Economic Issues, Vol. 4, Part 2, September 1999

Analysing Trends in UK Household Income Inequality: A New Approach

Derek Leslie and Alex Wharton[1]

Abstract

Using data derived from 30 years of the IFS Households Below Average Income Data Set, this paper examines the impact that the UK social security system had on UK household income inequality between 1961 and 1991. It asks whether rising inequality during the 1980s was due to a greater proportion of households being in receipt of social security or whether it was due to the fact that the system became less generous to households on social security because benefits were index linked rather than being linked to average living standards. Using a new, systematic procedure for decomposing changes in income inequality, strong evidence is found for a rising inequality 'numbers' effect, largely offset by an inequality-decreasing 'generosity' effect, during the 1970s and an uncompensated inequality-increasing generosity effect during the 1980s.

1. Introduction

Did UK household income inequality increase during the 1980s because of the fact that greater numbers of households were in receipt of social security benefit or did it increase because the social security system became less generous, relative to other income sources? Or did it increase for other reasons? The main aim of this paper is to develop a new, systematic method of answering this type of question. We show that there was a strong 'numbers' effect before 1980, but that rising inequality during the 1980s was largely the effect of decreasing 'generosity'.

Why is it necessary to take a systematic approach? The reason is that we can not simply assert how a known change in the social security system will affect income inequality. The following example illustrates this point. Suppose there are two types of household, households on social security (n_1) and households not on social security (n_2). Suppose further that the log of income from all sources (Y) is described by the equation

$$Y = \alpha + \gamma X + u, \qquad (1)$$

where X is a categorical variable, indicating that a household is either on (X=1) or not on (X=0) social security and u is a random disturbance term. Given equation (1), the variance of the log of income is

$$Var(Y) = \frac{n_2\gamma^2}{n} - (\frac{n_2}{n}\gamma)^2 + Var(u), \qquad (2)$$

where the ratio n_2/n gives the proportion of households on social security. This has a maximum value when n_2/n = 1/2. Equation (2) demonstrates that changes in the number of households on social security can either be inequality increasing or inequality reducing.

In this example, inequality is measured as

the variance of log income. However, in what follows the inequality measure will be a percentile range, specifically the 90th - 10th percentile of log income. It is clear from the example that the percentile range will also have a maximum when $n_2/n = 1/2$. Since, in reality, the distribution of household income is determined by a far more complex process than has been supposed here, it may be safely concluded that changes in the number of households on social security will have an unpredictable effect on income inequality, and that in general the impact on income inequality of a particular change in the social security system can only be determined by an ex-post analysis of the available data. Needless to say, a systematic approach, such as that taken here, is based on certain strong assumptions. Nevertheless, it does provide an intriguing new perspective on recent income inequality trends.

2. *The method of analysis*

Here the method is described in general terms, and the specific application will be outlined later. The technique was developed by Juhn, Murphy, and Pierce (JMP; 1991, 1993) to account for trends in US earnings inequality, but is used here to analyse UK household income inequality trends.[2] The basic idea is that changes in the distribution of income may be broken down into three constituent parts, accounting for changes in the distribution of observable household characteristics (the numbers effect), changes in the rewards associated with particular characteristics (the generosity effect), and changes in unobservables.

To be precise, let household income be described by the equation

$$Y_{it} = X_{it}\gamma_t + u_{it}, \qquad (3)$$

where Y_{it} is log weekly income for household i in year t, X_{it} is a vector of observed characteristics for household i in year t, γ is the corresponding vector of coefficients, and u_{it} is residual income accounted for by unobserved variables. For analytical reasons that will become clearer below, residual income may be broken down into two components, representing an individual household's percentile location in the residual distribution, θ_{it}, and the distribution function of the residuals, $F_t(\ )$. By definition of the cumulative distribution function,

$$u_{it} = F_t^{-1}(\theta_{it} \mid X_{it}), \qquad (4)$$

where $F_t^{-1}(\bullet \mid X_{it})$ is the inverse cumulative residual distribution for households with characteristics X_{it} in year t.

Next, construct two artificial distributions. The first distribution is given by the equation

$$Y_{it}^1 = X_{it}\bar{\gamma} + \bar{u}_{it}, \qquad (5)$$

where $\bar{\gamma}$ is a vector of average coefficients for the study period, and u_{it} is a vector of residuals taken for each household i from the average cumulative distribution of residuals for the same period. Since coefficients and residuals are effectively fixed, any change in the distribution of Y_{it}^1 is attributable only to changes in the distribution of X_{it}s.

The second distribution is constructed using the equation

$$Y_{it}^2 = X_{it}\gamma_t + \bar{u}_{it}. \qquad (6)$$

Changes in the distribution of Y_{it}^2 are now attributable to changes in the X_{it}s and γ_ts.

The assignment of the residuals u_{it} is central to the decomposition, and is made as follows. First, estimate equation (3) for every year t of

First, estimate equation (3) for every year t of the study period. Secondly, pool the residuals and arrange them in percentile order. Finally, assign an 'average' residual to each of the households in year t by matching the location of each household's estimated residual income in the distribution of year t residuals to the corresponding residual in the distribution of pooled residuals. For example, if the ith residual in equation (3) is located by the pth percentile of the u_t distribution, then the assigned residual u is located by the pth percentile of the distribution of pooled residuals. The JMP technique may be used to analyze changes at any chosen percentile difference. It is used here to analyze changes in the log household income differential between the 90th and 10th percentiles of the Y_{it} distribution. A change in this differential across any two years is denoted as ΔZ. If the corresponding change in the differentials of the Y_{it}^1 and Y_{it}^2 distributions are represented respectively by ΔZ^1 and ΔZ^2, then the decomposed change in Y_{it} is given by

$$\Delta Z \equiv \Delta Z^1 + (\Delta Z^2 - \Delta Z^1) + (\Delta Z - \Delta Z^2) \quad (7)$$

This identity breaks down the overall change in the spread of income into three parts. The first part is attributable to changes in the distribution of the Xs, the second part is attributable to changes in the γs, and the final part is attributable to changes in the distribution of residual income. No attempt is made to isolate the separate impacts of unobserved variables and their unobserved coefficients. Besides, Suen (1996) argues that this is a somewhat dubious exercise.

3. Application of the technique

Equation (3) is used to analyze trends in UK household income inequality in three separate ways. Model (1) is estimated in two ways and version (1) is

$$Y_{it} = \gamma_0 + \gamma_1 PHISE_{it} + \gamma_2 PHIPP_{it} + \gamma_3 PHII_{it} + \gamma_4 PHISSB_{it} + \gamma_5 PHOI_{it} + u_{it}, \quad (8)$$

where *PHISE* is the proportion of income obtained from self-employment, *PHIPP* is the proportion obtained from private pensions, PHII is the proportion obtained from investment income, *PHISSB* is the proportion obtained from social security, and *PHOI* is the proportion obtained from other sources. Income from employment is omitted because all the proportions would, by definition, sum to unity. The specification of this equation is based on the idea that income levels may be related to income composition. Here the proportions representhousehold characteristics and the coefficients represent the generosity with which particular sources are rewarded relative to earned income. This equation is admittedly *ad hoc* but turns out to fit the data quite well.

Equation (8) is simplified to give version (2) of model (1),

$$Y_{it} = \beta_0 + \beta_4 PHISSB_{it} + u_{it}. \quad (9)$$

It turns out that equation (9) fits the data almost as well as equation (8), suggesting that the social security system has been a major factor in income inequality changes. Equation (9) can then be used to answer the questions raised earlier. The coefficient β_4 represents the 'relative generosity' of the social security system. A fall through time in the value β_4 relative to β_0 indicates that the system has become 'less generous' as a source of income. *PHISSB*, representing the proportion of income obtained from social security, can range from zero to one for individual

indicates that either the numbers dependent on social security have increased (a change from zero to a positive number), or simply that a greater proportion of dependent households' income comes from the social security system. In short, it gives a summary measure of how changes in the importance of the social security system have impacted on income inequality, and underlies what is termed as the 'numbers' effect. The decomposition of equation (9) is a practical way of untangling these two separate influences (generosity and numbers) on income inequality trends.

Model (2) gives a different but complementary perspective on income inequality trends by linking household income to the kind of household characteristics (age, region, household size and so on) that might be found in a typical earnings function, rather than the income variables of model (1). The crucial difference between this and the usual earnings function is that here income refers to total after tax income from all sources and not just income from employment. Altogether we have 50 separate household characteristics, split into seven main categories. These can be seen in Table 3 and will be described more fully later.

A simple example describes how model (2) can provide a complementary picture about the 'numbers' and 'generosity' effect. Consider the variable 'age'. An ageing population, represented by an increase in the older age category, may imply that the number of pensioner households is increasing, with a consequent rise in inequality - a 'numbers' effect. Changes in the coefficient attached to the age variable show how well such households are rewarded relative to other groups. This would be the generosity effect. The same interpretation can be applied to the other variable groups that make up model (2).

4. The data

Data on household income and household characteristics are derived from 30 years of the Households Below Average Income (HBAI) data set, survey years 1961-1991.[3] Households need to be distinguished from benefit units. There are six benefit unit types in the HBAI data set, namely couple pensioner, single pensioner, couple with children, couple no children, single with children, and single no children. The typical household consists of a single benefit unit. However, there are a number of multiple benefit unit households. The HBAI data set assigns the entire household income to each of the benefit units within these households. This income is the total income received by all the separate benefit units that make up the household. Some analysts treat benefit units as the basic income unit. This is not possible here because the HBAI data set does not provide a way of assigning income to individual benefit units, other than the ad hoc method of equal division. Therefore, multiple benefit unit households are described by the Head of Household benefit unit. Model (2) which focuses on specific household characteristics always refers to the Head of Household's benefit unit in the case of multiple benefit households. The existence of multiple benefit unit households in a small number of cases is controlled for using additional explanatory variables for the number of benefit units in the household.

Household income is defined as the log of weekly disposable income in January 1994 prices. This is the total of after-tax receipts from employment, self-employment, private pensions, investment, social security benefits, and other sources across all the individuals within a household.[4] To maintain the consistency of total net income summing to its component parts, income is effectively equivalent to the HBAI before-housing-cost

measure of income. The only differences are that income here includes local taxes and is not MIRAS adjusted, the reason being that it is not clear from which income component these two categories should be subtracted. In practice, this makes little difference.

The following data are excluded from our inquiry; data on households whose total income is non-positive, and data on households for whom any one component of income is negative. Once again this is a very small number of observations. 1964 data are omitted for model (2) because region codes are not available. The advantage of the HBAI data set is its consistency, which enables a comparison to be made of results over time. Nevertheless, the sample sizes of Family Expenditure Survey data prior to 1968 were relatively small. Since HBAI data are derived from FES data, the pre-1968 results should be treated with caution. This problem is more apparent with model (2).

Due to non-response, the HBAI data set does not provide a representative sample of the population. This problem can be dealt with using grossing factors. However, as the aim of this decomposition is not to make statements about the proportion of the population that are in poverty, the issue of non-response is not a major difficulty. Only if the non-response changed in a systematic way through time could the decomposition produce potentially misleading results.[5]

Finally, income data are not equivalised. This is because of the specific question to be addressed. We want to explore how changes in the social security system impacted on the distribution of total (non-equivalised) household income. We are not specifically concerned with living standards. Anyway, according to Jenkins (1995), the general picture about trends is robust to the choice of inequality index, definition of income used, receiving unit, or equivalence scale used.[6]

5. Results

Models (1) and (2) were estimated for each of the years 1961 to 1991 (except 1964). An example of the first version of model (1) is shown in table (1) for 1991 data.

The model explains 48 per cent of the variance of income. This compares more than favourably with typical estimates of the popular human capital earnings model.[7] The intercept implies that the level of household income is positively related to the proportion of income coming from employment; on average, the greater the proportion of income from employment, the greater is the level of household income. In the absence of income from all other sources, households in 1991 would have expected to earn approximately £19,000 per annum from employment (January 1994 prices, calculated as 52*exp(5.8914)). The other coefficients indicate that households with income from investments, social security benefit, and other sources (but not self-employment income and private pensions), received significantly less income than those on earnings alone. For example, if all income came from social security recipients alone table (1) predicts this as approximately £5,000 (calculated as 52*exp(5.8914 - 1.3210)). The mean annual income of households in 1991 was approximately £11,500.

An example of the second version of model (1), which consists of the single explanatory variable, namely the proportion of income received from social security, is shown in table (2), again for 1991. Version (2) explains 47 per cent of the total variation of income and supports the earlier contention that the parsimonious specification works almost as well as version (1). Note that the coefficient PHISSB is very similar to that in table (1),

Table 1.

Estimates of the influence of income source on log (weekly) household income in 1991.

Explanatory Variable	Coefficient	Standard Error
Constant	5.8914*	0.0115
Proportion of household income from:		
Self-employment (HISE)	-0.0636	0.0379
Private pensions (HIPP)	-0.0834	0.0501
Investment income (PHII)	-0.3096*	0.0541
Social security benefit (PHISSB)	-1.3210*	0.0195
Other sources (PHOI)	-0.7101*	0.0705
N = 5139		
R^2 = 0.4838		

* Significant at 1 per cent level

Table 2.

Estimates of the influence of social security benefit on log (weekly) household income in 1991.

Explanatory Variable	Coefficient	Standard Error
Constant	5.8394*	0.0097
Proportion of household income from:		
Social security benefit (PHISSB)	-1.2907*	0.0191
N = 5139		
R^2 = 0.4703		

* Significant at 1 per cent level.

Table 3.

Estimates of the influence of observed household characteristics on log (weekly) household income in 1991.

Explanatory Variable	Coefficient	Standard Error
Constant	6.0943*	0.0254
Benefit unit type		
Couple pensioner	0.1503*	0.0471
Single pensioner	-0.3788*	0.0441
Couple, with or without children (BENCHMARK)		
Single, with or without children	-0.4817*	0.0208
Economic status		
Single or couple, at least one in full-time self-employment	-0.2894*	0.0275
Single or couple, in full-time employment (BENCHMARK)		
Couple, one in full-time employment, one in part-time employment	-0.2274*	0.0261
Couple, one in full-time employment, one not working	-0.3388*	0.0254
One or more in part-time work	-0.5225*	0.0322
Head or spouse aged 60 or over	-0.7638*	0.0377
Head or spouse unemployed	-1.0117*	0.0360
Others	-0.6753*	0.0302
Household tenure type		
Local authority rented unfurnished or housing association	-0.2807*	0.0197
Other rented unfurnished	-0.2423*	0.0375
Rented furnished	-0.1791*	0.0385
Owned with mortgage, including owned by rental purchase (BENCHMARK)		
Owned outright	-0.0995*	0.0207
Rent free	-0.2769*	0.0457

Region		
Northern	-0.1824*	0.0302
Yorks and Humberside	-0.1643*	0.0262
North Western	-0.1178*	0.0244
East Midlands	-0.1844*	0.0282
West Midlands	-0.1776*	0.0258
East Anglia	-0.1609*	0.0364
Greater London	0.0536**	0.0245
South East (BENCHMARK)		
South Western	-0.0966*	0.0262
Wales	-0.2004*	0.0330
Scotland	-0.0761*	0.0262
Northern Ireland	-0.3141*	0.0482
Age of head		
Aged 25 or less	-0.1337*	0.0294
Aged 26 to 35 (BENCHMARK)		
Aged 36 to 45	0.0142	0.0213
Aged 46 to 55	-0.0157	0.0236
Aged 56 to 60	0.0297	0.0319
Aged 61 to 65	0.0871**	0.0395
Aged 66 or more	-0.0338	0.0526
Number of children		
Zero (BENCHMARK)		
One	0.0520**	0.0226
Two	0.1163*	0.0243
Three	0.1514*	0.0357
Four or more	0.3215*	0.0584
Number of Benefit units		
One (BENCHMARK)		
Two	0.4374*	0.0188
Three or more	0.6439*	0.0347
N = 5139		
R^2 = 0.6008		

* Significant at 1 per cent confidence ** Significant at 5 per cent level level

further supporting the view that the simple specification is an adequate one. Given the well-known unreliability of self-employment income and a tendency to under-report its value, a simple model that ignores this income source is additionally useful. The decomposition analysis will subsequently confirm little difference between these two versions of model (1). The practical advantage of this is that, because it contains only one specific variable, it makes the JMP decomposition easy to interpret.

The corresponding example for model (2) is shown in table (3). Here income is explained by conventional household characteristics, rather than income sources. The first group of variables is called benefit unit type, with the default category shown as BENCHMARK. The other groups, again with the benchmark category indicated, are economic status, household tenure type, region, age, and number of children. The final group is number of benefit units in the household. This can be seen to have had an income-enhancing effect on household income, as might be expected.

Care must be taken in interpreting the results of model (2). For example, it would appear that a couple pensioner did rather better than the default category of working-age couples. However, most couple pensioners would be in the 'others' economic status category, mainly those economically inactive, which can be seen to have a large negative impact on household income.

The results imply that in 1991 the income of the benchmark household was significantly higher than the income of most other represented households. For example, the estimated intercept implies that the mean, annual, disposable income of the benchmark household (see above) was approximately £23,000 in January 1994 prices. By comparison, the income of a single, unemployed person, aged less than 26 years old, without children, and renting accommodation from a local authority in the North West, was about £3,000. The exceptions comprise variants of the benchmark household. Model (2) explains approximately 60 per cent of the variation in household income, which is once again rather better than typical estimates of the human capital model for earned income.

6. Decomposition analysis

A visual representation of the decomposition is the most informative way to identify broad trends. Figures (1) and (2) refer to versions (1) and (2) of model (1) and figure (3) refers to model (2). 1961 is the benchmark year indicated by 100. All comparisons are relative to 1961. Four series are shown, representing each of the four terms in equation (7). The first series is the overall change in inequality; a positive slope indicates rising inequality, and a negative slope indicates falling inequality. The parts accounted for by changes in characteristics, coefficients and unobservables are also shown. For any specific year, these components sum to the overall change in inequality.

Between 1961 and 1983 the overall income inequality series exhibits a clear, cyclical pattern. It shows income inequality fluctuating between about plus five per cent and minus ten per cent of its 1961 level, plotting roughly two peak-to-peak cycles, averaging eight years duration. It shows increasing income inequality in the periods 1961-1963, 1967-1972, 1976-1980, and decreasing income inequality in the periods 1963-1967, 1972-1976, and 1980-1984. Between 1984 and 1991 the series exhibits a strikingly different pattern. It shows a persistent annual increase in income inequality

Figure 1. The impact of changes in both the composition of household income and the generosity of employment relative to other sources of income.

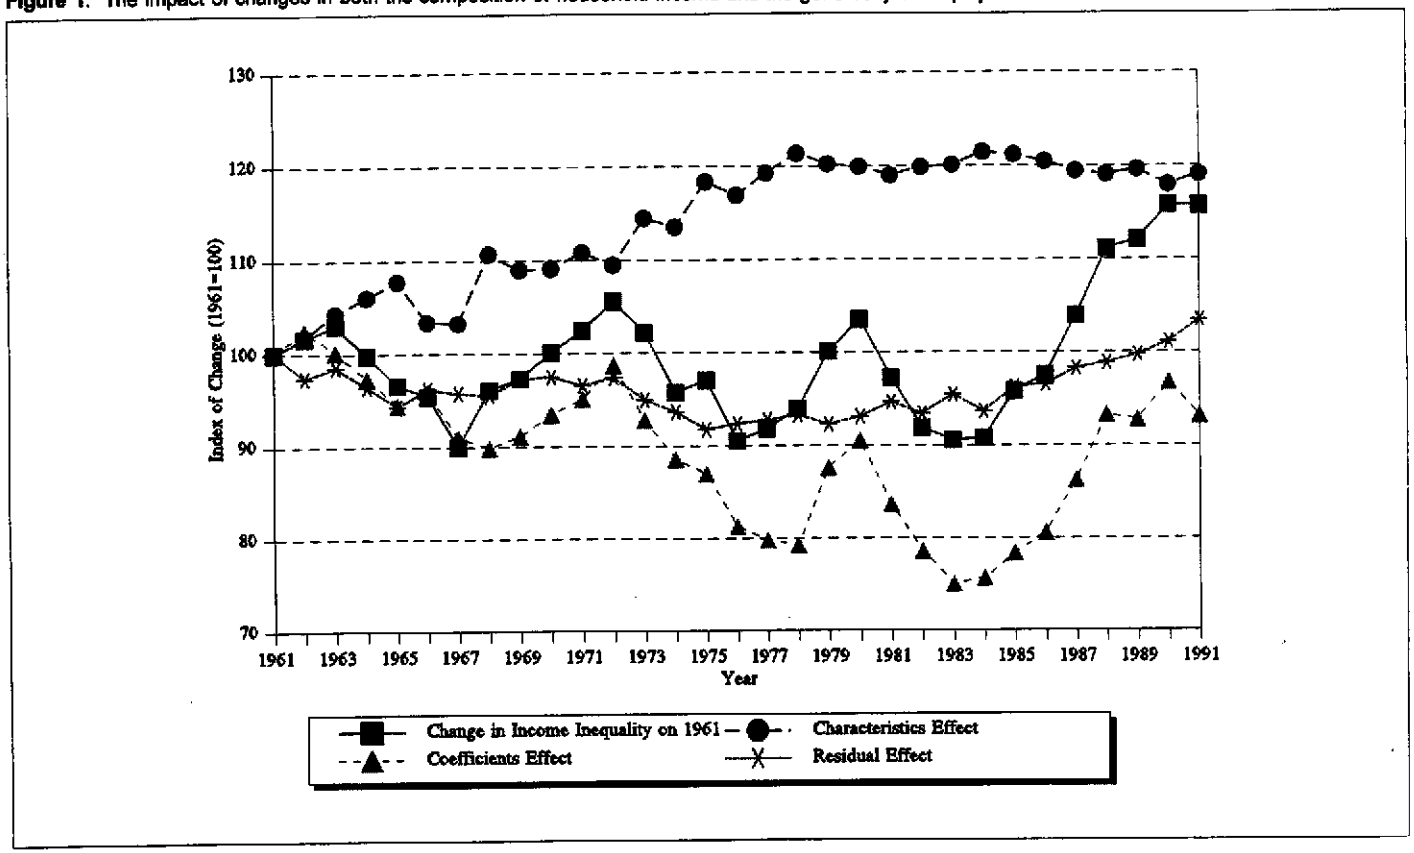

Figure 2. The impact of the social security system.

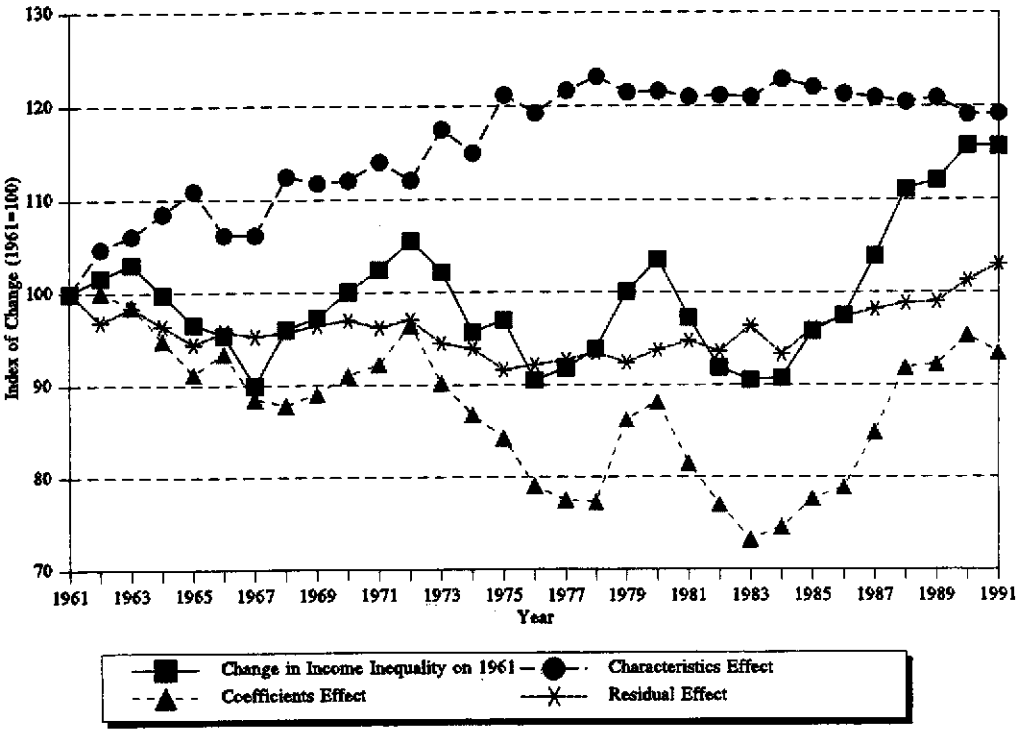

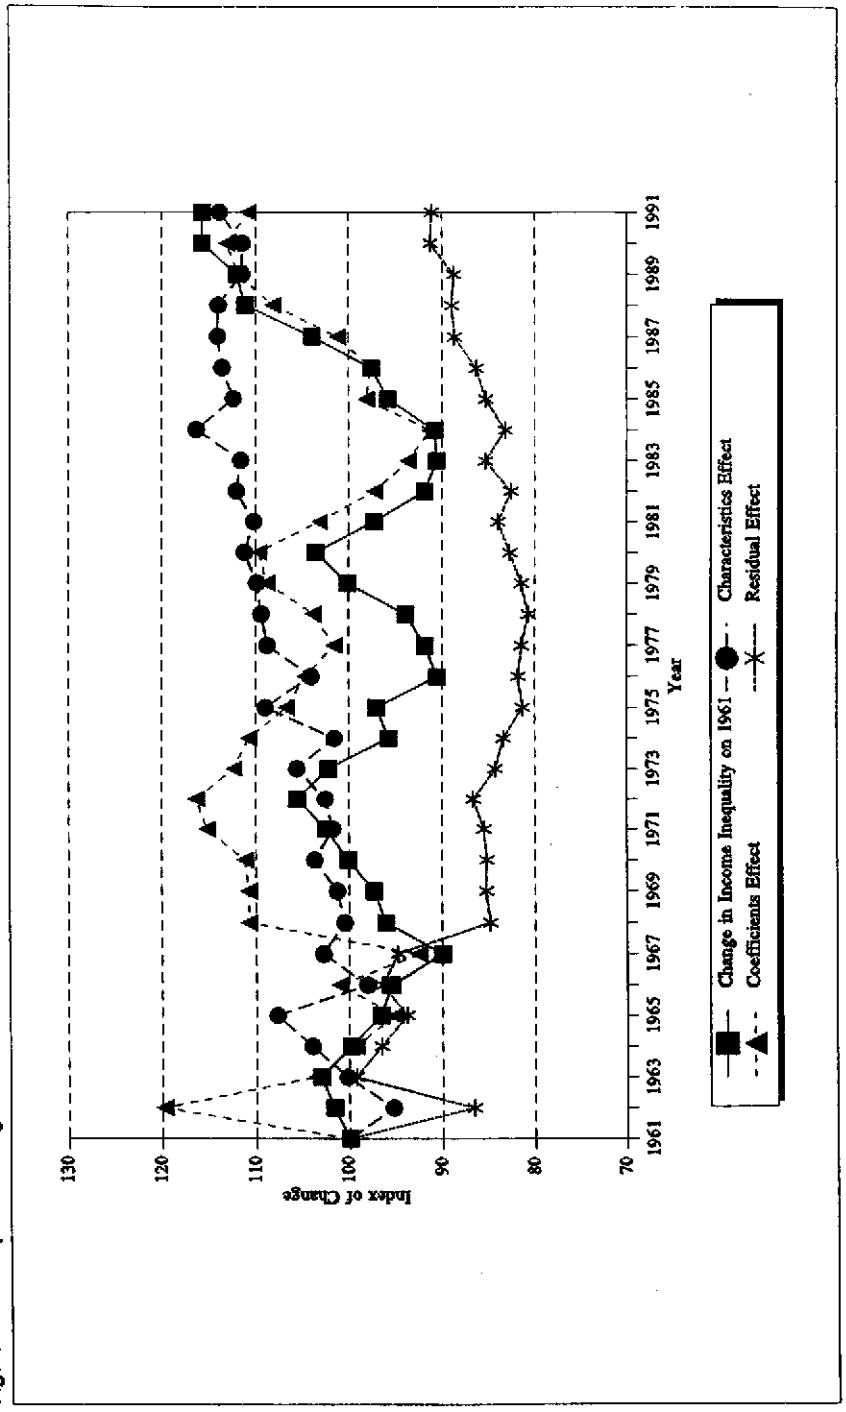

Figure 3. The impact of changes in both the distribution and relative incomes of household characteristics.

of approximately four per cent per annum, giving a six year increase of about 24 per cent. This mirrors trends identified by other inequality series, despite the several adjustments we made for the particular purpose to be considered here.[8]

Figures (1) and (2) are virtually identical, suggesting that we can safely ignore fig. (1), and focus instead on the questions addressed specifically by the second decomposition based on the parsimonious specification. In fig. (2), the characteristics effect shows how changes in the composition of income contributed to inequality changes (measured at the 90th-10th percentile). It suggests that there was an inequality enhancing characteristics effect up to 1978; but that after 1978 characteristics had a very modest effect on income inequality changes.

The coefficients series shows the component of the change in overall inequality accounted for by changes in the (relative) generosity of the social security system. As the plot of the series demonstrates, changes in the system's generosity had a considerable impact on income inequality throughout the study period, 1961-1991. Between 1961 and 1982, these changes tended to reduce inequality. However, between 1982 and 1991, they were clearly inequality increasing, accounting for roughly three-quarters of the total increase in income inequality during the 1980s.

The residual series describes the impact that changes in unobservables have had on income inequality. Up until 1984 changes in unobservables appear to have had a very modest effect. However, after 1984 they increased the differential by about seven percentage points. Jenkins (1996) suggests that this may be due to increasing inequality within the distribution of self-employment income.

Although figure (2) shows the impact of component parts on the overall change in inequality, it says nothing about how the actual values of these components changed. The example at the outset of the paper argued that it is not clear how a rising or decreasing proportion of income coming from social security will affect inequality.

Figure (4) sheds more light on this. It plots the average value of *PHISSB* year-by-year, changes in which cause the characteristics (numbers) effect. It also plots the ratio of the coefficients from equation (9). This ratio shows the relative generosity of the social security system, and is called the Generosity Index. An increase shows the system becoming more generous and a decrease indicates declining generosity.

Figure (4) exhibits distinct patterns that should be viewed in conjunction with figure (2). *PHISSB* shows two distinct phases. Up to 1983, the proportion of income obtained from the social security system steadily increased, before going into steady decline thereafter. Notice that this does not exactly translate into a similar change in inequality. For example, *PHISSB* increases in the 1975 1984 period, yet seems to have no effect on inequality. Only after 1984 when the share starts to fall does it have a modest inequality-reducing role. This supports the contention made at the outset, namely it is not possible *a priori* to predict the impact of changes in the social security system on income inequality.

The generosity line in figure (4) tells a similarly interesting story. There are two distinct phases in its development, indicating that the social security system became more generous up to 1983/84, and less generous thereafter. Note that the generosity effect was tending to have an inequality-reducing impact up to 1984, and an inequality-increasing

Figure 4. The impact of the social security system: beneath the numbers and generosity effects.

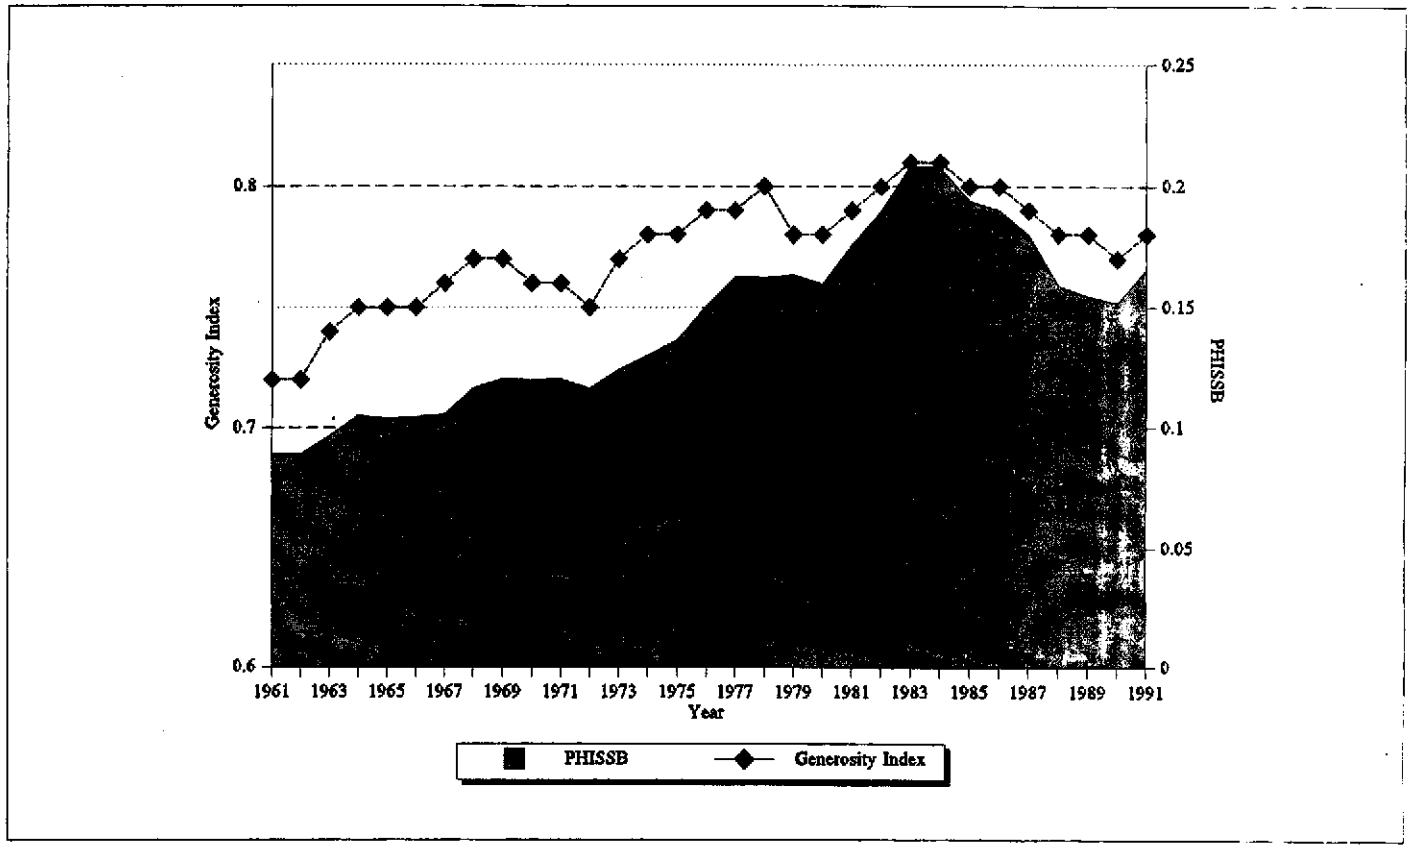

impact thereafter. So it seems that a less generous system is, in practice, associated with a more unequal income distribution. Note that in the example of equation (2), income becomes unambiguously more unequal if $ß_4$ declines in value relative to $ß_0$.

Figure (3) displays the decomposition of model (2). This shows three distinct periods, namely 1961 to 1968, 1968 to 1980, and 1980 to 1991. During the first period, falling income inequality appears to be best explained by changes in unobservables, but the results are fairly erratic reflecting the small sample sizes. Indeed, in this context, neither the subsequent sudden upswing of the relative incomes effect between 1967 and 1968 nor the corresponding downswing of the residual effect are surprising. During the second period, changes in income inequality are best explained by a combination of changes, both in the distribution of household characteristics and the relative incomes of those characteristics. The tendency for the two effects when combined to overstate the actual level of inequality change is explained by the inequality-reducing residual effect of 1968. Finally, the inequality-increasing tendency of the numbers effect seems to have been largely cancelled out by the tendency of the relative incomes effect to reduce inequality. After 1980, changes in income inequality are explained by changes in the relative incomes (generosity) effect and, to a lesser extent, by changes in unobservables.

7. Conclusion

The social security system was a major factor in income inequality changes throughout the period 1961 to 1991. During the 1960s and 1970s, changes in the importance of the social security system as a source of income (increasing numbers of dependents, increasing dependence) gave rise to a strong inequality-enhancing numbers effect, largely offset by an inequality-reducing generosity effect. More significantly, the substantial and extraordinary rise in UK household income inequality during the 1980s reported here and elsewhere is best explained by the fact that the social security system became less generous to households on social security.

According to Johnson (1996), the inequality-increasing generosity effect of the 1980s may be attributed in large part to the price indexation of social security benefits, as well as the tax changes which favoured those on high incomes. Since the average incomes of workers rose in real terms during this period, average index-linked benefits necessarily fell in relative terms. Consequently the income gap between households dependent on social security benefits and the rest of the population grew. This study has confirmed that the evolution of the social security system had an important influence on income inequality.

Endnotes

1. Manchester Metropolitan University and University of Manchester Institute of Science and Technology (UMIST), respectively. The IFS 'Households Below Average Income Data Set' is made available by the University of Essex ESRC Data Archive and is gratefully acknowledged. The authors wish to thank the editor and anonymous referees for useful comments. Remaining errors and omissions are their own.

2. Examples of previous applications of the JMP technique are Blau and Khan (1992,1993), Blackaby et al. (1997), Goldin and Margo (1992) and Margo (1995).

3. Goodman et al. (1995a). Households Below Average Income is the data set name; it does not mean that the sample consists of only those below average income.

4. See Goodman *et al.* (1994b), Appendix four. This gives definitions of income sources.

5. As well as using grossing factors the HBAI data set adjusts very high incomes for non-response. Once again, the consistency requirement precludes this adjustment, but only a very small number of households are involved.

6. The estimates were all repeated using the 75th - 25th percentile decomposition, and the same basic pattern was found.

7. Blackaby et al. (1998) is a good contemporary example using British data.

8. Goodman and Webb (1994) use the same HBAI data to analyse changes in a Gini coefficient measure of inequality. Their Figure (1) shows the same cyclical pattern as that shown here, with a sharp rise in inequality after 1984. They point to on-off incomes policies in the early period as one cause of the cyclical compression effect. Jenkins (1995,1996), Atkinson (1994), Johnson (1996), DSS (1994) also analyse the post-1980 rise in inequality.

References

Atkinson A B (1994) 'What is Happening to the Distribution of Income in the UK?' *Proceedings of the British Academy*, vol. 82, 317-51.

Blau F D and Kahn L M (1992) 'The Gender Earnings Gap: Learning from International Comparisons', *American Economic Review*, vol. 82, 533-38.

Blau F D and Kahn L M (1993) 'The Changing Composition of the U.S. Labor Force: Implications for Public Policy', Working Paper.

Blackaby D H, Clark K, Leslie D G and Murphy P D (1997) 'The Distribution of Male and Female Earnings, 1973 1991: Evidence for Britain', *Oxford Economic Papers*, vol. 49, 256-72.

Blackaby D H, Clark K, Leslie D G and Murphy P D (1998) 'The Ethnic Wage Gap and Employment in Britain in the 1990s: Evidence for Britain', *Economics Letters*, vol. 58, 97-103.

Department of Social Security (1994) *Households Below Average Income: A Statistical Analysis, 1979 1991/2*, London: HMSO.

Goldin C and Margo R A (1992) 'The Great Compression: The Wage Structure in the United States at Mid-Century', *Quarterly Journal of Economics*, vol. 107, 1-34.

Goodman A and Webb S (1994) 'For Richer, For Poorer: The Changing Distribution of Income in the UK, 1961-1991', *Fiscal Studies*, vol. 15, 29-62.

Goodman A *et al.* (1995a) *Households Below Average Income Dataset, 1961-1991*, [computer file]. Colchester: ESRC Data Archive.

Goodman A *et al.* (1995b) *Households Below Average Income Data Set, 1961-1991: User Guide*. Colchester: ESRC Data Archive.

Jenkins S P (1995) 'Accounting for Inequality

Trends: Decomposition Analysis for the UK, 1971-1986', *Economica*, vol. 62, 29-63.

Jenkins S P (1996) 'Recent Trends in the UK Income Distribution: What Happened and Why?', *Oxford Review of Economic Policy*, vol. 12, 29-46.

Johnson P (1996) 'The Assessment: Inequality', *Oxford Review of Economic Policy*, vol. 12, 1-14.

Juhn C, Murphy K M and Pierce B (1991) 'Accounting for the Slowdown in Black-White Wage Convergence', in M H Kosters, (ed) *Workers and Their Wages: Changing Patterns in the United States*, Washington: American Enterprise Inst. Press.

Juhn C, Murphy K M and Pierce B (1993) 'Wage Inequality and the Rise in the Returns to Skill', *Journal of Political Economy*, vol. 101, 410-42.

Margo R A (1995) 'Explaining Black-White Wage Convergence, 1940 1950: The Role of the Great Compression', *Industrial and Labour Relations Review*, vol. 48, 470-81.

Suen W (1996) 'Decomposing Wage Residuals: Unmeasured Skill or Statistical Artefact', Working Paper, Department of Economics, University of Hong Kong.