

HEDG

HEALTH, ECONOMETRICS AND DATA GROUP

THE UNIVERSITY *of York*

WP 13/11

Health and Wealth on the Roller-Coaster: Ireland, 2003- 2011

David Madden

May 2013

york.ac.uk/res/herc/hedgwp

Health and Wealth on the Roller-Coaster: Ireland, 2003-2011

David Madden

(University College Dublin)

May 2013

Abstract: This paper reviews developments in income and health poverty in Ireland over the 2003-2011 period using data from the Survey of Income and Living Conditions (SILC). It also examines developments in the correlation between the two. Income poverty fell up to and including 2009, after which this trend is reversed. Health poverty shows less of a trend over the period though there is some evidence of a reduction in health inequality from 2006. Movements in bi-dimensional poverty are mostly driven by income poverty, but there is evidence of a reduction in the correlation between health and income poverty over the period.

Keywords: multidimensional poverty, dominance.

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

I am grateful to Tim Callan and Brian Nolan for useful discussion. I would also like to thank Jean Yves Duclos and Abdelkrim Araar for use of their DASP package. The usual disclaimer applies.

Health and Wealth on the Roller-Coaster: Ireland, 2003-2011

1. Introduction

It seems fair to describe Ireland's recent macroeconomic experience as a roller-coaster. The years from 2003 to about 2007 saw the last period of the era starting around 1995 when Ireland became known as the Celtic Tiger.¹ However, from about 2008 onwards, Ireland experienced one of the sharpest declines in output in the OECD area. As table 1 shows, both GNP and GNP per head showed falls approaching 10 per cent in 2009, and output has continued to fall since then (we use GNP as opposed to GDP as net factor flows in Ireland are unusually large and so GNP is regarded as a more accurate measure of "National Income"). The fall in consumption per head has not been quite so dramatic but the turnaround since 2008 is still stark (as explained below our micro data only became available in 2003, so we date our analysis from that year). Unemployment has also increased dramatically from 4.6 per cent at end 2003 to 14.6 per cent at end 2011.

Given such dramatic developments in the macro aggregates, what has happened at a more micro level? In particular, how have indicators of living standards such as income poverty and health evolved? It is particularly interesting to look at developments in this area since the recent downturn began and since the first budgetary responses to the crisis were introduced in October 2008. That is the question which this paper addresses. We examine developments in income poverty over the period using nationally representative micro data. We also analyse developments in self-assessed health using the approach pioneered by Allinson and Foster (2004) and finally, acknowledging that poverty may be regarded as a multi-dimensional phenomenon, we apply sequential stochastic dominance to analyse income and health poverty *together*.

Before describing the overall layout of the paper, it is useful to review some other work in the area. The evidence appears to be that since the onset of the crisis in 2008, budgetary policy at least has, in relative terms, been progressive. Callan et al (2013) provide a review of developments in overall income inequality for the 2003-2011 period and also examine the specific contribution of budgetary policy. They find that overall inequality has changed

¹ While the phrase "Celtic Tiger" is not favoured by all commentators (see for example Walsh and Honohan, 2002, who prefer to label it the "Irish Hare") by this stage it has become the standard phrase used to describe this period of high growth.

very little over a long period (from 1994 up to the current period) and this also applies to the more recent recession. However they also find a drop in inequality specifically in 2009, which was reversed almost immediately in 2010. Over the recession period of 2008-2011 they find that budgetary policy has been broadly progressive with greater reductions in incomes for those in the upper half of the distribution and within that half, greater reductions again for the better-off. Developments in the lower half of the income distribution are driven to a large extent by the degree to which old-age pensions have not been cut during the recession, which has relatively favoured those in the 2nd and 3rd deciles of the distribution. The main focus of Callan et al's analysis is developments in inequality but they also provide some data on changes in poverty and their results are consistent with our findings below.

Callan, Nolan and Maitre (2011) look at the Irish experience with respect to distribution of household income during what they term the Great Recession (i.e. the downturn from about 2007 onwards). They conclude that while the macro figures as illustrated in table 1 mark Ireland out as one of the countries most affected by the downturn, in relative terms the principal part of the burden has fallen on higher income groups. They also point out that, in common with other OECD countries, though to a more pronounced extent, the impact of the recession on the household sector was considerably less than on the economy as a whole, with the company sector bearing the brunt. The combination of falling average incomes with a greater impact on richer households meant that relative poverty, as measured by the fraction below 60 per cent of median income, fell quite substantially between 2007 and 2009. Of course in the same way that a relative poverty measure can be misleading in times of rapid economic growth (where it can understate improvements in living standards), so too it can be misleading at a time of economic contraction. This is because relative poverty rates can fall even though overall living standards are also falling. Poverty measures based upon an absolute poverty line (i.e. fixed in real terms and not expressed as a fraction of a measure of central tendency) can give quite a different picture, and this is explored in more detail below.

The main conclusion of the Callan, Nolan and Maitre paper is that the initial budgetary responses to the recession and the fiscal deficit in Ireland were mainly concentrated on higher taxes. Since many in the lower part of the income distribution do not pay tax, in relative terms this impacted more on the better-off. The same was true of the pension levy charged on public sector workers and the public sector paycuts, since most public sector

workers are in the middle or upper part of the income distribution. Going into 2011 and 2012 the focus of adjustment turned to social welfare payments. While reductions in these payments would be expected to have a greater impact further down the income distribution, Callan, Nolan and Maitre speculated that, even allowing for this, in broad terms the recession would still have had an equalising effect on household incomes. As we have more recent data at our disposal, we explore this later on in the paper.

In a further contribution in this area Nolan et al (2012) cover much the same ground as Callan, Nolan and Maitre (2011) with the advantage of having an extra year of data. Their conclusions for the 2007-2009 period are very similar but they also find that the trend of declining relative poverty was reversed in 2010. It should be noticed that this analysis relies on data from the 2010 Survey of Income and Living Conditions (SILC), data which was subsequently revised. The revised data (which is used in this paper) typically shows that the reversal in the downward trend of poverty witnessed in 2010 was less pronounced than originally thought. However, the qualitative results remain more or less unchanged. Given that average incomes in this period were declining, this also implies that poverty as measured with respect to an absolute poverty line showed quite a sharp rise.

So far we have concentrated on developments in household income distribution. But household welfare and poverty are arguably multidimensional concepts and one of the contributions of this paper will be to examine what has happened to multidimensional poverty. Madden (2011) examined the development of the Bourguignon-Chakravarty bi-dimensional poverty index defined over income and health over the 2003-2006 period. He found that uni-dimensional income poverty fell, while health poverty rose and then fell. Movement in the bi-dimensional indices reflected movements in the individual indices and depended upon the relative weights assigned to income and health poverty. However, his analysis did not cover the recent recession.

More recently Whelan et al (2012) have examined a snapshot of multi-dimensional poverty in Ireland in 2009. They apply the recently developed Alkire and Foster (2011) methodology to a wide sample of European countries, including Ireland. Their measure of multidimensional poverty includes relative income poverty and also various measures of deprivation, including health. However, given that this is just a snapshot of a single year, it is not possible to see how this multidimensional measure has evolved over the boom and subsequent recession.

Finally, the Department of Social and Family Affairs (2013) have released a social impact assessment of the main tax and welfare changes introduced in the 2013 budget. As SILC data only runs as far as 2011, their analysis uses the microsimulation model SWITCH.

The analysis shows that the distributive impact of the combined direct tax and welfare measures depend upon assumptions made regarding the deferral or otherwise of the local property tax. If deferrals of property tax are treated like waivers or exemptions from the charge, then the impact of the welfare and direct tax measures in the 2013 budget could be regarded as broadly neutral as the impact was greatest (in percentage terms) on the middle quintile, with the least impact upon the highest and lowest quintiles. However, if no deferral of the tax is assumed, then the combined impact is regressive, with the greatest percentage loss for the poorest quintile and the least percentage loss for the richest quintile. It should be noted that deferral is not the same as exemption and that it must be assumed that at some stage the tax will be paid. Thus while timing of payment does complicate distributional analysis, it does seem fair to conclude that the impact of the combined measures of the 2013 budget was regressive. Those households who fared relatively the worst were those with children (reflecting changes in child benefit).

The layout of this paper is as follows: in section 2 we describe the income poverty and health measures we use as well as the approach of sequential stochastic dominance. In section 3 we discuss our data and present results, while section 4 offers some concluding comments.

2. Unidimensional and Multidimensional Poverty

The last ten years or so has seen substantial developments in the measurement of multidimensional poverty. This reflects 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.² For example, individuals may experience poverty with respect to housing or other assets, education, nutrition or health as well as income. The approach to multi-dimensional poverty analysis can also differ with some authors choosing to calculate multi-dimensional poverty indices (Bourguignon and Chakravarty, 2003) and others

² For a recent review of work in this area, see Lustig (2011).

looking for more robust multi-dimensional poverty orderings for broader classes of measures (Duclos et al, 2006). One important issue in analysing multidimensional poverty 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). However, this dichotomy has been criticised on the basis that the union approach leads to “too many” being classified as poor, whereas the intersection approach (particularly as the number of dimensions increases) leads to too few (see Lustig, 2011). One of the more influential recent contributions is that of Alkire and Foster (2011) which provides multidimensional indices of which the union and intersection approaches are special cases. Ravallion (2010) in contrast has questioned the need for multidimensional indices at all and suggests instead a “dashboard” of multiple indices. A review of this literature which advocates an eclectic approach can be found in Ferreira and Lugo (2012).

Our approach in this paper is to first of all search for robust income and health poverty orderings in one dimension. In the case of health, given that our health measure is ordinal and categorical, we will instead be looking for dominance relationships across the whole of the health distribution, following the approach of Allinson and Foster (2004). We then look for sequential poverty dominance following the approach of Duclos and Makdissi (2009) and Duclos and Echevin (2011). Sequential poverty dominance is suited to finding robust poverty orderings in two dimensions when one of the dimensions is an ordered categorical measure. Essentially it involves looking for income poverty dominance between two periods for that group with poorest health, followed by a search for dominance for the two poorest groups and so on. Along the way we also provide evidence on standard poverty indices and we also investigate the correlation between health and income for those who are poor and non-poor..

Univariate Poverty Dominance

We now give a brief account of dominance in the areas of income poverty and health, followed by an account of sequential dominance. We will first briefly run through the analysis of poverty dominance for the single dimension of income (or whatever measure of household resources we are using). Let x be the measure of household resources and let z be the poverty line. Following the exposition in Duclos et al (2006), when dealing with poverty in a single dimension the stochastic dominance curve for x is given by

$$P^\alpha(z) = \int_0^z (z-x)^\alpha dF(x).^3$$

where $F(x)$ is the cumulative distribution of x . For first order stochastic dominance, poverty does not increase for any possible choice of z when moving from a distribution A to a distribution B , if the incidence of poverty under distribution A is never greater than under distribution B (i.e. where $\alpha=0$). If this condition for first order dominance is not met, then second order dominance ($\alpha=1$) may be investigated which requires that normalised poverty deficits should not increase for any possible choice of z when moving from A to B . Thus, in general, orders of dominance $s=\alpha+1$ can be examined. It is also possible to search for dominance over a more restricted range of poverty lines.

Health Dominance

The approach in the previous section is suitable when the variable x is cardinal. However, when dealing with health, very often a cardinal measure will not be available. More typically, measures which address general health may come in the form of an ordered, categorical self-assessment of health. In that instance the analyst has two choices: either to employ a dominance approach which is specifically designed to deal with ordered categorical data, or else to transform the ordered measure into a cardinal measure. The transformation of ordered categorical data into cardinal data has been discussed extensively by Van Doorslaer and Jones (2003), but as Madden (2010) points out, results obtained can be sensitive to the approach adopted.

As we wish to retain our data in its original, ordered categorical format, we choose to adopt a dominance approach specifically designed to deal with this sort of data, the approach of Allinson and Foster (2004), henceforth AF. While this approach has the advantage of being specifically designed to deal with ordered categorical data, it carries a disadvantage in that it is ill-equipped to deal with poverty dominance. This is because of the difficulty in identifying a poverty line for categorical data. Thus the AF approach could be best described as a *welfare dominance* (with respect to health) approach.

³ In common with most of the literature in this area, we choose to work with the continuous as opposed to the discrete version of poverty measures when dealing with dominance issues.

The measure of self-assessed health we have is the following: individuals answer a question of the form: what is your general health status? The possible answers are: very bad, bad, fair, good and very good. While this measure appears to give a good indicator of overall health (Idler and Benyamini, 1997) it is not cardinal, and with only five categories, it is not suited to the application of standard poverty and inequality indices. This is because standard measures of the spread of a distribution which use the mean as a reference point, such as the Gini, are inappropriate when dealing with categorical data, since the inequality ordering will not be independent of the (arbitrarily chosen) scale applied to the different categories. In this instance a more appropriate reference point is the median category and the cumulative proportions of the population in each category is the foundation of the AF analysis of dominance and inequality with categorical data. This is because while changes in the scale used will affect the width of the steps of the cumulative distribution, the height of the cumulative distribution is invariant to the choice of scale, thus providing the crucial property of scale independence.

Thus suppose we have a measure of SAH with n different categories which can be clearly ordered $1, \dots, n$. Let m denote the median category and let P and Q denote two cumulative distributions of SAH with P_i and Q_i indicating the *cumulative* proportion of the population in category i , in each distribution, where $i=1, \dots, n$. In this case distribution P will dominate distribution Q if the cumulative frequency at each point on the ordinal scale (as we go from lower to higher) is always higher in Q than in P . This is equivalent to the first order stochastic dominance condition referred to above. For a recent example of application of this approach to a comparison of SAH between different social classes, see Dias (2009).

AF also provide a partial inequality ordering based on a median-preserving spread of the distribution (analogous to the partial ordering based on a mean preserving spread provided by say a Lorenz comparison). For the case where both P and Q have identical median states m then P has less inequality than Q if for all categories $j < m, P_j \leq Q_j$ and for all $j \geq m, P_j \geq Q_j$. What this is effectively saying is that distribution Q could be obtained from distribution P via a sequence of median-preserving spreads.

In section 3 we will present results for health and health inequality dominance using the AF approach. Before that we give a brief account of sequential poverty dominance.

Sequential Poverty Dominance

In terms of producing poverty indices, multidimensional analysis works best when attributes can be measured cardinally. Unfortunately, in some cases, one or more attributes are not available on a cardinal basis. As discussed above general health measures are usually only available on an ordinal basis, such as measures of self-assessed health. Since these measures typically give health status on the basis of a discrete number of categories e.g. very good, good, fair, bad, very bad, there is a clear ranking and it is this clear ranking which permits the application of multidimensional analysis to situations where one of the attributes is ordinal in nature.

Following the exposition in Duclos and Makdissi (2009) suppose that the population can be divided into K mutually exclusive and exhaustive subgroups with population share defined by

$\phi(k), k = 1, \dots, K$ and $\sum_{k=1}^K \phi(k) = 1$. The subgroups could be defined over a wide range of

characteristics but what is most important is that, for a given measure of the continuous attribute (income), certain groups can be reasonably viewed as having lower overall well-being than others. Thus, for a given level of income, an individual with, say zero literacy, has lower overall well-being than someone with 100% literacy, or someone with “very bad” health has lower well-being than someone with “very good” health.⁴ We assume that, for given alternative indicators of well-being, that the K subgroups can be ordered in decreasing value of needs, so that group 1 has greater needs than group 2 etc. Assuming that the poverty indices are additive then poverty in the groups can be denoted as

$$P(k; z(k)) = \int_0^{z(k)} \pi_k(x) f(x; k) dx$$

⁴ Note that it is possible that responses to the self-assessed health question may differ systematically, so that, for example, an older person may view a given health state as “very good” while a younger person might view it only as “good”. This issue is typically addressed via anchoring vignettes and is beyond the scope of this paper. For a general discussion, see King and Wand (2007).

where $f(x; k)$ is subgroup k 's density of living standards at x and $z(k)$ is the poverty line in x for subgroup k . $\pi_k(x; z(k))$ is the contribution to subgroup k 's poverty of an individual in that group with $\pi_k(x; z(k)) = 0$ if $x > z(k)$. Total population poverty is then given by

$$P(z(1), \dots, z(K)) = \sum_{k=1}^K \phi(k) \int_0^{z(k)} \pi_k(x) f(x; k) dx$$

$$= \sum_{k=1}^K \phi(k) P(k; z(k))$$

The poverty lines for the different subgroups can be ordered as follows:

$$z(1) \geq z(2) \geq \dots \geq z(K)$$

i.e. the poverty line for the most needy group must always be at least as great as that for the next most needy group and so on (in the application below, we assume the same income poverty line for each group). For first-order dominance then we need an ordering of the first-order derivatives of $\pi_k(x)$ with respect to x , which is denoted as $\pi_k^{(1)}(x)$, such that

$$\pi_1^{(1)}(x) \leq \pi_2^{(1)}(x) \leq \dots \leq \pi_K^{(1)}(x) \leq 0 \forall x.$$

What this essentially says is that poverty is effectively decreasing in income and the decrease in poverty following a one unit increase in income must be at least as great for the most needy group as it is for the next most needy group and so on.

In the case of continuous poverty indices we assume that the derivatives of the functions $\pi_k^{(1)}(x)$ are continuous at the poverty line up to the $(s-1)$ th order. For first order dominance this requires that $\pi_k(x; z(k)) = 0 \forall k = 1, \dots, K$. The class of multidimensional poverty indices $\Pi^1(z(1), \dots, z(K))$ then includes all of the P indices which satisfy the assumptions that $z(1) \geq z(2) \geq \dots \geq z(K)$, $\pi_1^{(1)}(x) \leq \pi_2^{(1)}(x) \leq \dots \leq \pi_K^{(1)}(x) \leq 0 \forall x$, and $\pi_k(x; z(k)) = 0 \forall k = 1, \dots, K$.

Duclos and Makdissi (2009) then establish the following condition for first order poverty dominance for heterogenous populations:

$$\begin{aligned} \Delta P(\zeta(1), \dots, \zeta(K)) &> 0 \\ \forall P(\zeta(1), \dots, \zeta(K)) &\in \Pi^1(\zeta(1), \dots, \zeta(K)) \\ \text{and } \forall \zeta(k) &\in [0, z(k)], k = 1, \dots, K \\ \text{iff } \sum_{k=1}^i \Delta P^0(k; \zeta) &> 0, \forall \zeta \in [0, z(i)], i = 1, \dots, K \end{aligned}$$

What this condition effectively requires is that if we start off with the neediest group ($k=1$) then for dominance to hold for distribution A over distribution B we first require that the headcount for this group should be higher in distribution A compared to B. We then require that the cumulative headcount for the two neediest groups ($k=1,2$) should be higher in A than B and so on. Thus the sequential cumulative headcount for all groups up to where we reach group K should be higher for A compared to B, where the sequence is carried out starting with the neediest group etc. Note that it does not require that each subgroup k have more poverty independently in A. Poverty for, say, group 3, could be lower in A compared to B, as long as cumulative poverty for groups 1-3 is higher in A. For a recent application of this to health and income data from Canada and the US, see Duclos and Echevin (2011).

For our example here we investigate income poverty for 2003 through to 2011 where the population is partitioned into four groups on the basis of self-assessed health. Even though we have five categories of self-assessed health, as the fraction of the population reporting “very bad” health is so small (typically less than 1%), we combine the two lowest categories.

3. Data and Results

Our data comes from nine consecutive cross-sectional surveys (2003-2011) which are the Irish part of the European Union Survey of Income and Living Conditions (EU-SILC).⁵ 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 13,000 and 14,000 for each year. However, in Ireland there was only six months data collection for 2003 (as

⁵ For details of the Irish part of EU-SILC see CSO (2007) and the documentation at <http://www.cso.ie/eusilc/default.htm>

opposed to twelve months 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 ordinal health measure 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. We confine our analysis to those aged 16 and over, as the health question was not put to those aged younger than 16. This reduces our sample size by around 2000-3000 each year. Average incomes are slightly higher for the reduced sample, and while the moments of the distributions appear to be quite close together, the Kolmogorov-Smirnov tests for equality of distributions rejects the null for each year.⁶ However, when we analyse the FGT P_α poverty indices for income only for the two samples, the trends over time are practically identical. On balance we do not believe that our qualitative results are greatly affected by working with the smaller sample, and if we wish to include health in our poverty analysis, then unfortunately we have no other option.

In table 2 we provide summary statistics for the frequencies of the categories of self assessed health and for mean equivalised income. Equivalised income is presented in 2010 prices. Note that in order to remove the influence of outliers we trim the data of the top and bottom 0.5% (by non-equivalised income). Table 2 shows that mean equivalised income rose continuously from 2003 to 2007. It was essentially constant between 2007 and 2009 but then dropped quite sharply in 2010 and 2011.

Health Dominance

The position with respect to self-assessed health is a little more complicated. Between 2003 and 2011 the fraction declaring very good health dropped from just under 47% to just under 40%. For the most part this has been offset by an increase in the fraction declaring good health. However as explained above, unless dominance is observed, it is not possible to state

⁶ Details of the comparison between the distribution of income with and without those aged 16 and younger are available on request.

that health has improved, without assigning an arbitrary scale to each health category. In table 3 we present a grid which shows whether dominance applies. As the frequencies in the category “very bad” are so small, we combine the two lowest categories together. Recall that if year A stochastically dominates year B then the cumulative frequency for year B at each point on the ordinal scale (as we go from lower to higher) is always higher in B than in A. Note also that entries which are above the main diagonal indicate a situation where the earlier year (the row year) dominates the later one (the column year), while entries below the diagonal indicate where the later year dominates the earlier one. In table 3 “F” refers to first-order dominance and we can see there are seven instances of first order dominance over the period. Four of these are above and to the right of the main diagonal, indicating that on balance health was deteriorating over time, between 2003 and subsequent years. The only instance where health improves moving forward in time are where 2007 dominates 2006 and where 2010 also dominates 2006 and 2009. In all other cases, first order dominance is not observed.

What about the spread measure introduced by AF? Recall that spread dominance, as denoted by “S” indicates greater spread in the row year compared to the column year. Thus the distribution for year A has greater spread than that for year B if (a) year A and year B both have the same median category, m , (b) for all categories below the median the cumulative frequency in A is at least as great as in B and (c) for all categories greater than or equal to the median the cumulative frequency in B is at least as great as that in A. Thus instances of S below the main diagonal indicate that health inequality has been decreasing over time, and this seems to be the case. For example, 2010 and 2011 show less health inequality than all years from 2003 to 2006 inclusive. 2010 and 2011 are also both more equal than 2009.

Thus to summarise, the grid in table 3 indicates first order dominance for 2007 and 2010 over 2006 and 2009. It also indicates that health inequality appeared to decline over the period. What is also noticeable is that health inequality did not increase as the recession began, as is indicated by the greater spread in 2005 and 2006 compared to the 2008-2011 period.

Poverty Dominance – Fixed Poverty Line

What about poverty dominance in the area of income? We will first of all look at the case where the poverty line is fixed in real purchasing power terms. We now need to fix what we could view as a “reasonable” range for the poverty line to lie within. We chose a range from

zero to 80% of median equivalised income for the year when equivalised income was highest (2007), noting that a typical value for a poverty line is 50-60% of median equivalised income. That gives a figure (in 2010 prices) of just over €170. We use the “dompov” command from the DASP package of Arrar and Duclos (2012) and provided we observe dominance for one distribution over another up to the upper limit of our poverty line, we regard that as poverty dominance.

Once again we analyse this using a grid. Table 4 shows the grid for poverty dominance. If poverty is *declining* over time, then we expect to see entries below the main diagonal. We use “F” to indicate first order dominance and “S” to indicate second order dominance. We note many instances of F below the diagonal up to about 2009. However, after 2009 we note that entries appear above the main diagonal, along the 2010 and 2011 columns. This implies that these years are poverty dominated by earlier years and indicates a pattern whereby poverty was falling up to and including 2009, but then this was reversed in 2010 and 2011.

Note that we indicate a number of entries as F*. By this we mean that poverty dominance applies across the range of income up to the poverty line except for a crossing at very low levels of income (below the first percentile), crossings which we believe are more likely to reflect measurement error rather than genuine crossings of the poverty incidence curves. We also include a category which we label WF which indicates weak first order dominance. This is the situation where the difference between poverty incidence curves was not statistically significant over some part of our poverty range, but that it was statistically significant over another part. We do not observe dominance if either there is no part of the poverty range where there is a statistically significant difference or if over our range we observe two (or more) instances of a statistically significant difference but of opposite sign (this is essentially the criterion suggested by Bishop et al, 1991).

Thus table 4 indicates falling poverty for the first two years of the recession, 2008 and 2009, but a reversal of this in later years. This is consistent with the findings of Callan et al (2013) and Nolan et al (2012).

For completeness sake in tables 5 and 6 we present poverty indices for the well known P_α measures of Foster, Greer and Thorbecke (1984, henceforth known as FGT) for fixed and relative income poverty lines. We show the results for three values of α , 0, 1 and 2. An α

value of 0 refers to a headcount measure, a value of 1 gives a measure which is sensitive to the depth of poverty while a value of 2 indicates a measure which places a greater weight on the poverty shortfalls of the very poorest. These tables bear out the previous results whereby poverty falls pretty consistently up to and including 2009, but then this trend is reversed quite sharply in 2010 and 2011, so that poverty indices return to values previously seen around 2005-2006, for the absolute poverty line. It is also worth noting that the rise in poverty starting in 2010 is proportionately greater for the P_2 measure, which is more sensitive to income distribution within the poor.

Thus to summarise so far in terms of univariate analysis: income poverty falls consistently up to and including 2009 and this is reversed in 2010 and 2011. Thus for the early part of the recession at least the evidence is that income poverty does not rise. Owing to the categorical nature of self-assessed health it is more difficult to come to any definitive conclusion here. There is a fall in the proportion reporting very good health. However there are also falls in the proportions reporting very bad and bad health. The proportions involved moving out of the lowest two categories are quite small compared to those moving out of the highest category, so unless there was a considerably higher weight attached to the improvements in health in the lower categories, the likelihood is that most scales would probably record a deterioration in health, although most of that deterioration had occurred by 2006. What also seems clear is that since about 2007 there has been a reduction in health inequality.

Sequential Poverty Dominance

We now move on to sequential poverty dominance for income and health. Recall that this involves looking for income poverty dominance between two years for the most needy group, followed by looking for dominance for the two neediest groups, then the three neediest groups and so on. In our case we identify “needy” with self-assessed health. However as indicated above, since the proportions in the neediest group (very bad health) are so small, typically less than one per cent, we combine the two lowest health categories. Thus our neediest group is those with very bad and bad health, the next neediest is those with fair health and so on.

As before, we present our dominance results in table 7 using a grid. The pattern is very similar to that in tables 4 and 5. We observe many instances of weak sequential dominance below the main diagonal and the only instances above the diagonal are for 2010 and 2011,

once again indicating that even when we look at developments exploiting the distribution of two dimensions, 2010 shows a major reversal of what had been happening in earlier years, and this continues in 2011.

The similarity in results between tables 4 and 7 is notable. It is helpful to examine in more detail one instance where the dominance results differ. This is in the comparison between 2005 and 2006. Table 4 indicates that for income dominance only, 2006 weakly dominates 2005. Figure 1 shows the difference in poverty incidence curves (along with 95% confidence intervals) between the two years. As can be seen, the difference is negative (indicating lower incidence in 2006 compared to 2005) and it is statistically significant for much of the range of the poverty line. Figure 2 however shows the difference between poverty incidence curves for the neediest group (those with very bad and bad health). There is no range of the poverty line where we observe statistically significant dominance and hence dominance does not apply. Thus while poverty dominance was to be observed for the population as a whole, it was not observed for the neediest, indicating that improvements in income between 2005 and 2006 were more concentrated amongst the healthier (or less needy). Analysis of income by health group show that average incomes for those with “very bad” and “bad” health (our neediest group) fell between 2005 and 2006, while average incomes for the three less needy groups all rose. A corollary of this is that polarisation of income by health group increased between 2005 and 2006, with the less healthy becoming poorer while the relatively healthier groups experienced income increases (this is confirmed by the change in the Duclos, Esteban and Ray polarisation index which increased between 2005 and 2006 from 0.199 to 0.203, a change which is significant at the 10% level).

Multidimensional Poverty Indices

As a complement to the sequential analysis above, we also calculate multidimensional poverty indices. As discussed above an important issue in the choice of index is whether to adopt an intersection or union approach. The advantage of the Alkire and Foster (2011) methodology is that the union and intersection choices are extreme cases of a more general approach, and a compromise or intermediate position is possible. If we have m possible dimensions of poverty then the compromise involves counting poverty in k dimensions where k lies between one (a union approach) and m (an intersection approach). However, when working with only two dimensions, no such k exists and so it is easier to present results for both the union and intersection approaches. Bearing in mind also the criticisms of Ravallion

(2010), in addition we present results showing the degree of dependency between the distributions of health and income. Given that we have already provided information on the marginal distributions of health and income, information on the dependence between these distributions could be regarded as constituting the final element in the “dashboard” (see Ferreira and Lugo, 2010).

Table 8 gives the headcount poverty rates for both intersection and union approaches for the period 2003-2011. Note we only present headcount measures as gap or weighted gap measures would not be appropriate given that our health measure has no cardinal interpretation. We also include two different income poverty lines, one representing a fixed amount in real purchasing power (60% of median 2007 income), and the other simply 60% of median income of the year in question. There is also a choice to be made concerning the health poverty line. Given the classifications available the two obvious candidates are those with health less than or equal to “bad health” or the category of “fair health”.

For the fixed income poverty line and intersection approach we observe sharp declines in poverty up to and including about 2009. We then see a levelling out for the lower of the health poverty lines (HPov=3) but an increase for the higher of the lines (HPov=4). This suggests an increase in income poverty for those with health in the “fair health” category. The union approach sees quite a sharp pick-up in poverty from 2010. Given the parallel movement in income poverty it seems likely that people who previously had neither income nor health poverty may now be falling into income poverty.

Results for the relative income poverty line are quite similar, except in the intersection case with the higher health poverty line, where multidimensional poverty continues falling through 2010 and 2011. What this seems to indicate is that the increase in relative income poverty which occurred between 2009 and 2011 was more concentrated amongst the *more healthy*, since the intersection measure shows a fall between the two years, suggesting that the households who moved into income poverty (as evidenced by table 6) were not, for the most part, households which were health poor.

We further explore this issue by looking at FGT poverty indices for each health grouping. One of the attractive features of the FGT index is that if the population can be decomposed into mutually exclusive and exhaustive subgroups then any of the FGT indices can be

expressed as the population weighted sum of the FGT indices for each subgroup. We divide our population into four subgroups corresponding to the self-assessed health categories (where once again we combine the categories “bad” and “very bad”). A correlation between health and income poverty would suggest that the FGT index for the less healthy subgroups should be higher than for the more healthy subgroups. Figures 3-5 show the FGT P_α measures for $\alpha=0, 1, 2$ for the four subgroups. As expected the index is higher for the less healthy groups, but what is also noticeable is how the difference in the indices between the subgroups has been narrowing over the years. This further confirms the conjecture that the correlation between health and income poverty is weakening.

As a final check on this weakening correlation between health and income poverty, in table 9 we present measures of dependence between income and health. The measures we use are the polyserial correlation coefficient, the Spearman rank correlation coefficient and the Kendall tau-b rank correlation coefficient. These measures were chosen in preference to the standard Pearson correlation coefficient, since the value of this measure would depend upon the (arbitrary) scale used in the ordinal health measure. We present them for the complete distribution and also for those observations below the income poverty line and below the health poverty line. For the income poverty line we use the same upper limit as was used in the dominance analysis (80% of median income in 2007) and for the health poverty line we define as health poor those with health less than or equal to “fair”.

Looking at the dependence measures for the distribution as a whole we see that dependence has been falling since about 2007. Between 2003 and 2007 there was some fluctuation but no real trend evident. However, since 2007 all three measures of dependence show a decline. This decline is consistent with developments below the poverty lines. Just looking at the subset of people who are income poor, we can identify three phases. From about 2003 to about 2005 the Spearman and Kendall tau correlations hover around about 0.1, and all are statistically significant. For 2006 and 2007 the correlation drops to about 0.05, still statistically significant, though arguably not economically very significant. Since 2008 however, the correlation has vanished and by 2010/2011 it had even turned negative. For the subset of people who are health poor, the series is somewhat more volatile, with generally lower significance levels. Nevertheless, the trend of declining dependence in recent years is still evident.

One possible explanation for the reduced correlation between health and income is the experience of pensioners (those aged 65 and over). As has been documented by Nolan et al (2012) the relative position of pensioners in income terms has improved significantly in recent years. Since this group in general have poorer health than the non-pension population, their relative improvement in income terms could explain the apparent decoupling between health and income poverty. In table 10 we re-calculated table 9 for the non-pension population and found that in qualitative terms the results were quite similar. This is illustrated in figures 6 and 7 (for ease of visual interpretation we just include the Spearman correlation coefficient). For the under-65 group as a whole (i.e. including the non-poor) there is a slight decline in the correlation. However for the health and income poor from about 2007 there is a clear fall in the correlation (albeit from a fairly low level). Perhaps the main difference compared to the population including older people is that the negative correlation between income and health below the income poverty line in 2010/2011 which is evident in table 9 is not present in table 10. Thus while the relative improvement of the pension population over the period under review may explain part of the reduced correlation between health and income poverty, it is not the complete story.

4. Summary and Conclusions

We have presented quite a lot of results in the previous section, so it is useful to try to draw together our conclusions so far. Looking at the univariate analysis, we see that income poverty fell quite consistently up to about 2009, but there was a sharp reversal of this trend in 2010, which continued in 2011. Developments in health poverty are more difficult to assess, as it is more difficult to arrive at a poverty line when data is ordinal. But the dominance analysis carried out suggest that health deteriorated up to about 2006, improved slightly in 2007 and also appears to have slightly improved in 2010. Health inequality has also narrowed since about 2007. Overall though, and acknowledging the difficulty in assessing changes over time in an ordinal measure, what changes we have observed in self-assessed health since about 2006 appear to have been relatively modest.

Turning now to the bi-dimensional analysis, the dominance analysis for the most part mirrors what happened with univariate income poverty analysis. Once again this is consistent with a situation where developments in health have been less dramatic. Perhaps the most interesting

development in this area has been the reduced correlation between income and health over the period, both for the population as a whole and also for those experiencing health and income poverty. Part of this can be explained by the experience of the 65 years and over age group, but only part. It is possible that developments in health will eventually follow those in income with a lag (though it is likely that causality runs both ways). It is also possible that budgetary changes which affect the provision of health care will also feed into health over time and that this may have different impacts across the income distribution. For the present however, given this apparent decoupling, the likelihood is that developments in bivariate poverty will for the most part be driven by developments in income poverty.

References

- Alkire, S., and J. Foster (2011):** “Counting and Multidimensional Poverty Measurement”, *Journal of Public Economics*, Vol. 95, pp. 476-487.
- Allison, R.A., and J. Foster, (2004).** Measuring Health Inequality Using Qualitative Data. *Journal of Health Economics* Vol. 23, 505-524.
- Ararr, A., and JY Duclos (2012):** *DASP: Distributive Analysis STATA Package*. Université Laval and World Bank.
- Bishop, J.A. J. Formby and W.J. Smith (1991):** “Lorenz Dominance and Welfare: Changes in the U.S. Distribution of Income, 1967-1986 “, *Review of Economics and Statistics*, Vol. 73, pp. 134-139.
- Bourguignon, F., and S. Chakaravarty (2003):** The Measurement of Multidimensional Poverty. *Journal of Economic Inequality*. 1. 25-49.
- CSO (2007):** *EU Survey on Income and Living Conditions*. Stationery Office, Dublin.
- Callan, T., B. Nolan, C. Keane, M. Savage and J. Walsh, (2013):** “Crisis, Response and Distributional Impact: The Case of Ireland”, *ESRI Working Paper*, No. 456.
- Department of Social Community and Family Affairs (2013):** *Social Impact Assessment of the Main Welfare and Direct Tax Measures of the 2013 Budget*. Available at http://www.welfare.ie/en/downloads/2013-03_SIABudget2013_Final.pdf.
- Dias, P.R. (2009):** “Inequality of Opportunity in Health: Evidence from a UK Cohort Study”, *Health Economics*, Vol. 18, pp. 1057-1074.
- Duclos, J.Y., D. Sahn and S. Younger (2006):** Robust Multidimensional Poverty Comparisons. *Economic Journal*. 116. 943-968. (2006)
- Duclos, JY., and P. Makdissi (2009):** “Sequential Stochastic Dominance and the Robustness of Poverty Orderings”, *mimeo*.
- Duclos, JY., and D. Echevin (2011):** “Health and Income: A Robust Comparison of Canada and the US”, *Journal of Health Economics*, Vol. 30, pp. 293-302.
- Ferreira, F., and M. Lugo (2012):** “Multidimensional Poverty Analysis: Looking for a Middle Ground”, *IZA Policy Paper*, No. 45.
- Foster, J., J. Greer and E. Thorbecke (1984):** A Class of Decomposable Poverty Measures”, *Econometrica*. Vol. 52, pp. 761-766.
- Honohan, P., and B. Walsh, (2002).** Catching up with the Leaders: The Irish Hare. *Brookings Papers on Economic Activity* Vol. 1, 1-77.
- 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.

King, G. and J. Wand. (2007): “Comparing incomparable survey responses: Evaluating and selecting anchoring vignettes”. *Political Analysis*, Vol. 15, pp. 46-66.

Lustig, N., (2011): “Multidimensional Indices of Achievement and Poverty: What Do We Gain and What Do We Lose? An Introduction to JOEI Forum on Multidimensional Poverty”, *Journal of Economic Inequality*, Vol. 9, pp. 227-234.

Madden, D., (2010): “Ordinal and Cardinal Measures of Health Inequality: An Empirical Comparison”, *Health Economics*, Vol. 19, pp. 243-250.

Madden, D.; (2011): 'Health and income poverty in Ireland, 2003-2006'. *Journal of Economic Inequality*, 9 (1):23-33.

Nolan, B., Maitre, B., Voitchovsky, S., and C. Whelan (2012): “Inequality and Poverty in Boom and Bust”, GINI Discussion Paper.

Ravallion, M., (2010): Troubling Tradeoffs in the Human Development Index”, *World Bank Policy Research Working Paper* 5484.

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

Table 1: Ireland, Key Economic Indicators, 2003-2011

Year	GNP % Change	GNP Per Head % Change	Consumption Per Head % Change	Unemployment Rate (%) – end year s.a.
2003	4.9	3.2	1.6	4.6
2004	4.0	2.3	2.3	4.4
2005	5.0	3.5	4.8	4.3
2006	5.8	3.9	4.6	4.4
2007	4.2	0.8	2.9	4.7
2008	-1.8	-4.2	-2.5	8.5
2009	-8.1	-9.1	-6.4	12.8
2010	0.9	0.5	0.5	14.4
2011	-2.5	-2.9	-2.9	14.6

Source: Central Statistics Office “National Income and Expenditure” and Census 2006, Labour Force Survey, Quarterly National Household Survey.

Table 2: Summary of Self-Assessed Health and Equivalised Income (mean, 2010 prices)

	Very Bad	Bad	Fair	Good	Very Good	Equiv Y
2003 (N=6101)	0.0097	0.0298	0.1557	0.338	0.4668	214.65
2004 (N=10896)	0.0103	0.0317	0.154	0.3623	0.4417	222.55
2005 (N=11915)	0.0089	0.0332	0.1631	0.3663	0.4285	228.97
2006 (N=11360)	0.0082	0.0318	0.1714	0.3843	0.4043	237.60
2007 (N=10778)	0.0069	0.0256	0.1716	0.3892	0.4068	254.24
2008 (N=10013)	0.0053	0.0272	0.1712	0.3984	0.3980	250.84
2009 (N=9800)	0.0059	0.0305	0.1645	0.3984	0.4007	254.12
2010 (N=8704)	0.0083	0.0267	0.1600	0.3977	0.4073	237.39
2011 (N=8116)	0.0063	0.0297	0.1529	0.4134	0.3977	224.65

Source: Central Statistics Office, Survey of Income and Living Conditions (SILC), 2003-2011.

Table 3: Health Dominance

	2003	2004	2005	2006	2007	2008	2009	2010	2011
2003		F	F	F					
2004			F						
2005									
2006									
2007			S	F					
2008			S	S	S				
2009			S	S					
2010	S	S	S	F			F		
2011	S	S	S	S			S		

Table 4: Equivalised Income Poverty Dominance (Fixed Poverty Line)

	2003	2004	2005	2006	2007	2008	2009	2010	2011
2003									
2004	WF								
2005	F	WF							WS
2006	F*	F*	WF						S
2007	F*	F*	F*	WF		S		WF	F
2008	F	F*	WF	WF				WF	F
2009	F	F*	F*	WF		WF		F	F
2010	WF	WF							F
2011									

WF: no statistically significant dominance of column over row, statistically significant dominance for row over column for some range of poverty line.

F: first order dominance

S: second order dominance

Table 5: Poverty Indices, 2003-2010 – fixed poverty line, 60% of median 2007 income

Year	P₀	P₁	P₂
2003	0.273480	0.076970	0.032063
2004	0.263301	0.063769	0.022254
2005	0.244228	0.056857	0.019733
2006	0.220867	0.047143	0.015028
2007	0.170925	0.034029	0.010609
2008	0.158499	0.032806	0.011404
2009	0.138416	0.030122	0.011621
2010	0.177784	0.042139	0.018765
2011	0.211248	0.057624	0.027463

Table 6: Poverty Indices, 2003-2010 – relative poverty line, 60% of median income

Year	P₀	P₁	P₂
2003	0.212015	0.052117	0.021452
2004	0.214403	0.045720	0.015244
2005	0.201265	0.044659	0.015002
2006	0.187219	0.036553	0.011260
2007	0.170925	0.034029	0.010609
2008	0.157737	0.032591	0.011330
2009	0.149943	0.031861	0.012209
2010	0.147275	0.035399	0.016432
2011	0.154679	0.042398	0.021561

Table7: Sequential Stochastic Dominance, Income and Health

	2003	2004	2005	2006	2007	2008	2009	2010	2011
2003									
2004									
2005	WF	WF							
2006	WF	WF							
2007	WF	WF	WF	WF					
2008	F	WF	WF	WF	WF			WF	WF
2009	F	WF	WF	WF	WF	WF		WF	WF
2010	WF								WF
2011									

WF: no statistically significant dominance of column over row, statistically significant dominance for row over column for some range of poverty line.

F: first order dominance

Table 8: Bi-dimensional Poverty Indices

	Fixed Income Poverty Line				Relative Income Poverty Line			
	Intersection		Union		Intersection		Union	
	HPov=3	HPov=4	HPov=3	HPov=4	HPov=3	HPov=4	HPov=3	HPov=4
2003	0.020	0.089	0.288	0.353	0.014	0.068	0.233	0.312
2004	0.021	0.086	0.280	0.350	0.018	0.071	0.235	0.316
2005	0.017	0.073	0.263	0.342	0.014	0.062	0.223	0.310
2006	0.015	0.063	0.237	0.326	0.013	0.054	0.206	0.302
2007	0.010	0.051	0.187	0.279	0.010	0.051	0.187	0.279
2008	0.006	0.038	0.178	0.276	0.006	0.038	0.177	0.275
2009	0.006	0.034	0.160	0.270	0.006	0.037	0.171	0.279
2010	0.007	0.040	0.203	0.306	0.006	0.032	0.173	0.283
2011	0.007	0.045	0.232	0.332	0.005	0.030	0.178	0.291

Table 9: Dependence Measures Between Health and Income

	Total Distribution			Below Income Pov Line			Below Health Pov Line		
	PS	Sp	K _τ	PS	Sp	K _τ	PS	Sp	K _τ
2003	0.293 (0.015)	0.275*	0.212*	0.058 (0.025)	0.137*	0.104*	0.105 (0.051)	0.045	0.036
2004	0.295 (0.015)	0.296*	0.228*	0.044 (0.019)	0.120*	0.092*	0.167 (0.04)	0.089*	0.071*
2005	0.294 (0.011)	0.299*	0.230*	0.039 (0.018)	0.115*	0.088*	0.059 (0.035)	0.051 ⁺	0.041 ⁺
2006	0.273 (0.013)	0.274*	0.212*	0.017 (0.020)	0.059*	0.045*	0.179 (0.043)	0.096*	0.077*
2007	0.306 (0.014)	0.296*	0.229*	0.009 (0.021)	0.05*	0.038*	0.149 (0.049)	0.067*	0.054*
2008	0.262 (0.013)	0.261*	0.201*	-0.008 (0.020)	-0.013	-0.01	0.013 (0.04)	0.001	0.001
2009	0.248 (0.013)	0.243*	0.188*	-0.017 (0.021)	-0.035 ⁺	-0.027 ⁺	0.060 (0.04)	0.028	0.023
2010	0.214 (0.013)	0.210*	0.162*	-0.015 (0.021)	-0.035*	-0.027*	0.030 (0.043)	0.007	0.006
2011	0.216 (0.013)	0.199*	0.153*	-0.020 (0.022)	-0.045*	-0.035*	-0.003 (0.044)	-0.007	-0.006

Table 10: Dependence Measures Between Health and Income for Population Aged Under 65

	Total Distribution			Below Income Pov Line			Below Health Pov Line		
	PS	Sp	K _τ	PS	Sp	K _τ	PS	Sp	K _τ
2003	0.229 (0.018)	0.197*	0.154*	0.067 (0.030)	0.164*	0.126*	0.137 (0.065)	0.083 ⁺	0.066 ⁺
2004	0.233 (0.014)	0.224*	0.175*	0.048 (0.022)	0.141*	0.108*	0.192 (0.049)	0.108*	0.087*
2005	0.237 (0.013)	0.232*	0.181*	0.041 (0.021)	0.131*	0.101*	0.058 (0.046)	0.067 ⁺	0.053 ⁺
2006	0.222 (0.015)	0.219*	0.171*	0.026 (0.023)	0.088*	0.067*	0.229 (0.059)	0.136*	0.109*
2007	0.259 (0.017)	0.245*	0.192*	0.020 (0.025)	0.092*	0.070*	0.169 (0.068)	0.089*	0.072*
2008	0.220 (0.015)	0.216*	0.169*	0.005 (0.025)	0.028	0.021	0.045 (0.053)	0.020	0.017
2009	0.221 (0.005)	0.203*	0.159*	-0.009 (0.024)	-0.014	-0.011	0.083 (0.058)	0.055 ⁺⁺	0.045 ⁺⁺
2010	0.195 (0.016)	0.198*	0.154*	0.007 (0.025)	0.030	0.023	0.028 (0.054)	0.011	0.009
2011	0.206 (0.015)	0.189*	0.146*	0.006 (0.025)	0.023	0.018	0.079 (0.057)	0.036	0.027

Figure 1: Difference between Poverty Incidence Curves 2005-2006

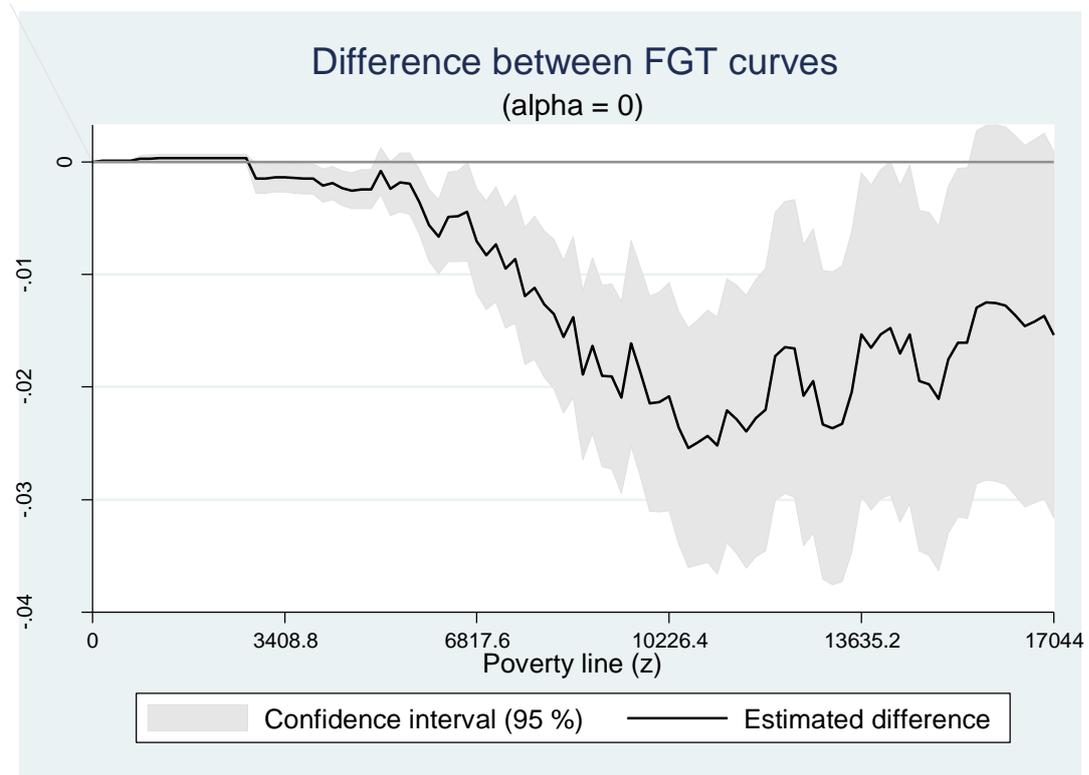


Figure 2: Difference between Poverty Incidence Curves 2005-2006 for neediest group

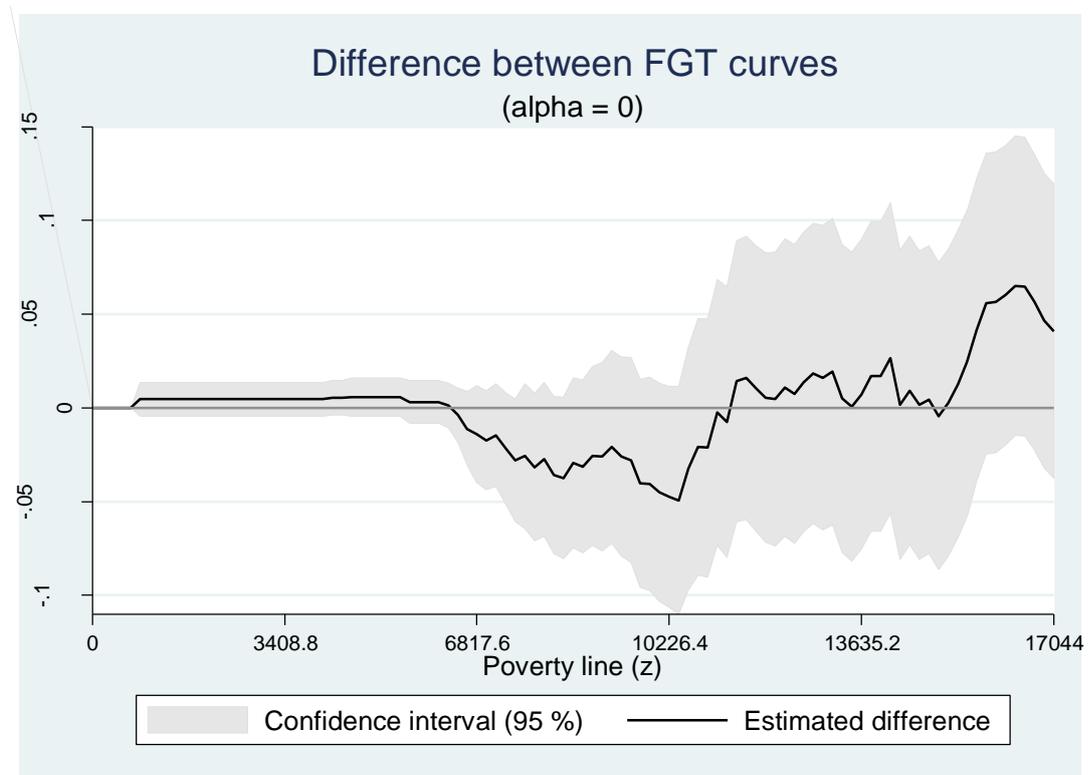


Fig 3: FGT, $\alpha=0$

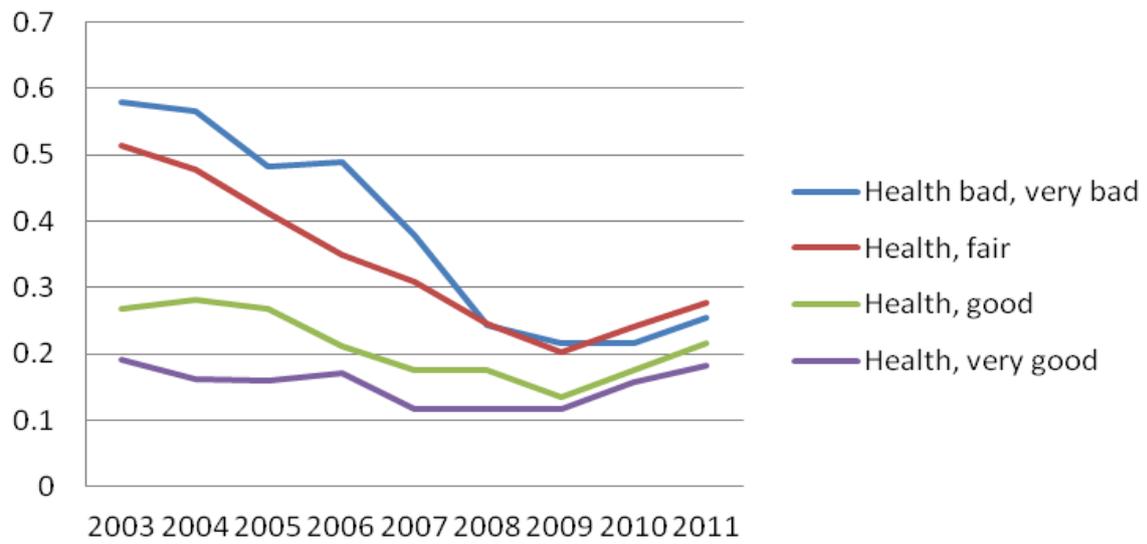


Fig 4: FGT, $\alpha=1$

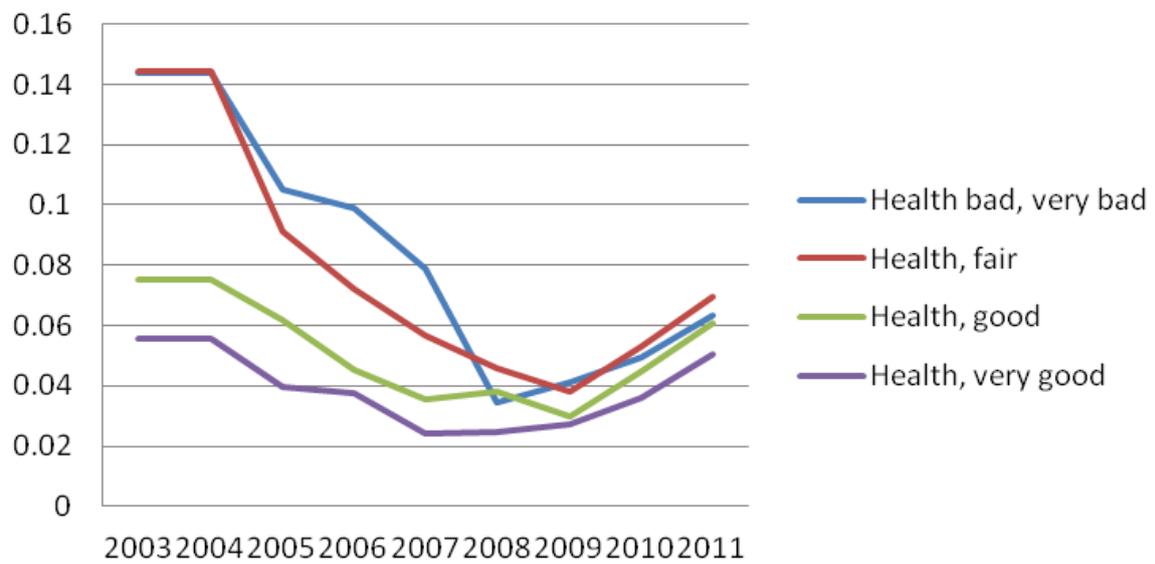


Fig 5: FGT, $\alpha=2$

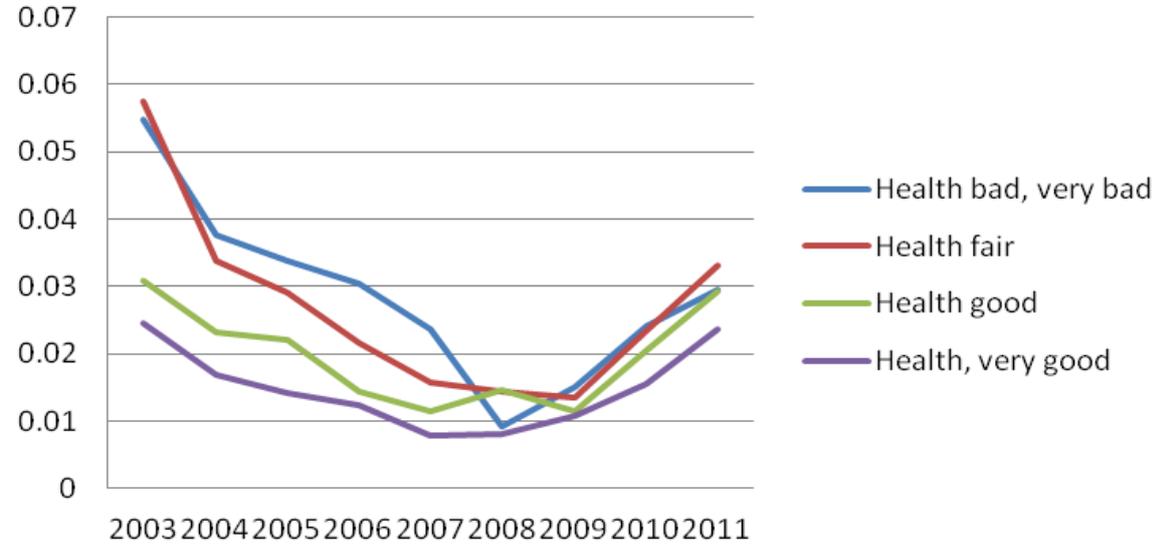


Fig 6: Measures of Dependence, All Ages

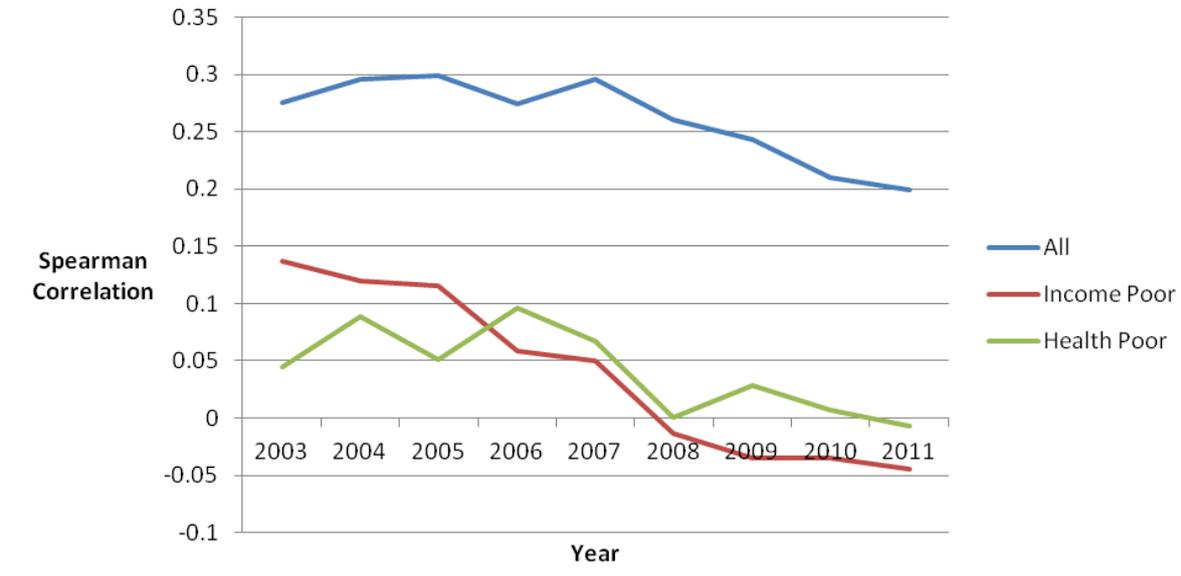
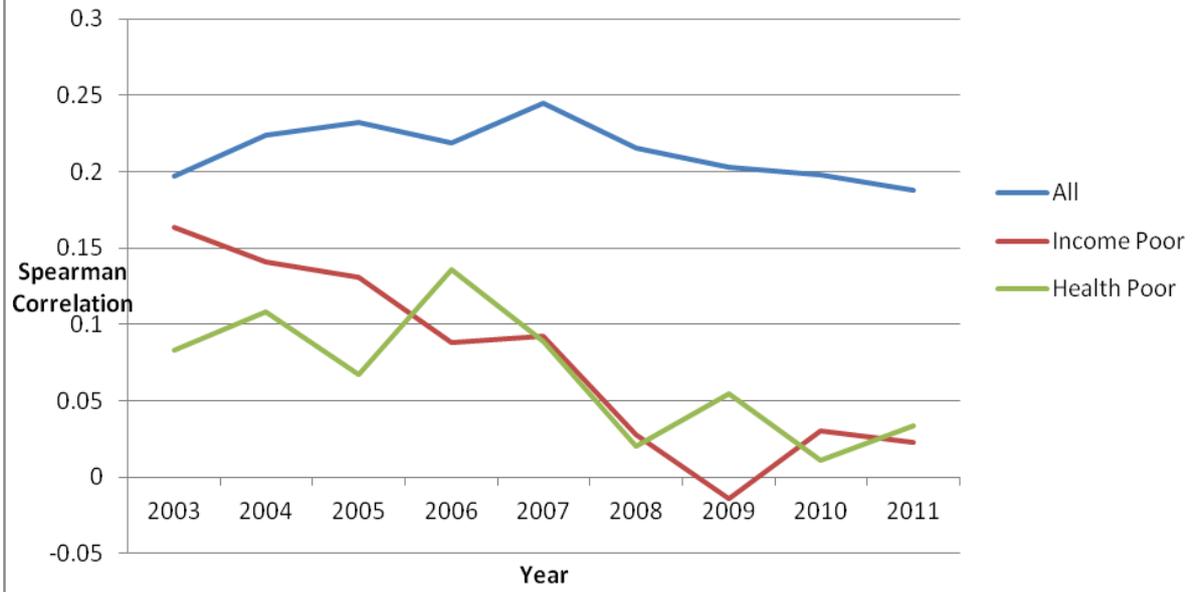


Fig 7: Measures of Dependence, Under 65s



Appendix 1: Definition of Income

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). The unit of analysis is all adults (i.e. those aged 16 and over) in the household.