Wage Structure and Inequality: First Lecture

LABOR ECONOMICS (ECON 385)

BEN VAN KAMMEN, PHD

Introduction

Recently the topic of inequality has been mentioned frequently and arguments over its severity and consequences have *occupied* a significant portion of our collective attention. The objective of this lecture is to convey the following.

- Several reasons that economists care about income inequality.
- How inequality can be measured.
- Why trends in inequality over time have occurred.

Terminology: poverty versus inequality

- Absolute Poverty: a situation where a population or section of a population is, at most, able to meet only its bare subsistence essentials of food, clothing, and shelter to maintain minimum levels of living.*
  - Absolute poverty is usually defined in terms of an official "line", e.g., the . . .
- International Poverty Line: An international real income measure, usually expressed in constant dollars (e.g., $1 per day), used as a basis for estimating the proportion of the world's population that exists at bare levels of subsistence.

*These definitions come from: Todaro, Michael, and Smith, Stephen. Economic Development, 10th ed. Boston: Pearson, 2009.

Relative poverty

- Absolute poverty contrasts the notion of relative poverty, which exists in any income distribution is not *perfectly* equal. For example, one may define the *relatively* poor as the subset of the population in the lowest 20% of the distribution (at least 80% of the population has higher incomes).
  - This 20% is not poor in the absolute sense, though, if they still reside above the absolute threshold, e.g., the international poverty line.
  - Similarly if the entire income distribution is below the absolute poverty line, it is possible for even the relatively rich subset to be absolutely poor.

Inequality: the variance and skewness of the income distribution

- When we study inequality, we ask questions like:
  - "how does the income at the 20th percentile compare to the median?" or
  - "how does the income at the 90th percentile compare to the median?",
  - taking the median as given (more on this in the measurement section).

- Studying the average and the degree of poverty is the focus of development economics.

Why do economists care about inequality (independent of poverty)?*

- Saving rates are lower when there is a lot of inequality.**
  - The middle of the income distribution has the highest "propensity to save", and if a small fraction of the population resides in this range (near the center) a larger fraction resides in the (lower saving rate) tails.
  - As you may already know, lower saving implies lower investment and slower income growth according to many models of macroeconomic growth.

*Again this borrows heavily from Todaro and Smith's Economic Development textbook.

**Mason, Andrew. 1988. "Savings, Economic Growth, and Demographic Change." *Population and Development Review*, Vol. 14: 113-144.

Why do economists care about inequality? (continued)

- High inequality undermines pro-growth public policy.
- Good government is under attack from both sides when there is high inequality.
  - The rich end of the distribution has disproportionate resources to spend seeking political favors and is inclined toward rent seeking.
  - And the poor end of the distribution is prone to "populist" political impulses that would forcibly redistribute incomes, e.g., some representatives of the "Occupy Wall Street" movement in the United States.
  - Both of these tendencies are bad for growth.

Why do economists care about inequality? (continued)

- Rent seeking—in the form of lobbying, corruption, and cronyism—is wasteful. The resources devoted to courting political favoritism could be reallocated to competitive, innovative, and productive purposes instead of toward stifling competition and raising prices.
- Redistribution of wealth eliminates the incentive to work and take risks to acquire it in the first place. This has adverse consequences for the economy as a whole; it slows the pace of technological progress and growth by removing the incentives to innovate.
- Todaro and Smith summarize the problem thusly:

  ". . . with high inequality, the focus of politics often tends to be on the redistribution of the existing economic pie rather than on policies to increase its size."

Why do economists care about inequality? (continued)

"The system under which people make their own choices—and bear most of the consequences of their decisions . . . produced overwhelming {fortunes} from developing new products or services, or new ways of producing products or services, or of distributing them widely. The resulting addition to the wealth of the community as a whole, to the well-being of the masses of the people, amounted to many times the wealth accumulated by the innovators. Henry Ford acquired a great fortune. The country acquired a cheap and reliable means of transportation and the techniques of mass production."

Friedman, Milton and Rose. Free to Choose. Harcourt: Orlando, 1990: pp. 138-39.

Why do economists care about inequality? (continued)

- Inequality is perceived as unfair. Substantial inequality is conferred at birth by the accident of one's nation of birth and parents. In light of this, many survey respondents state that they would prefer to be born into a world with less inequality than is observed in reality—assuming they did not know what part of the distribution they would be born into ahead of time.
- A related possibility that complicates matters: individuals' utilities are reference-dependent.
  - Sometimes called the "keeping up with the Joneses" characterization of subjective well-being.
  - In addition to the absolute level of one's standard of living, he also cares about where he ranks relative to his peers (a "reference group").
    - For example, even if he has a very comfortable existence, a person will still subjectively feel less well-off if all of his neighbors have larger homes and televisions.
- This seems to be the principle behind the quote in the front of Borjas's chapter 7: "What makes inequality such a difficult business is that we only want it with our superiors."

Why do economists care about inequality? (concluded)

- Perhaps economists' attitude toward inequality can be summarized as agreeing with most other observers that "less inequality is better" with the very predictable addition, *ceteris paribus*.
- Reducing inequality is desirable inasmuch as it does not harm the level and growth rate of income significantly.

How does one measure inequality?

- Ratios of two quantiles of the income distribution.
- The variance of individual incomes.
- The Lorenz curve—showing cumulative density of the income distribution.
- The Gini coefficient—indicating how an observed income distribution compares to a uniform distribution.

- Appendix about distributions.

Quantiles

- The incomes of the population can be divided into a finite number of “bins”.
  - One decides ahead of time how many (q) bins he wants to divide the observations into, a fixed proportion $(1/q)$ of the observations will fall into each bin.
  - The population is divided into *q* quantiles.

$q = 4 \rightarrow Quartiles; \frac{1}{4}$ of population in each quartile

$q = 5 \rightarrow Quintiles; \frac{1}{5}$ of population in each quintile

$q = 10 \rightarrow Deciles; \frac{1}{10}$ of population in each decile

$q = 20 \rightarrow Ventiles; \frac{1}{20}$ of population in each ventile

$q = 100 \rightarrow Percentiles; \frac{1}{100}$ of population in each percentile

Quantiles (continued)

- The ratio of the 90th percentile to the 10th percentile is a statistic that can be used to measure inequality.
  - So is the 75-25 ratio, but the 90-10 wage gap and the 50-10 wage gap are popular examples of easy-to-compute measures of inequality.
  - The 90-10 gap tells you the percentage wage difference between someone in the upper tail (90) of the distribution and someone in the lower tail (10). The greater is this ratio, the more unequal the incomes are.
- A single observation of the 90-10 wage gap may be revealing, but it doesn't necessarily mean much without context.
  - It's generally more interesting to compare two countries in terms of inequality or compare the same country to itself in an earlier time period.
  - If the gap rises over time, you can plausibly say that income has become more unequal in the intervening years.

Variance and standard deviation

- Every statistician's favorite measure of variability is one that uses all observations in the population—the variance, denoted σ^2. Variance basically relies on the deviations from the mean:

$$income_{person\ i} - (mean\ income) \equiv x_i - \mu$$

- A distribution with more variability (inequality) will have more people with deviations that are large in absolute value.

Variance and standard deviation (continued)

- Instead of literally using the absolute value, variance uses an operation that has the isomorphic consequence of making all deviations positive. It squares all the deviations from the mean, i.e.,

$$(x_i - \mu)^2$$

- Finally the squared deviations from the mean are added over all people in the population and then divided by the population size, N.

$$\sigma^2 \equiv \frac{1}{N}\sum_{i=1}^{N}(x_i - \mu)^2$$

Again a higher variance reflects a population in which many people have incomes either very far below or very far above the mean, i.e., an unequal distribution. So higher variance reflects more inequality.

The Lorenz curve

- Once the population is divided into quantiles, each quantile's share of the total can be calculated. These shares can then be used to approximate the cumulative density.
  - For example, say that income was distributed as follows:

Lowest Quintile: 3.6%

2nd Quintile: 8.9%

3rd Quintile: 15.0%

4th Quintile: 23.4%

Top Quintile: 49.1%

- The cumulative density (proportion making less than or equal to my level) looks like:

Lowest Quintile: 3.6%

2nd Quintile: (8.9+3.6)=12.5%

3rd Quintile: (15.0+12.5)=27.5%

4th Quintile: (23.4+27.5)=50.9%

Top Quintile: (49.1+50.9)=100%

The Lorenz curve: plotting the cumulative shares

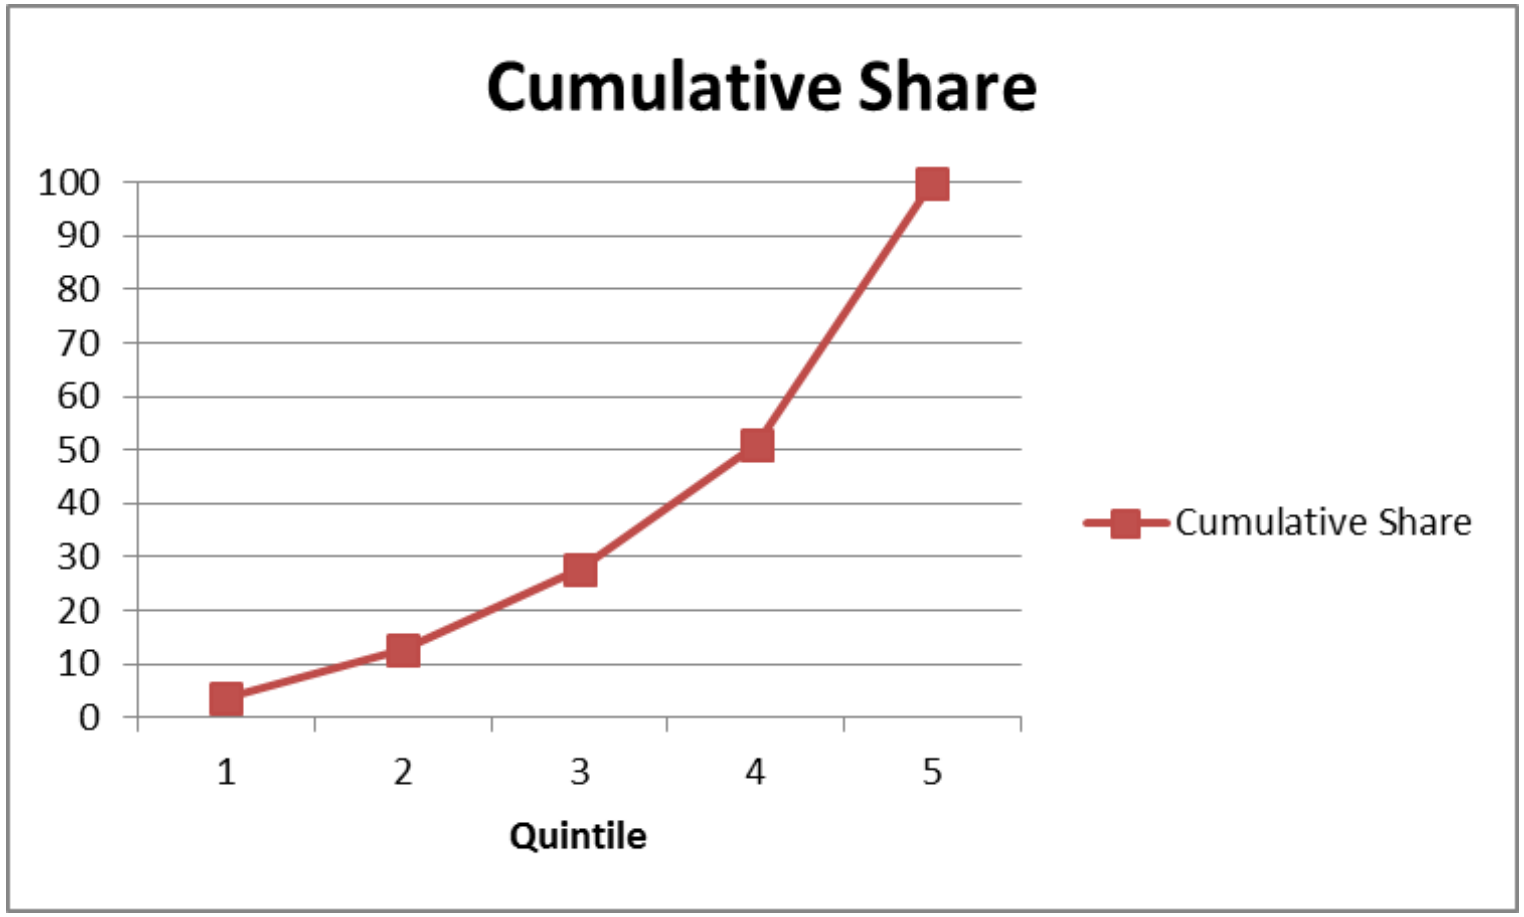

The Lorenz curve (continued)

- Again this is a fun curve to look at in isolation, but it doesn't tell the reader much without context.
- It is informative to compare two distributions according to how unequal they are.
  - To this end, consider an alternative distribution where, say, 10% of the income of the top quintile's income is redistributed (costlessly) to the other 4 quintiles equally.
  - Ignore the possibility that the redistribution could change the composition of the quintiles by moving some tops down into the 4th and some 4ths into the top.
- The new distribution would be:

Lowest Quintile: 4.828%

2nd Quintile: 10.128%

3rd Quintile: 16.228%

4th Quintile: 24.628%

Top Quintile: 44.19%

The Lorenz curve (continued)

- The new distribution would be:

  Lowest Quintile: 4.828%

  2nd Quintile: 10.128%

  3rd Quintile: 16.228%

  4th Quintile: 24.628%

  Top Quintile: 44.19%

- The new cumulative density would be:

  Lowest Quintile: 4.828%

  2nd Quintile: (10.128+4.828)=14.955%

  3rd Quintile: (14.955+16.228)=31.1825%

  4th Quintile: (24.638+31.1825)=55.81%

  Top Quintile: (44.19+55.81)=100%

New Lorenz curve

- The blue line that lies above the red line:

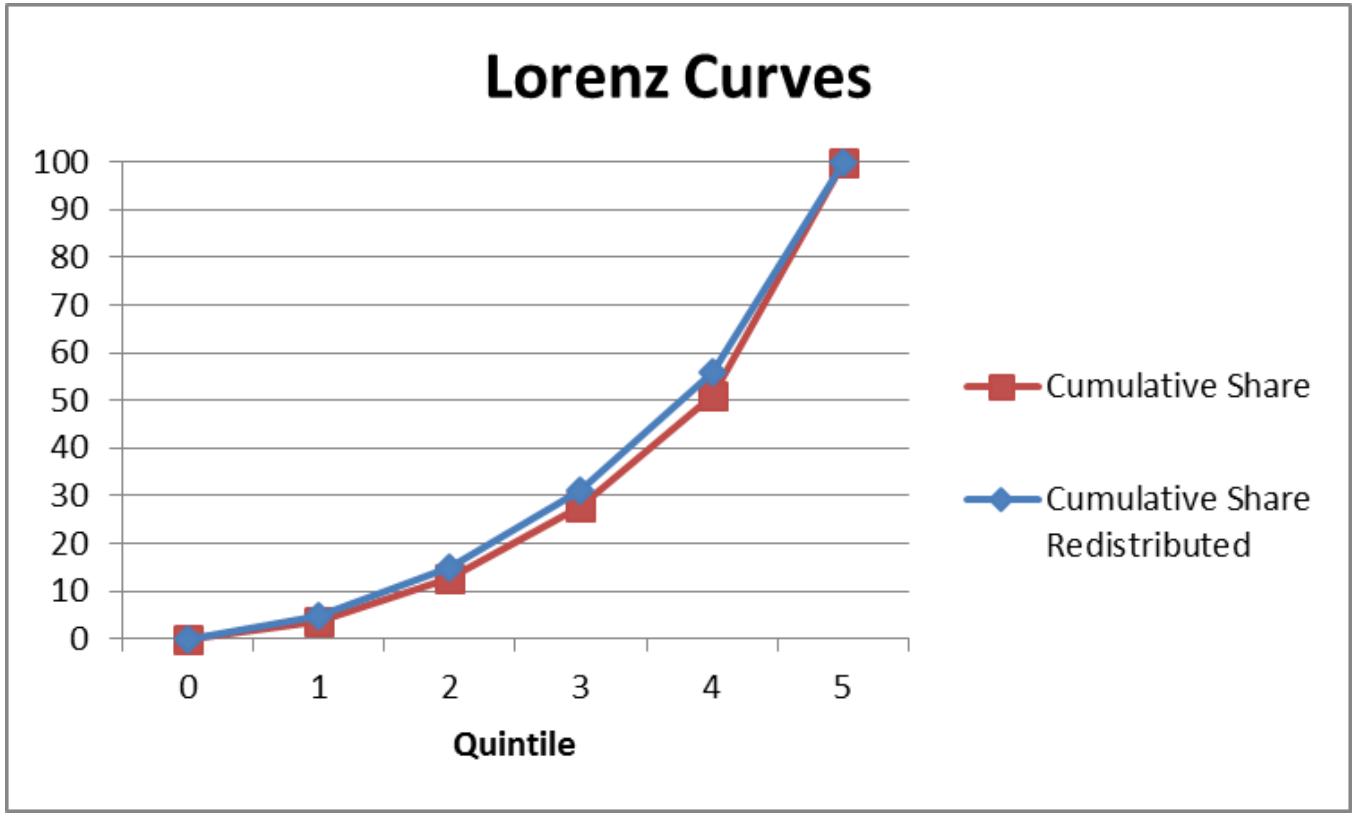

The Lorenz curve and inequality

- If you continue redistributing income from higher quantiles to lower ones, the Lorenz curve will keep "bowing" less and less until income is uniformly distributed. If the distribution looks like this,

  Lowest Quintile: 20%

  2nd Quintile: 20%

  3rd Quintile: 20%

  4th Quintile: 20%

  Top Quintile: 20%

  the Lorenz curve will look like the green line on the following graph.

Perfect equality Lorenz curve

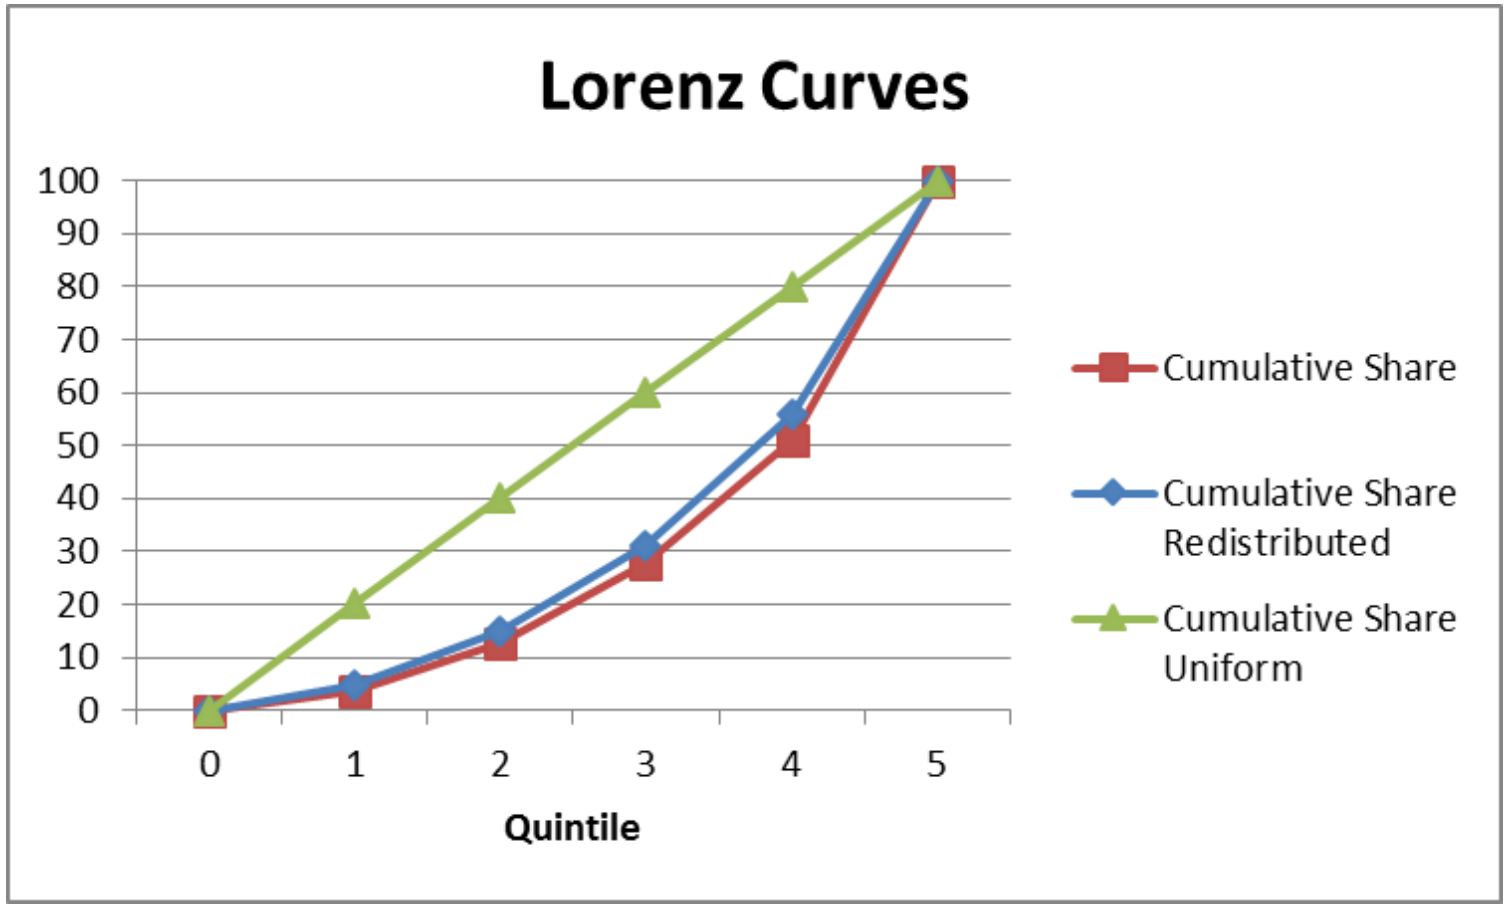

Highly unequal Lorenz curve

- Conversely if the top quintile has almost all of the income, the Lorenz curve will look like the purple line on the graph.

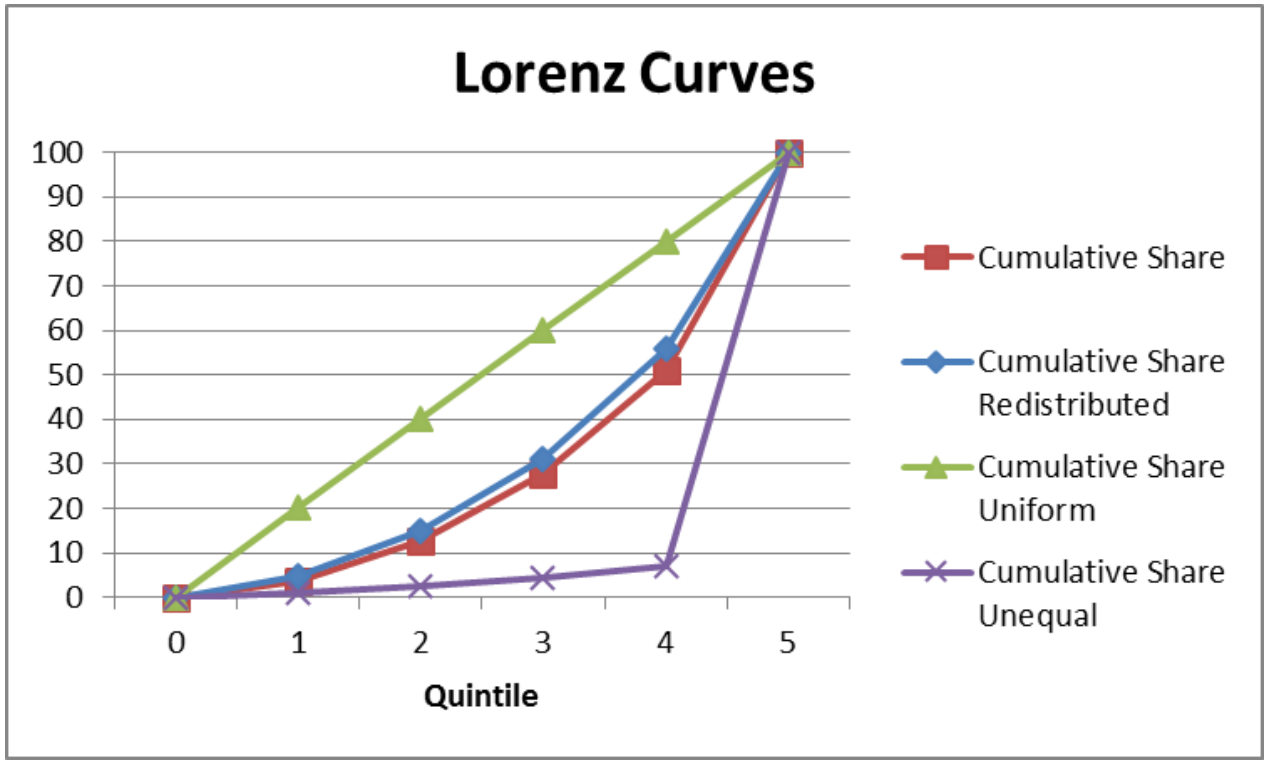

The Gini coefficient

- As the preceding examples illustrate, the more unequal a distribution is, the further below from the green line the Lorenz curve will lie.
- This immediately suggests another way of measuring inequality: the area between the actual Lorenz curve and the uniform distribution Lorenz curve. The orange shaded area on the next graph.

The Gini coefficient (continued)

- The larger this area is, the more unequal the distribution will be. This area is the basis for the last of our four measures of inequality: the Gini coefficient. This is the ratio of the orange area to the entire area below the green curve.

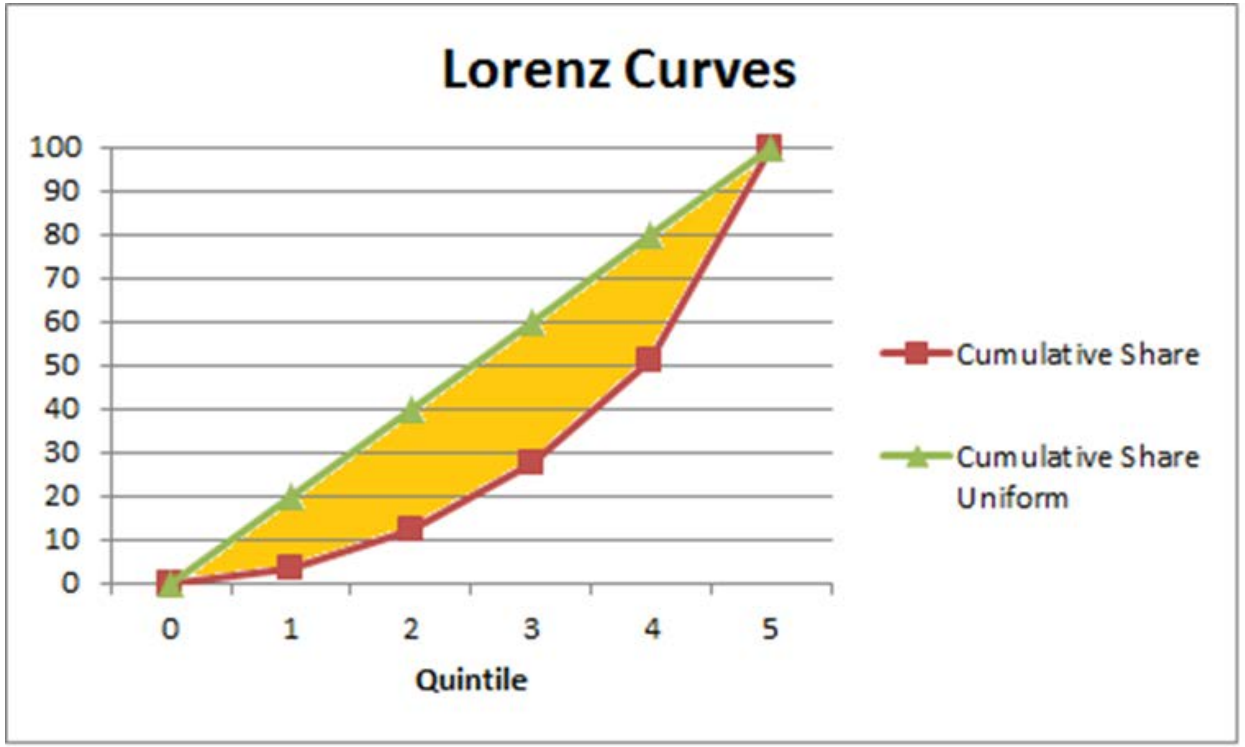

The Gini coefficient (continued)

- More inequality moves the Gini coefficient closer to 1.
  - If there is perfect equality, the area between the Lorenz curve and the green line will be 0.
  - Also if a single individual has all the income, the Lorenz curve will approach a right angle, and the area between will approach the area under the green curve—a ratio of 1:1.
- For the Lorenz curves in the preceding example, the Gini coefficients are respectively:

$$Gini = \begin{cases} 0.422\ (red; \text{actual statistics from U.S. Census 1994}) \\ 0.3729\ (\text{blue; with costless redistribution}) \\ 0.74\ (\text{purple; hypothetical high inequality}) \end{cases}$$

The Gini coefficient (continued)

- There is a way of approximating the Gini coefficient for these examples that requires only algebra, i.e., the area formulae for triangles and rectangles. The formula is as follows:

$$Gini = \frac{1}{50 * (number\ of\ quantiles)} * \left\{ 50q - \sum_{quantiles=i}^{q} \left[LC_{i-1} + \frac{1}{2}(LC_i - LC_{i-1}) \right] \right\}$$

$$\Leftrightarrow Gini = \left[1 - \frac{1}{100q} * \sum_{quantiles=i}^{q} (LC_i + LC_{i-1}) \right],$$

where “LC” stands for the height of the Lorenz curve for quintile, *i*.

Conclusion

- Inequality and poverty are not synonymous.
  - Inequality is reflected in the variance and skewness of an income distribution.
- There are various ways of measuring inequality.
  - Some are easier but less informative.
  - Some are more difficult but more informative.
- These measures enable comparisons across countries and over time in the level of income inequality.

Economics of poverty and development

- We do not concern ourselves with absolute poverty here because there is a separate subject in economics devoted to this (and related) subjects: development economics (also studied in international economics and macroeconomics classes).
- Nor are we concerned with the "first moment" of the income distribution (the mean or median which measures its location)—except inasmuch as it is influenced by the "second moment" (the variance which measures how spread out it is) and the "third moment" (the skewness which measures how symmetrical it is).
- Again, the evolution of the average income (think GDP per capita) falls within the province of growth models in macroeconomics and development. Here we focus on inequality *at a given average income* level—reflected in the dispersion and skewness of the income distribution.

Back.

The *distribution* concept

- This refers to the manner in which the observations (individual or household incomes) are *distributed* along a number line.
  - The numbers on that line indicate income levels.
  - The observations are usually represented by a bar or a plot (see figure 7-1 in the textbook) that tells you what proportion of the population earns that income level.
- Income is basically a continuous variable; it can take on any fraction of any non-negative number—infinitely many different values.

The *distribution* concept (continued)

- Many graphical representations of income distributions will divide the infinitely many possible values into “bins” of equal size.
  - Every observation that falls into a given bin, then, is counted and the frequency in each bin is plotted on the graph. This is how you end up with graphs with a finite number of plots like figure 7-1.
- When instructors are in a hurry, they might just draw a curve on the chalk board to approximate this kind of graph—instead of plotting a finite series of points and connecting them with a line.
  - When they illustrate distributions in this way, they are attempting to make the same picture that the software programs do using data.
- The height of the distribution is referred to as the density or the pdf (probability density function); it means “the probability that income takes the value, x”.

A pdf, illustrated

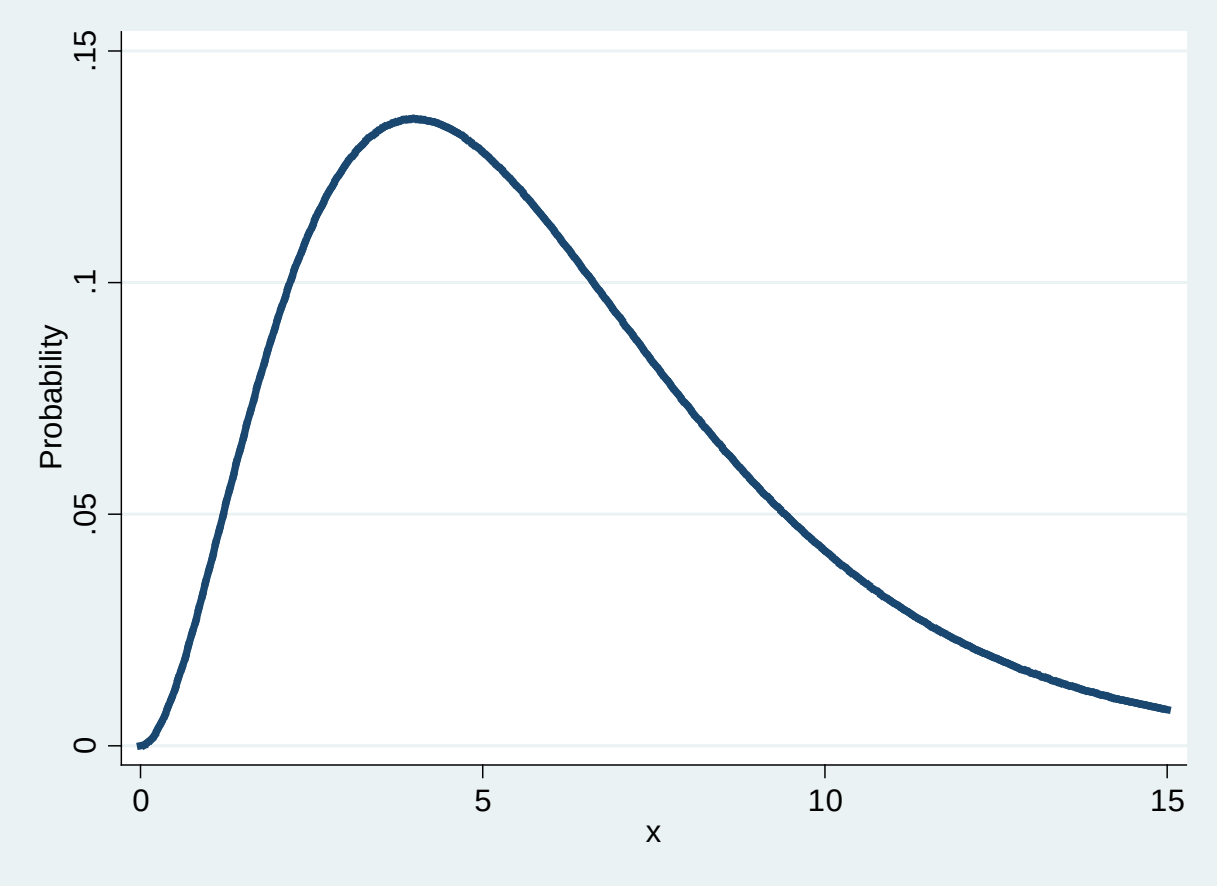

The *distribution* concept (concluded)

- The sum of the densities equals 1. This is represented graphically as the area under the distribution.
- The area under the distribution to the left of a certain point is called the cumulative density or the cdf. It means "the probability that income takes a value less than or equal to x".
- In most distributions, the most common values are in the middle, leaving two "tails" on each end. Discussion will refer to these as the upper and lower tails of the distribution. A highly skewed income distribution will be assessed to have a "long upper tail".

A cdf, illustrated

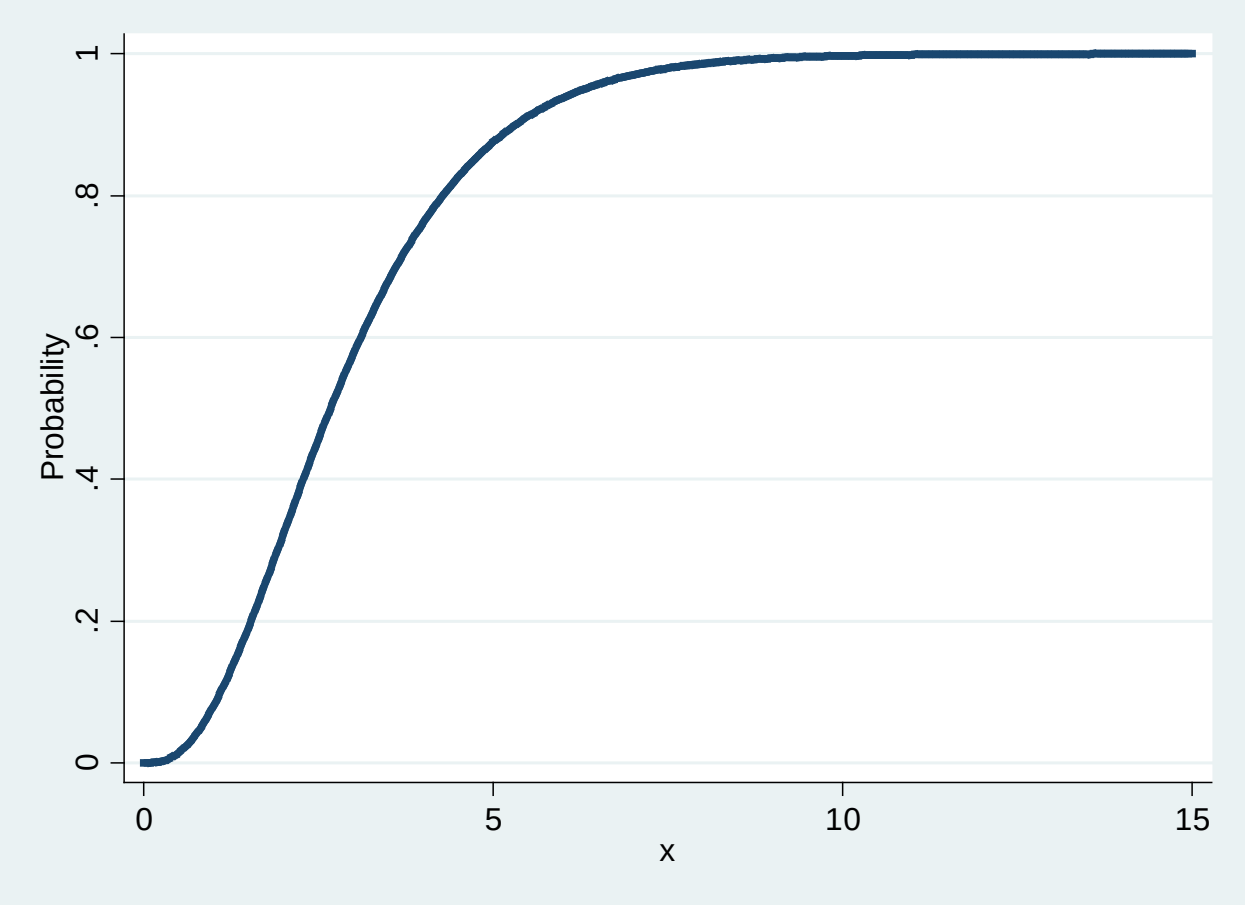

Back.