UCD CENTRE FOR ECONOMIC RESEARCH

WORKING PAPER SERIES

2008

Health and Income Poverty in Ireland, 2003-2006

David Madden,
University College Dublin

WP08/15

July 2008

**UCD SCHOOL OF ECONOMICS
UNIVERSITY COLLEGE DUBLIN
BELFIELD DUBLIN 4**

Health and Income Poverty in Ireland, 2003-2006

David Madden

(University College Dublin)

July 2008

Abstract: Recent advances in the measurement of bi-dimensional poverty are applied to a measure of poverty which incorporates income and health poverty. The correlation between income and poverty is examined using the Receiver Operating Characteristics curve. Following from this uni-dimensional and bi-dimensional poverty indices are calculated for Ireland for the years 2003-2006. Individual and bi-dimensional indices generally show a decline over the period with the biggest decline between 2003 and 2004. The results are generally not sensitive to the degree of poverty aversion or the substitutability between the different dimensions of poverty.

Keywords: receiver operating characteristic, multidimensional poverty.

JEL Codes: I12, I31, I32.

Corresponding Author: David Madden,
School of Economics,
University College Dublin,
Belfield,
Dublin 4,
Ireland.

Phone: 353-1-7168396
Fax: 353-1-2830068
E-Mail: david.madden@ucd.ie

Health and Income Poverty in Ireland, 2003-2006.

1. Introduction

Recent developments in the measurement of multidimensional poverty reflect the fact that poverty (and by corollary welfare) can be viewed as occurring in a number of different dimensions, apart from the most typically used ones of income or expenditure.[1] For example, individuals may experience poverty with respect to housing or other assets, education, nutrition or health as well as income. Empirical studies in multidimensional poverty have so far concentrated upon income poverty in association with poverty in education (Bourguignon and Chakravarty, 2003), height-for age (Duclos et al, 2006), asset poverty (Duclos et al, 2006) and survival probability (Duclos et al, 2006). The choice of poverty dimension is partly motivated by the nature of the data, as multidimensional measures typically work best with continuous variables. The approach to multidimensional poverty analysis can also differ with some authors choosing to calculate multidimensional poverty indices (Bourguignon and Chakravarty, 2003) and others looking for more robust multidimensional poverty orderings for broader classes of measures (Duclos et al, 2006).

The development of multidimensional poverty indices is partly motivated by the recognition that deprivation can occur in more than one dimension and that the correlation between different dimensions of deprivation may not be perfect.[2] Thus in the first part of this paper we empirically examine the correlation between income poverty and poverty in health, where we use a cardinal index of health derived from an ordinal measure of self-assessed health. We examiner the correlation by analysing the receiver operating characteristic (ROC) curve for the two dimensions. This approach provides a summary of the degree to which poverty in one dimension acts as a signal for poverty in the other dimension. Following on from this we then calculate multidimensional poverty indices for Ireland over the period 2003 to 2006.

---

[1] For recent work in this area see Atkinson (2003), Bourguignon and Chakravarty (2003), Duclos et al (2006) and Tsui (2002).

In the next section we briefly outline the approach lying behind the ROC curve. We also describe our data and present results on the degree of correlation between the two dimensions of poverty. The following section then briefly reviews some of the key issues in the measurement of multidimensional poverty before presenting results for bi-dimensional poverty indices.

2. The Correlation Between Different Dimensions of Poverty

As outlined above, one of the justifications for taking a multidimensional approach to poverty is that well-being (and deprivation) may have many dimensions (income, consumption, literacy, health etc) and that the correlation between poverty in one dimension (say income) and another dimension (say health) is unlikely to be perfect. By "poverty" here we mean that in the case of a continuous variable an individual is below some critical threshold, or in the case of a categorical variable an individual is in some critically identified category or categories.[3] If the correlation between all conceivable dimensions of poverty is perfect, then the choice of dimension upon which to measure the incidence of poverty is unimportant, since all dimensions will identify the same individuals as poor. In this case there would seem to be little advantage in moving from a uni-dimensional to a multidimensional measure, at least from the point of view of identifying the poor. However, as the degree of correlation between different dimensions of poverty decreases, then the case for a multidimensional approach becomes more persuasive. In the absence of perfect correlation then different dimensions will identify different individuals as being poor (though there is likely to be some, perhaps considerable, overlap). Reliance on a uni-dimensional measure such as income may then be misleading as it will ignore individuals who are not income poor, even though they may be poor in another dimension.

---

[2] For example, the European Union has adopted a common set of social indicators which includes measures such as income inequality, long-term unemployment and poor health, but as yet they have not been aggregated into a single index.

[3] In this paper we will refer to individuals as our data is individual-based. However most individuals are part of households and depending upon within-household distribution of resources there may be a divergence between poverty on an individual and household basis.

Ultimately the degree of correlation between different dimensions of poverty is an empirical issue. By correlation here we mean the extent to which income poverty overlaps with health poverty i.e. the extent to which one acts as an indicator for the other. One possible way of examining this is to define a poverty threshold for income and one for whatever measure of health is employed (we confine ourselves to bivariate comparisons though multivariate analysis is possible). We can then examine the proportion of people classified as poor under both thresholds and compare them to those classified as poor under only one threshold and those not classified as poor under either threshold. The results obtained however may be sensitive to the threshold chosen. To assess the degree of overlap in a way which is independent of the health poverty threshold we adopt the approach of Fusco (2006) who uses the ROC curve, a technique commonly used in disease diagnosis.[4]

The ROC curve is a procedure in signal detection theory which originated in the Second World War when it was used to recognise radar and sonar signals which were affected by noise. It provides a useful procedure for analysing the extent to which a given signal can detect an underlying condition. In the application here, the income poverty line initially partitions the population into the binary categories of poor and non-poor.[5] We then assess the degree to which the health poverty line would produce the "same" partition. If the health poverty measure assigns someone as poor who is also poor under the income poverty definition then this is called a "true positive" (TP). If it signals someone as poor who is not poor under the income definition it is a "false positive" (FP). If it signals someone as non-poor even though they are poor under the income definition it is a "false negative" (FN). Finally "true negatives" (TN) are those who are classified as non-poor under both definitions.[6]

The TP rate is sometimes called the sensitivity of the signal and is TP/(TP+FN), while the corresponding concept for the TN rate is known as specificity and is

---

[4] In using the ROC approach it is necessary that one of the thresholds be fixed. Since most uni-dimensional poverty studies are in the area of income we choose to fix the income poverty line.
[5] ROC curves have been used in poverty analysis by, among others, Wodon (1997), Baulch (2002) and in the application closest to this paper Fusco (2006).
[6] Note that since we are comparing different health measures with the one income indicator we are effectively assuming that income poverty is in some sense "true" poverty. This is an inevitable consequence of this type of analysis where some measure is by default regarded as the true underlying measure. Of course, income poverty itself may be subject to problems with measurement error etc.

TN/(FP+TN), which in turn is equal to one minus the FP rate. The ROC curve then graphs the TP rate (on the vertical axis) against the FP rate (one minus the specificity rate) for all possible values of the health poverty threshold. When the health poverty threshold takes on its lowest possible value then all people are above this threshold and hence no-one is considered health-poor and so none of the income-poor population are correctly identified (TP=0) while all of the income non-poor are correctly identified (TN=1 and hence FP=0). This is the point (0,0) in figure 1. As the health poverty threshold increases, some people will now be signalled as poor by the health measure. If they are also income poor, then the TP rate must increase and the curve traces up. However, if they are not income poor, they are regarded as FPs and so the curve traces to the right. When the health poverty threshold is at its highest then the whole population will be signalled as poor. Since all the income-poor are now signalled as poor the TP rate equals one, but since the non income-poor are also signalled as poor the FP rate also equals one and so we are at the (1,1) point.

Thus as the threshold goes from its lowest to its highest level the ROC curve traces out from (0,0) to (1,1) and the better the signal the further above and to the left (or north-west) of the 45^0 line will be the curve. The less accurate the signal the nearer the curve will be to the 45^0 line. If the curve lies below the 45^0 line then it is effectively acting as a contra-indicator and paradoxically the further to the south-east the curve lies the better, since the ROC curve for the negative of the indicator is simply the mirror image of the ROC curve for the original indicator.

Clearly if one ROC curve for one indicator always lies above and to the left of that of another then the former indicator acts as a better signal for all values of the threshold and can be said to "dominate". Just like Lorenz curves in inequality analysis however, there is no guarantee that dominance will be found when comparing any two indicators. In that case a summary index may be used. Probably the most popular one is the area under the ROC curve. If the ROC curve lies on the 45^0 line then this area equals 0.5 and this corresponds to the situation where the indicator effectively gives no signal. If the ROC curve corresponds to the vertical line from (0,0) to (0,1) and then across to (1,1) the area under the curve is one and the indicator gives a perfect signal. Intuitively the area under the curve corresponds to

the probability that health poverty for a randomly chosen income poor person is higher than the health poverty for a (randomly chosen) non income-poor person.

Before presenting results for the areas under the ROC curves we first discuss our data. Our data comes from four consecutive cross-sectional surveys which are the Irish part of the European Union Survey of Income and Living Conditions (EU-SILC).[7] This survey is the successor to the European Community Household Panel survey. After allowing for missing observations for certain variables the sample sizes are between 11,000 and 12,000 for each year. However, in Ireland there was only six months data collection for 2003 (as opposed to 12 month collection for the other years) hence the sample size for 2003 is only about half of that for the other years (see CSO, 2007).

As our income measure we use equivalised income after social transfers, using the EU definition of income (details of this measure are included in the appendix) and the modified OECD equivalence scale (1.0 for first adult, 0.5 for subsequent adults and 0.3 for children aged less than 14).

The cardinal health index we use is based on responses to a question concerning self-assessed health. The self assessed health question asks: in general, how good would you say your health is? The possible answers are: very bad, bad, fair, good and very good. While this measure appears to give a good indicator for overall health (Idler and Benyamini, 1997) it is not cardinal, and with only five categories, it may be difficult to find a plausible health poverty threshold. Various attempts have been made to translate this ordinal measure into a cardinal one (for a review, see van Doorslaer and Jones, 2003). We choose to estimate an ordered probit of self-reported health using a variety of plausible independent variables such as age, gender, education, marital status and principal economic status.[8] We then take the linear prediction of this ordered probit and re-scale it so that it takes a value from zero to one. This provides us with a cardinal measure of health and hence a wide range of possible health poverty thresholds.

---

[7] For details of the Irish part of EU-SILC see CSO (2007) and the documentation at http://www.cso.ie/eusilc/default.htm

As our poverty line we follow what is now relatively standard procedure and choose 60% of median income. We then investigate the extent to which the income poverty line acts as a signal for health poverty. In table 1 we present results for both a relative income line (i.e. 60% of median income for each year) and an absolute line (where we hold the poverty line fixed at 60% of 2003 median income).[9] The results show that the area under the ROC curve was about 0.73 in 2003 i.e. for any random income poor person there is a 0.73 probability that they will also be health poor. By 2006, this probability has fallen to 0.69. The decline in the area under the ROC curve is even more pronounced when the fixed income poverty lines is used. This is slightly surprising as the fixed income poverty line is lower than the relative income poverty line for 2004-2006, owing to income growth. Thus those who are identified as poor using the fixed income poverty line have on average *lower* incomes than those identified as poor under the relative income poverty line, and so we might expect a higher probability that a random poor person (as defined by a fixed income line) would also be health poor. But that does not appear to be the case. This indicates that health poverty is more pronounced among those who are just above the fixed income poverty line, but just below the relative income line.

Thus overall, the results for the ROC curve analysis, with areas under the curve of around 0.7, suggest that income poverty is a good but far from perfect indicator of health poverty. At the least, there appears to be sufficient justification to analyse multidimensional poverty measures incorporating income and health poverty, which we turn to in the next section.

3. Measuring Multidimensional Poverty

Extensive discussions of multidimensional poverty indices and multidimensional poverty orderings are available elsewhere in the papers referred to above. Many of the standard "desirable" properties of uni-dimensional poverty indices such as symmetry, replication invariance and monotonicity translate in a straightforward

---

[8] To economise on space we do not include details of the ordered probit regressions, but they are available on request.
[9] For a discussion of absolute and relative poverty lines see Madden (2000).

manner to the multidimensional case. Here we briefly discuss some key properties where the translation is not so clearcut.

One important issue when dealing with a multidimensional poverty index is whether the poor are identified as those who are poor in *any one* dimension of poverty (the so-called *union* approach) or those who are poor in *all* dimensions of poverty (the so-called *intersection* approach). Clearly the latter is a more restrictive condition for identification. In this paper we present results for both definitions.

The choice between union or intersection approaches can depend upon the actual dimensions of poverty which are chosen. For example, if the two dimensions are income and longevity, then it seems reasonable to define someone as poor if they are below the income threshold though above the longevity threshold. It is probably less reasonable if the dimensions are income and housing conditions, since presumably if someone is far enough above the income threshold, this begs the question of why they do not use this income to improve their housing conditions.

Clearly the extent to which a surplus in one dimension can compensate for a deficit in another dimension is crucial here and this issue of *substitutability* between poverty dimensions returns below in our discussion of transfer principles in a multidimensional setting. Substitutability is also central to the assumption of *focus* in a multidimensional setting, since it assumes that if the jth attribute of a poor person exceeds the poverty threshold for that dimension then giving them more of that attribute does not affect the level of poverty.

Following the contribution of Sen (1976), it has become standard for poverty measures to take account of the *distribution* of attributes among the poor. Thus poverty indices should not fall (rise) following a Pigou-Dalton regressive (progressive) transfer between two poor people. The analogous condition in a multidimensional setting builds upon the multidimensional transfers principle of Kolm (1977). This property holds that if we have a distribution of a set of attributes summarised by a matrix X, then this is more equal than that of another matrix Y if and only if $X=BY$ where B is a bi-stochastic matrix (and not a permutation matrix). Effectively what is happening here is that the original bundle of attributes in Y is

8

being replaced by a convex combination of them in *X*. Following from this Tsui (2002) introduced the multidimensional transfer principle whereby there is no more poverty in distribution *X* than in distribution *Y* if *X* is obtained from *Y* by re-distributing the attributes of the poor according to the bistochastic transformation.

There is another critical aspect to transfers in a multidimensional setting which brings us back to the substitutability issue which arose earlier. Suppose we have two people *i* and *t* in two-dimensional poverty with attributes *j* and *k*. Suppose *i* has more of *k* but less of *j*. If we interchange the amounts of *j* between the two people then the person who had more of *k* now has more of *j* as well and thus there is an increase in the correlation between the attributes. The effect of such a switch on poverty will depend upon the extent to which the attributes correspond to similar or different aspects of poverty. What Bourguignon and Chakravarty (2003) term the *non-decreasing poverty under correlation increasing switch* postulate (NDCIS) says that poverty must not increase following such a correlation increasing switch. The converse property is NICIS.

The property of NDCIS implies that poverty should not fall following a correlation increasing switch if the attributes involved in the switch are substitutes. In this case substitutability is defined in what Atkinson (2003) terms the Auspitz-Lieben-Edgeworth-Pareto (ALEP) sense as opposed to the perhaps more common sense suggested by Hicks. If attributes are substitutes then the marginal utility of one attribute decreases when the quantity of the other increases. Thus if we have a poverty function $\pi(x;z)$ where *x* is the vector of attributes and *z* is the vector of poverty threshold levels, then, presuming this function is twice differentiable, two attributes *j* and *k* are substitutes whenever $\pi^{jk}(x;z) > 0$ for all *x*. Thus say the two dimensions of poverty are income and health, the fall in poverty due to a unit increase in income is less important for people with health close to the health poverty threshold as opposed to those with very poor health. The drop in poverty is larger for those with health close to the health threshold if the two attributes are complements i.e. $\pi^{jk}(x;z) < 0$.

Arguments in favour of the NDCIS property could be made along the lines that resources should be directed towards those with multiple deprivation even though reaching those individuals may be more costly. Thus it is more desirable to improve the incomes of those in very poor health, as opposed to those in "only" marginally poor health. However, it is also possible to make arguments in favour of NICIS e.g. if the two dimensions of poverty are nutrition and education then what we could term the "lifeboat ethic" suggests that education resources should be directed at those best equipped to take advantage of them i.e. those who are nutritionally less deprived. This can be particularly true if there are increasing returns to scale in poverty reduction for certain attributes over certain ranges and may lie behind the apparently unequal division of resources within some very poor households (Stiglitz, 1976). Tsui (2002) argues in favour of NDCIS while Duclos et al (2006) in their analysis of dominance criteria point out that NICIS limits the scope for "poverty-frontier robust" orderings.

Bourguignon and Chakravarty (2003) provide a detailed discussion of various functional forms which can give rise to different poverty indices. We concentrate on indices which satisfy the multidimensional transfer principle (MTP) referred to above. As suggested by Bourguignon and Chakravarty (2003) the individual poverty function in the two-dimensional case, $\pi(x;z)$, can be represented by the following general functional form:

$$\pi(x;z) = I[Max(1-\frac{x_1}{z_1},0), Max(1-\frac{x_2}{z_2},0)]$$

where x_i and z_i represent the values and poverty threshold levels of attribute i respectively and $I(u_1,u_2)$ is an increasing, continuous, quasi-concave function with $I(0,0)=0$. The poverty index for a population of size n then becomes

$$P(X;z) = \frac{1}{n}\sum_{i}^{n} I[Max(1-\frac{x_{i1}}{z_1},0), Max(1-\frac{x_{i2}}{z_2},0)].$$

One possible specification for the *I(.)* function is the CES form which gives a poverty index of the following form:

$$P(X;z) = \frac{1}{n}\sum_{i=1}^{n} f\left[\left\{.[Max(1-\frac{x_{i1}}{z_1},0)]^{\theta} + b^{\frac{1}{\theta}}.[Max(1-\frac{x_{i2}}{z_2},0)]^{\theta}\right\}^{\frac{1}{\theta}}\right]$$

where *f(.)* is an increasing and convex function with *f(0)=0, b>0* reflects the relative weight attached to the two attributes and θ is a parameter which determines the elasticity of substitution between the shortfalls of the different attributes. Satisfaction of MTP requires that $\theta > 1$. However the issue of whether this function satisfies NDCIS or NICIS remains ambiguous. The particular version of this family of poverty indices which we calculate here is that where the *f(.)* function is obtained from the Foster-Greer-Thorbecke P_{α} index which gives

$$P_{\alpha}^{\theta}(X;z) = \frac{1}{n}\sum_{i=1}^{n}\left[\left\{.[Max(1-\frac{x_{i1}}{z_1},0)]^{\theta} + b^{\frac{\theta}{\alpha}}.[Max(1-\frac{x_{i2}}{z_2},0)]^{\theta}\right\}^{\frac{\alpha}{\theta}}\right]$$

where $\alpha > 0$ reflects the relative weight attached to the very poor.[10] When $\alpha = 0$ this measure corresponds to a multidimensional headcount, although care must be taken in the interpretation as it corresponds to the number of people who are poor in *at least* one attribute and thus is likely to be greater than either of the headcounts for individual attributes. When $\alpha = 1$, then the P_{α}^{θ} measure is a multidimensional poverty gap which is a form of average of the two individual gaps, with the precise form depending upon the values of *b* and θ. Whether this index satisfies NDCIS or NICIS depends upon the relative values of α and θ, with NICIS holding when $\alpha > \theta$.

The P_{α}^{θ} measure can be generalised to more than two attributes. However, if the formulation above is retained it implies that the elasticity of substitution between

---

[10] When α=0 the P_{α} index is effectively a headcount. When α=1 it becomes a poverty gap measure and when α>1 it becomes a weighted poverty gap measure with greater weight attached to the poorest.

any two attributes must be the same, which in turn implies that the poverty indices are NDCIS or NICIS for all pairs of attributes, which may not be desirable.

Table 2 first of all presents values of the unidimensional P_α poverty indices for income and health for both relative and fixed poverty lines. In both cases the poverty line is 60% of the median and for the fixed poverty line it is held at the relevant 2003 values. Table 3 presents the simple headcounts on the intersection and union basis i.e. those who were classified as poor in *both* dimensions (intersection) or poor in *either* dimension (union). Tables 4 to 6 then present values of the Bourguignon-Chakravarty index for various values of b, α and θ and for fixed and relative poverty lines.

The results in table 2 for income poverty show substantial declines in poverty for the fixed poverty line, consistent with the high economic growth in Ireland over that period. This is evident for all values of α but particularly for $\alpha=2$. For the case of the purely relative poverty line there is little change when $\alpha=0$ for the 2003-2005 period but 2006 sees a fall of a couple of per cent. For $\alpha>0$ we see steady falls in income poverty and once again the proportional fall is greatest when $\alpha=2$. This indicates that over the 2003-2006 period the fall in poverty was mainly concentrated amongst the poorest of the poor and that the distribution of income amongst the poor became more equal.

The situation with regard to health is quite different. In the first instance there is much less difference between the results for the fixed and relative poverty lines. Unlike the case of income, the median level of health (and hence the health poverty line) changes very little over the 2003-2006 period. This is not surprising given the nature of the health index and the method by which it is re-scaled so that each year the values must lie in the (0, 1) interval. The figures show a sharp drop in health poverty between 2003 and 2004 for all values of α. In 2005 and 2006 health poverty rises again and for the case of the fixed poverty line with $\alpha=0$ it regains its 2003 level. The increase in health poverty over the 2004-2006 period is not as sharp for higher values of α, suggesting that the observed rise in poverty arises more from greater numbers of

---

Higher values of α correspond to relatively greater weight being attached to the very poor. See Foster et al (1984).

individuals below the poverty line as opposed to a worsening situation for those people who are health poor.

The uneven pattern of health poverty over the 2003-2006 period mainly reflects an unusually sharp fall in 2004. In particular the variance of the cardinal health measure fell (from 0.046 in 2003 to 0.038 in 2004). This tightening of the distribution reduced the fraction of the population whose health fell below 60% of the median. In 2005 and 2006 the distribution returned to a shape more similar to that of 2003.

Tables 3 to 6 show various bi-dimensional indices. Since these indices are effectively weighted sums of the uni-dimensional indices in table 2 the pattern over the 2003-2006 period essentially reflects developments in the uni-dimensional indices. Table 3 shows the bi-dimensional indices for the case where $\alpha=0$, the pure headcount case. Here the choice of a parameter such as θ is not relevant and the choice between indices boils down to an intersection or union approach. The intersection approach shows a fall over the period, with a greater fall observed with fixed poverty lines, reflecting the fall in income poverty. The fall in the union based index is less pronounced, and in the case of the fixed poverty line the index rises between 2004 and 2006. This reflects the fact that while there is movement out of income poverty after 2004, there is movement back into health poverty. This is the counterpart of the fall in the area under the ROC curve in table 1, which shows a lower correlation between the two types of poverty.

Tables 4 to 6 show the Bourguignon-Chakravarty indices and the pattern over the period is mainly influenced by the value of b, the relative weight given to health versus income poverty. In all cases there is quite a sharp drop between 2003 and 2004, as poverty in both dimensions falls. For the years from 2004 to 2006 the indices move according to the value of b. When b=0.3, with a low weighting on health poverty then the bi-dimensional indices continue to fall, albeit not as sharply as between 2003 and 2004. When b=3, with the high weighting on health poverty the indices rise after 2004, though they do not regain their 2003 level. For the intermediate case of b=1 the indices stay fairly flat from 2004.

What about the sensitivity of the indices to α and θ? There is relatively little sensitivity to the value of α except for the case when b=3 i.e. a higher weight on health poverty. In this instance health poverty rises between 2005 and 2006 when α=1, but it stays level or falls slightly when α=2, indicating that the rise in health poverty was less concentrated amongst the poorest of the health poor. There is very little sensitivity of the indices to the value of θ.

4. Conclusion

This paper has calculated bi-dimensional poverty indices covering income and health for Ireland for the 2003-2006 period. First of all the correlation between income and health poverty was analysed by examining the area under the ROC curve. This analysis showed that income poverty was a good though not perfect indicator of health poverty and provided a justification for calculating indices which combined information about both dimensions of poverty.

The calculation of uni-dimensional and bi-dimensional poverty indices then showed that income poverty fell gradually over the period with a greater fall experienced amongst the very poor. Health poverty fell quite sharply in 2004 but then increased between 2004 and 2006 but did not regain its 2003 level. Once again relatively speaking the poorest amongst the health poor fared best. The bi-dimensional indices all showed a fall over the period though the degree of that fall was most influenced by the relative weights attached to income and health. The greatest fall was experienced between 2003 and 2004 with relatively little change after 2004. The results showed slight sensitivity to the relative weight attached to the very poor but very little sensitivity to the degree of substitutability assumed between the different dimensions of poverty.

References:

Atkinson, A., (2003): "Multidimensional Deprivation: Contrasting Social Welfare and Counting Approaches", *Journal of Economic Inequality*, Vol. 1, pp. 51-65.

Baulch, B., (2002): "Poverty Monitoring and Targetting Using ROC Curves: Examples from Vietnam", Institute of Development Studies, IDS Working Paper No. 161.

Bourguignon, F., and S. Chakaravarty (2003): "The Measurement of Multidimensional Poverty", *Journal of Economic Inequality*, Vol. 1, pp. 25-49.

CSO (2007): *EU Survey on Income and Living Conditions*. Stationery Office, Dublin.

Duclos, J.Y., D. Sahn and S. Younger (2006): "Robust Multidimensional Poverty Comparisons", *Economic Journal*, Vol. 116, pp. 943-968.

Foster, J., J. Greer and E. Thorbecke (1984): "A Class of Decomposable Poverty Measures", *Econometrica*, Vol. 52, pp. 761-766.

Fusco, A., (2006): "On the Overlap of Multidimensional and Income Poverty", prepared for 2006 Conference of the European Panel Users Network.

Idler, E., and Y. Benyamini (1997): "Self-rated Health and Mortality: A Review of 27 Community Studies", *Journal of Health and Social Behaviour*, Vol. 38, pp. 21-37.

Kolm, S., (1977): "Multidimensional Egalitarianisms", *Quarterly Journal of Economics*, Vol. 91, pp. 1-13.

Madden, D., (2000): "Relative or Absolute Poverty Lines: A New Approach", *Review of Income and Wealth*, June 2000, Vol. 46, pp. 181-200.

Sen, A., (1976): "Poverty: An Ordinal Approach to Measurement", *Econometrica*, Vol. 44, pp. 219-231.

Stiglitz, J., (1976): "The Efficiency Wage Hypothesis, Surplus Labour and the Distribution of Income in L.D.C.'s", *Oxford Economic Papers*, Vol. 28, pp. 185-207.

Tsui, K., (2002): "Multidimensional Poverty Indices", *Social Choice and Welfare*, Vol. 19, pp.69-93.

Van Doorslaer, E., and A. Jones (2003): "Inequalities in Self-Reported Health: Validation of a New Approach to Measurement", *Journal of Health Economics*, Vol. 22, pp. 61-87.

Wodon, Q., (1997): Targeting the Poor using ROC Curves", *World Development*, Vol. 25, pp. 2083-2092.

Figure 1: ROC Curve

Table 1: Area Under ROC Curve, 2003-2006 (SE in brackets)

	Relative Income Poverty Line	Absolute Income Poverty Line
2003 (N=6107)	0.7312 (0.0083)	0.7312 (0.0083)
2004 (N=11069)	0.7260 (0.0063)	0.7031 (0.007)
2005 (N=11950)	0.7129 (0.0063)	0.6657 (0.0074)
2006 (N=11352)	0.6924 (0.0068)	0.6408 (0.0085)

Table 2: Unidimensional P_α Indices, 2003-2006 (SE in brackets)

	Income		Health	
	Fixed Line	**Relative Line**	**Fixed Line**	**Relative Line**
$\alpha=0$				
2003 (N=6107)	0.208 (0.005)	0.208 (0.005)	0.189 (0.004)	0.189 (0.004)
2004 (N=11069)	0.169 (0.004)	0.209 (0.004)	0.108 (0.003)	0.108 (0.003)
2005 (N=11950)	0.142 (0.003)	0.201 (0.004)	0.145 (0.003)	0.145 (0.003)
2006 (N=11352)	0.114 (0.003)	0.182 (0.004)	0.189 (0.004)	0.163 (0.003)
$\alpha=1$				
2003 (N=6107)	0.054 (0.002)	0.054 (0.002)	0.046 (0.002)	0.046 (0.002)
2004 (N=11069)	0.035 (0.001)	0.044 (0.001)	0.026 (0.001)	0.026 (0.001)
2005 (N=11950)	0.030 (0.001)	0.042 (0.001)	0.034 (0.001)	0.032 (0.001)
2006 (N=11352)	0.023 (0.001)	0.035 (0.001)	0.041 (0.001)	0.035 (0.001)
$\alpha=2$				
2003 (N=6107)	0.025 (0.002)	0.025 (0.002)	0.019 (0.001)	0.019 (0.001)
2004 (N=11069)	0.014 (0.001)	0.017 (0.001)	0.011 (0.001)	0.011 (0.001)
2005 (N=11950)	0.013 (0.001)	0.016 (0.001)	0.016 (0.001)	0.015 (0.001)
2006 (N=11352)	0.009 (0.001)	0.013 (0.001)	0.017 (0.001)	0.015 (0.001)

Table 3: Bidimensional Union and Intersection Headcount Ratios, 2003-2006

	Intersection		Union	
	Fixed Line	**Relative Line**	**Fixed Line**	**Relative Line**
2003 **(N=6107)**	0.08 (0.003)	0.08 (0.003)	0.318 (0.006)	0.318 (0.006)
2004 **(N=11069)**	0.044 (0.002)	0.057 (0.002)	0.233 (0.004)	0.260 (0.004)
2005 **(N=11950)**	0.035 (0.002)	0.058 (0.002)	0.251 (0.004)	0.285 (0.004)
2006 **(N=11352)**	0.03 (0.002)	0.059 (0.002)	0.273 (0.004)	0.286 (0.004)

Table 4: Bourguignon-Chakravarty Indices, b=1 (equal weights on income and health poverty), 2003-2006 (SE in brackets)

	$\theta=1$		$\theta=5$	
	Fixed Line	**Relative Line**	**Fixed Line**	**Relative Line**
$\alpha=1$				
2003 **(N=6107)**	0.1 (0.003)	0.1 (0.003)	0.091 (0.002)	0.091 (0.002)
2004 **(N=11069)**	0.061 (0.001)	0.070 (0.002)	0.057 (0.001)	0.064 (0.001)
2005 **(N=11950)**	0.065 (0.002)	0.074 (0.002)	0.061 (0.001)	0.068 (0.001)
2006 **(N=11352)**	0.065 (0.001)	0.071 (0.002)	0.062 (0.001)	0.066 (0.001)
$\alpha=2$				
2003 **(N=6107)**	0.053 (0.002)	0.053 (0.002)	0.042 (0.002)	0.042 (0.002)
2004 **(N=11069)**	0.028 (0.001)	0.032 (0.001)	0.024 (0.001)	0.027 (0.001)
2005 **(N=11950)**	0.033 (0.001)	0.037 (0.001)	0.028 (0.001)	0.031 (0.001)
2006 **(N=11352)**	0.029 (0.001)	0.033 (0.001)	0.026 (0.001)	0.028 (0.001)

Table 5: Bourguignon-Chakravarty Indices, b=0.33 (three times higher weight on income poverty), 2003-2006 (SE in brackets)

	θ=1		θ=5	
	Fixed Line	**Relative Line**	**Fixed Line**	**Relative Line**
α=1				
2003 (N=6107)	0.068 (0.002)	0.068 (0.002)	0.063 (0.002)	0.063 (0.002)
2004 (N=11069)	0.043 (0.001)	0.052 (0.001)	0.041 (0.001)	0.049 (0.001)
2005 (N=11950)	0.040 (0.001)	0.052 (0.001)	0.038 (0.001)	0.048 (0.001)
2006 (N=11352)	0.035 (0.001)	0.046 (0.001)	0.033 (0.001)	0.043 (0.001)
α=2				
2003 (N=6107)	0.036 (0.002)	0.036 (0.002)	0.030 (0.001)	0.030 (0.001)
2004 (N=11069)	0.019 (0.001)	0.023 (0.001)	0.017 (0.001)	0.019 (0.001)
2005 (N=11950)	0.020 (0.001)	0.024 (0.001)	0.017 (0.001)	0.021 (0.001)
2006 (N=11352)	0.016 (0.001)	0.020 (0.001)	0.014 (0.001)	0.017 (0.001)

Table 6: Bourguignon-Chakravarty Indices, b=3 (three times higher weight on health poverty), 2003-2006 (SE in brackets)

	θ=1		θ=5	
	Fixed Line	**Relative Line**	**Fixed Line**	**Relative Line**
α=1				
2003 (N=6107)	0.192 (0.006)	0.192 (0.006)	0.179 (0.005)	0.179 (0.005)
2004 (N=11069)	0.113 (0.003)	0.122 (0.003)	0.108 (0.003)	0.114 (0.003)
2005 (N=11950)	0.135 (0.004)	0.139 (0.004)	0.130 (0.003)	0.132 (0.003)
2006 (N=11352)	0.151 (0.004)	0.144 (0.004)	0.147 (0.004)	0.137 (0.003)
α=2				
2003 (N=6107)	0.097 (0.004)	0.097 (0.004)	0.080 (0.004)	0.080 (0.004)
2004 (N=11069)	0.053 (0.002)	0.058 (0.002)	0.046 (0.002)	0.048 (0.002)
2005 (N=11950)	0.068 (0.003)	0.072 (0.003)	0.061 (0.002)	0.062 (0.002)
2006 (N=11352)	0.067 (0.003)	0.068 (0.003)	0.062 (0.002)	0.059 (0.002)

Appendix 1

Definition of Income: The income measure we use is equivalised income after social transfers using the EU definition of income and the modified OECD equivalence scale. The EU definition of income consists of:

- Direct income (employee cash and non-cash income)
- Gross cash benefits or losses from self-employment
- Other direct income (but not pensions from individual private plans, value of goods produced for own consumption, employer's social insurance contributions)
- All social transfers (e.g. unemployment benefits, housing allowances, sickness allowances etc).

Tax on income and contributions to state and occupational pensions are deducted from this to give disposable income, which is then adjusted to equivalised income by applying the modified OECD scale (1.0 first adult, 0.5 other adults, 0.3 children aged less than 14). For details see CSO (2007).