

December 2007

*Remember that you are an Englishman, and have
consequently won first prize in the lottery of life.*
Cecil Rhodes

Where in the world are you?

Assessing the importance of circumstance and effort in a world of
different mean country incomes and (almost) no migration

Branko Milanovic¹
Development Research Group, World Bank

Suppose that all people in the world are allocated only two characteristics: country where they live and income class within that country. Assume further that there is no migration. We show that 90 percent of variability in people's global income position (percentile in world income distribution) is explained by only these two pieces of information. Mean country income (circumstance) explains 60 percent, and income class (both circumstance and effort) 30 percent of global income position. We find that about 2/3 of the latter number is due to circumstance (approximated by the estimated parental income class under various social mobility assumptions), which makes the overall share of circumstance unlikely to be less than 75-80 percent. On average, "drawing" one-notch higher income class (on a twenty-class scale) is equivalent to living in a twelve-percent richer country. Once people are allocated their income class, it becomes important, not only whether the country they are allocated to is rich or poor, but whether it is egalitarian or not. This is particularly important for the people who "draw" low or high classes; for the middle classes, country's income distribution is much less important than mean country income.

JEL classification: D31; I3; O57

Key words: Global inequality, income distribution, migration

Number of words: about 8,700

¹ I am grateful to Patrick Conway, Jim Davies, Francisco Jose Goerlich, Francisco Ferreira, Misha Lokshin, Jaime Pozuelo Monfort, John Roemer, Xiaobo Zhang and the participants of the World Bank conference on Global distributive justice held on May 1-2, 2007 in Washington, International Association for Research in Income and Wealth held in Beijing on September 19-21, 2007, and students and faculty at Free University Amsterdam for helpful comments and suggestions. I am also thankful to Gouthami Padam for excellent help in collecting and processing the data.

1. Setting the stage

In Rawls's *Law of Peoples* individuals from various countries meet to organize a contractual arrangement regulating their relations in a metaphor similar to the one for the citizens of the same nation from his *Theory of Justice*. There are differences though since the global gathering is between representatives of each nation (people) rather than between all world individuals. And the outcome is different too, in two important respects. Rawls rejects the application of the global difference principle in favor of fairly limited aid to the "burdened peoples" that are hampered by poverty from achieving a "decent" society, and assumes that migration takes place only in response to egregious violation of human rights, famine, and political and religious oppression. In other words, global redistribution is minimal and with a clear cut-off point,² and economically-driven migration is not approved.³ Thus, peoples are basically separated entities.

We shall take Rawls's assumptions as a fair representation of the existing world situation. Indeed, they are. First, in 2004, aid from rich to poor nations amounted to one-quarter of one percent of rich nations' Gross Domestic Income.⁴ At the same time, these nations were spending between 3 and 8 percent of GDI for domestic welfare payments. Obviously, domestic and foreign poor are not treated equally: one "domestic poor" is worth, on average, about 100,000 "foreign poor (Milanovic, 2006). Similarly, using an optimal taxation framework, Kopczuk, Slemrod and Yitzhaki (2005), calculate that the implicit weight US policy makers place on a foreign poor is 1/2000 of the implicit weight they assigns to an American poor.

Second, in 2002, total migration from poor to rich countries was 4.3 million people which represented a tiny percentage (less than 1/10 of one percent) of more than

² Not having open-ended international transfers was one of key points explicitly stressed by Rawls (1999, p. 106 and p. 118).

³ See Rawls (1999, p. 39 and p. 74).

⁴ See <http://www.oecd.org/dataoecd/17/39/23664717.gif>, accessed February 9, 2007. This includes only Development Assistance Committee (DAC) members (basically, the "old" OECD countries).

5 billion people living outside the rich world.⁵ So, both of Rawls's assumptions (or desiderata) seem to hold.⁶

But we shall, for the sake of exposition, modify the Rawlsian metaphor in so much as we shall let the global assembly (i) be the one of all individuals in the world, and not of peoples' representatives, and (ii) not be designed so that the individuals reach a contractarian arrangement. As is customary (from *Theory of Justice*), individuals meet behind the veil of ignorance. In our original position, each of them is allocated two characteristics that will determine his fate: county and income class within that country.⁷ As we have just seen, assignment to country is "fate" since there is no inter-country movement of people. Things are a bit more complicated regarding assignment to income class. It can also be seen as "fate" if there is no social mobility within countries. At the other extreme, with perfect social mobility, assignment to income class would not matter as each individual would, through his own exertion and luck, find his "merited" position in society.

We know that differences between mean country incomes are large: about three-quarters of global inequality between individuals is due to between-national income differences.⁸ Consequently, to what nation one gets allocated is indeed of significant import for own life chances. By being allocated to a country, the individual receives two "public" goods that are unalterable by his own effort and that are basically fixed during the largest part of his life: mean income of the country (relative to the rest of the world) and national level of inequality (which determines what share of total income will be

⁵ See <http://www.oecd.org/dataoecd/61/37/34607274.pdf>, accessed February 9, 2007. The poor and rich countries are defined here conventionally: the rich as the European Economic Area (inclusive of Switzerland but without Greece), United States, Canada, Australia and Japan; the poor as everybody else.

⁶ The total stock of people living outside a country of their birth is estimated at 3 percent of world population (see Ozden and Schiff, 2006).

⁷ If there are N countries, the probability of being assigned a given country is 1/N. In other words, the probability does not depend on the country's population. One could of course envisage a different "lottery" where the probability of being assigned a country would be proportional to the country's population, or to its share of people born in a given year.

⁸ See Milanovic (2002, p.78 and 2005, p. 112), Sutcliffe (2004), Bourguignon and Morrisson (2002, p. 734), Berry and Serieux (2007, p. 84).

received by his class). This represents, of course, a somewhat strong assumption. While these parameters are unalterable by any one's individual effort, there are indeed many examples that within one's lifetime the relative position of a country has been transformed, whether by being improved, as in the case of China over the last quarter century, or worsened as in the example of Argentina after World War II, or many African and transition countries more recently. Even national inequality, measured by the Gini coefficient, which, as Li, Squire and Zou (1998) show, tends to be fairly sluggish, can experience, at times, violent swings. The increases in inequality during the first stage of transition from planned to market economy (including in China), or under the Thatcher-Reagan rule in the UK and the United States, are such examples. For simplicity, however, we shall assume that, for an individual, both mean country income and inequality in his country of assignment are given and unrelated to any effort or desert from his part. They are thus two "morally arbitrary" features allocated to him (see Pogge 1994, p. 197; Nagel 2005, p. 119). They will be referred to as "circumstance" (Roemer 1998).

The assignment to an income class is more ambiguous than the assignment to a country: on the one hand, the assignment to a low (or high) income class will determine to a large extent individual's life-time prospects and hence his life-time income. One may (almost) argue that there are no reasons for thinking that being assigned to a top or bottom income class may not be as much a position unalterable in one's life as being assigned to a country. Yet, there is some inter-class mobility in practically every society with some countries closer to one theoretical end of the spectrum (no social mobility at all) than to the other (full social mobility). With full mobility, if we find people in a given income class within their nation, we assume that their being there reflects only their work effort and luck. It is the second part of Roemer's dichotomy: the "effort." With no social mobility at all, assignment to income class is entirely a circumstance. However, because different countries display different levels of social mobility, the actual share of "circumstance" and "effort" will differ between the countries.^{9 10}

⁹ Assignment to an income class differs from the "assignment" of a Gini coefficient. Since individuals are allowed to move up and down along the income scale of their country, the first assignment has to do with mobility. The second (the Gini coefficient) has to do with inequality of distribution, or more exactly with

This issue can be set in the more explicitly Roemerian (1998) terms. Income of i -th individual in j -th country can be, in general as in (1), written as a function of country-specific circumstances α 's, running from 1 to m (e.g., average income of the country or its level of inequality), own specific circumstances γ 's, running from 1 to n (like gender or race),¹¹ person's effort E_{ij} , and a random shock which can also be called luck (u_{ij}):

$$y_{ij} = f(\alpha_{j, \dots}^1 \alpha_j^m; \gamma_{ij}^1 \dots \gamma_{ij}^n; E_{ij}; u_{ij}) \quad (1)$$

Our objective will be to find out how much of income can be explained by circumstances alone. In the empirical analysis, we shall focus on α 's (country-specific circumstances) and leave out individual circumstances on which we have no information with the exception of one: parental income class. In addition, we shall assume that effort is independent of circumstances, either α 's or γ 's: in other words, that circumstances affect income only directly, and not indirectly through effort.

Suppose that we observe cumulative distribution of individual effort in countries 1 (poor) and 2 (rich) (Figure 1). Take individuals that are at the same point of their countries' respective distributions of effort such as $F(E^*)$. Roemer's definition of equality of opportunity requires that people who are at the same percentile of their type's distribution of effort be rewarded equally.¹² Yet, reflected in the income space (y), the

the share of each income class in total income. Thus, a society can be very unequal—in the sense that the relative income of the poor is low—while at the same time it allows for high mobility (in the sense, that being born poor does not “condemn” one to remain in that class). It is often thought that the US, compared to Europe, exemplifies precisely such a society, even if recent studies (Blanden, Gregg and Machin, 2005) have cast doubt on the superior social mobility in the United States. See also the discussion in Jackson and Segal (2004, p.p. 29-30).

¹⁰ In that sense, mobility will be the third “public good” received through country allocation.

¹¹ Assuming, for the sake of simplicity, that how these individual circumstances affect income does not depend on country.

¹² In other words, conditional on circumstance, people at the same percentile of effort should be rewarded the same (or treated equally). Roemer (1998, Chapter 3) distinguishes between the relative effort (“degree of effort”) and the absolute effort (“level of effort”). The relative effort is the effort expended compared to what is expected with a given set of circumstances. Equality of opportunity requires that the outcomes be

effort of these two individuals will result in two different income levels, y_2 , for the individual living in a rich country, and y_1 , for the one living in a poor country.¹³ If we adjust for the advantage conferred by country's higher mean income, and still obtain incomes such that $y_2 > y_1$ (even if the gap between the two is smaller), we may conclude that there are other circumstances for which we have failed to adjust. They could be country-specific inequality or institutions and norms that limit social mobility; or they could be individual-specific circumstances that are independent of country. These additional factors also confer "advantages" and have to be included under the rubric of "circumstance". Eventually, after we have adjusted for all (reasonable) circumstances that may give advantage to one or the other individual, the same relative efforts should be rewarded equally in the income space: y_1 and y_2 will be equalized (see the two arrows in Figure 1). Our objective will be to find out how much of the original difference $y_2 - y_1$ can be explained by a relatively few circumstances that we can quantify empirically.

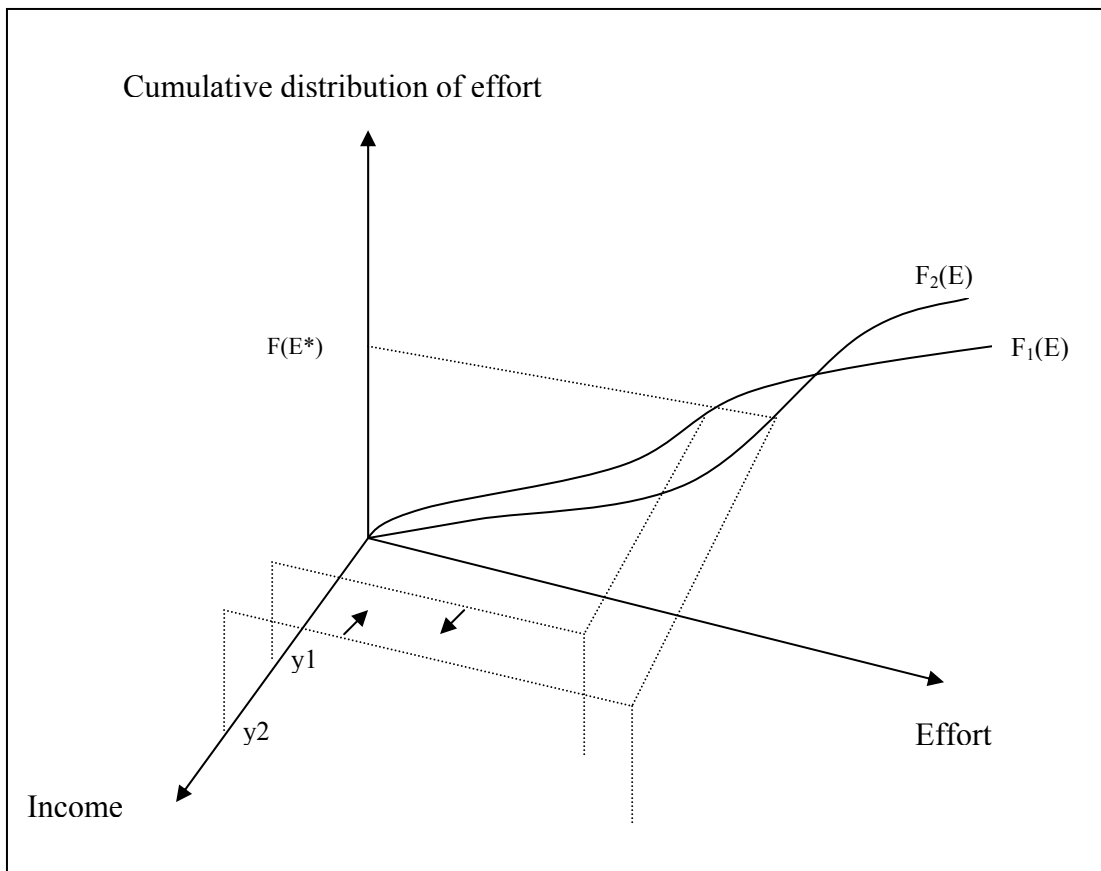
The questions we want to ask are the following: How much of one's life chances will be determined by his assignment to a given country vs. given income class? Does this "trade-off" systematically vary with income class? How much can one improve one's position in the world income distribution through his own effort (that is, by climbing income ladders in his country)? What does all of this tell us about the equality of opportunity across all individuals in the world?

We shall first (section 2) describe the source of global income distribution data that help us address these questions empirically and review our definitions of country and class. In Section 3, we present some broad regularities regarding the way global income is distributed between countries and income classes. Sections 4 and 5 are the core parts of the paper: they present the analysis that attempts the answers the questions posed above. The last part gives the conclusions.

the same for each percentile of the distribution of effort (that is, for each relative effort) allowing thus the same absolute effort to be rewarded differently.

¹³ The same would, of course, be true for all pairs of individuals who are at the same percentile of their type's effort distributions. The income distribution curve of the richer country, 2, will extend over the higher values of income than the income distribution curve of the poorer country, 1.

Figure 1. Equality of opportunity for two different types of individuals



2. Data and definitions

The data used in the paper come from the World Income Distribution (WID) database constructed to study the evolution of global inequality. The database is comprised almost entirely of micro data from representative household surveys from most of the countries in the world. For the year 2002, which is used here, the data come from 120 countries' household surveys representing 94 percent of world population and 98 percent of world dollar income.¹⁴ The geographical coverage is almost complete for all parts of the world except Africa (see Table 1).

Table 1. Population and income coverage of the surveys (in %)

	Africa	Asia	Latin America	E.Europe and CIS	WENAO	World
Population	77	96	96	97	99	94
Income	71	95	95	99	100	98
Number of countries	29	23	21	26	21	120

Source: World Income Distribution database. See <http://econ.worldbank.org/projects/inequality>.

Note: WENAO is Western Europe, North America and Oceania (Australia and New Zealand). CIS = Commonwealth of Independent States. Eastern Europe includes formerly Communist countries.

For the vast majority of surveys (115 surveys) we had access to micro data which means that any type of distribution (by decile, ventile, percentile; by households or individuals) could have been created. In order to limit the number of data points and make the analysis manageable we have limited the number of data points per country to 20 ventiles (each ventile contains 5 percent of country's population). All individuals in a survey are ranked from the poorest to the richest according to their household per capita income (or expenditures, depending on what welfare aggregate is used in the survey). Since not all countries produce annual surveys, we had to use a "benchmark" year (2002 in this case), that is, try to get 2002 household surveys for as many countries as possible, but where there were no surveys conducted in 2002, to use a year as close to 2002 as possible. In the event, 79 country surveys were conducted in the benchmark year or one

¹⁴ We cannot express the share of the included countries in terms of \$PPP income because for most of the countries for which we lack surveys, we also lack PPP data (e.g., Afghanistan, Iraq, Sudan etc.) The dollar incomes however are typically available.

year before or after it, and all but two surveys within two years of the benchmark. For the surveys conducted in non-benchmark years, we adjust reported incomes by the Consumer Price Index of the country so that all amounts are expressed in 2002 local currency units. These amounts are then converted into international (PPP) dollars using the 2002 estimates of \$PPP exchange rates provided by the World Bank. Thus, for each ventile, we calculate the average per capita amount of PPP dollars received as income (or spent in the form expenditures).¹⁵

The fact that each country is divided into 20 groups of equal size (ventiles) is extremely helpful. This allows us to compare the positions of say, the third ventile of people in China with the seventh ventile of people in Nigeria etc. It also allows us to define income classes in the same way across countries. To fix the terminology, we shall call each ventile an “income class”. Income classes thus run from 1 to 20 with 20 being the highest.

Income class and country of residence pin down a person’s position in global income distribution. That position is expressed by his percentile rank in world income distribution. A person can be, say at the 72nd percentile in the world—implying that his income is higher than incomes of 72 out of each 100 people in the world. This will be referred for simplicity simply as “position” or “position in the world.” Since we divide the world into one hundred percentiles according to per capita income, the position runs from 1 (lowest) to 100 (highest). Each percentile contains in principle 1/100th of the world population included in the analysis here, *i.e.*, approximately 57 million people.¹⁶

We now move to some descriptive issues showing how the world thus “partitioned” into countries and income classes really looks.

¹⁵ As mentioned, the household surveys we use are either income- and expenditure (consumption)-based. For the simplicity of presentation we speak throughout of “income” distribution and “income” position in the world.

¹⁶ Unfortunately, not all percentiles are of the exactly the same size. This is due to the “lumpiness” of the ventiles, particularly those for China and India. Thus as one Chinese ventile takes more than one percentile of the world income distribution, the next world percentile is accordingly reduced in size.

3. Diversity of the world

Figure 2 combines the two aspects of within-national and international distributions. Income of each country ventile is shown in the global distribution. Consider Germany. Since Germany is a rich country, and its income inequality is moderate, most of its population will be highly placed in the world income distribution. The poorest German ventile is at the 73rd percentile of the world income distribution. All other ventiles are obviously higher, and the richest ventile belongs to the top world percentile. The same interpretation is for other countries. We call such curves “the position curves”. Unlike Germany, where the span between the richest and the poorest ventile is 27 percentiles, in China, the distribution covers a much wider range, from the third to the 85th percentile. Brazil, with its unequal income distribution, covers practically the entire global spectrum, from the lowest percentiles to the richest. India, in contrast, is shown to be fairly poor with the poorest ventile belonging to the 4th poorest percentile in the world and the richest ventile to the 70th. This last position shows that the richest people in India (as a group—admittedly a large one since it contains more than 50 million people) have lower per capita income than the poorest people (as a group) in Germany.


The graph can also be read as a type of generalized Lorenz curve where instead of the income level on the vertical axis, we have income position in the world. The advantage of this “positional” approach is that it reduces the measurement error,¹⁷ but since position is bounded from above these specific generalized Lorenz curves will in many cases be concave rather than convex. The interpretation however is the same as with the generalized Lorenz curves.¹⁸ From Figure 2 we can easily conclude that Sri Lanka’s distribution is first-order dominant with respect to India’s, and that Germany’s distribution is first-order dominant compared to any other country save Brazil. No first-

¹⁷ Household surveys do not measure income or expenditures perfectly. They are less likely however to make large mistakes that may result in the misplacing of individuals into “wrong” world percentiles.

¹⁸ The first-order positional dominance *must* imply the first-order income dominance. The reverse may not hold because a distribution may be income dominant but the difference in income may be so small as to place the ventiles from both countries into the same global percentile.

order dominance can be established between Brazil, China and India because of the situation at the bottom where the poorest Brazilians are shown to be poorer than the poorest people in India and China. Of course, the middle class Brazilians (approximately people in the ventiles 7 through 15) are better off than the middle classes in China, Sri Lanka and India. One may also note that the biggest difference in the positions holds for the poorest ventiles: while in Germany, the poorest ventile is at the 73rd world percentile, in the other four countries, the poorest ventiles are close to the bottom of global income distribution. The positional difference for the top ventiles is much less.

Figure 2. The position curves: inequality in the world—by countries and by income class



Source: World Income Distribution (WYD); benchmark year 2002.

World income distribution can be conventionally broken down into that part of inequality which is due to the differences between mean country incomes, and that part of inequality due to inequality within countries. Using 2002 data, Table 2 shows that, depending on the inequality measure, between 66 and 87 percent of global inequality is due to differences in countries' mean incomes. Taking the Gini coefficient, which is the most frequently used measure in global inequality studies, income differences between world citizens amount to 65.5 Gini points out of which 55.7 points are due to the between-country component.

Table 2. Global income inequality and the between-country component
(benchmark year 2002)

	(1)	(2)	(3)
	Global inequality between individuals	The between- country component of global inequality	Share of (2) in (1) (in percent)
Relative mean deviation	0.517	0.450	87
Coefficient of variation	1.751	1.278	73
Standard deviation of log of incomes	1.234	0.982	80
Gini coefficient	0.655	0.557	85
Mehran measure	0.783	0.683	87
Piesch measure	0.591	0.494	84
Kakwani measure	0.357	0.274	77
Theil entropy measure	0.835	0.579	69
Theil mean log deviation	0.846	0.562	66

Source: World income distribution (WYD) database. All incomes expressed in 2002 international dollars.

4. The relative importance of country vs. income class

Predicting global income position based on knowledge of country and class (in the aggregate)

As we have seen, one's position depends on two factors: allocation to country and allocation to income class. We can write for i -th individual living in j -th country:

$$P_{ij} = b_0 + b_1 m_j + b_2 G_j + b_3 C_{ij} + \varepsilon_{ij} \quad (2)$$

where P_{ij} = income position (percentile) in the world income distribution, m_j = mean country income, G_j = national inequality (say, Gini coefficient), and C_{ij} = person's income class in country j , and ε_{ij} = the error term.¹⁹

¹⁹ As mentioned, because of the large size of ventiles for India and China, not all percentiles contain exactly 1/100th of world population. As an alternative, we ran regressions where P_{ij} is the exact position in the world income distribution (e.g., at the 63.4 percentile, not as a part of the 64th percentile). The results are almost identical. They are available from the author on request.

The results of estimation of (2) are shown in Table 3.²⁰ We begin by asking how much of one's global income position is explained by country's mean income alone (regression 1). The answer is 60 percent. Note that each increase of 10 percent in mean country income raises a person's position in the world by about 2.3 percentiles on average. But when individuals are allocated a country, they are not only allocated its mean income but also its inequality. Including both of them in the regression however does not make much of a difference (regression 2).²¹

By putting together country and income class (regression 3), we are able to explain more than 90 percent of the variation in people's positions in the global income distribution. As before, each 10 percent increase in the mean country income lifts a person, on average, by 2.23 percentage points. On the other hand, being placed in a higher income class increases one's position by 2.8 percentiles on average. Thus, in the aggregate, belonging to one-notch higher income class in one's country is, on the average, equivalent to residing in a country whose mean income is just over 12 percent higher. The trade-off between income class, that is, what we may consider to be a partial reflection of one's effort, and the morally arbitrary placement in a rich country is now clear. If one were, through his effort and luck, to climb eight income classes in his country, he would have "traversed" the road equivalent to being born in a country about twice as rich.

²⁰ The regressions are run unweighted implying that each country (regardless of its population) matters equally. This makes sense from the point of view of the original position where, for an individual, the probability of being assigned to any given country is the same. The Rawlsian lottery would be different if probabilities of country assignment were proportional to the population sizes of the countries.

²¹ Each Gini point increase will, on average, lower person's position by about 0.33 percentage point. This, of course, holds only in the aggregate. If we break individuals by income class, then living in a more unequal country (and controlling for the mean income) would be advantageous for higher-class individuals. And the reverse for the people allocated to low social classes. This point is pursued below.

Table 3. Explaining one's position in world income distribution (dependent variable: percentile in world income distribution)

	Country only		Country and own income class	Including parent's income class					Hypothetical
	1	2	3	Base case	Optimistic	Pessimistic	Super optimistic	Super pessimistic	9
Mean per capita income (in ln)	22.92 (0)	22.32 (0)	22.32 (0)	22.33 (0)	22.33 (0)	22.42 (0)	22.32 (0)	22.32 (0)	---
Gini index (in %)		-0.33 (0)	-0.33 (0)	-0.34 (0)	-0.34 (0)	-0.35 (0)	-0.33 (0)	-0.33 (0)	-0.14 (0)
Own income class (ventile)			2.80 (0)						4.77 (0)
Parents' income class (ventile)				2.39 (0)	2.30 (0)	2.46 (0)	2.09 (0)	2.52 (0)	
Constant term	-126.2 (0)	-108.2 (0)	-137.6 (0)	-132.5 (0)	-131.5 (0)	-134.0 (0)	-130.2 (0)	-134.4 (0)	23.08 (0)
Number of observations	2300	2300	2300	230,000	230,000	230,000	230,000	230,000	2300
R ²	0.60	0.61	0.91	0.80	0.78	0.82	0.76	0.83	0.96
F value	4254 (0)	1799 (0)	1202 (0)	31,539 (0)	28,467 (0)	34,155 (0)	24,705 (0)	37,816 (0)	3353 (0)

Note: The regressions are run with the cluster option to adjust for the correlation of within-country observations. Regressions are unweighted. There are 115 countries times 20 ventiles = 2300 observations in regressions 1-3 and 9. The expanded regressions have 100 times more observations. *p* values between brackets. Income class ranges from 1 (lowest) to 20 (highest).

When we break the importance of country and income class in explaining one's position in the global income distribution, we find that 63 percent is due to the country of residence, and 31 percent to income class.²² However, income class can be fully treated as “effort” only if we are willing to argue that (1) income class a person is assigned at birth and income class he is in now are totally orthogonal, and that (2) the latter is dependent on his effort (and luck) alone. More formally, we can express that situation as the one where the correlation between one's current income and his parents' income is zero. At the other extreme, with no social mobility at all, one's income class at birth determines his current income class (*i.e.*, the one observed in the surveys). In that case, the entire income class variable has to be “ascribed” to circumstance.

The situation in the real world will, of course, differ between the countries and will lie somewhere between the two extremes. Ideally, if we had the data for the correlation of children's and parents' income by country, we could use these country-specific coefficients to estimate the actual role of one's inherited position. Unfortunately, we have such data for only a dozen, mostly rich, countries. They show that intergenerational mobility is relatively high in the Nordic European countries and Canada, that it is less in the United States and the UK, and (arguably) even less in the continental Europe (see Solon, 1999, pp. 1784-89; Checchi et al. 1999; Bjorklund and Jantti 1997)). The coefficients of intergenerational earnings' elasticity, ρ , in rich countries range between 0.2 in Nordic countries (and in some studies only), and 0.6.²³ For a few Third World countries in Asia and Latin America where estimates are available, the coefficients tend to be high and range from just under 0.5 to 0.7 (Lam and Schoeni, 1993; Grawe 2001; Ferreira and Veloso 2006). Based on a survey of the literature, we have incorporated these results into our “base case” scenario on mobility shown in Table 4. To

²² This is obtained by the analysis of the variance and is independent of the order with which the regressors are introduced.

²³ The coefficients are calculated from the regressions of children's on parental earnings with both earnings expressed in logarithms. We can treat the elasticities as the correlation coefficients if we assume that the standard deviations of parental and children's earnings are approximately the same. Note finally that the estimates of intergenerational elasticity used here apply to earnings and not to income as we would ideally like.

see how the results may be sensitive to different mobility assumptions, we introduce several additional scenarios: optimistic and pessimistic, where social mobility is respectively greater and less than in the base case (see Table 4), and then the two extreme scenarios, a super-optimistic one—which serves more as a blueprint of an ideal world—where social mobility is high and equal in all parts of the world, and a super-pessimistic scenario, where mobility is very low in all countries.

Table 4. Coefficients of intergenerational elasticity between parents' and children's income used in the simulations

	Base case	Optimistic (high mobility)	Pessimistic (low mobility)	Super optimistic	Super pessimistic
Nordic countries	0.2	0.15	0.3	0.2	0.9
Rest of WENAO	0.4	0.3	0.5	0.2	0.9
Eastern Europe/CIS	0.4	0.3	0.5	0.2	0.9
Asia	0.5	0.4	0.6	0.2	0.9
Latin America	0.66	0.5	0.9	0.2	0.9
Africa	0.66	0.5	0.9	0.2	0.9

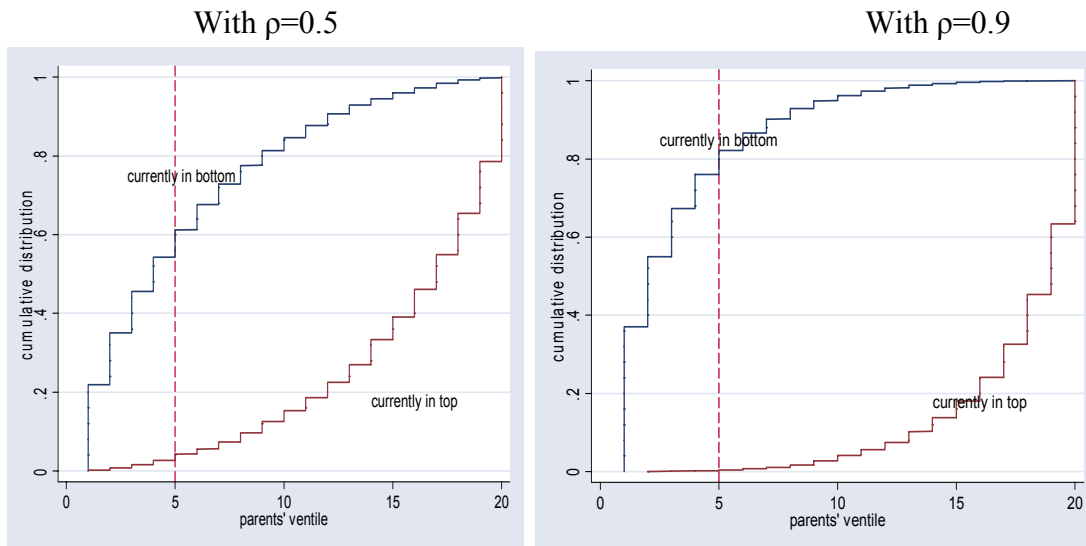
Once we have assumed the correlations for all the countries in the sample, we proceed to the following simulation exercise. Take a country j with its correlation coefficient ρ_j . We do not know in what ventile of income distributions have been the parents of people whom we observe in (say) the bottom ventile. To estimate this, we run a random data generation process

$$y_{ij} = \rho_j y_{ij}^* + e_j \quad (3)$$

where y_{ij} = income (in logs) of i -th individual drawn from a normal distribution, y_{ij}^* = income (in logs) of i -th individual's parents (the asterisk denotes parents) and e_j = the error term drawn from a normal distribution $N(0,1)$. After generating incomes of parents and children, we partition both parents' and children's incomes into twenty ventiles, and for each children's ventile calculate the conditional distribution of parents' ventiles. Figure 3 shows such cumulative conditional distributions for the bottom and the top ventile when ρ takes the values of 0.5 and 0.9. As can be easily seen, with a high ρ , people whom we currently observe in the bottom (top) ventile are very likely to have

come from the parents who were also in the bottom (top) ventile. But as ρ decreases, that probability lessens. For example, with $\rho=0.9$, people who are currently in the bottom ventile come with the probability of 80 percent from the parents who have been themselves located in the bottom five ventiles (see the vertical line at $x=5$ in the right panel of Figure 3). But with greater social mobility ($\rho=0.5$), such probability is just over 60 percent (see the left panel in Figure 3). If eventually ρ were to be 0, the distribution of parents' income (or more accurately, the distribution of parents' ventiles) will be the same for each income ventile of children.

Figure 3. Cumulative distribution functions of parents' ventile position for the children in bottom and top ventile of income distribution



Note: simulations based on equation (3). Children's ventiles are labeled "current".

Using thus generated parental ventiles, we proceed as follows. Each currently observed (children's) ventile is expanded by a factor of 100, and to each child in ventile i is assigned an estimated parental income ventile. For example, if for a given value of ρ and a given children ventile, the distribution of parental incomes is such that 30 percent of parents come from the first ventile, 40 percent from the second ventile, and 30 percent from the third ventile, then 30, 40 and 30 children in this ventile will be assigned respectively the first, second and third parental ventile. We thus achieve heterogeneity of

parental ventiles within a given children ventile (which would not be possible if we were to assign to all children in a given ventile the same *expected* parental ventile). Over such expanded sample, we run regression

$$P_{ij} = \beta_0 + \beta_1 m_j + \beta_2 G_j + \beta_4 C_{ij}^*(\rho) + \varepsilon_{ij} \quad (4)$$

where C_{ij}^* is the estimated income class of i -th individual's parents, which of course varies in function of ρ . The fact that parental income classes differ between individuals (children) that belong to the *same* observed income ventile and live in the *same* country, brings us closer to isolating the effect of circumstance. This is because in addition to inter-ventile variability of parents' position, we introduce variability in parental positions also within each income ventile of children. Now, the entire explained part of the regression can be treated as "circumstance". Columns (4)-(8) of Table 3 show the results for the five scenarios delineated above: the base case, optimistic and pessimistic, as well as the two extreme ones.

Comparing regressions (3) and (4), we note that the substitution of own ventile by parental ventile (in the base case scenario) reduces total "explained" variability of income position in the world from 0.91 to 0.80. Parents' income class is statistically significant and its absolute value is smaller than that of own income class: on average, having parents' ventile position go up by one notch rises one's position in the world by about 2.4 percentage points, some 0.4 percentage points less than if own income class is one notch higher.

The importance of circumstance decreases in the optimistic scenario (see regression 5 in Table 3), and goes even further down to $R^2=0.76$ in the super-optimistic scenario when we assume an equally high social mobility in all parts of the world. With a pessimistic (and super-pessimistic) scenario of very low social mobility, the role of circumstance increases to about 82 to 83 percent (regressions 6 and 8 in Table 3). The importance of higher parental income class is, as expected, greater when we assume lower social mobility.

In conclusion, between 60 and more than 80 percent of variability in the global income position can be attributed to the circumstances beyond individual control. Sixty percent represents the lower bound where only the mean income of the country and country's inequality are allowed to play a role. Eighty percent or more is obtained when we include person's parental income as part of circumstance, and use either base-case or pessimistic assumptions regarding income mobility in various parts of the world. In any case, the part which remains for effort and "episodic luck" (to use John Roemer's felicitous phrase) remains relatively small.

Finally, we compare the actual role of location to a hypothetical case where all countries' mean incomes are equal.²⁴ We still "allocate" people to different countries and income classes in our Rawlsian lottery, but now location implies only a difference in the income distributions between the countries (different Gini coefficients), not the difference in the average income. The results are shown in column (9) Table 3. The coefficient on income class more than doubles compared to regression (3), and when we decompose the two effects, income class is found to explain more than 90 percent of variability in the global income position, while location (through its specific inequality) accounts for less than 5 percent.²⁵ This counterfactual also allows us to conclude that location really matters through its mean income effect, not through its specific (national) inequality.

Median global position and its variability when income class is given


A different way to look at effort is to consider by how much one's position in the world improves, on average, if he is able to move up the income ladder within his country. For example, for a person in the bottom income class, the median position in the world is the 7th percentile. Suppose now that he manages to climb up to the 5th income class (out of 20). His median position will have improved to the 39th percentile. Another

²⁴ This is the situation referred by Roemer (2007) as Equality of opportunity of degree 1. Incidentally, if all mean incomes were equalized the global Gini would be only 37.4 vs. the actual Gini of 64.2 (based on the World Income Distribution 2002 dataset).

²⁵ Historically, something similar might have obtained in the early 19th century when, according to the Bourguignon and Morrisson (2002) study of the long-run global inequality, class (within-national inequality) explained about 90 percent of world inequality between individuals.

equivalent climb of five income classes will place him at the 56th percentile. Figure 4 shows the results for each of the twenty income classes. The marginal gains are very significant at the bottom (e.g., a move from the lowest to the second income class improves one's median position by 14 percentiles), taper off in the middle, and increase again at the very top: going from the 19th to the highest income class improves one's median position by ten percentiles (from the 82nd to the 92nd).

Figure 4. Median position in the world as function of one's income class




Note: unweighted data, each country's ventile represents one observation.

So far we have considered only the *median* position of a person if his national ventile is given. What is important to take into account also is that the variability of one's position in the world income distribution is not the same regardless of the income class. In other words, the *distribution* of global positions for various income classes is different. Figure 5 illustrates this for the two extremes, the top and the bottom ventiles. The distributions are of different shapes, in addition to covering obviously different parts of the global income distribution. The overlap between the two distributions is small but the very fact that it exists illustrates how unequal national mean country incomes are because

in some cases people belonging to the top national ventile are poorer than the people who are in the lowest ventile of another country. Overall, if one belongs to the lowest income class, he is very likely (probability of more than 60 percent) to be placed in the bottom quintile of the world income distribution. But he can—at the extreme—if he lives in a rich country, rank as high as the 84th world percentile (this is the case if he lives in Luxembourg). On the other hand, if he belongs to the highest national ventile, his range of possible outcomes, although wide, is narrower than in the previous case: in the worst case scenario (if he lives in Tanzania), his position in the world would be at the 37th percentile while in the majority of cases he would be placed above the 90th percentile.

Figure 5. Density function of one's position in the world as function of one's national income class



Note: Unweighted data, each country's ventile represents one observation.

A slightly different, and a more complete, way to look at this is shown in Figure 6. There we plot percentile ranks in the global income distribution for each income class against the mean country income. The upward sloping curves show that, for any given income class, the increase in the mean country income is associated with a higher

position of that income class in the global income distribution. The relationship is sharper as we move from low to high income classes. This means that the variability of outcomes, due to nationally idiosyncratic factors, will be greater among the nationally poor than among the nationally rich.

Figure 6. Income class, country mean income and position in global income distribution




Note: Each graph for one income class, running from 1 to 20. Mean incomes in logs.

In effect, the variability of the global positional outcomes, measured by the standard deviation, steadily decreases (with one exception) as income class goes up (see Figure 7). For the low income classes (below the fifth), the standard deviation is about 30 percentiles; for the top income classes, the standard deviation is less than 20 percentiles. A significant exception to this regularity is the lowest income class whose variability of position is less than that of the second, third and the few following classes.

To summarize: if one is in the top income class of his country, the median position in the world that he can expect to attain is the 92nd percentile and the standard

deviation is only about 12 percentiles. If he belongs to the bottom income class in his country, his median position in the world is the 7th percentile but the standard deviation is much larger: about 26 percentiles. In other words, for those who belong to the low income classes (*i.e.*, the “nationally poor”), location matters even more than to those who are “nationally rich”. To this issue we turn next.

Figure 7. Standard deviation of one’s position in the world income distribution as function of one’s income class



5. Varying importance of location for different income classes

If income class is given, how well can we predict global position with knowledge of country income alone?

When people are allocated income class in our Rawlsian lottery, it is not a matter of indifference, as we have seen, what country they get allocated to. Location, if one “draws” a rich country, can more than compensate for a low income class. But the impact of location is not uniform across all ventiles. This is because when a person is allocated a country, he is allocated two relevant features of that country: its mean income, and also its income distribution. Table 5 shows the results of regressions similar to (2) but with income class being held constant. That is, for each income class, we regress person’s position in the world income distribution on country’s characteristics alone: its mean income and a measure of its inequality (the ventile’s share of total income).²⁶ These two characteristics always explain more than 90 percent of the variability in person’s global income position. For example, looking at the people in the lowest income class, the R^2 is about 0.9, and each 10 percent increase in the mean country income is worth 2.3 percentiles climb in the global income distribution. But for a person belonging to the highest income class, each 10 percent increase in the mean country income is worth only 1.2 percentiles increase in the global income distribution. We find again that location matters more to nationally poor than to nationally rich people.

Trade-off between country’s mean income and country’s distribution across income classes

The two country characteristics (mean income and inequality) can also be seen as substitutes: given his income class, a person might prefer to be “allocated” into a more equal society even if its mean income is less. He could benefit more (if he is poor) by the first than lose by the second. Intuitively, we can also see that if a person is allocated to a top income class, then the gain from belonging to a more equal society will be negative. Thus, the trade-off between mean income and inequality is not the same across income classes. If we consider the bottom income class (as in regression 1 in Table 5), we see

²⁶ Percentage of total income received by a given ventile in a given country.

that each point increase in the bottom group's ventile share (distributional gain) is worth a huge climb of 23 percentage points in world income position. Now, to achieve the same increase of 23 points in the global position, a person in the bottom ventile would need to be located in a country twice as rich (see the same regression). This is the shape of the trade-off for those in the lowest income class. Contrast this with the fact that if the ventile share of the people in the richest class goes up by 1 percentage point their position in the world improves by only 0.6 percentile which is an increase equivalent to living in a country that is only 5 percent richer (regression 20 in Table 5).

Table 5. Explaining a person's position in the world income distribution—given his national income class (ventile)

Income class	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Mean country income	23.06 (0)	24.77 (0)	25.22 (0)	25.32 (0)	25.4 (0)	25.17 (0)	24.98 (0)	24.66 (0)	24.23 (0)	23.88 (0)	23.39 (0)	22.92 (0)	22.11 (0)	21.31 (0)	20.33 (0)	19.17 (0)	17.86 (0)	15.99 (0)	11.75 (0)	
Ventile share (in %)	22.89 (0)	20.2 (0)	17.79 (0)	15.59 (0)	13.34 (0)	11.35 (0)	9.49 (0)	8.00 (0)	6.70 (0)	5.64 (0)	4.58 (0)	3.72 (0)	3.07 (0.01)	2.39 (0.05)	2.01 (0.05)	2.43 (0.01)	3.03 (0)	3.08 (0)	1.48 (0)	0.62 (0)
Constant	-182.5 (0)	-195.3 (0)	-196.3 (0)	-193.7 (0)	-189.5 (0)	-184.9 (0)	-178 (0)	-172.0 (0)	-164.9 (0)	-157.4 (0)	-150.0 (0)	-141.9 (0)	-134.5 (0)	-123.7 (0)	-114.2 (0)	-107.7 (0)	-102.1 (0)	-92.7 (0)	-63.74 (0)	-18.31 (0)
Adj. R ²	0.902	0.951	0.963	0.968	0.971	0.969	0.966	0.966	0.963	0.960	0.958	0.954	0.952	0.948	0.946	0.941	0.939	0.938	0.933	0.906
No of obs	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115
F value	385.1 (0)	814 (0)	1052 (0)	1168 (0)	1310 (0)	1351 (0)	1255 (0)	1246 (0)	1103 (0)	983 (0)	879.9 (0)	750.1 (0)	648.2 (0)	581.9 (0)	510.1 (0)	468.4 (0)	413.5 (0)	343.8 (0)	291.2 (0)	173.9 (0)

Note: Ventile share expressed in percent of total country income. Mean per capita income in \$PPP per annum. *p*-values between brackets

However, the reasonable trade-off has to allow for the fact that the increase of 1 percentage point in the ventile share is, in relative terms, much greater (and much less likely to obtain) for the poor people than for the top income class. For the poor, such an increase would mean a doubling of their average share: for the richest, an increase of less than 1/20 (see Table 6). To normalize for this and make the analysis more realistic, we consider a trade-off where a person is, in each case (that is, *given* the income class he belongs), placed in a country whose ventile share is one standard deviation above the world average. This means that for the poorest income group, his gain would be 0.52 percent of total income and for the richest group 7.35 percent (see Table 6). Now, the relative “worth” of national income distributions thus defined is contrasted to the “worth” of higher mean country income. The results are shown in Figure 8. The importance of national distributions is, as expected, very high for the poor: “getting” a country whose bottom class’s share is one standard deviation above the world mean is equivalent to “drawing” a country that is 50 percent richer. The trade-off then gradually weakens before picking up for the richest three classes. There too, “drawing” a (very unequal) country such that, for example, the highest class has a ventile share that is one standard deviation higher than the world mean ventile share of that class, is equivalent to living in a 40 percent richer country. We therefore have to modify our earlier conclusion: for both the people who are “assigned” to be nationally poor and nationally rich, “drawing” respectively more equal and more unequal country will matter a lot. But for the people in the middle of national income distributions, “drawing” a more or less equal country has very little value compared to being placed in a mean-richer country.²⁷

²⁷ The same analysis can be done by using national Ginis instead of ventile shares in all regressions. We find that for the ventiles ranging from the poorest to the twelfth, greater country inequality reduces their position in the global income distribution; for the ventiles 13 to 16, Gini is not significant, and for the ventiles 17th and above greater inequality is “good” (that is, it rises their global income position). The results are available from the author on request.

Table 6. Average share of total income received by each ventile of national income distributions (unweighted average calculated from 115 household surveys)

Ventile	Average ventile share in total income (in %)	Standard deviation of ventile share (in %)
First	1.00	0.52
Second	1.50	0.60
Third	1.80	0.63
Fourth	2.06	0.64
Fifth	2.31	0.65
Sixth	2.54	0.65
Seventh	2.78	0.64
Eighth	3.02	0.63
Ninth	3.28	0.62
Tenth	3.55	0.60
Eleventh	3.85	0.58
Twelfth	4.18	0.56
13 th	4.55	0.53
14 th	4.99	0.48
15 th	5.52	0.45
16 th	6.18	0.41
17 th	7.07	0.48
18 th	8.36	0.75
19 th	10.72	2.01
Twentieth (top)	20.74	7.35
Total	100	

Note: Distributions for the benchmark year 2002. Source: WYD database. Each country is one observation.


These results have implications for migration. If low income class people migrate to richer countries, and expect that they would end up there too among low income classes, then equality of the receiving country's income distribution must be quite important for them. A very large increase indeed in the mean country income is needed to compensate for this "distributional premium". But differently, if the nationally rich people (say, highly skilled) migrate from a poor to a rich country, and expect to be among high income groups in their new country too, then they might prefer to select highly unequal societies, even if their mean income is less than the mean income of an alternative migration destination.²⁸

²⁸ An interesting example is provided by Bustillo (2007, pp. 21-22). His results show that the percentage of immigrants monotonically decreases as one moves from poorer to richer deciles in Spain. But in the United States, the share of immigrants charts a U curve: it is very large in the bottom and top deciles.

Given mean income of the recipient country, and given expectations on where one might be placed in the income structure of the new country, we would expect low-skilled people to migrate into more equal countries and more skilled people to migrate into more unequal countries. This parallels the idea underlying Borjas's (1987, 1999) self-selection hypothesis. However, note that the picture here is a bit more complex, in the sense that while the increase in the mean income has to be high at both ends of the income distribution to compensate for either unequal income distribution (for the poor) or equal income distribution (for the rich), the offsetting increase in the mean country income is rather minimal for the middle-income groups (e.g., ventiles 11 through 18 in Figure 8). This means that for the middle classes, the distribution in the receiving country will not matter much: country's mean income will be much more important.²⁹ In turn, this result implies that for most people with moderate skill levels, or with people with high skill levels who do not expect to be able to make it to the top of the income ladder in the receiving country, mean income of the receiving country would trump other considerations.

²⁹ The finding parallels Palma's (2006) recent emphasis on the share constancy of the middle deciles regardless of how equal or unequal the overall distribution is. In other words, inequality of distributions is determined by high or low shares of the top or bottom fractiles, not by the shares of the middle groups.

Figure 8. Value of one standard deviation increase in the ventile share at different points of national income distribution (measured in terms of mean country income)



Note: Calculated from Tables 5 and 6.

6. Conclusions

This paper allows us to make three key conclusions.

First, with only two characteristics, person's country (which in a world with no significant migration, essentially means his place of birth), and income class, we are able to explain more than 90 percent of the variability in global income position (global percentile rank). The first characteristic (location) is clearly a "circumstance", or a morally inconsequential feature. It explains about 60 percent of one's position in the global income distribution. The second characteristic, to the extent that social mobility is not absolute, also has a share of circumstance in it. When we approximate that part by the estimated parental income class, we obtain—depending on the assumptions about social mobility in different parts of the world—that between 76 and 83 percent of variability in the global income position is accounted by circumstances. Other features (gender, race, or ethnicity), which are not included in the analysis, may further increase this share.

Second, the ability to "predict" well one's location in the global income distribution from only two characteristics, holds, not only in the aggregate, but for each income class separately. Thus, for any given income class, the knowledge of the country where a person lives is sufficient to explain 90 percent or more of that person's global income position: the predictive power of the country mean income is strong, not only in the aggregate, but for each income class. Living in a richer country is particularly important for low income classes, where each 10 percent increase in the country's mean income, lifts person's global income rank by 2.3 percentiles on average. The "location premium" is significant but smaller for top income groups where it amounts to between 1 and 1.5 global percentiles on average. In other words, the value of living in a richer country is shown to hold for the entire national income distribution spectrum, but to be particularly strong for the "nationally" poor.

Third, given person's income class, there is also a trade-off between wealth of the country (reflected in its mean income) and its income distribution. Thus, a person who is allocated a low class might prefer to be allocated to a more egalitarian country even if

that country's mean income is less. The opposite, of course, holds for a person allocated to a high class: he might benefit from country's inegalitarian distribution more than from its high mean income. The trade-off is such that being placed in a country that is one standard deviation more *egalitarian* than the world average is equivalent, for a person belonging to the lowest income class, to living in a 50 percent richer country. For a person who belongs to the highest income class, getting a one standard deviation more *inegalitarian* country is equivalent to living in a 40 percent richer country. But these sharp trade-offs between the internal income distribution of a country and its mean income hold mostly for the extreme income classes. For the middle classes, distribution is relatively unimportant—because income shares of the middle groups do not vary much across nations.

The last point has clear implications for migration. If people who migrate expect to be placed in the middle of the national income distribution of the receiving country, they will be focused primarily on country's mean income. But if people who migrate expect to end up in the bottom of the recipient country's income distribution, whether the recipient country is egalitarian will be of significant importance in their decision-making. And the reverse if they expect to end up in the top of income distribution of the recipient country.

REFERENCES

Berry Albert and John Serieux (2007), “World economic growth and income distribution, 1980-2000” in *Flat World, Big Gaps* (2007), edited by K.S. Jomo and Jacques Baudot, Hyderabad, India, London; UK and New York, US; Penang, Malaysia: Orient Longman, Zed Press and Third World Network, pp. 1-23.

Blanden, Joe, Paul Gregg and Stephen Machin (2005), “International mobility in Europe and North America”, Centre for Economic Performance, London, April.

Borjas George J. (1987), “Self-selection and the earnings of immigrants”, *American Economic Review*, vol. 77, pp. 531-553.

Borjas, G. J. (1999). “The economic analysis of immigration”, in O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, vol. 3A, Chapter 28, pp. 1697-1760. Elsevier Science North Holland.

Bustillo, Rafael Munoz de (2007), “Migration and inequality: What do we know?”, Provisional draft, May 2007.

Bourguignon, Francois and Christian Morrisson (2002), “The size distribution of income among world citizens, 1820-1990”, *American Economic Review*, September, pp. 727-744.

Bjorklund, Anders and Markus Jantti (1997), “Intergenerational income mobility in Sweden compared to the United States”, *American Economic Review*, vol. 87, No. 5, pp. 1009-1018.

Checchi, D., A. Ichino and A. Rustichini (1999), “More equal but less mobile? Education financing and intergenerational mobility in Italy and in the US”, *Journal of Public Economics*, vol. 74, No. 3, pp. 351-394.

Ferreira, Sergio Guimaraes and Fernando A. Veloso (2006), “Intergenerational mobility of wages in Brazil”, *Brazilian Review of Econometrics*, vol. 26., No. 2 (November), pp. 181-211.

Grawe, Nathan (2001), “Intergenerational mobility in the US and abroad: Quantile and mean regression measures”, Ph.D. Thesis, Department of Economics, University of Chicago.

Jackson, Ben and Paul Segal (2004), “Why inequality matters”, A Catalyst working paper, Catalyst: London, UK.

Kopczuk, Wojciech, Joel Slemrod and Shlomo Yitzhaki (2005), “The limitations of decentralized world redistribution: An optimal taxation approach”, *European Economic Review*, vol. 29, pp. 1051-1079.

Lam, David and Robert Schoeni (1993), “Effects of family background on earnings and returns to schooling: Evidence from Brazil”, *Journal of Political Economy*, vol. 110, pp. 710-740.

Li, H., L. Squire and Heng-fou. Zou (1998), “Explaining international and intertemporal variations in income inequality,” *The Economic Journal*, 108, 26-43.

Milanovic, Branko (2002), “True world income distribution, 1988 and 1993: First calculations based on household surveys alone”, *Economic Journal*, vol. 112, No. 476, January, pp. 51-92.

Milanovic, Branko (2005), *Worlds Apart: Global and International Inequality 1950-2000*, Princeton University Press.

Milanovic, Branko (2006), “Ethical and economic case for global transfers”, mimeo, October, forthcoming in *Ethics and International Affairs*.

Nagel, Thomas (2005), “The problem of global justice”, *Philosophy and Public Affairs*, vol. 35, No. 2, pp. 113-147.

Ozden, Caglar and Maurice Schiff (2006), “Overview” in Ozden and Schiff (eds.), *International migration, remittances and brain drain*, Washington D.C.: World Bank and Palgrave.

Palma, Jose Gabriel (2006), “Globalizing inequality: “centrifugal” and “centripetal” forces at work”, United Nations DESA Working paper No. 35, September 2006.

Pogge, Thomas (1994), “An Egalitarian Law of Peoples”, *Philosophy and Public Affairs*, vol. 23, No. 3 (Summer), pp. 195-224.

Rawls, John (1971), *A theory of justice*, Cambridge, Mass.: Harvard University Press. Revised edition published in 1999.

Rawls, John (1999), *The law of peoples*, Cambridge, Mass.: Harvard University Press.

Roemer, John E. (1998), *Equality of opportunity*, Harvard University Press.

Roemer, John E. (2007), “What portion of global income inequality is due to country of residence”, mimeo. February 5, 2007.

Solon, G. (1999), “Intergenerational mobility in the labor market”, in O. Ashenfelter and D. Card (eds.), *Industrial concentration and economic inequality: Essays in honor of Peter Hart*, Edward Elgar, Aldershot, pp. 3-21.

Sutcliffe, Bob (2004), “World Inequality and Globalization”, *Oxford Review of Economic Policy*, vol. 20, No. 1, pp. 15-37.