



UNITED NATIONS
UNIVERSITY

UNU-WIDER

World Institute for Development
Economics Research

Research Paper No. 2005/57

Poverty and the Welfare Costs of Risk Associated with Globalization

Ethan Ligon*

September 2005

Abstract

‘Globalization’ implies change, and uncertainty over future change may affect household welfare. We use data on Lorenz curves over the last fifty years for a sample of 53 (mostly developing) countries. Treating each country-quintile-year as an observation, we first account for variation in consumption expenditures, finding that global shocks are of less importance than country-level shocks in explaining variation in country-quintile consumption growth. While poorer households experience more rapid consumption growth than do wealthier households, they also bear much more risk. However, we find no evidence that this greater risk is related to globalization.

Keywords: globalization, risk, growth, inequality, shocks, consumption

JEL classification: F02, O47, E21

Copyright © UNU-WIDER 2005

* Department of Agricultural and Resource Economics, University of California, Berkeley
email: ligon@are.berkeley.edu

This is a revised version of the paper originally prepared for the UNU-WIDER project meeting on The Impact of Globalization on the World’s Poor, directed by Machiko Nissanke and Erik Thorbecke, held in Helsinki 29-30 October 2004.

UNU-WIDER acknowledges the financial contributions to the research programme by the governments of Denmark (Royal Ministry of Foreign Affairs), Finland (Ministry for Foreign Affairs), Norway (Royal Ministry of Foreign Affairs), Sweden (Swedish International Development Cooperation Agency—Sida) and the United Kingdom (Department for International Development).

ISSN 1810-2611 ISBN 92-9190-737-5 (internet version)

Tables and Figures appear at the end of the paper.

The World Institute for Development Economics Research (WIDER) was established by the United Nations University (UNU) as its first research and training centre and started work in Helsinki, Finland in 1985. The Institute undertakes applied research and policy analysis on structural changes affecting the developing and transitional economies, provides a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and promotes capacity strengthening and training in the field of economic and social policy making. Work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

www.wider.unu.edu

publications@wider.unu.edu

UNU World Institute for Development Economics Research (UNU-WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

1. INTRODUCTION

Though there's great disagreement over its costs and benefits, "globalization" embodies the notion of change—greater financial integration, reductions in barriers to trade and increases in the exchange of information are expected to lead to changes in income, prices, patterns of employment, and so on. Broadly speaking, proponents of these changes tend to focus on the value of improvements in aggregate quantities, such as GDP predicted by theory, while opponents fear increases in inequality both within and across countries, and tend to disparage the benefits associated with exposing poor households to the "discipline" of free, competitive markets.

These latter concerns have to do with the risks households bear. Even if one accepts the argument that aggregate measures of income will increase, if a second consequence of globalization is increased uncertainty then risk-averse households may resist globalization's onslaught; poor households in particular may be greatly harmed by increases in variation in income and expenditures.

Closely related to the issue of what risks households bear is the issue of what risks they share. Indirect evidence of risk-sharing can be had by examining the extent to which households' consumptions co-move. In this paper we don't have household-level data to exploit. Instead our basic data has to do with the aggregate consumption of households within a given consumption quintile within a given country, or what we'll term a "country-quintile." Using a panel of such observations on five quintiles in 53 countries over 50 years, we infer the aggregate risk borne by the households within a country-quintile risk by first measuring the shocks experienced by these households, and then using the empirical distribution of these shocks to estimate risk. What is a shock? For our present purposes, we infer that a shock has occurred when we see deviations from trend in country-quintile consumption growth, where the trend in consumption growth is estimated separately for each quintile.

We distinguish among three different categories of shocks to consumption. First among these are *global shocks*. These can be inferred by measuring the common component of changes in consumption growth across all country-quintiles. Though we don't attempt to identify the cause of any particular shocks in this paper, instances of shocks which have a global effect might include technological changes; changes in global temperature or precipitation; major wars; or changes in U.S. monetary policy. The second category of shocks we consider are *country-specific* shocks. These are shocks which have a common effect on consumption growth for all the country-quintiles within a given country. Examples of country-specific shocks might include changes in national policies; regional wars or civil insurrection; unusual weather patterns; or a currency crisis.

The third category of shocks we consider we'll term *globalization shocks*. These are shocks which effect households in different consumption quintiles differently, but which are nonetheless shared across countries, like global shocks. Thus globalization shocks allow us to measure the extent to which the fortunes of, e.g., poor households are linked across national boundaries. Examples of globalization shocks might include things like skill-biased technological change, which might have differential effects on wages and consumption growth across different parts of the consumption distribution; or alternatively more or less local shocks when markets for credit or insurance are segmented by wealth, rather than by location. We employ the term "globalization shock" rather than, say, 'distributional shock' or 'quintile-specific shock' because we wish to emphasize the point that such shocks affect the distribution of consumption within countries, but do so *globally*, since the distribution of consumption within each country changes in similar ways across all countries. A prerequisite for any shock to have a common effect on a given consumption quintile across all countries consumption across countries is that there be some kind of economic linkages of households within a quintile across countries.

Our aim in this paper is to show how changes over time in generalized Lorenz curves (or the inverse consumption distribution) for expenditures can be used to estimate the welfare costs associated with risk. We then use data on such Lorenz curves to actually estimate a lower bound on these costs in selected countries where adequate data are available. We're able to show how risk is borne across expenditure quantiles, and so are able to indicate what share of risk is borne by poor households, and to estimate the "globalization risk" which may affect differently situated households across national boundaries.

The remainder of this paper is organized as follows. In Section 2, we first summarize the predictions of theory for household-level consumption when markets are complete and households have time-separable preferences exhibiting constant relative risk aversion. We provide a precise definition of risk used in this paper; the measure we provide is cardinal, and is consistent with the ordinal approach taken by Rothschild and Stiglitz (1970). We next provide methods to estimate the risk borne by households within a particular consumption quintile, and show that this measure can be easily decomposed into risk from various sources. Section 3 describes the data on consumption and inequality we employ, while Section 4 reports the results of applying the methods described in Section 2 to the data described in Section 3. Section 5 concludes.

2. METHODS

2.1. Consumption with Complete Markets. Consider a population of households indexed by $i = 1, 2, \dots, n$, and denote household i 's realized consumption (per adult equivalent) in year t by c_{it} . Household i derives utility from c_{it} according to the momentary utility function

$$U(c) = \begin{cases} \log(c) & \text{if } \gamma = 1 \\ \frac{c^{1-\gamma}-1}{1-\gamma} & \text{otherwise.} \end{cases}$$

Here the parameter γ has the interpretation of being equal to the household's relative risk aversion (Arrow, 1965; Pratt, 1964).

We assume that the household may be subject to shocks which affect the resources it has available at any particular date. Remarkably, so long as $\gamma > 0$ and markets are complete, it's straightforward to show that the household's consumption won't depend on these shocks, beyond the effects of these shocks on the *aggregate* resources available to all households. In particular, household i 's realized consumption at date t can be written as the product of a fixed household factor λ_i and an aggregate quantity μ_t which may vary over both dates and (aggregate) states, but which is common to all households, so that when markets are complete we have

$$(1) \quad c_{it} = \lambda_i \mu_t$$

(Wilson, 1968). Importantly for the exercise of this paper, a similar relationship holds for any subset of households (say, all the households within a particular country). In particular, for any set of household indices H we can use (1) to write

$$\bar{c}_{Ht} \equiv \sum_{i \in H} c_{it} = \mu_t \sum_{i \in H} \lambda_i,$$

so that, for example, complete markets implies that aggregate consumption for the population of any entire country will depend only *global* shocks which have an equal effect on the populations of all nations.¹

Of course, the notion of complete markets is a hypothesis, not a statement of fact, and happens to be a hypothesis which has been rejected in a variety of (though not all) environments.² Still, it seems to be true that income is much more variable than consumption in most environments, and regardless of the level of aggregation. This suggests, as Mace (1991) observes, that the predictions of the complete markets model are the appropriate “benchmark” for research on risk.

2.2. Predicting Consumption. Following Mace's suggestion, we proceed by initially adopting the hypothesis of global complete markets, so that every households'

consumption varies over time only in response to truly global shocks. In particular, let \hat{c}_{it}^0 be the consumption predicted by a model with global complete markets. Then by taking the logarithm of both sides of (1), we obtain

$$(2) \quad \log c_{it} = \log \hat{c}_{it}^0 + \epsilon_{it}^0 = \alpha_i + \eta_t + \epsilon_{it}^0,$$

where $\alpha_i = \log \lambda_i$ is a variable related to household wealth (unchanging because perfectly insured) and $\eta_t = \log \mu_t$ is a measure of the global resources available at t . With household panel data we could estimate the latent variables α_i and η_t ; any systematic deviation from the predictions of the global complete markets model would show up in the residuals ϵ_{it}^0 .

As an example of the kind of deviation we might expect, suppose that markets for risk were complete within countries, but not across them. Then we would expect the residuals ϵ_{it}^0 to reflect this, exhibiting a high correlation across households within a country, but zero correlation across countries. A simple way to test this would be to add an additional set of latent variables to the estimating equation (2). In particular, let the set of countries be indexed by $\ell = 1, 2, \dots, C$, with ℓ_i denoting the index of the country in which household i is resident. Let \hat{c}_{it}^1 be the consumption predicted when we have complete markets *within* countries, but not necessarily across countries, so that

$$\hat{c}_{it}^1 = \lambda_i \mu_t^{\ell_i},$$

and (using notation analogous to that above)

$$\log c_{it} = \log \hat{c}_{it}^1 + \epsilon_{it}^1 = \alpha_i + \eta_t + \eta_t^{\ell_i} + \epsilon_{it}^1,$$

with the additional restriction that $\sum_{c=1}^C \eta_t^c = 0$. Estimating this regression would give us a simple test of the hypothesis of complete global risk markets—if the latent variables $\{\eta_t^\ell\}$ are jointly significant, then we can reject the null hypothesis of complete global risk markets. If *some* risk is shared at the global level, then estimates of η_t will be significant, and an analysis of the variance accounted for by the global and

country-specific latent variables will provide a measure of the importance of global versus country-level risk sharing.

For our present exercise there’s nothing particularly special about grouping households into countries—any grouping based on some observed fixed characteristics could also serve to construct tests of risk sharing. For example, one could extend the argument above to test country-level risk sharing by adding latent variables to capture, say, province-level shocks, or add indicators based on initial landholdings as in Townsend (1994), initial wealth as in Jalan and Ravallion (1996), or network membership as in Dercon and de Weerd (2004). However, we lack household-level panel data, and so these options are foreclosed. One question of interest for us revolves around the amount of risk borne by households of different wealth levels, but if risk sharing is imperfect then these wealth levels will change over time, and we won’t be able to track *how* they change.

We proceed as follows. First, we let q index consumption quintiles, and let (ℓ_i, q_{it}) denote the country and consumption quintile containing household i at time t (so that, e.g., if household i is in the bottom consumption quintile of country 2 at time t , then $(\ell_i, q_{it}) = (2, 1)$). Adopt the null hypothesis that there’s perfect risk sharing within each country, but not necessarily across countries. As a consequence, households won’t move across quintiles within a country, and so initial quintile membership is a fixed characteristic under the null hypothesis.³ Let $Q_q = \{i | q_{i0} = q\}$ denote the collection of sets of country-quintile household indices at time zero. Using the aggregation result described above, we have the prediction that

$$\hat{c}_t^{(\ell, q)} \equiv \sum_{j \in \{Q_q, L_\ell\}} c_{jt} = \mu_t^{(\ell, q)} \sum_{j \in \{Q_q, L_\ell\}} \lambda_j,$$

yielding the estimating equation

$$\log \bar{c}_t^{(\ell, q)} = \alpha^{(\ell, q)} + \eta_t + \eta_t^\ell + \bar{\epsilon}_{(\ell, q)t}^1.$$

Here the terms $\{\alpha^{(\ell,q)}\}$ are country-quintile “fixed effects,” and capture variation in the expected level of log consumption across countries and quintiles, but not across time. The terms $\{\eta_t\}$ capture the average impact of common (and hence “global” in our terminology) shocks on all country-quintiles, while the terms $\{\eta_t^\ell\}$ capture country-level variation over time in the log consumption of quintiles within the country.

2.3. Tests for ‘Globalization’. One encounters claims in the globalization literature of the form: ‘Globalization [however defined] is responsible for increasing/decreasing poverty/inequality.’⁴ Regardless of the direction of the supposed effect, any of these claims seems to involve there being some correlation in household level consumption within a given consumption quantile across countries. With the foregoing, we have available a simple test. The question is whether or not a household’s position in the consumption distribution has any power to explain consumption outcomes for that household after controlling for global and country-level shocks. This question suggests the estimating equation

$$(3) \quad \log \bar{c}_t^{(\ell,q)} = \alpha^{(\ell,q)} + \eta_t + \eta_t^\ell + \phi_t^q + \bar{\epsilon}_{(\ell,q)t}^2,$$

where $\{\phi_t^q\}$ is a collection of latent variables which captures the effects of shocks which affect, e.g., all poor households across the entire globe. If these are collectively significant, then we have evidence of global shocks which are specific in their effect on households of differing wealth levels. If the apparent magnitude of these shocks is increasing over time, then we may regard this as a measure of the progress of one sort of globalization.

2.4. Defining Risk. We take a utilitarian approach to defining a measure of the risk households face in any given period. As before we suppose there to be a finite population of households indexed by $i = 1, 2, \dots, n$. Following Ligon (2004) and Ligon and Schechter (2003) we define the risk faced by the household at t by the

function

$$R_{it} = U(\mathbb{E}c_{it}) - EU(c_{it}).$$

Taking expectations of a concave utility function has the effect of making risk depend not only on the mean of a household's consumption, but also on variation in consumption.

The measure R_{it} , which measures the risk faced by household i , is consistent with the ordinal measures of risk proposed by Rothschild and Stiglitz (1970). Further, this risk measure can usefully be further decomposed into distinct measures of risk corresponding to different sources. For example, let $\mathbb{E}(c_{it}|\mu_t)$ denote the expected value of consumption c_{it} conditioned on knowledge of the aggregate global shock μ_t . Then we can rewrite the risk facing household i as

$$R_{it} = [U(\mathbb{E}c_{it}) - EU(\mathbb{E}(c_{it}|\mu_t))] + [EU(\mathbb{E}(c_{it}|\mu_t)) - EU(c_{it})].$$

Here the first term expresses the *global risk* facing the household, while the second filters out the global component of risk to leave an unexplained component of risk borne by the household.⁵

An obvious next step follows our discussion in Section 2.2 to extend this decomposition to obtain measures of the risk due to country-level shocks, and to global shocks peculiar to different parts of the wealth distribution. Let μ_t^ℓ denote the shocks which affect country ℓ at date t , and let ν_t^k denote the globalization shocks which affect quintile k (across countries) in year t . Then we have

$$\begin{aligned} (4) \quad R_{it} = & [U(\mathbb{E}c_{it}) - EU(\mathbb{E}(c_{it}|\mu_t))] \\ & + [EU(\mathbb{E}(c_{it}|\mu_t)) - EU(\mathbb{E}(c_{it}|\mu_t, \mu_t^{\ell_i}))] \\ & + [EU(\mathbb{E}(c_{it}|\mu_t, \mu_t^{\ell_i})) - EU(\mathbb{E}(c_{it}|\mu_t, \mu_t^{\ell_i}, \nu_t^{k_i}))] \\ & + [EU(\mathbb{E}(c_{it}|\mu_t, \mu_t^{\ell_i}, \nu_t^{k_i})) - EU(c_{it})]. \end{aligned}$$

The first line of this expression gives us an estimate of global risk, the second line an estimate of country risk, the third an estimate of what we'll call "globalization risk", and the final term remaining risk.

2.5. Estimating Risk. Now, to operationalize our measure of risk we need to devise an estimator of the conditional expectations operators $E(c_{it}|\cdot)$ which appear in (4). For this we turn to the estimators described in Section 2.2; using least squares to estimate the regressions suggested there gives us predictions of the value of the log of country-quintile level consumptions. However, these equations are estimated with error, so to construct predictions on the *level* of consumption expenditures we also need to characterize the distribution of the errors of these regressions.

To estimate these conditional expectations using household level data, we would proceed in three steps. First, we estimate the innermost conditional expectations. Let $Z_{it}(k)$ be a matrix of variables pertaining to household i in an information set z_k^t , with z_k^t monotonically increasing in both k and t (in the sense that $z_k^t \subseteq z_{k'}^{t'}$ whenever $k' \geq k$ and $t' \geq t$). Then, using least squares, we estimate a sequence of parameters $\{\delta_k\}$ from

$$\log c_{it} = Z_{it}(k)\delta_k + v_{it}(k),$$

$k = 1, \dots, K$ where $v_{it}(k)$ are disturbances associated with the k th estimating equation; note that these will include the negative of the logarithm of any multiplicative measurement error associated with c_{it} . Letting $\hat{\delta}_k$ denote the estimated parameters from the k th regression, and $\hat{v}_{it}(k)$ the residuals, it follows that

$$E(c_{it}|Z_{it}(k)) = \exp(Z_{it}(k)\hat{\delta}_k)E(e^{\hat{v}_{it}(k)}|Z_{it}(k)).$$

We estimate the factor $E(e^{\hat{v}_{it}(k)}|Z_{it}(k))$ via a second regression, again using least squares,

$$(5) \quad e^{\hat{v}_{it}(k)} = Z_{it}(k)\phi_k + w_{it}(k),$$

thus constructing an estimate of the inner conditional expectations of c_{it}

$$(6) \quad \hat{c}_{it}^{(k)} = \exp(Z_{it}(k)\hat{\delta}_k)Z_{it}(k)\hat{\phi}_k,$$

where $\hat{\phi}_k$ denotes the least square estimates of ϕ_k in (5).⁶

Now, of course we don't have household level data. But note that if households within some subgroup (such as country-quintiles) have perfect within group risk sharing, then a similar set of restrictions will be satisfied. If, however, risk sharing within this group *isn't* perfect, then risk not explained by the group level variables will appear in the final "remainder" term of the risk decomposition. Our approach here is simply to ignore this final term (as we must given our aggregate data), but to remember that households are likely to bear additional risk beyond that we can account for using our aggregate data.

3. DATA

We construct an unbalanced panel dataset comprising data on consumption expenditures, where the unit of observation is a quintile-country-year. To construct this panel we use data from two sources.

The first is the "World Income Inequality Dataset" (henceforth "WIID" UNU/WIDER, 2005), which (despite its name) contains data on the distribution of consumption expenditures for a variety of countries and years. The WIID is itself descended from a dataset compiled and described by Deininger and Squire (1996), who use the data to explore the connection between income growth and inequality. We are unfortunately not able to use more than a fraction of the available WIID data. The WIID is a compilation of data from many different household level surveys. We first restrict our attention to data drawn from surveys that sampled the entire population of the country, without any *ex ante* restrictions on geographical coverage, or on the age or other demographic characteristics of the respondents. Much of the database consists of data on the distribution of income rather than consumption, though it's the latter that matters for household welfare (and the distinction between the two is of

particular importance when considering risk, since a household with highly variable income may nonetheless have very smooth consumption). We further limit the data we employ by using only data for which the WIID records either quintiles or deciles of consumption expenditures. Data recorded in deciles is aggregated up to quintiles, to give us consistent units of observation. Finally, because our methods require looking at *changes* in the distribution of consumption, we limit ourselves to using data for countries which have expenditure surveys in at least two different years.

The WIID reports only *shares* of expenditures, and not levels. Because we want to be able to track changes in the levels of expenditures over time, we augment the WIID data with data on aggregate consumption from version 6.1 of the Penn World Tables (henceforth “PWT”) (Summers and Heston, 1991). To construct our aggregate measure of consumption we use the data on chain-weighted GDP per adult-equivalent multiplied by the consumption share (variables labelled “RGDPEQA” and “CC” in the PWT), so that all consumption data is expressed in terms of 1996 US dollars. It’s worth noting that the estimates of consumption’s share reported in the PWT are constructed from aggregate national income and product accounts, and that estimates of consumption so constructed are likely to vary from estimates derived from household-level data. Though the latter sort of estimates might be better suited to the present study, we know of no reason to think that differences between the two will be related to risk, and so use the PWT estimates without further apology.

[Figure 1 about here.]

Though the WIID includes data satisfying our criteria for Algeria in 1988 and 1995 and for Cambodia in 1994, the PWT doesn’t have data on the variable RGDPEQA for these countries in these years. Rather than dropping these countries from our analysis, we use data from the PWT on real per capita GDP (RGDPCH), adjusting for the difference between adult equivalents and per capita measures of GDP by assuming that the ratio between RGDPEQA and RGDPCH data is unchanging over time for each country.

[Table 1 about here.]

[Figure 2 about here.]

Finally, after correcting some obvious minor errors in the WIID data (inconsistent units for shares and missing single quintiles for some country-years), we are left with more-or-less consistent data for 53 countries over varying periods. This is not a random sample of countries—the WIID is much more likely to report income (rather than expenditure) data for wealthy countries, and so the bulk of our countries are relatively poor. These countries and the years for which we have usable data are reported in Table 1, while Figure 2 presents a histogram which reveals the distribution of the number of surveys per country in the usable sample.

4. RESULTS

In this section we follow the pattern laid down by Section 2, and turn our attention first to results having to do with the prediction of consumption expenditures, and then subsequently to the characterization of risk.

We implement a version of an estimator for the estimating equation (3). We draw upon the experience of Ligon and Schechter (2004), who find that estimating equations of this form using household level panel data can lead to very severe biases if the consumption process is in fact non-stationary. We adopt the weaker assumption that the consumption process is difference-stationary,⁷ and estimate a differenced form of the equation

$$\Delta \log \bar{c}_t^{(\ell,q)} = \Delta \bar{\phi}^q + \Delta \eta_t + \Delta \eta_t^\ell + \Delta \phi_t^q + \Delta \bar{\epsilon}_{(\ell,q)t}^2.$$

with various combinations of the right hand variables.

Here the latent variables $\Delta \bar{\phi}^q$ capture the common trends in consumption growth shared within each world consumption quintile—thus, if the poorest twenty per cent of households in each country in the sample have growth rates of consumption which are significantly different from the richest twenty per cent, that should be evident

from estimates of $\{\Delta\phi^q(1)\}$ in the regression

$$\Delta \log \bar{c}_t^{(\ell,q)} = \Delta \bar{\phi}^q(1) + \Delta \bar{\epsilon}_{(\ell,q)t}(1).$$

Table 2 reports results from this regression for our sample. The first row of this table simply gives estimates of the average growth rate for each of the world quintiles (20%, 60%, 80%, 100%). These are all positive, and average 2.6% across quintiles. The most interesting pattern is that growth rates fall with wealth—while the poorest quintile has an average growth rate of 4.3%, the wealthiest has an average growth rate of only 1.5%, significantly less than the growth rate of the poorest quintiles. This is the kind of evidence relied upon by Sala-i-Martin (2002) when he advances the claim that global economic growth has reduced poverty. Our result is consistent with this claim, albeit for a rather different sample. The fact that the poorest households within each country experience the highest rates of growth suggests that (averaged across countries) inequality will fall over time.

[Table 2 about here.]

[Table 3 about here.]

[Table 4 about here.]

Table 3 presents results akin to a traditional analysis of variance exercise. The row labeled “World Quintile” reports the proportion of variation in country-quintile consumption growth over time which can be explained simply by a set of 5 dummy variables indicating the (within-country) consumption quintile into which these observations fall. The first column shows the proportion of variance in consumption growth which can be explained by different sets of latent variables, each group in isolation. Thus, the first entry of the table is equal to the R^2 statistic of the regression the results of which are reported in Table 2. It’s worth noting that though that table reveals that there are no significant differences in average consumption growth by global quintile, the proportion of variance explained by these factors is only 0.5 per cent, far from significant. The contribution of global shocks (corresponding to

a complete set of year dummies) to an explanation of variance is much greater, at 4.4 per cent. However, this statistic is insignificantly different from zero, suggesting that this sample of mostly poor countries is not very well integrated into the world economy. are both important and are to some extent shared across the consumption distribution. This point receives further reinforcement when we add a complete set of country-year effects. These country-year effects account for a very significant 36 per cent of the total observed variation in consumption expenditures. The combined variation explained by quintile effects, global shocks, and country-level shocks amounts to 36.5 per cent. This leaves only a maximum of 63.5 per cent of variation in consumption growth which can be peculiar to within-country factors.

The final row of Table 3 reports the amount of variation explained by world quintile-year effects, or what we've termed "globalization shocks." By themselves these latent variables account for 11 per cent of the total variation in consumption growth. However, after *controlling* for country-shocks, these variables account for only a further 6 per cent of the variation in growth, an amount not significantly different from zero.

We next turn our attention to a consideration of the risk associated with the changes in consumption which we've tried to predict with the regressions reported above. Table 4 decomposes the total average risk borne by different country-quintiles into four pieces. With our preference parameter $\gamma = 2$ the figures reported in Table 4 can be interpreted not only as (100 times) the loss of utility associated with different sources of risk, but *also* as the percentage of expected consumption which households would be willing to sacrifice in exchange for elimination of that source of risk. Measured this way, the welfare loss for each of the quintiles ranges from 2.6 to 18.8 per cent—that is, that if all risk could be eliminated so that consumption expenditures grew at a constant rate equal to the figures reported in Table 2, then households across all these quintiles would be willing to sacrifice these proportions of all current and future consumption.

Consistent with the evidence in Table 2, we see in the present table that the effect of *global* shocks on risk is quite small. Much more important is country-specific risk. However, the variation in consumption expenditures explained by country-specific factors turns out to be *negatively* correlated with other kinds of shocks. We interpret this as evidence that in fact country-level shocks provide some measure of insurance against the other kinds of shocks households experience. The value of this insurance varies considerably across quintiles, with the poorest quintiles benefitting most (average value equal to 4.3 per cent of consumption), and the wealthiest least (average value equal to 2.6 per cent of consumption).

Globalization risk associated with common shocks to quintiles across countries (after controlling for country level shocks) is very small, with only the poorest quintile bearing any risk greater than 1 per cent. Thus, while there's some slight evidence that a small share of the shocks borne by poor households are shared across countries, beyond this there's no evidence that any given quintile has integrated into the world economy to any great extent, since neither global shocks (common to all quintiles across countries) nor globalization shocks (common to particular quintiles across quintiles) seem to explain much variation in either the rate of growth in consumption or in the risk borne by quintiles in different countries.

5. CONCLUSION

In this paper we've sought to account for variation in the consumption growth experienced by households in different parts of the consumption distribution across countries and time, and then to estimate the welfare loss associated with different sources of shocks.

We construct a (unbalanced) panel of countries for which data on the distribution of consumption expenditures is available. We find, first, that the poorest country-quintile has an expected rate of consumption growth which is considerably greater than that of the top twenty per cent. However, while significant, these constant growth trends account for only a very small proportion of observed variation in growth

rates across quintiles. More variation (roughly one twentieth) is explained by global shocks. Over a third is explained by country-specific shocks, which affect all quintiles within that country more or less equally. Only 6 per cent of the remaining variation is accounted for by factors which are common to all households within a given global consumption quantile. We regard this as evidence against the hypothesis that the forces of globalization have had an important effect on global inequality, even conditioning on changes in inequality *across* countries.

We exploit these results to predict levels of consumption for different country-quintiles, and follow Ligon and Schechter (2003) in exploiting these predictions to estimate the *risk* borne by households within a country-quantile. We find that the total risk is large, with households willing to sacrifice on average 8.4 per cent of their expected consumption in exchange for eliminating all risk. Interestingly, country-quintile specific risk (which also will include measurement error) turns out to be *negatively* related to country-level shocks; one can interpret this as evidence that country-level shocks actually provide some measure of insurance against other kinds of risk.

We close the paper with an important caveat. The estimates reported here are no better than the data used as inputs. The WIID project which compiles data on inequality should be commended for making systematic empirical research on inequality possible, but the data available on the distribution of consumption within countries is still sadly incomplete. Small changes in the data provided in the WIID or changes in e.g., decisions about whether to weight observations by population can have dramatic effects on results. Through the efforts of WIID and others we've come a long way since the days in which one had to rely on an incomplete *cross-section* of countries to draw inferences about changes in inequality over time (e.g., Kuznets, 1955), but attempting to treat the available data on country-level inequality as a representative panel as we've done here, is certainly bold and possibly foolhardy. Hopefully this is a problem that the passage of time will likely solve.

References

REFERENCES

- Arrow, K. J. (1965). *Aspects of the Theory of Risk Aversion*. Helsinki: Yrjö Jahnssoonin säätiö.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy* 99, 957–976.
- Deininger, K. and L. Squire (1996, September). A new data set measuring income inequality. *The World Bank Economic Review* 10(3), 565–591.
- Dercon, S. and J. de Weerd (2004). Risk-sharing networks and insurance against illness. Manuscript.
- Dollar, D. and A. Kraay (2002). Growth is good for the poor. *Journal of Economic Growth* 7, 195–225.
- Grimard, F. (1997). Household consumption smoothing through ethnic ties: Evidence from Cote d’Ivoire. *Journal of Development Economics* 53, 391–422.
- Jalan, J. and M. Ravallion (1996). Transient versus chronic poverty: Evidence for rural China. Manuscript.
- Kuznets, S. (1955). Economic growth and income inequality. *AER* 45(1), 1–28.
- Ligon, E. (2004). Targeting and informal insurance. In S. Dercon (Ed.), *Insurance Against Poverty*. Oxford University Press.
- Ligon, E. and L. Schechter (2003). Measuring vulnerability. *Economic Journal* 113(486), C95–C102.
- Ligon, E. and L. Schechter (2004, May). Evaluating different approaches to estimating vulnerability. Social Protection Discussion Paper 0210, World Bank, Washington, DC.
- Lundberg, M. and L. Squire (2003, April). The simultaneous evolution of growth and inequality. *The Economic Journal* 113(487), 326–344.
- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99, 928–956.

- Milanovic, B. (2004, March). Can we discern the effect of globalization on income distribution? evidence from household surveys. Manuscript.
- Obstfeld, M. (1994). Risk-taking, global diversification, and growth. *American Economic Review* 84(5), 1310–1329.
- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica* 32(1–2), 122–136.
- Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory* 2, 225–243.
- Sala-i-Martin, X. (2002, April). The disturbing “rise” of global income inequality. Working Paper 8904, NBER.
- Summers, R. and A. Heston (1991). The Penn World Tables (Mark 5): An expanded set of international comparisons. *Quarterly Journal of Economics* 106(2), 327–368.
- Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica* 62(3), 539–591.
- UNU/WIDER (2005). World income inequality database. Technical report, UNU/WIDER. Version 2a.
- Wilson, R. (1968, January). The theory of syndicates. *Econometrica* 36, 119–132.

NOTES

¹It's worth noting that in this complete markets setup consumption growth and inequality will be simultaneously determined even when utilities aren't Gorman-aggregable as assumed here, providing an additional theoretical rationale for the empirical exercise of Lundberg and Squire (2003), who argue that treating growth and inequality independently may be misleading.

²See e.g., Cochrane (1991), Townsend (1994), and Grimard (1997) for tests which rely on household level data, and Obstfeld (1994) for tests of country-level aggregates.

³This leaves open the issue of how quintile membership may change under alternative hypotheses. If any household's position within the wealth distribution can change then the estimates of total risk we provide later will tend to be biased downward.

⁴Within the economics literature, the most common measure of 'globalization' is some measure of trade openness. For conflicting claims among economists, contrast (Dollar and Kraay, 2002) with (Lundberg and Squire, 2003) or (Milanovic, 2004).

⁵Note that these are measures of risk borne by households which implicitly take into account any behavioral response households may have adopted in order to deal with anticipated shocks. It may be, for example, that poorer households engage in precautionary savings to a greater extent than do wealthier households; this would have the effect of increasing measured inequality relative to what would be observed in the absence of this behavioral response. Our measure of risk remains (estimation issues aside) the correct measure of the welfare cost of any uncertainty in any given period remaining after any such adjustments made by the household.

⁶The measure of risk described in Section 2.4 is, in principle, a measure of risk in any given period. However, in estimating risk we're relying on time-series variation, so that we instead estimate the risk in the *average* period.

⁷In principle one could construct a test of the hypothesis of difference stationarity, but we're unaware of a suitable test one could use with a (short) unbalanced panel.

REFERENCES

- Arrow, K. J. (1965). *Aspects of the Theory of Risk Aversion*. Helsinki: Yrjö Jahanssonin säätiö.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy* 99, 957–976.
- Deininger, K. and L. Squire (1996, September). A new data set measuring income inequality. *The World Bank Economic Review* 10(3), 565–591.
- Dercon, S. and J. de Weerd (2004). Risk-sharing networks and insurance against illness. Manuscript.
- Dollar, D. and A. Kraay (2002). Growth is good for the poor. *Journal of Economic Growth* 7, 195–225.
- Grimard, F. (1997). Household consumption smoothing through ethnic ties: Evidence from Cote d'Ivoire. *Journal of Development Economics* 53, 391–422.
- Jalan, J. and M. Ravallion (1996). Transient versus chronic poverty: Evidence for rural China. Manuscript.
- Kuznets, S. (1955). Economic growth and income inequality. *AER* 45(1), 1–28.
- Ligon, E. (2004). Targeting and informal insurance. In S. Dercon (Ed.), *Insurance Against Poverty*. Oxford University Press.
- Ligon, E. and L. Schechter (2003). Measuring vulnerability. *Economic Journal* 113(486), C95–C102.
- Ligon, E. and L. Schechter (2004, May). Evaluating different approaches to estimating vulnerability. Social Protection Discussion Paper 0210, World Bank, Washington, DC.
- Lundberg, M. and L. Squire (2003, April). The simultaneous evolution of growth and inequality. *The Economic Journal* 113(487), 326–344.

- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99, 928–956.
- Milanovic, B. (2004, March). Can we discern the effect of globalization on income distribution? evidence from household surveys. Manuscript.
- Obstfeld, M. (1994). Risk-taking, global diversification, and growth. *American Economic Review* 84(5), 1310–1329.
- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica* 32(1–2), 122–136.
- Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory* 2, 225–243.
- Sala-i-Martin, X. (2002, April). The disturbing “rise” of global income inequality. Working Paper 8904, NBER.
- Summers, R. and A. Heston (1991). The Penn World Tables (Mark 5): An expanded set of international comparisons. *Quarterly Journal of Economics* 106(2), 327–368.
- Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica* 62(3), 539–591.
- UNU/WIDER (2005). World income inequality database. Technical report, UNU/WIDER. Version 2a.
- Wilson, R. (1968, January). The theory of syndicates. *Econometrica* 36, 119–132.

LIST OF FIGURES

- 1 Inverse distribution of consumption across country-quintile-years for our sample, with consumption measured in 1996 US dollars. 22
- 2 Number of Years of Surveys per Country in Sample. 23




FIGURE 1. Inverse distribution of consumption across country-quintile-years for our sample, with consumption measured in 1996 US dollars.




FIGURE 2. Number of Years of Surveys per Country in Sample.

LIST OF TABLES

1	Coverage of Sample	25
2	Decreasing Global Inequality? Average growth rates of consumption expenditures for different quintiles of the world consumption distribution appear in the first row of the table, with t statistics for these point estimates immediately below. The lower panel of the table presents t -tests of differences among the growth rates of different global quintiles.	26
3	Analysis of Variance of Consumption Growth	27
4	Decomposition of Risk Across Country-Quintiles	28

Country Name	Sample Years
Algeria	1988 1995
Armenia	1996 1998
Bangladesh	1983 1986 1988 1992 1996 2000
Belarus	1995 1996 1998 1999
Bulgaria	1995 1997
Burkina Faso	1994 1998
Cambodia	1994 1997 1999
Cote d'Ivoire	1985–1988 1993 1995 1998
Ecuador	1988 1994 1995
Egypt	1975 1991 1995 1997
Estonia	1992 1993
Ethiopia	1995 1997 2000
Gambia	1992 1998
Ghana	1987 1989 1992 1993 1997 1998 1999
Greece	1974 1981 1988
Guinea	1991 1994
Guinea-Bissau	1991 1994
Guyana	1993 1999
Hungary	1993 1997
India	1951–1970 1973 1974 1977 1983 1986–1992 1999
Indonesia	1976 1978 1980 1981 1984 1990 1993 1996
Jamaica	1975 1988–1993 1995–2000
Jordan	1980 1986 1992 1997
Kenya	1992 1994 1997
Kyrgyz Republic	1996–1999
Latvia	1996 1998
Lesotho	1986 1993 1995
Lithuania	1996 1998
Madagascar	1980 1993 1997 1999
Mali	1989 1994
Mauritania	1987–1989 1992 1993 1995
Mexico	1992 1998
Morocco	1985 1991 1995 1999
Niger	1992 1994 1995
Nigeria	1980 1985 1992 1996 1997
Pakistan	1969–1971 1979 1984–1988 1990 1991 1993 1996
Peru	1991 1994 1997
Philippines	1985 1988 1991 1994 1997
Poland	1990 1992
Russian Federation	1994 1996 1998
Senegal	1991 1994
Spain	1973 1980 1985–1990 1998–2000
Sri Lanka	1970 1973 1986 1991 1996 2000
Tanzania	1991–1993
Thailand	1981 1986 1988 1990 1992 1994 1996 1998 1999
Tunisia	1965 1985 1990 2000
Turkey	1987 1994 2000
Uganda	1989 1992 2000
Ukraine	1995 1996
Viet Nam	1993 1998
Yemen, Republic of	1992 1998
Zambia	1991 1993 1996 1998
Zimbabwe	1990 1995

TABLE 1. Coverage of Sample

Quintile	20%	40%	60%	80%	100%
Growth Rate	0.043	0.028	0.024	0.021	0.015
<i>t</i> -stat.	4.603	2.942	2.565	2.249	1.626
20%	0.000	1.175	1.441	1.665	2.105
40%	-1.175	0.000	0.267	0.490	0.930
60%	-1.441	-0.267	0.000	0.223	0.664
80%	-1.665	-0.490	-0.223	0.000	0.440
100%	-2.105	-0.930	-0.664	-0.440	0.000

TABLE 2. Decreasing Global Inequality? Average growth rates of consumption expenditures for different quintiles of the world consumption distribution appear in the first row of the table, with *t* statistics for these point estimates immediately below. The lower panel of the table presents *t*-tests of differences among the growth rates of different global quintiles.

Variables	Individual Contribution (R^2)	Cumulative Contribution (R^2)	p -value
World Quintile	0.005	0.005	0.290
Global Shocks	0.044	0.049	0.689
Country Shocks	0.360	0.365	0.000
Globalization Shocks	0.109	0.425	1.000

TABLE 3. Analysis of Variance of Consumption Growth

<i>Sources of Risk, Averaged over Country-Quintiles</i>					
Quintile	Global +	Country +	Globalization +	Quintile Risk =	Total
20%	-0.549	-4.291	1.211	22.454	18.825
40%	0.220	-2.283	-0.013	11.618	9.541
60%	0.297	-1.156	-0.046	7.553	6.648
80%	0.310	-0.603	-0.057	5.012	4.663
100%	0.235	-0.258	-0.066	2.676	2.587

TABLE 4. Decomposition of Risk Across Country-Quintiles