

GLOBAL INEQUALITY OF OPPORTUNITY: HOW MUCH OF OUR INCOME IS DETERMINED BY WHERE WE LIVE?

Branko Milanovic*

Abstract—Suppose that all people in the world are allocated only two characteristics over which they have (almost) no control: country of residence and income distribution within that country. Assume further that there is no migration. We show that more than one-half of variability in income of world population classified according to their household per capita in 1% income groups (by country) is accounted for by these two characteristics. The role of effort or luck cannot play a large role in explaining the global distribution of individual income.

I. Setting the Stage

IN Rawls's *Law of Peoples* (1999b), 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* (1999a). They meet behind the veil of ignorance. Imagine now a Rawls redux similar meeting of all individuals in the world where each is handed only one characteristic that will influence his or her economic fate: country of residence. We shall ask: How much of this person's income will be determined by this factor, unrelated to individual effort or luck? Is one's position in global income distribution largely decided by country where one lives?

Assignment to country is fate, decided at birth, for approximately 97% of the people in the world: less than 3% of the world's population lives in countries where they were not born.¹ Moreover, as the differences between mean country incomes are large—more than two-thirds of global inequality between individuals is due to national income differences—to what nation one gets “allocated” is indeed of significant import for one's life chances.² By being “allo-

cated” to a country, a person receives at least two “public” goods—average income of the country and inequality of income distribution—that are unalterable by one's own effort. They will be referred to as circumstances (Roemer, 1998). To be more precise and to account for the fact that citizenship at birth is not necessarily the same as citizenship over the rest of the person's life and, moreover, that citizenship and residence may not coincide, we speak of “residence” rather than of “citizenship.”

This issue can be set in more explicitly Roemerian (1998) terms. Income (y) of i th individual in j th country can be, in general as in equation (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 (e.g., parental income, gender, or race) whose effect also depends on country (hence subscripted by j); person's own effort E_{ij} , and a random shock which can also be called luck (u_{ij}):

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

We focus on two circumstances: mean income of country j (m_j) and the Gini coefficient of country j (G_j). Our objective will be to find out how much of income can be explained by them. The formulation as written in equation (1) assumes that effort is independent of circumstances; in other words, circumstances affect income only directly, and not indirectly through effort. We could also formulate equation (1) in such a way that effort appears as an argument in each individual circumstance, $\gamma_{ij}(E_{ij})$. However, as we shall show, which way effort enters equation (1) does not matter for our estimation because in the regressions, we shall have on the right-hand side only country-specific circumstances, α_j s, which are clearly exogenous.³ We are thus agnostic as to how effort and individual circumstances interact.

Why do we study this topic at all? We have to explain the importance of the topic in stark terms not only because the rationale for studying global, as opposed to national, inequality of opportunities is new, but because the topic itself is poorly understood. We study it because we want to find out whether, globally, effort pays off—or not. The topic of inequality of opportunity is traditionally studied at the national level, and recently both the sophistication of

³ According to Roemer (1998), conditional on circumstance, people at the same percentile of effort should be rewarded the same (or treated equally). Roemer distinguishes between relative effort (“degree of effort”) and absolute effort (“level of effort”). 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 the same for each percentile of the distribution of effort (i.e., for each relative effort), thus allowing the same absolute effort to be rewarded differently.

Received for publication January 3, 2010. Revision accepted for publication October 1, 2013. Editor: Philippe Aghion.

* University of Maryland.

I am grateful to Igvild Almås, Patrick Conway, Jim Davies, James K. Galbraith, Francisco Jose Goerlich, Elise Huillery, Francisco Ferreira, Michael Lokshin, John Roemer, Xiaobo Zhang, two anonymous referees, and the editor of the *Review* for very helpful comments. The paper has been improved by the comments of participants at various meetings where it was presented: World Bank, International Association for Research in Income and Wealth, Free University Amsterdam, University of Texas, SciencesPo in Paris, and University of Maryland. I am grateful to Gouthami Padam for excellent help in collecting and processing the data.

¹ The stock of migrants around the year 2000 was estimated at 165 million (see Ozden et al., 2011). The annual flow of people who move between countries (excluding tourism and very short visits) is estimated at 11 million, about 1/7 of 1% of the world's population (see Pritchett, 2006).

² See Milanovic (2002, 2005), Sutcliffe (2004), Bourguignon and Morisson (2002), and Berry and Serieux (2007). This result is obtained using the standard (Pyatt) Gini decomposition, which is appropriate in this case because it calculates the between-country component assuming that everybody in a given country has the mean income of that country. An alternative decomposition is proposed by Yitzhaki and Lerman (1991) and Frick et al. (2006). Its between component is always equal to or smaller than Pyatt's (see Frick et al., 2006). For the world, however, the differences are minimal. Using 2005 global data (see more on this below), the Pyatt between-country component is 61.5 Gini points (out of total global Gini of 70), while the Yitzhaki and Lerman is 58 Gini points.

TABLE 1.—POPULATION AND INCOME COVERAGE OF THE SURVEYS, 2008

	Africa	Asia	Latin America	East Europe and CIS	WENAO	World
Population	78%	98%	97%	92%	97%	94%
Dollar income	71%	93%	98%	98%	97%	96%
Number of countries	23	27	18	27	23	118

WENAO: Western Europe, North America, and Oceania (Australia and New Zealand). CIS = Commonwealth of Independent States. East Europe includes formerly Communist countries. World Income Distribution database. The data are available at <http://econ.worldbank.org/projects/inequality>.

the analysis (Bourguignon, Ferreira, & Menéndez, 2007) and the coverage of the countries have expanded.⁴ Suppose, in a given country, that one's income is entirely determined by one's parents' income. Not only would we deem this unjust, but economically the rationale for working hard would be lacking.

In a globalized world, the same question may be asked as well. Let income depend entirely on country of birth, thus implying that inequality within each country is zero. A person who is born in a poor country cannot by her own efforts improve her lot domestically or globally because she alone cannot influence her country's growth rate. It thus makes no sense for her to expend effort that will lead to no improvement in income. The only avenue that remains is migration. Posed in such extreme terms, it is easy to see why the question is important: it raises not only ethical issues (Is it fair that income should be decided at birth?) but because it has clear economic implications, where should the efforts of people in poor countries be directed: to work or to migrate?

In section II, we describe global income distribution data that help address these questions empirically and show some broad regularities about the way global income is distributed among countries and income classes. Sections III and IV are the core parts of the paper: they present the analysis that seeks to answer the questions we pose. Section V gives the conclusions.

II. Data and Definitions

The data used in the paper come from World Income Distribution (WYD) database constructed to study the evolution of global inequality. The database is composed almost entirely of microdata from representative national household surveys from most of the countries in the world. For the year 2008, which we use here, the data come from 118 countries' household surveys representing 94% of the world's population and 96% of world dollar income.⁵ (The list of countries, surveys, and other information about the database is available from the author on request.) The geo-

⁴ This is partially due to data availability, but probably more importantly to the unstated view that equality of opportunity is something that ought to hold at the national level or for which only national governments can be held responsible. But if we extend our consideration to the world as a whole, should not equality of opportunity apply to all individuals regardless of their nationality?

⁵ I cannot express the share of the included countries in terms of \$PPP income because for most of the countries for which we lack surveys (e.g., Afghanistan, Iraq, Sudan), we also lack PPP data. The dollar incomes, however, are typically available.

graphical coverage is almost complete for all parts of the world except Africa (see table 1).⁶

For all countries but one (Singapore) we have microdata, which means that any type of distribution (by decile, ventile, percentile; by household or individual) could have been created. 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). In order to provide precise income estimates while making the analysis manageable, we combine individuals into corresponding income percentiles and use a relatively dense distribution of 100 data points (percentiles) per country. The percentiles range from the poorest (percentile 1 or bottom percentile) to the richest (percentile 100 or top 1 percent).

Since not all countries produce annual surveys, we had to use a benchmark year (2008 in this case); that is, we sought 2008 household surveys for as many countries as possible. Where there are no surveys conducted in 2008, we used a year close to 2008. In the event, 89 of 118 household surveys were conducted in the benchmark year or one year before or after it, and all but 2 surveys were done within two years of the benchmark year. For the surveys conducted in nonbenchmark years, we adjust reported incomes by the consumer price index of the country so that all amounts are expressed in 2008 local currency units. These amounts are then converted into international (PPP) dollars using the 2008 estimates of \$PPP exchange rates for household private consumption provided by the newest round of the International Comparison Program.⁷ The PPPs are calculated by the Eltôte-Köves-Szulc method. For each percentile of population, we calculate the average annual per capita amount of PPP dollars received as disposable income.⁸


The fact that each country is divided into 100 groups of equal size (percentiles) is very helpful. This allows us to compare the positions of, for example, the 23rd percentile

⁶ An earlier version of this paper used the same WYD database but with the benchmark year of 2005. The results are quasi-identical to the ones reported in sections III and IV of this paper. We can thus argue that the results hold for at least two annual cross-sections of household surveys across the world.

⁷ This new round of the International Comparison Program has led to a sharp upward revision of China's and India's price level and, consequently, the sharp downward revision of their incomes (World Bank, 2008; Milanovic, 2012). These new results have been incorporated in the World Bank's more recent global poverty calculations. For the data, see <http://siteresources.worldbank.org/ICPINT/Resources/icp-final-tables.pdf>.

⁸ Household surveys are either income or expenditure (consumption) based. For simplicity of presentation, we refer to *income* distribution and *income* position in the world throughout this paper.

FIGURE 1.—INCOME LEVELS IN THE WORLD BY COUNTRY AND INCOME CLASS, 2008



World Income Distribution (WYD). Horizontal axis: country percentiles that run from 1 (poorest) to 100 (richest). Vertical axis: annual household per capita disposable income in 2008 international dollars (in logarithms). The horizontal line is drawn at the income level of the poorest German percentile.

of people in China with the 75th percentile in Nigeria. It also allows us to define income classes in the same way across countries. To fix the terminology, we call each percentile an income class. Income classes thus run from 1 to 100, with 100 being the highest. Incomes within all percentiles except the very highest one, and sometimes the poorest one, are extremely homogeneous. Gini coefficients for within-percentile individual incomes are generally less than 1 or 2 (i.e., less than 0.01 or 0.02), and it is only for the very top percentile that Gini takes two-digit values.⁹ Thus, except for the top 1%, within all other country/percentiles we deal with individuals whose household per capita incomes are practically undistinguishable from one another.

Figure 1 shows the average percentile incomes for five countries (obviously the same figure could be done for any of the 118 countries). Consider Germany. Because Germany is a rich country and its income inequality is moderate, most of its population is highly placed in the world income distribution. The poorest German population percentile has a per capita disposable income of about PPP \$2,200 per year (see the horizontal broken line). All other percentiles' incomes are obviously greater, and the richest percentile has an income per capita of about PPP \$104,000, which places it also in the top world percentile. The same interpretation holds for other countries. Unlike relatively egalitarian Germany, where the ratio between the richest and the poorest percentile is less than 50 to 1, the ratio in China between the top and bottom percentile is 66 to 1, with the poorest section of the Chinese population having an annual per capita income just under PPP \$300 and the richest percentile earning almost PPP \$20,000. Only about 40% of the Chinese population is richer than the poorest

Germans. This percentage is even smaller in the case of India. Brazil, with its unequal income distribution, covers almost the entire global spectrum, with the poorest people at less than PPP \$300, and the richest percentile at PPP \$60,000.

III. Predicting Income from Knowledge of Country of Residence Only

Using the just-discussed data, we can express the income level of people belonging to percentile i living in country j as follows:

$$y_{ij} = b_0 + b_1 m_j + b_2 G_j + \varepsilon_{ij}, \quad (2)$$

where y_{ij} is the annual average household per capita income in \$PPP, m_j is the country's GDP per capita in PPP terms, G_j is inequality in income distribution obtained from household surveys and measured by the Gini coefficient, and ε_{ij} is the error term. Both variables on the right-hand side are strictly exogenous to an individual effort: by her efforts, a person cannot affect, in any meaningful way, her country's level of GDP per capita or change her country's Gini.¹⁰ This is both substantively important for our analysis and econometrically convenient. It should also be noted that our objective here is not a precise explanation of the income level of each country or percentile (which could be improved if we used more explanatory variables) but rather to find out how much of global income variability can be accounted for by an extremely parsimonious formulation where just a few undeniable country-specific variables are used. We use GDP per capita instead of mean income from household surveys in order to avoid the reflexivity problem whereby the coefficient on the mean (b_1) would be biased toward 1.¹¹ This would happen because the arithmetic average of percentile values, our dependent variables, is equal to the mean.¹²

Two specification issues need to be addressed. First, we need to decide whether the regression will take into account countries' population sizes. (Notice that figure 1 implicitly treats the population sizes of all countries as the same.) Two different points of view are possible. If a person looks at, say, her own income only and asks, "How well would I have fared had I been born or lived in a different country?" then population sizes of countries do not matter. A person simply looks at her current income and compares it with the income that she might have if she were in a given percentile of income distribution in the United States or China or somewhere else. From that individual viewpoint, the popu-

¹⁰ If the regression contained individual circumstances (e.g., gender, race, age) that can plausibly be correlated with effort, the assumptions regarding how efforts enters individual income-formation equation such as equation (1) would matter.

¹¹ Thanks to a referee for pointing this out.

¹² The coverage of income or consumption in household surveys is much narrower, since it pertains to the household sector, than the coverage of GDP. In almost all cases, the household survey mean is lower than GDP per capita (Deaton, 2005).

⁹ Calculated from microdata from 25 surveys (5 from each region) used in our 2008 benchmark year (available from the author).

TABLE 2.—HOW INCOME DEPENDS ON CIRCUMSTANCES
DEPENDENT VARIABLE: NATURAL LOG OF HOUSEHOLD PER CAPITA INCOME IN \$PPP FOR EACH COUNTRY OR PERCENTILE

	Individual Viewpoint (unweighted regressions)			World as It Is (population-weighted regressions)		
	GDP per Capita (in logs)	Average Number of Years of Schooling	Country Dummies	GDP per Capita (in logs)	Average Number of Years of Schooling	Country Dummies
	1	2	3	4	5	6
Proxy for mean country income	0.868 (0)	0.335 (0)	—	1.011 (0)	0.408 (0)	—
Gini index (in %)	−0.015 (0)	−0.013 (0)	—	−0.012 (0)	−0.015 (0.14)	—
Constant term	0.800 (0.02)	5.779 (0)	5.220 (0)	−0.711 (0.28)	5.250 (0)	5.288 (0)
Country dummies	—	—	Yes	—	—	Yes
Number of observations	11,483	9,083	11,683	11,483	9,083	11,683
Number of countries (clusters)	115	91	117	115	91	117
Population weight (in million)	—	—	—	6,133	5,711	6,139
R^2	0.660	0.481	0.733	0.610	0.534	0.657

The regressions are run with the cluster option to adjust for the correlation of within-country observations. In the base-case regressions, (1) and (4), there are 115 countries, each with 100 percentiles, with the exception of Switzerland, where the bottom 13 percentiles are missing, and Lithuania, where 4 percentiles are missing. For Palestine and Kosovo, there are no data on GDP per capita in PPP terms, and they are not included in regressions (1) and (4). For Singapore, we do not have microdata, and the country is not included. p -values are in parentheses. (0 indicates significance at a level smaller than 0.000). All income variables are in 2008 international dollars. Coefficients on country dummies in regressions 3 and 6 are not shown here.

lation size of China, the United States, or any other country is immaterial. We shall call this approach the *individual viewpoint* (IV). But if we want to look at how the world is actually structured, then clearly population size matters. We call this second approach *the world as it is* (WAI).

Second, in order to test the robustness of the results—in particular, the importance of the country of residence for determining individual incomes—we shall use several proxies in addition to our base case variable of GDP per capita. Thus, we replace GDP per capita by the average number of years of education of the population over the age of 15 (AYOS).¹³ Education is a strong proxy of average income but is, of course, a distinct variable. An alternative is to run a simple LSDV (least square dummy variable) regression where country dummies replace both the mean income variable and Gini. Here we do not retrieve an overall income coefficient valid across the countries but a coefficient on each country dummy, which provides a country's location premium or penalty (with respect to one worldwide comparator country). We also focus on whether such a simple regression explains enough of the variability of individual income percentiles across the world.

Table 2 shows the results for the two scenarios (unweighted, IV, and population-weighted, WAI) and three specifications.

We consider first the individual viewpoint scenario. In the base case (regression 1), elasticity of own income with respect to country's GDP per capita is 0.866. We can call this the locational premium. The Gini coefficient enters with a negative sign, indicating that living in a more unequal country on average reduces one's income. A 1 Gini point increase is associated with a 1.5% decrease in own income. This reflects the fact that higher inequality numeri-

cally benefits fewer people than it harms.¹⁴ Overall, these two circumstances explain two-thirds of the variability of individual percentile incomes across the world.

When we use the average number of years of schooling (regression 2), the results change: R^2 drops to 0.48 (the number of countries for which we have data also drops from 115 to 91). The increase of a country's average educational level by one additional year of schooling is associated with an increase of individual incomes of more than 30%. When we include both AYOS and AYOS², the results (available from the author on request) show that the coefficient on AYOS decreases to about 0.25 (25%), while the one on AYOS² is positive but not statistically significant.

Finally, in regression 3, country dummies alone explain almost three-quarters of the variability of individual percentile incomes across the world. There are 117 countries in the regression. The omitted country dummy belongs to the Democratic Republic of Congo (DR Congo, formerly Zaire), the poorest country in the sample, so that the coefficients on individual country dummies show the locational premium that a person on average obtains by being a resident of a country other than DR Congo. For example, the U.S. locational premium is 355%, Sweden's is 329%, Brazil's is 164%, and Russia's is 230%, but Yemen's is only 32%.

In regressions (4) to (6), we look at the results within the context of the world as it is. This shows the importance of circumstances as actually experienced by the people in the world, and thus more populous countries will matter more. The results of population-weighted regressions are not very

¹³ Data obtained from the 2012 version of World Bank World Development Indicators, in turn partially based on the Barro and Lee data set (2013).

¹⁴ As we shall see below, greater inequality has a differential impact depending on where one is in one's country's distribution: it benefits higher-income classes (whose income goes up) and harms lower-income classes (whose income declines). Overall, there are more of the latter, and that is why in regressions such as equation (1), the coefficient on Gini is negative.

TABLE 3.—GLOBAL INEQUALITY OF OPPORTUNITY, 1988–2008
BETWEEN-COUNTRY COMPONENT OF GLOBAL INTERPERSONAL INEQUALITY

	Benchmark Year					
	1988	1993	1998	2002	2005	2008
(1) Between-country Gini	62.4	62.1	61.7	63.0	61.5	59.8
(2) Between-country Theil	86.2	79.2	76.4	77.3	76.6	67.7
(3) Global interpersonal Theil	107.0	104.9	103.5	104.9	104.2	98.3
(4) Share (in percent) of between-country in total interpersonal inequality (2)/(3)	81	76	74	74	74	70
(5) Population included in calculation (in <i>m</i>)	4,477	5,146	5,427	5,795	5,917	6,142
(6) Included population as % of world population	88	93	92	93	92	92
(7) Number of countries	100	119	121	119	119	118

The between-country component is population weighted. Theil is Theil (0) or mean log deviation index. Both Gini and Theil are given in percent.

different from the unweighted. The elasticity of own income with respect to GDP per capita is close to 1, and greater inequality (controlled for income) reduces more people's incomes than it raises. When we use *AYOS*, the locational premium, measured as returns to an additional year of (nationwide) education, is now higher than in the individual viewpoint scenario, indicating that more populous countries seem to benefit more from a given increase in the average educational level. When we use country dummies, the individual countries' locational premiums (computed with respect to being a resident of DR Congo) are unchanged. The key result, overall role of country circumstances, measured by R^2 , ranges between 53% and 66%. In the individual viewpoint case, the bounds were wider, from 48% to 73%.

Is the importance of country of residence for one's income increasing? To answer that, we calculate the standard inequality of opportunity index where people in the world differ by only one characteristic: country of residence.¹⁵ Global inequality of opportunity is then equal to the between-country component of an inequality statistic. Luckily, we have comparable data on global interpersonal inequality for six benchmark years in the period 1988 to 2008 and can calculate the between-country component. The sample size (countries included in the calculation) is basically the same for all benchmark years but the first. The population covered by the household surveys accounts for more than 90% of the world's population (in all years except 1988). All incomes are expressed in international (PPP) dollars.

Table 3 shows the results of calculation of the global inequality of opportunity using the Gini and Theil (0) indexes. Both indexes display more or less steady decline (the exception is the period between 1998 and 2002), although the decline, measured by Theil, is more significant (almost 20% over the entire period) than measured by Gini (about 4%). This is due to the greater dependence of the Gini on the mode of the distribution. A between-country component as a share of global interpersonal inequality has gone down over the same period, from accounting for 81%

of the total to 70%.¹⁶ Global inequality of opportunity due to place of residence is, as seen in the importance of the between-country component, huge but decreasing. The decrease is driven by rapid growth of relatively poor and populous countries, in particular China and India.¹⁷

In conclusion, whatever scenario or specification we select, at least around half of the variability in real (\$PPP) personal percentile incomes in the world can be attributed to two circumstances beyond individual control: level of development of one's country of residence, proxied by its GDP per capita or average number of years of education, and inequality of distribution within that country. When we replace both the income proxy and the Gini coefficient by country dummies, reflecting all unobservable country characteristics, the R^2 ranges between 66% and 73%. The part that remains for effort and "episodic luck" (to use John Roemer's felicitous phrase) is, within the worldwide context, relatively limited. This was true in 2008 despite a steady erosion of the importance of between-country components in global interpersonal inequality.

IV. Locational Premium across Income Classes

By construction, the location premium was so far assumed to be equal across all percentiles of income distribution, that is, the same for a given country regardless of person's place in his or her country's income distribution. But the locational premium need not be uniform across the entire distribution: it may vary between different parts of the distribution. To make the analysis more manageable, we use income ventiles (each ventile contains 5% of population, ranked from the poorest to the richest).

Table 4 shows the results of regressions similar to equation (1) but with the person's own income ventile held constant.¹⁸ For each ventile separately, we regress ventile

¹⁶ The share is shown only in terms of Theil. As is well known, the advantage of Theil over Gini in such a decomposition is that it is additively decomposable. In addition, the Theil (0) index has a feature, pointed out by Anand and Segal (2008), that it is internally consistent: thus, eliminating all between-country inequality would leave the within-country inequality component unchanged, which is not the case with Theil (1) index.

¹⁷ See, for example, Milanovic (2012).

¹⁸ Table 4 gives the results only for the unweighted regressions. Population-weighted regressions ("the world as it is") are given in the table appendix.

¹⁵ For a review of measurement of inequality of opportunity, see Brunori, Ferreira, and Peragine (2013), in particular their section 3.

TABLE 4.—EXPLAINING A PERSON'S POSITION IN THE WORLD INCOME DISTRIBUTION GIVEN HER NATIONAL INCOME CLASS (VENTILE)
DEPENDENT VARIABLE: NATURAL LOG OF HOUSEHOLD PER CAPITA INCOME IN \$PPP, 2008 (UNWEIGHTED REGRESSIONS; "INDIVIDUAL VIEWPOINT")

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
GDP per capita (logs)	0.769 (0)	0.830 (0)	0.845 (0)	0.857 (0)	0.865 (0)	0.871 (0)	0.873 (0)	0.878 (0)	0.881 (0)	0.880 (0)	0.885 (0)	0.883 (0)	0.886 (0)	0.886 (0)	0.886 (0)	0.886 (0)	0.885 (0)	0.882 (0)	0.876 (0)	0.862 (0)
Gini	-0.058 (0)	-0.044 (0)	-0.038 (0)	-0.033 (0)	-0.030 (0)	-0.027 (0)	-0.024 (0)	-0.021 (0)	-0.019 (0)	-0.016 (0)	-0.014 (0)	-0.011 (0)	-0.008 (0)	-0.006 (0)	-0.003 (0)	-0.0001 (0)	0.003 (0.36)	0.007 (0.04)	0.013 (0)	0.029 (0)
Constant	1.879 (0)	1.294 (0)	1.098 (0)	0.975 (0)	0.882 (0.01)	0.809 (0.02)	0.773 (0.03)	0.719 (0.04)	0.680 (0.05)	0.663 (0.06)	0.615 (0.08)	0.611 (0.08)	0.588 (0.09)	0.578 (0.10)	0.571 (0.10)	0.562 (0.11)	0.564 (0.10)	0.575 (0.10)	0.401 (0.25)	0.764 (0.03)
Adjusted R ²	0.90	0.91	0.91	0.91	0.91	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.89	0.89	0.89	0.89	0.87
Number of observations	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115
F value	517	550	45	541	528	520	519	509	502	507	501	503	494	492	487	483	478	473	450	376

GDP per capita in \$PPP. *p*-values are in parentheses.


income on country's GDP per capita and Gini coefficient. These two characteristics explain about 90% of the variability of income. To be clear on what it means, if we take all people who are in a given ventile of their country's income distributions (say, third or tenth ventile) some 90% of the variability of their incomes will be explained by GDP per capita and Gini coefficients of the countries where they live. In other words, the average income of the people in each ventile will largely depend on the mean income of their country and its distribution.¹⁹ The locational premium now varies across ventiles: it is relatively low for the bottom ventile (0.769), after which it rises; at the maximum, it reaches around 0.88. This means that while the locational premium holds for everyone (people in any ventile are better off if they live in a richer than in a poorer country), the premium is less for those in the lowest ventiles of income distribution.

The two country characteristics (mean income and inequality) can also be seen as substitutes: given her income class, a person might gain more by being "allocated" into a more equal society even if its mean income is less. 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 country income and inequality is not the same across income classes. If we look at the bottom income class (as in regression 1 in table 4), we see that each Gini point increase is associated with a 5.75% loss of income (greater inequality, with given mean income, will harm the poor). To exactly offset this, a person in the bottom ventile would have to be located in a country whose GDP per capita is about 7.5% higher (5.75 divided by the coefficient on mean income in regression 1, which is 0.769). This is the shape of the trade-off faced by those in the lowest ventile. For the second lowest ventile, the GDP per capita increase needed to offset 1 Gini point higher inequality is 5.3%.

The equivalent GDP per capita increases gradually decline as we move toward higher ventiles and become close to 0 around the sixteenth ventile (see figure 2). People belonging to the seventeenth ventile and richer benefit from increased inequality. The trade-off for them works in the opposite direction. To leave them with the same income, more unequal national distribution (from which they gain) has to be combined with lower GDP per capita (from which they, like everybody else, lose). For the top ventile, we see that each Gini point change is offset by about a 3.2% change of GDP per capita in the opposite direction. In other words, for the nationally rich, national distribution matters, but it matters less than for the poorest.

¹⁹ Note that we expect R^2 to be higher when we hold ventiles fixed than when, as in table 2, percentiles at once. In the former case, we compare only (say) the poorest people in poor and rich countries, and in such a case, country characteristics will be of overwhelming importance. When we have all country/percentiles together, the rich people from poor countries will "mix" with poor people from rich countries, and the role of circumstances will be less.

FIGURE 2.—HOW MUCH 1 GINI POINT CHANGE IS WORTH



Measured in terms of mean country income. Calculated from table 4.

The importance of change in national income distribution (represented by an increase or decrease of 1 Gini point) displays a U-shaped pattern with a substantially higher left end, as shown in figure 2. Those for whom national inequalities are important are either those at the bottom who gain from lower inequality or those at the very top who gain from higher inequality. For those around the middle (ventiles 13 to 18), equality or inequality of national income distributions matters very little because their income shares are about the same in both equal and unequal countries (on this point, see also Palma, 2011). Thus for them, the mean income of the country where they live is of crucial importance.²⁰

All of this leads us to two conclusions. First, while everybody (the poor, middle class, and the rich) benefits from higher mean income, that benefit is proportionately greater for the rich classes. Second, distributional change matters to the poor and to the rich (in the opposite directions, of course), while it is of little importance to the middle class. What seems to matter to the income of the middle class is whether the county is getting richer or poorer, not whether it is becoming more or less equal.

V. Conclusion

We present the conclusions first in a form of a metaphor and then list them specifically.

Imagine global income distribution as a long pole, similar to a flagpole, on which income levels (percentiles) are marked from the bottom, around the subsistence minimum of some PPP \$200–300, to the maximum household per capita income of almost PPP \$200,000. Imagine then each country's distribution to be given by a plaque, running along the pole, and covering the range of that country's income distribution by percentiles. India's plaque, for example, will run from PPP \$250 to PPP \$7,000. Korea's

²⁰ This can be also seen from the fact that the regression coefficient on Gini for ventiles 14 to 18 is not statistically significantly different from 0 (see table 4).

from PPP \$1,600 to PPP \$80,000, and that of the United States from PPP \$2,500 to PPP \$180,000. When a person is born, she gets assigned to a place on her country's plaque, which not only gives her position in national income distribution but also locates her in global income distribution. How can she improve her position? Effort or luck may push her up the national plaque. But while effort or luck can make a difference in individual cases, they cannot, from a global perspective, play a very large role because more than half of variability in income globally is explained by circumstances given at birth. She can hope that her country will do well: the country's plaque will then move up along the global pole, carrying, as it were, the entire population with it. If she is lucky enough so that her effort (movement higher up along the plaque) is combined with an upward movement of the plaque itself (increase in national mean income), she may perhaps substantially climb up in the global income distribution. Or, as a last possibility, she might try to move from a lower plaque (poorer country) to a higher one (richer country). Even if she does not end up at the high end of the new country's income distribution, she might still gain significantly. Thus, own efforts, hope that one's country does well, and migration are three ways in which people can improve their global income position.

Let us now go back to a more conventionally stated conclusion. First and most important, with only one or two circumstances, GDP per capita and income inequality of country of residence, or simply with country dummy variables, we are able to account for more than half of the variability in personal percentile incomes around the world (in only one formulation is R^2 just marginally less than half). The finding holds whether we run regressions simply across countries or use population weighting. Other features (gender, race, or ethnicity), which are not included in this analysis, would increase the share of circumstance. The role of place of residence calculated here is therefore a lower bound to global inequality of opportunity. The locational premium is very large: compared to living in the poorest country in the world (DR Congo), a person gains more than 350% if she lives in the United States, more than 160% if she lives in Brazil, but only 32% if she lives in Yemen.

Second, the ability to predict well a person's income from only these two country characteristics holds also for each income class separately. Thus, given the income class of a person (her country and income ventile), the knowledge of the country where that person lives is sufficient to explain about 90% of the variability of incomes globally. The locational premium is positive for the entire spectrum of national income distributions.

Third, again given the person's own income class, there is a trade-off between the GDP per capita of the country and its income distribution. Thus, a person who is in a low-income class might prefer to live in a more egalitarian country even if that country's GDP per capita is less. The opposite, of course, holds for a person in a high-income class: she might benefit from a country's inegalitarian dis-

tribution more than from its high GDP per capita. But these sharp trade-offs between national inequality and mean income hold mostly for the extreme income classes. For the middle classes, national distribution is relatively unimportant because income shares of the middle ventiles do not vary much across nations, whether the nations are equal or not. For the middle classes, therefore, the mean income of the country where they live will be the key factor in determining their own income level.

REFERENCES

- Anand, Sudhir, and Paul Segal, "What Do We Know about Global Inequality?" *Journal of Economic Literature* 46:1 (2008), 57–94.
- Barro, Robert, and Jong-Wha Lee, "A New Data Set of Educational Attainment in the World," *Journal of Development Economics* 104 (2013), 184–198.
- Berry, Albert, and John Serieux, "World Economic Growth and Income Distribution, 1980–2000" (pp. 1–23), in K.S. Jomo and Jacques Baudot, eds., in *Flat World, Big Gaps* (London: Orient Longman, Zed Press and Third World Network, 2007).
- Bourguignon, François, Francisco H. B. Ferreira, and Marta Menéndez, "Inequality of Opportunity in Brazil," *Review of Income and Wealth* 53 (2007), 585–618.
- Bourguignon, François, and Christian Morrisson, "The Size Distribution of Income among World Citizens, 1820–1990," *American Economic Review* (2002), 727–744.
- Brunori, Paolo, Francisco Ferreira, and Vito Peragine, "Inequality of Opportunity, Income Inequality and Economic Mobility: Some International Comparisons," World Bank Policy Research working paper 6304 (January 2013).
- Deaton, Angus, "Measuring Poverty in a Growing World (or Measuring Growth in a Poor World)," *this REVIEW* 87 (2005), 1–19.
- Frick, Joachim, Jan Goebel, Edna Schechtman, Gert G. Wagner, and Shlomo Yitzhaki, "Using Analysis of Gini (ANOGI) for Detecting Whether Two Subsamples Represent the Same Universe: The German Socio-Economic Panel Study (SOEP) Experience," *Sociological Methods and Research* 34 (2006), 427–468.
- Milanovic, Branko, "True World Income Distribution, 1988 and 1993: First Calculations based on Household Surveys Alone," *Economic Journal* 112 (2002), 51–92.
- , *Worlds Apart: Global and International Inequality 1950–2000* (Princeton, NJ: Princeton University Press, 2005).
- , "Global Inequality Recalculated and Updated: The Effect of New PPP Estimates on Global Inequality and 2005 Estimates," *Journal of Economic Inequality* 10 (2012), 1–18.
- , "Evolution of Global Inequality: From Class to Location, from Proletarians to Migrants," *Global Policy* 3 (2012), 124–133.
- Ozden, Caglar, Christopher Parsons, Maurice Schiff, and Terrie Walmsley, "Where on Earth Is Everybody? The Evolution of Global Bilateral Migration 1960–2000," *World Bank Economic Review* 25 (2011), 12–56.
- Palma, José Gabriel, "Homogeneous Middles vs. Heterogeneous Tails, and the End of the 'Inverted-U': The Share of the Rich Is What It's All About," *Development and Change* 42 (2011), 87–153.
- Pritchett, Lant, *Let Their People Come* (Washington, DC: Center for Global Development, 2006).
- Ravallion, Martin, "Should Poverty Measures Be Anchored to the National Accounts?" *Economic and Political Weekly* 34 (2000), 3245–3252.
- Rawls, John, *A Theory of Justice*, rev. ed. (Cambridge, MA: Harvard University Press, 1999a).
- , *The Law of Peoples* (Cambridge, MA: Harvard University Press, 1999b).
- Roemer, John E., *Equality of Opportunity* (Cambridge, MA: Harvard University Press, 1998).
- Sutcliffe, Bob, "World Inequality and Globalization," *Oxford Review of Economic Policy* 20 (2004), 15–37.
- World Bank, *Global Purchasing Power Parities and Real Expenditures: 2006 International Comparison Program* (Washington, DC: World Bank, 2008).
- Yitzhaki, Shlomo, and Robert I. Lerman, "Income Stratification and Income Inequality," *Review of Income and Wealth* 37 (1991), 313–329.

Table Appendix

The table explains a person's position in the world income distribution given her own national income class (ventile). The dependent variable is the natural log of household per capita income in \$PPP, is 2008. The regressions are population weighted—the "world as it is".

Income Class	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
GDP per capita (logs)	0.849 (0)	0.949 (0)	0.981 (0)	1.000 (0)	1.011 (0)	1.023 (0)	1.029 (0)	1.040 (0)	1.038 (0)	1.042 (0)	1.045 (0)	1.045 (0)	1.045 (0)	1.043 (0)	1.042 (0)	1.039 (0)	1.035 (0)	1.028 (0)	1.017 (0)	0.980 (0)
Gini	-0.061 (0)	-0.045 (0)	-0.038 (0)	-0.034 (0)	-0.030 (0)	-0.026 (0)	-0.023 (0)	-0.016 (0)	-0.016 (0)	-0.013 (0)	-0.009 (0)	-0.005 (0)	-0.002 (0)	0.001 (0)	0.005 (0)	0.008 (0)	0.011 (0)	0.014 (0)	0.018 (0)	0.025 (0)
Constant	1.216 (0.01)	0.148 (0.78)	-0.190 (0.75)	-0.405 (0.51)	-0.556 (0.39)	-0.680 (0.30)	-0.772 (0.26)	-0.950 (0.21)	-0.945 (0.18)	-1.020 (0.15)	-1.100 (0.13)	-1.156 (0.12)	-1.191 (0.11)	-1.212 (0.10)	-1.224 (0.10)	-1.211 (0.10)	-1.169 (0.10)	-1.077 (0.12)	-1.155 (0.08)	-0.349 (0.56)
R ²	0.92	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93
Number of Observations	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115	115
F value	390	368	324	303	287	276	267	267	257	266	268	267	261	253	246	243	239	239	231	205
	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)

GDP per capita in \$PPP. *p*-values in parentheses, with 0 indicating significance at the level smaller than 0.000.