Measuring inequality

LECTURE 13

C4D2 TRAINING

1

Outline for final lectures

- Once datasets have been finalized, it is time to produce results, with the aim of representing the patterns emerging from the data.
- In practice?
- Inequality — this lecture
- Poverty — next lecture
- Basic summary statistics on household demographics, education, access to services, etc.
- Average expenditures and incomes — final lecture

C4D2 TRAINING

2

Inequality and poverty measurement

1) a measure of living standards
2) high-quality data on households' living standards
3) a distribution of living standards (inequality)
4) a critical level (a poverty line) below which individuals are classified as "poor"
5) one or more poverty measures

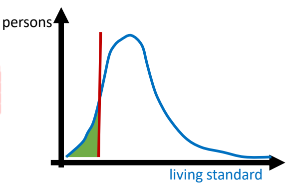

C4D2 TRAINING

3

Cowell (2011)

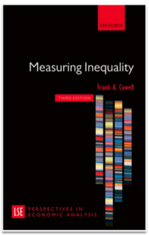

99.9% of this lecture is explained with better words in Cowell's work: this book and other (countless) journal articles

Warning

- During the course we paid attention to distinguish between different concepts: living standard, income, expenditure, consumption, etc.
- In this lecture we make an exception and use these terms interchangeably – we focus on measuring inequality of "a distribution"
- Similarly, I will not make a distinction between income per household, per capita, or per adult equivalent
- For once, and for today only, we will be (occasionally) inconsistent

Basic concepts

- Economists make a distinction between:
  - Functional distribution of income
    distribution among factors of production
    land (rent), labor (wages), and capital (profits)
  - Personal (or size) distribution of income
    distribution among persons, irrespective of their economic function
- We focus on the latter.

Functional vs Personal distribution of income

Average factor shares in Indian Economy, 1960-61 to 1991-1992

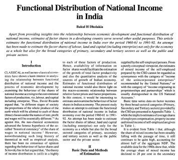
Functional Distribution of National Income in India

Bakul H Dholakia

Apart from providing insights into the relationship between economic development and functional distribution of national income, estimates of factor shares in a developing country serve several other useful purposes. This article estimates the functional distribution of national income in India over the period 1960-61 to 1991-92. An attempt has been made to estimate the factor shares of labour, land and capital (including enterprise) not only for the economy as a whole but also for the broad categories of primary, secondary and tertiary sectors as well as the public and private sectors.

Sector	Labour	Land	Capital
Primary sector	56.42	30.30	13.28
Secondary sector	67.68	3.47	28.85
Tertiary sector	61.57	3.74	34.69
All sectors	60.69	15.21	24.10
Public sector	86.15	0.83	13.02
Private sector	56.53	17.76	25.71

Dholakia (1996), Functional Distribution of National Income in India, Economic and Political Weekly

Focus on the term 'inequality'

- "When we say income inequality, we mean simply differences in income, without regard to their desirability as a system of reward or undesirability as a scheme running counter to some ideal of equality" (Kuznets 1953: xxvii)
- In practice, how can we appraise the inequality of a given income distribution? Three main options:
  1. Tables
  2. Graphs
  3. Summary statistics

The distribution of income and taxable income

2018 Tax statistics, National Treasury and the South African Revenue Service

Tax year	2017			
Income group	Number of taxpayers	Income before deductions (R million)	Deductions allowed (R million)	Taxable income (R million)
<= 0	116 998	-15 096	11	-15 107
1 – 70 000	401 447	14 835	565	14 270
70 001 – 350 000	2 689 263	543 389	58 380	485 009
350 001 – 500 000	764 197	317 965	44 483	273 483
500 000 +	926 660	898 846	109 690	789 157
Total	**4 898 565**	**1 759 939**	**213 128**	**1 546 811**

Tables: an assessment

- In general, tables are not recommended when the focus is inequality
- Difficult to get a clue of the extent of inequality in the distribution by looking at a table, plus income brackets are arbitrary
- Does putting income distribution into a graph (diagram) help to represent inequality?

Histograms

- Let the interval $[x^-, x^+]$ denote the range of the data.
- Partition $[x^-, x^+]$ into m^* non-overlapping bins (intervals) of equal width $\mathrm{h} = (x^+ - x^-)/m^*$.
- A histogram estimate of the density $f(x)$ is the fraction of observations falling in the bin containing x, divided by the bin width h:

$$\hat{f}(x) = \frac{\text{(fraction of sample obs. in same bin as x)}}{h}$$

- The area of each bar (= $h \times \hat{f}(x)$) is interpreted as the fraction of sample observations within the bin. All bar areas sum up to unity.

Histograms?

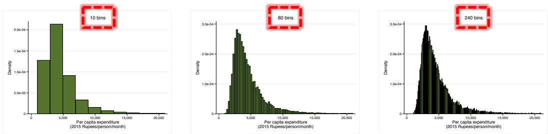

C4D2 TRAINING

Histograms: an assessment

- The position and number of bins is arbitrary
- Inherently lumpy: discontinuities at the edge of each bin
- Can provide very different pictures of the same distribution
- Read Cowell, Jenkins and Litchfield (1996) for more.

C4D2 TRAINING

Beyond histograms

- A probability density function (PDF) is the 'continuous version' of a histogram
- A convenient way to introduce the PDF is by starting from the cumulative distribution function (CDF)

C4D2 TRAINING 15

The cumulative distribution function (CDF)

- The cumulative distribution function (CDF) is defined as follows:

$$F(x) = \int_0^x f(x)dX$$

- $F(x)$ is the proportion of individuals having X less than or equal to x.
- If X is income and, say, $x = 2{,}000$ Rps., then $F(x) = \Pr(X < 2{,}000)$, that is the fraction of people with less than 2,000 Rps.

C4D2 TRAINING

The empirical cumulative distribution function (ecdf)

Mongolia HSES 2016, Cumulative distribution of per capita consumption (p.10)

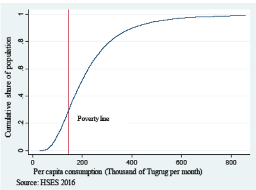

- Pick up any income level on the **x**-axis, and the curve **F**(**x**) will tell you the percentage of individuals in the population having a level of income <u>lower than</u> **x**.

C4D2 TRAINING

The probability density function (pdf)

- The probability distribution function (pdf) is the derivative of the CDF:

$$f(x) = \frac{dF(x)}{dx}$$

- By definition of derivative:

$$f(x) = \lim_{h \to 0} \frac{F(x+h) - F(x)}{h}$$

C4D2 TRAINING

The probability density function (pdf)

- Now drop the limit (and replace = by ≈):

$$f(x) \approx \frac{F(x+h) - F(x)}{h}$$

$$h.f(x) \approx F(x+h) - F(x)$$

$$h.f(x) \approx \Pr(X \leq x+h) - \Pr(X \leq x) \approx \Pr(x \leq X \leq x+h)$$

- The PDF f(x) is not a probability measure, but a scaled version of it: it is the probability of X falling in the interval $(x, x+h)$ divided by the length h of such an interval.

Mongolia, 2016

PDF of per capita consumption (p.11)

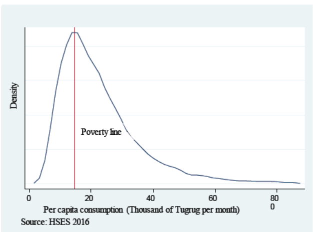

Interpretation

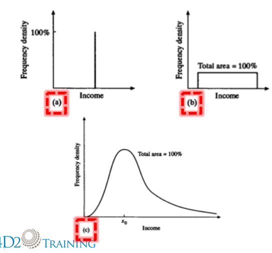

Two extremes:

a) perfect equality: everyone is concentrated at one particular income value

b) uniform density: income is spread uniformly from the poorest to the richest individual – significant inequality

c) in-between, typical case.

Pdfs: an assessment

- The bandwidth is arbitrary
- In most cases, they require some trimming of top values to avoid looking "squished" and being unreadable
- In general, it does not show what is going on in the upper tail very clearly

C4D2 TRAINING

The quantile function

- Let p = F(x) be the proportion of people in the population with income lower than x.
- The quantile function Q(p) is defined as:

$$F[Q(p)] = p \quad \text{or} \quad Q(p) = F^{-1}(p)$$

cdf — inverse cdf

- Q(p) is the income level below which we find a proportion p of the population.

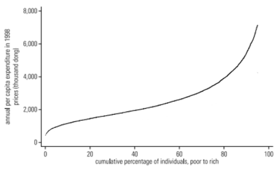

Source: Haughton and Khandker (2009).
Pen's Parade (Quantile Function) for Expenditure per Capita, Vietnam, 1998

C4D2 TRAINING

*The Parade of Dwarfs

Pen (1971)

- Assume that everyone in the population has height proportional to income.
- Line people up in order of height, and let them march.
- After some time, the shape of such a parade will be represented by the curve called Parade of Dwarfs (and a Few Giants).

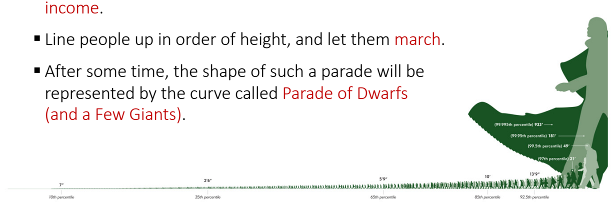

C4D2 TRAINING
Graphic from "The Atlantic"

The Lorenz Curve (1905)

Picture and intuition

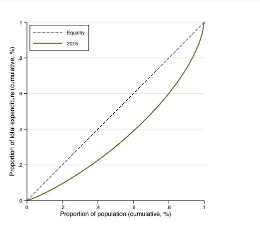

- Horizontal axis: cumulative % of population (individuals ordered poorest to the richest)
- Vertical axis: cumulative % of income received by each cumulative % of population.
- 45-degree line: Lorenz curve if perfect equality.
- The overall distance between the 45-degree line and the Lorenz curve is indicative of the amount of inequality present in the population.

C4D2 TRAINING

The Lorenz Curve (1905)

Mathematically

- The Lorenz curve L(p) is defined as follows:

$$L(p) = \frac{\int_0^p Q(p)\,dq}{\int_0^1 Q(p)\,dq}$$

- The numerator sums the incomes of the poorest p% of the population;
- The denominator sums the incomes of all.
- The ratio L(p) indicates the cumulative % of total income held by a cumulative proportion p of the population.
- Example: if L(0.5) = 0.3, then we know that the 50% poorest individuals hold 30% of the total income in the population.

C4D2 TRAINING

Quantile function and Lorenz curve: an assessment

- These graphical tools emphasize the ranking of shares of the population on the basis of income
- The Lorenz curve clearly shows how far the distribution is from perfect equality
- Still, no graph is as straightforward and easily comparable as a scalar measure of inequality

C4D2 TRAINING 27

Recap and next steps

- Not all graphs are OK to represent inequality
- Lorenz curve is the most popular
- A better conceptual understanding comes from constructing inequality measures from first principles.
- The most straightforward approach: inequality measures as pure statistical measures of dispersion.

C4D2 TRAINING

Inequality indicators

C4D2 TRAINING

Measures of dispersion

- Range $R = x_{max} - x_{min}$

▲PRO: Easy to compute and communicate

▼CON: Insensitive to changes between extremes (can we really know min and max?)

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

▲ PRO: Easy to compute, additively decomposable

▼ CON: not robust (outliers), depends on the scale of measurement

- Coefficient of Variation $CV = \frac{\sqrt{\sigma^2}}{\mu}$

▲PRO: Scale invariant

▼CON: not robust (outliers), properties?

C4D2 TRAINING

Quantiles, Quintiles, Quartiles, …

- The p-quantile of a distribution of values is a number x_p such that a proportion p of the population values are less than or equal to x_p.
- For example, if p = 0.5, then the 0.5-quantile $x_{0.5}$ is any value such that $F(X < x_{0.5}) = 0.5$.
- Certain quantiles have special names:
  - The 0.5-quantile $x_{0.5}$ is the median, or 50-th percentile.
  - The 0.1-quantile is the first decile, or 10-th percentile.
  - The 0.2-quantile is the first quintile, or 20-th percentile.
  - The 0.25-quantile is the first quartile Q_1, or 25-th percentile.
  - etc. etc.

C4D2 TRAINING

Quantile ratios

- A quantile ratio measures the gap between the rich and the poor.
- It is defined as the ratio of two quantiles, Q(p2)=Q(p1) using percentiles p1 and p2.
- Three popular indices are:
  - The quintile ratio (p2 = 80 and p1=20):

    QR = Q(p80)/Q(p20)
  - the decile ratio (p2 = 90 and p1=10):

    DR = Q(p90)/Q(p10)

C4D2 TRAINING

The decile ratio

per adult equivalent income (OECD def.)

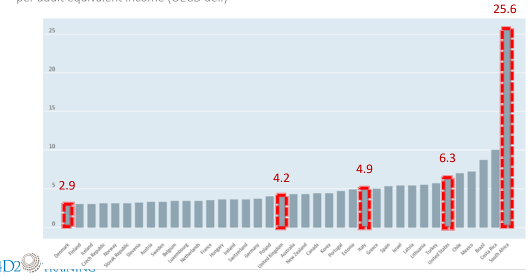

C4D2 TRAINING

Quantile share ratios

- Let S_{20} denote the share of (equivalised disposable) income received by the bottom 20% of the population, and S_{80} the income share received by the top 20% of the population.
- The quintile share ratio is defined as follows:
  $S_{80\text{-}20} = S_{80}/S_{20}$
- The quintile share ratio is the level-1 Laeken indicator, chosen by the EU to monitor income distribution.

C4D2 TRAINING

QSR around the world

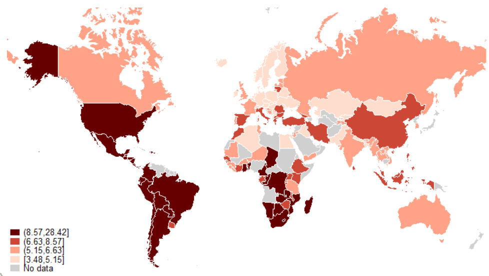

Source: (WDI) Income share held by highest 20% over Income share held by lowest 20%, last available 2010-2017

C4D2 TRAINING

S_{80}/S_{20} in Sub Saharan Africa

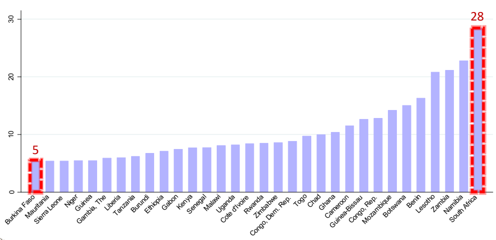

Source: (WDI) Income share held by highest 20% over Income share held by lowest 20%, last available 2010-2017

C4D2 TRAINING

The Gini Coefficient

A definition

- Yitzhaki (1997) counts more than a dozen formulas available for the Gini index.
- A classic definition of the Gini coefficient:

$$G = \frac{1}{2n^2\mu}\sum_{i=1}^{n}\sum_{j=1}^{n}\left|x_i - x_j\right|$$

- The Gini coefficient ranges from 0 (all recipients have the same income: full equality), to 100 (all income is received by one recipient: maximum inequality).

C4D2 TRAINING

The Gini Coefficient

Interpretation – Pyatt 1976: 244

- The Gini index "is the average gain to be expected, if each individual has the choice of being himself or some other member of the population drawn at random, expressed as a proportion of the average level of income"
- *E.g.*, if the Gini index for an Italian is 0.30, we can say that the expected gain from playing the experiment of exchanging income with someone else randomly chosen in the Italian population, is 30% of average income.

C4D2 TRAINING

The Gini Coefficient

graphical interpretation

- The Gini index is two times the area A between the Lorenz curve and the equality diagonal:

$$Gini = \frac{A}{(A+B)}$$
$$= 2A$$
$$= 2\left(\frac{1}{2} - B\right) = 1 - 2B$$

$$Gini = 2\int_0^1 [p - L(p)]dp$$
$$= 1 - 2\int_0^1 L(p)dp$$

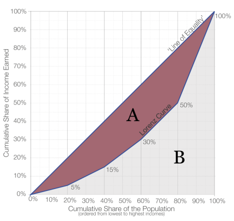

C4D2 TRAINING

Gini index around the world

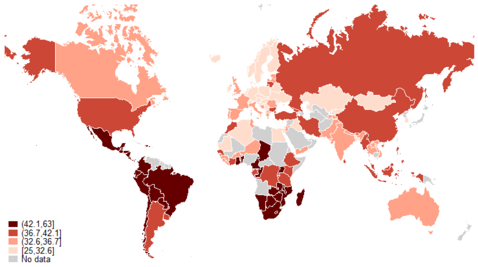

The Gini index in Sub-Saharan Africa

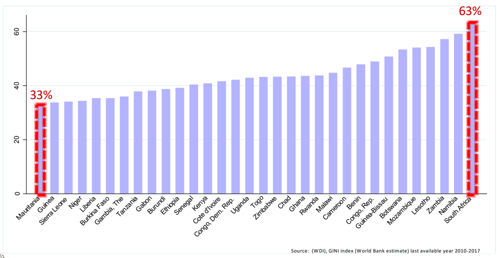

Gini around the World

A selection of countries

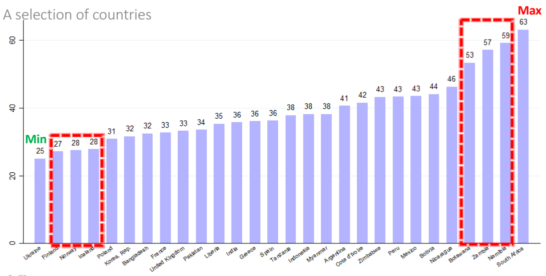

*Atkinson's paper

The paper

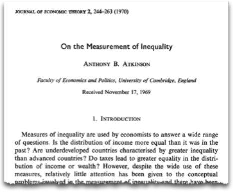
JOURNAL OF ECONOMIC THEORY 2, 244–263 (1970)

On the Measurement of Inequality

ANTHONY B. ATKINSON

Faculty of Economics and Politics, University of Cambridge, England

Received November 17, 1969

1. INTRODUCTION

Measures of inequality are used by economists to answer a wide range of questions. Is the distribution of income more equal than it was in the past? Are underdeveloped countries characterised by greater inequality than advanced countries? Do taxes lead to greater equality in the distribution of income or wealth? However, despite the wide use of these measures, relatively little attention has been given to the conceptual problems involved in the measurement of inequality and there have been

Tony Atkinson (1944-2017)

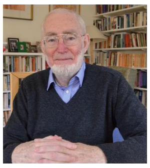

C4D2 TRAINING

Recap

- Quantile ratios, quantile share ratios, Gini, are all popular inequality measures
- They do a fine job at representing inequality with a number
- Problem
  they do not always have all the properties that we would want for an inequality measure
- Solution
  solve the problem backwards. First lay out some desirable properties, then construct a measure that complies with them

C4D2 TRAINING 44

Deriving inequality measures from axioms

- Axiom: a statement accepted as true as the basis for argument or inference.
- The axiomatic approach allows us to "custom-build" inequality measures that fit our needs:
  1. We define a set of elementary properties (axioms) that we think inequality measures ought to have
  2. We obtain a mathematical formula that delivers a class of inequality measures satisfying the axioms

C4D2 TRAINING

Five axioms of inequality measures

A. Anonymity (or Symmetry)
   Who is earning the income does not matter
B. The Population Principle
   Population size does not matter
C. Scale Invariance (or Relative Income Principle)
   Income *levels* do not matter
D. The (Pigou-Dalton) Principle of Transfers
   Rank-preserving rich-to-poor transfers reduce inequality
E. Decomposability (or Subgroup Consistency)
   The measure is additively decomposable

C4D2 TRAINING 46

*Five axioms of inequality measures

(A) Anonymity (or Symmetry)

- If income distribution X is any permutation of income distribution Y, then I (X) = I (Y).
- In short, it does not matter who is earning the income.

(P) The Population Principle

- When one income distribution is an n-fold replication of another, the two are distributionally equivalent.
- The population size does not matter: all that matters are the proportions of the population who earn different levels of income.

C4D2 TRAINING

*Five axioms of inequality measures

(S) Scale Invariance (or Relative Income Principle)

- If everyone's income changes by the same proportion, then inequality does not change.
  - $X = (x_1, x_2, \ldots, x_n)$
  - $Y = (\lambda x_1, \lambda x_2, \ldots, \lambda x_n)$
  - $I(X) = I(Y)$
- Inequality should not depend on whether income is measured in PKR or €.
- Income levels, in and of themselves, have no meaning as far as inequality measurement is concerned.

C4D2 TRAINING

*Five axioms of inequality measures

(T) The (Pigou-Dalton) Principle of Transfers

- If one distribution is obtained from another by transferring a positive amount of income δ from a relatively rich person to a relatively poor person, without altering their ranks in the distribution, then inequality must decrease.
- $X = (x_1, x_i, \ldots, x_j, \ldots, x_n)$
- $Y = (x_1, x_i + \delta, \ldots, x_j - \delta, \ldots, x_n)$, with $\delta > 0$
- $I(Y) \leq I(X)$

C4D2 TRAINING

*Five axioms of inequality measures

(D) Decomposability (or Subgroup Consistency)

- An **additively decomposable** inequality measure is one which can be expressed as a **weighted sum** of the inequality values calculated for population groups **plus** the contribution of differences between group means.

$$I = \sum_{k=1}^{K} \omega_k I_k + I(\bar{x}_1, \ldots, \bar{x}_k), \qquad \sum_{k=1}^{K} \omega_k = 1$$

where I_k is the inequality index calculated within the k-th group, and ω_k are the population shares.

C4D2 TRAINING

Generalized Entropy Indices (GEI)

Shorrocks (1980)

- Inequality measures that satisfy all axioms (A to E), must have the following form:

$$GE(\theta) = \frac{1}{\theta^2 - \theta}\left[\frac{1}{n}\sum_{i=1}^{n}\left(\frac{x_i}{\bar{x}}\right)^{\theta} - 1\right]$$

where θ is a parameter that may be given any value (positive, zero or negative).

C4D2 TRAINING

The Generalized Entropy Indices

- Depending on the value of the θ parameter:
  - $\theta = 0$ → Mean Logarithmic Deviation

    $$GE(0) = MLD = \frac{1}{n}\sum_{i=1}^{n} \log\left(\frac{\bar{x}}{x_i}\right)$$

  - $\theta = 1$ → Theil Index

    $$GE(1) = THEIL = \frac{1}{n}\sum_{i=1}^{n} \frac{x_i}{\bar{x}} \log\left(\frac{x_i}{\bar{x}}\right)$$

  - $\theta = 2$ → Half Coefficient of Variation Squared

    $$GE(2) = \frac{\kappa^2}{2}$$

C4D2 TRAINING

Uganda

National Household Survey 2012/2013

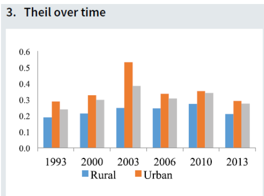

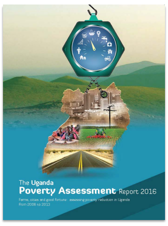

C4D2 TRAINING

53

Inequality decomposition

- Inequality decompositions are typically used to estimate the extent to which the heterogeneity of the population affects overall inequality. Two popular techniques are:
  1. Decomposition by population sub-group
  2. Decomposition by income source
- We focus on the former:
  - Societies can often be partitioned into groups (e.g. North-South). We would like to be able to decompose total inequality into two components, namely the inequality within the constituent groups, and inequality between the groups:

$$I_{TOTAL} = I_{WITHIN} + I_{BETWEEN}$$

C4D2 TRAINING

Inequality decomposition

- The most popular additively decomposable inequality index is the Mean Logarithmic Deviation.
- Partition the population into $k = 1, \dots, K$ groups. Then:

$$MLD = \underbrace{\sum_{k=1}^{K} v_k MLD_k}_{WITHIN} + \underbrace{\sum_{k=1}^{K} v_k \log\left(\frac{x}{x_k}\right)}_{BETWEEN}$$

where v_k are population shares.

Botswana, 2009/10

household income and expenditure survey

	2010		
	GE(0)	GE(1)	GE(2)
Total	0.669	0.823	3.206
Urban / rural			
Between-group inequality	0.063	0.068	0.077
Between as a share of total	0.094	0.083	0.024
Within-group inequality	0.606	0.754	3.130
Region			
Between-group inequality	0.060	0.068	0.081
Between as a share of total	0.090	0.083	0.025
Within-group inequality	0.609	0.754	3.125

Lessons learned

- Many ways to describe inequality, some more effective than others
- Graphs: most notable are quantile functions and Lorenz curves
- Measures: different inequality measures lead to different results. Based on their properties, the recommended choice is GEI (generalized entropy indices), and in particular the MLD (mean log deviation). However, Gini remains extremely popular in practice

References

Required readings

Cowell, F. (2011). Measuring inequality. Oxford University Press. (Chapter 1 & 2)

Suggested readings

Atkinson, A. B. (1970). On the measurement of inequality. Journal of economic theory, 2(3), 244-263.

Cowell, F., Jenkins, S. P., and Litchfield, J. A. (1996). The changing shape of the UK income distribution: kernel density estimates (pp. 49-75). Cambridge University Press.

Farris, F. A. (2010). The Gini index and measures of inequality. The American Mathematical Monthly, 117(10), 851-864.

Haughton and Khandker (2009). Handbook on Poverty and Inequality, Chapter 6.

Pen, J. (1971). Income Distribution: facts, theories, policies. Praeger.

Pyatt, G. (1976). On the interpretation and disaggregation of Gini coefficients. The Economic Journal, 86(342), 243-255.

Shorrocks, A. F. (1980). The class of additively decomposable inequality measures. Econometrica: Journal of the Econometric Society, 613-625.

Thank you for your attention

Homework

Exercise 1 – Engaging with the literature

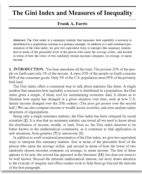
The Gini Index and Measures of Inequality

Frank A. Farris

Considering equations (10) to (12) in Farris (2010) give a brief interpretation of a Gini index of 63% for South Africa

Exercise 1 - Engaging with the literature

A solution

- The Gini index shows "how the lower of two randomly chosen incomes compares, on average, to mean income".
- E.g., if the Gini index for South African family income is 0.63, "we conclude that the lower of two South African family incomes, chosen at random, is about 37% *[=(1-0.47)*100]* of the mean; on the average, the poorer of two families earns only over one third of the national mean".

Exercise 2 - Inequality in South Asia

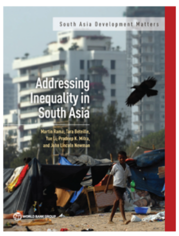

- Turn to page 2 of this report (see next slide)
- What criticisms would you make to this chart?

FIGURE 1 **Based on standard monetary indicators, South Asia has moderate levels of inequality**

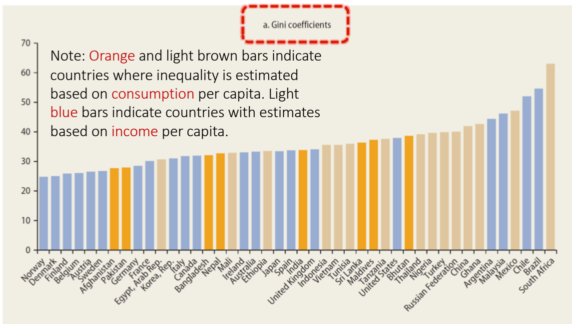

Exercise 3 - Functional vs Personal distribution of income

- The nature of the relationship that links the evolution of income shares to income inequality is complex and still widely debated among researchers.
- In that context, comment on Figure 19 of the ILO Global Wage Report 2016/2017.

C4D2 TRAINING 65