DISTRIBUTIONAL ASPECTS OF INFLATION

Ian Crawford
Zoë Smith

THE INSTITUTE FOR FISCAL STUDIES
Commentary 90

Distributional aspects of inflation

Ian Crawford
Zoë Smith

Institute for Fiscal Studies

Copy-edited by Judith Payne

The Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE

Published by

The Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
Tel: +44-20-7291 4800
Fax: +44-20-7323 4780
Email: mailbox@ifs.org.uk
Internet: www.ifs.org.uk

ISBN 1-903274-26-5

Printed by

KKS Printing
The Printworks
12–20 Rosina Street
London E9 6JE

Preface

Financial support from the Nuffield Foundation (grant reference OPD/00101/G) is gratefully acknowledged. Material from the Family Expenditure Survey, made available by the Office for National Statistics (ONS) through the ESRC Data Archive, has been used by permission of the Controller of HMSO. Neither the ONS nor the ESRC Data Archive bears responsibility for the analysis or the interpretation of the data reported here. The authors are very grateful to James Banks, Erich Battistin, Laura Blow, Tom Clark, Andrew Dilnot and Sharon Witherspoon for helpful advice and comments. The views expressed here are those of the authors and not of the Institute for Fiscal Studies, which has no corporate view. The authors are responsible for all errors.

Contents

Summary

The main measure of inflation used in the UK is the retail price index (RPI). One way to think of the RPI is as a measure of the changing cost of buying a very large shopping basket containing all of the purchases of a typical UK household. There is, of course, no such thing as a typical household. As a result, while the RPI may or may not provide a good measure of the average rate of inflation, it would be remarkable indeed if it were a good measure of inflation for everybody. Nevertheless, according to *The RPI Technical Manual*, 'The RPI, while not applying to any one household or person, will be close to the experience of inflation for the great majority of households'. So how true is this claim?

We find that

- The headline average rate of inflation is not necessarily a good guide to the actual rates of inflation faced by individual households and it is certainly not always close to the experience of inflation for 'the great majority of households'. On average over the period from 1976 to 2000, only about a third of households at a point in time faced inflation rates within 1 percentage point of the average rate. However, this proportion has been as low as 9 per cent in 1989 and the highest it has been is 65 per cent in 1994. The representativeness of the average rate tends to be lower when inflation is high.

- Over the period from 1976 to 2000, the average annual inflation rate for the poorest 10 per cent of households was 6.8 per cent, while the average annual inflation rate for the richest 10 per cent was 7.1 per cent. In any one year, differences in inflation rates can be much larger although the ranking often changes.

- Over the whole period, non-pensioners, mortgagors, the employed and childless households experienced higher-than-average inflation.

- The overall effect of average price indexation on the real value of the pension since 1987 has been a real increase of around 3.5 per cent. This means that the basic state pension is around £2.27 higher than it would have been if indexation had been carried out using an inflation rate for pensioners.

We analyse the impact of ignoring differential inflation on the measurement of income inequality and find that

- Not taking into account differential inflation can lead to misleading conclusions about the growth in inequality from one year to the next. For the period from 1977 to 2000, we find that not allowing for differential inflation could lead to the annual growth rate in inequality being overstated or understated by as much as 6 percentage points.

1. Introduction

The main measure of inflation used in the UK is the retail price index (RPI). The RPI is defined as an 'average measure of change in the prices of goods and services bought for the purposes of consumption by the vast majority of households in the UK'.[1] One way to think of the RPI is as a measure of the changing cost of buying a very large shopping basket containing all of the purchases of a typical UK household. There is, of course, no such thing as a typical household. As a result, while the RPI may or may not provide a good measure of the average rate of inflation, it would be remarkable indeed if it were a good measure of inflation for everybody. Nevertheless, according to *The RPI Technical Manual*, 'The RPI, while not applying to any one household or person, will be close to the experience of inflation for the great majority of households'.[2] The variation in inflation rates, if it is significant, is a matter of interest and importance because of the vital role of the RPI in, for example, the uprating of pensions, benefits and tax allowances and in wage settlements.

The aim of this Commentary is to analyse the differences in the inflation rates experienced by different households in the UK, to see exactly how close the average rate of inflation is to the experience of inflation for the great majority of households, and to consider the implications of the dispersion in inflation rates across households for the measurement of economic inequality.

[1] Baxter, 1998, p. 3.

[2] Baxter, 1998, p. 1.

2. Measuring inflation

A price index is supposed to be representative, in some way, of all the price changes observed between two periods.[3] There are a number of general approaches that can be used to calculate price index numbers.[4] In order to understand why we have taken the approach that we have, it is necessary to know a little about the data that we will be using in this study (we will give more details in due course). The price data are the monthly section indices of the RPI. These are the most disaggregated, comprehensive set of price data published in the UK. They only vary over time; no regional disaggregation, for example, is published. Our quantity data are from the UK Family Expenditure Survey (FES). The FES is the primary source of data used in the compilation of the RPI. It is an annual random cross-sectional survey of around 7,000 households a year. As well as collecting information about the characteristics of each household, the FES collects detailed expenditure information. This is recorded mainly by members of each household keeping a diary record of what they spend over two weeks. Households are surveyed only once.

For each household in the FES, we observe $\mathbf{q}_t^h$ (the list of quantities purchased by household h observed in period t). We also observe, in the section indices, the corresponding list of prices prevailing both in the month in which they were surveyed ($\mathbf{p}_t$) and in all other months. So, to compute a rate of inflation between two periods (t and, let us say, t+1) for this household, we have three pieces of information at our disposal: $\mathbf{q}_t^h$, $\mathbf{p}_t$ and $\mathbf{p}_{t+1}$. The simplest measure of inflation in the period to t+1 for this household (denoted by π_{t+1}^h) is

$$1+\pi_{t+1}^h = \frac{\mathbf{p}_{t+1}'\mathbf{q}_t^h}{\mathbf{p}_t'\mathbf{q}_t^h}.$$

This is the Laspeyres price index number formula and is the measure we use in this Commentary. It compares the cost of buying the observed set of commodities ($\mathbf{q}_t^h$) at two different sets of prices: the contemporaneous set ($\mathbf{p}_t$) and the set prevailing in the following period ($\mathbf{p}_{t+1}$).

The Laspeyres price index is not a true cost-of-living index. A cost-of-living index measures the average change in prices with reference, not to a fixed list of demands, but to a fixed standard of living. The Laspeyres price index is the upper bound of a true cost-of-living index in which the base-period welfare is the reference level. The reasoning behind this is that there may be a number of combinations of commodity demands that all make a household equally well off (in economic terms). Buying the first-period

[3] In fact, a price index that represents price changes observed between two periods is called a bilateral price index – the alternative is a multilateral price index, which is supposed to be representative of price changes observed in a number of periods.

[4] A complete review of these alternative approaches is not necessary here – we refer interested readers to the papers collected in Diewert and Nakamura (1993) and the references therein.

combination of goods is one way to achieve the base-period level of welfare, but the household may be able to achieve this level of welfare at lower cost by buying a different combination. Since the household will choose the combination (given the prices it faces) that minimises[5] the cost of achieving the reference level of welfare, the true cost-of-living index must be smaller than or equal to the Laspeyres price index.

There is a great deal of debate over the relative merits of fixed-base price indices and cost-of-living indices (see, for example, Triplett (2001) and Schultze and Mackie (2002)). Ultimately, the use of a cost-of-living index instead of our fixed-base price index rests on the assumption that households will always minimise the cost of reaching a given level of economic welfare. *This assumption may not be right*, even approximately. In principle, it is possible to test whether consumers do cost-minimise,[6] but in order to do this we would need to observe each household's demands at least twice and we only observe each household once; hence we would not be able to test the model so it would have to be taken on trust. Even if we do believe the principle that households always minimise the cost of reaching a given level of economic welfare, in order to use a true cost-of-living index we would have to choose an underlying model of consumer behaviour. There are many models of consumer behaviour to choose from and some of them rest on assumptions that are usually considered to be very restrictive and unrealistic. More general models do exist, but the cost-of-living indices that are associated with them[7] require us to observe households' consumption patterns in two different time periods, which we do not.[8] Using a Laspeyres index to measure inflation for each household is simple, it can be applied directly to our data and it can be regarded as an upper bound on the true, idiosyncratic, cost-of-living index for each household if we want to assume that such a notion is meaningful.

So far, we have discussed our individual, household-specific inflation measure. Given that in our data all households face the same prices, the differences in inflation rates across different households are generated by differences in their commodity demands (and the quality of the approximation to their true indices will depend upon the heterogeneity in their substitution responses). We now consider how to aggregate these household measures together into group indices. There are two broad approaches to calculating average inflation rates: democratic and plutocratic. Democratic indices weight sample households equally and give straightforward means (arithmetic and geometric being the most popular). The plutocratic approach weights households according to their share of total sample expenditure, which means that richer households receive more weight. Again, mostly arithmetic and geometric versions of this are considered.[9] We

[5] If the household did not do this, then its behaviour would be wasteful in the sense that it could improve its economic welfare simply by reallocating its consumption patterns.

[6] See, for example, Afriat (1967 and 1973), Houthakker (1950), Samuelson (1938), Varian (1982) and Blundell, Browning and Crawford (forthcoming).

[7] For example, one possibility is the Tornqvist index, which is exactly the same as the true cost-of-living index based at a certain level of welfare if consumer behaviour can be described by the translog cost function.

[8] Alternatively, we could predict the second-period demands using the model of consumer demand, but we would have to assume that household demands were conditionally similar. Also, the precision of our results would be affected.

[9] See Diewert (1981).

concentrate on the arithmetic versions of both indices and note that the UK RPI is a plutocratically weighted index.

The difference between the plutocratic and the democratic mean is often referred to as plutocratic bias. The size of this bias depends on the relationship between households' inflation rates and their total expenditure. The plutocratic mean gives more weight to the inflation rate experienced by richer households, so if richer households experience a higher inflation rate than poorer households, the plutocratic mean will be larger than the democratic mean. Similarly, the plutocratic mean will be smaller than the democratic mean whenever richer households tend to experience a lower inflation rate than poorer households. If there is very little relationship between households' inflation rates and their total expenditures (i.e. the covariance between them is small), the magnitude of the bias will be small because when richer households do not systematically experience different inflation rates from poorer households, the extra weight that they receive in the aggregation will not matter.

3. Data

The data we use in this study come from the UK Family Expenditure Survey from 1975 to 1999. The FES is an annual random cross-sectional survey of around 7,000 households a year. As well as collecting information about the characteristics of each household, the FES collects detailed expenditure information for each household. This is recorded mainly by members of each household keeping a diary record of what they spend over two weeks, but for some durable goods the information is recorded via retrospective recall and the time frame is longer.[10] In the FES, the information is aggregated to the household level and averaged across the two-week period to give weekly expenditure figures for over 300 different goods and services.

The FES has much to recommend it as a data source on household spending; in particular, the coverage of goods is comprehensive, and it excludes expenditures by businesses. Indeed, it is heavily used by government statisticians and academics. However, it does have a number of drawbacks. For example, it does not measure spending by all households: it does not cover the institutional population of people living in retirement homes, military barracks or student halls of residence or the residents of hostels and temporary homes. Also, up until 1995, the FES ignored spending by household members under the age of 16. There may also be a problem of non-response, as nearly one-third of households that are initially approached do not respond to the survey, and these non-respondents may be different in a systematic way from households that take part. In particular, non-response is highest amongst richer households, amongst very young households and among the very old.[11]

These problems may not be terribly serious, but there are other potential problems in the FES that might be more substantive. In particular, there may be problems of under- or over-reporting of expenditures either through genuine forgetfulness (e.g. food consumed outside the home) or through active concealment (e.g. receipts from massage parlours) or through a combination of forgetfulness and guilt (e.g. alcohol). Problems of under-reporting in relation to alcohol and tobacco are thought, by the Office for National Statistics, to be so severe that the FES data are supplemented with data from other sources (clearances from bonded warehouses, for example) for use in national accounting. Tanner (1998) shows that under-reporting of alcohol spending compared with the National Accounts is such that the FES captures about 60 per cent of the National Accounts total, and that this under-reporting has been relatively stable over time (1978–92). Tobacco under-reporting has increased, with the FES capturing around two-thirds of National Accounts spending in 1992 compared with three-quarters in 1978. Another problem is the extent to which the two-week diary period in the FES means that large infrequent purchases (of durables, for example) may be underestimated. Data on durables from the FES are bolstered by data from other sources in the computation of the RPI.

[10] Note that this means that our index is not *exactly* a Laspeyres index. However, the coverage of retrospective recall questions is very minor in terms of the share of total household spending, so we expect our measure to correspond *very closely* to a Laspeyres.

[11] Tanner, 1998.

The price data used here are the published section indices of the RPI. Because these data are collected from national sources, there is no regional variation and as a result this Commentary ignores regional issues and also issues to do with whether or not the prices actually paid by rich and poor households for ostensibly the same goods may have changed differentially over the period. Differences in cost-of-living indices between population groups are thus generated entirely by differences in their spending patterns, scaled by relative price movements. These data are the most disaggregate and the most comprehensive price data available in the UK. In order to calculate the corresponding demands from the household data, we have to aggregate up from our 300 commodities to 69 categories, and we must therefore ignore demand heterogeneity below this level of aggregation.

One important feature of our data to note is our treatment of housing. We are interested in measuring inflation. A change in the price of a good or service (e.g. housing) can only affect our measure of inflation if the household pays for that good or service. If, for example, a household does not consume alcohol, increases in the price of alcoholic drinks will not affect its rate of inflation. Similarly, if a household receives a good without making a payment for it, any increase in the price of that good will not affect the household's rate of inflation if the household is compensated for it. This is important when we think about housing payments because of the way in which housing benefit is paid to households differently depending on whether the household is a private or public renter. For public renters, housing benefit (either full or partial) is used to pay rents directly without the household ever actually receiving a payment. But usually in the case of private renters, households receive a cash payment and this cash payment may not cover the rent in full. When private rental prices increase, housing benefit does not necessarily increase fully in line with the increase, leaving private renters worse off, whereas it is assumed that any increases in public rents are met by housing benefit increases, leaving public renters no worse off. For this reason, we include rent as an item of spending for private tenants even if they receive housing benefit, but we do not include rent as expenditure for public tenants. For similar reasons, if a household owns its house outright, a zero payment on housing is recorded since a change in the price of housing will not affect the household's rate of inflation. On the other hand, mortgagors who pay mortgage interest at a variable rate will be affected by changes in the price of this payment, or, in other words, they will be affected by changes in interest rates. One thing to note is that households with an interest rate that is fixed for a long period of time will not be affected by changes in interest rates, but since we cannot identify the type of mortgage that people have, we use the RPI mortgage interest payment section index to calculate changes in the price of housing for all mortgagors.[12] The definition of income that we use is the one used to calculate official low-income statistics[13] and we use a 'before-housing-costs' measure, which means that housing costs are not deducted from income.

[12] See Baxter (1998) for details.

[13] Department of Social Security, 1994.

4. Results

4.1 The distribution of inflation rates

In order to get some overview of the degree of variation in inflation rates across households at a point in time, consider Figure 4.1. This shows estimates of the densities of the distribution of inflation rates by year from 1976 to 2000 (the mountain range) along with the unweighted mean for each year (shown by the bold line at 'sea level'). What emerges is a rather mixed picture. Sometimes, the distribution is quite 'peaky', indicating relatively small differences across households; in other years, the densities of the distribution are very flat, showing a very wide range in the experience of general price inflation across households; and sometimes the densities even extend into negative inflation rates. There are also some examples of 'twin peaks'.

Figure 4.1. Densities of the distribution of inflation rates, 1976–2000

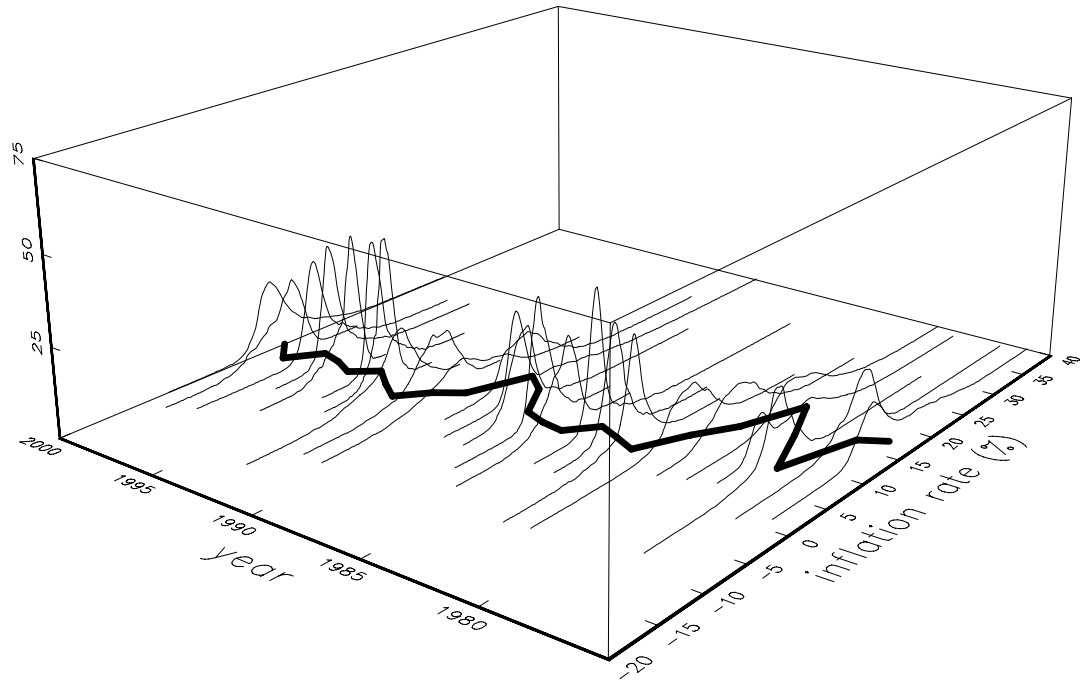

Figure 4.1 indicates that there is a fair degree of variation around the average rate of inflation. The annual average rate of inflation over the period is about 7 per cent (standard error 0.033 percentage points), with an annual standard deviation that varies between 1.21 per cent (standard error 0.01 percentage points) and 4.622 per cent (standard error 0.039 percentage points) and an overall average standard deviation of 2.8 per cent (standard error 0.021 percentage points). The semi-interquartile range varies between 0.675 percentage points (standard error 0.021 percentage points) and 3.66 percentage points (standard error 0.077 percentage points), with the average semi-interquartile range being about 1.62 percentage points (standard error 0.05 percentage points). This means that the difference between the first quartile of the inflation

distribution and the third quartile is, on average, over 3 percentage points, can be as much as 7.3 percentage points and is never less than about 1.4 percentage points. There is no clear pattern or trend in this measure of dispersion over time; however, it seems to be the case that inflation rates are more widely dispersed in the population when the average level of inflation is high. Figure 4.2 shows the relationship between a measure of location (the median) and a measure of dispersion (the semi-interquartile range) using the data from each of our 25 years of cross-sections. There is a (statistically significant[14]) positive relationship between the two (illustrated by the solid line) which says that, on average, a 1 percentage point increase in the median rate of inflation is associated with a 0.1 percentage point widening of the semi-interquartile range.

Figure 4.2. Relationship between the location of inflation and its dispersion

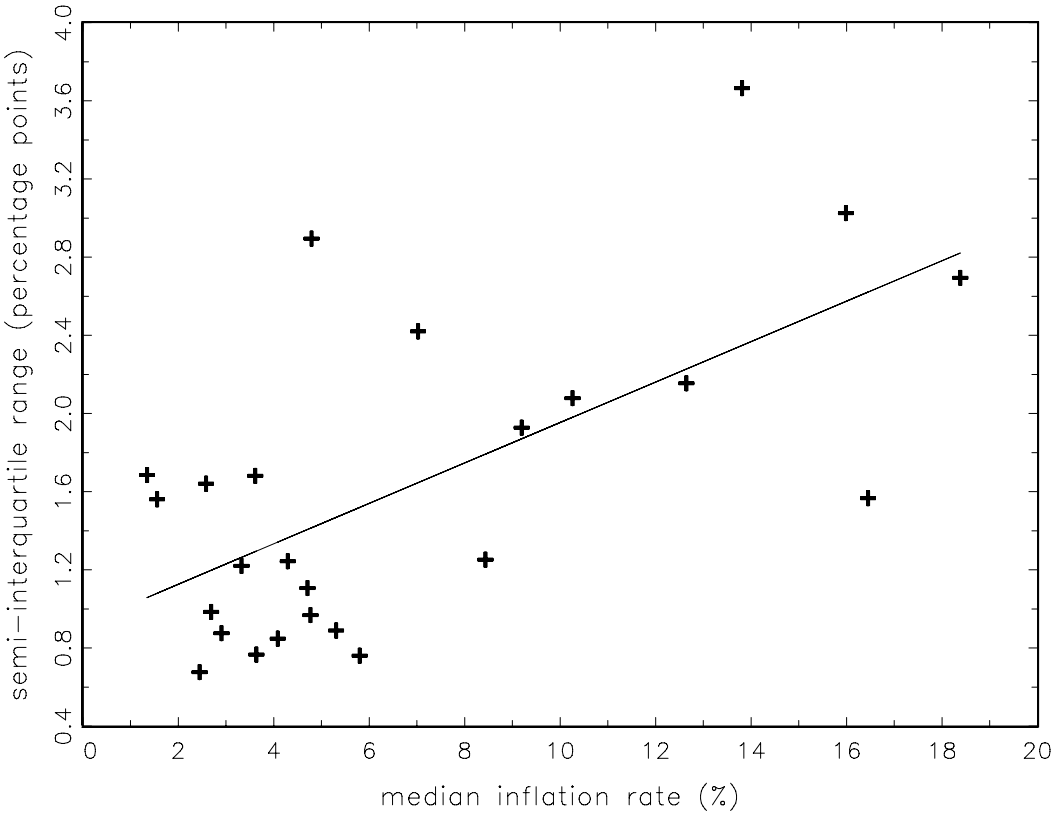

So far, we have seen that there is sometimes wide variation in the dispersion of inflation rates and that this dispersion is typically greatest when inflation is high. We now turn to the question of the extent to which the average rate is typical for the population of households. We take as our measure the proportion of households whose inflation rates are within 1 percentage point of the mean; that is, if the mean inflation rate is 10 per cent, we ask how many households' inflation rates are between 9 per cent and 11 per cent.

[14] Note that we have ignored the effects of sampling variation on the medians and semi-interquartile ranges themselves in this regression. The standard errors on these are very small due to the large sample sizes used to generate them and so this will not affect our conclusion.

The proportions are shown in Figure 4.3 by year. The small vertical lines at each data point are the 95 per cent confidence intervals, which are clearly rather small.[15] One feature that emerges is that the average level of the line is not all that high: on average over the entire period, only 35 per cent (standard error 0.5 percentage points) of households have inflation rates that are within 1 percentage point of the average. The proportion is also highly variable – as high as 65 per cent in 1994 but as low as 9 per cent in 1989. Year-to-year changes are also quite marked on occasion (the 1988 figure is 44 per cent, for example). As we might expect, given the evidence that dispersion is, in general, higher when the average level of inflation is high, Figure 4.4 shows that, on average, the percentage of households close to the mean declines as the average rate increases. The solid line in the figure shows the average relationship and indicates that a 1 percentage point increase in the mean rate of inflation is associated with a 1.7 percentage point (standard error 0.5 percentage points) fall in the percentage of households whose rates of inflation are close to the average.

Figure 4.3. Percentage of households within 1 percentage point of the mean rate of inflation

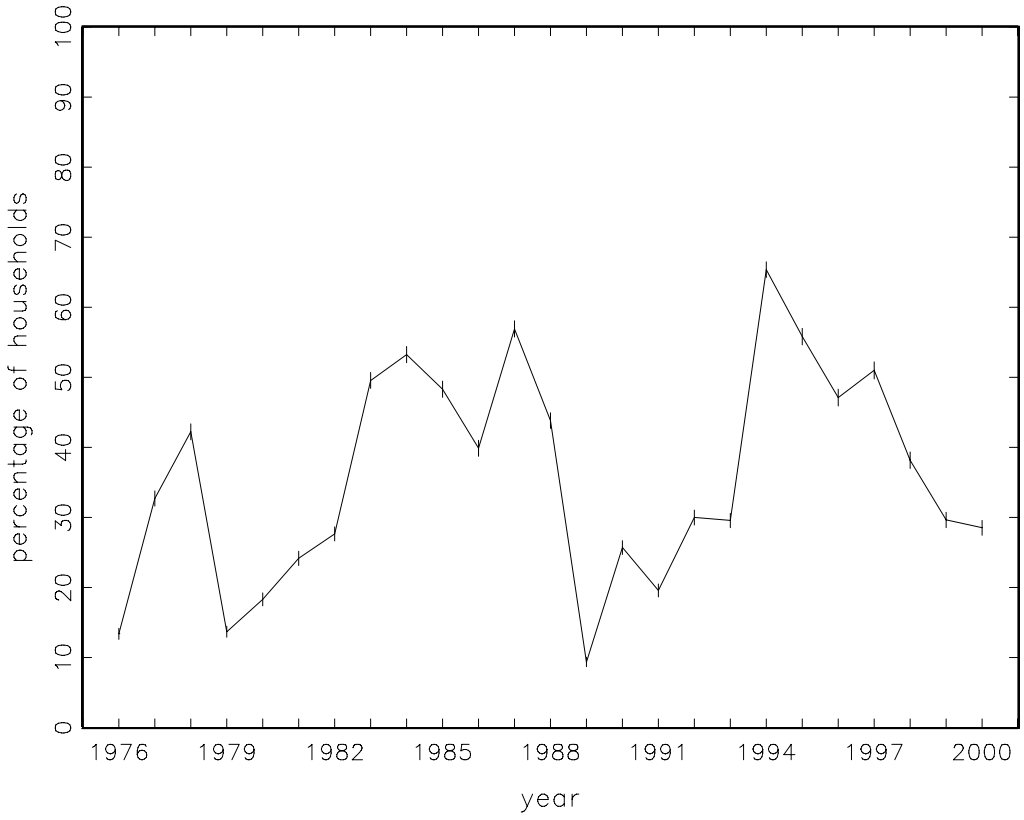

[15] Again, in this calculation, we have ignored the effect of sampling variation on the mean. This is very small and its inclusion would not affect our results. Bootstrap methods which account for this turn out to give confidence intervals that are even smaller than those shown here.

Figure 4.4. Relationship between the mean rate of inflation and the percentage of households within 1 percentage point of the mean

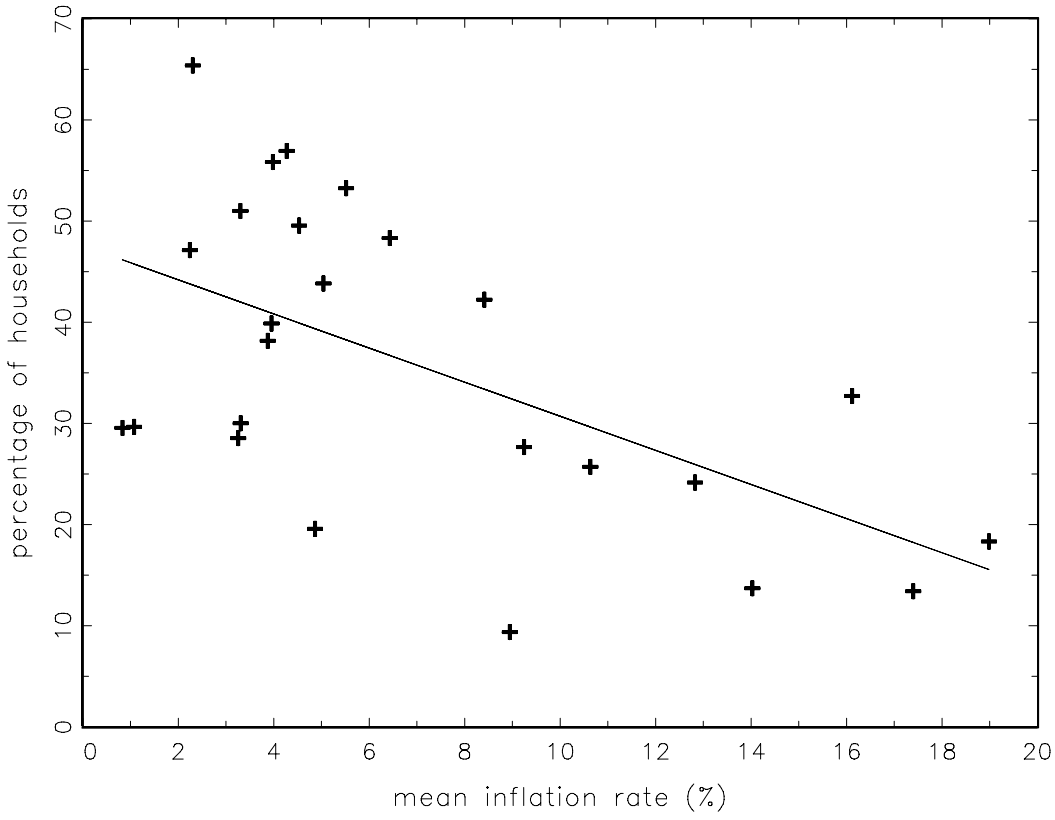

4.2 Group inflation rates

One of the benefits of calculating household-specific cost-of-living indices is that it enables us to group these across households according to various characteristics that we might be interested in. Given the household-specific cost-of-living index, we need some way to aggregate this information across the relevant groups. In Chapter 2, we noted that there are various schemes that we could use, two of which are the unweighted arithmetic mean (the democratic mean) and the weighted arithmetic mean with the weights being equal to the households' shares out of total expenditure (the plutocratic mean).

Figure 4.5 shows both the plutocratic mean and the democratic mean across all households. A 95 per cent confidence interval is plotted around the two mean values. The alternative weighting schemes result in statistically significantly different rates of inflation overall in 18 of the 25 years from 1976 to 2000. This can be seen from Figure 4.6, which shows the difference between the plutocratic mean and the democratic mean and the 95 per cent confidence interval (shown by the dashed lines) plotted around this difference. Whenever the confidence interval does not overlap zero, we can conclude that the difference is significantly different from zero. Even when the difference between the two measures is statistically significant, it is quantitatively small. The greatest difference between the plutocratic and democratic means occurs in 1981, when the democratic mean inflation rate is 12.8 per cent and the plutocratic mean is 11.9 per cent

Figure 4.5. Plutocratic and democratic inflation rates, 1976–2000

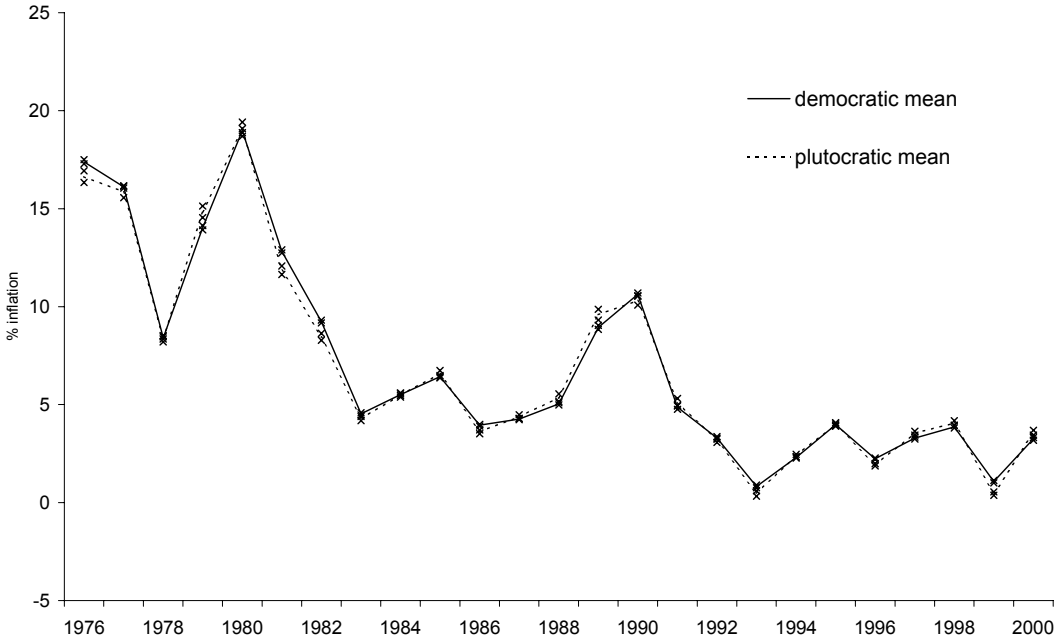

Figure 4.6. Plutocratic bias, 1976–2000

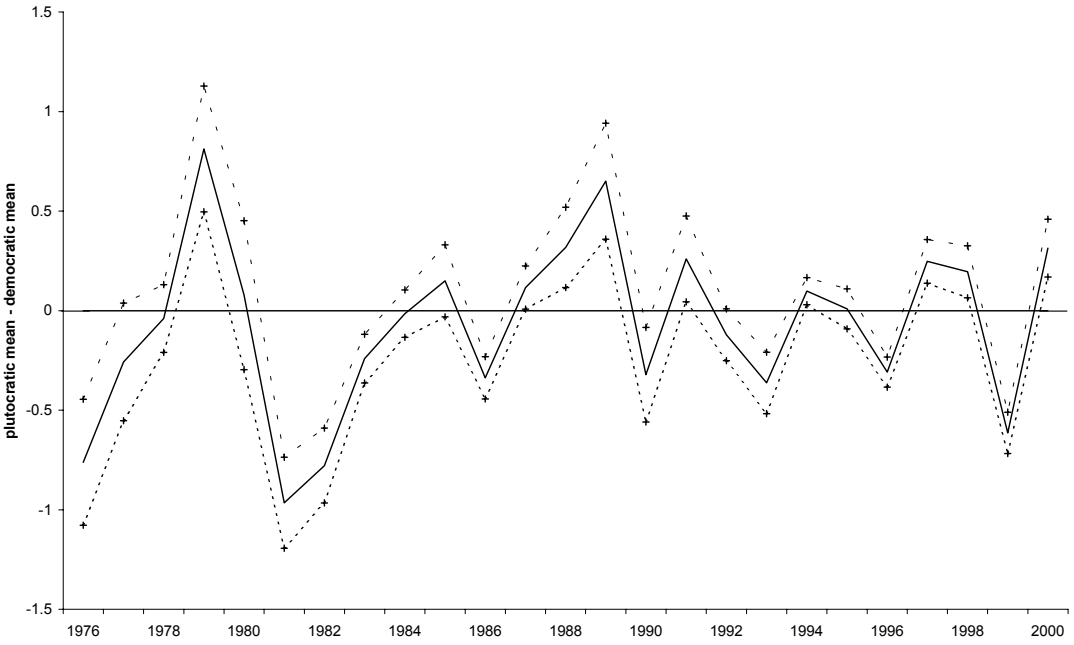

so that the difference between the two is just under 1 percentage point. As discussed in Chapter 2, the difference between the plutocratic and democratic means depends on the relationship between total household expenditure and household-specific inflation rates. On average, higher-spending households had above (unweighted) average inflation in, for example, 1989, and pulled the plutocratic mean (which gives them more weight) above the democratic mean. Because high-spending households tend to be those with high

equivalised incomes, we can relate the direction of the plutocratic bias to what we find in Figure 4.7, which shows the inflation rates for the richest and poorest 10 per cent of the population defined by before-housing-costs household income that has been adjusted for household size using the McClements equivalence scale.[16] The point estimate of the inflation rate for the richest 10 per cent of the population in each period is shown by the dashed line, while the point estimate of the rate for the poorest 10 per cent is shown by the solid line. In each year, a 95 per cent confidence interval is plotted around the mean. In 1989, when the plutocratic mean is higher than the democratic mean, we see that households in income decile 10 had a higher inflation rate on average than households in decile 1. Conversely, in 1981, when the plutocratic mean is lower than the democratic mean, we find that the mean inflation rate in the poorest income decile is higher than that in the richest income decile. In what follows, we will use the democratic mean.

Figure 4.7. Inflation rates by income, 1976–2000

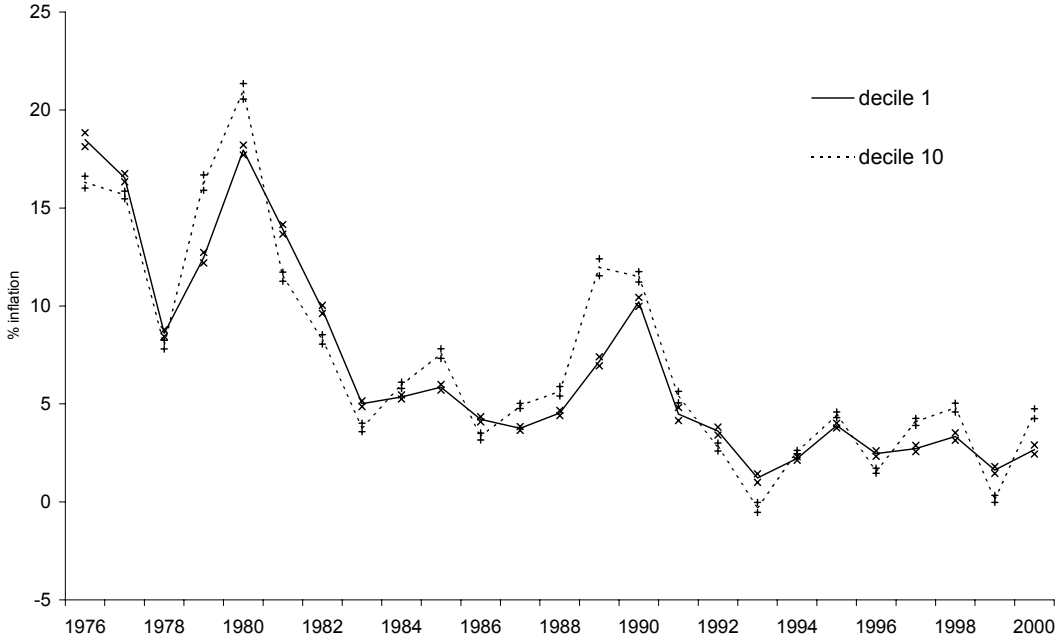

The average inflation rates for the richest and poorest income deciles are significantly different from each other in most years but the ranking of the two inflation rates experienced by the two groups changes frequently over the period. The rich experience the higher rate of inflation only slightly more often than the poor, but when this is the case, the difference tends to be larger than when the inflation rate is anti-poor. This shows up in the average inflation rates for each group, which are shown in Table 4.1 later. The average annual inflation rate for the poorest 10 per cent of households was 6.77 per cent, while the average rate for the richest 10 per cent of households was 7.14

[16] There is a slight complication when thinking about income deciles within a year. Our price data are monthly, so households that are sampled in a different month in the same calendar year will face different changes in relative prices. This is not a problem as long as households with the particular characteristic that we are interested in are evenly distributed throughout the year. This is true of all characteristics that we will group by in this section except income. The problem with income is that within a year, there will be some income growth, so households sampled at the end of a year will be more likely to be higher up the income distribution simply because of income growth. To overcome this problem, we create our income deciles within a quarter (we are unable to define them within a month as the FES is not random within a month).

per cent, and the difference between the two rates is statistically significant. Figure 4.8 shows the annual average (the geometric mean) inflation rate from 1976 to 2000 for all income deciles. The 95 per cent confidence intervals are also shown and the solid line shows the overall mean inflation rate (which was about 6.9 per cent). The poorest four income deciles and the sixth income decile have experienced inflation that, on average, has been less than the inflation rate for everyone, whereas the richest three income deciles have experienced inflation that is higher than the mean rate. These annual average inflation rates show the average inflation rate that households occupying (say) the bottom income decile experienced over time from 1976 to 2000. This is not the same as asking what the inflation rate from 1976 to 2000 was for people who occupied the bottom income decile in 1975. The former inflation rate allows the people who occupy the bottom income decile to change over time and the types of goods that they buy to change over time.

Figure 4.8. Average inflation rates by equivalised income decile, 1976–2000

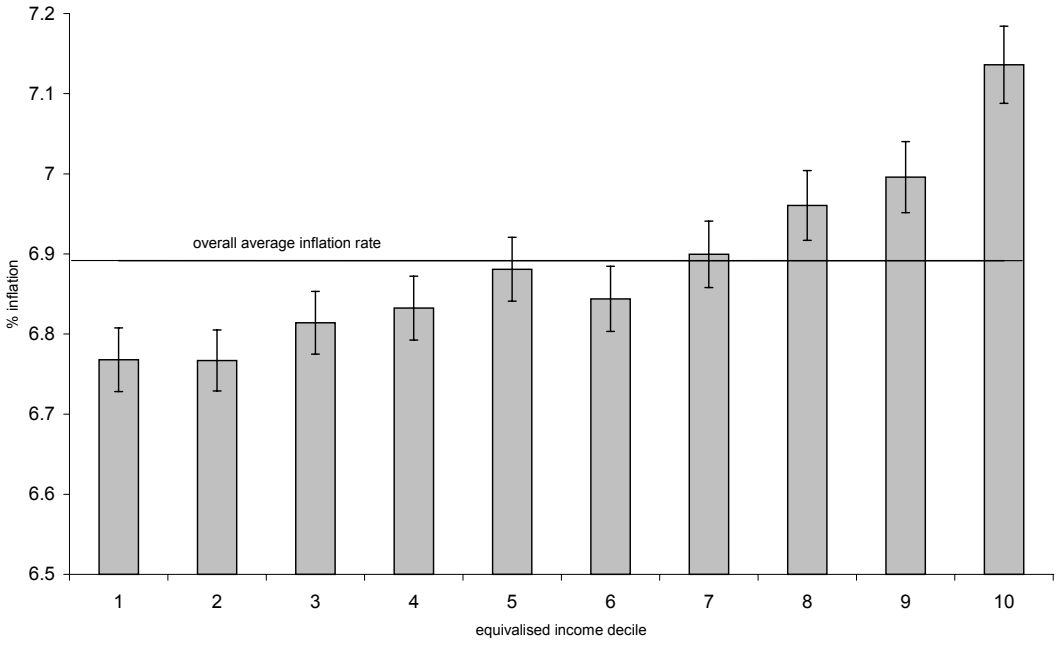

Figure 4.9 shows inflation rates for households with and without children. As for households with different incomes, the inflation rates experienced by households with and without children differ in a number of years but the ranking changes frequently. However, there is no significant difference between the average cumulative inflation rate experienced by households with and without children from 1976 to 2000.

One characteristic that we might expect to generate large differences in inflation rates is tenure. Since housing makes up a large proportion of total expenditure, any price changes will have a large effect on a household's inflation rate. For example, in 1999, for people who rented their property, rent payments made up 14 per cent of total household expenditure, and for mortgagors, mortgage payments made up about 16 per cent of total household expenditure. We group households according to whether they pay rent for

Figure 4.9. Inflation rates for households with and without children

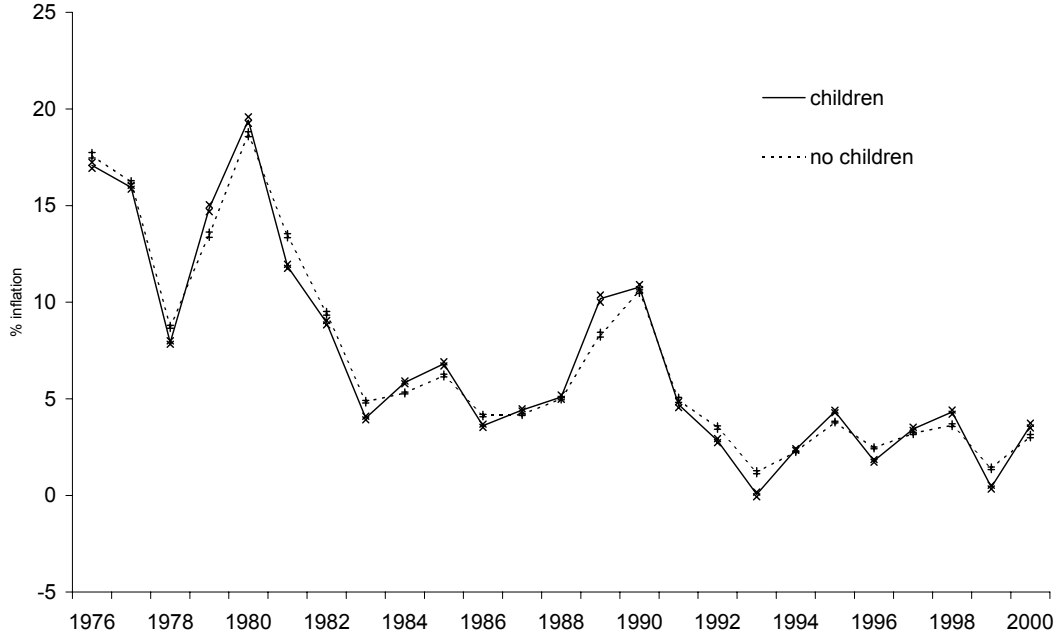

Figure 4.10. Inflation rates by tenure, 1976–2000

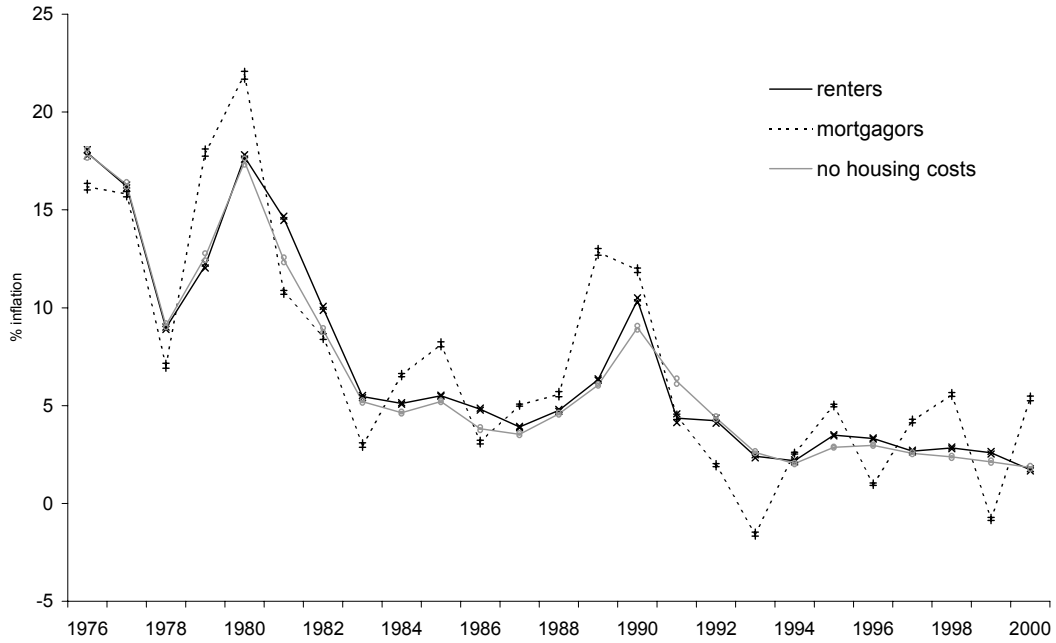

their accommodation, whether they own it with a mortgage or whether they do not pay any housing costs at all (as described in Chapter 3, this last group includes outright owners and also public renters who receive full housing benefit). Figure 4.10 shows inflation rates across these three groups of households and, as expected, there are large differences between them. The biggest difference occurs between mortgagors and the other two groups, which tend to follow each other more closely. In 1989, for example,

the inflation rate for both renters and people who pay no housing costs was 6.3 per cent and the inflation rate for mortgagors was more than double that, at 12.9 per cent. In the 1990s, the inflation rate for mortgagors was much more variable than the inflation rate for the other two groups: mortgagors experienced changes in the inflation rate of as much as 6 percentage points, whereas inflation was much more constant for renters, with changes in the rate of inflation never being more than 2 percentage points. Over the entire period, mortgagors experienced an average rate of inflation of 7.13 per cent, while renters experienced an average rate of 6.83 per cent and households that pay no housing costs experienced an average rate of 6.56 per cent.

Figure 4.11 shows the inflation rates for households grouped by whether the head is of pensionable age. Overall, pensioners have experienced slightly lower inflation than non-pensioners (an average of 6.78 per cent compared to 6.93 per cent for non-pensioners), although the ranking changes frequently, particularly in the 1990s.

Figure 4.11. Inflation rates for pensioner and non-pensioner households, 1976–2000

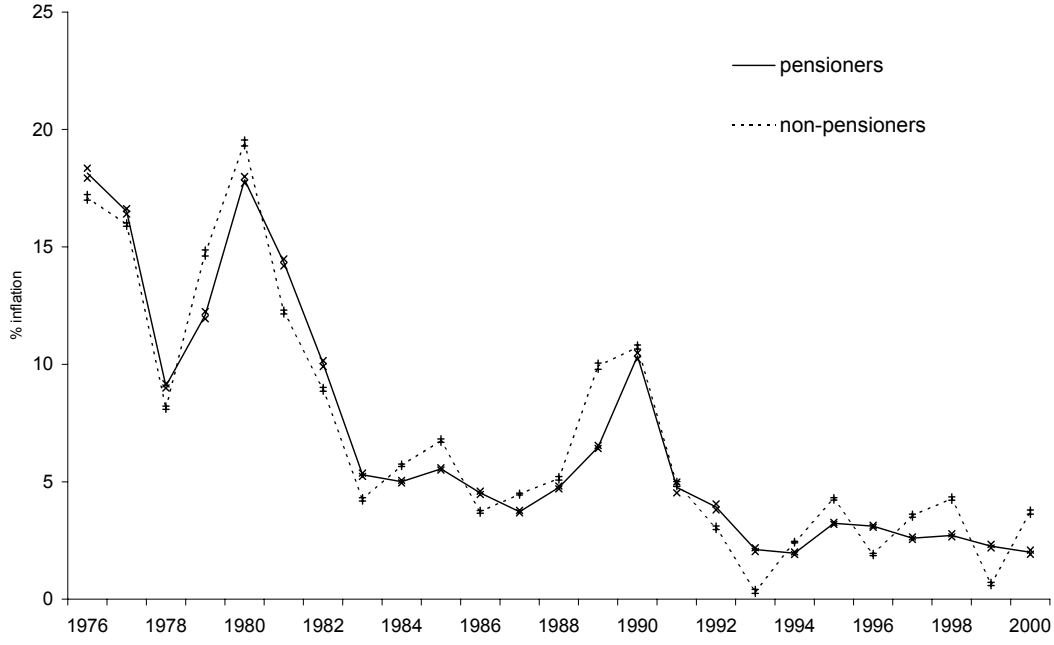

Since 1981, increases in the basic state pension have been formally linked to the average level of price inflation. Perfect price indexation would mean that the real value of the pension to its recipients should have remained flat since this time. But there are two reasons why price indexation is not perfect. First, the adjustment is done using a retrospective inflation figure (for example, in April 1999, the pension was adjusted using a figure for inflation from September 1997 to September 1998). Second, the different inflation experiences of different households mean that using an average rate, such as the RPI, to uprate the pension may in fact have left pensioners worse off if they experienced a higher level of inflation than the rest of the population. In Figure 4.12, we analyse the extent to which the increases in the basic state pension have compensated pensioner

Figure 4.12. Real growth in the basic state pension, 1988–99

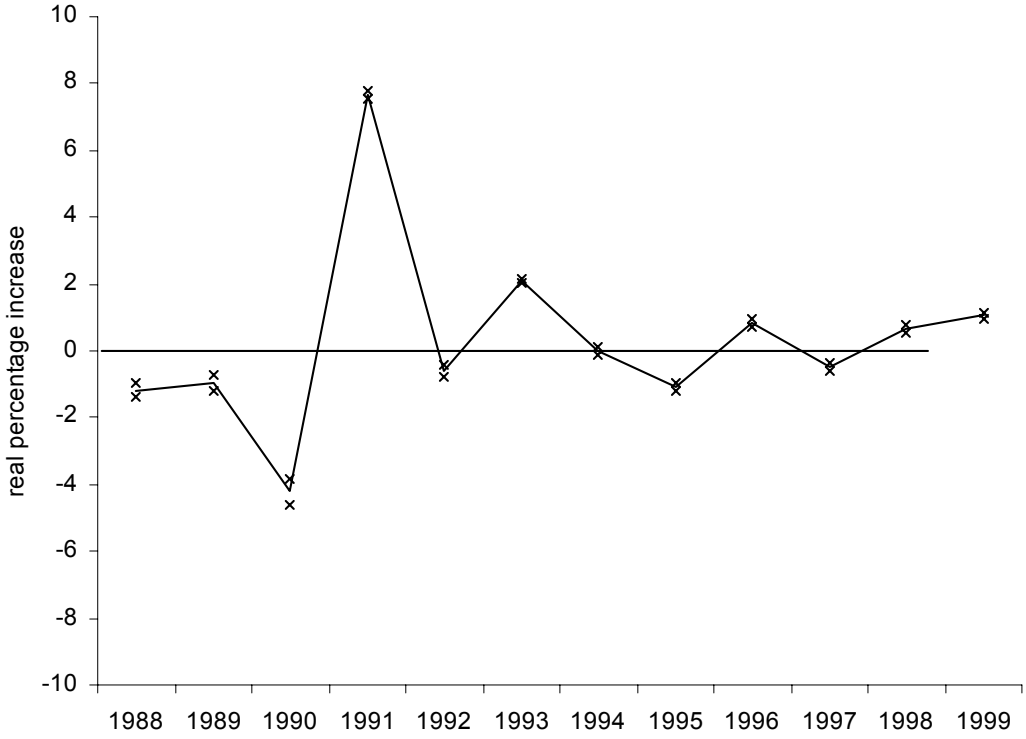

Figure 4.13. Uprating the basic state pension, 1987–99

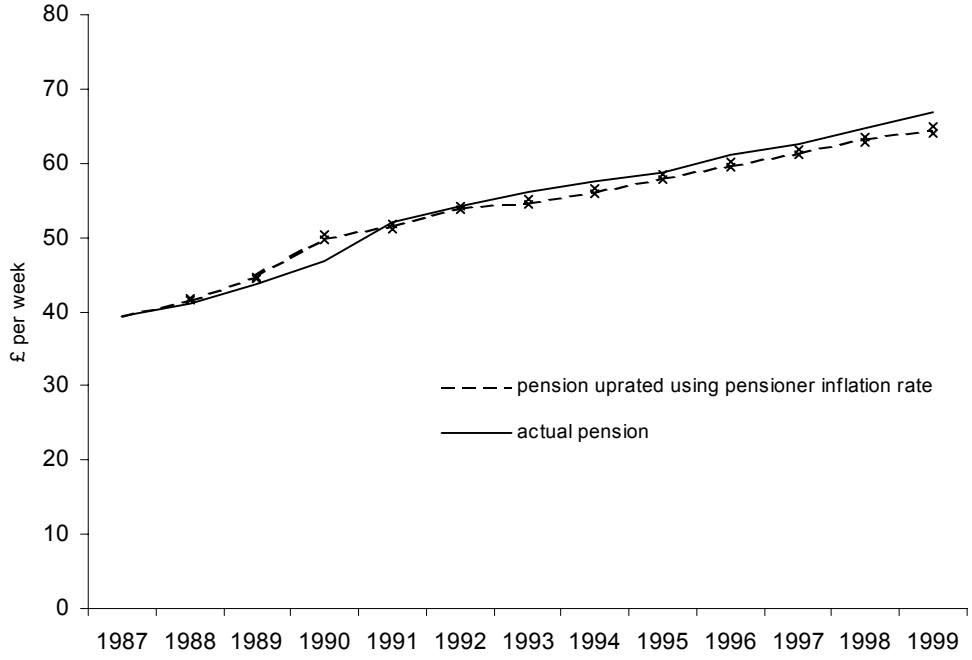

Table 4.1. Average annual inflation rates, 1976–2000, by group

Group	Average inflation	Standard error
All	6.89%	0.0069
Income decile 1	6.77%	0.0203
Income decile 10	**7.14%**	0.0246
Pensioners	6.78%	0.0108
Non-pensioners	**6.93**%	0.0083
Renters	6.83%	0.0900
Mortgagors	**7.13**%	0.0118
No housing costs	6.56%	0.0100
Employed	**6.97**%	0.0091
Unoccupied	6.79%	0.0103
Children	6.86%	0.0115
No children	**6.91**%	0.0085
Age ≤25	**7.00**%	0.0300
Age 26–35	**7.04**%	0.0174
Age 36–45	**6.96**%	0.0162
Age 46–55	6.85%	0.0159
Age 56–65	6.77%	0.0147
Born 1931–35	6.78%	0.0223
Born 1936–40	6.82%	0.0227
Born 1941–45	**6.95**%	0.0232
Born 1946–50	**7.05**%	0.0225
Born 1951–55	**7.03**%	0.0263
Lone parents	6.68%	0.0317
Couples with children	**6.92**%	0.0128
Couples without children	6.82%	0.0113
Single adults (below pension age)	**7.26%**	0.0274

Note: Bold type indicates groups with means that are above the all-household average.

households for the changes in prices that they faced.[17] Figure 4.12 shows the percentage real increase in the basic state pension in all years from 1988. For example, the basic state pension fell in real terms by around 4 per cent in 1990, but in the next year it rose in real terms by over 7.5 per cent. The overall effect over the period from 1987 to 1999 of using an average inflation rate on the real value of the pension can be seen from Figure 4.13. The solid line shows the actual (nominal) value of the pension since 1987 and the dashed line shows what the value of the pension would have been if it had been uprated using

[17] The inflation rates used in Figure 4.12 are different from those used in Figure 4.11. We want to know whether the increase in pension that each household faces accurately reflects its inflation rate over the past year. This means that we need a backwards-looking inflation rate (for example, the inflation rate from t–1 to t for households observed in period t), whereas in Figure 4.11 we used a forwards-looking inflation rate (for example, the inflation rate from t to t+1 for households observed in period t). The first year of data that we use is 1987 because this was the year when increases in the pension began to take effect from April (whereas they used to take effect from the November of each year). This gives us a 1988 growth rate.

the inflation rate that we have calculated for pensioners (with a 95 per cent confidence interval shown by the crosses). The late 1980s and early 1990s saw the real value of the pension fall, but from 1991 onwards the value of the pension has been higher than the amount that would have resulted from using our pensioner inflation rate. By the end of our period of data, the difference between the two values is £2.27. This means that over the whole period, the basic state pension has increased by around 3.5 per cent in real terms.

There are many more characteristics that we could group households by, and the average inflation rates over the 25-year period are shown for some additional characteristics in Table 4.1.

In the next section, we turn to the implications of our results.

4.3 The distributional effects of inflation

The measurement of changes in real living standards requires that nominal values[18] are converted into real values. This allows the living standards of households facing different prices to be compared. The idea is to get a reliable measure of households' command of real resources per head. Typical approaches use common price indices (and for that matter equivalence scales) to adjust incomes for changes in purchasing power. Often when official poverty indicators are published, the data are converted into constant prices[19] so that the growth in real equivalised incomes can be measured. Of course, if the correct adjustment for prices varies across households, then it should be important to allow for this in the calculation.

In general, inequality measures are invariant to the typical treatment of incomes, which is to deflate all households by a common average price index. The Gini coefficient and the family of generalised entropy measures (which include the Theil mean log deviation, the Theil index and the (half the squared) coefficient of variation as special cases) are all independent of such common rescaling. Indeed, this is almost an axiomatic requirement of inequality measures[20] which tends to have the effect of focusing attention on relative rather than absolute inequality and poverty. All inequality measures, however, are affected by the use of observation-specific deflators. In this section, we use our household-specific inflation rates to adjust household income so that we can make comparisons between real incomes in different time periods. We contrast the results from this approach with the results that are derived by using a common inflation adjustment for all households observed within the same period.

The effect of using household-specific inflation rates on the distribution of real incomes depends on the covariance between income and the price index.[21] A positive covariance between the price index and incomes in year t means that higher incomes are associated

[18] Whatever the nominal measure we are using – for example, equivalised nominal household income, or total expenditure.

[19] Department of Social Security (1994), for example.

[20] See Maasoumi (1997), for example.

[21] Ignoring, for the moment, the effect of the variance of the price index.

with higher inflation rates. This tends to disperse the income distribution in year *t*. If the price index deflates incomes to an earlier year, say *t*–1, a positive covariance between the price index and incomes means that *poorer* households have had *higher* inflation than richer households between *t*–1 and *t*. If we compare inequality in the real income distribution in years *t*–1 and *t* (in *t*–1 prices), we will see that growth in inequality between these years is *higher* than is apparent when the effects of differential inflation are ignored. If, however, the price index inflates incomes to a later year, say *t*+1, a positive covariance between the price index and inflation means that *richer* households have had *higher* inflation than poorer households between *t* and *t*+1. If we compare inequality in the real income distribution between *t* and *t*+1 (in *t*+1 prices), we will see that inequality between these two years has grown by *less* than is apparent under the common price index assumption.

We now turn to our results. The inequality measure that we use is the variance of the log.[22] Figures 4.14 and 4.15 show the evolution of inequality of real equivalised household income in constant prices. Figure 4.14 bases prices in 1975 and so deflates all households' incomes to 1975 prices. Figure 4.15 bases prices in 1999 and so inflates all households' incomes to 1999 prices. Each figure shows two measures of income inequality. The dashed line in each figure shows inequality in real equivalised household income calculated using a household-specific price index P_s^h:

$$\operatorname{var}\ln\left(y_t^h P_s^h\right),$$

where y_t^h is the nominal equivalised household income of household *h* observed in period *t* and

$$P_s^h = \frac{\mathbf{p}_s' \mathbf{q}_t^h}{\mathbf{p}_t' \mathbf{q}_t^h}$$

is the household-specific price index that revalues the household's period-*t* income into period-*s* prices. The solid line in each figure is inequality in real equivalised household income calculated by applying a common price index to each household; that is,

$$\operatorname{var}\ln\left(y_t^h \overline{P}_s\right),$$

where

$$\overline{P}_s = \frac{\sum_h P_s^h}{H}$$

is the average of all of the household-specific price indices. Since the price index does not vary cross-sectionally,

[22] We have also calculated all of the results in this section using alternative measures such as the Gini and members of the generalised entropy family. None of the results is affected by the choice of inequality measure, so we use the variance of the log as our object of interest as this allows for a fairly simple treatment of sample variation and it also decomposes readily.

$$\text{var}\ln\left(y_t^h \overline{P}_s\right) = \text{var}\ln y_t^h$$

and the variance of the log of equivalised real income is equal to the variance of the log of equivalised nominal income. This means that inequality as measured by nominal

Figure 4.14. Evolution of inequality, 1975 prices

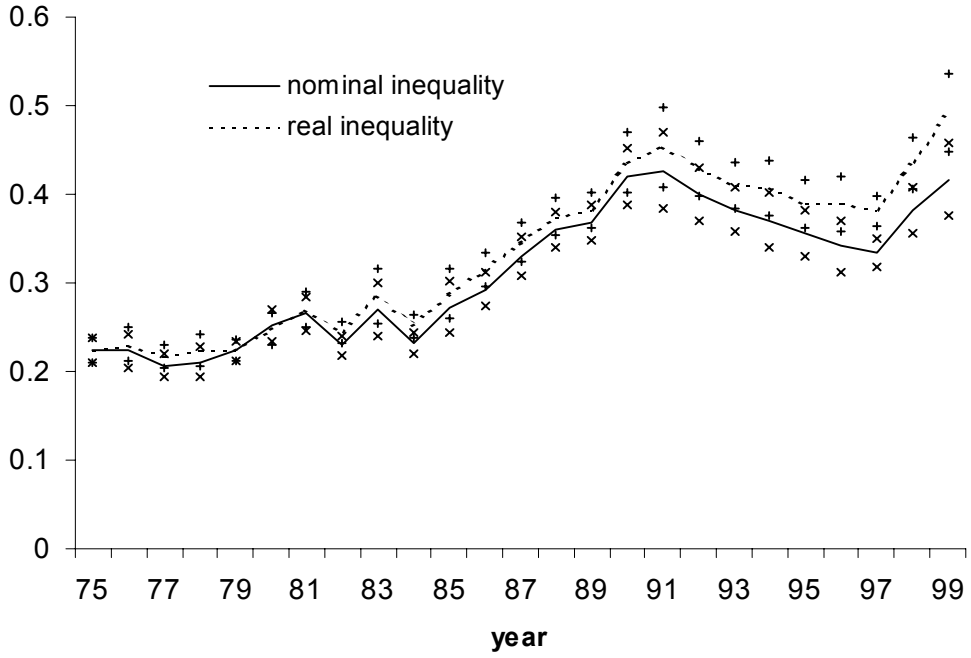

Figure 4.15. Evolution of inequality, 1999 prices

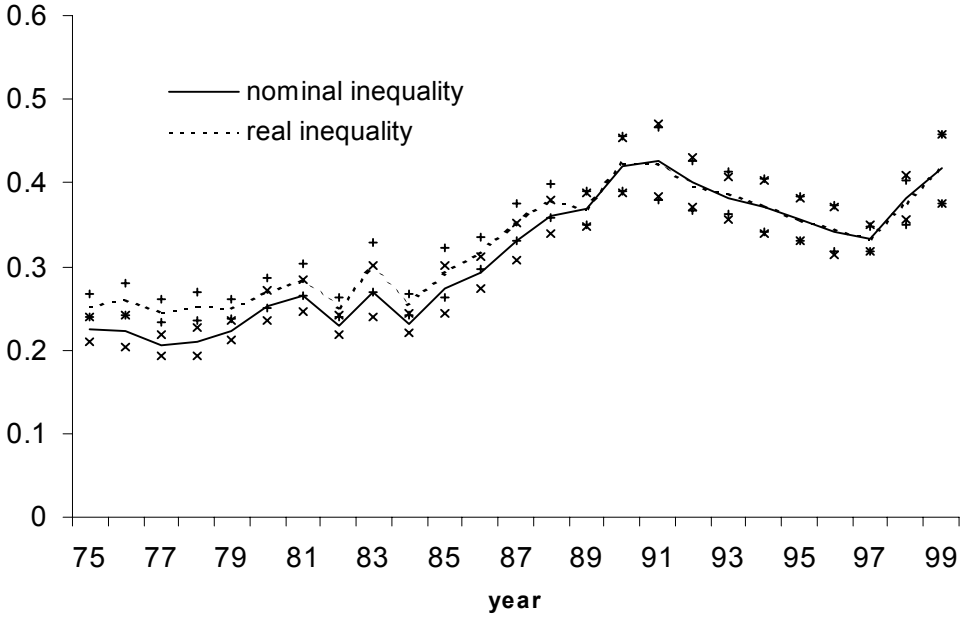

income is exactly the same as inequality as measured by real income when we use a common price index. For this reason, we will often refer to real income inequality calculated using a common price index as nominal income inequality and we will refer to real income inequality calculated using a household-specific price index as real income inequality. The pluses and crosses on Figures 4.14 and 4.15 denote the 95 per cent confidence interval on each inequality measure. It is not obvious from *these* confidence intervals whether the two series are really different from each other statistically (because to establish this it is necessary to allow for the covariance between them). Figures A.1 and A.2 in the appendix show the difference between the series in Figures 4.14 and 4.15 respectively and the 95 per cent confidence interval on this difference. From these, it can be seen that there is a statistically significant difference between the two series in nearly all years for both figures. What Figures 4.14 and 4.15 show is quite striking.

Firstly, using household-specific price indices gives a different picture of the evolution of inequality in living standards over the period from the one that emerges if these inflation effects are ignored. For instance, Figure 4.14 shows that inequality increased between 1978 and 1979 on the measure that assumes common inflation rates, whilst the measure that allows for differential inflation effects shows a decline in inequality over the same period. Figure 4.15 shows the same divergence. Another example occurs in the change from 1988 to 1989 in Figure 4.15. There are several examples in the two figures of the two series going in different directions.

Secondly, the picture that emerges depends upon the period in which prices are based. When incomes are expressed in 1975 prices, Figure 4.14 shows that the growth in inequality over the whole period was greater once the differential effects of inflation on real incomes are accounted for than when these effects are ignored. However, when incomes are expressed in 1999 prices, Figure 4.15 shows that the growth in inequality was lower once the differential effects of inflation on real incomes are accounted for. At first sight, this seems odd. That the two pictures are different is not surprising and follows from the fact that the two measures of real income that we have for each household in each year are different. Consider the real income of household h observed in year t. Our two measures of its real income are

$$y_t^h P_{99}^h = y_t^h \left(\frac{\mathbf{p}_{99}' \mathbf{q}_t^h}{\mathbf{p}_t' \mathbf{q}_t^h} \right) \qquad \text{and} \qquad y_t^h P_{75}^h = y_t^h \left(\frac{\mathbf{p}_{75}' \mathbf{q}_t^h}{\mathbf{p}_t' \mathbf{q}_t^h} \right).$$

The first measure is nominal income for household h in year t multiplied by (1 plus) the inflation rate from year t to 1999. The second measure is nominal income for household h in year t multiplied by (1 plus) the deflation rate from year t to 1975. These are clearly different and so it is not surprising that Figures 4.14 and 4.15 are not exactly the same. However, it is surprising that the two pictures tell completely different stories. If inflation has affected rich households more than poor, we would expect this to have a beneficial effect on the real incomes of poorer households and so lead to lower growth in inequality than when we ignore differential inflation. This is what we find in Figure 4.15. But if inflation has had the opposite impact and affected poor households more than rich, we would expect this to have a beneficial effect on the real incomes of richer households and so lead to higher growth in inequality than when we ignore differential

inflation. This is what we find in Figure 4.14. So, from Figure 4.14 it would appear that inflation has been anti-poor, but from Figure 4.15 it would appear that it has been anti-rich.

To see what is going on, consider the decomposition of our inequality measure,

$$\text{var}\ln\left(y_t^b P_s^b\right) = \text{var}\ln\, y_t^b + \text{var}\ln P_s^b + 2\,\text{cov}\left(\ln\, y_t^b, \ln P_s^b\right).$$

The variance of the log of real income decomposes into three elements: the variance of log equivalised nominal income ($\text{var}\ln\, y_t^b$), the variance of the log price index ($\text{var}\ln P_s^b$) and twice the covariance between the two ($2\,\text{cov}(\ln\, y_t^b, \ln P_s^b)$). The solid lines in Figures 4.14 and 4.15 are identical and show the variance of log equivalised nominal income. Figures 4.16 and 4.17 plot the remaining two elements: the variance of the log price index is the solid line in each, and the covariance term is the dashed line (the points on either side of each are 95 per cent confidence intervals).

Consider first the variance of the log price index (the solid line in each picture). This variance grows as the comparison period gets further away from the base period. This is not surprising since the variance of the price index (and the variance of the log price index) tends to zero as the change in relative prices of all the commodities becomes smaller. This is simply because budget shares sum to one. As the comparison period gets closer to the base period, the vector of relative prices tends towards a vector of ones and so the variance gets small and is actually zero in the base period.

The second thing to note about these figures is that the covariance term is not trending in a similar way to the variance term. The covariance is a simple measure of linear association between incomes and inflation. If richer households have higher inflation rates than poorer households, then this covariance term is positive. This covariance changes in sign over the period, a feature that is to be expected given Figure 4.7 above, which showed how the average rate of inflation in the first and tenth income deciles varied over the period. The net effect of all of these components on the overall inequality of real incomes will depend upon the relative sizes of each and the sign of the covariance. For example, the covariance term in Figure 4.16 in 1980 is negative (rich households in 1980 have a lower deflation rate (i.e. a higher inflation rate) and so have to be deflated by more than poorer households to get them to 1975 prices) and is absolutely larger than the variance of the log price index term. As a result, the dashed line in Figure 4.14 is pulled down just below the solid line and hence the differential effects of inflation act to reduce the inequality in living standards in 1980 relative to 1975. However, in many years, we tend to get counter-intuitive results. Take, for example, 1989 in Figure 4.16. The covariance between income and inflation in this year is negative and absolutely larger than in 1980. This means that we would expect real income inequality in 1989 to be lower than nominal income inequality, as we found in 1980 (leading us to conclude that inequality growth between 1975 and 1989 was less steep than was apparent when a common price index was used). But, instead, as we see in Figure 4.14, real income inequality in 1989 was higher than nominal income inequality in this year. The reason for this is that because we are further away from the base period in 1989 than in 1980, the

variance of log prices is much larger and so outweighs the effect of the negative covariance.

Figure 4.16. Variance of the log price index and twice the covariance between log equivalised income and the log price index, 1975 prices

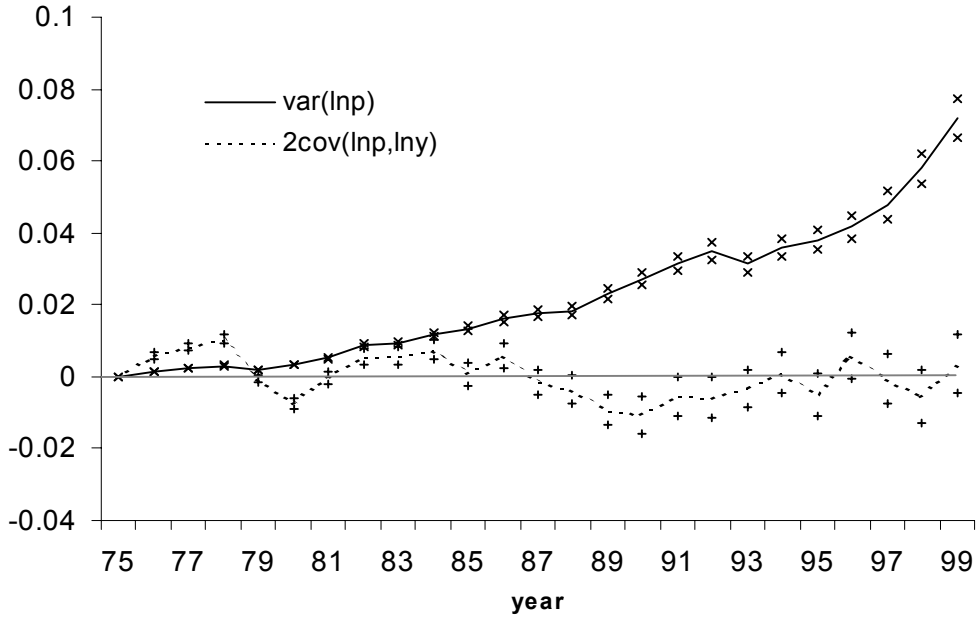

Figure 4.17. Variance of the log price index and twice the covariance between log equivalised income and the log price index, 1999 prices

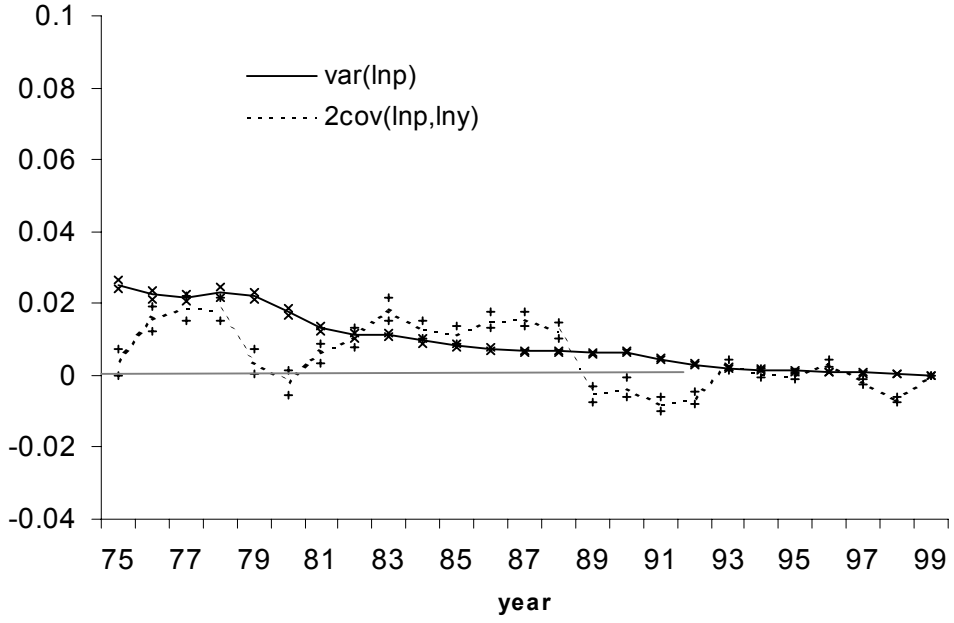

Similarly, in Figure 4.17, we see that in 1991 there is a negative covariance between income and the inflation rate (which means that in 1991 richer households had a lower inflation rate than poorer households) and this covariance just outweighs the variance of the log price index, giving an overall effect that reduced the inequality in living standards in this period relative to 1999, as shown in Figure 4.15. Again we also get some counter-intuitive results. In 1980 in Figure 4.17, for example, we get a very small covariance between income and inflation.[23] This means that there was very little relationship between income and inflation in this year: even though some households experienced higher inflation rates than others, on average rich and poor households were affected fairly equally. Hence we would expect real income inequality not to differ too much from nominal income inequality. However, this is not what we find in Figure 4.15, where there is a statistically significant difference between the two measures. Again this is because of the variance of log prices, which works to increase real income inequality.

Figures 4.14 and 4.15 are both therefore understandable and reasonable. But this does not make interpreting them particularly easy, and an answer to the question about what has happened to inequality in real living standards over the period does not seem nearer. What can we do about this dependence on the base in the results?

We have seen that the base dependence arises largely as a result of the variance of the log price index term. The contribution to the overall variability of real incomes that comes from the variance of the log price index may not be, in itself, informative; it must be positive and so could be thought of as simply the increase in variance we would expect if we multiplied income by *any* uncorrelated random variable. The covariance term, however, reflects the correlation between income and inflation (which can go in either direction) and is economically interesting because it shows how the living standards of rich and poor households are differentially affected by inflation. One approach to this measurement problem could be simply to ignore the contribution of the variance of the log price index, arguing that its evolution is mechanical, purely statistical and not particularly meaningful.

This is a little like the axiomatic invariance-to-scale-changes property of most inequality indices which renders them all invariant to the rescaling that common price indexation gives. In this case, we want our inequality measure to be invariant to multiplication of each observation by an uncorrelated random variable. If we do ignore the contribution of $\operatorname{var}\ln P_s^b$, then we get Figures 4.18 and 4.19. The results are still base dependent[24] but there is much closer agreement between the two sets. This is because the difference between the common price adjustment measure (solid line) and the measure that uses the household-specific price index (dashed line) is now due to just the covariance term and, as could be seen in Figures 4.16 and 4.17, these terms are quite small and there is (rather)

[23] It is not statistically significantly different from zero.

[24] As explained above, this is because the measures of real income for each household are different when different base years are used.

Figure 4.18. Evolution of inequality stripping out the variance of the log price index, 1975 prices

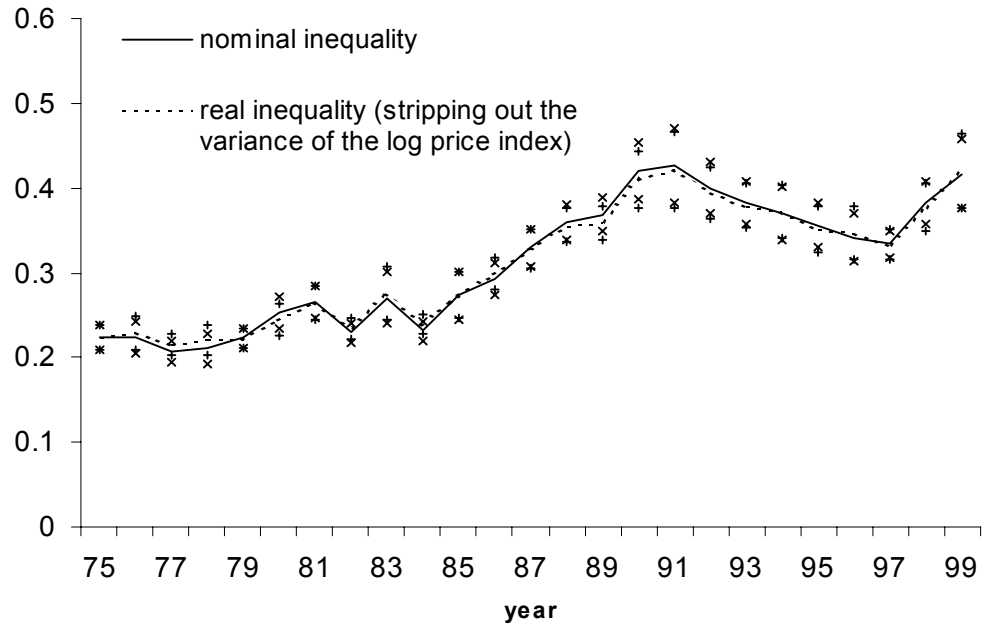

Figure 4.19. Evolution of inequality stripping out the variance of the log price index, 1999 prices

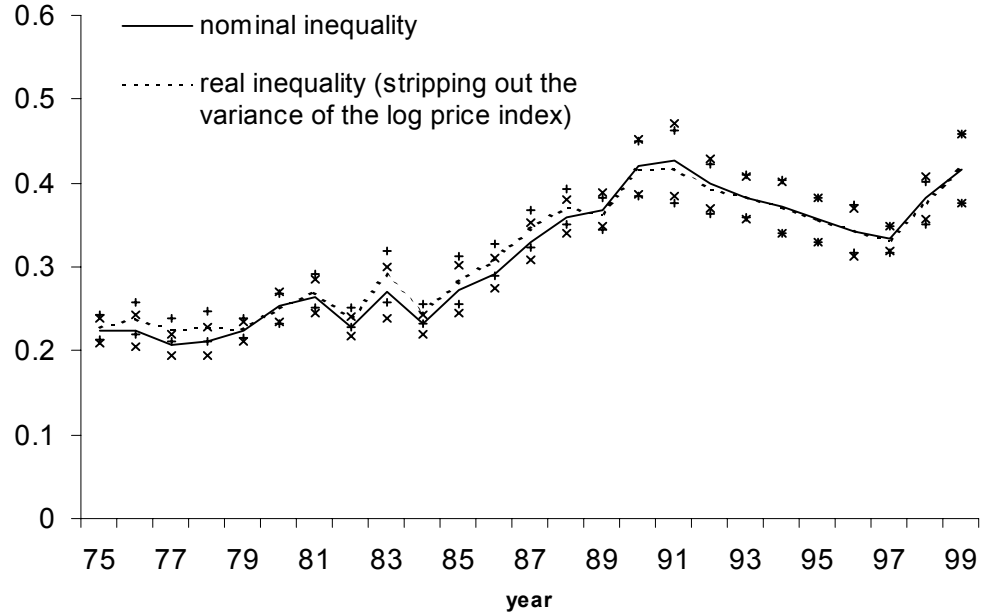

rough agreement between the covariance terms.[25] Note that the counter-intuitive results that we found in Figures 4.14 and 4.15 no longer hold. In 1989 in Figure 4.18, we now find what we expected to find, i.e. lower real income inequality due to the negative relationship between income and inflation; and in 1980 in Figure 4.19, we find that there is no statistically significant difference between real income inequality using a common price index and real income inequality using a household-specific price index, which again is what we expected given the very small covariance term in that year. Note that, even without the contribution of the price index variance, the covariance is enough to make the two series (common and household-specific price indices) go in different directions between 1978 and 1979 regardless of the base period chosen: the use of the household-specific inflation adjustment makes measured inequality fall between these periods for both base periods, whereas the common adjustment shows rising inequality. Figures A.3 and A.4 in the appendix show the difference between the series in Figures 4.18 and 4.19 respectively and the 95 per cent confidence interval on this difference. They take into account the covariance and from them it can be seen that the differences between the series in Figures 4.18 and 4.19 are often statistically significant.

The conclusion from all of this is that, given that it appears to be sensible to adjust different households' purchasing powers by the inflation rate that they specifically face in order to make comparisons of real incomes, then base-period dependence is introduced to the measurement of inequality. Removal of the effects due to the way in which uncorrelated variability in the price index will tend to increase dispersion as we move away from the base period may be one way of reducing this, but base-period dependence will, in general, remain and it is not clear what should be done about this.

In the price index literature, one way of tackling base dependence is to *chain* annual price indices together. That is, calculate a price index linking periods 1 and 2 (using, say, either or both period's demands as weights) and another price index linking periods 2 and 3 (using, say, either or both period's demands as weights); the final index linking periods 1 and 3 is then the product of the two bilateral indices and is known as a chained index. The chained index has a number of advantages but the main ones are that it can allow new goods to be introduced to the index in a timely manner and that it cuts down the differences between different fixed-base indices which will tend to emerge over time (cutting the spread, for example, between the Paasche and Laspeyres indices). Chained price indices are very often used in practice.

A similar approach may be taken to the description of changes in inequality. Instead of holding the base period constant over the 25 years of data and looking at the level of inequality, we could just look at the (bilateral) annual change in inequality, i.e. the change from period 1 to 2, then the change from period 2 to 3, and so on. Overall changes in inequality over longer time periods could then be constructed (if required) by chaining these annual bilateral changes. In each bilateral comparison, there is only the choice of

[25] The covariance terms will not, of course, be identical because they measure the association between income and two different price indices (for a household observed in period *t*, one price index measures the (lower) cost of buying its current set of goods when prices were those in 1975 compared with those in period *t*, while the other measures the increased cost of buying these goods when prices are those in 1999 compared with those in period *t*).

either of the two years as the base period, and the differential effects of base dependence should be small.

Figure 4.20. Annual growth in inequality, base-year prices

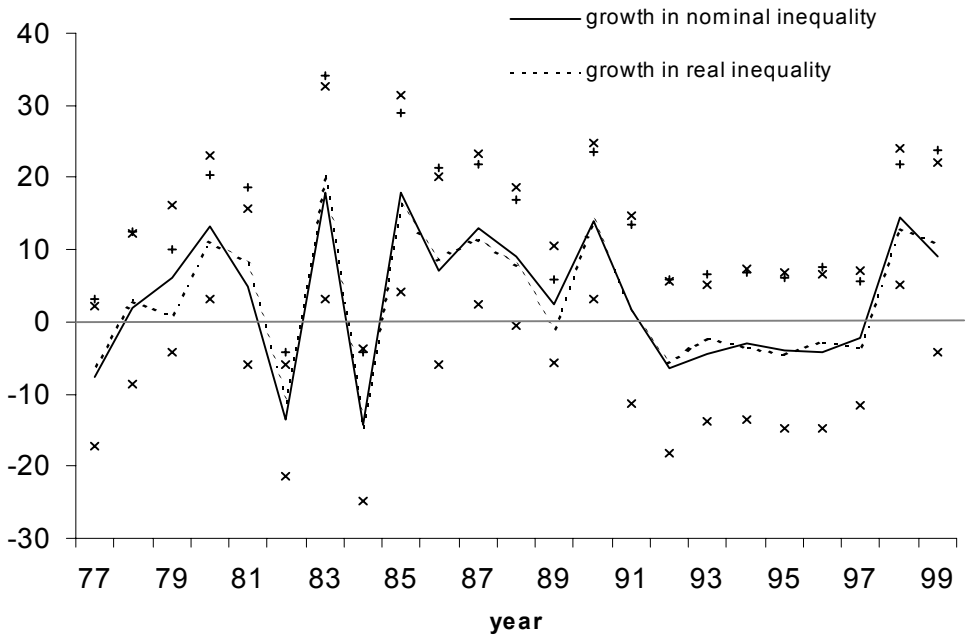

Figure 4.21. Annual growth in inequality, end-year prices

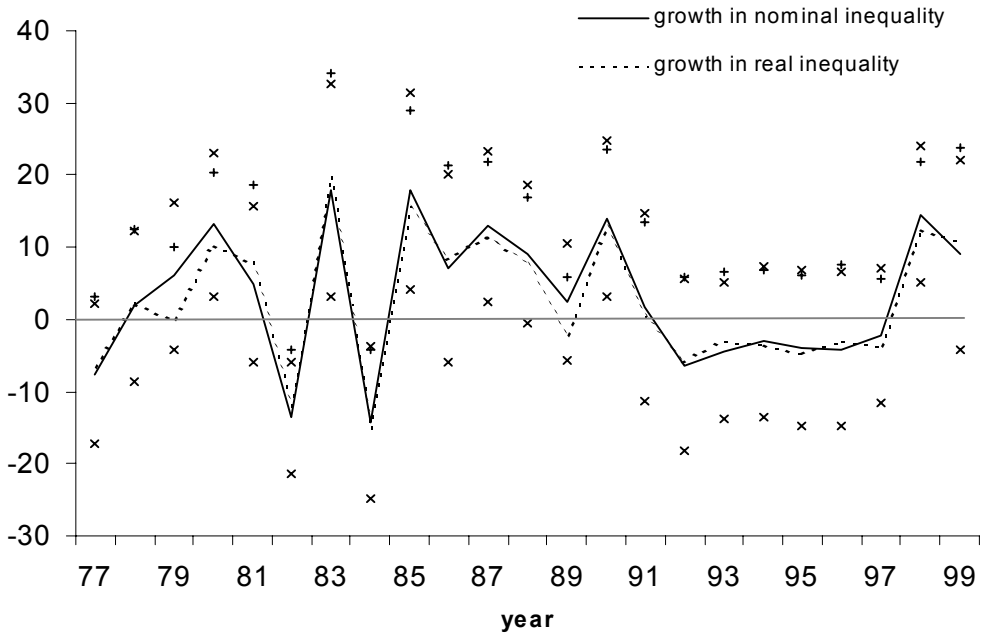

Figure 4.20 shows the annual growth in our inequality measure in each pair of years.[26] In this case, incomes are expressed in base-year prices. For example, the data point for 1977 shows the change in inequality between 1976 and 1977 expressed in 1976 prices, and the data point for 1999 shows the change between 1998 and 1999 expressed in 1998 prices. The solid line is the annual growth rate in nominal income inequality (equivalent to real income inequality using a common deflator). The 95 per cent confidence intervals are shown too. Figure 4.21 also shows the annual growth in our inequality measure in each pair of years. This time, incomes are expressed in final-year prices. For example, the data point for 1977 shows the change in inequality between 1976 and 1977 measured in 1977 prices, and the data point for 1999 shows the change between 1998 and 1999 measured in 1999 prices. Figures A.5 and A.6 in the appendix show the difference between the two series in Figures 4.20 and 4.21 respectively. They also show the 95 per cent confidence intervals on these differences, and from these it can be seen that despite the apparent closeness of the series, the difference between them is, in fact, in general statistically significant.

As can be seen, the series for the annual change in inequality using the household-specific inflation measure are similar in both pictures. This is because the effect of base dependence is minimised by the frequent rebasing as the contribution to inequality of the variance of the log price index is very small in each year. The important point to note from Figures 4.20 and 4.21 is that not allowing for differential inflation can lead to a very different story about year-to-year comparisons in inequality. Sometimes, the differences in the growth rates are small – for example, in 1990 or 1992 (in both figures). However, in 1979, the difference is more than 6 percentage points in Figure 4.21 and more than 5 percentage points in Figure 4.20; in many other years, the difference between the growth rates is fairly large: in most years, the difference is greater than 1 percentage point, and in seven out of the 23 years in Figure 4.20 and in six out of the 23 years in Figure 4.21, the difference is greater than 2 percentage points.

If we do not allow for differential inflation and we make comparisons of income inequality over longer time periods, we may or may not be reaching misleading conclusions, depending on the period that we are looking at. If we take the period from 1987 to 1991 as an example, we can see from Figure 4.7 that this is a five-year period when the rich experienced higher inflation than the poor in each year. It is well known and can be seen from Figure 4.14 that this was a time when nominal income inequality grew rapidly. But if we correctly allowed for the fact that richer households experienced higher inflation over some of this period than the poor, we would conclude that inequality grew by less than what we see when we do not allow for differential inflation. If, however, we take a period when the rich and poor are differentially affected by inflation but the ranking changes frequently, not taking account of this would be less serious. This is because we may reach a misleading conclusion in any one year, but sometimes we would overstate the growth rate of inequality and at other times we would understate it, and these effects would cancel each other out to some extent.

[26] Note that our earliest year of data is now 1976 (giving us a 1977 growth rate) because in order to calculate inequality in 1975 using an inflation rate from 1974 to 1975, we need 1974 prices, and we do not have prices that correspond to our commodity groups in this year.

5. Conclusions

In this Commentary, we have looked at the distribution of inflation rates for UK households in the period 1976–2000. We have looked at the extent and pattern of variations in inflation rates across individual households and groups of households. We have also looked at some of the implications of these variations. We have found the following:

- There is a high degree of variation in inflation rates amongst different households at the same point in time. The annual average standard deviation in inflation is around 2.8 per cent compared with an annual mean rate of inflation of about 7 per cent. The average annual semi-interquartile range over the period is 1.62 percentage points.

- There is a fairly robust pattern that indicates that whenever inflation rates are, on average, high, the dispersion in the rates amongst individual households tends also to be high.

- The headline average rate of inflation is not necessarily a good guide to the actual rates of inflation faced by individual households and it is certainly not always close to the experience of inflation for 'the great majority of households'. On average, only about a third of households at a point in time face inflation rates within 1 percentage point of the average rate. The representativeness of the average rate tends to be lower when inflation is high.

- The practice of plutocratically weighting the headline inflation measure (i.e. giving greater weight to the inflation rates experienced by high-spending, typically richer households) often gives a measure that is statistically significantly different from a democratic measure, although these differences are not quantitatively significant. There is no persistent 'bias' in either direction in the plutocratic index compared to the democratic index over the period studied.

- Average inflation rates defined for different population subgroups are often significantly different from each other. The average rates of inflation of the richest and poorest households are often quite different, although their rankings swap over throughout the period studied. Overall, the average annual inflation rate for the poorest 10 per cent of households over the period was about half of a percentage point lower than that for the richest 10 per cent of households. On occasion, the difference between rich and poor was as much as 5 percentage points.

- Over the period, households with children have experienced inflation little different from that experienced by other households. Pensioners have tended to have a lower rate of inflation than non-pensioners over the period. However, there has been variation in this and, in fact, in one year (1990), as a result of average price indexation, the real value of the pension fell by 4 per cent compared with the level necessary to maintain living standards. In other years, the real value of the pension rose – for example, in 1991, the real value of the pension rose by over 7.5 per cent. The overall effect of average price indexation on the real value of the pension since 1987 has been a real increase of around 3.5 per cent.

- Unlike the case in which a common price index is used to measure real living standards, the measurement of real economic inequality is sensitive to the choice of reference prices when heterogeneity in inflation rates is allowed for. Since this heterogeneity is important, we suggest taking a chained approach to the analysis of inequality and concentrating on annual changes in inequality. This is a standard way of measuring inflation and is an answer to an analogous and related issue of base dependence in fixed-base price indices. This mitigates much of the base-dependence problem and allows for a clearer interpretation of changes in economic inequality.

- It is important to allow for the differential effects of inflation when measuring changes in the distribution of real living standards. We show that the effects of differential price adjustment on the incomes of a population of heterogeneous households can sometimes give strikingly different results from the standard approach which uses an inflation adjustment that is common to all households. The resulting effect on measured inequality can, on occasion, cause these two measures to move in opposite directions. Even when the two measures move in the same direction, not allowing for differential inflation could lead to the annual growth rate in inequality being overstated or understated by as much as 6 percentage points.

Appendix. Further figures

Figure A.1. Evolution of inequality, 1975 prices: difference between the two inequality measures

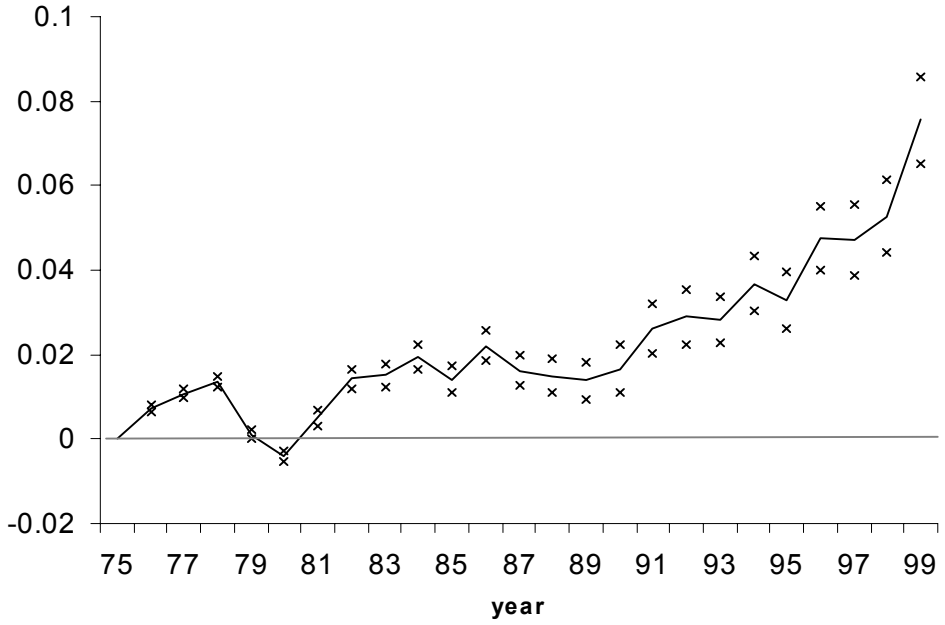

Figure A.2. Evolution of inequality, 1999 prices: difference between the two inequality measures

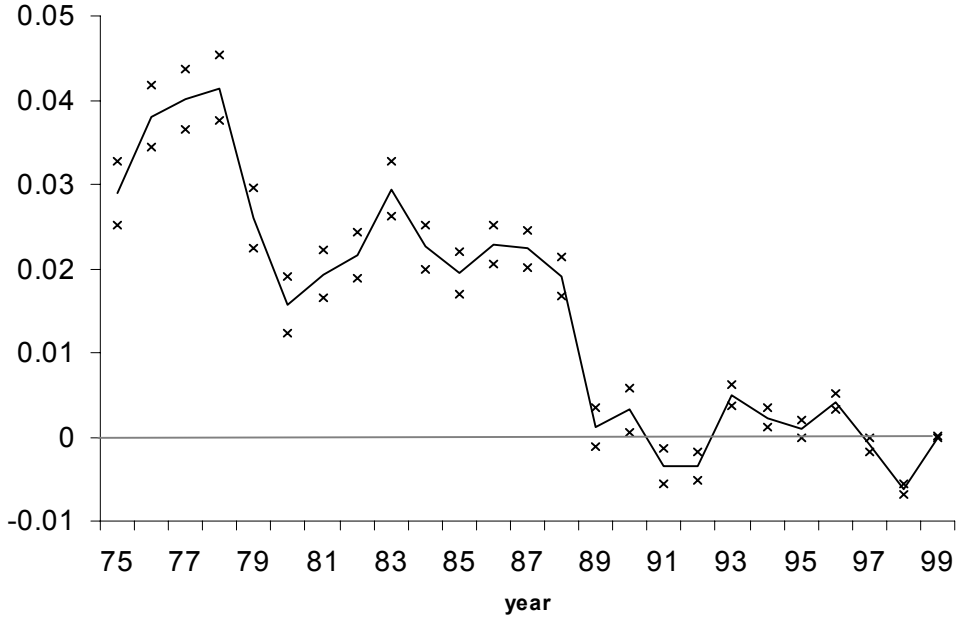

Figure A.3. Evolution of inequality stripping out the variance of the log price index, 1975 prices: difference between the two inequality measures

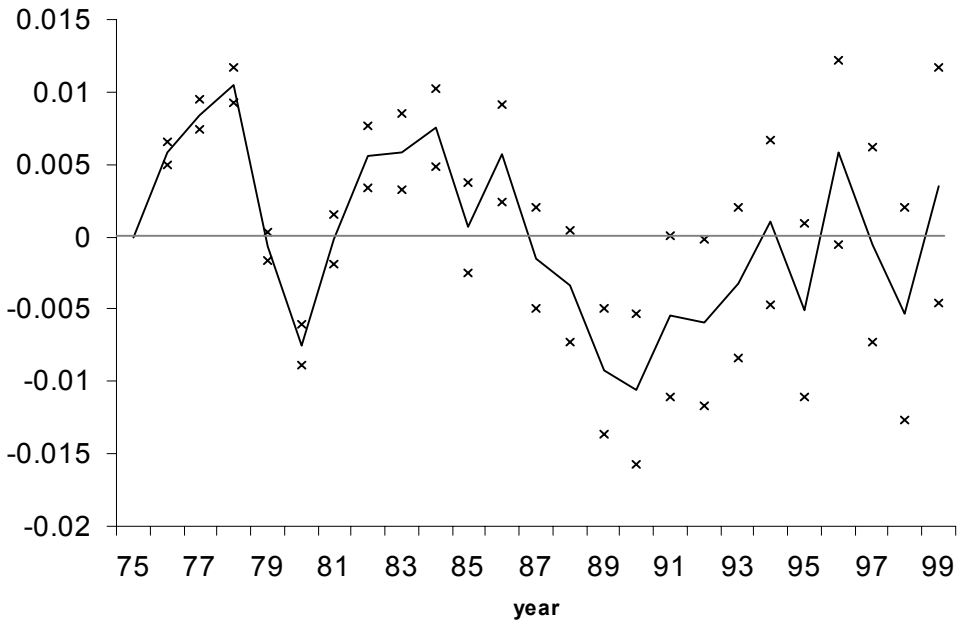

Figure A.4. Evolution of inequality stripping out the variance of the log price index, 1999 prices: difference between the two inequality measures

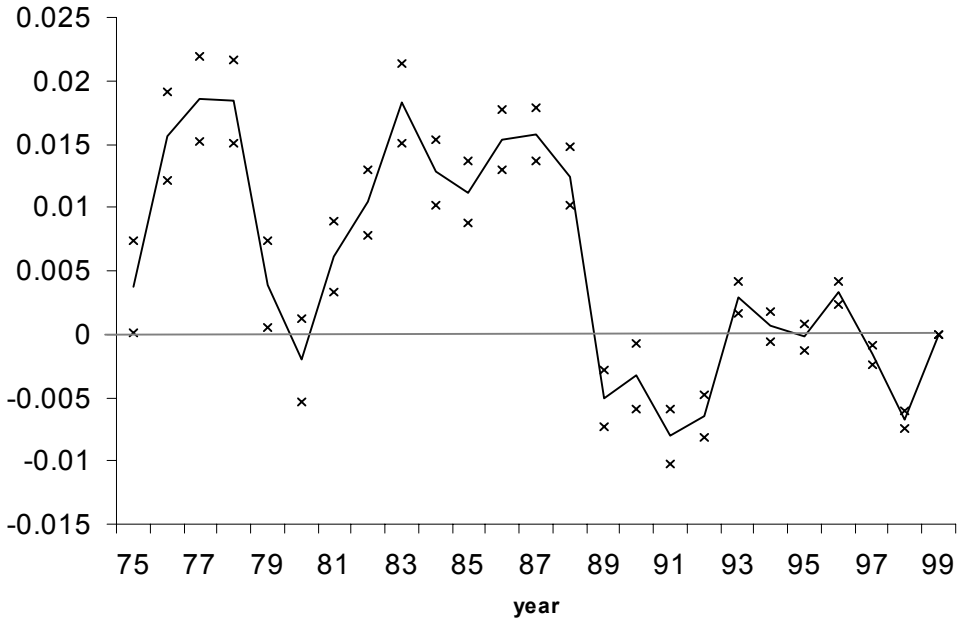

Figure A.5. Annual growth in inequality in base-year prices: difference between the two inequality measures

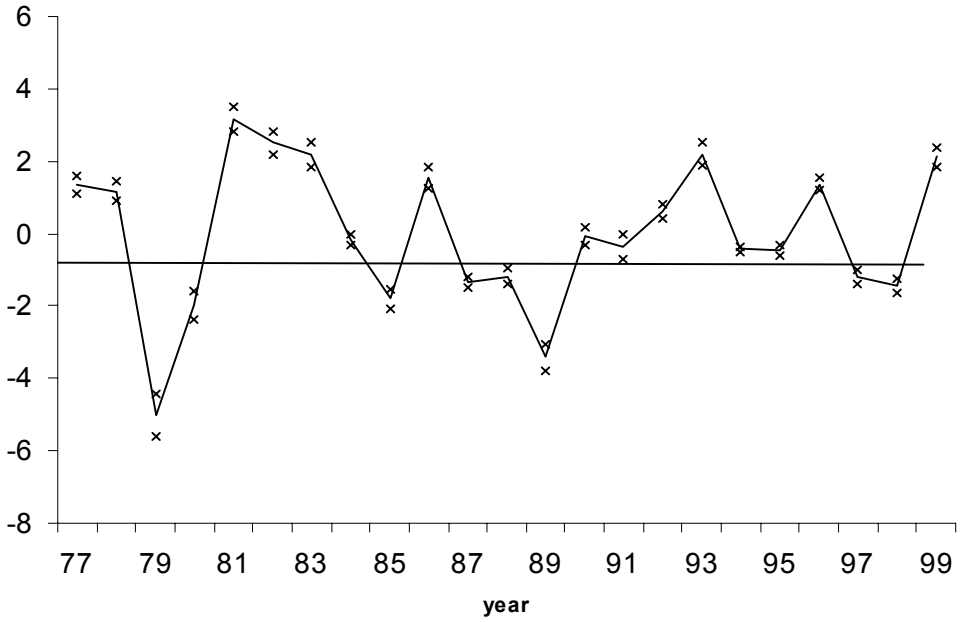

Figure A.6. Annual growth in inequality in end-year prices: difference between the two inequality measures

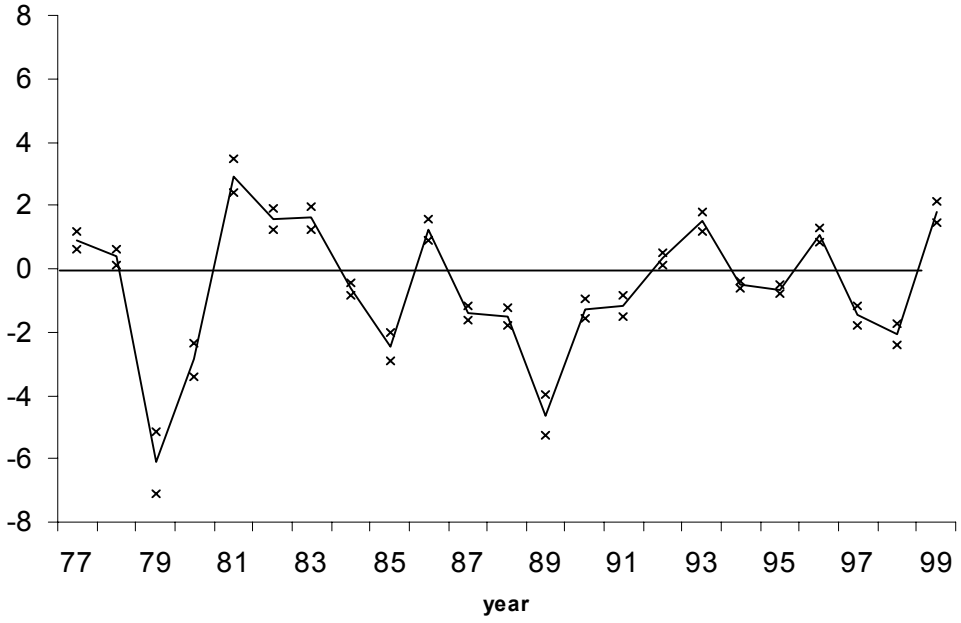

References

Afriat, S. N. (1967), 'The construction of utility functions from expenditure data', *International Economic Review*, vol. 8, pp. 67–77.

Afriat, S. N. (1973), 'On a system of inequalities in demand analysis: an extension of the classical method', *International Economic Review*, vol. 14, pp. 460–72.

Baxter, M. (1998), *The Retail Prices Index Technical Manual*, London: The Stationery Office.

Blundell, R., Browning, M. and Crawford, I. (forthcoming), 'Nonparametric Engel curves and nonparametric demands', *Econometrica*, 2003.

Department of Social Security (1994), *Households Below Average Income: A Statistical Analysis 1979–1991/92*, London: HMSO.

Diewert, W. E. (1981), 'The economic theory of index numbers', in A. Deaton (ed.), *Essays in the Theory and Measurement of Consumer Behaviour in Honour of Sir Richard Stone*, Cambridge: Cambridge University Press.

Diewert, W. E. and Nakamura, A. O. (eds) (1993), *Essays in Index Number Theory*, vol. 1, Amsterdam: North-Holland.

Houthakker, H. S. (1950), 'Revealed preference and the utility function', *Economica*, New Series, vol. 17, pp. 159–74.

Maasoumi, S. (1997), 'Empirical analyses of inequality and welfare', in H. Pesaran and P. Schmidt (eds), *Handbook of Applied Econometrics, Vol. II: Microeconomics*, Oxford: Blackwell.

Samuelson, P. A. (1938), 'The empirical implications of utility analysis', *Econometrica*, vol. 6, pp. 344–56.

Schultze, C. L. and Mackie, C. (eds) (2002), *At What Price: Conceptualising and Measuring Cost-of-Living and Price Indexes*, Washington, DC: National Academy Press.

Tanner, S. (1998), 'How much do consumers spend? Comparing the FES and National Accounts', in J. Banks and P. Johnson (eds), *How Reliable Is the Family Expenditure Survey?*, London: Institute for Fiscal Studies.

Triplett, J. E. (2001), 'Should the cost of living index provide the conceptual framework for a consumer price index?', *Economic Journal*, vol. 111, pp. F311–34.

Varian, H. (1982), 'The nonparametric approach to demand analysis', *Econometrica*, vol. 50, pp. 945–74.