



Pooling risk among countries [☆]

Michael Callen ^a, Jean Imbs ^{b,c,*}, Paolo Mauro ^d

^a Kennedy School of Government, Harvard University, United States

^b Paris School of Economics (CNRS), France

^c CEPR, United Kingdom

^d Peterson Institute for International Economics, United States



ARTICLE INFO

Article history:

Received 4 September 2014

Received in revised form 28 December 2014

Accepted 20 January 2015

Available online 31 January 2015

JEL classification:

E21

E32

E34

F41

Keywords:

Risk sharing

Diversification

Growth-indexation

GDP-indexed instruments

ABSTRACT

Suppose that international sharing risk—worldwide or with large numbers of countries—were costly. How much risk-sharing could be gained in small sets (or “pools”) of countries? To answer this question, we compute the means and variances of poolwide gross domestic product growth, for all possible pools of any size drawn from a sample of 74 countries, and compare them with the means and variances of consumption growth in each country individually. From the difference, we infer potential diversification and welfare gains. As much as two-thirds of the first best, full worldwide welfare gains can be obtained in groupings of as few as seven countries. The largest potential gains arise from pools consisting of countries in different regions and including countries with weak institutions. We argue that international risk-sharing fails to emerge because the largest potential gains are among countries that do not trust each other’s willingness and ability to abide by international contractual obligations.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Despite major strides in lifting capital controls around the world and impressive increases in cross-border holdings of financial assets, international financial integration is still far from complete: Individual countries’ consumption remains more volatile than what would result from complete risk-sharing with the rest of the world. In this paper, we conjecture that practical obstacles make international risk sharing costly, and that they depend on the number and characteristics of potential partner countries to share risk with. Then, it can become desir-

able to share risk within a well chosen subset of countries only.¹ This paper evaluates the potential gains associated with such limited risk sharing contracts. How large would such groups need to be for the gains to be sizable? And which groups would yield the largest gains?

Our main contribution consists in running a systematic search on all possible groupings, or “pools” of countries, using the variance-covariance matrix of output and consumption growth rates observed in standard data for 74 countries. We compare the observed volatility of consumption for each country with the volatility of poolwide output,

[☆] We are grateful to two anonymous referees, and to Eric van Wincoop, Tamim Bayoumi, Mick Devereux, Michael Kremer, Raghuram Rajan, Jaime Ventura, and seminar participants at U.C. Berkeley, Cambridge U., the International Monetary Fund, U. Pompeu Fabra, the San Francisco Fed, and Tufts U. for insightful suggestions; to Nicolas Metzger, and José Romero for superb research assistance; and to Jean Salvati and especially Huigang Chen and Alin Mirastean for invaluable help in programming and computational support. Financial support from the Banque de France Chair at the Paris School of Economics is gratefully acknowledged. All errors are our own.

* Corresponding author at: Paris School of Economics, 106 Bvd de l’Hopital, 75013 Paris, France.

E-mail address: Jean.Imbs@parisschoolofeconomics.eu (J. Imbs).

¹ We do not observe worldwide risk sharing, but we do not observe sharing of GDP risk within groups of countries either. How could risk sharing be achieved in practice for a subset of countries? Various schemes have been proposed, including Robert C. Merton’s (1990, 2000) networks of bilateral swaps of GDP-linked income streams. Elements of risk sharing among groups of countries are also present, for example, in pooling arrangements for international reserves, such as the Chiang Mai initiative, the Latin American Reserve Fund (FLAR), or networks of bilateral swap arrangements (e.g., among the G7 in the 1960s–70s, among the European countries during the run up to the establishment of the Euro, and among several countries during the global financial and economic crisis that began in 2007–8). While these existing real-world arrangements do not seek to share GDP risk explicitly, they do imply some degree of sharing of macroeconomic risks among their member countries. On FLAR, see Eichengreen (2007) and www.flar.net; on the Chiang Mai initiative, see Park and Wang (2005), and <http://aric.adb.org>; on the earlier European experience, see Eichengreen and Wyplosz (1993). On the sharing of GDP risks more generally see Shiller (1993); and Borensztein and Mauro (2004) for a review of the literature.

for each possible pool.² Consumption data are relevant because they reflect the insurance mechanisms that already exist in each country. Under risk sharing within the pool, consumption growth in each country equals poolwide output growth. Therefore, the comparison quantifies the potential for additional diversification gains—and ultimately welfare gains—that would accrue to each country from moving to complete risk-sharing within each pool considered. For any possible pool size, we identify the country groupings that minimize poolwide GDP volatility, and maximize welfare gains from international diversification.

We find that pools of fewer than ten countries can provide the bulk of the potential first best, worldwide risk sharing gains. In those well chosen groupings, the marginal gains decline quickly for groups beyond six or seven members. Many small pools yield large risk-sharing gains. Unsurprisingly, such pools involve relatively volatile economies.³

If the gains from international risk sharing are so significant, and could in principle be attained in relatively small groups of countries, why do such arrangements not emerge more often? One possibility is that they face particularly costly obstacles. As we show in the paper, the largest potential gains are often attained by sharing risk with distant countries characterized by weak institutional quality and a history of default on international debt obligations. In that light, the observed reluctance to engage in international risk-sharing becomes less surprising: it may just stem from insufficient information about the trustworthiness of potential partners or difficulties related to international law. Then, risk sharing is especially costly precisely in those small groups of countries where they would carry maximal gains.

To explore this issue further, we report the potential gains that would result from risk-sharing arrangements if they were constrained to countries selected from a universe with certain characteristics—the same geographic region, or relatively strong institutions. In those potential country groupings, the costs associated with sharing GDP risk are presumably lower, and so a local arrangement is more likely to emerge. But as we show, these are also countries whose GDP risks tend to be strongly correlated, and so where the gains from risk sharing are smaller in the first place.

The rest of the paper is organized as follows. Section 2 provides the theoretical background to our empirical exercise—adapted from existing work. The section also outlines how we handle the considerable combinatorial problem involved in manipulating a sample of 74 countries. Section 3 presents the general results on risk-sharing gains across subsets of countries. In Section 4, we estimate how much risk sharing gains are reduced when countries are constrained to share risk only with specific partners. Section 5 concludes.

2. Methodology

This section first outlines the theory motivating the paper's empirics. Then it describes the algorithms used to compute the risk sharing gains, in any subset of countries, of any size, drawn from a universe of 74 countries with available consumption and output data.

2.1. Risk-sharing, volatility, and welfare

The argument relies on a well known framework, based on Lewis (2000) and Obstfeld (1994). As they do, we abstract from non tradability and non separability in utility, and from the possible impact

of uncertainty on growth. These simplifications enable us to compute the welfare gains for risk sharing among a large set of 74 countries.⁴

Following Epstein and Zin (1989), utility at time t in country j is given by

$$U_t^j = \left\{ \left(C_t^j \right)^{1-\theta} + \beta \left[E_t \left(U_{t+1}^j \right)^{1-\gamma} \right]^{(1-\theta)(1-\gamma)} \right\}^{1/(1-\theta)} \tag{1}$$

where C_t^j is consumption at time t in country j . The process for endowment income at time t in country j is

$$y_t^j = y_{t-1}^j + \mu_j - \frac{1}{2} \sigma_j^2 + \varepsilon_t^j \tag{2}$$

where $y_t^j = \ln Y_t^j$ and $\varepsilon_t^j \sim N(0, \sigma_j^2)$. $0 < \beta < 1$ denotes the subjective discount rate, $\gamma \geq 0$ is the coefficient of relative risk aversion and θ is the inverse of the elasticity of intertemporal substitution in consumption. μ_j denotes the long run growth rate of output in country j , and σ_j^2 its variance around trend growth. Eq. (2) assumes permanent shocks to income, which is well known to magnify the welfare gains from diversification, for any pool size. But the assumption does not affect how quickly welfare gains increase with the number of countries sharing risk, the paper's main question. The assumption of permanent shocks is maintained for tractability.

As in Lewis (2000), the analysis focuses on the welfare gains afforded by international diversification. This assumes away alternative sources of consumption smoothing, such as self insurance and saving. It is consistent with the purpose of evaluating the potential from international risk sharing. Following Lewis (2000) $C_t^j = Y_t^j$ under autarky, and time t welfare in country j is given by

$$U_t^j = C_t^j \left(1 - \beta M_j^{1-\theta} \right)^{-1/(1-\theta)} \tag{3}$$

where $M_j = \exp \left(\mu_j - \frac{1}{2} \gamma \sigma_j^2 \right)$. If instead country j enters a risk sharing agreement, it will have a claim on poolwide income, \bar{Y}_t , which we assume is distributed log-normally, with mean $\bar{\mu}$ and volatility $\bar{\sigma}^2$: $\bar{y}_t = \bar{y}_{t-1} + \bar{\mu} - \frac{1}{2} \bar{\sigma}^2 + \bar{\varepsilon}_t$.⁵ In the pool, time t welfare in country j is therefore given by

$$\bar{U}_t^j = \bar{C}_t^j \left(1 - \beta \bar{M}^{1-\theta} \right)^{-1/(1-\theta)} \tag{4}$$

where $\bar{M} = \exp \left(\bar{\mu} - \frac{1}{2} \gamma \bar{\sigma}^2 \right)$, and \bar{C}_t^j is consumption in country j at time t if the country is part of the risk sharing pool. How much of a claim does country j have on poolwide output? The only asset available in country j is the security that pays Y_t^j , whose price is denoted by p_t^j . Entering the pool means acquiring the security that pays poolwide output \bar{Y}_t , whose price is denoted by \bar{p}_t . Therefore, as in Lewis (2000) country j 's claim on poolwide output at time $t > 0$ is given by

$$\bar{C}_t^j = \frac{p_t^j}{\bar{p}_t} \bar{Y}_t. \tag{5}$$

We now introduce the possibility that a cost has to be paid by any country j willing to participate in the risk sharing agreement. The cost τ_j is paid once and for all, at the time the agreement is contracted: it could for instance be paid outside of the pool, to a supra-national agency, whose remit is to monitor that the agreement is subsequently honored,

² In this regard, our approach differs from previous studies: by Obstfeld (1994) and Lewis (2000), who used consumption data only, and by Tesar (1995), who modeled explicitly the saving / investment decisions that determine production in general equilibrium. However, the paper's results are unchanged if we use consumption data only, or production data only.

³ Pallage and Robe (2003) show that the welfare cost of economic fluctuations is far larger in developing countries than in advanced economies.

⁴ With non-separabilities, the literature usually considers two countries only, see Cole and Obstfeld (1991) or Coeurdacier (2009). Lewis and Liu (forthcoming) consider up to eight countries, but have to deal with issues of existence and uniqueness.

⁵ Lewis (2000) shows the sum of log-linear processes can be approximated by a log-linear process.

Download English Version:

<https://daneshyari.com/en/article/962936>

Download Persian Version:

<https://daneshyari.com/article/962936>

[Daneshyari.com](https://daneshyari.com)