

Microeconomic Theory

Lecture 7: Poverty and Inequality

Marcel Fafchamps
Oxford University

"Money is better than poverty, if only for financial reasons." Woody Allen.

Readings: Jean-Yves Duclos, Poverty and Equity: Measurement, Policy and Estimation With DAD, available on the internet at http://www.idrc.ca/en/ev-97152-201-1-DO_TOPIC.html

1. The measurement of poverty and welfare

Economists have two ready concept for the measurement of welfare: utility and its sister, equivalent variation. Utility provides a ranking among bundles of goods. Equivalent variation translates these rankings into an easily understandable money metric measure. These concepts are very general and can be applied to many different situations. In practice, using these concepts to make judgements about welfare is tricky. This is because we do not observe individual utility functions. We discuss here some of the issues raised by the measurement of poverty and welfare.

1.1. Public goods

When all goods and services are purchased through the market, making welfare judgement is easy. Not so when goods are publicly available (or imposed) so that individuals do not decide how much of a good they consume. Think of pollution, road congestion, or biodiversity. These goods enter utility but are not priced.

It is also possible that certain goods are purchased but not at a price that clears the market. Think of health care, for instance. It is common for many developing countries to set user fees below actual cost. This is a simple and effective way of providing health insurance without having to set an explicit insurance system. What people pay does not represent the true cost of what they receive. This typically leads to rationing. The market then needs to clear in other ways, such as queues, waiting lists, loteries, and the like. In such an environment, what people pay for health care is less than the value they derive from it. Actual health care expenditures underestimate the true utility value of the services received.

When certain goods are rationed or publicly provided, private consumption expenditures no longer represent the value of the entire consumption vector. Some of the things that are consumed are not priced. They may be very important to the individual, but they are not measured by private income or consumption expenditures.

If we knew the shape of the direct utility function, we could calculate the equivalent variation of non-priced or subsidized items and add it to y_i . But in many cases, we do not know how these goods are valued by consumers. In those cases, it may be necessary to look at quantities consumed. If consumption itself is not measurable, we may be forced to use a proxy, such as the nutritional or health status of the individual.

It is not necessary that a good be purchased on the market for it to appear in the utility function. For instance, people derive utility from the air they breathe even though fresh air is not, as a rule, for sale (unless you are a scuba diver!). There are many freely available or publicly provided goods and services that generate utility without being acquired through the market. Think of parks, roads, access to health care, etc. Of course, some of these goods and services may only be accessible to those who reside nearby, in which case rents and property values probably reflect, in an implicit way, the value of these services.

A good example is access to a good school. It has been shown, for instance, that property values are higher in good school districts – think of Summertown with its access to Cherwell compared to Cowley. Although in this case parents do not pay an explicit school fee to send their child to a private school, they pay an implicit school fee via higher property prices (higher rents, higher mortgage payments). So much for the idea that government schools are less elitist...

Utility can also be applied to goods that are not rival, such as the knowledge that blue whales roam the ocean. If people value this reality, then it appears in their utility function. If something appears in the utility function, we can in principle compute its equivalent variation: it is the additional income someone would need to be equally satisfied – to have the same utility – in a world without blue whales. Contingent valuation is the name of the method that seeks to elicit such value. It has been applied by economists to measure the value of various environmental goods, such as old growth redwood forests or iconic wildlife species.

1.2. Choice

There is one thing, however, that cannot be accommodated by the utility function: choice itself. The direct utility function is defined over outcomes – goods and services consumed. It places no value on the process by which outcomes are achieved. For instance, suppose your child always orders a chocolate ice cream. Knowing this, when the ice cream vendor shows up you may simply order a chocolate ice cream for your child. As a parent, you will quickly discover, however, that this is not a good idea: your child values the process by which the ice cream is selected. Even if your child chooses a chocolate ice cream, he values having had the choice to order something else.

This issue arises all the time in development economics. Compare women in Northern and Southern Nigeria, for instance. Many women in the North are kept in purdah, which means that they are not allowed out of the house. As a result they do not work in the fields. In contrast, Southern Nigerian women toil in the fields all year long but have more individual freedom – including in many cases financial independence. Who is better off? Some women choose not to work and prefer to stay at home with their children. Others want to work. What if society makes the choice for them? What if society is paternalistic and decides what is good for people – wearing seat belts, refraining from alcohol, cigarettes and drugs? What of free food and shelter for the homeless, many of whom would rather spend the money on drugs or alcohol? What of development aid that focuses on clean water and schooling when people would rather buy new clothes or a television?

These issues of process and choices are unescapable for any economist, but especially for development economists. As development aid increasingly takes the form of welfare programs, this raises the issue of paternalistic choices and patronizing attitudes. Should aid be conditional? Should aid agencies behave like, say, the Salvation Army that provides food and shelter to those who put up with the moralizing speech and the lack of booze? These are difficult ethical and political issues which we will not solve here. But they underlie much of the debate surrounding the actual measurement of poverty and inequality.

1.3. Consumption and income

Suppose for a moment that all goods and services are purchased on the market. In this case, measuring individual welfare is fairly straightforward. The indirect utility function of individual i is $V_i(y_i, p)$. Since V_i increases monotonically with income, we can use y_i as measure of i 's welfare: the higher income is, the higher utility is.¹ We do not need to know on what goods and services i spends his or her income. Income measures the consumption possibility set of individual i , the range of consumption choices that i can make. When all goods and services can freely be purchased on the market at a given price, income is a welfare measure that describes all the choices open to an individual. Put differently, it measures choice even though direct utility, per se, only values outcomes. This is quite a remarkable result.

We have not yet introduced time in our model of individual behavior. But it should be clear that, in any given period, people may not spend all their income – or they may spend more than their income. The difference is saving and dissaving. Since utility depends on consumption, not on income directly, the welfare of an individual over a given period should be measured by total consumption expenditures (although this ignores the sense of security that wealth may provide for future consumption).

We do not need to observe consumption of individual items. Some people may consume more food, others more entertainment. It is their choice. Income – or total consumption expenditure – is all we need to measure whether the utility of an individual rises or falls. If we decided to judge people's welfare on the basis of their consumption of individual goods – e.g., calorie intake – we would be making a value judgement on their choice of consumption basket. Focusing on total consumption expenditure recognizes individual freedom of choice.

1.4. Capabilities and functionings

In an attempt to deal with these difficulties, Sen proposed a generalization of the utility approach to multiple dimensions. Capabilities are defined as the range of choices open to the individual. Examples include the freedom to choose where to work, to seek a healthy lifestyle, to pursue knowledge, to be free of the fear of discrimination, crime, political persecution, etc. These freedoms or capabilities are but a generalization of the idea of personal choice embedded in the utility framework and measured by income. In a multidimensional context, income is but one form of capability.

Functionings are defined as outcomes, such as being in good health, being educated, employed, well fed, clothed, housed, rested, etc. Functionings are the result of individual choices. Some people may choose to become educated, others prefer to work and start a family when they are young. Similarly, some people may opt for a healthy diet and exercise, while others prefer to watch television with a bag of chips. Judging someone's welfare by examining functionings may give an incorrect picture of capabilities: a rich person may elect to lead a frugal life while a poor person turns up obese. Comparing their calorie intake would yield a distorted view of their capabilities.

It is not difficult to agree that capabilities are good things. Unfortunately in many cases they are not measurable. Functionings, on the other hand, can be observed and much data exists on them. Paradoxically, Sen's name is often invoked by those who pass judgement on individual welfare by using measurable functionings such as literacy, nutritional status, infant mortality, or female participation in the labor force. So doing, they forget Sen's main message, which is that

¹Here I ignore price variations. These were discussed in the lecture on welfare and need not be discussed further here.

when we use functionings to estimate welfare, we implicitly make a patronizing or paternalistic judgement on how other people choose to live their life. This is particularly questionable when those being judged are poor and those passing judgement are enjoying the comfort provided by developed economies.

1.5. Subjective well-being

A different approach to the measurement of individual welfare is simply to ask people whether they are happy or satisfied with the life they live, or whether they judge their consumption or income level adequate for their needs. This is the approach taken by many psychologists and by an increasing number of economists.

There is noise in the data – answers are affected by the weather, the time of the day, or the mood of the respondent. But the evidence has nevertheless shown that answers to subjective questions about happiness and satisfaction are strongly correlated with income, consumption and various capabilities or functionings – such as being married, employed, and free of physical disability. Experiments have also shown that answers to these subjective questions are correlated with objective measures of happiness, such as suicide rates or clinical depression.

The evidence has also shown, however, that answers to subjective well-being questions are calibrated on the standards of living prevailing in the surveyed population at the time of the survey. Average subjective happiness, for instance, did not increase noticeably in the US since the 1950's in spite of a massive increase in average consumption. Many interpret such findings as implying that people answer these questions in a relative way, as if they derive utility from their consumption relative to that of others, not from the absolute level of their consumption. Psychologists think this is due to a reference effect: people derive satisfaction from their achievements in life, and achievements are judged relative to a peer group. For evidence in developing countries, see for instance recent papers by Lokshin and Ravallion (Malawi), Kingdon and Knight (South Africa) and myself (Nepal).

1.6. Interpersonal comparisons

So far we have discussed how to judge someone's welfare level. Now we discuss how to make interpersonal comparisons. What we want is a criterion that does not depend on individual preferences. If all consumption goods are purchased on the market, interpersonal comparisons appear easy: it seems that we just need to look at income. People with more income can acquire more goods and thus ought to reach a higher utility level.

There several conceptual – or even philosophical – difficulties with this approach. First, as we just discussed, individuals may have different reference points. This means that while the average Indian household may be quite happy with being offered an income of \$10,000 a year, the average English household probably would not (even correcting for price differences). This is because, in India, many people live with much less than \$10,000 a year, while in England the opposite is true. Poverty analysis typically chooses to ignore this issue and sets universal absolute standards instead, such as the \$1/day and \$2/day criteria used in the MDGs.

Second, even if we choose to ignore reference points – and there are good ethical and philosophical reasons for doing so – we may want to correct for individual specific features, such as disability. The idea is that a disabled person may need more income to reach a level of consumption comparable to that of an able person. Similar principles may apply to pregnant women, children, the elderly, etc. We will not delve further in these issues here. But it is important to be aware of them, because in some policy debates the issue is precisely how to make amendments to poverty criteria that account for such differences.

Third, human beings normally live in households consisting of several individuals. The reason, of course, is that they can share household public goods, etc. This raises the issue of how to compare individuals living in households with a different composition. In the literature, this issue appears under the name of adult equivalence or consumption equivalence scale. This is a contentious issue that can have a profound effect on interpersonal comparisons. To my knowledge, there is no clear accepted way of dealing with this issue.

If we abstract from these difficulties and are willing to assume that all individuals share the same utility function $V(y)$, there remains the issue of how to measure how different people are. For instance, say individuals A, B, C and D have incomes 100, 110, 1000, 1010. Is the welfare difference between A and B the same as between C and D ? The income difference is the same but, intuitively, we suspect that the additional income is more valuable to A and B than it is to C and D . This is because of satiation: as people get more stuff, we expect them to value additional units less than initial units. Applying this principle to consumption as a whole, it leads us to believe that $U(y)$ is concave, that is, $U''(y) < 0$.² This implies that:

$$U(110) - U(100) > U(1010) - U(1000)$$

What remains unclear is how curved $U(\cdot)$ should be.

Now that we are better understand the pitfalls and shortcomings of the approach, we are now ready to delve into the measurement of income distribution as an approximation for the interpersonal distribution of welfare. We begin with an overview of the tools developed by economists to describe income distribution. This will take us over to the next lecture. We will then see how these tools can be applied to measure poverty.

2. Measuring income distribution

We begin by developing some notation and by describing how income distribution can be characterized in a simple way.

2.1. Lorenz curve and Gini index

Let $p = F(y)$ be the proportion of individuals with income $\leq y$. This is the cumulative distribution function. We have $F(0) = 0$ and $F(\infty) = 1$.

The *Lorenz curve* $L(p)$ is a useful visual tool to represent income distribution. The Lorenz curve gives the share of total income earned by the poorest proportion p of the population.

The *Gini index of inequality* is simply the area between the Lorenz curve and the line of perfect equality. To ensure that the index takes values between 0 (perfect equality) and 1 (perfect inequality), the area is multiplied by 2. Formally, the Gini index is defined as:

$$Gini \equiv 2 \int_0^1 (p - L(p)) dp$$

2.2. Welfare function

The Gini index has been generalized by introducing weights. The purpose of these weights is to give more importance, in the definition of inequality, to poor individuals. A more direct approach is to use explicit welfare functions.

²Of course $V'(y) > 0$ since more income raises welfare.

Define the *quantile function* as the level of income of the p th percentile individual. Put differently, $Q(p)$ is the income level below which we find of proportion p of the population. A social welfare function W is nothing but a weighted average of a utility function defined over individual incomes $U(Q(p))$ with:

$$W \equiv \int_0^1 U(Q(p))\omega(p; \rho)dp$$

Function $U(y)$ may be the utility function of a representative consumer. Or it could be the welfare function of the social planner/policy maker, i.e., the weight the policy maker puts on inequality. Function $\omega(p; \rho)$ is a yet-to-be-specified weighting function. Here we only consider the case where $\omega(p; \rho) = 1$. In that case, W is simply the average value of $U(\cdot)$ over the total population.

In general we wish W to be *homothetic* so that the ratio of marginal utilities of any two individuals is unchanged if we multiply all incomes by a constant – if we express incomes in pounds or pennies. For this to happen, we need $U(y)$ to take a specific functional form:

$$\begin{aligned} U(Q(p); \varepsilon) &= \frac{Q(p)^{1-\varepsilon}}{1-\varepsilon} \text{ when } \varepsilon \neq 1 \\ &= \log Q(p) \text{ when } \varepsilon = 1 \end{aligned}$$

This functional form has been used extensively in economics to model decisions over time and with uncertainty. For now we just note that they ensure that the ratio of marginal utilities is homogeneous of degree 0 in all incomes.

Using an equivalent variation concept, let us define the *equally distributed equivalent (EDE) income* $\xi(\rho, \varepsilon)$ as the income that satisfies:

$$U(\xi(\rho, \varepsilon)) = \int_0^1 U(Q(p); \varepsilon)\omega(p; \rho)dp$$

The index of inequality $I(\rho, \varepsilon)$ corresponding to welfare function $W(\rho, \varepsilon)$ can then be defined as:

$$I(\rho, \varepsilon) \equiv \frac{\mu - \xi(\rho, \varepsilon)}{\mu}$$

The advantage of this measure is that it has a money metric interpretation since it is expressed as a percentage of average income.

3. Measures of poverty

Two approaches have been used to devise poverty indices. The first transforms income and the poverty line into poverty gaps. The second applies the concept of equally distributed equivalent (EDE) income to those below the poverty line.

3.1. The poverty gap approach

Let z be a poverty line. Define

$$Q^*(p; z) \equiv \min(Q(p), z)$$

The *poverty gap* g at percentile p is defined as the distance from the poverty line to the income of percentile p :

$$g(p; z) \equiv z - Q^*(p; z)$$

Note that for individuals over the poverty line, $g(p; z) = 0$. The *average poverty gap* is simply the average of the poverty gap over all individuals:

$$\mu^g(z) \equiv \int_0^1 g(p; z) dp$$

Individuals above the poverty line do not enter into the calculation of the average poverty gap. The average poverty gap represents the amount of money that would be needed, on an average basis or per capita basis, to take all poor people to the poverty line.

The average poverty gap belongs to a popular class of poverty gap indices called the *Foster-Greer-Thorbecke (FGT) class*. These indices use an ‘ethical’ parameter α and are defined as:

$$\begin{aligned} \bar{P}(z; \alpha) &\equiv \int_0^1 \left(\frac{g(p; z)}{z} \right)^\alpha dp \text{ (normalized version)} \\ P(z; \alpha) &\equiv \int_0^1 g(p; z)^\alpha dp \text{ (un-normalized version)} \end{aligned}$$

This class of indices includes some important special cases:

1. When $\alpha = 0$, the FGT index yields the *poverty headcount ratio*, which is simply the proportion of the population in poverty:

$$P(z; \alpha = 0) = \bar{P}(z; \alpha = 0) = F(z)$$

The MDG’s, for instance, are expressed in terms of poverty headcount ratio. Many World Bank-sponsored poverty assessment exercises also focus on the headcount ratio.

Note that the poverty headcount ratio does not put any weight on how poor people are. It just counts the number of people below the poverty line. Taken literally, this can have some disturbing policy implications. For instance, if the policy objective is the reduction of the poverty headcount, then we may consider taking money away from the very poor to give to the less poor so as to push them just above the poverty line. For instance, one may get a long way towards satisfying the MDG’s by getting many somewhat poor Chinese above the poverty line while at the same time letting Africa sink further below the poverty line. Ethically, this does not seem right. This is where the average poverty gap comes into the picture.

2. When $\alpha = 1$, the FGT index yields the *average poverty gap* defined earlier:

$$P(z; \alpha = 1) = \int_0^1 g(p; z) dp = \mu^g(z)$$

The average poverty gap has an easy intuitive interpretation. If you multiply this number by the size of the population, you get the total amount of income transfer that would be required to take all the poor out of poverty.

The poverty gap is a more satisfying welfare measure since, unlike the headcount ratio, it cannot be improved by taking from the very poor to give to the marginally poor. Doing so would simply leave the poverty gap unchanged. But the poverty gap still does not capture inequality among the poor: taking from the very poor to give to the less poor leaves the poverty gap unchanged. To get inequality aversion among the poor, we need to introduce more curvature, which requires setting a value of α greater than 1. Unfortunately, FGT measures are less intuitive for values of α other than 0 or 1. They also have strange properties that are discussed in Duclos.

3.2. The EDE approach

The FGT class of poverty indices is relatively easy to use and to explain to non-economists. But it is based largely on ad hoc assumptions. The EDE approach, in contrast, builds directly on welfare economics. The poverty index is defined as:

$$P(z; \rho, \varepsilon) \equiv z - \xi^*(z; \rho, \varepsilon)$$

where ξ^* is the number that satisfies the following equality:

$$U(\xi^*(\rho, \varepsilon)) = \int_0^1 U(Q^*(p; z); \varepsilon) \omega(p; \rho) dp$$

The advantage of this approach is that we can apply the welfarist approach to a description of inequality among the poor. Intuitively, from the satiation principle, welfare goes down if there is more inequality among the poor. This is because, other things being equal, more inequality among the poor means some people are really very poor and probably suffer from extreme destitution. The EDE approach is an intellectually elegant and theoretically consistent way of describing these concerns.

If we choose $\varepsilon = 0$ in $U(Q(p); \varepsilon) = \frac{Q(p)^{1-\varepsilon}}{1-\varepsilon}$ and set $\omega(p; \rho) = 1$ we obtain a *Gini index of poverty* which is also the average poverty gap:

$$\begin{aligned} P(z; \rho, \varepsilon = 0) &= z - \int_0^1 Q^*(p; z) dp \\ &= \int_0^1 g(p; z) dp \\ &= \mu^g(z) \end{aligned}$$

3.3. Poverty rankings

Ultimately, we hope to use poverty indices to make comparisons across groups – e.g., to decide who deserves a targeted intervention – or across counter-factuals when making policy decisions. Unfortunately, poverty indices do not always yield unequivocal rankings of alternatives.

To illustrate this, I borrow an example from Duclos, reproduced in the table below.

	Distribution A	Distribution B
Individual 1	4	6
Individual 2	11	9
Individual 3	20	20
$F(5)$	0.33	0
$F(10)$	0.33	0.66
$\mu^g(10)$	2	1.66

We have three individuals and two possible poverty lines, $z = 5$ and $z = 10$. With the first poverty line, the headcount ratio $F(5)$ is 0.33 (1 person out of three) in distribution A, and 0 in distribution B. There is thus less poverty in distribution B. Results are reversed if we use another poverty line, however: with $F(10)$, distribution B is worse than A. Rankings are also sensitive to the choice of FGT measure. If we use the average poverty gap instead of the headcount but keep the poverty line unchanged, the ranking again is reversed. What this illustrates is the difficulty of making poverty rankings that are robust to changes in poverty line definition or to choice of poverty measure.

4. Policy implications of poverty analysis

Counting the poor tells us nothing on how to eliminate poverty. Let me illustrate this with some examples.

1. Say we find that there are lots of poor people in region A and few in region B. Does it mean that we have to set up special programs to assist region A? Or should we help people migrate from region A to region B?
2. Say we find lots of poor people in occupation A and few in occupation B. Does it mean we have to set up a special program to assist occupation A? Or should we help people retrain so that they can join occupation B?
3. Say we find lots of people will lose their job in firm A because it is going bankrupt. Does it mean that we should save firm A from bankruptcy? Or should we make sure the labor market is flexible and free of discrimination so that laid-off workers can easily find a job in another firm?
4. Say we find that blind people are poor. Does this mean we have to set up financial assistance to help them get by? Or should we help them get a job by ensuring that the workplace is organized in such a way that blind people can contribute productively? Or should we encourage research on blindness, hoping for a technological solution? Or should we stive to reduce the risk of blindness, e.g., by (successfully) fighting river blindness in West Africa?
5. Say we find that divorced women are poor. Does this mean we should prevent divorce? Or provide financial assistance to divorced women? Would that encourage divorce? What of the effect on children? Or should we instead seek to strengthen the status of women within marriage so as to reduce the likelihood that they are abandoned by an opportunistic husband? Etc.

The point of all this is that counting the poor and identifying who they are is a good thing. But it falls far short from suggesting a way of helping them. For this we need a better understanding of the sources of poverty and prosperity, the different reasons why the poor are unable to rise above poverty. Answering these questions is much harder than counting the poor.