

Pooling Risk Among Countries*

Michael Callen[†] Jean Imbs[‡] Paolo Mauro[§]

August 2014

Abstract

Suppose that international sharing risk—worldwide or with large numbers of countries—implied sizable costs. How much risk-sharing could be gained in small sets (or “pools”) of countries? To answer this question, we compute the moments of poolwide gross domestic product growth, for pools of any size, and compare them with the moments of consumption growth in each country individually. From the difference, we infer the 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 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.

Keywords: Risk Sharing, Diversification.

JEL Classification: E21, E32, E34, F41.

*We are grateful to an anonymous referee, and to Eric van Wincoop, Tamim Bayoumi, Mick Devereux, Michael Kremer, Raghuram Rajan, Jaume 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. All errors are our own. The opinions expressed are those of the authors and do not necessarily represent the views of the IMF. Corresponding author: Jean Imbs - Paris School of Economics - 106 Bvd de l’Hopital - 75013 Paris, France. Jean.Imbs@parisschoolofeconomics.eu.

[†]Kennedy School of Government, Harvard University

[‡]Paris School of Economics (CNRS) and CEPR.

[§]International Monetary Fund

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 desirable 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, for each possible pool. Under risk sharing within the pool, consumption growth in each country equals poolwide output growth. Therefore, the comparison quantifies the potential for 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.²

¹We 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.

²Pallage and Robe (2003) show that the welfare cost of economic fluctuations is far larger in developing countries than in advanced economies.

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 II provides the theoretical background to our empirical exercise—adapted directly from existing work. The section also outlines how we handle the combinatorial problem. Section III presents the general results on risk-sharing gains across subsets of countries. In Section IV, we estimate how much risk sharing gains are reduced when countries are constrained to share risk only with specific partners. Section V 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 refinements tend to augment the welfare implications of a given amount of risk sharing. As the same would presumably occur in our setup, our conclusions would only be strengthened.

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

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

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

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

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. Equation (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. Under autarky, $C_t^j = Y_t^j$, and time t welfare in country j is given by

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

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 $\bar{\sigma}^2$: $\bar{y}_t = \bar{y}_{t-1} + \bar{\mu} - \frac{1}{2}\bar{\sigma}^2 + \bar{\varepsilon}_t$.³ In the pool, time t welfare in country j is therefore given by

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

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 ,

³Lewis (2000) shows the sum of log-linear processes can be approximated by a log-linear process.

whose price is denoted by \bar{p}_t . Therefore, country j 's claim on poolwide output at time $t > 0$ is given by

$$\bar{C}_t^j = \frac{p_t^j}{\bar{p}_t} \bar{Y}_t \quad (5)$$

Suppose now 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 the agreement is subsequently honored, or to help provide relevant information about the pool's members.⁴ If the risk sharing agreement is contracted at time 0, this implies

$$\bar{C}_0^j = \frac{p_0^j - \tau_j}{\bar{p}_0} \bar{Y}_0 \quad (6)$$

The cost τ_j affects only the level of initial consumption in country j , at the time the pool is contracted: In particular, it does not affect the subsequent growth rates of consumption for all countries in the pool. They are all equal to the growth in \bar{Y}_t .⁵

Once it is in the pool, consumption in country j obeys the following Euler equation for all $t > 0$:

$$1 = \beta^\phi E_t \left[\bar{R}_{t+1}^{\phi-1} R_{t+1}^j \left(\frac{\bar{C}_{t+1}^j}{\bar{C}_t^j} \right)^{-\theta\phi} \right] \quad (7)$$

with $\phi = \frac{\gamma-1}{\theta-1}$, $\bar{R}_{t+1} = \frac{\bar{Y}_{t+1} + \bar{p}_{t+1}}{\bar{p}_t}$, and $R_{t+1}^j = \frac{Y_{t+1}^j + p_{t+1}^j}{p_{t+1}^j}$. Analogously, poolwide consumption is governed by the following Euler equation for all $t > 0$:

$$1 = \beta^\phi E_t \left[\bar{R}_{t+1}^\phi \left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\theta\phi} \right] \quad (8)$$

Equations (5), (7) and (8) constitute a system of differential equations in the asset prices \bar{p}_{t+1} and p_{t+1}^j . Since in the pool consumption grows at the constant rate $\bar{\mu}$, Lewis (2000) shows that prices are given by

$$p_t^j = Y_t^j \frac{\beta \bar{M}^{-\theta} H_j}{1 - \beta \bar{M}^{-\theta} H_j} \quad (9)$$

$$\bar{p}_t = \bar{Y}_t \frac{\beta \bar{M}^{1-\theta}}{1 - \beta \bar{M}^{1-\theta}} \quad (10)$$

where $H_j = \exp \left[\mu_j + \frac{1}{2} \gamma \bar{\sigma}^2 - \gamma \text{cov} \left(\varepsilon_t^j, \bar{\varepsilon}_t \right) \right]$.

⁴See for instance Brennan and Cao (1997), or Kang and Stulz (1997).

⁵We are grateful to an anonymous referee for this suggestion.

Consider now the welfare gains associated with the decision to form a pool of J countries at time 0. The percentage poolwide welfare gains associated with the formation of the pool are given by $W = \frac{\sum_{j=1}^J \bar{U}_0^j}{\sum_{j=1}^J U_t^j} - 1$. Since in autarky $C_0^j = Y_0^j$, combining equations (3), (4) and (6) gives

$$W \equiv W^* - T$$

where W^* is the percentage welfare gain associated with the same pool of size J in the absence of any cost τ_j , and $T = \frac{\frac{Y_0}{\bar{P}_0} (1 - \beta \bar{M}^{1-\theta})^{-1/(1-\theta)}}{\sum_{j=1}^J Y_0^j (1 - \beta M_j^{1-\theta})^{-1/(1-\theta)}} \sum_{j=1}^J \tau_j$. For sufficiently large costs T , it is clear that the welfare gains W can be negative, so that a pool of the corresponding size J is actually not optimal.

With asset prices defined in equations (9) and (10), it is possible to compute a value for W^* that only depends on the parameters of the model. In particular, we have

$$W^* = \frac{\frac{(1 - \beta \bar{M}^{1-\theta})^{-\theta/(1-\theta)}}{\bar{M}} \sum_{j=1}^J Y_0^j \frac{H_j}{1 - \beta \bar{M}^{-\theta} H_j}}{\sum_{j=1}^J Y_0^j (1 - \beta M_j^{1-\theta})^{-1/(1-\theta)}} - 1$$

In the rest of the paper, we compute W^* letting J vary for all possible subsets

of countries in our sample. All that is needed for this purpose are calibrated values of β , γ and θ , and estimates of M_j , H_j , Y_0^j and \bar{M} for all countries j in any considered pool of size J .

Importantly, the paper measures W^* , the *potential* percentage welfare gains associated with a given pool. There is no need to calibrate a value for τ_j , or for T . The costs' existence explains why full worldwide risk sharing may not be optimal, and thus justifies considering risk sharing over subsets of countries only. For legibility, the paper also presents the welfare gains W^* as a share of first best, worldwide gains.

How could risk sharing arrangements be limited to subsets of countries? First, selective capital controls vis-à-vis nonmembers could ensure that only the residents of countries in the pool have access to GDP indexed securities. Alternatively, GDP swaps could be introduced along the lines proposed by Merton (1990, 2002), either as a network of bilateral swaps, or as swaps intermediated by a central entity for the pool. Each period, each country uses swap contracts to pay the others the net difference between its current output and its share in poolwide output. Participation in the network of swaps defines the pool membership.

2.2 Implementation

This paper computes *potential* welfare gains W^* for subsets of countries of any size J . For each size J , we seek to report the grouping of specific countries that maximizes poolwide welfare (or, alternatively, that minimizes poolwide volatility).

Proceeding incrementally, we use three approaches to computing diversification gains. First, we focus on volatility reduction. We report the standard deviation of the growth rate for individual country consumption, σ_j^2 in the model, and compare it with the variance of the growth rate in poolwide GDP, $\bar{\sigma}^2$ in the model. We use consumption in autarky and output in the pool because autarky consumption reflects other insurance schemes that already exist in each country. This simple approach conveys most of the key economic intuition: $\bar{\sigma}^2$ falls quickly as J increases. Second, we compute W^* as implied by these volatility reductions, imposing at first that growth rates remain unchanged, i.e. $\mu_j = \bar{\mu}$, at a level that is calibrated. This follows exactly from Obstfeld (1994). As was the case there, welfare is a monotonic transformation of volatility. Third and finally, the assumption that growth rates are the same for all countries is relaxed, and we use historical data to estimate μ_j and $\bar{\mu}$ in any pool of size J .

Searching for pools of countries of any size, that yield minimal variance (or maximal welfare) is not straightforward, in light of the vast number of possible combinations of countries. We consider the $N = 74$ countries in our sample individually, then all of their possible combinations 2 countries at a time (of which there are C_2^N), then 3 at a time (given by C_3^N), and so on, where $C_J^N = \frac{N!}{J!(N-J)!}$. The total number of partitions is $\sum_{J=1}^N C_J^N = 2^N - 1$. It quickly reaches astronomical levels as N rises.

We implement a computational algorithm whose details are provided in a Technical Appendix available upon request. We are able to keep track of all possible combinations for any pool size J , for a sample containing up to 31 countries, i.e. 2.1×10^9 combinations. This algorithm can handle, for example, the universe of 26 emerging market countries—about 6.7×10^7 combinations. But when the universe consists of all 74 countries, the same algorithm only allows us to analyze all combinations of pools of size $J = 7$ or less ($C_7^{74} = 1.8 \times 10^9$). Since $C_J^N = C_{N-J}^N$, we can also draw the inventory of all combinations of $J = 67$ or more.

Outside of these pool sizes, we need to resort to an approximation algorithm. When $N = 74$ for instance, the total number of groups to consider increases to $2^{74} = 1.9 \times 10^{22}$, too large for existing computing power. For each group, one needs to sum GDP for all countries in the pool, to compute an aggregate growth rate, the corresponding standard deviation, and covariances. Even if each operation took a nanosecond to complete, running an exhaustive search over all possible pools amongst 74 countries would take hundreds of centuries.

For sample sizes where exhaustive inventories are out of reach, we implement recursive searches. We first run exhaustive searches up to the maximum pool size where it is feasible. This includes all pools of maximum size $J = 1, \dots, 7$ (or $J = 67, \dots, 74$) drawn from the universe of 74 countries. For each J , we order pools by decreasing volatility (or increasing welfare). Then, for each J , we save the best pool, but also the following S pools *that include each country in the sample*. In other words, when drawing the inventories of groups of size J , we collect the one group with lowest volatility, but also the next S groups that contain country 1, the next S groups that contain country 2, and so on. We therefore collect $S \times N$ additional pools for each J . We choose $S = 1,351$, so that $1,351 \times 74 \simeq 100,000$. We call these “seed” pools. We can collect seed pools no matter the value of J . But for $J \leq 7$ (or $J \geq 67$) we know the universe of *all* pools.

Then, for each pool size J , we isolate all groups that include the members of the optimal pool of size $J - 1$, plus one of the $N - J + 1$ remaining countries. Among these, we find the best pool of size J , as well as the best new $S \times N$ seed pools of size J . The procedure is iterated. There is a recursive aspect to the approach, but at each stage we consider the best S pools for each of the N countries: This gives plenty of opportunities for countries that are in the best pool of size $J - 1$ to drop out at the next increment.

We have verified the reliability of this approximation in three different ways. First, we ran exhaustive searches for all possible combinations for $J = 1, \dots, 7$ and $J = 67, \dots, 74$. We compared the groupings implied by these exhaustive inventories to the results of our approximation. They were always identical. Second, we have experimented with different values for S , as low as 2, and have found systematically the same results as with $S = 1,351$. Third, for each pool size J , we have compared the results implied by the approximation to large numbers of alternative pools, composed of countries drawn randomly. We have not found a single instance in which a pool drawn randomly was preferable to those identified as the best through the approximation procedure.⁶

⁶Combinatorial problems similar to those we are tackling are the object of a large literature in computer sciences. It revolves around the so-called “Traveling Salesman” problem, for which well-established approximated solution methods exist. To our knowledge however none can be applied to our baseline setup. For instance, Han, Ye and Zhang (2002) propose an approximation algorithm that can be applied to minimize the variance of a sum. But we minimize the variance of a *weighted* sum, where the weights themselves depend on the group’s membership. In Imbs and Mauro (2007), we use the Han, Ye and Zhang (2002) algorithm to identify risk diversification benefits for a given absolute size of the risk-sharing contract (for example, a US\$1 contract). That exercise involves an unweighted average of GDP growth rates. Our conclusions are virtually identical to what we get here with recursive searches.

3 Data and Results

This section begins with a description of the data sources. We then build intuition through a simple, single country example, and generalize it to the main results. We describe a “global envelope” of the groupings that achieve maximal risk-sharing gains for all sizes J . We first focus on volatility reduction holding growth constant, then infer welfare gains holding growth constant, and finally allow for growth rates to differ.

3.1 Data

Data on yearly real GDP and consumption are drawn from the World Bank’s World Development Indicators. They are evaluated in purchasing power parity (PPP) U.S dollars, for the period 1974–2004.⁷ The sample consists of 25 advanced, 26 emerging market, and 23 developing economies with complete coverage and data of reasonable quality. The full country list is provided in the Appendix. Advanced economies are defined following the International Monetary Fund’s World Economic Outlook. The remaining economies are considered emerging if they are included in either the stock-market-based International Financial Corporation’s Major Index (2005), or JPMorgan’s EMBI Global Index (2005), which includes countries that issue bonds on international markets. The rest are classified as developing.

We use these data to compute autarkic growth and volatility in consumption, and poolwide growth and volatility in GDP. The approach makes use of the variance-covariance matrix in GDP growth to infer expected gains arising from international risk sharing. A key assumption is that this matrix is relatively stable over time, which seems to be the case in the data. Doyle and Faust (2005), for instance, show that there is no evidence of significant changes in the correlation of output growth rates or other macroeconomic aggregates, despite claims that rising integration among the G-7 economies has increased cycles synchronization. And a large empirical literature has documented the cross-sectional properties of international business cycles, which appear to have highly persistent determinants, such as trade linkages or patterns of production (see Frankel and Rose, 1998 or Baxter and Kouparitsas, 2005).

3.2 A Simple Example

To develop intuition, we first report the results corresponding to the pools that minimize risk from the standpoint of an individual country, Chile. For each pool size J , Figure 1 plots the standard deviation of the growth rate of poolwide

⁷Previous studies (for example, Backus and Smith, 1993; and Ravn, 2001) have established that real exchange rate fluctuations worsen the case for international risk sharing. Indeed, GDP data at market exchange rates would imply far higher volatility—harder to hedge through international risk sharing. Using PPP-adjusted data amounts to assuming that all countries value the same consumption bundle in utility

GDP for the group of countries containing Chile chosen to minimize poolwide volatility. This is not a welfare measure, which explains its non-monotonicity with respect to J . The envelopes are displayed for various sample restrictions. There are four cases: pools drawn from the whole sample of 73 countries; pools with other emerging markets; with other emerging markets in Latin America; or with advanced economies. To give a sense of the importance of choosing well one's risk-sharing partners, a fifth line also plots the *maximum* value of $\bar{\sigma}$ for all pools containing Chile, and for all J .

Several results deserve mention. First, the lowest possible standard deviation for poolwide GDP growth in a group that includes Chile is 0.61%, far below 4.41% for Chile itself. The minimum obtains for a group of 20 countries. Second, a small number of carefully chosen partners is sufficient to yield the bulk of available diversification benefits. With just one well-chosen partner (France), poolwide standard deviation falls to 1.26%. For the best pool of seven members, the standard deviation of GDP growth reaches 0.72%, barely above the absolute minimum. Not surprisingly, this obtains for a motley set of economies: Austria, Cameroon, Chile, New Zealand, Nicaragua, Sweden, and Syria.

As will become apparent, the finding that most diversification gains are attained in relatively small pools holds in general. In the United States for instance, despite the large size of the U.S. economy, pooling with another five or six well-chosen economies implies a near halving of US volatility. This result is reminiscent of the well-known finding in finance that a small set of stocks is sufficient to provide most of the diversification opportunities available from a market portfolio (Solnik, 1974). Marginal diversification gains quickly become small, and even negative for $J > 30$. Beyond a certain pool size, hedging opportunities are exhausted, and the pool starts including countries with high volatilities, and positive covariances. A contrario, the (upper) envelope that traces the worst possible pools of each size highlights the importance of choosing one's partners carefully. For small J , a poorly chosen pool can deliver higher volatility than in autarky.

Various types of (economically relevant) constraints reduce maximum diversification benefits. For example, Figure 1 reports the extent to which potential gains decline when the universe of countries is constrained. The lowest attainable standard deviation is 0.61% for unconstrained pooling in Chile, but 0.87% when Chile must pool with advanced countries only, 1.07% within the universe of emerging markets, and 1.91% when pooling within Latin America. Risk-sharing agreements that restrict the heterogeneity of the pool membership have substantial consequences on the potential diversification gains.

3.3 The Global Envelope

We now generalize the approach. Pools with minimum volatility are not constrained anymore to include any given country. Figure 2 reports the envelope

of minimal output volatility for all pool sizes J using the recursive approach described in section 2.2. The Figure reports minimal poolwide volatilities across four different samples: pools of up to 74 countries, which traces a “global envelope”, and pools within developing, emerging or advanced economies. Across all four samples, the bulk of possible diversification gains is attained with relatively small pools. The global envelope implies the lowest possible poolwide volatility is 0.50%, and it is obtained in a pool of 17 countries. But volatility is already as low as 0.62% for $J = 7$ along the global envelope. Thus, diversification gains continue to be achieved within groups consisting of a small number of countries in this general setup.⁸

The list of countries involved in minimum-volatility pools confirms that heterogeneity is key. Interestingly, the list overlaps with that obtained for Chile, which suggests that the sample of countries providing the best hedging properties within a universe of 74 economies is relatively small and robust. The variance-covariance matrix of GDP growth rates contains a few countries with systematically negative off-diagonal elements, i.e. desirable hedging properties.

Figure 2 also reports minimum volatilities for three sub-samples: advanced, emerging, and developing countries. Diversification gains remain substantial within each sub-sample, but in all cases the minimum levels of poolwide volatility are substantially higher than in the unconstrained sample. Advanced economies achieve smaller gains, consistent with their lower volatility and internationally correlated business cycles. All four envelopes display the same non-monotonicity as in Figure 1, with high marginal diversification gains for $J < 10$. They turn negative for $J > 15$, when countries with poor hedging properties start being included.

Figures 1 and 2 help quantify the diversification gains that could accrue to the representative consumer living in a pool of a given size. It captures the decrease in output volatility as pool size increases. This is an input in the welfare gains from pooling risk, which we next evaluate.

3.4 Welfare Gains

This sub-section turns to computing welfare W^* . This requires calibrating the preference parameters β , γ and θ , and estimating M_j , H_j , Y_0^j and \bar{M} for all countries j in any considered pool of size J . The subjective discount rate β is set at 0.95, $\theta = 2$ and $\gamma = 5$. These are later subjected to robustness checks. For each country and each pool, M_j is computed using consumption data, but \bar{M} , H_j , and Y_0^j are computed using GDP: In autarky, consumption data are used

⁸The value reported for $J = 1$ corresponds to the standard deviation of the individual GDP growth rate for the least volatile country during the sample period, namely France. Diversification gains for specific countries cannot be easily read off the figure, because the identities of countries involved in pools of different sizes change.

because they reflect existing sources of insurance, for instance inter-temporal. But in the pool, it is GDP data that must be used, because they reflect the potential for risk sharing of a given pool of size J .

Growth rates are still constrained to be equal across countries, $\mu_j = \bar{\mu} = 3\%$. With this constraint, U_t^j and \bar{U}_t^j are never close to zero, and welfare is well behaved: This will no longer be the case any longer when growth rates are allowed to vary by country. Figure 3 reports the highest value of W^* for every pool size J . In the Figure, W^* is expressed in percentage of the first best welfare gains obtained with full risk sharing worldwide. The results are reported once again for four different samples. Welfare gains increase monotonically with J , as they should, with a maximum for $J = N$. Just as for volatility, the marginal increases in W^* are largest for $J < 10$. Marginal gains begin to peter out as soon as $J > 7$ or 8. This holds true for the whole sample, but also across the four sub-groups of countries considered, among advanced, emerging, or developing countries. Interestingly, the gains are largest for emerging markets, and smallest for developing countries. This reflects the relative heterogeneity in business cycles in the former group.

Table 1 reports summary statistics for the key variables serving as inputs to the welfare computations. The Table reports the extrema, mean and median values for σ_j , μ_j , $cov(\varepsilon_t^j, \bar{\varepsilon}_t)$, and p_t^j / \bar{p}_t , for the universe of countries in our dataset, $J = 74$. The volatility in GDP growth ranges from 0.01% to 4.32%, with median 0.20 and mean 0.38. The distribution of country volatilities is slightly skewed, with a few countries displaying exceptionally volatile business cycles. These are the countries for which risk insurance is particularly desirable. GDP growth is distributed symmetrically, ranging from 1.18 to 8.35%, with mean and median values close to each other, around 3.5%. The distribution of the covariances between idiosyncratic and poolwide growth rates is heavily skewed, with minimum -0.026 and maximum 0.048 . There are very few countries with negative covariances, i.e., very few countries that represent especially good hedges worldwide. Finally, the relative prices p_t^j / \bar{p}_t are heavily skewed, with a median percentage share of 3.22, but a mean equal to 1.34%. This reflects that most countries have very small claims to the worldwide mutual fund, while very few have large claim, up to 26.15% for $J = 74$. The distribution of welfare gains is therefore highly unequal.

Table 2 illustrates the composition of W^* for various pools. It does so reporting individual countries' welfare gains upon entry in the pool, δ_j , defined as:

$$U_0^j [C_0^j (1 + \delta_j), \mu_j, \sigma_j^2] = U_0^j [\bar{C}_0, \bar{\mu}, \bar{\sigma}^2] \quad (4)$$

which can be computed for all countries j in a given pool, under the same parametrization as the one used to estimate W^* .

The Table reports the minimum value of δ_j , its median, and the value of W^* for four specific pools. The first pool is formed by the entire 74-country

sample: there, δ_j equals 1.9% of permanent poolwide consumption on average (unreported). But the distribution of δ_j across the membership is skewed, with a minimum for the US at 0.50%, and a median for Hong Kong at 6.21%. The welfare consequences of pools are heterogeneous across their membership, reflecting the very heterogeneity of the constituent countries.

Individual country's welfare gains decrease with per capita income, because they increase with autarkic output volatility. Among advanced economies, median δ_j equals 0.89%, whereas it is 4.97% in emerging markets and 7.17% in developing countries. The distribution of δ_j is also less dispersed in homogeneous pools: for advanced countries, its minimum value is 0.40%, and its median 0.89%. Contrast this with the dispersion of δ_j among developing countries, with a minimum of 1.00%, and a median of 7.17%. The heterogeneity in output volatility is much more pronounced there.

How important is carefully chosen diversification in reaping the benefits described in Figure 3? How do the optimal pools described in the Figure differ from alternative groupings of countries, drawn randomly? In Figure 4 we report the welfare gains implied by pools drawn at random. We consider random draws of 10,000 pools for each value of J . The Figure reports the maximum value of W^* , normalized by the first best welfare gains, as implied by the 10,000 draws, along with its 99th and 95th percentiles. For ease of comparability, it also reproduces the recursive maximum reported in Figure 3. The random maximum displays a trend increase in J , with some convexity. As J increases, the gains fluctuate quite considerably around the trend, reflecting the importance of a few countries in delivering the smooth welfare gains from Figure 3. But the 99th and 95th percentiles lose any similitude with the recursive measure: they are almost linear in J , and miss most of the convexity apparent from previous Figures. In other words, while high values of W^* can be reached for low values of J , this only happens in extremely rare groupings of countries. Diversification gains do not accrue quickly between countries selected at random.

The welfare gains from risk insurance obviously depend on the calibration of preferences, and more specifically on the intertemporal elasticity of substitution. Figures 3 and 4 correspond to specific values of the rate of time preference $\beta = 0.95$, the coefficient of risk aversion $\gamma = 5$, and $\theta = 2$, which corresponds to a relatively low elasticity of intertemporal substitution of 0.5. Unsurprisingly, the welfare gains of international risk sharing depend on all parameter values: they increase in risk aversion, fall in the intertemporal elasticity, and fall in the rate of time preference. But varying these parameters has symmetric effects on the first best, worldwide welfare gains, and on the poolwide welfare gains. Since our measures of W^* are always normalized, the proportional increases we document in Figures 3 and 4 are in fact invariant to alternative calibrations of these three important parameters.⁹

⁹These results are available upon request.

3.5 Pooling Growth Rates

The previous section computes W^* holding growth rates constant. In principle, countries with relatively high expected growth rates should be able to obtain a higher share of poolwide consumption, with higher values for p_0^j . In practice however, the challenges involved in predicting growth rates make it difficult to incorporate differences in expected growth into risk-sharing contracts. For instance, country rankings with respect to growth rates change dramatically from one decade to the next, as shown by Easterly and others (1993). To estimate expected economic growth, this section simply considers the average of historical growth rates over the entire period under consideration. We also assume that individual countries' growth rates are unaffected by risk sharing. We estimate μ_j as the 1975-2004 average of GDP growth in country j , and $\bar{\mu}$ as the 1975-2004 average of poolwide GDP growth, for any pool of size J . A possible concern might be that lower volatility in a pool may create lower mean growth. But this seems unlikely in light of the evidence that lower-volatility countries tend to have relatively high mean growth (Ramey and Ramey, 1995). So our estimates of W^* on the basis of historical growth rates represent if anything a lower bound.¹⁰

Figure 5 reproduces Figure 3, using historical values for μ_j and $\bar{\mu}$. It confirms the broad pattern of our results: Most marginal gains in W^* accrue for $J < 10$. More than 80% of maximum W^* are reached in pools of 10 countries or fewer. In all samples the shape of the envelopes remains the same.¹¹

4 Pooling Risk Within Sub-Samples

The previous section establishes substantial welfare gains are accessible to small pools of countries. Yet risk sharing agreements, or GDP contingent securities,

¹⁰An issue with using historical growth rates is the possibility of zero utility in the definition of U_t^j . We ensure this does not happen by constraining $\beta M_j^{1-\theta}$ away from 1. In particular, we drop combinations of parameters that result in $0.98 < \beta M_j^{1-\theta} < 1.2$. With these constraints, there are no instances of unbounded utility, or of non-positive prices for any value of $cov(\varepsilon_t^j, \bar{\varepsilon}_t)$. This reduces the sample to a maximum of 59 countries, listed in the Appendix, where 3 advanced countries, 7 emerging markets, and 5 developing economies were dropped. There is no reason to expect the dropped countries would be precisely the ones that invalidate the convexity of the welfare function apparent in Figure 5.

¹¹It is not apparent from Figure 5, which is normalized, but when growth rates are country-specific, the levels of W^* shift up for all J - and so do the gains from worldwide risk sharing. Such a systematic rise is reminiscent of Obstfeld (1994), who argued international financial integration gives access to high returns, which can have large welfare consequences. In unreported work, we obtained poolwide weighted averages of μ_j , computed for all J . We compared them with $\bar{\mu}$, once again computed for all pools. We found little difference amongst developing countries, some improvement for advanced economies, and largest increases for emerging markets. Historical growth rates are such that it is among emerging markets that growth increases would be most pronounced for countries participating in a pool. As a result, the increase in W^* is largest in this sample.

are rarely observed. The reasons for this are not clear (see Borensztein and Mauro, 2004, for a review). Some authors have suggested that this absence may be due, for example, to difficulties in acquiring information about other countries or in enforcing international contracts, perhaps because of weak institutions. Such difficulties are summarized by T in Section 2.1; it stands to reason that the costs implied by such issues should depend on the size and the composition of a pool of risk-sharing countries.

To further explore these issues, we create samples determined by institutional quality or regional proximity, which is often synonymous with strong trade linkages. Goods trade creates incentives to honor international commitments towards trade partners - often a neighbor in a given region. We compute W^* for pools of countries that are constrained to belong to the same category, e.g. with high institutional quality.

Then existing regional agreements are considered. These include free-trade agreements, such as the European Union, the ASEAN or Mercosur. We also include pools that were built to provide some degree of cross-country insurance (but clearly fall short of complete risk sharing). We consider the Chiang-Mai Initiative and the Latin American Reserve Fund (FLAR). We illustrate the reduction in consumption volatility that would occur under complete sharing of GDP risk in the various samples considered. Potential welfare gains W^* are computed in each case, then compared with the type of welfare gains that could be attained in a pool of identical size, but whose membership would be entirely unconstrained.

4.1 Institutional Quality

We explore the effects of restricting the sample on the basis of default history and scores on the measures of institutional quality that capture contract enforcement. Two definitions are considered. The first, labeled “excellent enforceability” includes all countries that were in the top half of the distribution of the institutional quality index compiled by Kaufmann et al (2005), and that never experienced severe international repayment difficulties between 1970 and 2004. The second, “above-average institutional quality” is based on the institutional quality index only. In addition to advanced countries, the former sample includes four emerging markets and developing countries, whereas the latter includes eight emerging markets and three developing countries.

For each sub-sample, Table 3 reports the median value of aurtakic consumption volatility σ_j across the countries j in the sample. Column (2) reports the minimum variance $\bar{\sigma}$ that can be obtained in a pool formed by countries in the indicated sample, for any pool size. Column (3) reports the median δ_j for each indicated sample, and column (4) reports the corresponding value for W^* . The latter two measures are computed assuming $\mu_j = \bar{\mu} = 3\%$.

Column (1) confirms output volatility is higher in countries with poor institutions. It is 5.24% with low enforceability vs. 2.25% with excellent enforceability, and 5.45% with institutional quality below average vs. 3.04% if it is above average. Column (2) suggests the diversification gains from volatility reduction are larger in groups with poor institutions: a decrease of around 4 percentage points for such groups, as against 1.5 percent for pools with good institutions. This translates in drastically larger welfare gains for countries with poor institutions in column (3). The last column in Table 3 reports estimates of W^* . Countries with institutions that ensure excellent contract enforceability can only achieve 32% of first-best worldwide risk-sharing gains. But the complementary set of countries, below excellent enforceability, can achieve more than 60% of first-best welfare gains.

The lower panel of Table 3 reports the same statistics, for pools selected on the basis of their regional proximity. We consider three economic zones: the European Union, Asian emerging countries, and Latin American emerging economies. Geographical constraints turn out to be very important for advanced European countries. With a median volatility σ_j of 1.91%, individual country's welfare gains δ_j are well below 1%. As a result, these pools can hope to achieve at best 10.15% of first-best welfare gains. The result presumably reflects high cycle synchronization between rich economies in general, with little diversification to be gained.

The same is not true of emerging markets, with median volatilities σ_j of 3.93% in Asia and 5.14% in Latin America. Column (2) illustrates that pooled volatility $\bar{\sigma}$ can be reduced to just below 2%. For both groups, individual country's median welfare gains are between 3% and 5% of permanent consumption, and W^* doubles to around 20% of first-best welfare. In other words, pooling risk among homogeneous European countries carries little gain. Potential gains are substantially larger in (relatively heterogeneous) emerging Asia or emerging Latin America. But either still fall considerably short of what is possible within emerging markets as a whole, or within countries with poor institutions in general.

4.2 Existing Groupings

Finally, we consider the potential welfare gains that would arise if complete risk-sharing were achieved within groups of countries that already exist for other purposes. We include free-trade agreements, whose participants have long-established cooperation, as well as arrangements that involve an element of macroeconomic risk-sharing, such as the Chiang-Mai Initiative and the FLAR. At present, none of these groups explicitly seek to share GDP risk. But we consider them interesting, because their previous history of cooperation and trade and financial relations would presumably make it easier for these groups

to attain the information and mutual trust that would be required to move to complete risk sharing.¹²

Table 4 reports the results for several such agreements. Row (1) reports the median value of country-level autarkic volatility σ_j for each group. Row (2) is the value of each pool's diversified volatility, $\bar{\sigma}_{Pool}$, which is then compared with row (3), the minimum value of $\bar{\sigma}$, $\bar{\sigma}_{Min}$, that can be achieved in *any, unconstrained* pool of similar size J . Row (4) reports the value of W^* corresponding to the volatility reduction implied by the existing pool, $\bar{\sigma}_{Pool}$, assuming $\mu_j = \bar{\mu} = 3\%$.

Median volatility is highest in groupings that involve emerging markets, in Asia (APEC, ASEAN, Chiang Mai), or in Latin America (FLAR), at around 4.5%. It is much smaller in Europe, at 1.19% in the EMU, or 1.91% in the EU. Poolwide volatilities decrease substantially for the APEC, Chiang Mai, and the FLAR, with values between 1.29 and 2.50%. In contrast, volatility reduction is minimal in Europe. For instance, in the EMU poolwide volatility is equal to its median value across the Union, so that risk sharing would lower consumption volatility for half the membership only. This happens because of the homogeneity in these countries business cycles, and their relatively low volatility.

The values of W^* mirror these volatility reductions, along with the relative sizes of each agreement. ASEAN and FLAR only involve five and seven countries, respectively. With relatively small reduction in volatility, the values of W^* are around five percent of the potential first-best welfare gains. EMU, EU, and APEC deliver much larger values of W^* , around 30 percent of first-best gains. This reflects large values of J , equal to 12, 18, and 16, respectively. These values are consistent with the convexity in Figure 3. The Chiang-Mai grouping is interesting for a different reason: there are nine participating countries, but W^* is almost equal to 15 percent of first-best gains. This must be because of the large volatility reduction afforded by this country grouping.

To isolate the relative importance of pool size vs. volatility reduction, the third row in Table 4 reports the *potential* volatility reductions that would be possible in carefully chosen pools of identical sizes, drawn from the full sample of 74 countries. Interestingly, these alternative groupings always deliver substantially larger diversification gains in row (3) than in row (2), including for the EMU. Of course, these pools are constituted of countries with much more heterogeneous business cycles than the groupings in the headings of Table 4. That such small groups of countries can deliver large diversification gains, provided the membership be chosen carefully, is the paper's main point.

¹²For free-trade agreements, on the one hand, trade linkages imply a mutual interest in each other's economic performance and in honoring of international obligations; on the other hand, trade partners are well known to have synchronized cycles, thus reducing risk-diversification benefits (see, for example, Frankel and Rose, 1998.)

5 Conclusion

Despite a trend toward greater international financial integration, the full gains from complete international risk sharing have hitherto remained elusive. In this paper, we have shown that the bulk of these gains could effectively be reaped in carefully chosen groupings of a small number of economies, often less than ten. This makes it even harder to understand why risk sharing does not arise endogenously. We conjecture it does not because of difficulties in contracting risk sharing agreements between countries. Where such difficulties are likely to be minimal, between countries with good quality institutions, or established trade partners, we show that the potential welfare gains are minimal. In contrast, those small pools that deliver large welfare gains are constituted of economies that are often far away from each other, and include countries with weak institutions or a history of default. Thus, a possible interpretation of our empirical results is that complete risk sharing arrangements have remained elusive because the potential welfare gains may be insufficient to offset obstacles such as concerns about the difficulty of enforcing international contracts.

Appendix Table 1: Country Samples

Advanced Economies [25]	Emerging Markets [26]	Developing Countries [23]	Advanced Europe [18]	Emerging Market Latin America [11]	Emerging Market Asia [8]	Excellent Enforceability [29]	Above Average Institutional Quality [37]
Australia	Argentina*	Algeria*	Austria	Argentina	China	Australia	Australia
Austria	Brazil	Bangladesh	Belgium	Brazil	India	Austria	Austria
Belgium	Chile	Benin	Denmark	Chile	Indonesia	Belgium	Belgium
Canada	China	Bolivia	Finland	Colombia	Korea	Botswana	Botswana
Hong Kong SAR	Colombia	Botswana	France	Dom. Rep.	Malaysia	Canada	Brazil
Denmark*	Cote d'Ivoire*	Cameroon	Germany	Ecuador	Pakistan	Denmark	Canada
Finland	Dom. Rep.	Congo Rep.	Greece	El Salvador	Philippines	Finland	Chile
France	Ecuador	Costa Rica	Iceland	Mexico	Thailand	France	Costa Rica
Germany	Egypt	Gabon	Ireland	Peru		Germany	Denmark
Greece	El Salvador*	Gambia	Italy	Uruguay		Greece	Finland
Iceland	Hungary	Ghana	Luxembourg	Venezuela		Hong Kong	France
Ireland	India	Guatemala	Netherlands			Hungary	Germany
Italy	Indonesia	Kenya	Norway			Iceland	Greece
Japan	Korea	Lesotho	Portugal			Ireland	Hong Kong SAR
Luxembourg	Malaysia	Madagascar*	Spain			Italy	Hungary
Netherlands	Mexico	Malawi	Sweden			Japan	Iceland
New Zealand*	Morocco	Nicaragua	Switzerland			Luxembourg	Ireland
Norway	Pakistan	Paraguay	United Kingdom			Malaysia	Italy
Portugal	Peru*	Rwanda*				Netherlands	Japan
Singapore	Philippines	Senegal				New Zealand	Korea
Spain	South Africa	Syria				Norway	Luxembourg
Sweden	Thailand	Togo*				Portugal	Malaysia
Switzerland*	Tunisia	Trin. and Tob.*				Singapore	Morocco
United Kingdom	Uruguay*					South Africa	Netherlands
United States	Venezuela*					Spain	New Zealand
	Zimbabwe*					Sweden	Norway
						Switzerland	Portugal
						United Kingdom	Singapore
						United States	South Africa
							Spain
							Sweden
							Switzerland
							Thailand
							Trinidad and Tobago
							United Kingdom
							United States
							Uruguay

Notes: Advanced countries are defined as in the International Monetary Fund's World Economic Outlook. The remaining countries are emerging if they are included in either the stock-market based International Financial Corporation's Major Index (2005) or JPMorgan's EMBI Global Index (2005), which includes countries that issue bonds on international markets. The remaining countries are classified as developing. Above average institutional quality is according to the index of Kaufmann, Kraay, and Mastruzzi (2005). Excellent enforceability is defined as above average institutional quality and no defaults on international debt in 1970-2004 according to Detragiache and Spilimbergo (2001) and Reinhart, Rogoff and Savastano (2003). GDP data in current U.S. dollars are from the IMF's World Economic Outlook.) GDP data at PPP are from the World Bank's World Development Indicators. Stars denote countries that are excluded to preserve stationarity in Section 3.5.

Appendix Table 2: Additional Country Samples

APEC [16]	ASEAN [5]	CHIANG MAI [9]	EMU [12]	EU [18]	FLAR [7]
Australia	Indonesia	China	Austria	Austria	Bolivia
Canada	Malaysia	Hong Kong	Belgium	Belgium	Colombia
Chile	Philippines	Indonesia	France	Denmark	CostaRica
China	Singapore	Japan	Germany	Finland	Ecuador
Hong Kong	Thailand	Korea	Greece	France	Peru
Indonesia		Malaysia	Ireland	Germany	Uruguay
Japan		Philippines	Italy	Greece	Venezuela
Korea		Singapore	Luxembourg	Iceland	
Malaysia		Thailand	Netherlands	Ireland	
Mexico			Portugal	Italy	
New Zealand			Spain	Luxembourg	
Peru			Switzerland	Netherlands	
Philippines				Norway	
Singapore				Portugal	
Thailand				Spain	
United States				Sweden	
				Switzerland	
				United Kingdom	

References

- Backus, David K., and Gregor W. Smith, 1993, "Consumption and real exchange rates in dynamic economies with non-traded goods," *Journal of International Economics*, November, Vol. 35, No. 3-4, pp. 297–316.
- Baxter, Marianne, and Michael A. Kouparitsas, 2005, "Determinants of business cycle comovement: a robust analysis," *Journal of Monetary Economics*, January, Vol. 52, No. 1, pp. 113–157.
- Borensztein, Eduardo R., and Paolo Mauro, 2004, "The Case for GDP-Indexed Bonds," *Economic Policy*, April, Vol. 19, No. 38, pp. 165–216.
- Brennan, Michael and H. Henry Cao, 1997, "International Portfolio Investment Flows," *Journal of Finance*, December, Vol. 52, No. 5, pp. 1851–1880.
- Doyle, Brian and Jon Faust, 2005, "Breaks in the Variability and Co-Movement of G-7 Economic Growth", *Review of Economics and Statistics*, Vol. 87, No. 4, pp. 721–740.
- Easterly, William, Michael Kremer, Lant Pritchett, and Lawrence H. Summers, 1993, "Good policy or good luck? Country growth performance and temporary shocks," *Journal of Monetary Economics*, Vol. 32, pp. 459–483.
- Eichengreen, Barry, 2007, "Insurance Underwriter or Financial Development Fund: What Role for Reserve Pooling in Latin America?" *Open Economies Review*, Vol. 18, No. 1, pp. 27–52.
- Eichengreen, Barry, and Charles Wyplosz, 1993, "The Unstable EMS," *Brookings Papers on Economic Activity*, No. 1, pp. 51–143.
- Epstein, Larry and Stanley Zin, 1989, "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", *Econometrica*, Vol. 57, No. 4, pp. 937–969.
- Frankel, Jeffrey A., and Andrew K. Rose, 1998, "The Endogeneity of the Optimum Currency Area Criteria," *Economic Journal*, July, Vol. 108, No. 449, pp. 1009–25.
- Han, Qiaoming, Yinyu Ye, and Jiawei Zhang, 2002, "An Improved Rounding Method and Semidefinite Programming Relaxation for Graph Partition," *Mathematical Programming*, Vol. 92, No. 3, pp. 509–535.
- Imbs, Jean and Paolo Mauro, 2007, "Pooling Risk Among Countries", IMF Working Paper 07/132.

Kang, Jun-Koo, and Rene Stulz, 1997, "Why is there a home bias? An analysis of foreign portfolio equity ownership in Japan," *Journal of Financial Economics*, 46(1), pp. 3-28.

Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi, 2005, "Governance Matters IV: Governance Indicators for 1996–2004," The World Bank, <http://www.worldbank.org/wbi/governance/govdata/>.

Lewis, Karen K., 2000, "Why Do Stocks and Consumption Imply Such Different Gains from International Risk Sharing?" *Journal of International Economics*, Vol. 52, pp. 1–35.

Lucas, Robert E., Jr., 1987, *Models of Business Cycles* (Oxford: Blackwell Publishers).

Merton, Robert C., 1990, "The Financial System and Economic Performance," *Journal of Financial Services Research*, Vol. 4, No. 4, pp. 263–300.

Merton, Robert C., 2002, "Future Possibilities in Finance Theory and Finance Practice." In *Mathematical Finance - Bachelier Congress 2000*, edited by H. Geman, D. Madan, S. Pliska, and T. Vorst. Berlin: Springer-Verlag.

Obstfeld, Maurice, 1994, "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability," *European Economic Review*, Vol. 38, No. 7, pp. 1471–86.

Pallage, Stephane, and Michel Robe, 2003, "On the Welfare Cost of Economic Fluctuations in Developing Countries," *International Economic Review*, Vol. 44, No. 2, pp. 677–698.

Park, Yung Chul, and Yunjong Wang, 2005, "The Chiang Mai Initiative and Beyond," *The World Economy*, Vol. 28, No. 1, pp. 91–101.

Ramey, Garey, and Valerie A. Ramey, 1995, "Cross-Country Evidence on the Link between Volatility and Growth," *American Economic Review*, December, Vol. 85, No. 5, pp. 1138–51.

Ravn, Morten O., 2001, "Consumption Dynamics and Real Exchange Rate," CEPR Discussion Papers 2940.

Shiller, Robert J., 1993, *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks* (Oxford: Oxford University Press, Clarendon Series).

Solnik, Bruno, 1974, "Why Not Diversify Internationally Rather Than Domestically?" *Financial Analysts Journal*, pp. 48–54.

Table 1: Summary Statistics for J=74

	Minimum	Mean	Median	Maximum
Volatility σ_j	0.016	0.381	0.202	4.326
Growth μ_j	1.179	3.552	3.349	8.355
Covariance $cov(\epsilon_j, \bar{\epsilon})$	-0.0262	0.0030	0.0024	0.0479
Price p_j/\bar{p}	0.0046	1.345	3.229	26.150

Notes: Summary statistics for the universe of 74 countries. $cov(\epsilon_j, \bar{\epsilon})$ denotes the covariance between individual countries and poolwide growth rates. The price ratio p_j/\bar{p} is defined in equations (9) and (10). All values are in percentages. Real GDP and consumption data are from the World Bank's World Development Indicators.

Table 2: Welfare Gains

	Minimum δ_i	Median δ_j	W* (Global risk-sharing = 100)
Full Sample	0.55 United States	4.10 Dominican Republic	100.00
Advanced Countries	0.40 United States	0.89 Australia	28.33
Emerging Markets	1.67 Colombia	4.97 Malaysia	57.51
Developing Countries	1.00 Bangladesh	7.17 Madagascar	6.58

Notes: Minimum, median, and total welfare gains for sub-samples indicated, assuming the same expected growth (3 percent) across countries. The results assume $\theta=2$, $\gamma=5$, and $\beta=0.95$. Country lists are provided in the Appendix. GDP data are from the World Bank's World Development Indicators.

Table 3: Gains from Risk Pooling Among Countries

	(1)	(2)	(3)	(4)
	Median σ_j individual country	Minimum poolwide variance	Median δ_j	W^* (% worldwide risk-sharing)
All countries (pooling with any country)	4.49	0.54	6.21	100.00
Costs of weak enforcement				
Excellent enforceability countries (pooling only with excellent enforceability countries)	2.25	0.87	1.73	32.21
Below excellent enforceability countries (pooling only with below excellent enforceability countries)	5.24	1.30	8.77	60.89
Above average inst. countries (pooling with only above average inst. countries)	3.04	0.82	2.65	44.67
Below average inst. countries (pooling with only below average inst. countries)	5.45	1.50	9.96	48.64
Costs of regional constraints				
European Union (pooling only with EU)	1.91	1.05	0.68	10.15
Asian emerging (pooling only with Asian emerging)	3.93	1.84	2.98	18.18
Latin American emerging (pooling only with Latin American emerging)	5.14	1.90	4.96	21.56

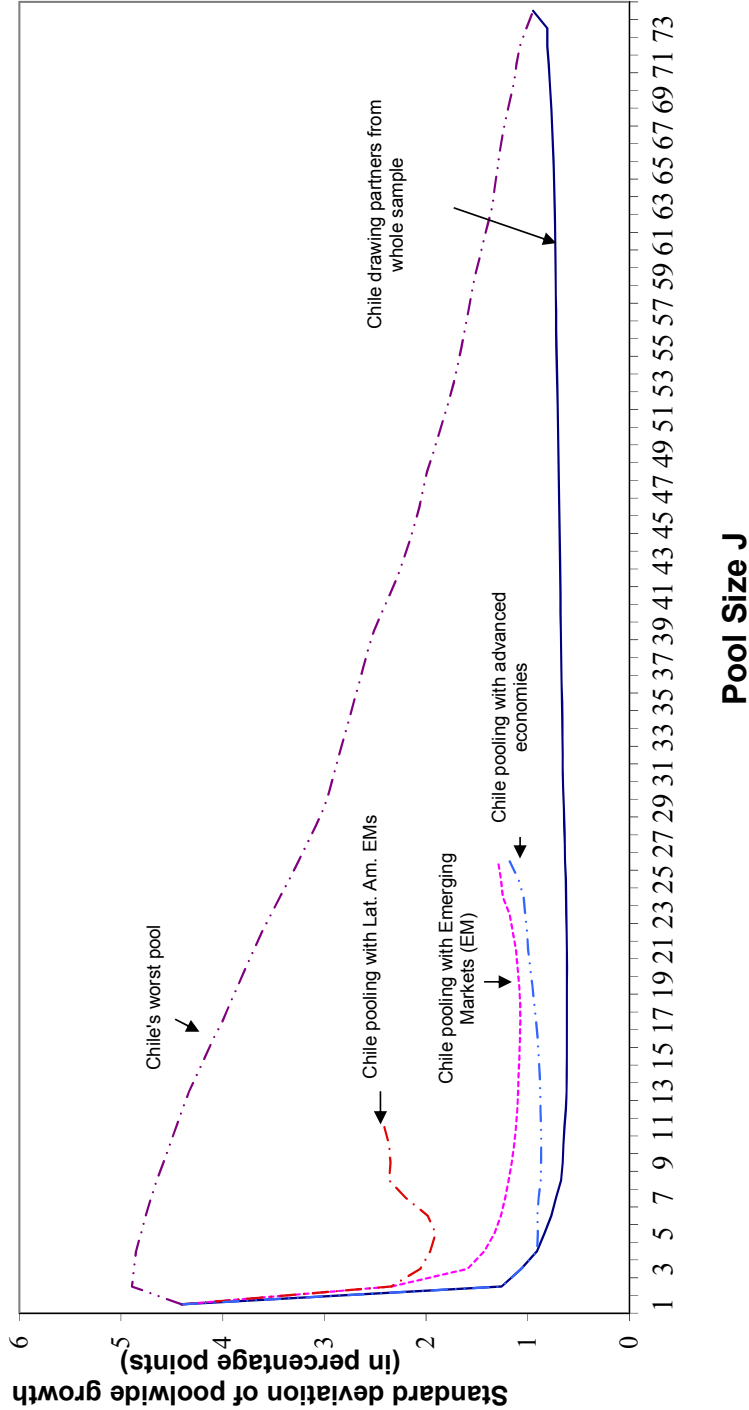
Notes: Column (1) reports the median of σ_j (across countries in the indicated sub-sample). Column (2) reports the median of the lowest possible standard deviation of poolwide growth (across countries in the indicated sub-sample). Column (3) reports the median δ_j (across countries in the indicated sub-sample) assuming growth rates fixed at 3 percent per annum. Column (4) reports W^* , assuming growth fixed at three percent per annum.

Table 4. Poolwide Volatility and Welfare for Selected Groups

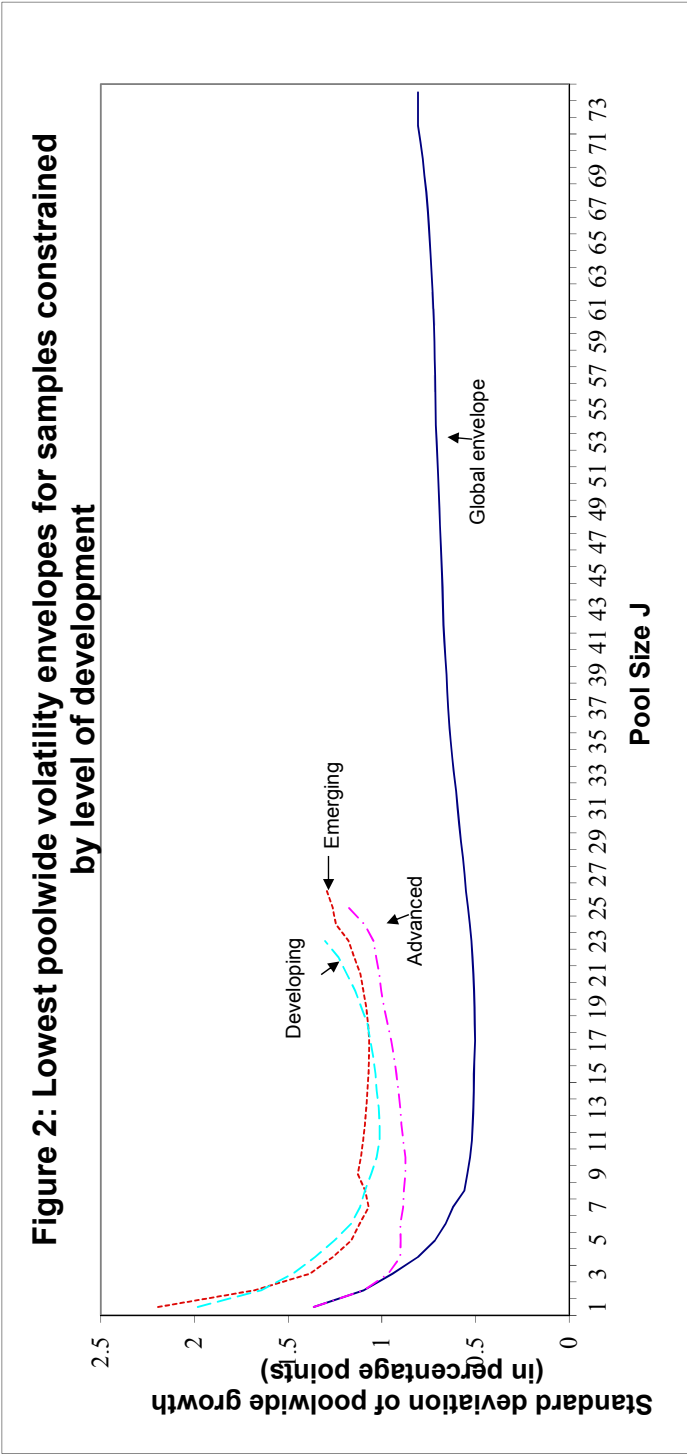
	APEC	ASEAN	CHIANG MAI	EMU	EU	FLAR
Median σ_j	4.44	4.53	4.48	1.19	1.91	4.62
$\bar{\sigma}_{Pool}$	1.29	3.47	1.40	1.19	1.09	2.50
$\bar{\sigma}_{Min}$	0.68	0.86	0.79	0.66	0.65	0.71
W^* (worldwide = 100)	29.05	5.25	14.50	29.71	32.17	5.67

Notes: Median country volatility σ_j , poolwide volatility $\bar{\sigma}_{Pool}$, and minimum poolwide volatility $\bar{\sigma}_{Min}$ that can be achieved in any, unconstrained pool with the same number of countries J . The welfare gains W^* use $\bar{\sigma}_{Pool}$. Growth is set at 3 percent, $\theta = 2$, $\gamma = 5$, and $\beta = 0.95$. The list of countries in each sample is in the Appendix. GDP data are from the World Bank's World Development Indicators.

**Figure 1. Benefits of diversification under various restrictions:
The case of Chile**

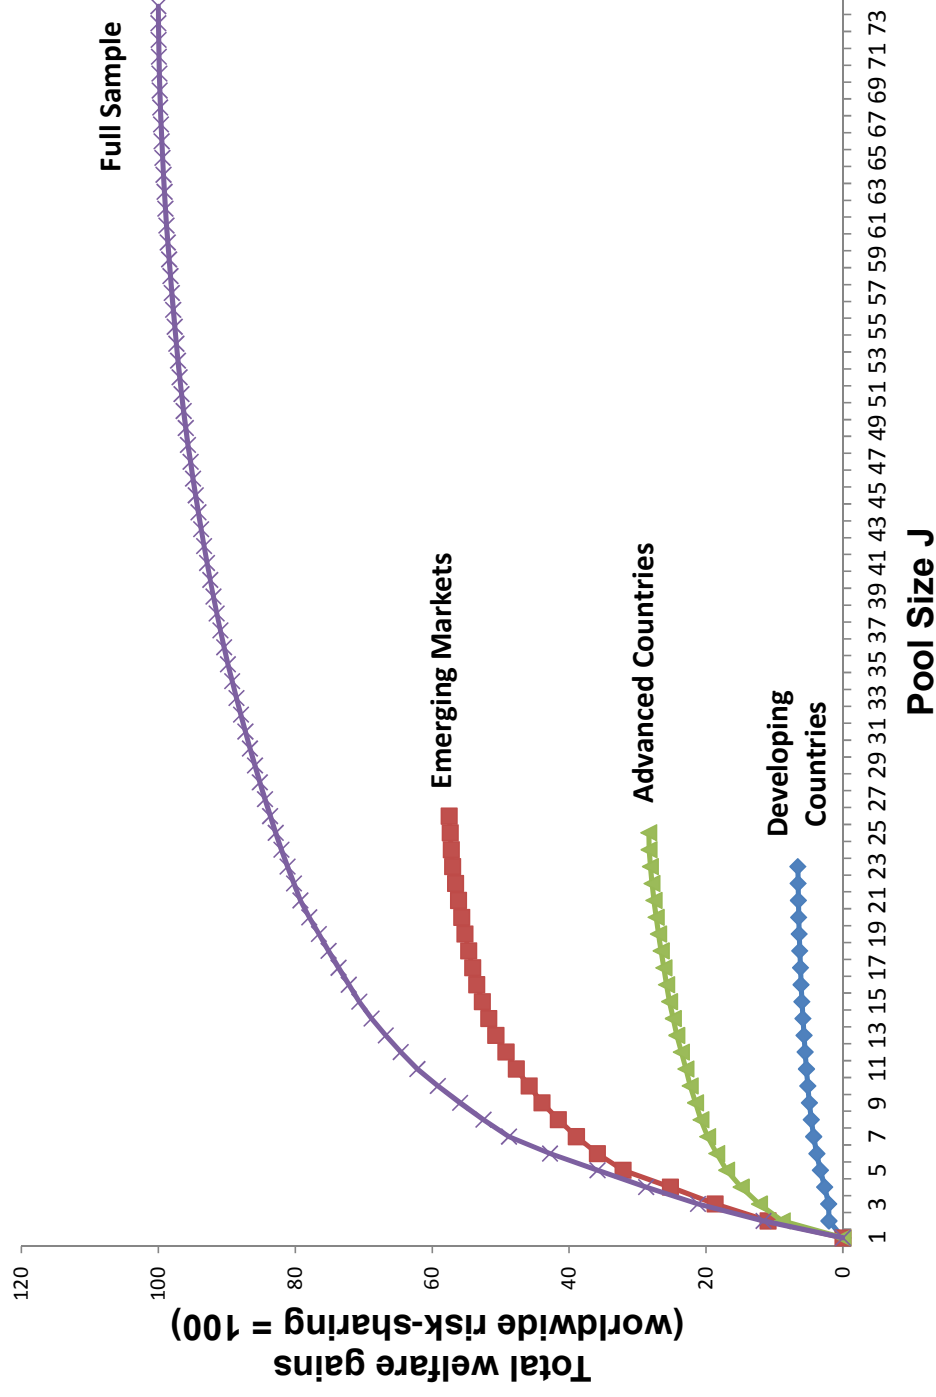


Notes: The figure reports the standard deviation of the growth rate of aggregate (poolwide) GDP for the pool (of each size) that yields the lowest standard deviation (or the highest, in the case of the top line). Each group is constrained to contain Chile. GDP data at purchasing power parity are drawn from the World Bank's World Development Indicators. The pool yielding the lowest (or highest) standard deviation is found by checking all possible combinations of countries for each J and each sample, and using the approximation procedure described in the text.



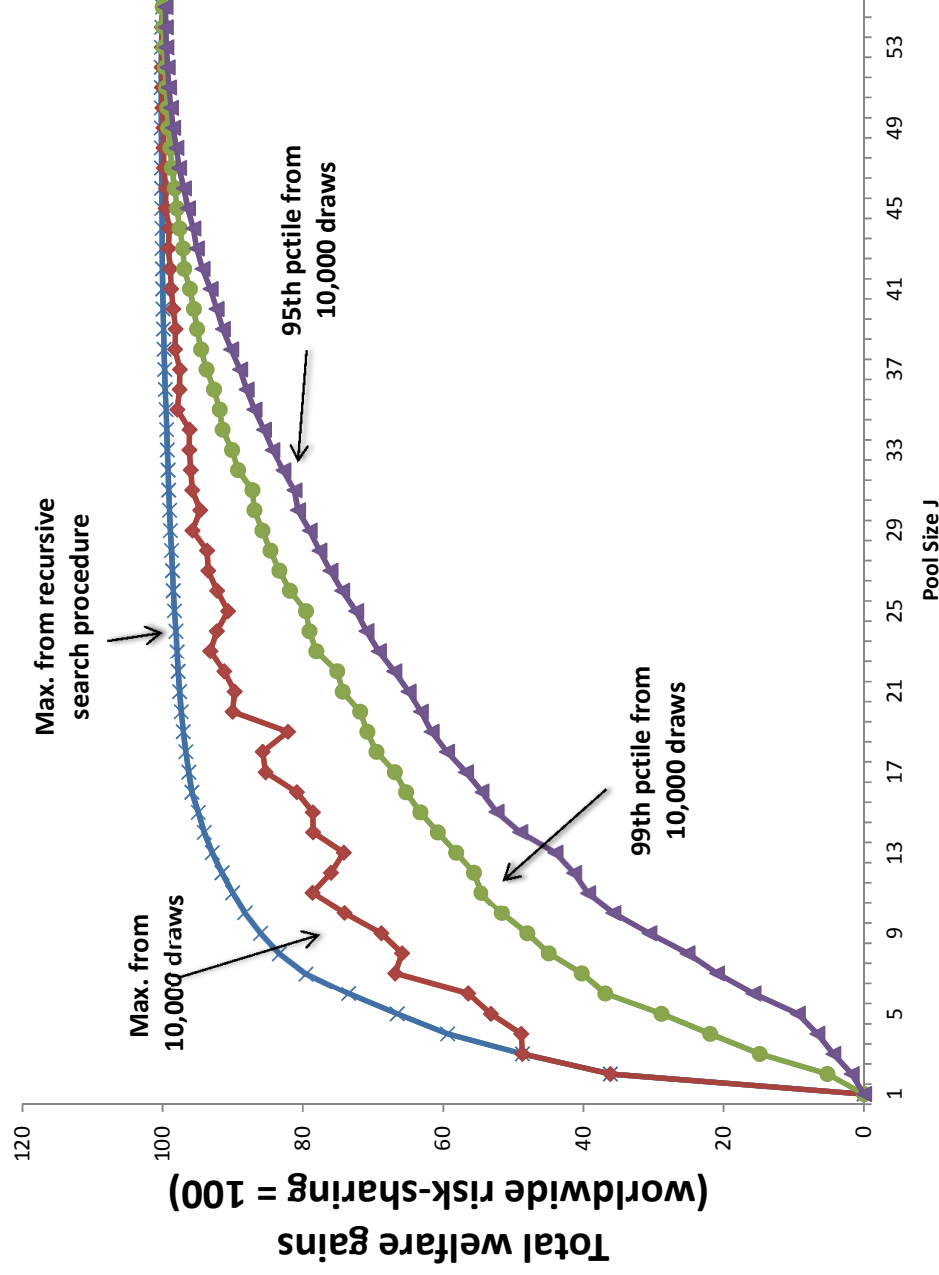
Notes: The figure reports the standard deviation of the growth rate of poolwide GDP for the pool (of each size) that yields the lowest standard deviation for each sub-sample. The pool yielding the lowest standard deviation is found by checking all possible combinations of countries and by running the approximation procedure described in the text. GDP data are from the World Bank's World Development Indicators.

Figure 3: Pooling Gains



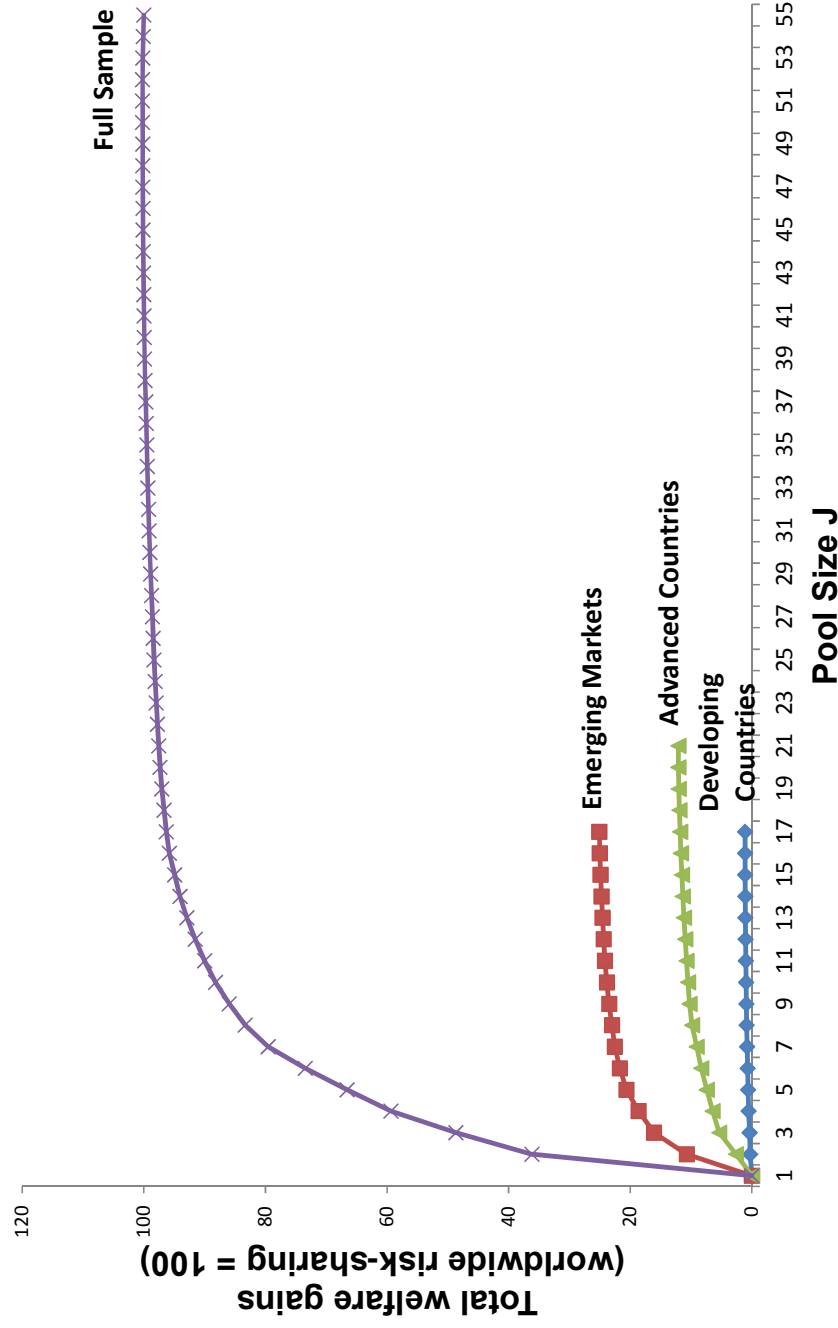
Notes: For each pool size, the figure reports the highest possible value for W^* using the procedure described in the text. Welfare gains are computed assuming $\gamma = 5$, $\beta = 0.95$, $\theta = 2$, and a constant growth rate of three percent for all countries and pools. GDP data are from the World Bank's World Development Indicators.

Figure 4: Random Pools and Actual Gains



Notes: For each pool size, the figure reports the highest possible value for W^* using the procedure described in the text. In addition, the figure reports the maximum, 99th percentile, and 95th percentile obtained from 10,000 random pool draws. Welfare gains are computed using the W^* formula given in the text assuming $\gamma = 5$, $\beta = 0.95$, $\theta = 2$ and a constant growth rate of three percent for all countries and pools. GDP data are from the World Bank's World Development Indicators.

**Figure 5: Pooling Gains
(Country-Specific Growth Rates)**



Notes: The figure reports the highest possible values for W^* using the search procedure and the formula described in the text. The results assume $\gamma = 5$, $\beta = 0.95$, $\theta = 2$, and country-specific growth rates, given by their historical means. GDP data are from the World Bank's World Development Indicators.