used in the literature poses additional econometric challenges too. For each country, equation (10) involves a nonlinear function of the risk-sharing parameters of all countries. Thus, empirical implementation of the model requires more complex econometric techniques than those found in earlier literature. This is the topic of the next section.

expectations as in Hansen and Sargent (1980). In addition, any classical measurement error in consumption can also be included in the error term $u_{i,t}$.

[15] This assumption is maintained in subsequent empirical applications of Crucini's framework (with the exception Ho and Ho (2015)).

3 Empirical implementation

The core of our empirical analysis is the estimation of (10) using a cross-country time-series dataset. However, the equation involves the quasi-innovations to the present value of income growth, which are not observed. We proceed on two stages. First, we construct the innovations based on an estimated time-series model of income growth. In the second stage, we use these innovations to estimate (10).[16] In this section we first discuss the econometric issues surrounding the estimation of the parameters of the risk-sharing model, and then we turn to those related to the construction of the quasi-innovations required for such task. Finally, we briefly summarize our data sources.

3.1 Implementing the risk-sharing model

Estimation of (10) in a cross-country setting 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 (10) as

$$\Delta c_{i,t} = \mu_i + (1-\lambda_i)\hat{\epsilon}_{i,t} + \lambda_i \sum_{j=1}^{N} \frac{\lambda_j H_j \omega_{j,0}}{\sum_{k=1}^{N} \lambda_k H_k \omega_{k,0}} \hat{\epsilon}_{j,t} + u_{i,t}, \tag{11}$$

where $u_{i,t}$ is potentially heteroskedastic and possibly correlated across countries and over time.

To express the empirical model in more compact form, we introduce the following notation. Let $\Delta \mathbf{c_i} = (\Delta c_{i,1}, ..., \Delta c_{i,T})'$ and $\boldsymbol{\epsilon}_i = (\epsilon_{i,1}, ..., \epsilon_{i,T})'$ denote the (column) vectors of consumption growth and quasi-innovations to the present discounted value of income growth of country i, where T denotes the number of time series observations per country, and define the $T \times N$ matrix $\mathbf{Z} = (\boldsymbol{\epsilon}_1, ..., \boldsymbol{\epsilon}_N)$. In addition, let $\boldsymbol{\lambda} = (\lambda_1, ..., \lambda_N)'$ denote the vector of risk-sharing parameters for all countries, and define the $N \times N$ diagonal matrices $\mathbf{H}$ and $\boldsymbol{W_0}$ which have along their main diagonal the population of each country H_j and its per capita GDP share $\omega_{j,0}$, respectively, with all off-diagonal elements equal to zero. After some straightforward manipulations, the T observations of (11) corresponding to the i-th

[16] In principle, one could combine both steps into a single one, and estimate jointly the parameters of the income process and the risk-sharing model. The estimation problem would be nonlinear due to the presence of the permanent income innovations, even if the risk-sharing model were linear in its parameters. However, the risk-sharing model is itself heavily nonlinear, and this would greatly add to the complexity of the joint estimation approach. For this reason, we proceed in two stages, as described in the text.

country can be compactly written

$$\Delta \mathbf{c}_i = \boldsymbol{\iota}_T \mu_i + \mathbf{Z}\left(\boldsymbol{e}_i \boldsymbol{e}_i'(\boldsymbol{\iota}_N - \boldsymbol{\lambda}) + \mathbf{H}\boldsymbol{W}_\mathbf{0}\boldsymbol{\lambda}(\boldsymbol{\iota}_N'\mathbf{H}\boldsymbol{W}_\mathbf{0}\boldsymbol{\lambda})^{-1}\boldsymbol{e}_i'\boldsymbol{\lambda}\right) + \mathbf{u}_i, \tag{12}$$

where $\boldsymbol{\iota}_N$ denotes an $N \times 1$ column vector of ones and $\boldsymbol{e}_i$ is an $N \times 1$ vector of zeros with a 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

$$\Delta \mathbf{c} = \boldsymbol{\mu} \otimes \boldsymbol{\iota}_T + (\mathbf{I}_N \otimes \mathbf{Z})\left[\left(\boldsymbol{\lambda} \otimes \mathbf{H}\boldsymbol{W}_\mathbf{0}\boldsymbol{\lambda}(\boldsymbol{\iota}_N'\mathbf{H}\boldsymbol{W}_\mathbf{0}\boldsymbol{\lambda})^{-1}\right) + \mathsf{vec}\left(\mathbf{I}_N - \mathsf{diag}(\boldsymbol{\lambda})\right)\right] + \mathbf{u}, \tag{13}$$

where $\otimes$ denotes the Kronecker product, and $\mathsf{diag}(\boldsymbol{\lambda})$ 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 quasi-innovations to the present value of income growth in all N countries, contained in the matrix $\mathbf{Z}$) are the same in all equations. In this context, we use system NLS to estimate the parameters of (13). 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

$$\min_{\boldsymbol{\lambda}} \left(\Delta \tilde{\mathbf{c}} - \boldsymbol{\phi}(\tilde{\mathbf{Z}}, \mathbf{H}, \boldsymbol{W}_\mathbf{0}, \boldsymbol{\lambda})\right)' \left(\Delta \tilde{\mathbf{c}} - \boldsymbol{\phi}(\tilde{\mathbf{Z}}, \mathbf{H}, \boldsymbol{W}_\mathbf{0}, \boldsymbol{\lambda})\right), \tag{14}$$

where $\boldsymbol{\phi}(\tilde{\mathbf{Z}}, \mathbf{H}, \boldsymbol{W}_\mathbf{0}, \boldsymbol{\lambda}) = (\mathbf{I}_N \otimes \tilde{\mathbf{Z}})\left[\left(\boldsymbol{\lambda} \otimes \mathbf{H}\boldsymbol{W}_\mathbf{0}\boldsymbol{\lambda}(\boldsymbol{\iota}_N'\mathbf{H}\boldsymbol{W}_\mathbf{0}\boldsymbol{\lambda})^{-1}\right) + \mathsf{vec}\left(\mathbf{I}_N - \mathsf{diag}(\boldsymbol{\lambda})\right)\right]$. 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).[17]

This procedure yields a set of unrestricted estimates of the risk-sharing coefficients $\boldsymbol{\lambda}$. However, we are also interested in learning about the covariates of risk sharing. To do this, we restrict the risk-sharing coefficients to be a function $\boldsymbol{\lambda} = \mathbf{m}(\mathbf{X}\boldsymbol{\delta})$, 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 (13), estimation proceeds along the same lines as

[17] In the formula of the covariance matrix, the usual matrix of regressors is replaced in our case by the matrix of partial derivatives $\nabla_{\boldsymbol{\lambda}}\boldsymbol{\phi}$.

above, with the parameter vector now given by δ. That is, we solve

$$\min_{\delta} \left(\Delta\tilde{\mathbf{c}} - \boldsymbol{\phi}(\tilde{\mathbf{Z}}, \mathbf{H}, \mathbf{W_0}, \mathbf{m}(\mathbf{X}\delta))\right)' \left(\Delta\tilde{\mathbf{c}} - \boldsymbol{\phi}(\tilde{\mathbf{Z}}, \mathbf{H}, \mathbf{W_0}, \mathbf{m}(\mathbf{X}\delta))\right), \tag{15}$$

Like in the unrestricted case, we use the Driscoll and Kraay (1998) robust covariance matrix estimator to conduct inference on δ.

3.2 Income prediction

Empirical implementation of the risk-sharing model requires suitable forecasts of the future income growth. To construct them, we estimate a simple time-series model of per capita GDP growth. In order to allow for correlated shocks affecting growth in multiple countries, we use a factor approach. Specifically, we assume that per capita GDP growth can be expressed as a dynamic factor model:

$$\Delta y_{i,t} = \alpha_{0,i} + (\alpha_{i,1}\Delta y_{i,t-1} + ... + \alpha_{i,p}\Delta y_{i,t-p}) + (\boldsymbol{\beta}_{i,0}\mathbf{f}_t + ... + \boldsymbol{\beta}_{i,s}\mathbf{f}_{t-s}) + \nu_{i,t},$$

for $i = 1, ..., N$ and $t = 1, ..., T$. Here, $\mathbf{f}_t$ is a $q \times 1$ vector of (unobserved) common factors and we assume that all parameters can vary freely across i. Stacking for given t, we have

$$\Delta\mathbf{y}_t = \boldsymbol{\alpha}_0 + \mathbf{A}(L)\Delta\mathbf{y}_{t-1} + \boldsymbol{\beta}(L)\mathbf{f}_t + \boldsymbol{\nu}_t, \tag{16}$$

for $t = 1, ..., T$, where L represents the lag operator, $\mathbf{A}(L)$ is block-diagonal, and the covariance matrix of $\boldsymbol{\nu}_t$ is assumed diagonal. This specification nests three cases of interest. First, if $\mathbf{A}(L) = \mathbf{0}$ and $\boldsymbol{\beta}(L) = \mathbf{0}$, for each i per capita GDP growth is iid (as in Ho and Ho (2015)); equivalently, per capita income is a random walk with drift. Second, if $\boldsymbol{\beta}(L) = \mathbf{0}$ but $\mathbf{A}(L) \neq \mathbf{0}$, growth follows an independent AR process in each country (Crucini, 1999). Lastly, if $\mathbf{A}(L) \neq \mathbf{0}$ and $\boldsymbol{\beta}(L) \neq \mathbf{0}$ we have a factor model, dynamic if $s > 0$.

To make this setting operational, assume that the unobserved factors follow the AR(h) process $\mathbf{f}_t = \mathbf{d}(L)\mathbf{f}_{t-1} + \boldsymbol{\eta}_t$. Letting $\mathbf{F}_t = (\mathbf{f}'_t, \mathbf{f}'_{t-1}, ...\mathbf{f}'_{t-s+1})'$, the model can be rewritten (ignoring $\boldsymbol{\alpha}_0$ for simplicity) as a static factor model:

$$\Delta\mathbf{y}_t = \mathbf{A}(L)\Delta\mathbf{y}_{t-1} + \mathbf{B}\mathbf{F}_t + \boldsymbol{\nu}_t, \tag{17}$$

$$\mathbf{F}_t = \mathbf{D}(L)\mathbf{F}_{t-1} + \mathbf{G}\boldsymbol{\eta}_t, \tag{18}$$

or, equivalently, as a factor-augmented VAR with coefficient restrictions

$$\begin{pmatrix} \mathbf{F}_t \\ \Delta \mathbf{y}_t \end{pmatrix} = \begin{pmatrix} \mathbf{D}(L) & \mathbf{0} \\ \mathbf{B}\mathbf{D}(L) & \mathbf{A}(L) \end{pmatrix} \begin{pmatrix} \mathbf{F}_{t-1} \\ \Delta \mathbf{y}_{t-1} \end{pmatrix} + \begin{pmatrix} \mathbf{G}\boldsymbol{\eta}_t \\ \mathbf{B}\mathbf{G}\boldsymbol{\eta}_t + \boldsymbol{\nu}_t \end{pmatrix}. \tag{19}$$

For the general case $\mathbf{A}(L) \neq \mathbf{0}$ and $\mathbf{B} \neq \mathbf{0}$, estimation proceeds in two steps (see Stock and Watson, 2005; Song, 2013). We first estimate $\mathbf{A}(L)$, $\mathbf{B}$, and $\mathbf{F}_t$ doing OLS iteratively on (17). Specifically, we start with an initial guess for $\mathbf{F}_t$ and update it with the first r principal components of $\left(\Delta \mathbf{y}_t - \hat{\mathbf{A}}(L)\Delta \mathbf{y}_{t-1} - \hat{\mathbf{B}}\hat{\mathbf{F}}_t\right)$. We use the information criteria of Bai and Ng (2002) and Choi (2013) to determine the number of static factors r (with $r = q(s+1)$). To perform inference on $\mathbf{A}(L)$ and $\mathbf{B}$, we use the results of Bai and Ng (2006, 2013).

At the second step, we estimate $\mathbf{D}(L)$ in (18) by OLS using the estimated $\mathbf{F}_t$ from the previous step. We use standard information criteria to determine the order of $\mathbf{D}(L)$ (generally given by $\max(1, h-s)$; see Bai and Ng, 2007).

Using these estimates, the quasi-innovations to the present value of income growth can be constructed in a straightforward manner using income growth forecasts computed recursively with (19).

3.3 Data

The empirical sample is dictated by data availability. We work with the 50 largest countries in the world in terms of their respective GDP (in U.S. dollars in the year 2005) for which complete annual data on consumption and income, as well as measures of de jure and de facto financial openness (as described below) could be assembled. Table B.1 in the appendix lists the countries included in the analysis. Taken together, they account for almost 90 percent of world GDP in the year 2005.

Data on real GDP growth, aggregate consumption growth, and total population are taken from the Penn World Tables 7.1 (PWT 7.1). Along the time dimension, our sample runs from 1970 to 2010. We use the 1970-71 averages to construct the initial shares $\boldsymbol{W}_0$ and population $\mathbf{H}$.[18] The regression sample therefore runs from 1972 to 2010, and comprises 1,950 country-year observations. As explained in Section 2, for the calculation of the quasi-innovation to the discounted present value of income growth, we set $\rho = 0.97$.

For our analysis of the covariates of risk sharing, we measure countries' international financial integration using data on financial openness drawn from the Chinn-Ito dataset, updated to 2012, as well as international asset and liability positions from the Lane and

[18] Estimation results change very little if we use instead the averages over 1970-74 or the year 1970 only.

Milesi-Ferretti dataset updated up to 2011.[19] In addition, we also consider trade openness (as measured by imports plus exports divided by GDP, taken from PWT 7.1) 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, taken from the World Development Indicators.

Figure 1 illustrates the trends in de jure international financial openness in the sample countries, as captured by the cross-country average of the Chinn-Ito openness measure. The figure shows both the overall mean as well as the means for advanced and developing countries separately. All three exhibit a rising trend, although with varying timing and magnitude.[20] The advanced-country mean rises steadily since the late 1970s, from under 0.50 to almost the maximum admissible value of 1, with little change over the final decade of the sample. In turn, the developing-country mean also displays an upward trend, but starting more than a decade later (in the early 1990s), and the extent of the rise – from under 0.30 at the beginning of the sample to just over 0.50 at the end – is also more modest than in advanced countries. The time path of the overall mean lies between those of the two group means.

4 Empirical results

We proceed in two stages. First, we estimate the income process given by (19) and construct the quasi-innovations to the present value of income growth; then we use the latter to estimate the risk-sharing model (13).

4.1 Estimation of the income process

As already explained, we estimate three different specifications of (16). The first one imposes $\mathbf{A}(L) = \mathbf{0}$ and $\boldsymbol{\beta}(L) = \mathbf{0}$ (which in turn implies that $\mathbf{B} = \mathbf{0}$ in (19)), so that each country's per capita GDP growth rate is an iid process. The second specification allows

[19] Specifically, from the former dataset we take the KAOPEN financial openness indicator, while from the latter we use the ratio of total foreign assets and liabilities to GDP.

[20] Foreign asset and liability positions, relative to GDP, provide an alternative way to assess the trends in de facto international financial integration. They also exhibit a marked rising trend over the sample period and, like in Figure 1, the rise is steeper among advanced than among emerging countries.

Figure 1: Financial openness

The figure displays the evolution of the average Chinn-Ito index of financial openness for all countries and for the group of advanced and emerging countries.

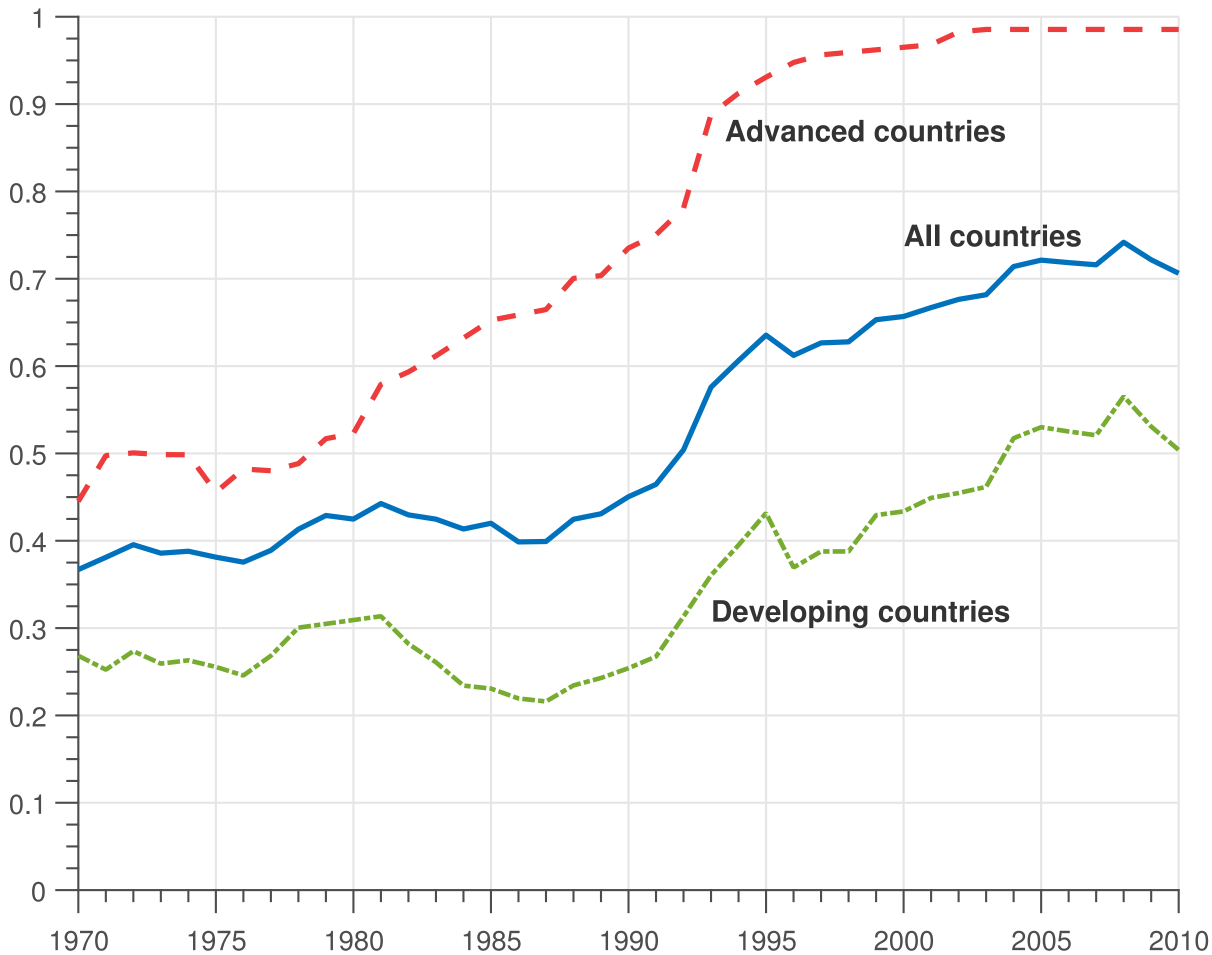

$\mathbf{A}(\mathrm{L}) \neq \mathbf{0}$ but still imposes $\boldsymbol{\beta}(\mathrm{L}) = \mathbf{0}$, so that each country's GDP growth is described by an AR process; experiments with both first- and second-order processes proved the former to be sufficient.[21] The third specification is a factor-augmented AR model with $\mathbf{A}(\mathrm{L}) \neq \mathbf{0}$ and $\boldsymbol{\beta}(\mathrm{L}) \neq \mathbf{0}$. All three specifications allow for unrestricted cross-country heteroskedasticity and parameter heterogeneity.

Table 1 reports the estimation results and forecasting performance of the three specifications. The top half of the table summarizes the parameter estimates (except for the iid specification). With the AR(1) model, shown in the second column, the average of the country-specific autoregressive parameters equals 0.24, and is significantly different from zero. Twenty-five of the fifty individual estimates are significantly different from

[21]Standard tests could not reject the joint null that all the coefficients on twice-lagged growth equal zero, so the results are not reported here.

zero at the 10 percent confidence level. The majority of the estimates (all but six) are positive, and all are smaller than one in absolute value (the largest one equals 0.65), so the estimated growth dynamics are stable. Across countries, the median adjusted R-squared of the AR(1) models exceeds 0.4. Lastly, a Wald test resoundingly rejects the null that all the AR parameters equal zero.

Next we turn to the factor-augmented AR(1) model. Most of the information criteria of Bai and Ng (2002) and Choi (2013) suggest one single static factor. This in turn implies that only the contemporaneous value of the factor enters the growth equation (16). The common factor itself exhibits little persistence; we use a first-order autoregressive process to characterize its dynamics, but the empirical results are virtually unchanged if we instead assume the common factor to be serially uncorrelated.[22]

The third column of Table 1 summarizes the parameter estimates of the factor-augmented model. The average value of the country-specific autoregressive parameters increases slightly, to 0.27. As before, all the individual estimates are smaller than one in absolute value; all but four are positive, and thirty-three of them are significantly different from zero at the 10 percent level.

We scale the global factor so that it has unit variance, and choose its sign so that the loading of the largest country (the U.S.) is positive. Across countries, the estimated loadings average 0.08. The vast majority (all but four) are positive, and thirty five out of fifty are significantly different from zero at the 10 percent level. A Wald test shown in Table 1 confirms that the factor loadings are jointly highly significant. The global factor contributes a major fraction of the overall variance of growth – over one-fourth on average.[23] The contribution is especially large among advanced countries: the 12 countries where it exceeds 50 percent comprise the U.S., Canada, and 10 European Union members. Finally, inclusion of the global factor raises substantially the explanatory power of the growth equation: across countries, the median adjusted R-squared rises to 0.56.

As a further check on model selection, the bottom half of Table 1 reports the out-of-sample forecast performance of the three growth specifications. For this exercise, we re-estimated them setting aside the last 5 annual observations, which were then used for forecasting. According to the RMSEs in the table, the factor-augmented specification exhibits the best performance at all but the 5-year horizon, while the iid specification shows the worst.

[22] The reason is that the autoregressive parameter is very small and falls short of statistical significance. However, the first-order lag improves marginally the forecasting performance of the factor model over that of its serially uncorrelated version, so we opt for keeping it in the model.

[23] The value refers to the share of the long-run variance attributable to the global factor, calculated as $\mathsf{Var}(\boldsymbol{\beta}_i(L)\mathbf{f}_t)/\mathsf{Var}(1 - L\boldsymbol{A}_i(L)\Delta y_{i,t-1})$.

Table 1: **Estimates of the income prediction model**

The table reports the estimates and the forecasting performance of different models for the income process: IID assumes that each country's per-capita GDP growth is an iid process; AR(1) imposes a first order autoregressive process for per-capita GDP growth independent across countries; and FAPAR(1) refers to a factor augmented model in which each country's per-capita GDP growth depends on its own lagged value and a global factor. Robust t-statistics appear in parentheses. The Wald test statistics are distributed as chi-square with 50 degrees of freedom.

	IID	AR(1)	FAPAR(1)
Lagged GDP growth			
Mean of the country-specific AR estimates	–	0.24	0.27
		(11.04)	(13.66)
Number of significant (10%) AR estimates	–	25	33
Global factor			
Mean of the country-specific factor loadings	–	–	0.08
			(17.47)
Number of significant (10%) factor loadings	–	–	35
Average long-run variance contribution	–	–	0.27
Median $\bar{R}^2$		0.42	0.56
Wald test of joint significance (p-values)			
Lagged GDP growth	–	0.00	0.00
Global factor loadings	–	–	0.00
Out of sample forecast performance (% RMSE)			
1-step ahead	2.62	2.26	2.21
2-step ahead	2.56	2.46	2.46
3-step ahead	3.09	3.03	2.99
4-step ahead	6.49	6.44	6.39
5-step ahead	3.15	3.15	3.17

Using the parameter estimates summarized in Table 1, the quasi-innovations to the present discounted value of income growth can be readily constructed employing (19) recursively. We construct three sets of innovations, corresponding to the three model specifications in Table 1. The three sets of innovations are not very different from each other: their pairwise correlations exceed 0.90. In particular, the correlation between the innovations obtained from the AR(1) specification and those obtained from its factor-augmented version equals 0.99, which indicates that the two sets of innovations are virtually identical. Still, as the results shown in the table give a slight advantage to the factor-augmented model, it is the primary focus of the analysis below. However, the estimates of the risk-sharing model show very little change if the AR(1) income growth specification is used instead.

4.2 Estimation of the risk-sharing model

Table 2 summarizes the estimates of the risk-sharing parameters obtained using the three sets of quasi-innovations constructed in the previous step. They share some common features. First, although the estimation imposes no restrictions on the parameters whatsoever, the vast majority do lie between zero and one, as predicted by theory. In particular, when using our preferred specification of the income process (the factor-augmented AR(1) in column 3), all the country-specific parameters except one fall in the admissible zone. The remaining one is negative but not significantly so, according to a one-sided t-test. Likewise, the estimates based on the AR(1) model (column 2) yield two negative parameters,[24] but again standard tests show that they are not significantly smaller than zero, neither individually nor jointly.[25] In contrast, the least-preferred iid specification of the income growth process (column 1) yields seven negative parameters, of which two are significantly negative at the 10 percent level.

The other common feature is the heterogeneity of the individual-country risk-sharing parameter estimates, and in particular the fact that, for all three specifications in Table 2, the estimates are higher on average for advanced countries than for developing countries. The Wald tests at the bottom of the table confirm that these differences across estimates, as well

[24]The negative estimate in column 3 corresponds to Turkey, to which column 2 adds Argentina. Both countries are in the top decile of the sample in terms of consumption growth volatility.

[25]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 all parameters were explicitly constrained to lie within the unit interval. To investigate this issue, we re-estimated the two specifications in columns 2 and 3 of 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 the restricted and unrestricted estimates is around 0.98 under both the AR(1) and the common-factor growth prediction models. Thus, to save space we do not report the restricted estimates.

Table 2: **NLS estimates of the country-specific risk-sharing parameters**
Robust t-statistics in parentheses, 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.

Income forecasting models	(1) IID	(2) AR(1)	(3) FAPAR(1)
Summary statistics of risk sharing estimates			
Number of risk sharing parameters between 0 and 1	43	48	49
of which significantly positive	29	43	40
Number of risk sharing parameters greater than 1	0	0	0
of which significantly greater than 1	0	0	0
Number of risk sharing parameters smaller than 0	7	2	1
of which significantly smaller than 0	2	0	0
Median	0.22	0.52	0.53
Maximum	0.87	0.90	0.90
Minimum	-0.23	-0.03	-0.07
Average estimates			
All countries	0.26	0.49	0.50
	(13.2)	(27.7)	(27.7)
Advanced countries	0.30	0.57	0.59
	(10.0)	(23.3)	(23.9)
Developing countries	0.24	0.42	0.43
	(9.6)	(18.2)	(17.2)
Wald test of difference in means (p-value)	0.00	0.00	0.00
Wald test of joint significance of all risk sharing parameters	0.00	0.00	0.00

as between the two group means, are statistically highly significant. Additional tests show that even within these two groups the individual-country coefficients display significant variation: 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.

Figure 2 plots the estimates of the individual-country risk-sharing parameters under the three specifications (Table B.2 in the appendix reports their actual values). It can be seen that the AR(1) income growth process and its factor-augmented version yield very similar estimates for almost all countries – indeed, the correlation between both sets of estimates is 0.96. In contrast, the estimates based on the iid process are fairly different from the rest – for example, their correlation with those obtained with the factor-augmented income process is just 0.70.

Figure 2: Estimates of risk-sharing parameters

The left panel of the figure displays a scatter plot with the estimated coefficient of risk sharing λ assuming iid income growth (vertical axis) and imposing a factor model on income growth (horizontal axis). The right panel displays a similar scatter plot but assuming that income growth follow independent AR(1) processes for each country (vertical axis) and using the factor model (horizontal axis).

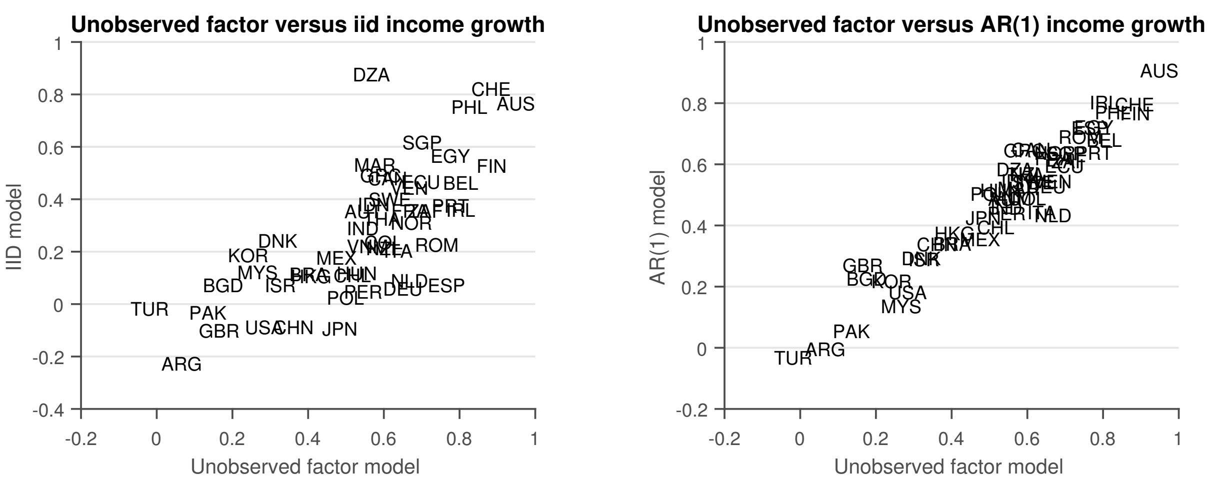

With our preferred specification, shown in column 3 of Table 2, the country-specific risk-sharing parameters are centered around 0.50. Figure 3 shows histogram plots of the distributions of risk-sharing coefficients of the two sets of countries we consider. The advanced-country distribution lies to the right of its developing-country counterpart. The mean for industrial countries is just under 0.60, while for developing countries it equals 0.43. These average estimates are not too different from those obtained by Crucini (1999) using the homogeneous 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 also 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. Kose et al. (2009b), using conventional risk-sharing regressions, and Flood et al. (2012), using a different measure of imperfect risk sharing, also find that advanced countries exhibit higher degrees of risk sharing than developing countries.

Figure 3: Histogram of estimated λ's

The figure displays histograms of the estimated coefficient of partial risk sharing $\hat{\lambda}_i$ for the groups of advanced and developing countries, using the unobserved factor model to forecast the income processes.

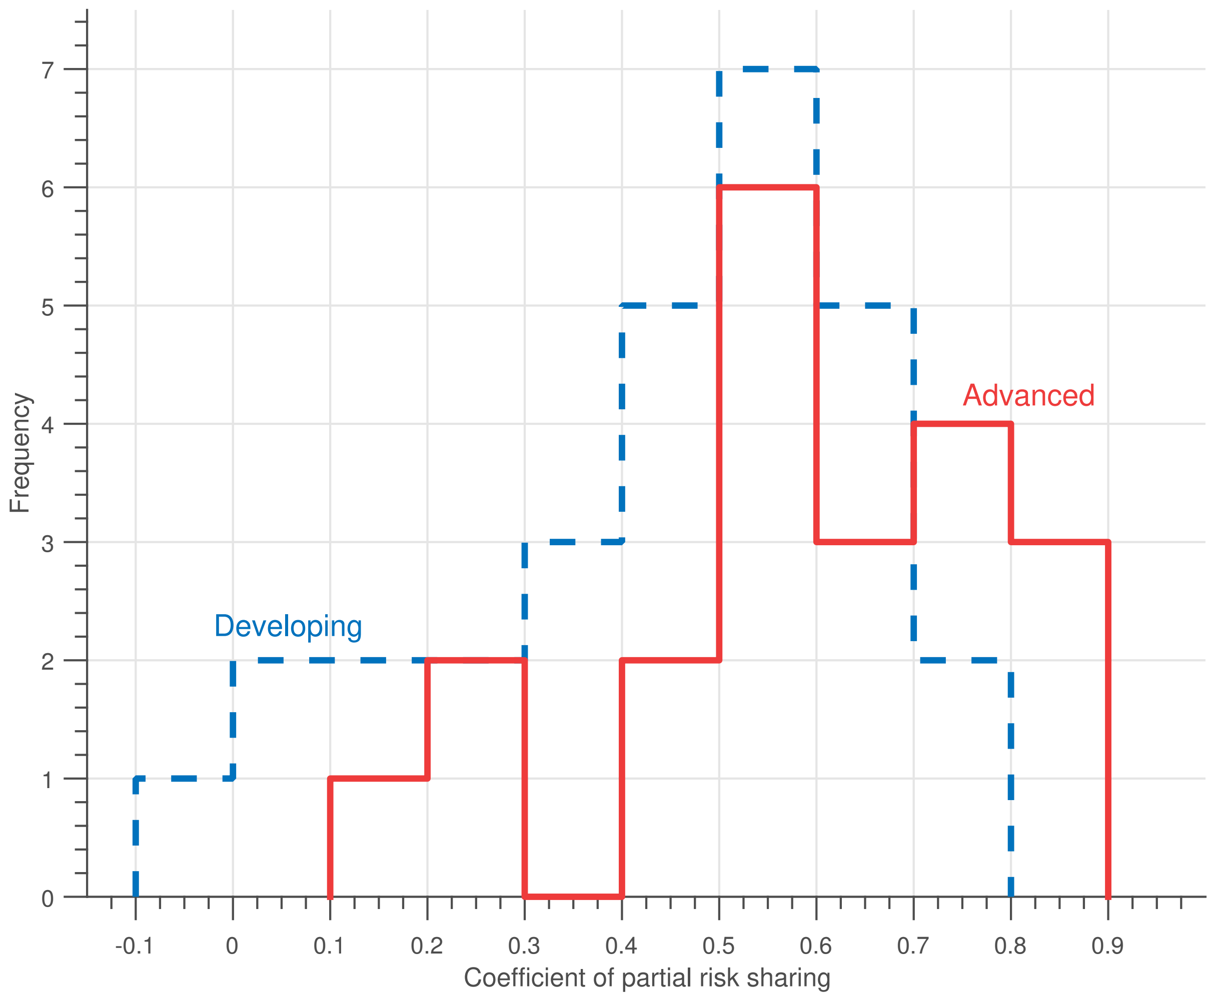

4.3 Patterns of consumption risk sharing over time

Since the mid-1980s, the world has seen a large increase in the degree of financial integration, facilitated by the removal of barriers to international capital movements (as illustrated in Figure 1), 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 20 years each. Importantly, prior to estimation of the risk sharing model we also recalculate the innovations to the present value of income growth in the same way, by estimating the income process over matching rolling samples of 20 years of data each; this ensures that both construction of the innovations and estimation of the risk sharing parameters use data from the same time period.[26]

The results from this exercise, based on the factor-augmented specification of the income growth process, are summarized in Figure 4. The figure plots the mean of the risk-sharing parameter estimates obtained in each estimation window, along with their two standard-error bands. The top panel reports the overall mean, while the lower panels report the means for industrial and developing countries.

The overall mean for all countries displays a rising trend, especially marked in the windows beginning around the early 1980s. From that point on, the average coefficient of risk sharing rises without interruption until the global crisis enters the estimation window. Thereafter, the average estimate shows a relatively modest decline in the last two estimation windows. Still, the mean estimates obtained in the last windows lie 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. This is confirmed by a quick glance at the individual-country risk sharing parameter estimates: a majority of countries (33 out of 50) exhibit higher coefficients in the final estimation window than in the initial one, and for 19 of them the increase is significant, in the sense that the final estimate lies outside the 95 percent confidence interval of the initial one. Note that the increase in the average coefficients of risk sharing over time matches the upward trend of financial openness, as displayed in Figure 1.

The bottom panels of Figure 4 plot the time path of the average risk sharing coefficients for industrial and developing countries, along with their respective two standard-error bands. There is a clear contrast between the mean estimates of the two country groups. The mean for industrial countries displays a rising trend over time, especially steep across the

[26] Similarly, for each window we recalculate population $\mathbf{H}$ and the initial income shares $\boldsymbol{W_0}$ using the average of the two annual observations preceding the initial year of the window.

Figure 4: Evolution of estimated risk-sharing coefficients

The figure displays the average coefficient of risk sharing for all countries (top panel), advanced countries (bottom left panel), and developing countries (bottom right panel) including 90 percent confidence bands. The estimates were computed using rolling samples of 20 years each and using the unobserved factor model to perform forecasts of future output growth. The first window corresponds to the years 1972-1991; the second, the years 1973-1992; and so forth until the years 1991-2010.

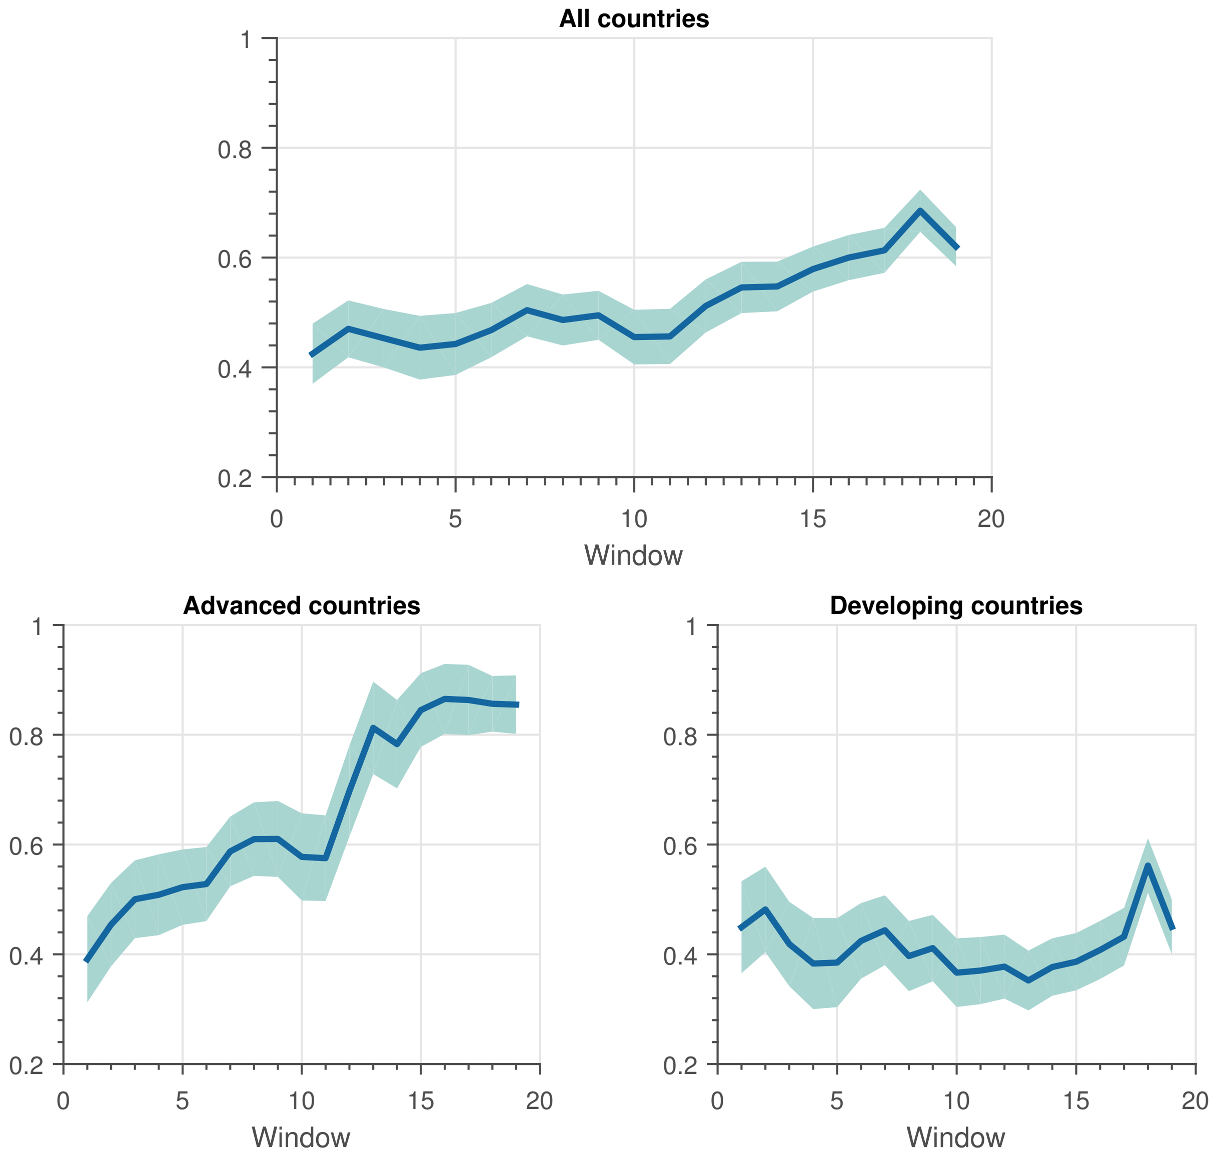

estimation windows that begin in the mid-1980s. The rise eventually tapers off over the last estimation windows. In contrast, there has been no discernible trend among developing countries. The figure suggests, if anything, a slight initial decline in their mean risk sharing, followed by a recovery. In fact, for the developing country mean the 95-percent confidence region of the final window is contained almost in full in that of the initial one, which suggests that no significant change has taken place. For advanced countries, however, the

95-percent confidence region corresponding to the final window lies fully outside that obtained in the initial window, pointing to a significant rise in mean risk sharing among industrial countries. Thus, this rising trend of average risk sharing among rich countries is the force behind the trend increase in overall mean risk sharing found in the top panel of the figure. This conclusion is confirmed by the individual country estimates: 22 of the 33 countries whose risk-sharing coefficients are higher in the final estimation window than in the initial one belong to the industrial country group, while 16 of the 17 countries whose coefficients are lower belong to the developing country group.

As already noted, the estimates in Figure 4 are based on our preferred specification of the income growth process – the factor-augmented AR(1) model. However, using instead the AR(1) specification without common factors yields very similar qualitative results – i.e., advanced-country estimates display a rising trend over time (although somewhat less marked than with the factor-augmented model) while developing-country estimates remain flat. In contrast, however, the upward trend of the rich-country estimates disappears if the rolling estimation is based on the iid specification of the income growth process. Recall that, as shown in Table 1, the empirical performance of the iid specification is inferior to those of the other alternatives considered, both in the estimation sample as well as in out-of-sample forecasting. However, it is the specification employed by Ho and Ho (2015), who find little change in risk-sharing estimates over time. Hence, the results in this section suggest that such conclusion is not robust to the use of improved growth forecasting models.

4.4 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. [27]

[27] For example, using this approach, Kose et al. (2009b) 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.

Formally, to explore the covariates of consumption risk sharing we re-estimate the model as specified in (15), that is, with the vector of country-specific risk-sharing coefficients given by $\boldsymbol{\lambda} = m(\mathbf{X}\boldsymbol{\delta})$, where $\mathbf{X}$ is a matrix of time-invariant risk-sharing determinants, and $\boldsymbol{\delta}$ now is the vector of parameters to be estimated.[28]

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 quasi-innovations constructed using the factor-augmented specification of income growth (that is, the model in the last column of Table 1); results using instead AR(1)-based forecasts are very similar but are not reported to save space. For each estimate of $\boldsymbol{\delta}$ we can compute the implied risk-sharing coefficients of the different countries as $\hat{\boldsymbol{\lambda}} = \exp(\mathbf{X}\hat{\boldsymbol{\delta}})$; these are summarized in the bottom half of the table.

The first column of Table 3 relates the degree of risk sharing to the Chinn-Ito measure of financial openness. The parameter estimate is positive and highly significant, implying that higher degrees of openness come along with higher consumption risk sharing. The individual risk-sharing coefficients implied by this specification are centered around 0.55. On average, they are higher for industrial countries than for developing countries. However, their correlation with the unrestricted estimates (shown in the first column of Table B.2) is just 0.27.

As already argued, larger countries should be expected to exhibit lower degrees of risk sharing, as their income shocks have a bigger effect on the overall income pool than do shocks to the income of smaller economies – they are, to a larger extent, common shocks – and thus they derive a smaller benefit from contributing to the income pool. To verify this conjecture, Column 2 of Table 3 relates each country's degree of risk sharing to its aggregate GDP. The parameter estimate is negative and highly significant. However, the implied risk-sharing coefficients are, on average, higher for developing countries than for industrial countries, and their correlation with unrestricted estimates in Table 2 is less than 0.2.

Column 3 turns to the total foreign asset and liability stock, as percent of GDP, as a measure of international financial integration. Its role in conventional regressions of consumption risk sharing has been explored, for example, by Kose et al. (2009b). The parameter estimate is positive but insignificant. Further, like in the preceding column, the implied risk-sharing coefficients are on average higher in industrial countries than in developing countries. They also show a very weak correlation with the unrestricted

See also Sørensen et al. (2007), Fratzscher and Imbs (2009), and Holinski et al. (2012).

[28]The variables in $\mathbf{X}$ are measured as time averages. The functional form we use is $\boldsymbol{\lambda} = \exp(\mathbf{X}\boldsymbol{\delta})$, which offers the advantage of preventing negative values of $\hat{\boldsymbol{\lambda}}$.

Table 3: NLS estimates of the covariates of risk sharing

Robust t-statistics in parentheses, using the spatial-correlation consistent standard errors of Driscoll and Kraay (1998). 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 the estimated unrestricted coefficients of partial risk sharing.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Capital account openness (Chin-Ito)	0.221				0.181	0.280	0.250
	(2.61)				(2.38)	(3.29)	(2.80)
Log real GDP		-0.072			-0.064	-0.067	-0.079
		(2.67)			(2.50)	(2.44)	(2.94)
Total foreign assets+foreign liabilities/GDP			0.000			0.004	0.005
			(1.37)			(2.39)	(2.59)
Imports + exports / GDP				0.000		-0.004	-0.005
				(0.84)		(2.55)	(2.93)
Private sector credit / GDP							0.001
							(2.81)
Wald test of joint significance					0.000	0.000	0.000
Implied risk sharing coefficients							
Median	0.561	0.525	0.521	0.525	0.530	0.526	0.549
Maximum	0.645	0.628	0.639	0.593	0.650	0.814	0.822
Minimum	0.446	0.260	0.505	0.515	0.305	0.293	0.255
Mean-All countries	0.548	0.502	0.528	0.529	0.521	0.521	0.533
Advanced countries	0.601	0.472	0.523	0.526	0.538	0.555	0.604
Developing countries	0.509	0.524	0.532	0.530	0.509	0.496	0.482
Correlation with unrestricted risk sharing coefficients	0.270	0.189	0.126	0.098	0.352	0.436	0.433

estimates. Similar considerations apply to the results in column 4, which focuses on trade openness, measured as total imports plus exports over GDP. The estimation finds no significant association with risk sharing.

Columns 5-7 report estimation results combining these variables. In column 5, size (as measured by real GDP) and financial openness (as captured by the Chinn-Ito index) are both highly significant, and their coefficients show little change relative to those in columns 1-2. The correlation of the implied risk-sharing coefficients with their unrestricted counterparts rises considerably. Column 6 adds to the specification the total foreign asset and liability position as well total foreign trade, both as percent of GDP. The former variable carries a positive coefficient, while the latter carries a negative one; both are statistically significant. The positive sign of the foreign asset and liability position echoes the result found by Kose et al. (2009b) for advanced countries (but not for developing countries). In turn, the negative sign of the trade openness measure suggests that, once financial openness is taken into account, increased trade openness reduces the estimated coefficient of risk sharing. One possible explanation for this result is that, by amplifying the impact of global terms of trade shocks, increased trade openness raises the correlation between domestic income and global output, reducing the benefits from participating in the risk-sharing agreement (see e.g., Loayza and Raddatz, 2007). The risk-sharing coefficients implied by the specification in column 6 exhibit a fairly high correlation (0.44) with the unrestricted estimates in column 3 of Table 2. The range of the individual estimates also widens considerably relative to that found in the other columns of Table 3.

Finally, column 7 adds domestic financial development, as measured by the stock of credit to the private sector over GDP. Its parameter estimate is positive and significant, while the parameter estimates of the other regressors show only modest changes. The range of the implied individual-country estimates widens further, as does the gap between the implied advanced- and developing-country averages, while the correlation with the unrestricted risk-sharing coefficients remains virtually unchanged relative to the previous column.

To conclude this section, 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 financial openness and the magnitude of countries' external portfolios are positively related to their performance in terms of consumption risk sharing. In contrast, the size of the economy, and its degree of real openness, show the opposite pattern.

5 Concluding remarks

A considerable literature is concerned with assessing the extent to which countries share their consumption risk. Much of it, however, uses 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 present discounted value of income growth 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 has 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, and more developed domestic financial sectors exhibit larger degrees of risk sharing. In contrast, the size of the economy and its degree of trade openness (after controlling for financial openness) seem to hinder consumption risk sharing.

References

Artis, Michael and Mathias Hoffmann, "Financial Globalization, International Business Cycles and Consumption Risk Sharing," *Scandinavian Journal of Economics*, 2008, *110* (3), 447–471.

Asdrubali, Pierfederico and Soyoung Kim, "Incomplete Intertemporal Consumption Smoothing and Incomplete Risk Sharing," *Journal of Money, Credit and Banking*, October 2008, *40* (7), 1521–1531.

Bai, Jushan and Serena Ng, "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, January 2002, *70* (1), 191–221.

_ and _ , "Confidence Intervals for Diffusion Index Forecasts and Inference for Factor-Augmented Regressions," *Econometrica*, 07 2006, *74* (4), 1133–1150.

_ and _ , "Determining the Number of Primitive Shocks in Factor Models," *Journal of Business & Economic Statistics*, 2007, *25*, 52–60.

_ and _ , "Principal components estimation and identification of static factors," *Journal of Econometrics*, 2013, *176* (1), 18–29.

Bai, Yan and Jing Zhang, "Financial integration and international risk sharing," *Journal of International Economics*, 2012, *86* (1), 17–32.

Bengui, Julien, Enrique G. Mendoza, and Vincenzo Quadrini, "Capital mobility and international sharing of cyclical risk," *Journal of Monetary Economics*, 2013, *60* (1), 42–62.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, December 2008, *98* (5), 1887–1921.

Byrne, Joseph P., Giorgio Fazio, and Norbert Fiess, "The Global Side of the Investment-Saving Puzzle," *Journal of Money, Credit and Banking*, 2009, *41* (5), 1033–1040.

Campbell, John and N. Gregory Mankiw, "Permanent Income, Current Income, and Consumption," *Journal of Business & Economic Statistics*, 1990, *8* (3), 265–79.

Campbell, John Y. and N. Gregory Mankiw, "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," in "NBER Macroeconomics Annual 1989, Volume 4" NBER Chapters, National Bureau of Economic Research, Inc, 1989, pp. 185–246.

Canova, Fabio and Morten O. Ravn, "International Consumption Risk Sharing," *International Economic Review*, 1996, *37* (3), 573–601.

Choi, In, "Model Selection for Factor Analysis: Some New Criteria and Performance Comparisons," Working Papers 1209, Research Institute for Market Economy, Sogang University 2013.

Coakley, Jerry, Farida Hasan, and Ron Smith, "Saving, Investment, and Capital Mobility in LDCs," *Review of International Economics*, 1999, *7* (4), 632–640.

Cochrane, John H, "A Simple Test of Consumption Insurance," *Journal of Political Economy*, October 1991, *99* (5), 957–76.

Crucini, Mario J., "On International and National Dimensions of Risk Sharing," *The Review of Economics and Statistics*, 1999, *81* (1), 73–84.

Crucini, Mario J and Gregory D Hess, "International and Intranational Risk Sharing," in Gregory D Hess and Eric van Wincoop, eds., *Intranational Macroeconomics*, NY: Cambridge University Press, 2000, chapter 2, pp. 37–59.

Deaton, Angus, "On risk, insurance, and intra-village consumption smoothing," 1990.

Devereux, Michael B. and Changhua Yu, "International financial integration and crisis contagion," Working Paper, Vancouver School of Economics 2016.

Driscoll, John C. and Aart C. Kraay, "Consistent Covariance Matrix Estimation With Spatially Dependent Panel Data," *The Review of Economics and Statistics*, November 1998, *80* (4), 549–560.

Flood, Robert P., Nancy P. Marion, and Akito Matsumoto, "International risk sharing during the globalization era," *Canadian Journal of Economics*, 2012, *45* (2), 394–416.

Fratzscher, Marcel and Jean Imbs, "Risk sharing, finance, and institutions in international portfolios," *Journal of Financial Economics*, 2009, *94* (3), 428 – 447.

Hansen, Lars Peter and Thomas J. Sargent, "Formulating and estimating dynamic linear rational expectations models," *Journal of Economic Dynamics and Control*, 1980, *2*, 7 – 46.

Ho, Chun-Yu and Wai-Pin Alex Ho, "Dynamics and Heterogeneity of Inter- and Intranational Risk Sharing," *Journal of Money, Credit and Banking*, 2015, *47* (1), 123–142.

Holinski, Nils, Clemens J.M. Kool, and Joan Muysken, "The impact of international portfolio composition on consumption risk sharing," *Journal of International Money and Finance*, 2012, *31* (6), 1715 – 1728.

Kose, M Ayhan, Eswar Prasad, Kenneth Rogoff, and Shang-Jin Wei, "Financial Globalization: A Reappraisal," *IMF Staff Papers*, April 2009, *56* (1), 8–62.

Kose, M. Ayhan, Eswar S. Prasad, and Marco E. Terrones, "Does financial globalization promote risk sharing?," *Journal of Development Economics*, July 2009, *89* (2), 258–270.

Lane, Philip R. and Gian Maria Milesi-Ferretti, "The external wealth of nations mark II: Revised and extended estimates of foreign assets and liabilities, 1970-2004," *Journal of International Economics*, November 2007, *73* (2), 223–250.

Lewis, Karen K., "What Can Explain the Apparent Lack of International Consumption Risk Sharing?," *Journal of Political Economy*, 1996, *104* (2), 267–297.

Loayza, Norman and Claudio Raddatz, "The Structural Determinants of External Vulnerability," *World Bank Economic Review*, 2007, *21* (3), 359–387.

Obstfeld, Maurice, "Are Industrial-Country Consumption Risks Globally Diversified?," in Leonardo Leiderman and Assaf Razin, eds., *Capital mobility: the impact on consumption, investment and growth*, Cambridge UK: Cambridge University Press, 1994, chapter 2, pp. 13–44.

_ and Kenneth Rogoff, "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?," *NBER Macroeconomics Annual*, 2001, *15*, 339–390.

Song, Minkee, "Asymptotic theory for dynamic heterogeneous panels with cross-sectional dependence and its applications," Working Paper, Korea Institute of Finance 2013.

Sørensen, Bent E., Yi-Tsung Wu, Oved Yosha, and Yu Zhu, "Home bias and international risk sharing: Twin puzzles separated at birth," *Journal of International Money and Finance*, 2007, *26* (4), 587 – 605. Financial Globalization and Integration.

Stock, James H. and Mark W. Watson, "Implications of Dynamic Factor Models for VAR Analysis," NBER Working Papers, National Bureau of Economic Research, Inc July 2005.

Townsend, Robert M, "Risk and Insurance in Village India," *Econometrica*, May 1994, *62* (3), 539–91.

Vogelsang, Timothy J., “Heteroskedasticity, autocorrelation, and spatial correlation robust inference in linear panel models with fixed-effects,” *Journal of Econometrics*, 2012, *166* (2), 303–319.

A Appendix: Proofs

In this appendix we derive expression (7). Dividing equation (2) by $Y_{i,t}$ and using the definition of $\omega_{j,0}$ we obtain

$$\frac{\bar{Y}_{i,t}}{Y_{i,t}} = 1 - \lambda_i + \lambda_i \sum_{j=1}^{N} \left(\frac{Y_{i,0}}{Y_{j,0}}\right) \theta_j \frac{Y_{j,t}}{Y_{i,t}}.$$

Taking logs on both sides of the equation gives

$$\bar{y}_{i,t} - y_{i,t} = \log\left(1 - \lambda_i + \lambda_i \sum_{j=1}^{N} \left(\frac{Y_{i,0}}{Y_{j,0}}\right) \theta_j \exp\left(y_{j,t} - y_{i,t}\right)\right).$$

Linearizing the expression on the right side around $Y_{j,0}/Y_{i,0} = \exp(y_{j,0} - y_{i,0})$ gives

$$\bar{y}_{i,t} - y_{i,t} \approx \log\left(1 - \lambda_i + \lambda_i \sum_{j=1}^{N} \theta_j\right) + \frac{\lambda_i \sum_{j=1}^{N} \theta_j \left[y_{j,t} - y_{i,t} - (y_{j,0} - y_{i,0})\right]}{1 - \lambda_i + \lambda_i \sum_{j=1}^{N} \theta_j}.$$

But $\sum_{j=1}^{N} \theta_j = 1$ implies $1 - \lambda_i + \lambda_i \sum_{j=1}^{N} \theta_j = 1$, hence the approximation becomes

$$\bar{y}_{i,t} \approx (1 - \lambda_i) y_{i,t} + \lambda_i \sum_{j=1}^{N} \theta_j y_{j,t} + \lambda_i (y_{i,0} - \sum_{j=1}^{N} \theta_j y_{j,0}).$$

Differentiating this expression with respect to time gives equation (7).□

B Appendix: Additional tables

Table B.1: **List of countries**

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

Table B.2: **Estimated coefficients of partial risk sharing**

The table reports the country-specific estimates of the coefficients of partial risk sharing λ_i for different models for the income process: IID assumes that each country's per-capita GDP growth is an iid process; AR(1) imposes a first order autoregressive process for per-capita GDP growth independent across countries; and FAPAR(1) refers to a factor augmented model in which each country's per-capita GDP growth depends on its own lagged value and a global factor.

Advanced countries				**Developing countries**			
	FAPAR(1)	AR(1)	IID		FAPAR(1)	AR(1)	IID
Australia	0.898	0.904	0.762	Algeria	0.518	0.582	0.874
Austria	0.496	0.485	0.352	Argentina	0.012	-0.007	-0.230
Belgium	0.757	0.677	0.460	Bangladesh	0.122	0.224	0.070
Canada	0.558	0.647	0.476	Brazil	0.351	0.338	0.111
Denmark	0.267	0.291	0.239	Chile	0.467	0.392	0.107
Finland	0.845	0.764	0.525	China	0.308	0.337	-0.090
France	0.620	0.618	0.354	Colombia	0.548	0.487	0.233
Germany	0.598	0.526	0.056	Ecuador	0.645	0.592	0.463
Greece	0.537	0.643	0.488	Egypt, Arab Rep.	0.723	0.719	0.564
Ireland	0.765	0.800	0.358	Hong Kong SAR, China	0.354	0.373	0.103
Italy	0.600	0.442	0.201	Hungary	0.475	0.512	0.114
Japan	0.437	0.422	-0.096	India	0.502	0.456	0.287
Netherlands	0.619	0.430	0.086	Indonesia	0.530	0.544	0.378
New Zealand	0.553	0.566	0.211	Israel	0.285	0.287	0.069
Norway	0.618	0.635	0.307	Korea, Rep.	0.188	0.216	0.183
Portugal	0.727	0.636	0.372	Malaysia	0.213	0.133	0.119
Spain	0.716	0.717	0.069	Mexico	0.422	0.353	0.175
Sweden	0.560	0.538	0.398	Morocco	0.522	0.520	0.529
Switzerland	0.831	0.794	0.818	Pakistan	0.086	0.053	-0.036
United Kingdom	0.112	0.268	-0.104	Peru	0.495	0.437	0.045
United States	0.233	0.179	-0.090	Philippines	0.778	0.767	0.749
				Poland	0.449	0.501	0.023
				Romania	0.682	0.688	0.224
				Singapore	0.649	0.634	0.615
				South Africa	0.663	0.608	0.353
				Thailand	0.546	0.564	0.323
				Turkey	-0.069	-0.034	-0.019
				Venezuela, RB	0.616	0.542	0.442
				Vietnam	0.502	0.489	0.219