Chapter 4. Measures of Poverty

Summary

Assume that information is available on a welfare measure such as income per capita, and a poverty line, for each household or individual. This chapter explains how one may then construct summary measures of the extent of poverty.

The *headcount index* (P_0) measures the proportion of the population that is poor. It is popular because it is easy to understand and measure. But it does not indicate how poor the poor are.

The *poverty gap index* (P_1) measures the extent to which individuals fall below the poverty line (the poverty gaps) as a proportion of the poverty line. The sum of these poverty gaps gives the minimum cost of eliminating poverty, if transfers were perfectly targeted. The measure does not reflect changes in inequality among the poor.

The squared *poverty gap ("poverty severity") index* (P_2) averages the squares of the poverty gaps relative to the poverty line. It is one of the Foster-Greer-Thorbecke (FGT) class of poverty measures that may be written as

$$P_\alpha = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{G_i}{z}\right)^{\alpha}$$

where N is the size of the sample, z is the poverty line, G_i is the poverty gap and α is a parameter; when α is larger the index puts more weight on the position of the poorest.

The *Sen-Shorrocks-Thon index* is defined as

$$P_{SST} = P_0 P_1^P (1 + \hat{G}^P)$$

where P_0 is the headcount index, P_1^P is the poverty gap index for the poor only, and G^P is the Gini index for the poverty gaps for the whole population. This measure allows one to decompose poverty into three components and to ask: Are there more poor? Are the poor poorer? And is there higher inequality among the poor?

Other measures of poverty are available. The *time taken to exit* measures the average time it would take for a poor person to get out of poverty, given an assumption about the economic growth rate; it may be obtained as the Watts Index divided by the growth rate of income (or expenditure) of the poor.

Learning Objectives

After completing the module on *Measures of Poverty*, you should be able to:

c. Describe and explain the headcount index, indicate why it is popular, and explain why it is an imperfect measure of poverty.
d. Describe and compute the poverty gap and poverty severity indexes, and evaluate their adequacy as measures of poverty.
e. Explain and evaluate the FGT (Foster-Greer-Thorbecke) family of poverty measures.
f. Compute the Sen and Sen-Shorrocks-Thon indexes of poverty, and show how the latter may be decomposed to identify the sources of changes in poverty.
g. Compute the Watts index and the related Time-Taken-To-Exit measure.
h. Argue that there is no single best measure of poverty.

Given information on a welfare measure such as per capita consumption, and a poverty line, then the only remaining problem is deciding on an appropriate summary measure of aggregate poverty. There are a number of aggregate measures of poverty that can be computed. The formulas presented here are all based on the assumption that the survey represents a simple random sample of the population, which makes them relatively easy to understand. Where the sampling is more complex – the typical situation in practice – weighting is needed, and the relevant formulas and associated programming are somewhat more difficult, but can be handled fairly easily by most major statistical packages such as Stata and SPSS.

4.1 Headcount index

By far the most widely-used measure is the *headcount index*, which simply measures the proportion of the population that is counted as poor, often denoted by P_0. Formally,

(4.1) $$P_0 = \frac{N_p}{N}$$

where N_p is the number of poor and N is the total population (or sample). If 60 people are poor in a survey that samples 300 people, then P_0 = 60/300 = 0.2 = 20%. For reasons that will be clearer below, it is often helpful to rewrite (4.1) as

(4.2) $$P_0 = \frac{1}{N}\sum_{i=1}^{N} I(y_i < z),$$

Here, $I(.)$ is an indicator function that takes on a value of 1 if the bracketed expression is true, and 0 otherwise. So if expenditure (y_i) is less than the poverty line (z), then $I(.)$ equals to 1 and the household would be counted as poor. N_p is the total number of the poor.

The greatest virtues of the headcount index are that it is simple to construct and easy to understand. These are important qualities. However the measure has at least three weaknesses:

First, the headcount index does not take the intensity of poverty into account. Consider the following two income distributions:

Headcount Poverty Rates in A and B, assuming poverty line of 125					
	Expenditure for each individual in country				Headcount poverty rate (P_0)
Expenditure in country A	100	100	150	150	50%
Expenditure in country B	124	124	150	150	50%

Clearly there is greater poverty in country A, but the headcount index does not capture this. As a welfare function, the headcount index is unsatisfactory in that it violates the transfer principle – an idea first formulated by Dalton (1920) that states that transfers from a richer to a poorer person should improve the measure of welfare. Here if a somewhat poor household were to give to a very poor household, the headcount index would be unchanged, even though it is reasonable to suppose that poverty overall has lessened.

Some argue that if it is to meaningful, the headcount index should imply that there is a "jump" or discontinuity in the distribution of welfare at about the poverty line, so it makes sense to speak of the poor and the non-poor. In practice, such a jump is not found (Ravallion 1996, p.1330).

Second, the head-count index does not indicate how poor the poor are, and hence does not change if people below the poverty line become poorer. Moreover, the easiest way to reduce the headcount index is to target benefits to people just below the poverty line, because they are the ones who are cheapest to move across the line. But by most normative standards, people just below the poverty line are the least deserving of the poor.

Third, the poverty estimates should be calculated for individuals and not households. If 20% of households are poor, it may be that 25% of the population is poor (if poor households are large) or 15% are poor (if poor households are small); the only relevant figures for policy analysis are those for individuals.

But survey data are almost always related to households, so in order to measure poverty at the individual level we must make a critical assumption that all members of a given household enjoy the same level of well-being. This assumption may not hold in many situations. For example, some elderly members of a household, or girls, may be much poorer than other members of the same household. In reality, not all consumption is evenly shared across household members.

4.2 Poverty gap index

A moderately popular measure of poverty is the *poverty gap index*, which adds up the extent to which individuals on average fall below the poverty line, and expresses it as a percentage of the poverty line. More specifically, define the poverty gap (G_i) as the poverty line (z) less actual income (y_i) for poor individuals; the gap is considered to be zero for everyone else. Using the index function, we have

(4.3) $$G_i = (z - y_i).I(y_i < z).$$

Then the poverty gap index (P_1) may be written as

(4.4) $$P_1 = \frac{1}{N}\sum_{i=1}^{N}\frac{G_i}{z}.$$

This table shows how the poverty gap is computed, divided by the poverty line, and averaged to give P_1, the poverty gap index.

Calculating the Poverty Gap Index, assuming poverty line of 125					
	Expenditure for each individual in country				Poverty Gap Index (P_1)
Expenditure in country C	100	110	150	160	
Poverty gap	25	15	0	0	
G_i/z	0.20	0.12	0	0	0.08 [= 0.32/4]

This measure is the mean proportionate poverty gap in the population (where the non-poor have zero poverty gap). Some people find it helpful to think of this measure as the cost of eliminating poverty (relative to the poverty line), because it shows how much would have to be transferred to the poor to bring their incomes or expenditures up to the poverty line (as a proportion of the poverty line). The minimum cost of eliminating poverty using targeted transfers is simply the sum of all the poverty gaps in a population; every gap is filled up to the poverty line. However this interpretation is only reasonable if the transfers could be made perfectly efficiently, for instance with lump sum transfers, which is implausible. Clearly this assumes that the policymaker has a lot of information; one should not be surprised to find that a very "pro-poor" government would need to spend far more than this in the name of poverty reduction.

At the other extreme, one can consider the maximum cost of eliminating poverty, assuming that the policymaker knows nothing about who is poor and who is not. From the form of the index, it can be seen that the ratio of the minimum cost of eliminating poverty with perfect targeting (i.e.G_i) to the maximum cost with no targeting (i.e. z, which would involve providing everyone with enough to ensure they are not below the poverty line) is simply the poverty gap index. Thus this measure is an indicator of the potential saving to the poverty alleviation budget from targeting: the smaller is the poverty gap index, the greater the potential economies for a poverty alleviation budget from identifying the characteristics of the poor – using survey or other information – so as to target benefits and programs.

The poverty gap measure has the virtue that it does not imply that there is a discontinuity ("jump") at the poverty line. To see this, consider the following example:

Poverty Gap Poverty Rates in A and B, assuming poverty line of 125						
	Expenditure for each individual in country				Poverty gap rate (P_1)	Headcount index (P_0)
Expenditure in country A	99	101	150	150	0.10	50%
Expenditure in country B	79	121	150	150	0.10	50%

For both of these countries, the poverty gap rate is 0.10, but most people would argue that country B has more serious poverty because it has an extremely poor member. Alternatively, one could think of the distribution in A as being generated from that in B by transferring 20 from the poorest person to the next poorest person – hardly an improvement in most people's eyes, yet one that has no effect on the poverty gap rate!

4.3 Squared poverty gap ("poverty severity") index

To construct a measure of poverty that takes into account inequality among the poor, some researchers use the squared poverty gap index. This is simply a weighted sum of poverty gaps (as a proportion of the poverty line), where the weights are the proportionate poverty gaps themselves; a poverty gap of (say) 10% of the poverty line is given a weight of 10% while one of 50% is given a weight of 50%; this is in contrast with the poverty gap index, where they are weighted equally. Hence, by squaring the poverty gap index, the measure implicitly puts more weight on observations that fall well below the poverty line. Formally:

(4.5)
$$P_2 = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{G_i}{z}\right)^2.$$

This table shows how the poverty gap is computed, divided by the poverty line, squared, and averaged to give P_2, the squared poverty gap index.

Calculating the Squared Poverty Gap Index, assuming poverty line of 125					
	Expenditure for each individual in country				Squared Poverty Gap Index (P_2)
Expenditure in country C	100	110	150	160	
Poverty gap	25	15	0	0	
G_i/z	0.20	0.12	0	0	
$(G_n/z)^2$	0.04	0.0144	0	0	0.0136 [= 0.0544/4]

The measure lacks intuitive appeal, and because it is not easy to interpret it is not used very widely. It may be thought of as one of a family of measures proposed by Foster, Greer and Thorbecke (1984), which may be written, quite generally, as

(4.6)

$$P_\alpha = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{G_i}{z}\right)^\alpha, \qquad (\alpha \geq 0)$$

where α is a measure of the sensitivity of the index to poverty and the poverty line is z, the value of expenditure per capita for the i-th person's household is x_i, and the poverty gap for individual i is $G_i=z-x_i$ (with $G_i=0$ when $x_i>z$) When parameter $\alpha=0$, P_0 is simply the head-count index. When $\alpha=1$, the index is the poverty gap index P_1, and when α is set equal to 2, P_2 is the poverty severity index. For all $\alpha > 0$, the measure is strictly decreasing in the living standard of the poor (the lower your standard of living, the poorer you are deemed to be). Furthermore, for $\alpha > 1$ it also has the property that the increase in measured poverty due to a fall in one's standard of living will be deemed greater the poorer one is. The measure is then said to be "strictly convex" in incomes (and "weakly convex" for $\alpha=1$). Another convenient feature of the FGT class of poverty measures is that they can be disaggregated for population sub-groups and the contribution of each sub-group to national poverty can be calculated.

Although the Foster, Greer and Thorbecke measure provides an elegant unifying framework for measures of poverty, it leaves unanswered the question of what is the best value of α. Moreover some of these measures also lack emotional appeal.

The measures of poverty depth and poverty severity provide complementary information on the incidence of poverty. It might be the case that some groups have a high poverty incidence but low poverty gap (when numerous members are just below the poverty line), while other groups have a low poverty incidence but a high poverty gap for those who are poor (when relatively few members are below the poverty line but with extremely low levels of consumption). Table 4.1 provides an example from Madagascar. According to the headcount measure (P_0), unskilled workers show the third highest poverty rate, while the group is in the fifth rank according to the poverty severity index (P_2). Compared to herders, they have a higher risk of being in poverty, but their poverty tends to be less severe. The types of interventions needed to help the two groups are therefore likely to be different.

Table 4.1: Poverty Indices By sub-groups, Madagascar, 1994

	Head count: % P_0	Rank	Poverty gap: % P_1	Rank	Poverty severity: × 100 P_2	Rank
Small farmers	81.6	1	41.0	1	24.6	1
Large farmers	77.0	2	34.6	2	19.0	2
Unskilled workers	**62.7**	**3**	**25.5**	**4**	**14.0**	**5**
Herders/fishermen	51.4	4	27.9	3	16.1	3
Retirees/handicapped	50.6	5	23.6	5	14.1	4

Source: Coudouel, Hentschel and Wodon (2001)

4.4 Sen Index.

Sen (1976) has proposed an index that sought to combine the effects of the number of poor, the depth of their poverty, and the distribution of poverty within the group. The index is given by

(4.7) $$P_s = P_0(1-(1-G^P)\frac{\mu^P}{z}),$$

where P_0 is the headcount index, μ^P is the mean income (or expenditure) of the poor, and G^P is the Gini coefficient of inequality among the poor. The Gini coefficient ranges from 0 (perfect equality) to 1 (perfect inequality), and is discussed in chapter 5 in the context of measuring inequality. The Sen Index can also be written as the average of the headcount and poverty gap measures, weighted by the Gini coefficient of the poor, giving:

(4.8) $$P_s = P_0 G^P + P_1(1-G^P)$$

It can be shown (Osberg and Xu 2002) that the Sen Index may also be written as

(4.9) $$P_S = P_0 P_1^P (1+G^{PP}),$$

where G^{PP} is the Gini coefficient of the poverty gap ratios of only the poor and P_1^P is the poverty gap index *calculated over poor individuals only*.

The Sen index has been widely discussed, and has the virtue of taking the income distribution among the poor into account. However the index is almost never used outside of the academic literature, perhaps because it is lacks the intuitive appeal of some of the simpler measures of poverty, but also because it "cannot be used to decompose poverty into contributions from different subgroups" (Deaton, 1997, p.147).

4.5 The Sen-Shorrocks-Thon index.

The Sen index has been modified by others, and perhaps the most compelling version is the Sen-Shorrocks-Thon (SST) index, defined as

(4.10) $$P_{SST} = P_0 P_1^P (1+\hat{G}^P),$$

which is the product of the headcount index, the poverty gap index (*applied to the poor only*), and a term with the Gini coefficient of the poverty gap ratios (i.e. of the G_n's) *for the whole population*. This Gini coefficient typically is close to 1, indicating great inequality in the incidence of poverty gaps.

Example. In 1996, 12.4% of the population of Quebec province (Canada) was in poverty. The poverty gap index, *applied to the poor only*, stood at 0.272. And the Gini coefficient of the poverty gap ratios was 0.924. Thus the Sen-Shorrocks-Thon index was 0.065 (=0.124 × 0.272 × (1+0.924)).

Application. Osberg and Xu (1999) use the SST index to compare poverty across the 10 Canadian provinces for 1984, 1989, 1994, 1995 and 1995, as well as to put the degree of Canadian provincial poverty into an international context. A number of graphs from their study are reproduced below. Figure 4.1 provides an international comparison, using the SST index, and shows that the US is an outlier with its relatively high poverty rate (as measured by the SST). A comparison of the US and Canada over time (figure 4.2) shows that while poverty was similar in the two countries a generation ago, it is now clearly higher in the US than in Canada. Figure 4.3 provides information on some Canadian provinces: Newfoundland was the poorest in 1984, but by 1996 had become much less of an outlier.

One strength of the SST index is that it can help give a good sense of the sources of change in poverty over time. This is because the index may be decomposed into

(4.11) $$\Delta \ln P_{SST} = \Delta \ln P_0 + \Delta \ln P_1^P + \Delta \ln(1+\hat{G}^P),$$

which may be interpreted as, % change in SST index = % change in headcount index + % change in poverty gap index (among poor) + % change in (1+Gini coefficient of poverty gaps).

In plain English, this allows us to decompose poverty into three aspects: are there more poor? are the poor poorer? and is there higher inequality among the poor?

FIGURE 1

1990s – Country Rankings by SST Poverty Index – Poverty Line Based on Half the Median Equivalent Income

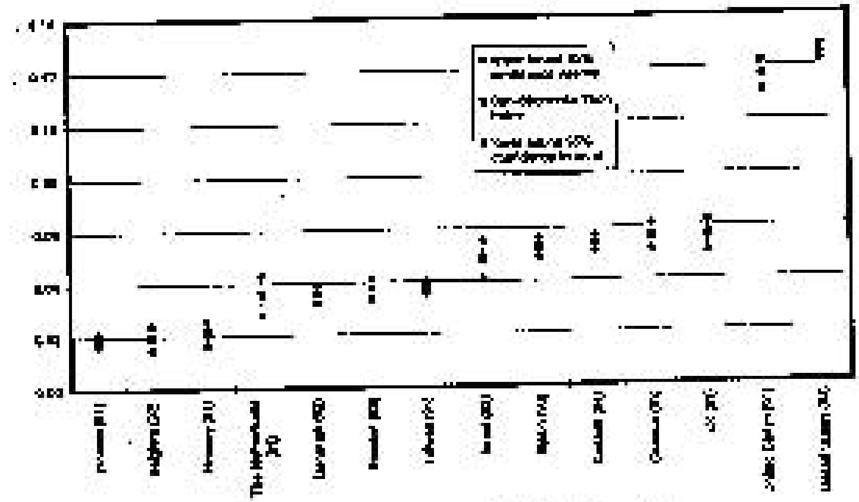

Note: [95% confidence interval = mean +/-2 standard deviations] of 500 bootstraps.
Source: Luxembourg Income Study, Osberg and Xu (1997).

FIGURE 2

Canada–United States Comparison – The SST Index from 1971 to 1994 – Poverty Line Based on Half the Median Equivalent Income

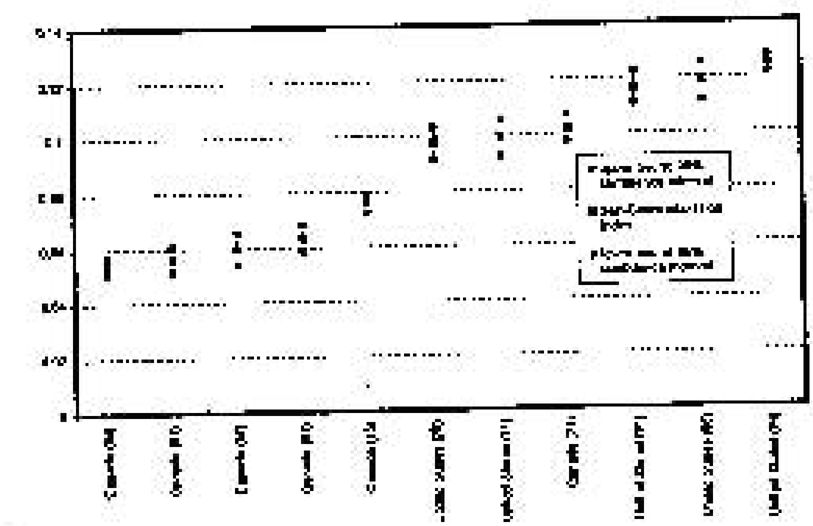

Note: [95% confidence interval = mean +/-2 standard deviations] of 500 bootstraps.
Source: Luxembourg Income Study, Osberg and Xu (1997).

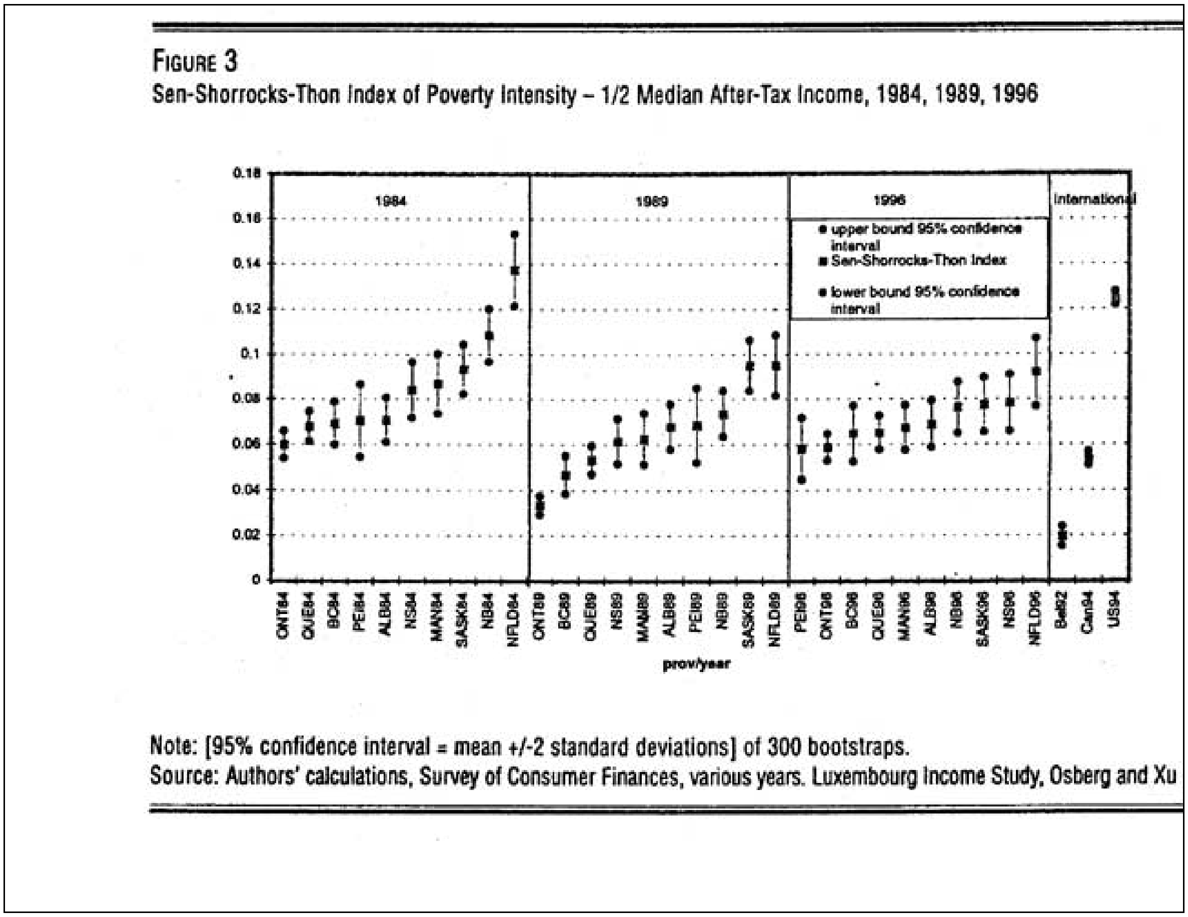

Example. The information in table 4.2 comes from Osberg and Xu, and traces the evolution of poverty in the Canadian province of Newfoundland between 1984 and 1996. It is clear that most of the change in the poverty rate over time was due to variations in the number of people in poverty (P_1), rather than in the size of the poverty gap per poor person (P_1^P) or the distribution of poverty among the poor (G^P).

Table 4.2: Decomposition of poverty, and changes in poverty, in Newfoundland, 1984-1996

	SST index	P_0	P^P_1	$1+G^P$	ΔlnSST index	ΔLnP_0	ΔlnP^P_1	$\Delta ln(1+G^P)$
1984	.137	.245	.304	1.844				
1989	.095	.169	.296	1.897	-.370*	-.372*	-.027	.028
1994	.105	.184	.304	1.884	.104	.086	.026	-.007
1995	.125	.212	.316	1.864	.168	.141	.038	-.010
1996	.092	.164	.294	1.897	-.307	-.254	-.071	.018

Notes: * denotes statistically significant at the 95% level. Poverty line is half of median equivalent income, using the "OECD scale" (i.e. equivalent income = 1 + 0.7(N_{adults}-1)+0.5($N_{children}$).
Source: Osberg and Xu, 1999.

Note that the values of the Sen-Shorrocks-Thon index provided by Osberg and Xu do not give just a single point estimate for each province; they also provide a confidence interval. Because the SST index is complex, it is not possible to compute these confidence intervals analytically. Instead, they are computed artificially using *bootstrapping*. The basic idea behind the bootstrap is straightforward and clever. Suppose we have a survey sample of 2,000 households. Now pick a sample of 2,000 from this sample *with replacement* – i.e. pick a household, then put it back into the sample, pick another household, put it back into the sample, and so on, until you have picked 2,000 households. Some households will be chosen more than once, but that's fine. Now compute the SST index using this artificial sample. Then repeat the process many times; Osberg and Xu use 300 repetitions. The result is a distribution of values of the SST, from which it is easy to find (say) the 95% confidence interval. Sample Stata code to generate confidence intervals for the SST index is given in the Appendix, in the exercises associated with Chapter Five.

4.6 The Watts Index

The first distribution-sensitive poverty measure was proposed in 1968 by Watts (see Zheng 1993), and in its discrete version takes the form:

(4.13)
$$W = \frac{1}{N}\sum_{i=1}^{q}[\ln(z) - \ln(y_i)]$$

where the N individuals in the population are indexed in ascending order of income (or expenditure), and the sum is taken over the q individuals whose income (or expenditure) y_i falls below the poverty line z.

This table shows how the Watts index is computed, by dividing the poverty line by income, taking logs, and finding the average over the poor. The Watts index is attractive in that it satisfies all the theoretical properties that one would want in a poverty index, and is increasingly used by researchers in generating such measures as the poverty incidence curve (see chapter xxx).[6] However, it is not a particularly intuitive measure, and so is rarely seen in practical field work.

Calculating the Watts Index, assuming poverty line of 125					
	Expenditure for each individual in country				Watts Index
Case 1 (poor)					
Expenditure in country C	100	110	150	160	
z/y_i	1.25	1.14	0.83	0.78	
log (z/y_i)	0.223	0.128	-0.182	-0.247	**0.351**
Case 2 (less poor)					
Expenditure in country C	110	120	150	160	

[6] Ravallion and Chen (2001) argue that three axions are essential to any good measure of poverty. Under the focus axiom the measure should not vary if the income of the non-poor varies); under the monotonicity axiom, any income gain for the poor should reduce poverty; and under the transfer axiom, inequality-reducing transfers among the poor should reduce poverty. The Watts index satisfies these three axioms, but the headcount (P_0) and poverty severity (P_1) measures do not.

z/y$_i$	1.14	1.04	0.83	0.78	
log (*z/y*$_i$)	0.128	0.041	-0.182	-0.247	**0.169**
Case 3 (deeper poverty)					
Expenditure in country C	90	120	150	160	
z/yi	1.25	1.10	0.83	0.78	
log (z/yi)	0.329	0.041	-0.182	-0.247	**0.369**

4.7 Time taken to exit

Most poverty profiles for Cambodia, and indeed for most countries, rely on the three basic classes of Foster Greer Thorbecke poverty statistics discussed above. But when thinking about poverty reduction strategies, it may be useful to show how long it would take, at different potential economic growth rates, for the average poor person to exit poverty. A poverty statistic with this property is derived by Morduch (1998); the statistic is decomposable by population sub-groups and is also sensitive to how expenditure (or income) is distributed among the poor. For the *j*th person below the poverty line, the expected time to exit poverty (i.e., to reach the poverty line), if consumption per capita grows at positive rate *g* per year is:

(4.12)
$$t_g^j \approx \frac{\ln(z) - \ln(x_j)}{g} = \frac{W}{g}.$$

In other words, the time take to exit is the same as the Watts index divided by the expected growth rate of income (or expenditure) of the poor.

What effect can economic growth have on the elimination of poverty? Figure 4.4 shows the average time it would take to raise the consumption level of a poor person in Cambodia to the poverty line, for various hypothetical growth rates. It is assumed that this growth rate is continuous, is in real terms, and is distributionally neutral *among* the poor. If the economic growth rate enjoyed by the poor were only one percent per year, it would take over 20 years for the average poor person to exit poverty. But at a growth rate of four percent per year it would take less than six years for the average poor person to exit poverty. Hence, economic growth that acts to raise the real consumption levels of the poor can have a powerful effect on the elimination of poverty.

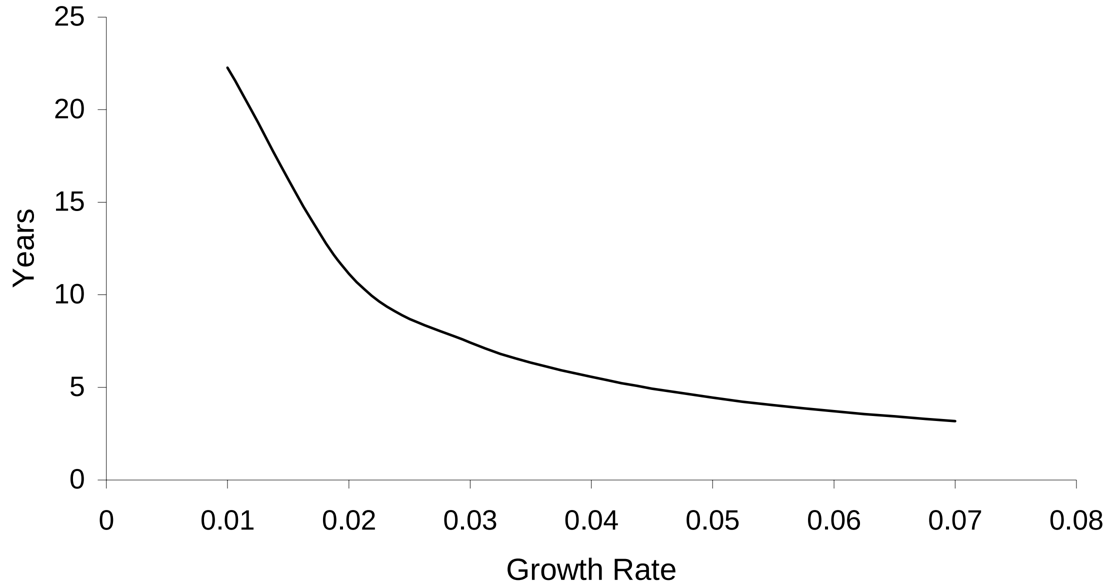

Figure 4.4: ***Average exit time from poverty***

Despite the potency of economic growth, it will generally take more than just growth to rapidly improve the lives of the very poor. The expected time to exit poverty for those people who are so poor that they are below the food poverty line in Cambodia – i.e. they cannot afford enough food, even if they were to devote all their consumption spending to food – is more than 15 years, even at a three percent continuous annual growth rate. Thus, targeted programs are needed to deliver benefits to the poor, for instance in the form of improvements in their human and physical assets or through interventions (e.g., infrastructure, markets) that improve the returns they get from those assets.

4.8 Other Measures

There are other additive poverty measures that are distribution sensitive. Following Atkinson (1987), one can characterize a general class of additive measures, encompassing W, the FGT (Foster, Greer and Thorbecke) class of measures, and some other measures (such as the second measure proposed by Clark, Hemming and Ulph, 1981), as taking the following form:

(4.14)
$$P = \frac{1}{N}\sum_{i=1}^{N} p(z, y_i)$$

where p(z, yi) is the individual poverty measure, taking the value zero for the non-poor (y_i>z) and some positive number for the poor, the value of which is a function of both the poverty line and the individual living standard, non-decreasing in the former and non-increasing in the latter.

Given the wide variety of aggregate measures of poverty that are available, which ones should one use? We turn to this question in chapter 5.

References

Vu Tuan Anh. 1997. "Criteria for Rural Poverty Monitoring," *Vietnam's Socio-Economic Development*, Spring.

Atkinson, Anthony. 1987. "On the measurement of poverty," *Econometrica*, 55: 749-764.

Clark, Hemming and Ulph. 1981.

Dalton, Hugh. 1920. "The measurement of the inequality of incomes," *Economic Journal*, 30: 384-361.

Deaton, Angus. 1997. *The Analysis of Household Surveys*, Johns Hopkins University Press, Baltimore MD.

Foster, James, J. Greer and Eric Thorbecke. 1984. "A class of decomposable poverty measures," *Econometrica*, 52(3): 761-765.

Morduch, Jonathan. 1998. "Poverty, economic growth, and average exit time," *Economics Letters*, 59: 385-390.

Lars Osberg and Kuan Xu. 1999. *Poverty Intensity: How Well Do Canadian Provinces Compare?" Canadian Public Policy/Analyse de Politiques*, XXV(2) :

Osberg and Xu. 2002.

Martin Ravallion. 1996. "How well can method substitute for data? Five experiments in poverty analysis," *The World Bank Research Observer*, 11(2), 199-221, August.

Ravallion, Martin and Shaohua Chen. 2001. *Measuring Pro-Poor Growth*, Development Research Group, World Bank.

Sen, Amartya. 1976. "Poverty: an ordinal approach to measurement," *Econometrica*, 46: 437-446.

Zheng, B. 1993. "An axiomatic characterization of the Watts poverty index," *Economics Letters*, 42(1): 81-86.