ECON 133: Global Inequality and Growth
Lecture Review #3: Inequality between Individuals

Jakob Brounstein *

1 Lorenz curve and Gini coefficient

Lorenz curve for a distribution of income is a curve formed by ranking people according to their income (or whatever object of interest), and then plotting their cumulative share of total income as one moves up the distribution. The curve starts at 0 and ends at 100 per cent; if all incomes were identical, the curve would follow the diagonal joining these end points (the line of equality, the 45 degree line: $y = x$).

The **Gini coefficient** is a measure of relative inequality with values lying between 0 (complete equality; everyone gets the same income) and 100 per cent (complete inequality; one person gets all the income and everyone else gets nothing). Geometrically, the Gini coefficient is the area between the Lorenz curve and the line of equality, divided by the area of the whole triangle: Gini $= \frac{A}{A+B} = 2A = 1 - 2B$, because the triangle defined by x-axis from 0 to 1,[1] the line $y = x$, and the vertical line $x = 1$ has area $A + B = 0.5$. Computationally, multiplying by two just represents a scale normalization, and the Gini coefficient is typically seen expressed between 0 and 1 or between 0 and 100.

Figure 1: Examples of Gini and Lorenz curves

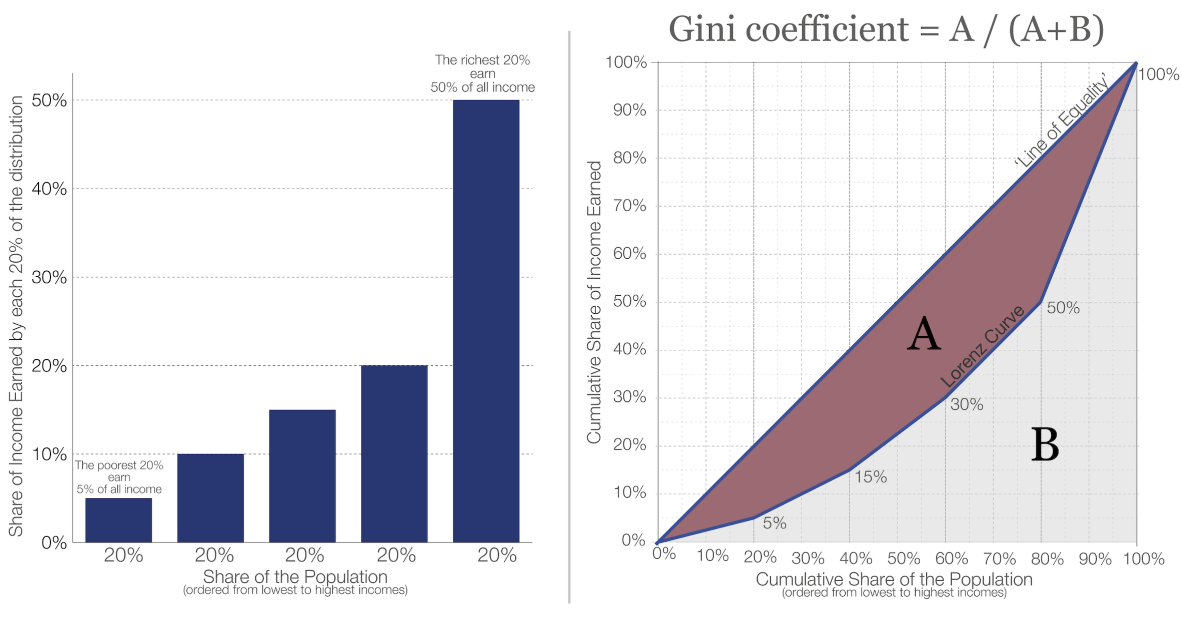

*These notes borrow from past notes by Cristóbal Otero, Nina Roussille, Juliana Londoño-Vélez, Marcelo Milanello, and John Schellenberg. All mistakes are our own.

[1] NB: The axes are axes are scaled from 0 to 1.

Figure 2: Gini Index World Map

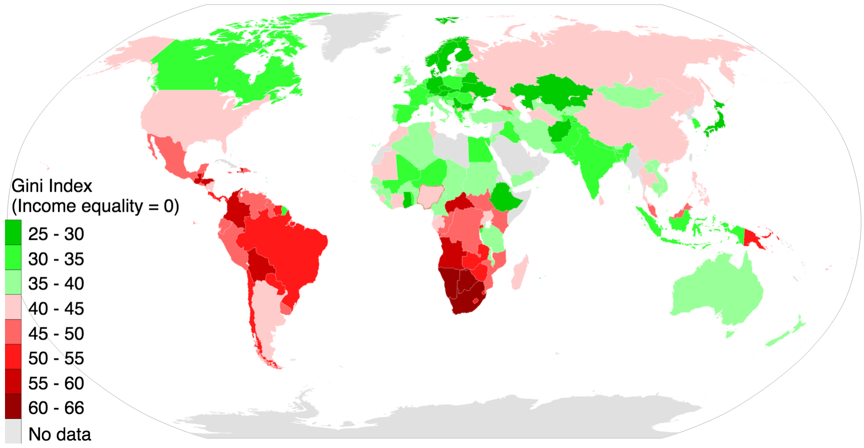

Source: The World Bank (2014) World Development Indicators: Distribution of income or consumption (Table 2.9). Available online here: http://wdi.worldbank.org/table/2.9/

One critique of the Gini coefficient is that it does not lend to incredibly straightforward intuition. How can we intuitively parse the difference between Gini coefficients of .25, .4, .55? One way for us to build some intuition will involve a highly stylized example.

Another critique: computing a Gini coefficient either requires considerable data or requires potentially important income homogeneity assumptions within discrete groups.

Suppose that we have two homogeneous groups, one below percentile p_1 with income share s_0 and one *above* percentile p_1 with income share $s_1 = 1 - s_0$. When we say *homogenous*, we mean that the income distribution is identical *within each group*. Visually, that would correspond to a piecewise, linear Lorenz curve with two total parts, as in Figure 1 (following the two connected diagonal line segments in the triangle's interior). In this case, it is easy to show that the Gini coefficient is as follows:

$$G = s_1 + p_1 - 1$$

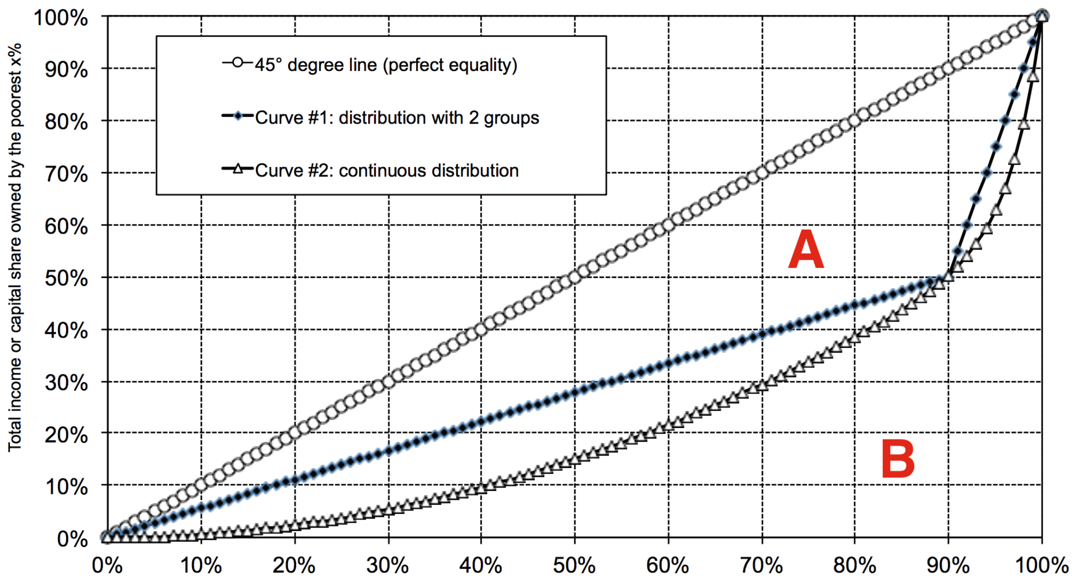

Curve 1 assumes that the poorest 90% and the richest 10% own 50% of total income or capital each, and that both groups are homogenous (hence a linear curve); curve 2 assumes a continuous distribution

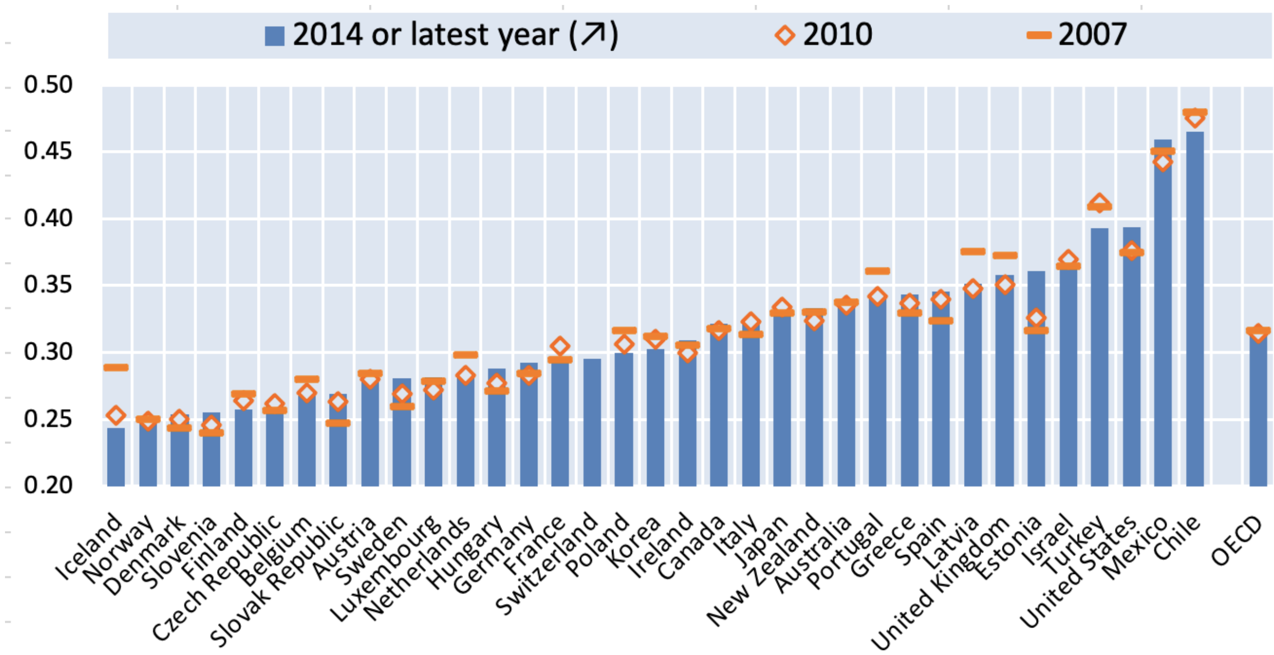

$$
\begin{aligned}
G &= 2A \\
&= 2\left(\frac{1}{2} - B - C\right) \\
&= 2\left(\frac{1}{2} - \frac{p_1 s_0}{2} - \frac{(1+s_0)(1-p_1)}{2}\right) \\
&= 1 - p_1 s_0 - 1 - s_0 + p_1 + p_1 s_0 \\
&= p_1 - s_0 \\
&= s_1 + p_1 - 1.
\end{aligned}
$$

What happens when s_0 or s_1 increases/decreases? What happens when p_1 increases/decreases?

Example to give intuition: two-group societies where 1) top 5% have 50% of income; 2) top 10% have 20% of income? What does the Gini coefficient look like in your ideal distribution of income?

1.1 Alternatives

Due to critiques around intuition/interpretability and intensive data requirements, some people prefer other parameterizations to talk about inequality: "90-10 share", "50-10" share, etc. Any ideas about alternatives?

2 Pareto Coefficients

The Pareto law is usually considered a good approximation of the upper tail of the distribution (e.g. top 10%).[2] Basically, an earlier economist, Vilfredo Pareto, documented that empirical fact that within countries, above some high income y, the mean income could be expressed as

[2] Pareto distributions have a probability density function $f(y) = \frac{ac^a}{y^{1+a}}$.

some scalar multiple of that income y: $\bar{y}_+ \approx y \cdot b = y \cdot \frac{a}{a-1}$.

The Pareto law is given by the following cumulative distribution function $F(y)$ for income y: $1 - F(y) = (c/y)^a$ where c is a constant and a is the Pareto coefficient.

Figure 3: Empirically, the Pareto coefficient a is very stable above $y^* = \$400,000$

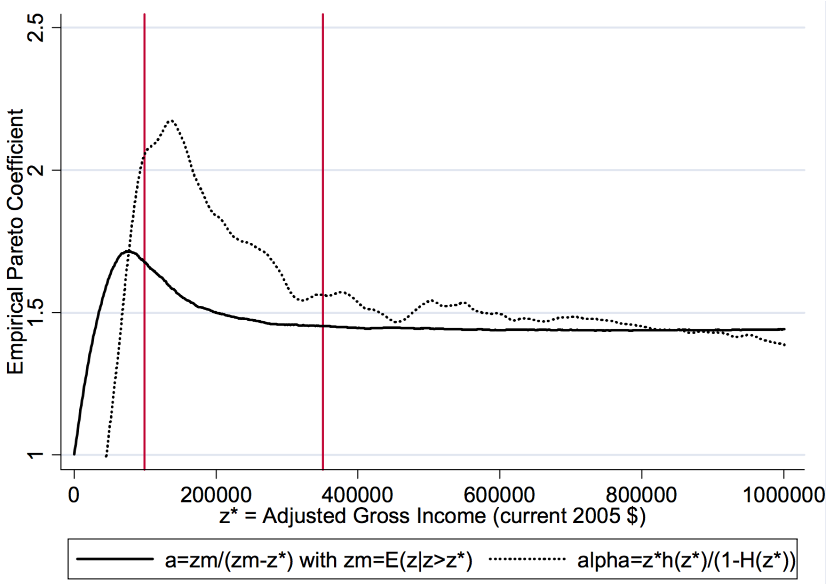

Source: Diamond and Saez, JEP 2011

More explicitly stated, the key property of Pareto distributions is that the ratio of average income $y^*(y)$ of tax units with income above y to y does not depend on the income threshold y: $y^*(y) = y\frac{a}{a-1}$. Therefore, $b = \frac{y^*(y)}{y} = \frac{a}{a-1}$ (the "inverted Pareto coefficient") measures how concentrated incomes are at the top and is arguably more intuitive than a:[3] for instance, a coefficient of $b = 2$ means that the average income above \$100,000 is equal to \$200,000, and the average income above \$1 million is equal to \$2 million, and so on. Similarly, a coefficient of $b = 3$ means that the average income above \$100,000 is equal to \$300,000, the average income above \$1 million is equal to \$3 million, etc. I.e., we might think of settings with greater inverted Pareto coefficients are more unequal.

Pareto laws provide a very useful statistical approximation technique to study the upper tails of the income and wealth distributions. In particular, tax data – which is often available in the form of tabulations reporting the numbers of taxpayers and the amounts of income for a certain number of tax brackets – can be easily used to estimate a and b within the top 10% or the top 1%.

Empirically, the standard Pareto coefficient a is 1.7 in the US and 3 in Denmark. (Does this mean incomes are more concentrated in Denmark or in the US?)

Important note: on a fundamental level, there is no rigorous proof that by force of nature top incomes need behave in this manner. There are plenty of theoretical works which make assumptions on how wealth and income can accumulate and demonstrate how these environments can reproduce Pareto distributions in the upper tails of resource distributions. However, for the most part, work that uses Pareto interpolation relies on the empirical observation that top income levels tend toward this well-behaved manner.

[3] Note that a and b are intrinsically related: $a = b/(b-1)$ and $b = a/(a-1)$.