

Pooling Risk Among Countries*

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

July 2011

Abstract

We present a model where the enforcement of international risk sharing contracts is costly. With non diversifiable enforcement costs, welfare is not necessarily maximized for perfect worldwide risk sharing. Some groupings, or “pools” of countries can deliver higher welfare, with higher diversification gains net of enforcement costs. We construct an exhaustive list of such pools of countries. For each pool, we compute the volatility of poolwide consumption and Gross Domestic Product growth, and compare it with the volatility in each country individually. From the difference, we infer the diversification and welfare gains — gross of enforcement costs— associated with risk sharing for each pool. Welfare gains increase quickly with the size of the pool. Groupings of as few as seven countries deliver two-thirds of the maximum obtained with full worldwide integration. They are composed of heterogeneous countries, and are typically not observed, presumably because enforcement costs are large. A contrario, the few risk sharing agreements we do observe usually have a regional dimension, high trade, and presumably low diversification gains, but low enforcement costs. We conclude large welfare gains remain untapped because the enforcement of international risk sharing is costly.

Keywords: Risk Sharing, Diversification, Enforceability.

JEL Classification: E21, E32, E34, F41.

*The first draft of this paper was written while Imbs was a resident scholar in the IMF's Research Department. Financial support was provided by the National Center of Competence in Research “Financial Valuation and Risk Management”. The National Centers of Competence in Research (NCCR) are a research instrument of the Swiss National Science Foundation. We are grateful to an anonymous referee, 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.

[†]University of California at San Diego, Department of Economics

[‡]Paris School of Economics and CEPR. Corresponding author: Maison des Sciences Economiques, 106-112 bvd de l'Hopital, 75013 Paris France. +33 1 44 07 82 01. jean-imbs@gmail.com. www.jeanimbs.com.

[§]International Monetary Fund

1 Introduction

Under perfect international risk sharing, country-specific risk is insured away as households consume out of an identical portfolio of state-dependent assets. Full diversification entails payments going from booming economies into ones in recession, and requires contractual arrangements to be enforced internationally. If such enforcement is costly, the question of who to share risk with acquires key importance. The first best is not necessarily full worldwide integration, which can entail large costs. Choosing a membership involves a tradeoff between diversification benefits and enforcement costs, and may thus result in groups consisting of a limited number of countries.

We introduce a simple theory of international risk sharing with costly enforcement. We use it to motivate a study of risk sharing among groupings, or “pools” of countries. We consider all possible groupings that exist in a sample of 74 countries. For each pool, we compare the volatility of poolwide output (or consumption) with its value in autarky. This quantifies the potential for diversification gains—and ultimately welfare gains—for all possible pools of countries in our sample. Interestingly, well-chosen pools of fewer than ten countries can provide the bulk of the potential gains from international risk sharing that would accrue under worldwide risk sharing.

Our main methodological innovation consists in running a systematic search on all possible country groupings, using the variance-covariance matrix of output and consumption growth rates observed in standard data for 74 countries. For any possible pool size, we isolate the specific country groupings that minimize poolwide GDP (or consumption) volatility, and maximize welfare diversification gains.

Up to two-thirds of the maximum welfare gain can obtain in pools consisting of a handful of countries. The potential marginal gains decline quickly for groups beyond six or seven members. We find that many small pools (not surprisingly, involving relatively volatile economies) yield large risk-sharing gains—more than ten times what Lucas (1987) found for the United States.¹ But if large welfare gains can be attained by pooling few countries, why do these arrangements not emerge spontaneously more often? The largest gains are attained among countries that are heterogeneous not just in terms of business cycle characteristics, but also institutional quality, income level, and geographic location. We show that the potential diversification gains are far smaller when pools are formed within sub-samples of countries characterized by high institutional quality and an unblemished repayment record. We conjecture that enforcement may be costly for heterogeneous groupings, or for groupings that involve countries whose institutional quality and perceived creditworthiness are lower. Indeed,

¹Pallage and Robe (2003) show that the welfare cost of economic fluctuations is far larger in developing countries than in advanced economies.

the large untapped gains we identify can be interpreted as the welfare costs of poor enforcement or, more generally, weak institutions.

A few schemes have sought to foster international sharing of macroeconomic risks within clubs consisting of a limited number of countries. These schemes include, for example, pooling arrangements for international reserves, such as the Chiang Mai initiative, the Latin American Reserve Fund (FLAR), or networks of bilateral swap arrangements among the G-10 in the 1960s-70s and among the European countries during the run up to the establishment of the Euro. In fact a number of schemes have been proposed, which seek to achieve international sharing of GDP risk among small groups of countries, including Robert C. Merton's (1990, 2000) suggestions regarding networks of bilateral swaps of GDP-linked income streams.²

The few pooling arrangements observed in practice often involve a regional element, reflecting trade linkages, or a mutual interest in each other's economic performance. The honoring of international obligations is more likely, and risk sharing less costly to enforce. But trade partners are well known to have synchronized cycles. See, for example, Frankel and Rose, 1998. For the actual pools we observe, it has to be the case that the positive impact of trade linkages on contract enforceability dominates their negative impact on diversification gains. We estimate the risk-sharing gains that would accrue if existing reserve-pooling arrangements or free trade areas were to fully share GDP risk. We compare them with the benefits that could be provided by pools of similar size chosen in an unconstrained manner from the whole sample. The results are consistent with the view that contract enforceability is an important consideration. Existing arrangements seem to forgo large diversification gains in exchange for greater contract enforceability.

The paper is related to two strands of the literature. First, we build on the extensive work evaluating the gains from international risk sharing (see, for example, Cole and Obstfeld, 1991; Tesar, 1993; Lewis, 1996; van Wincoop, 1999 or Athanasoulis and van Wincoop, 2000). Our welfare analysis introduces an incomplete market version of the models in Lewis (2000) and Obstfeld (1994). Second, an important ingredient in our framework relates to the international comovement of GDP, the object of a large empirical literature (including Backus, Kehoe and Kydland, 1994; Kehoe and Perri, 2002; Imbs, 2004; and Baxter and Kouparitsas, 2005). This literature has documented an impressive degree of persistence in the correlation matrix of GDP fluctuations across countries.

The paper is organized as follows. Section II introduces a theory of international risk sharing with non-diversifiable enforcement costs, and outlines how

²On 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.

we handle the combinatorial problem. Section III presents our general results on the potential for risk-sharing gains in the sample of countries with data. In Section IV, we estimate the extent to which the potential for risk sharing is reduced when countries can only choose their partners within a constrained universe. We focus on regional constraints, and on the need for countries to have sufficiently strong institutional quality in order to be trusted. Section V concludes.

2 Methodology

A model of international risk sharing with non-diversifiable enforcement costs is first described. Then we discuss the algorithms involved in our search for optimal pools of countries.

2.1 Risk-Sharing, Volatility, and Welfare

We introduce a model of risk sharing where the purchase of a foreign security entails non-diversifiable enforcement costs. These costs capture the difficulty in ensuring dividends will be paid internationally. Perfect international risk sharing requires the availability of Arrow-Debreu securities with payment contingent on the realization of each country's idiosyncratic risk. We argue the purchase of such foreign securities comes with a premium, that can reflect the risk associated with non-payment of the contingent dividend, the necessity for an international agency whose role is to enforce contractual commitments, or simply uncertainty about the quality of local institutions. In the presence of such costs, markets are incomplete and diversification gains have to be weighted against the cost of enforcing a given risk-pooling area.

Full, worldwide risk sharing can entail large enforcement costs, and thus imply negative *net* gains. In such instance, full risk sharing does not obtain in equilibrium, and the corresponding securities do not emerge endogenously. It is plausible that enforcement costs increase with the size of the pool, so that net gains can be positive in smaller groups of countries. In such circumstances, a small-scale risk sharing pool is optimal, but only inasmuch as it excludes parties that would increase poolwide enforcement costs. Depending on the feasibility of such selective capital controls, this can be achieved either via financial markets, or via official, government-sponsored risk sharing arrangements. Given the virtual absence of a liquid market for GDP-indexed securities along the lines suggested by Shiller (1993), the latter seems to be more realistic.

We rely on a well known framework, based on Lewis (2000) and Obstfeld (1994). We abstract from non tradability and non separability in utility, and from the possible impact of uncertainty on growth. These refinements tend to

boost 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.

Utility follows Epstein and Zin (1989). Output at time t in country j , Y_t^j , is log-normally distributed. 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 , and the process for endowment is

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

where $y = \ln Y_t^j$ 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.³

We abstract from self insurance and saving. As in Lewis (2000), we focus the analysis on the welfare gains afforded by international diversification. This assumes away alternative sources of consumption smoothing. But it is consistent with our purpose of evaluating the potential from *international* risk sharing. Under autarky, $C_t^j = Y_t^j$, and welfare in country j at time $t = 0$ is given by

$$U_0^j = C_0^j \left\{ 1 - \beta \exp \left[(1-\theta) \left(\mu_j - \frac{1}{2}\gamma\sigma_j^2 \right) \right] \right\}^{-1/(1-\theta)} \quad (3)$$

We can now ask the question of the welfare gains in country j associated with moving from autarky to pooling. Risk sharing within the pool ensures that country j 's consumption grows with poolwide output at rate $\bar{\mu}$ and fluctuates with poolwide volatility $\bar{\sigma}^2$. Define δ_j , the compensating differential that makes country j indifferent between autarky and pooling at time 0. We have

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

where \bar{C}_0 is consumption at time 0 in the pool. Rearranging

$$\delta_j = \frac{\bar{C}_0}{C_0^j} \left(\frac{1 - \beta M_j^{1-\theta}}{1 - \beta \bar{M}^{1-\theta}} \right)^{1/(1-\theta)} - 1 \quad (5)$$

³We assume throughout that endowments have a unit root. This tends to magnify the welfare consequences of diversifying risk, since shocks are permanent. But it also tends to lower the measured volatility of output and consumption, and thus the welfare gains from diversification.

For notational convenience $M_j = \exp(\mu_j - \frac{1}{2}\gamma\sigma_j^2)$ and $\bar{M} = \exp(\bar{\mu} - \frac{1}{2}\gamma\bar{\sigma}^2)$ denote the certainty equivalent growth paths of endowments in autarky and in the pool. Under risk sharing, consumption in country j shares the features of group-wide, pooled output. The welfare gains from risk sharing have three components. First, pure diversification gains, i.e. the difference between individual and poolwide volatilities, σ_j^2 and $\bar{\sigma}^2$. Second, growth differentials, i.e., the difference between growth rates within and without the pool, μ_j and $\bar{\mu}$. Third, the ratio between initial consumption in autarky, C_0^j , and initial consumption in the pool, \bar{C}_0 . This term reflects an “entry transfer”, expressed in terms of initial consumption C_0^j , that country j pays to other members in exchange for being allowed into the pool and consume \bar{C}_0 instead. Entry transfers that are greater than one mean country j increases its level of consumption at the time it enters the pool. It must reflect the desirability of country j in the pool, from the standpoint of its diversification potential and growth prospects. Of course, it also reflects the compensation needed by country j to share risk with others whose growth rates and volatilities are unattractive from its point of view.⁴

The entry transfers are a key component of welfare. They depend endogenously on the parameters of the model, and in particular on enforcement costs. In autarky, $C_0^j = Y_0^j$. In the pool, initial consumption \bar{C}_0 is given by

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

where p_0^j is the time 0 poolwide price of the security issued by country j and contingent on the realization of Y_t^j . Similarly, \bar{p}_0 is the time 0 price of the poolwide security contingent on the realization of poolwide endowment $\bar{y}_t = \bar{y}_{t-1} + \bar{\mu} - \frac{1}{2}\bar{\sigma}^2 + \bar{\varepsilon}_t$.⁵ Enforcement costs are denoted by τ . The ratio $\frac{p_0^j}{\bar{p}_0 + \tau}$ represents the endogenous share of poolwide income at time 0, \bar{Y}_0 , that country j can claim. That share depends on the price of the securities country j sells, relative to the purchase price of the poolwide security, inclusive of enforcement costs τ . Enforcement costs are not proportional to the price of the security, but they are specific to each pool.

Asset prices are endogenous to the characteristics of the economy. Lewis (2000) shows that simple expressions for both prices satisfy the optimality conditions arising from the maximization of utility (1), and in particular the Euler equations corresponding to the autarkic and the integrated economies. She shows that

$$p_0^j = Y_0^j \frac{\beta \bar{M}^{-\theta} H_j}{1 - \beta \bar{M}^{-\theta} H_j} \quad (7)$$

⁴Scale effects are implicitly embedded in all three components. A country so large that its pooled growth and volatility are virtually identical to what obtains in autarky, will stand to gain very little from international risk sharing.

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

where $H_j = \exp \left[\mu_j + \frac{1}{2} \gamma \bar{\sigma}^2 - \gamma \text{cov} \left(\varepsilon_t^j, \bar{\varepsilon}_t \right) \right]$. The price of country j 's security decreases with $\text{cov} \left(\varepsilon_t^j, \bar{\varepsilon}_t \right)$, as it worsens its hedging benefits for the rest of the pool. Moreover

$$\bar{p}_0 = \bar{Y}_0 \frac{\beta \bar{M}^{1-\theta}}{1 - \beta \bar{M}^{1-\theta}} \quad (8)$$

Combining equations (6)-(7)-(8) yields an expression for entry transfers arising from country j entering a pool, given by

$$\frac{\bar{C}_0}{C_0^j} = \frac{\beta \bar{M}^{-\theta} H_j}{1 - \beta \bar{M}^{-\theta} H_j} \frac{\bar{Y}_0 (1 - \beta \bar{M}^{1-\theta})}{(\bar{Y}_0 - \tau) \beta \bar{M}^{1-\theta} + \tau} \quad (9)$$

which can be substituted in equation (5) to obtain the welfare gains associated with the pool from the point of view of country j .

Consider now the welfare gains W associated with a pool as a whole. Define

$$W = \sum_{j=1}^N \left(\frac{Y_0^j}{\sum_{j=1}^N Y_0^j} \right) \delta_j \equiv \sum_{j=1}^N \omega_j \delta_j \quad (10)$$

where N denotes the number of countries in the proposed pool. The welfare criterion weighs each participating economy by its initial share in poolwide endowment, at the time the pool is initiated. In the Appendix we show that, in the absence of large outliers in country average growth rates or volatility, total poolwide welfare simplifies into

$$\begin{aligned} W &\simeq \frac{\bar{Y}_0 \beta \bar{M}^{1-\theta}}{(\bar{Y}_0 - \tau) \beta \bar{M}^{1-\theta} + \tau} \sum_{j=1}^N \omega_j \left(\frac{1 - \beta M_j^{1-\theta}}{1 - \beta \bar{M}^{1-\theta}} \right)^{1/(1-\theta)} - 1 \\ &\equiv \frac{\bar{Y}_0 \beta \bar{M}^{1-\theta}}{(\bar{Y}_0 - \tau) \beta \bar{M}^{1-\theta} + \tau} GW - 1 \end{aligned} \quad (11)$$

Total welfare has two components. First a weighted sum of country-specific diversification gains, GW , which abstracts from enforcement costs. GW represents the poolwide welfare that would obtain in the absence of any enforcement costs, as can be seen by setting $\tau = 0$ in equation (11). It is a measure of gross welfare gains, which would be maximized under full worldwide diversification, where the certainty equivalent growth path of poolwide endowment \bar{M} would be maximized. Second, gross welfare gains are scaled by a term that is strictly decreasing in the enforcement costs. This term reflects the entry transfers introduced in equation (5).

Equation (11) crystallizes the tradeoff between gross diversification gains and poolwide enforcement costs. The potentially large diversification gains, i.e. high

values for \bar{M} , have to be weighted against potentially large enforcement costs. It is possible for W to be strictly negative, which happens as soon as

$$GW < \frac{(\bar{Y}_0 - \tau) \beta \bar{M}^{1-\theta} + \tau}{\bar{Y}_0 \beta \bar{M}^{1-\theta}} \quad (12)$$

since $(\bar{Y}_0 - \tau) \beta \bar{M}^{1-\theta} + \tau$ is always positive. The right-hand side of inequality

(12) increases in τ , but enforcement costs leave the left-hand side unaffected. Some combinations of \bar{M} and τ can therefore be such that $W < 0$. We do not observe τ directly, but can conjecture that it tends to take high values for large pools of countries. If so, equation (12) can explain why GDP-indexed securities are virtually absent: it is because the risk sharing arrangements they imply are too costly to enforce. It may also explain why the few risk sharing arrangements we do observe have a strong regional dimension: it is because τ tends to be low in small, regional groups of countries. That way *net* welfare W can remain positive even if GW tends to be low amongst such homogeneous sets of economies.

Concretely, two types of arrangements can implement the type of risk sharing consistent with our setup. Under the first, countries in the pool issue claims on their output as proposed by Shiller (1993). Capital controls vis-à-vis non-members then ensure that only the residents of countries in the pool have access to such securities. This is important under the premise that enforcement costs τ depend on the actual membership of the pool. Net welfare gains W can fall as the representative portfolio in the pool starts including securities contingent on the risk of a country with, for instance, poor institutions. If this happens, the issuance of the security in the first place can prove not to be optimal. A second type of arrangement consists of GDP swaps, along the lines proposed by Merton (1990, 2000), either as a network of bilateral swaps, or as swaps intermediated by a central entity for the pool. Under the swaps, each period, each country pays the others the net difference between its current output and its share in poolwide output, as implied by its long-run share of poolwide wealth. Participation in the network of swaps defines the pool membership, and, we argue, the level of τ . The latter type of arrangement is practical, as it controls membership in the pool and thus the level of enforcement costs. The former requires selective capital controls, possibly harder to implement.

2.2 Implementation

We compute *gross* welfare gains GW . We do so because τ is not observable. Our purpose is to evaluate the *potential* for risk sharing in pools of any size, which makes sense in the presence of enforcement costs. Equation (11) shows gross poolwide welfare does not depend on the covariances between each country's idiosyncratic risk. H_j drops from GW , which simplifies considerably our

exercise. We obtain values for GW corresponding to observed realizations of μ_j , $\bar{\mu}$, σ_j , $\bar{\sigma}$, and for given preference parameters.

We proceed incrementally. First, we report the standard deviation of the growth rate for individual country consumption and GDP, vs. their poolwide counterparts. This simple approach, focused on pure diversification gains, conveys most of the key economic intuition. In particular, we show $\bar{\sigma}$ falls quickly as the number of countries pooling risk increases. Second, we compute the poolwide welfare gains implied by these changes in volatility, assuming expected (GDP or consumption) growth is the same for all countries ($\mu_j = \bar{\mu}$). This follows exactly from Obstfeld (1994). As was the case there, welfare is a monotonic, non-linear transformation of volatility. Third, we relax the assumption that growth rates are the same for all countries. We project μ_j and $\bar{\mu}$ using past observed growth rates.

Searching for pools of countries that yield the lowest possible variance in the growth rate of poolwide GDP (or consumption) 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 (given by C_2^N), then 3 at a time (given by C_3^N), and so on, where $C_p^N = \frac{N!}{p!(N-p)!}$. The total number of partitions is $\sum_{p=1}^N C_p^N = 2^N - 1$, which 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 p , 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 $p = 7$ or less ($C_7^{74} = 1.8 \times 10^9$). Since $C_p^N = C_{N-p}^N$, we can also draw the inventory of all combinations of $p = 67$ or more. Outside of these pool sizes, we need to resort to an approximation algorithm. We need an approximation. 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 (or consumption) 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 obtain all possible combinations up to the maximum pool size where it is feasible through an exhaustive search. This includes all pools of maximum size $p = 7$, drawn from the universe of 74 countries. We save not only the best pool of each size, but also the best S pools that include each of the countries in the universe under consideration. For groups drawn from

the worldwide sample, we impose $S = 1,351$. When drawing the inventories of groups of size 3, for instance, we collect the grouping with lowest volatility, but also the next 1,351 that contain country 1, country 2, etc. In other words, we collect $1,351 \times 74 \simeq 100,000$ additional pools. We call these “seed” pools. We can collect seed pools no matter the value of p . The difference is that for $p \leq 7$ we know the universe of *all* pools.

For each pool size p , we collect $S \times N$ seed pools. We isolate all groups that include the members in the optimal pool of size $p - 1$, plus one of the $N - p$ remaining countries. Among these, we find the best pool of size p , as well as the best new $S \times N$ seed pools of size p . The procedure is iterated. Although there is a recursive aspect to this, the fact that at each stage we consider the best S pools for each of the N countries gives plenty of opportunities for countries that are in the best pool of size $p - 1$ 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 of 7 (or less) and 67 (or more) countries selected amongst 74. We compared the groupings implied by an exhaustive inventory 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 p , we have checked large numbers of random samples of countries. 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

3 Data and Results

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

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.

3.1 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. Compared with the widely used Penn World Tables (PWT), the World Bank database has similar quality, and in fact builds from largely identical information. But it has greater country coverage and provides PPP-adjusted data until 2004 rather than 2000. We cross-checked the two databases over the period covered by both. The results are largely unaffected if we use PWT.

The sample includes 25 advanced countries, 26 emerging market countries, and 23 developing countries with complete coverage and data of reasonable quality. (The full country list is provided in Appendix B). Advanced countries are defined as in the International Monetary Fund’s World Economic Outlook. The remaining countries 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.

In line with the bulk of the literature on international risk sharing, we assume that PPP holds. This corresponds to the notion that risk sharing is contracted at a pre-agreed exchange rate, one that is expected to prevail in the long run. While standard, it is an important assumption. Previous 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.

We use these data to compute autarkic and poolwide growth and volatility for both GDP and consumption. The approach makes use of the variance-covariance matrices of both variables across countries, to infer expected gains arising from international risk sharing. Both matrices are assumed to be 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 work out the pools that minimize risk from the standpoint of an individual country. We focus on Chile, viewed by interna-

tional investors as a relatively safe emerging market. Chile is not participating in existing or prospective reserve-pooling arrangements and its economy is not overwhelmingly linked to a single or a few other countries. For each pool size p , Figure 1 plots the standard deviation of the growth rate of poolwide GDP for the groups of countries containing Chile, chosen to minimize poolwide volatility. This is not a welfare measure, which explains its non-monotonicity with respect to p . Equation (5) illustrates that δ_j is not proportional to the difference between σ_j and $\bar{\sigma}$. The envelopes are displayed for various restrictions on the universe of potential partners. We present four cases: pools with the whole sample of 73 countries, and pools with other emerging markets, developing countries, or advanced economies. To give a sense of the importance of choosing well one's risk-sharing partners, we also plot the maximum value of $\bar{\sigma}$ for all pools containing Chile.

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 often 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 $p > 30$. Beyond a certain pool size, hedging opportunities are exhausted, and the pool starts including countries with high volatilities. 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 p , 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 possible gains decline when the universe of countries is constrained. The lowest possible 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 large 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. We also compute diversification gains for both GDP and consumption. Figure 2 reports the envelope of minimal output volatility for all pool sizes p 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 $p = 7$ along the global envelope. Diversification gains continue to be achieved within groups consisting of a small number of countries in this general setup.⁷

The list of countries involved in minimum-volatility pools confirms that heterogeneity is key. Interestingly, the list overlaps with that obtained for Chile. This is unlikely to be an artifact of our approximation method, despite its recursive structure, because the procedure leaves plenty of opportunities for countries to drop out of the best pool as p increases. Rather, the evidence 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 covariances, i.e. desirable hedging properties.

Figure 2 also reports minimum volatilities for sub-samples constrained to advanced economies, emerging markets, and developing countries. Diversification gains remain substantial within each sub-sample, but they are a full percentage point above the global envelope. For $p > 3$, even low volatility advanced countries can diversify substantially more risk in pools that involve emerging or developing countries. Advanced economies achieve somewhat smaller gains, which is consistent with their lower volatility and internationally correlated business cycles. All four envelopes display the same non-monotonicity as was apparent in Figure 1, with high marginal diversification gains for $p < 10$, that turn negative for $p > 15$, when countries with poor hedging properties start being included.

Figures 1 and 2 help quantify the potential diversification gains that would

⁷The value reported for $p = 1$ corresponds to the standard deviation of the individual 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. But we know the identities of the relevant groupings, and can thus assess the gains that optimal pooling would provide to member countries. For example, for $p = 7$, the standard deviations of growth rates range from 1.44% for Sweden to 8.97% for Nicaragua. The diversification gains are distributed unequally, with far larger gains accruing to countries with more volatile individual growth rates. This asymmetry has implications for entry transfers.

accrue to the representative consumer living in a pool of a given size. It captures the decrease in output volatility, which we ultimately interpret in welfare terms. But it is entirely possible some of these gains have already been reaped. In that case, the volatility of output would be what we document in Figures 1 and 2, but that of consumption—which is what matters for welfare—would be lower. Does the volatility of consumption imply different conclusions? We verify this in Figure 3, where we reproduce the exercise on the basis of consumption data.

The results are virtually identical. Consumption volatilities are highest amongst developing and emerging markets, where maximum diversification gains are in fact slightly higher than in Figure 2. This reflects the well known fact that consumption volatility tends to exceed output volatility in the developing world. The envelopes for advanced countries, and the global envelope, are slightly below their counterparts in Figure 2. We interpret the finding as a reflection of the financial development in the developed world, where output fluctuations are smoothed somewhat in consumption. But the main results stand: the diversification of consumption volatility accrues at low values of p , below 10. Pools with fewer than 10 countries deliver the bulk of volatility reduction, just as was the case for output data. The absence of any differences in results reflects the “quantity puzzle” coined by Backus, Kehoe and Kydland (1994). Contrary to the predictions of a model with complete markets, international correlations in consumption are in fact lower than business cycles synchronization.

3.4 Welfare Gains

We now turn to welfare. We compute GW , but still constrain the expected growth rate to be the same for all countries, $\mu_j = \bar{\mu}$. The subjective discount rate β is set at 0.95, for a 5% annual discount rate, $\theta = 2$ and $\gamma = 5$. Figure 4 reports the highest value of GW for any pool size p , and once again for four different sub-samples of countries. *Gross* welfare gains increase monotonically with p , with a maximum when all N candidate countries are sharing risk. Just as for volatility, the marginal increases in GW are largest for $p < 10$. Gross welfare gains are large for small pools. Marginal gains peter out as soon as $p > 7$ 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.

Table 1 illustrates the composition of GW for various pools. The Table reports the minimum value of δ_j , its median, and the total gross welfare gain corresponding to pools drawn from four samples of countries. The first pool is formed by the entire 74-country sample. GW reaches its maximum for $p = 74$, at 1.9% of permanent poolwide consumption. But the distribution of δ_j across the membership is skewed, with a minimum for the US at 0.55%, and a median for the Dominican Republic at 4.10%. The welfare consequences of pools are heterogeneous across their membership, reflecting the very heterogeneity of the constituent countries.

Gross welfare gains decrease with per capita income, because they increase with autarkic output volatility. Among advanced economies, GW reaches a maximum of 0.7%, whereas it is 4.4% in emerging markets and 7.4% in developing countries. The distribution of δ_j is also less dispersed in homogeneous pools, ranging from 0.4 to 0.9% among advanced economies. It is more dispersed among emerging markets or developing countries, where the heterogeneity in output volatility is more pronounced.

How important is active diversification in reaping the benefits described in Figure 4? How do the optimal pools described in the Figure differ from alternative groupings of countries, drawn randomly? In Figure 5 we report the welfare gains implied by pools drawn at random. We consider purely random draws of 10,000 pools for each value of p . The Figure reports maximum GW for each draw, its 99th and 95th percentiles, along with the recursive maximum reported in Figure 4. The random maximum displays a trend increase in p , with some convexity. As p increases, GW does fluctuate quite considerably around the trend, reflecting the importance of a few countries in delivering the smooth welfare gains from Figure 4. But the 99th and 95th percentiles in the distribution of simulated GW lose any similitude with the recursive measure: they are almost linear in p , and miss most of the convexity apparent from previous Figures. In other words, while high values of GW can be reached for low values of p , 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. Figure 4 corresponds to $\theta = 2$, i.e. an elasticity of substitution of 0.5. This represents a relatively low value of the parameter, and therefore tends to minimize welfare gains. In Figure 6 we investigate the robustness of our results to alternative calibrations. We compute GW for all p using intertemporal elasticities of 0.9 or 2. Unsurprisingly, the envelopes shift up for higher values of the elasticity. But crucially, the convexity of all envelopes is preserved. Even for $\theta = 0.5$, the bulk of maximum gross welfare is reached for $p < 10$. The welfare gains from international risk sharing are larger for lower values of θ , but they accrue rapidly as p increases irrespective of the calibration choice.

The measures of welfare discussed so far compare growth and volatility of output in autarky with their value within a risk sharing pool. Output in autarky may have drastically different properties from consumption in autarky, which is what welfare is concerned with. If some risk sharing gains are reaped already by the countries in our sample, using output rather than consumption data will bias upwards our measures of GW . In Figure 7, we compare the properties of *consumption* in autarky with that of output in the pool. The comparison isolates precisely the welfare gains associated with a risk sharing pool given observed consumption. Reassuringly, Figure 7 is virtually identical to Figure 4, where we only used output data. Maximum values of GW are the same, and so are

the rankings of welfare envelopes across sub-samples of countries. Consumption data suggest little of the risk sharing gains afforded by the variance covariance matrix of output fluctuations are effectively reaped.

3.5 Pooling Growth Rates

The previous section computes GW holding growth rates constant across the pool. 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 more than a few years ahead make it difficult to incorporate differences in expected growth into risk-sharing contracts. As shown by Easterly and others (1993), country rankings with respect to growth rates change dramatically from one decade to the next. Similarly, Jones and Olken (2008) document that most countries experience both growth miracles and failures at some point in their history. Such uncertainty is liable to complicate the enforcement of risk sharing agreements, with end effects on τ and on net welfare W .

We now investigate the behavior of gross welfare GW when expected growth is allowed to vary within the pool. To estimate expected economic growth, we simply consider the naïve averaging of historical growth rates over the entire period under consideration. We also assume that individual countries' growth rates are unaffected by pooling arrangements. 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. Thus $\bar{\mu}$ depends on the identity of the pool, which μ_j does not. 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 GW on the basis of historical growth rates represent if anything a lower bound.

Figure 8 reproduces Figure 4, allowing for $\mu_j \neq \bar{\mu}$. All envelopes shift up. For instance, the maximum value of GW in the full sample is now above 5% of consumption, vs. 1.9% with constant growth rates. Such a large rise is reminiscent of the literature started by Obstfeld (1994), arguing international financial integration gives access to high returns, which can have large welfare consequences. Welfare increases, especially for emerging markets, but also in the sample formed by advanced countries. The change in GW is least pronounced for pools drawn from developing countries. In unreported work, we obtained poolwide weighted averages of μ_j , computed for all p . We compared them with $\bar{\mu}$, once again computed for all pools and across all four samples. We found little difference amongst developing countries, some improvement for advanced economies, and largest increases for emerging markets. From a pure accounting standpoint therefore, past 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 GW is largest in this sub-sample.

Most importantly, Figure 8 confirms the broad pattern of our results. Most marginal gains in GW accrue for $p < 10$. More than four fifths of maximum GW are reached in pools of 10 countries or fewer. The welfare gains are larger in most instances, but the shape of the envelopes remains the same.

4 Pooling Risk Within Sub-Samples

The previous section establishes substantial gross welfare gains are accessible to small pools of countries. Yet risk sharing agreements are rarely observed, and GDP contingent securities almost never arise endogenously. It must be therefore that $W < GW$ in most instances, i.e. enforcement costs τ render most potential pools Pareto inferior. In this section we take a stab at quantifying the forgone welfare gains implied by enforcement costs τ . This is done in two ways. First, we focus on countries where enforcement costs have observable reasons to be low. 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 GW for pools of countries that are constrained to belong to the same category, e.g. high institutional quality. We reason τ is low in these hypothetical pools. Gross welfare there is contrasted with the value of GW in pools with open membership, e.g. inclusive of countries with low institutional quality. The forgone welfare associated with an unconstrained membership is indicative of the costs of τ .

Second, existing schemes are considered. These include free-trade agreements, such as the European Union, ASEAN or Mercosur. The European Monetary Union is of special interest, given the increasingly pressing relevance of insurance schemes within the Eurozone. We also include schemes that were built to provide some pool-wide insurance. Such schemes are rare. We consider the Chiang-Mai Initiative and the Latin American Reserve Fund (FLAR). We illustrate the reduction in output volatility possible in the various considered sub-samples. Each pool's GW is computed, and compared with the type of gross welfare gains that could be attained in a pool of identical size whose membership would be entirely unconstrained. Assuming existing arrangements ultimately purport to share risk internationally, the difference is indicative of τ , since it represents the forgone *gross* welfare gains.

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 2 reports the median value of σ_j across the countries j in the sample, the median value of minimum variance $\bar{\sigma}$ across all p , the median value of δ_j across the countries in the sample, and finally GW . The latter two measures are computed assuming $\mu_j = \bar{\mu}$, calibrated at 3%. Column (1) confirms output volatility is higher in countries with poor institutions. It is 4.41% with low enforceability vs. 2.11%, and 4.26% with institutional quality below average vs. 2.61%. The median welfare gain δ_j reported in column (3) suggests countries with poor institutions benefit sizeably more from unconstrained diversification. The median value of δ_j is almost 6% for countries below excellent enforceability, provided they can share risk with partners that have good institutions. This is five times larger than the gains accruing to countries with excellent track record, 1.17% if they pool with any partner in the sample, and 1.03% if they only pool with their kin. The magnitudes are analogous when institutions are ranked according to Kaufmann et al (2005).

The last column in Table 2 reports estimates of GW . The only groups with large gross welfare gains from diversification are countries with low quality institutions. And in order to reap above 5% of consumption, they need to be able to pool risk with partners that have high quality institutions. If pooling is constrained to only include countries with good institutions, GW collapses to below 1%. Inasmuch as τ is close to zero in such pools, the discrepancy we document in Table 2 means the welfare costs associated with τ must at least represent 4% of consumption.⁸

The lower panel of Table 2 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 do not turn out to be very important for advanced European countries. With a median volatility σ_j of 1.84%, individual country’s welfare gains δ_j are well below 1%, and so are estimates of GW . This holds true irrespective whether European countries pool within Europe or with advanced economies in general. 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 3.62% in Asia and 4.41% in Latin America. Column (2) illustrates that pooled

⁸ τ cannot be zero amongst countries with good institutions. If it were, we would observe such pooling and/or such securities, since then $W = GW = 1\%$.

volatility $\bar{\sigma}$ can be reduced to just above 1%, provided each country is allowed to share risk beyond its own region. If Asian countries are constrained to pool within the region, $\bar{\sigma} = 1.84\%$, while it is $\bar{\sigma} = 1.90\%$ for Latin American emerging economies. This represents a sizeable decrease in volatility, even if constrained within the region. Hence, the estimated values of GW vary between 2.94 and 3.54% for Asia, and between 4.08 and 5.39% for Latin America. Risk sharing agreements do not exist that involve all of emerging Asia, or of emerging Latin America, or indeed all of the emerging world. Securities that are contingent on these countries GDP are not being traded. And, with the possible exception of FLAR, no official agreement exist either. Such absence suggest that, even in these relatively homogeneous regions, the enforcement of international contracts is costly enough to deter from agreements that could generate a 4 to 5% permanent increase in consumption.

4.2 Existing Arrangements

We consider the potential welfare gains arising from existing international agreements. We include free-trade agreements, whose participants have long-established cooperation and thus presumably low τ . We also discuss actual risk-sharing arrangements: the Chiang-Mai Initiative and the FLAR. For several such agreements, Table 3 reports the median and minimum values of σ_j in each agreement, and the pool's diversified volatility, $\bar{\sigma}_{Pool}$. We also compute the minimum value of $\bar{\sigma}$ that can be achieved in *any, unconstrained* pool of similar size p , labeled $\bar{\sigma}_{Min}$. The discrepancy reflect the forgone diversification gains, which must be ascribed to τ in the considered pools.

Median volatility is highest in West Africa, followed by ASEAN countries and Mercosur. Those are also the groups with highest minimum volatility σ_j . In contrast, both median and minimum volatility are low in developed countries, e.g. in the EU, EMU, and NAFTA. The third line in Table 3 confirms that all the considered pools (except FLAR) deliver diversification, as pooled volatility $\bar{\sigma}$ is smaller than the minimum value of σ_j . Unsurprisingly, the decrease is small for developed economies, from 1.8% down to 1.1 – 1.2% in the EMU and EU, and from 2.11 to 1.81% in NAFTA. Decreases in volatility are more substantial for groupings of countries in the developing world. For instance, median individual volatility in the Chiang-Mai group is 3.93%, but the pool's volatility is down to 1.4%. In the FLAR, volatility decreases from a median value of 4% down to 2.48%. These are substantial decreases. Nevertheless, for all agreements considered in Table 3, there exist alternative risk pools that, for identical p , deliver substantially larger reduction in poolwide volatility. For instance, a pool of the size of the European Union could potentially display poolwide volatility as low as 0.65%. A pool of the size of FLAR would have volatility diversified down to 0.71%. All values of $\bar{\sigma}_{Min}$ in the last row of Table 3 are below 1%, and substantially below $\bar{\sigma}_{Pool}$ for all considered arrangements. Such comparisons between $\bar{\sigma}_{Pool}$ and $\bar{\sigma}_{Min}$ can be interpreted as a measure of

forgone diversification gains, because of high τ in unconstrained, heterogeneous groupings.

5 Conclusion

We develop a model of international risk sharing with non-diversifiable enforcement costs. In such an environment, the welfare gains from international risk sharing must be understood net of enforcement costs. Net welfare gains do not necessarily reach their highest value for perfect, worldwide financial integration, even though gross welfare does. If enforcement costs are large, net welfare can even fall with risk sharing, which can explain why traded securities contingent on the realization of country risk do not arise endogenously.

With such market incompleteness, it can be optimal to share risk within pools of countries, rather than globally. We provide a methodology that calculates the *gross* welfare gains associated with *any* pool comprising *any* number of economies, selected among a maximum of 74 countries. We find that the bulk of the diversification gains, and ultimately of gross welfare gains, can be achieved in pools involving fewer than 10 carefully chosen countries. Even though they consist of few members, some of these pools deliver large gross welfare gains, well above 5% of consumption. But their constituent economies display considerable heterogeneity in terms of business cycle characteristics, institutions and default history. Thus they continue to entail substantial enforcement costs, which explains why they do not arise endogenously.

We find enforcement costs are likely to be large. We compute the gross welfare gains arising from risk sharing between countries with good institutions and pristine default history. We compare it with gross welfare arising from risk sharing involving countries with poor contract enforcement. The difference equals up to 4% of consumption. Similarly, we show the diversification gains afforded by existing pools of countries fall dramatically short of what could be reaped in alternative, carefully chosen pools of identical size, but with weaker institutions. Both pieces of evidence point to considerable welfare costs of low institutions.

6 Appendix A: Poolwide Welfare

From the definitions of W and δ_j , we have

$$W = \bar{Y}_0 \frac{\beta \bar{M}^{-\theta} (1 - \beta \bar{M}^{1-\theta})^{\theta/(\theta-1)}}{(\bar{Y}_0 - \tau) \beta \bar{M}^{1-\theta} + \tau} \sum_{j=1}^N \omega_j \frac{H_j (1 - \beta M_j^{1-\theta})^{1/(1-\theta)}}{1 - \beta \bar{M}^{-\theta} H_j} - 1 \quad (\text{A1})$$

The summation can be rearranged into

$$\sum_{j=1}^N \omega_j \frac{H_j (1 - \beta M_j^{1-\theta})^{1/(1-\theta)}}{1 - \beta \bar{M}^{-\theta} H_j} = \sum_{j=1}^N \omega_j \frac{(1 - \beta M_j^{1-\theta})^{1/(1-\theta)} \prod_{k=1}^N H_k}{\prod_{k \neq j}^N H_k - \beta \bar{M}^{-\theta} \prod_{k=1}^N H_k} \quad (\text{A2})$$

From the definition of H_j ,

$$\begin{aligned} \prod_{k=1}^N H_k &= \exp \left[\sum_{k=1}^N \mu_k + \frac{N}{2} \gamma \bar{\sigma}^2 - \gamma \sum_{k=1}^N \text{cov}(\varepsilon_t^k, \bar{\varepsilon}_t) \right] \\ &= \exp \left[\sum_{k=1}^N \mu_k + \left(\frac{N}{2} - 1 \right) \gamma \bar{\sigma}^2 \right] \end{aligned}$$

since by definition $\sum_{k=1}^N \text{cov}(\varepsilon_t^k, \bar{\varepsilon}_t) = \bar{\sigma}^2$. Rearranging,

$$\prod_{k=1}^N H_k = \bar{M} \exp \left[\sum_{k=1}^N \mu_k - \bar{\mu} + \frac{N-1}{2} \gamma \bar{\sigma}^2 \right] \equiv \bar{M} \Gamma \quad (\text{A3})$$

Analogously, by definition,

$$\prod_{k \neq j}^N H_k = \exp \left[\sum_{k=1}^N \mu_k - \mu_j + \frac{N-1}{2} \gamma \bar{\sigma}^2 - \gamma \bar{\sigma}^2 + \gamma \sigma_j^2 + \gamma \sum_{k \neq j}^N \text{cov}(\varepsilon_t^k, \bar{\varepsilon}_t) \right]$$

Now in the absence of large outliers in average output growth or volatility,

$\mu_j \simeq \bar{\mu}$ and $\bar{\sigma}^2 \simeq \sigma_j^2 + \sum_{k \neq j}^N \text{cov}(\varepsilon_t^k, \bar{\varepsilon}_t)$, so that

$$\prod_{k \neq j}^N H_k \simeq \Gamma \quad (\text{A4})$$

Substituting equation (A3) and (A4) into (A2) yields

$$\sum_{j=1}^N \omega_j \frac{H_j (1 - \beta M_j^{1-\theta})^{1/(1-\theta)}}{1 - \beta \bar{M}^{-\theta} H_j} = \frac{\bar{M}}{1 - \beta \bar{M}^{1-\theta}} \sum_{j=1}^N \omega_j (1 - \beta M_j^{1-\theta})^{1/(1-\theta)} \quad (\text{A5})$$

which can be used in the expression for welfare in equation (A1) to obtain the result in the main text.

References

- Athanasoulis, Stefano G., and Eric van Wincoop, 2000, "Growth Uncertainty and Risksharing," *Journal of Monetary Economics*, Vol. 45, pp. 477–505.
- Backus, David K., Patrick J. Kehoe, and Finn E. Kydland, 1994, "Dynamics of the Trade Balance and the Terms of Trade: The J-curve?" *American Economic Review*, pp. 84–103.
- 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.
- Cole, Harold L., and Obstfeld, Maurice, 1991, "Commodity trade and international risk sharing: How much do financial markets matter?" *Journal of Monetary Economics*, August, Vol. 28, pp. 3–24.
- 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, 2004, "Trade, Finance, Specialization and Synchronization," *Review of Economics and Statistics*, August, Vol. 86, No.3.
- Imbs, Jean and Paolo Mauro, 2007, "Pooling Risk Among Countries", IMF Working Paper 07/132.
- Jones, Benjamin, and Benjamin Olken, 2008, "The Anatomy of Start-Stop Growth," *Review of Economics and Statistics*, August, Vol. 90, No. 3

- 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/>.
- Kehoe, P., and F. Perri, 2002, "International business cycles with endogenous incomplete markets," *Econometrica*, Vol. 70, pp. C907–C28.
- Lewis, Karen K., 1996, "Consumption, Stock Returns, and the Gains from International Risk-Sharing," NBER Working Paper W5410 (Cambridge, Massachusetts: National Bureau of Economic Research).
- 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., 2000, "Future Possibilities in Finance Theory and Finance Practice," Harvard Business School Working Paper No. 01–030.
- 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.
- Tesar, Linda L., 1993, "International risk-sharing and nontraded goods," *Journal of International Economics*, Vol. 35, pp. 69–89.
- Van Wincoop, Eric, 1999, "How Big Are Potential Welfare Gains from International Risksharing?" *Journal of International Economics*, Vol. 47, pp. 109–135.

Table 1. Welfare Gains

	Minimum δ_j	Median δ_j	GW for Sub-Sample Indicated
Full Sample	0.55 United States	4.10 Dominican Republic	1.87
Advanced Countries	0.40 United States	0.89 Australia	0.69
Emerging Markets	1.67 Colombia	4.97 Malaysia	4.39
Developing Countries	1.00 Bangladesh	7.17 Madagascar	7.37

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 $\beta = 0.95$. The list of sub-samples is provided in Appendix Table 1. GDP data are from the World Bank's World Development Indicators.

Table 2. Gains from Risk Pooling Among Countries

	(1)	(2)	(3)	(4)
	Median σ_j	Min $\bar{\sigma}$ for all p	Median δ_j for all p	GW
All Countries (Pooling with any country)	3.62	0.54	3.88	1.87
<i>Costs of Weak Enforcement</i>				
Excellent Enforceability Countries (pooling with any country)	2.11	0.61	1.17	0.85
Below Excellent Enforceability Countries (pooling with any country)	4.41	0.52	5.90	5.12
Excellent Enforceability Countries (pooling only with excellent enforceability countries)	2.11	0.87	1.03	0.76
Above Average Institutional Quality (pooling with any country)	2.61	0.61	1.94	1.06
Below Average Institutional Quality (pooling with any country)	4.26	0.52	5.32	5.42
Above Average Institutional Quality (pooling only with above average institutional quality countries)	2.61	0.82	1.74	0.99
<i>Costs of Regional Constraints</i>				
European Union (pooling only with advanced)	1.84	0.89	0.75	0.77
European Union (pooling only with EU)	1.84	1.05	0.68	0.55
Asian Emerging (pooling only with emerging)	3.62	1.09	3.60	3.54
Asian Emerging (pooling only with Asian emerging)	3.62	1.84	2.98	2.94
Latin American Emerging (pooling only with emerging)	4.41	1.07	5.73	5.39
Latin American Emerging (pooling only with Latin American emerging)	4.41	1.90	4.96	4.08

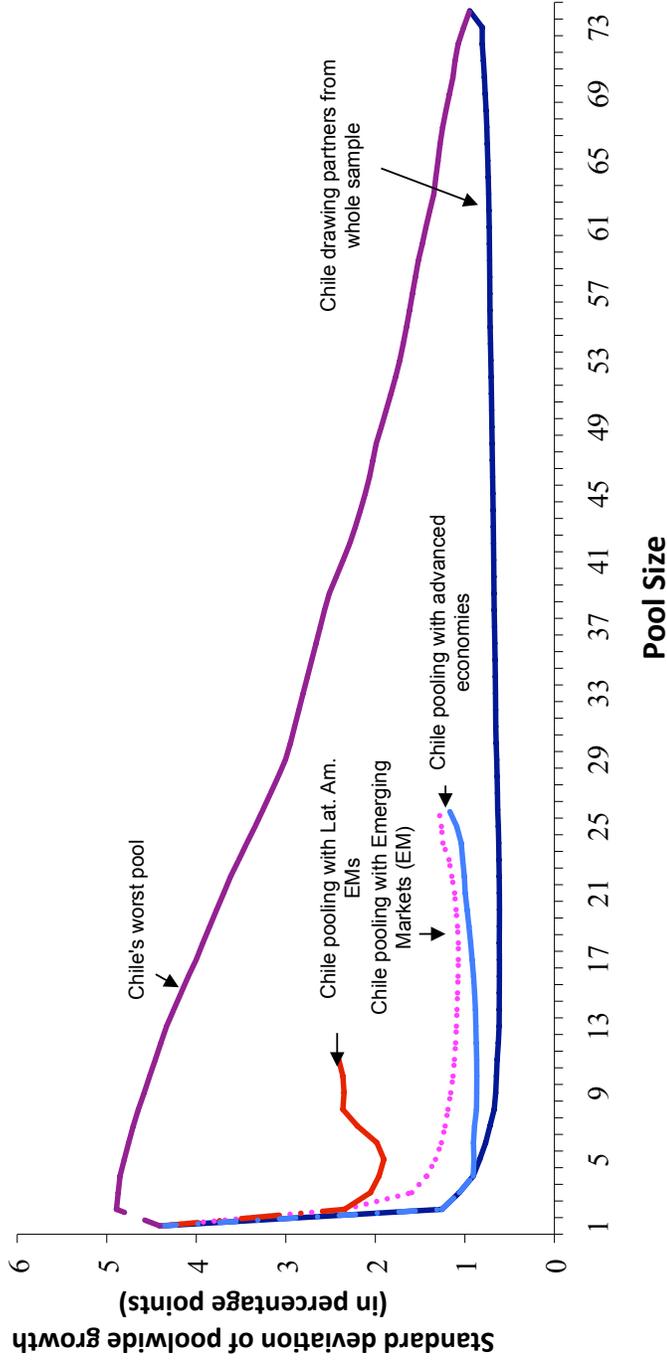
Notes: Column (1) reports the median of σ_j (across countries in the indicated sub-sample). Column (2) reports the median (across countries in the indicated sub-sample) of the lowest possible $\bar{\sigma}$. Column (3) reports the median δ_j (across countries in the indicated sub-sample) assuming growth rates fixed at 3 percent per annum. Column (4) reports GW assuming growth fixed at three percent per annum.

Table 3. Poolwide Volatility for Selected Groups

	APEC	ASEAN	CHIANG MAI	ECOWAS	EMU	EU	FLAR	MERCO SUR	NAFTA
Median σ_j	3.73	4.36	3.93	4.58	1.78	1.77	4.00	4.33	2.11
Minimum σ_j	2.00	3.65	2.01	3.07	1.36	1.36	2.19	3.51	2.00
$\bar{\sigma}_{Pool}$	1.29	3.47	1.40	2.19	1.19	1.11	2.48	3.09	1.81
$\bar{\sigma}_{Min}$	0.68	0.86	0.79	0.72	0.66	0.65	0.71	0.74	1.16

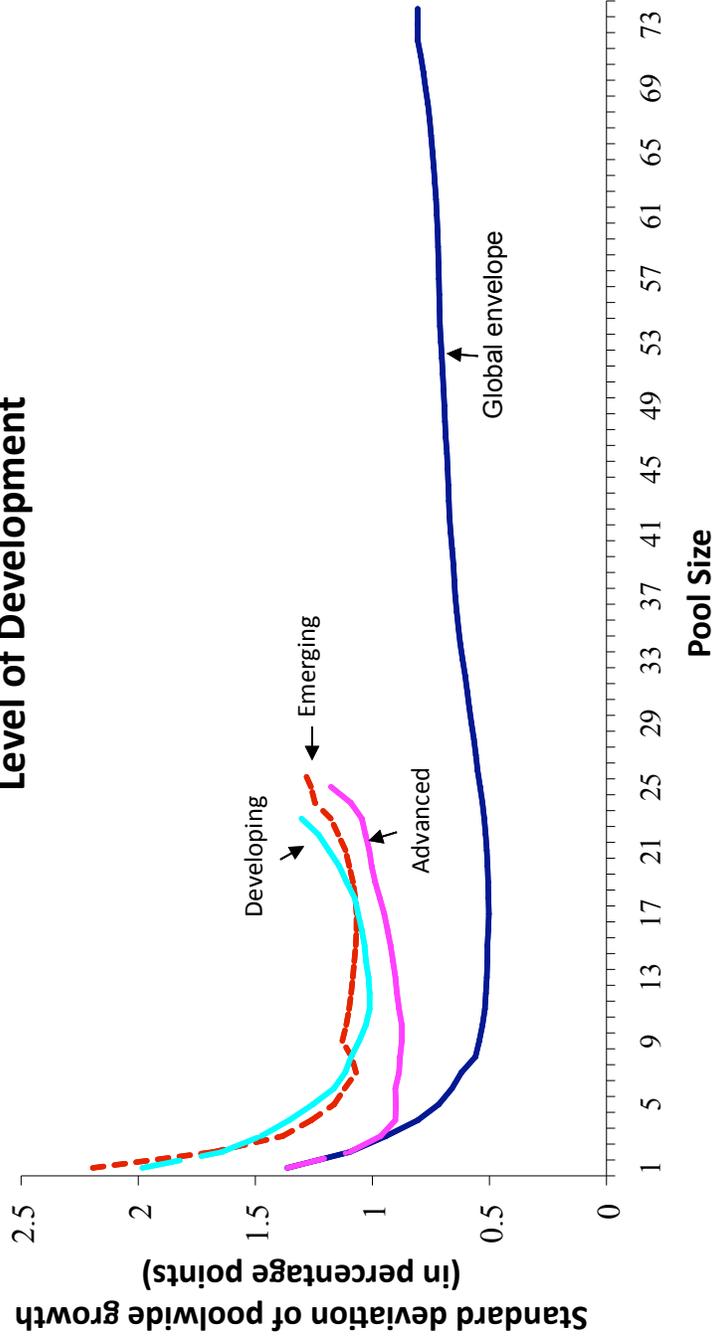
Notes: Minimum, median, and total welfare gains for sub-samples indicated, assuming the same expected growth (3 percent) across countries. The results assume $\theta = 2$, $\gamma = 5$, and $\beta = 0.95$. The list of sub-samples is provided in Appendix Table 1. GDP data are from the World Bank's World Development Indicators.

Figure 1. Chile - Benefits of Diversification Under Various Restrictions



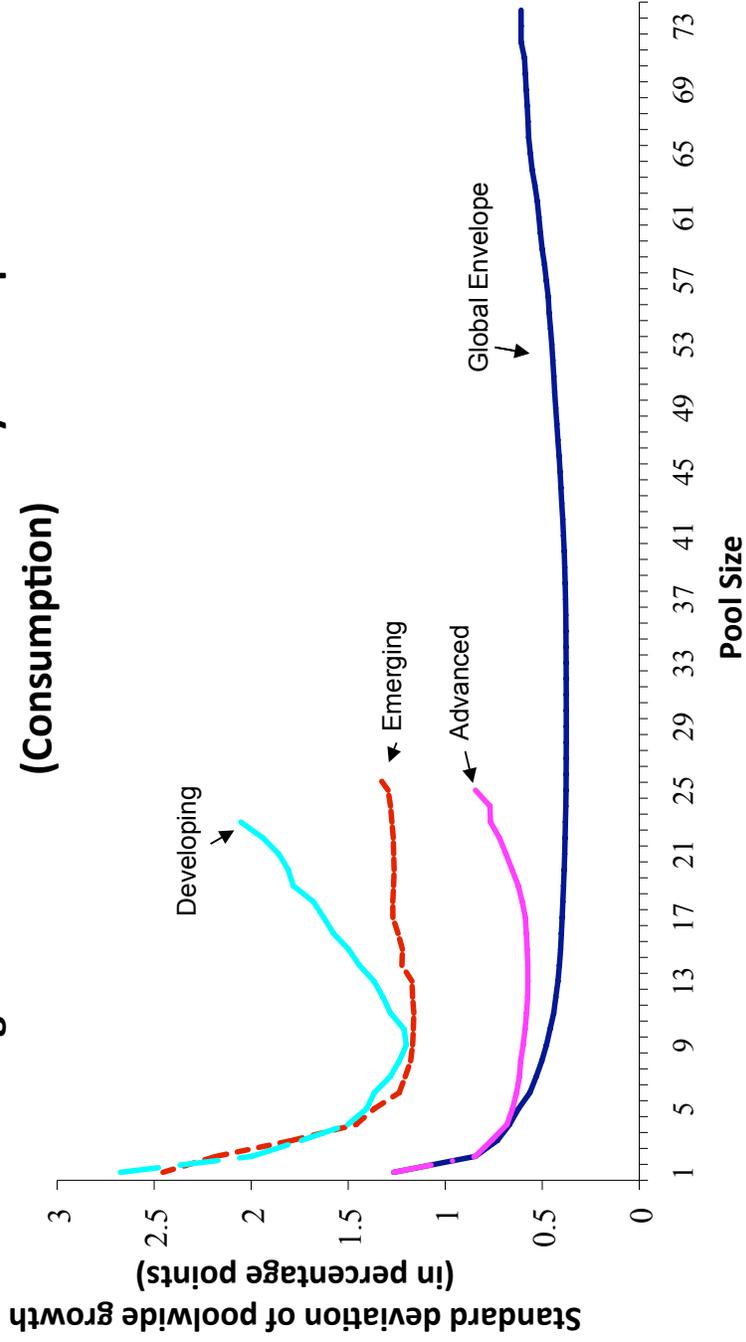
Notes: The figure reports the standard deviation of the growth rate of aggregate (poolwide) GDP for the pool (of each size) that yields the lowest standard deviation (or the highest, in the case of the top line). Each group is constrained to contain Chile. GDP data at purchasing power parity are drawn 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 the sub-samples, and by the approximation procedure described in the text for the full sample.

Figure 2. Lowest Poolwide Volatility Envelopes for Samples by Level of Development



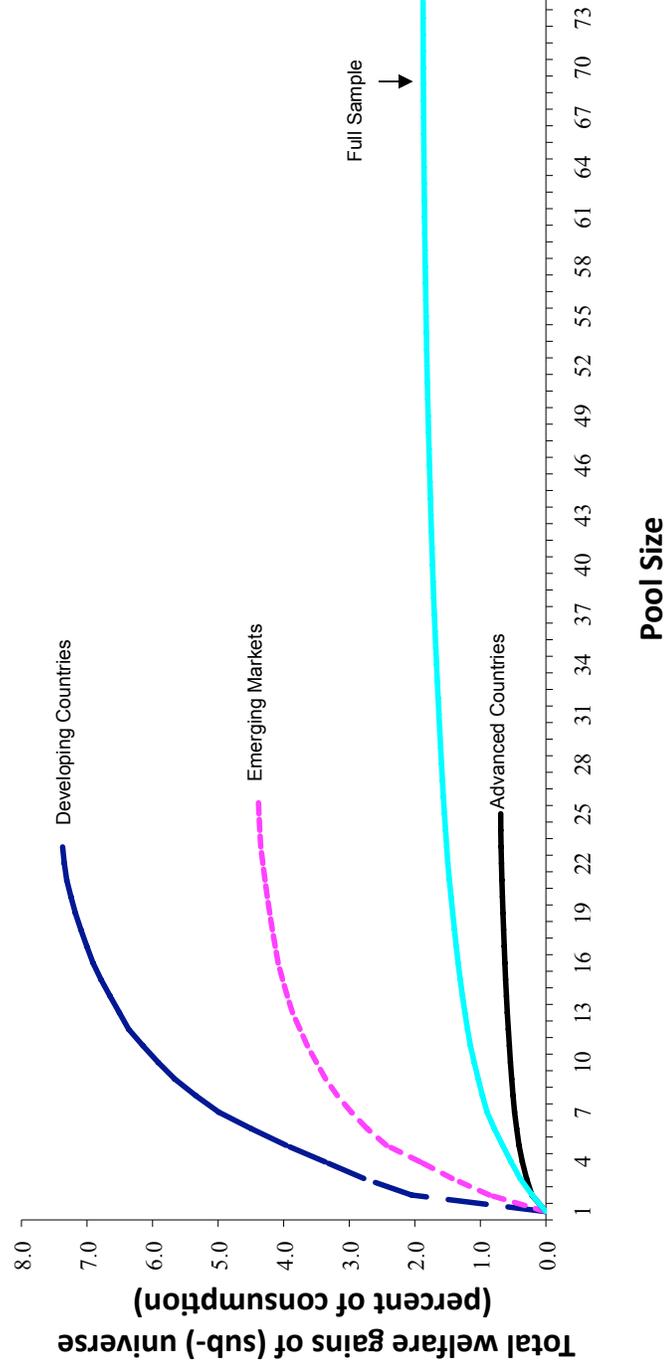
Notes: The figure reports the standard deviation of the growth rate of poolwide GDP for the pool (of each size) that yields the lowest standard deviation for each sub-sample. The pool yielding the lowest standard deviation is found by checking all possible combinations of countries and by running the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

**Figure 3. Lowest Poolwide Volatility Envelopes
(Consumption)**



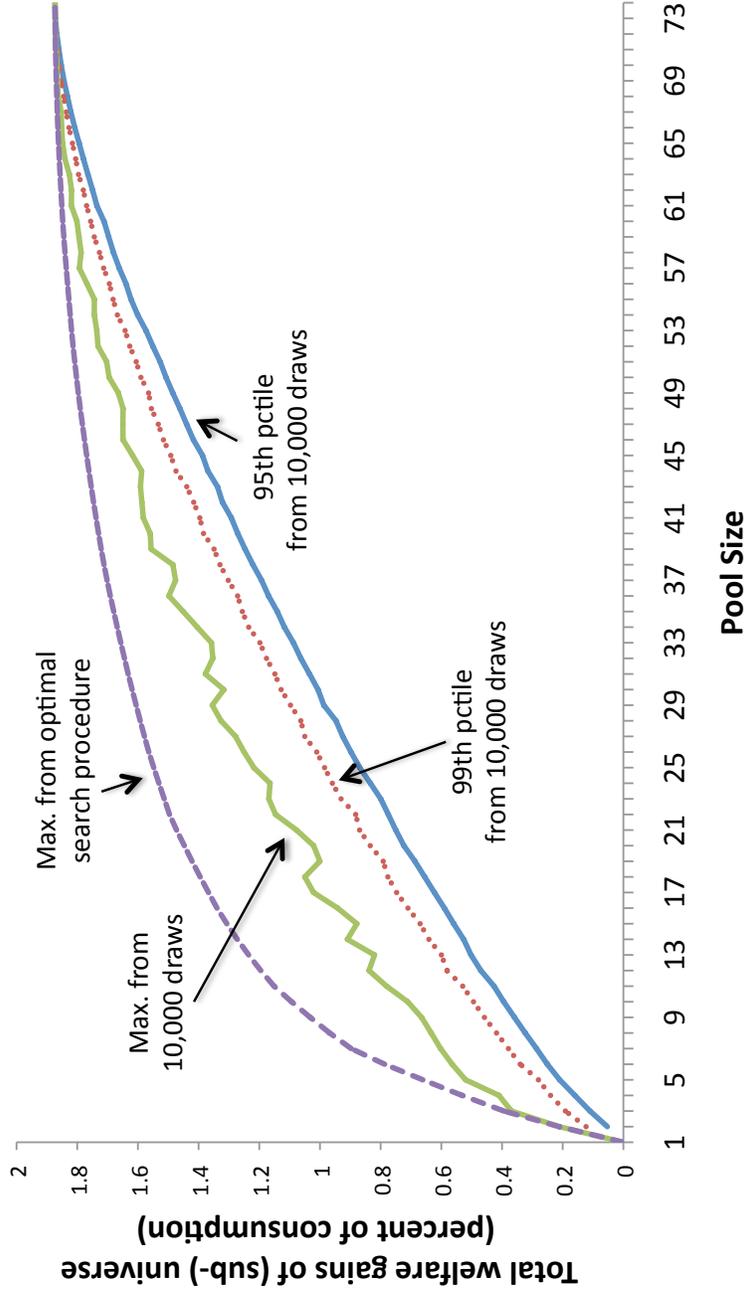
Notes: The figure reports the standard deviation of the growth rate of poolwide consumption for the pool (of each size) that yields the lowest standard deviation for each sub-sample. The pool yielding the lowest standard deviation is found by checking all possible combinations of countries and by running the approximation procedure described in the text. Consumption data are calculated using consumption shares from the World Bank's World Development Indicators. The consumption share for some years in Malawi is based on the IMF World Economic Outlook.

Figure 4. Pooling Gains



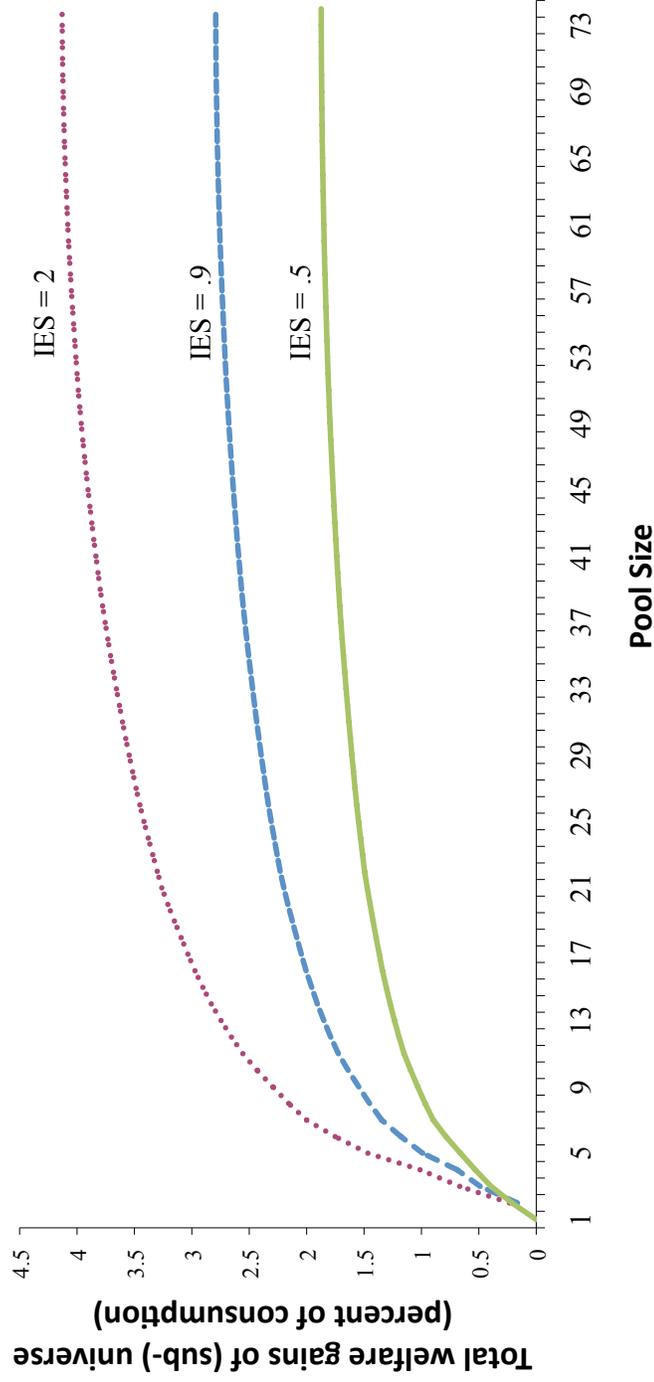
Notes: For each pool size, the figure reports the highest possible value for GW using the search procedure and GW formula described in the text. Welfare gains are computed assuming $\gamma = 5$, $\beta = 0.95$, $\theta = 2$, and a constant growth rate of three percent for all countries and pools. GDP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

Figure 5. Estimated and Actual Gains



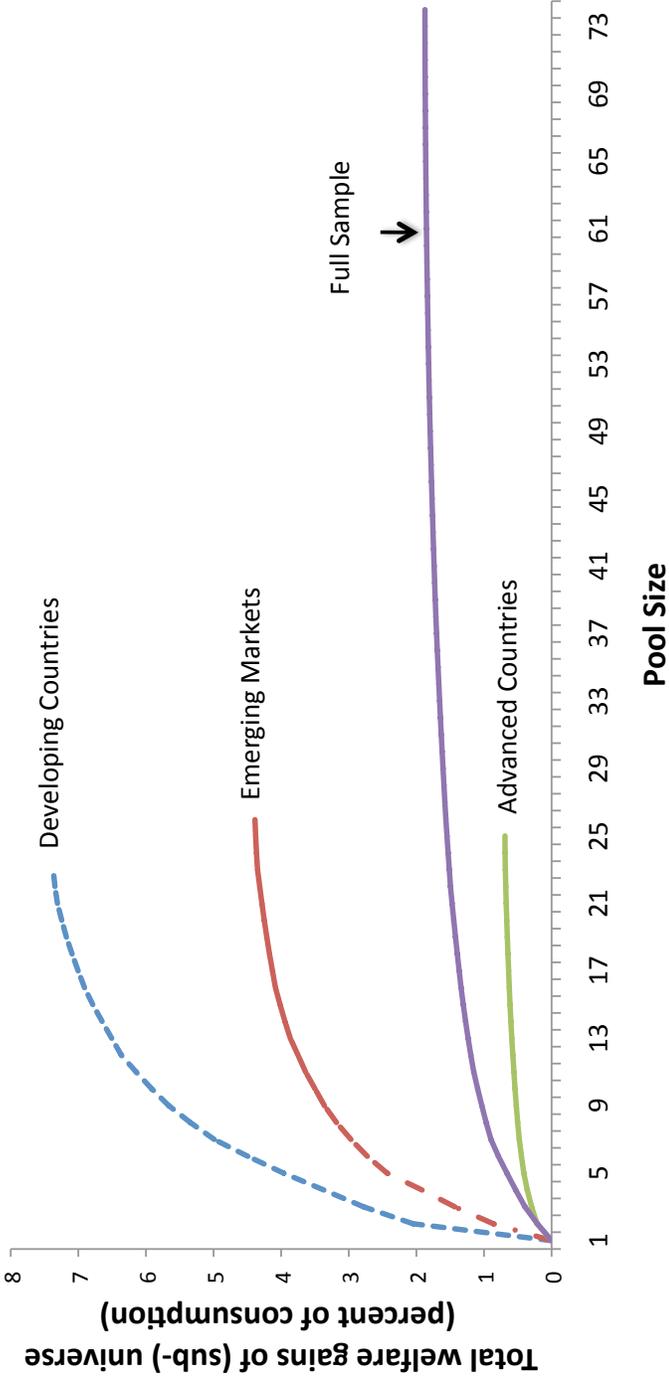
Notes: For each pool size, the figure reports the highest possible value for GW using the search procedure described in the text. In addition, the figure reports the maximum, 99th percentile, and 95th percentile obtained from 10,000 random draws. Welfare gains are computed using the GW formula given 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 6. Varying the IES

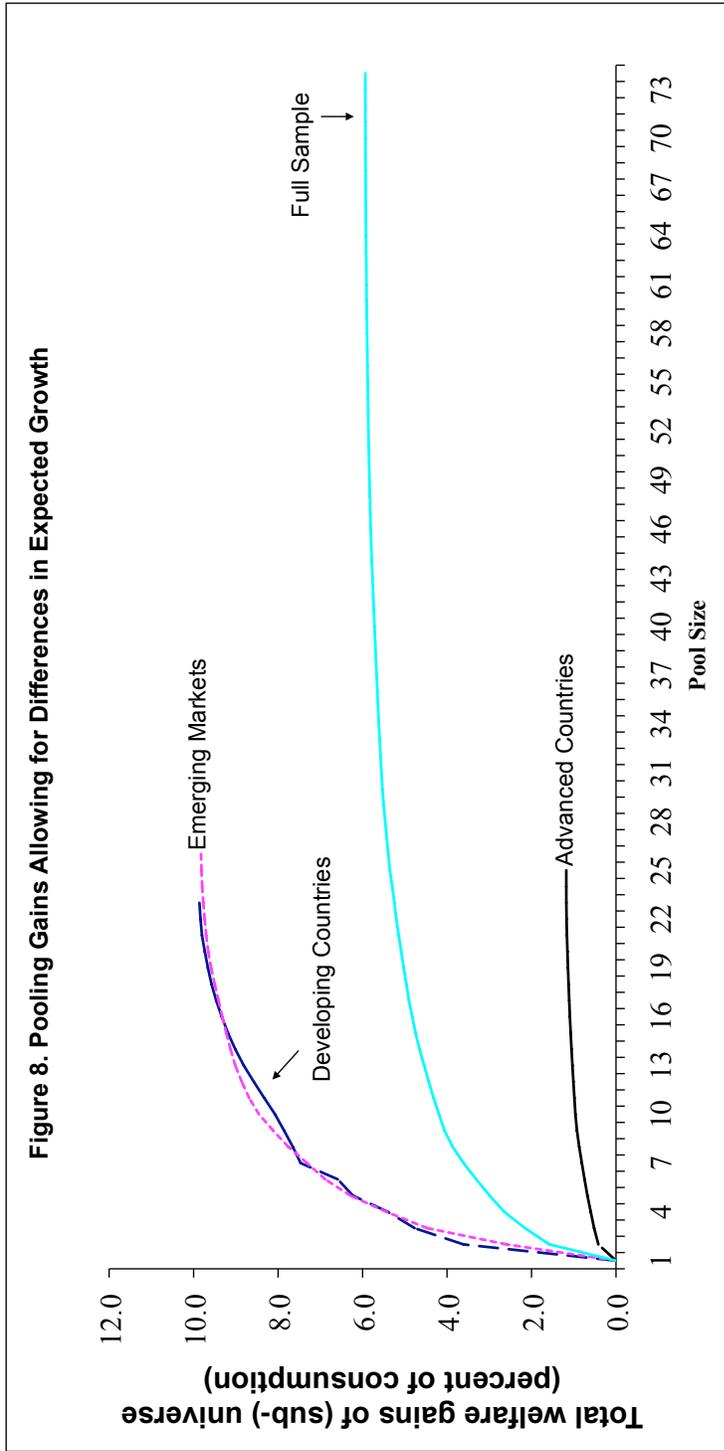


Notes: The figure reports the highest possible values for G^W using the search procedure and the G^W formula described in the text. The results assume $\gamma = 5$, $\beta = 0.95$, $\theta \in \{\frac{1}{2}, \frac{1}{0.9}, \frac{1}{0.5}\}$ as indicated, and a constant growth rate of three percent for all countries and pools. GDP data are from the World Bank's World Development Indicators.

**Figure 7. Pooling Gains
(Autarky Consumption vs. Pooled Output)**



Notes: For each pool size, the figure reports the highest possible value for GW. The welfare gains are computed using the GW formula in the text. The results assume $\gamma = 5$, $\beta = 0.95$, $\theta = 2$, and a constant growth rate of three percent for all countries and pools. Consumption data are calculated using consumption shares and GDP data from the World Bank's World Development Indicators. GDP data are from the World Bank's World Development Indicators.



Notes: For each pool size, the figure reports the highest possible value for GW using the search procedure and GW formula described in the text. The results assume $\gamma = 5$, $\beta = 0.95$, $\theta = 2$ and observed average growth rates over the period 1975–2004. GDP data are from the World Bank’s World Development Indicators.