Econ 5700 slides
Inequality

September 26, 2017

The distribution of individual incomes

Consider two ways of calculating *mean* GDP/capita of the world

- Total world GDP over total world population

- Unweighted mean per-capita GDP across countries of the world

Correspondingly different ways of calculating the *dispersion* (e.g., standard deviation) of the world income distribution

Important for understanding trends in world income inequality over time: different measures can give different results

Hypothetical example

- Africa: three poor countries, each with population $= 1$, GDP/capita $= 2$
- China: population $= 10$, GDP/capita $= 4$ (initially)
- USA: population $= 3.8$, GDP/capita $= 10$

No income dispersion within each country (for now)

Country	Population	GDP/capita	GDP
Africa 1	1	2	2
Africa 2	1	2	2
Africa 3	1	2	2
China	10	4	40
USA	3.8	10	38
Sum or mean	16.8 (sum)	5 or 4 (mean)	84 (sum)

World mean GDP/capita could be:

- Total world GDP over total world population $= 84/16.8 = 5$
- Unweighted mean of GDP/capita across five countries $= 4$

Similarly different ways to calculate variance

Suppose China grows from GDP/capita of 4 to 5 (sort of what happened)

What happens the variance in GDP/capita?

Result depends on method

Method A: non-weighted cross-country measures

Each country is one observation regardless of population size

Initial mean $= (2+2+2+4+10)/5 = 4$

Initial variance $= 9.6$

$$\frac{3 \times [2-4]^2 + 1 \times [4-4]^2 + 1 \times [10-4]^2}{5} = \frac{12+0+36}{5} = 9.6$$

Mean after China has grown $= (2+2+2+5+10)/5 = 4.2$

Variance after China has grown $= 9.76$

$$\frac{3 \times [2-4.2]^2 + 1 \times [5-4.2]^2 + 1 \times [10-4.2]^2}{5}$$

$$= \quad \frac{14.52 + 0.64 + 33.64}{5} = 9.76$$

Intuition: China grows from 4 to 5, and thus departs from the mean (4), which raises inequality

Method B: weighted cross-country measures

If no inequality within countries, same as one person being one observation

Initial mean $= 84/16.8 = 5$

Initial variance ≈ 7.86

Mean after China has grown ≈ 5.6

Variance after China has grown ≈ 6.91

Intuition: China grows from 4 to 5, but now approaches the mean (5), which lowers inequality

	No pop weighting	With pop weighting
Variance before China growth	9.60	7.86
Variance after China growth	9.76	6.91
Change	$(+)$	$(-)$

Conclusion: whether dispersion has increased or decreased depends on how we measure it

Example stylized, but somewhat realistic: initially poor, populous countries (China, India) have grown above average last few decades

If no income dispersion within countries, then population-weighting method gives the variance of the world distribution of individual incomes

But of course there *is* dispersion within countries

Moreover, this could be quantitatively important when looking at changes over time, because within-country income inequality has increased a lot in fast-growing countries

So we would like to estimate the distribution of GDP/capita in which each person is one observation

Sala–Martin (2002) starts by estimating distributions of per-capita income within countries, and then aggregates to one world income distribution

Two-country example:

Let $F(y)$ be the cumulative density function (cdf) for the whole world

$F(y) =$ fraction of world population with incomes below y

Let $F_A(y)$, $F_B(y)$ be the cdf's for countries A and B, with population levels P_A and P_B

$P_A + P_B =$ world population

Then

$$\underbrace{(P_A + P_B)F(y)}_{\text{pop below } y \text{ in the world}} = \underbrace{P_A F_A(y)}_{\text{pop below } y \text{ in A}} + \underbrace{P_B F_B(y)}_{\text{pop below } y \text{ in B}}$$

Or:

$$F(y) = \left(\frac{P_A}{P_A + P_B}\right) F_A(y) + \left(\frac{P_B}{P_A + P_B}\right) F_B(y)$$

If we have cdf's and population shares for each country, we can find cdf of the world

Note: generalizes to many countries; cdf's and population shares can be time dependent, and so can the set of countries

Let F_t be the world cdf at time t, $F_{i,t}$ the cdf of country i at time t, N_t the set of countries at time t, then:

$$F_t(y) = \sum_{i \in N_t} \pi_{i,t} F_{i,t}(y)$$

$\pi_{i,t}$ easy to find, $F_{i,t}$ harder

Estimating income distributions

Some survey data on within-country income distributions available from World Bank for 138 countries, various years

Information on quintile income shares = income shares of the 20%, 40%, 60%, 80% poorest in the population (in given country and year)

For missing countries/years: interpolate (disregard for now)

Want to use income shares to estimate the whole distribution

Tricky: looking for a function; not just a single parameter (e.g., the mean)

Different techniques: parametric and non-parametric

Below: some discussion

Parametric: specific parametric distribution assumed

Example:

- Suppose we know income levels, y, for a sample of individuals

- Assume these are drawn from a *log-normal* distributed: $\ln(y) \sim N(\mu, \sigma)$ where μ and σ are mean and standard deviation, which can be estimated from the sample

- Since these two parameters characterize the whole distribution we get an estimate of cdf $F(y)$ and pdf $f(y) = F'(y)$

- Now we can calculate any moment of the distribution, e.g., the ratio of the income shares of the richest 5% over the poorest 5%

Non-parametric: no specific parametric distribution assumed

Idea:

- Consider first a simple histogram over sampled data
- The shape of the histogram depends on the number of bins, and the number of observations in each bin, which determines the height of each bar
- The result is typically a "jagged" structure, due to locally varying bar heights
- Instead: estimate a smooth function within each bin, using a *kernal density function*

Illustration: look at some examples based on simulated data

- Three variables: one drawn from a normal distribution, one from a log-normal distribution, one constructed in a messy way to have a bimodal (twin-peaked) histogram

- For each variable, ask Stata to estimate the distribution in two ways: (1) assuming the distribution is normal (even though it may not be); and (2) using a kernal density function

Results: see plots

Which method is best?

The log-normal distribution indeed proxies income distributions of many developed countries well; often used in macro applications

Works less well here, because many countries have bimodal income distributions; log-normal distribution is unimodal (cf. Footnote 15)

Sala-i-Martin uses non-parametric estimate

Results

Once Sala-i-Martin has the distribution from each country, he makes them comparable across countries by "anchoring"

Means the distribution is shifted to make the mean for each country-year coincide with GDP/capita data from the PWT

Then he aggregates, using the time-varying populations of each country

Figure II shows how distributions for some countries have evolved over time

- Most have grown, but not all (e.g. Nigeria); some have bimodal (or multimodal) distributions, including the US

- Note: populations on vertical axis (rather than densities); surface under curves equals total world population

Figure III aggregates all distributions into one world distribution

- In 1970 almost bimodal, in 1990 less dispersed

- China and India have moved in where there was less mass in 1970

World trends in poverty

Figure IV shows world income distributions in different years in one single diagram. Note how the distribution shifts to the right

Same in Figure V but with cdf's

Interested in how the fraction of the population below certain income thresholds have changed over time

- Declines almost regardless of which threshold is used (Figure VI, Table I)

- These are measures of poverty *rates.* Note also decline in "head counts" despite population growth

Regional trends in poverty

- See Table II, Figure VII

- Decline in poverty rates (here $570/year) found across the world, except Africa; largest declines in East and South Asia

- "The great Asian success contrasts dramatically with the African tragedy." (p. 377)

- Africa accounted for almost 75% of the world's poor in 2000, but only 15% in 1970 (p. 380)

Trends in other measures of inequality

Multiple ways to measure income inequality (other than poverty rates)

- Gini coefficient
  - Computed as the surface between 45-degree line and the Lorenz curve
  - The Lorenz curve plots cumulative income shares (on vertical axis) against cumulative population shares (on horizontal axis); e.g., poorest x% of population together earn z% of total income
  - With total equality the Lorenz curve coincides with the 45-degree line, implying a Gini coefficient of zero

- Atkinson index with *inequality aversion* parameter (ε) equal to 1 and 0.5
  - Letting μ be the mean income across N individuals with income levels y_1, $y_2 ...,y_N$, the Atkinson index equals

$$A(\varepsilon) = \begin{cases} 1 - \frac{1}{\mu}\left[\frac{1}{N}\sum_{i=1}^{N} y_i^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}} & \text{for } \varepsilon \geq 0 \text{, and } \varepsilon \neq 1 \\ 1 - \frac{1}{\mu}\left[\frac{1}{N}\prod_{i=1}^{N} y_i^{1-\varepsilon}\right]^{\frac{1}{N}} & \text{for } \varepsilon = 1 \end{cases}$$

  - If $\varepsilon = 0$, then $A = 0$ (zero inequality aversion)
  - If $y_1 = y_2 = ... = y_N = \mu$, then $A = 0$ (no inequality)

- Variance in log per-capita incomes

- The ratio of average income among the richest 20% over that of the poorest 10%; same for 10% richest/poorest

- Mean Log Deviation and Theil Index (skip for now)

All measures show declines in world income inequality from 1970 to 2000

Declines not monotonic; see e.g. trends in the Gini coefficient in Figure VIII

What drives these trends? Short answer: China

- If we exclude China, then Gini coefficient has increased rather than decreased (Figure IX)

- But if we exclude China, US and Africa ("main convergers" and "main divergers"), then still decrease in Gini

Broader questions

Why does inequality matter? What causes it?

Discussion in Ray (2010) about these broader questions

Central idea: when a country starts to grow not all sectors will grow equally fast

Growth tends to be (initially) *uneven* (or unbalanced)

Starting point: the "tunnel parable" by Hirschman and Rothschild (sadly car-centric)

- Cars stuck in tunnel in different lanes; initially no lane is moving

- Then one lane starts to move. Possible reactions among drivers in other lanes:

  - Hope your own lane will start moving

  - If it doesn't: resentment, desire to change lane, and/or "redistribute" from faster lanes to slower

Uneven growth thus has implications for, e.g., occupational choice, political economy, and more

Sources of uneven growth

In most off-the-shelf growth models, economy gravitates toward a balanced growth path where all sectors (production factors) grow at the same rate

(In this course, only one-sector growth models)

Why may the balanced growth path not be a good approximation of reality?

Composition-of-demand explanations

- If workers in growing manufacturing sector buy goods from non-growing (agricultural) sector, say food or services, then there should be spillovers

- With homothetic preferences (e.g., log utility) that always happens

- With non-homothetic preferences it may not: e.g., spending on food may not be proportional to income

- Extreme case: when manufacturing incomes are spent only on manufacturing goods, agricultural incomes only on agricultural goods, and there is no migration of workers across sectors; called a *dual economy*

Obstables to factor mobility

- Workers in non-growing sector should want to switch to growing sector

- What might prevent them from doing that? One example is borrowing constraints together with high costs of education (Galor and Zeira 1993)

- Even when educational levels do rise it can take a generation, at which point technologies may have changed

Reactions to uneven growth

Political-economy mechanisms

- Lobbying to protect or subsidize non-growing sector
- Reversal of fortune effects: previously prosperous sectors may have resources, or political influence, to lobby to prevent new sectors from emerging

- Conflict, civil war
  - Mixed evidence: often ethnic and religious in nature; but economic shocks do matter (eg., rainfall)
  - "Uneven growth" could be relevant for conflict