ECON 2456: Measuring Poverty

Robert Szarka

October 14, 2010

1 Introduction

Why measure poverty? Presumably, we want to know how big of a problem we have and whether our efforts to eliminate, or at least ameliorate, poverty are working. Thus, we seek a statistic (or set of statistics) that consistently measures both the *incidence* (how many people are poor?) and *intensity* (how poor are they?) of poverty.

Let's begin by supposing that for each of the M households (or, equivalently, families) in our population of N people we have a measure Y_i of household income and some poverty threshold Z_{pi} against which we can compare each household's income, and that q households fall below their respective thresholds. There may be issues with these measures, of course. For example, our choice of threshold Z_{pi} will have a profound effect on our measurement of both the incidence and intensity of poverty, as well as the change in our statistic over time.[1] But let's ignore that problem here.

[1] As Wolff has calculated (p. 115), the choice of the CPU-U rather than the CPI-U-RS as the deflator resulted in a measured increase of 1.5 percentage points in the US poverty rate over 1973-2005, rather than a 2.8 point decrease.

2 The Head Count Ratio

The most widley used measure is the *poverty rate* or *head count poverty ratio*. The idea is simple: what percentage of the population is poor? To be more precise: since we're working with household data, it measures the percentage of the population living in poor households. Given a weight w_i that gives the number of people in each household, the head count ratio can be expressed as $HCR = (1/N)\sum_{i=1}^{q} w_i$. Alternatively, following the general approach of Foster, Greer, & Thornbecke[2], we can write it in the form

$$HCR = (1/N)\sum_{i=1}^{q}[w_i(g_i/Z_{pi})^0] \quad (1)$$

where $g_i = Z_{pi} - Y_i$ (this is the *poverty gap* for household i; see below).

A similar measure $P_0 = q/M$ could be defined as in Wolff (p.101) and written as

$$P_0 = (1/M)\sum_{i=1}^{q}(g_i/Z_{pi})^0 \quad (2)$$

This would give the proportion of the households that are poor, which will in general be different than HCR as defined above.

Whether we measure the poverty rate as a percentage of individuals living in poor households, or as a percentage of households that are poor, the head count ratio captures only the incidence of poverty, not its intensity. Thus, a policy that ameliorates but doesn't eliminate poverty for a poor household would not impact this measure; nor would a policy that redistributes income from a poor household to a poorer household.

[2] Foster, James, Joel Greer, and Erik Thorbecke. A Class of Decomposable Poverty Measures. Econometrica 52, no. 3 (May 1984): 761-766.

3 The Poverty Gap Ratio

In order to get some idea of the intensity of poverty, then, we need to turn to another kind of poverty measure. The simplest such measure may be the *aggregate poverty gap*, defined as $G = \sum_{i=1}^{q} g_i$. That is, G is the total money income we would need to transfer to poor households in order to completely eliminate poverty for the period in question. Of course, this assumes that we could target the transfers precisely, so that only poor households received transfers and each received exactly the amount needed to raise their income to their respective poverty threshold. But, in principle at least, this number answers a tangible question.

The aggregate poverty gap lacks context, though. How large is this gap relative to the resources of the economy at large? And, if the number increases from one year to the next, is this because poverty has increased or because there are simply more households/people in the economy?

Therefore, a common approach—taken by several different measures sharing names like *poverty gap ratio*, *poverty gap index*, *income gap ratio*, and *income deficiency index*—is to normalize the poverty gap (by individual or household, or in aggregate) so it ranges between 0 (when there is no poverty gap at all) and 1 (when the poverty gap is as large as possible, though the interpretation varies for each statistic). For example, the simplest such measure might divide the aggregate poverty gap by the aggregate income poor households would need to reach their respective poverty thresholds:

$$AR = G / \sum_{i=1}^{q} Z_{pi} \tag{3}$$

While this measure does capture something about the intensity of poverty, and could be useful in conjunction with the poverty rate, consider the following situation: suppose we transfer \$1 from a rich household to poor household i with $g_i = 1$? Since that household would no longer be counted as poor, and assuming that the remaining poor households have a larger poverty gap, AR will actually increase! To address this problem, we may normalize by dividing by the total number of households rather than only the poor ones. We could also construct similar measures by using weights to convert our household data into individuals, children, "adult equivalents", etc.

A more complex version of a poverty gap measure is given by Wolff (p. 100), which we can write as

$$R = (1/q)\sum_{i=1}^{q}(g_i/Z_{pi})^1 \tag{4}$$

The interpretation of R is the mean distance separating poor households from the poverty line. As with AR above, R gives us a measure of the intensity of poverty, but is subject to the problem that improvements at the margin can increase R. (We could, as above, use appropriate weights to convert this into a measure of individual poverty, rather than household poverty. But it would be subject to the same problem.)

Therefore, one widely-used poverty gap index measure is as defined in Wolff (p.101):

$$P_1 = (1/M)\sum_{i=1}^{q}(g_i/Z_{pi})^1 \tag{5}$$

The interpretation is the same (the mean distance of a household to the poverty line), but here we normalize by using all households and assigning a distance of 0 to households above their respective poverty thresholds. (Again, we could use appropriate weighting to convert this to a measure of individual poverty.)

4 Foster-Greer-Thorbecke Measures in General

In general, we can take different values of α to produce a variety of measures of the form

$$P_\alpha = (1/N)\sum_{i=1}^{q}(g_i/Z_{pi})^\alpha \tag{6}$$

where N can be the number of households, families, or individuals, as appropriate, and q is the number of corresponding units classified as poor. We have already seen that when $\alpha = 0$, we obtain a head count ratio; while, when $\alpha = 1$ we obtain a poverty gap ratio. Another common measure, the *squared poverty gap ratio* takes $\alpha = 2$. This measure captures not only the intensity of poverty, but also the inequality of the distribution of poverty among poor households.

Why might we want to use the squared poverty gap ratio, or another measure that captures inequality among poor households? If helping the poorest of the poor is a policy priority, using a measure like P_2 will ensure that the statistic reflects progress toward that goal. As the saying goes, "You can't manage what you don't measure."

5 Lessons

What lessons can we draw from our encounter with the confusing world of poverty measures? (We've only scratched the surface!) Here are a few suggestions:

1. If you want to use a poverty measure, it's important to know how it's constructed. The name alone is not enough to tell you whether it measures what you think it measures. This is especially important when comparing statistics from different sources, which may use different methodologies.

2. Conversely, if you're publishing a poverty measure, you need to detail your methodology so consumers of your statistics can use them appropriately.

3. If you're lucky enough to have the data and the resources to construct your own measure from scratch, choose the one that will best measure what you want policy to accomplish.

4. One measure may not be enough.

Now, go forth and measure!