IFS

Permanent differences?

Income and expenditure inequality in the 1990s and 2000s

Alissa Goodman
Zoë Oldfield

The Institute for Fiscal Studies

IFS Report No. 66

Permanent Differences?
Income and Expenditure Inequality in the 1990s and 2000s

Alissa Goodman
Zoë Oldfield

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 1-903274-38-9

Printed by

Patersons, Tunbridge Wells

Preface

The authors would like to thank Tom Clark, Andrew Leicester and Andrew Shephard for their contributions to this Report, and James Banks, Ian Crawford and Jonathan Shaw for their very helpful comments on earlier drafts. Financial support from the Nuffield Foundation ('Inequality in the 1990s', grant number OPD/00111/G) and from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy at IFS (grant number M535255111) is also very gratefully acknowledged. The Nuffield Foundation is a charitable trust established by Lord Nuffield. Its widest charitable objective is the advancement of social well-being. The Foundation has long had an interest in social welfare, and it has supported this project to stimulate public discussion and policy development. The views expressed are, however, those of the authors and not necessarily those of the Foundation. Data from the Family Resources Survey and the Households Below Average Income data-sets were made available by the Department for Work and Pensions, which bears no responsibility for the interpretation of the data in this Report. Material from the Family Expenditure Survey was made available by the Office for National Statistics through the ESRC Data Archive and has been used by permission of the Controller of HMSO. Any errors and all opinions expressed are those of the authors.

Contents

Executive Summary

This Report sets out what has happened to income and expenditure inequality in the 1990s and early 2000s, comparing the changes with those of previous decades.

Income and expenditure are alternative measures of living standards. Measures of income inequality provide us with a snapshot of income differences across the population. Expenditure inequality tells us more about the longer-run, or lifetime, differences in living standards between people.

Throughout the Report, the income measure we use is household income from all sources (such as earnings, social security benefits, self-employment income and investment income), and it is measured net of direct taxes. The expenditure measure we use captures spending on all goods, including housing and durables. Both income and spending are equivalised to take into account family size and composition.

Income Inequality

1. The most commonly used measure of inequality is the Gini coefficient. The Gini coefficient for income fluctuated over the 1990s, but was higher at the end of the 1990s than at the start.
2. The growth in inequality as measured by the Gini coefficient over the 1990s was small compared with its growth over the 1980s. Between 1980 and 1990, the Gini coefficient for income rose by 8.4 percentage points (or 33 per cent). Between 1990 and 2000, there was just a 1.2 percentage point (or 4 per cent) growth in this measure of inequality.
3. What increase in the Gini coefficient there was over the 1990s mostly occurred in the second half of the decade. It was driven by rapidly rising incomes at the top of the income scale, particularly amongst the top half-a-million people, as well as by incomes at the very bottom of the scale that have not kept pace with general income growth across the population.
4. Rising incomes at the top of the scale are also reflected in a growing share of total income held by the top 1 per cent. By the start of the 2000s, the top 1 per cent held around 8 per cent of the total income, compared with income shares of around 5.8 per cent in 1990 and 3.5 per cent in 1980.
5. Among the large majority, incomes became a little more equal over the 1990s. For example, the 90/10 ratio — a measure of inequality that does not take into account the incomes of the very poorest or very richest — fell back over the early 1990s, and then remained roughly flat over the later part of the 1990s. By the end of the decade, the 90/10 ratio was around 95 per cent of its 1990 level.
6. Over the early 2000s, the Gini coefficient for income has fallen slightly, though not by a statistically significant amount.
7. These findings suggest that the sustained period of rising inequality during the 1980s has been halted. However, the level of inequality inherited after the big inequality rise of the 1980s has not been much reversed, and income inequality remains near a 40-year high.

Expenditure Inequality

1. The Gini coefficient for expenditure fluctuated over the 1990s, but ended the decade around 2.5 percentage points (or 7 per cent) below its 1990 level.

2. This fall in expenditure inequality did little to reverse the much bigger rise in inequality seen over the 1980s. At the end of the 1990s, expenditure inequality was still at the level experienced in 1987 and 3 percentage points (or 10 per cent) above the level at the beginning of the 1980s.
3. Comparing the paths of expenditure and income inequality, we find that the Gini coefficients for income and spending diverged during the 1980s, with income inequality growing faster than spending inequality. This divergence continued in the first half of the 1990s, when expenditure inequality fell but income inequality was fairly stable.
4. One explanation for the divergence — that income has become more volatile — has been studied in some detail in previous research for the 1980s. This explanation makes sense because expenditure inequality tends to reflect longer-run, or lifetime, differences in people's circumstances but income inequality also reflects short-term income volatility. More work is needed to see if this explanation continues to hold into the mid-1990s, when expenditure inequality was falling.
5. In the second half of the 1990s, income and expenditure inequality have moved in parallel. As well as explaining the continued divergence in the first half of the 1990s, future research might also focus on why the two series have moved in parallel in recent years.

CHAPTER 1
Introduction

The growth in income inequality in Britain over the 1980s was exceptional; the large and sustained increase in inequality — amongst the largest in the industrialised world (Atkinson, 1999) — spawned an equally large literature seeking to chart it, explain it, decompose it into its various parts and assess its implications (see, for example, Johnson and Stark (1989), Johnson and Webb (1993), Goodman, Johnson and Webb (1997), Hills (1996), Jenkins (1991 and 1995) and Jenkins and Cowell (1994), to name but a few). Studies of the distribution of expenditure in Britain also led to some important conclusions about the extent to which the growth in inequality in the 1980s was driven by lifetime or 'permanent' differences in living standards, and how much was related to more transitory income changes (Blundell and Preston, 1998). Relatively little has been said about the changes to income and spending inequality over the 1990s and beyond. This Report aims to fill this gap by setting out what has happened to incomes and income inequality, focusing in particular on the experience of the 1990s. We also consider the distribution of household spending and what we can learn from the differences between the distributions of income and expenditure over the 1990s.

What we show is that, contrary to popular perception perhaps, the 1990s were quite different from the 1980s in terms of both income and spending inequality. Although inequality in after-tax incomes rose on some measures, it fell on others; what rising inequality there was occurred during some years of the decade, whilst other periods saw inequality fall. Focusing on household expenditure, rather than income, we find inequality falling and then rising over the decade. The level of inequality remains historically high, but the overall picture we are left with is of a halt in the rapid inequality growth of the 1980s, and instead a return to a fluctuating pattern of inequality.

The purpose of our Report is mainly descriptive: to chart the trends in inequality in recent years. However, we do give some tentative explanations for the trends we find.

Chapter 2 sets out the measurement issues we contend with in this Report. Chapter 3 describes the main changes in the distribution of income in the 1990s and beyond, contrasting these changes with those in previous decades. Chapter 4 shows how the distribution of household expenditure has changed and looks to reconcile the different patterns in income and expenditure inequality. Chapter 5 concludes.

CHAPTER 2
Measurement Issues and Data

When we chart changes in inequality, the concept that most people are ultimately concerned about is well-being and the extent to which there are differences between individuals in their well-being. Our interest in this Report is in economic, or material, well-being. We focus on inequality in two different measures, both of which attempt to capture economic well-being: household income and household expenditure. Of course, economic well-being depends on many more things than just these simple measures of material circumstances. However, these two measures do take us some way towards capturing material living standards and how they differ across the population.

2.1 Income or Expenditure?

The most commonly used measure of material circumstances for charting changes in inequality in Britain is income. For most people, income will fluctuate, possibly both in the short term (due to unemployment spells, job changes, changes in family circumstances, variable sources of income, etc.) and over their lifetime (with a typical income profile showing income rising and then falling over an adult's lifetime). Although short-term income inequality may be important, we are often interested in more 'permanent', or lifetime, income differences between people, rather than in differences due to these short-term fluctuations or changes over the life cycle. In this respect, the most appropriate measure of living standards would be so-called 'permanent' or lifetime income. Unfortunately, the data required to compute lifetime income do not exist; surveys of income usually record income over a period of a week, a month or perhaps a year.

Data on household expenditure, on the other hand, can give us a useful insight into the longer-term resources available to individuals. Over a lifetime, people can counteract fluctuations in income by saving during periods of higher income and borrowing or running down savings (in order to finance consumption) at times when income is lower. For this reason, consumption is perhaps a better measure of 'permanent', or lifetime, living standards than current income. Although consumption itself is fundamentally hard to measure, we can use expenditure as a proxy.

Using expenditure to proxy consumption is imperfect because many goods provide a flow of benefits over time and the purchase price does not accurately reflect that benefit. Large durable goods such as televisions and freezers are good examples.

In this Report, we analyse inequality in material living standards using both income and expenditure. In general, the income measure will reflect both 'permanent' and transitory (or short-term) circumstances, whereas spending will be a better gauge of longer-term circumstances. For each of these measures, there are a number of further specific issues to consider.

2.2 The Definition of Income and Expenditure

The definition of income that we use throughout the Report is that used for the Households Below Average Income (HBAI) statistics produced annually by the Department for Work and Pensions. This includes income from all sources including earnings, self-employment income, social security benefits, private pensions,

investments and other incomes.[1] The measure of income that we use is net of direct tax (income tax, National Insurance and council tax). We do not attempt to calculate the value of benefits in-kind, such as health or education spending (but see Lakin (2004) and Sefton (2002) for attempts to do this).

One important issue is the treatment of housing costs in the measurement of income. There are two broad possibilities: a measure of income before housing costs have been paid for ('before-housing-costs' or BHC income) and a more disposable measure of income, after housing costs have been paid for ('after-housing-costs' or AHC income). In this Report, we only consider incomes on a BHC basis: this is the most commonly used measure when considering incomes across the whole population. When focusing specifically on low incomes and poverty, AHC measures are also frequently considered (see Brewer et al. (2004) for a discussion of the issues surrounding the choice of measure).

The definition of expenditure that we use is total expenditure including expenditure on durables and housing but excluding any consumption in-kind such as home-grown food. There are also important issues to consider in the treatment of housing costs in measuring expenditure. In particular, although we observe housing costs, it is very difficult to measure housing *consumption.* For those who own their homes outright, or live rent-free, measuring housing consumption is very difficult because they show zero spending on housing but are certainly consuming housing services. Measuring housing consumption is also not straightforward for households that rent or pay mortgages.[2] The approach taken here is to measure housing expenditure rather than consumption. The exact approach taken is to measure housing costs on the same basis as that used by the Department for Work and Pensions to calculate AHC income in official poverty statistics (see Department for Work and Pensions (2004)). This includes rent (gross of housing benefit), water rates and mortgage interest payments.

These issues highlight once again the imperfect nature of using expenditure as a measure of consumption.

Income and consumption sharing

The extent of sharing of income and consumption within households or families will be important in determining the extent of inequality across individuals. However, finding out how much sharing takes place is very difficult. For this reason, we use the same approach that is used in the calculation of the HBAI statistics, which is to use *household* measures of income and expenditure, scaled appropriately using 'equivalence scales' to take account of differing family size and economies of scale associated with living with other people rather than alone. This approach implicitly assumes that both income and spending are shared equally across household members. (The implications of using individual incomes instead of household incomes have been considered in Women and Equality Unit (2003).)

There are many different equivalence scales that could be used. The one we use — again following HBAI methodology — is based on work by McClements (1977); we scale incomes and expenditures to the value that a childless couple would need in order to be equally well off (in income or expenditure terms).

[1]See Department for Work and Pensions (2004) for a fuller statement of the income definition.

[2]For a more detailed discussion of the issues involved in measuring housing costs, see Blow, Leicester and Oldfield (2004).

Sources of data

We have detailed data on household incomes from two different surveys. From 1961 onwards, we have the Family Expenditure Survey (FES). In most years, this has information on around 7,000 households, but between 1961 and 1966, sample sizes were just 3,000–3,500 per year. Since the financial year starting in 1994/95, in addition to the FES we have the much larger Family Resources Survey (FRS). This has around 25,000 households each year. Unlike the FES, from its inception until 2002, the FRS excluded Northern Ireland, so we exclude FES households from that province throughout this Report. This leaves us with a consistent focus on Great Britain. Both the FES and the FRS record household and individual demographic details.

In addition to income data, the FES also contains a comprehensive measure of expenditure. It asks household members to complete a diary of expenditure over a two-week period. Our measure of spending begins in 1974 and ends in 1999/2000.[3]

While the majority of expenditure is recorded through the completion of the two-week diary, expenditure on some items is recorded retrospectively. This means that respondents are asked to recall how much they have spent on some items in a given time period. These items are usually those where payment is infrequent but occurs at regular intervals (for example, utility bills) or large durable items (such as cars). In the late 1980s, the number of items on which spending is recorded retrospectively was expanded. To avoid any spurious change in inequality in expenditure caused by this change, we use a measure of household spending that, for each item, uses either diary information throughout or retrospective information throughout.

Participation in both the FES and the FRS surveys is voluntary, so, inevitably, a proportion of people who are surveyed do not take part. Data from the census and other sources show us that in certain dimensions (family type and region, for example), the sample we are left with is non-representative — some groups of families are more likely to agree to take part in the survey than others. To try to counteract this bias, we weight the data so that classes of families that are under-represented in the sample survey have their weights increased relative to types of families that are more likely to be in the survey. In this way, the weighted sample is calibrated to match the overall population as closely as possible in a range of dimensions.

Adjusting the incomes of the very rich

One important adjustment to the survey data affects households with the very highest incomes. The high proportion of incomes held by these families means that slight over- or under-sampling, or a slight bias in the incomes of this subsample, can have significant bearing both on the mean income and on the calculated value of many inequality summary statistics.

This problem is confronted by using the Inland Revenue's administrative data, the Survey of Personal Incomes (SPI). From this source, we establish the number of individuals above specific high-income thresholds in each year. We then calibrate the weights of those households in our sample that contain such high-income individuals so that our weighted sample contains the correct proportion of households with such incomes. Finally, the average income of all those above the threshold in the SPI is

[3]Although the FES data are available from 1961 to 2001/02, our consistent spending data, which have been adjusted in a number of ways to be comparable to our household income series, have only been calculated from 1974 to 1999/2000.

assigned to everyone above the threshold in our sample. In general, this 'SPI adjustment' is applied to roughly the richest ½ per cent of individuals in our data. The nature of the SPI adjustment means that certain measures of inequality are affected while others are not: any measure that takes into account the entire distribution may be biased downwards (because the variation at the very top of the distribution is reduced by the SPI adjustment), while any measure that does not (such as the 90/10 ratio) will not be affected. However, as long as the SPI adjustment is applied consistently over time, the conclusions we reach about changes in inequality should not be affected.

2.3 Measures of Inequality

Inequality is not a concept with a single, precise mathematical definition. In this Report, we present a range of measures of inequality, including

- the Gini coefficient;
- the 90/10 ratio;
- income shares of different decile and percentile groups of the distribution;
- half the squared coefficient of variation.

We have a number of reasons to present a range of measures of inequality. First, we will sometimes be interested in one concept of inequality but not another. For example, the Gini coefficient is a measure of inequality that takes into account incomes at all points in the income scale, measuring differences in income between all individuals, including those at the very bottom and very top of the distribution. (See Box 2.1.)

BOX 2.1

The Gini coefficient

A widely used measure of income inequality is the Gini coefficient. The Gini coefficient benefits from an intuitive geometric interpretation in the form of the Lorenz curve. In Figure 2.1, the horizontal axis corresponds to the cumulative percentage of individuals in the population lined up in ascending order on the basis of their household income, while the vertical axis gives the cumulative percentage of income. The Lorenz curve then shows the relationship between the percentage of income recipients and the percentage of income actually received. Figure 2.1 shows the Lorenz curve drawn using actual data from 2002/03. If income were equally distributed across households, then 10 per cent of the population would have exactly 10 per cent of total income, 20 per cent of the population would have 20 per cent of total income, and so on. The line of perfect equality is therefore given by the 45-degree line, AB. Note that the further is the Lorenz curve from the line of perfect equality, the greater is the degree of inequality.

The Gini coefficient is obtained by taking the ratio of the shaded area to the area of triangle ABO. When there is perfect equality, the shaded area will have zero measure so that the Gini coefficient will be 0. Conversely, when there is complete inequality (a single household having command over the entire income of the economy), the shaded area will coincide with triangle ABO so that the Gini coefficient will equal 1.

FIGURE 2.1

The Lorenz curve and Gini coefficient, 2002/03

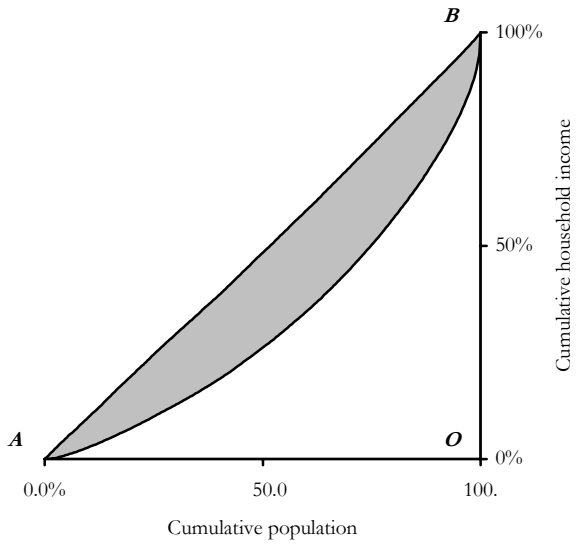

Source: Brewer et al., 2004.

By contrast, the 90/10 ratio (the ratio of the income at the 90th percentile point of the income distribution to the income at the 10th percentile point) measures the scale of overall inequality by looking at the gulf in incomes that occurs over the middle 80 per cent of the distribution, comparing the typical income of an individual 'near' the top of the income distribution with the typical income of an individual 'near' the bottom. Thus the 90/10 ratio would not register any inequality, for example, in a society where everyone had the same income apart from a tiny oligarchy — say the top 1 per cent — who owned a huge proportion of overall income.

A second reason for looking at different measures of inequality relates to the decomposability of some measures. For example, in Chapter 4, we present some results on expenditure inequality using half the squared coefficient of variation, because this measure can be decomposed into inequality between different sources of income or spending. As a measure of inequality, it is particularly sensitive to high income or expenditure outliers. (See Box 2.2.)

BOX 2.2

Half the squared coefficient of variation

This measure of inequality is simply a measure of the variance, normalised by the square of the mean. Although it is very sensitive to outliers, and hence not often used for simple descriptions of trends in inequality, it is frequently used for its ready decomposability into contributions from different sources (in our case, we use it to break down inequality in expenditure into contributions from different types of expenditure).

The measure can be written as

$$I_{cv} = \frac{\sigma^2}{2\bar{y}^2}.$$

Its decomposability has been shown by Shorrocks (1982) and also by Cowell (1995).

CHAPTER 3
Income Inequality

This chapter brings up to date what has happened to various measures of income inequality in Britain. In doing this, we pay particular attention to the experience of the 1990s and early 2000s, and how this differs from the changes taking place in the preceding decades. For our analysis, the time horizon is constrained to the last 40 years, because this is the time period over which we can use the detailed micro-level survey data on the incomes and characteristics of households in Britain described in the last chapter. But much can also be gained from considering the picture over a longer period: in Section 3.4, we set our findings in the context of research looking at changes in incomes over the whole of the twentieth century.

Before continuing to our main analysis of income inequality over time, we begin by setting out a picture of the income distribution in the most recent year for which we have data, 2002/03, and how income levels at different points in the income scale compare.

3.1 A Picture of the Distribution of Income

A good place to start in understanding the distribution of income is to look at how many people are to be found at different income levels. Figure 3.1 presents such a picture, showing the income distribution in 2002/03, the latest year for which complete data are available. This graph shows the number of people living in households with different equivalised income levels, grouped into £10 income bands. The height of each bar represents the number of people in that income band. As can be seen, the current distribution is highly skewed, with 65 per cent of individuals having household incomes below the mean, of £396 per week. The median income is considerably lower, at £323 per week. Remember that incomes are measured at the household level, net of direct taxes, and are expressed as the equivalent income for a childless couple (see Chapter 2).

It should also be noted that the distribution shown in Figure 3.1 has been truncated at income levels in excess of £1,100 per week; 1.2 million individuals (out of a private household population of approximately 57 million) have incomes above this amount.

The graph also shows how the distribution of income divides into 10 approximately[4] equally sized groups, known as decile groups. The alternately shaded sections represent these different decile groups. The first decile group contains the poorest 10 per cent of the population, the second decile group contains the next poorest 10 per cent, and so on. As can be seen, the distribution is particularly concentrated within a fairly narrow range of incomes in decile groups 2 to 4. As we move further up the income distribution, a widening of the decile group bands can be seen. Note that the 10th decile group band is much wider than is shown in Figure 3.1 because of the graph being truncated at incomes of £1,100 and above.

[4]See Note to Figure 3.1.

FIGURE 3.1

The income distribution in 2002/03

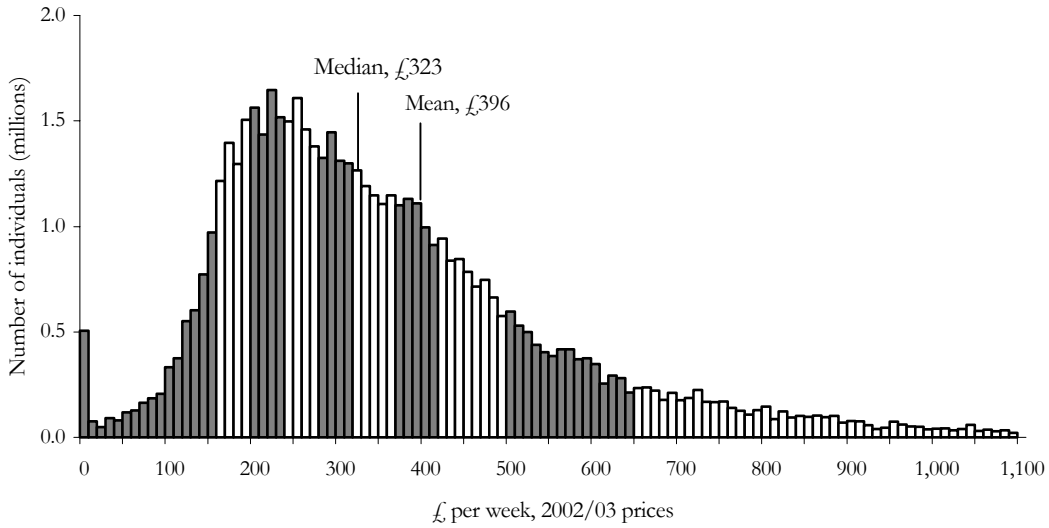

Note: The alternately shaded sections represent the decile groups, marked to the nearest £10 income band.

Source: Brewer et al., 2004.

3.2 Changes in Income Levels over Time

The value of the incomes of people at different points in the income distribution has changed quite markedly over time. Average income, as measured by the mean, has trended upwards in a cyclical pattern over time. Since 1961, the mean equivalised income has more than doubled in real terms, from under £200 per week in today's prices in the early 1960s, to nearly £400 in 2002/03 (Figure 3.2).

FIGURE 3.2

The real value of incomes

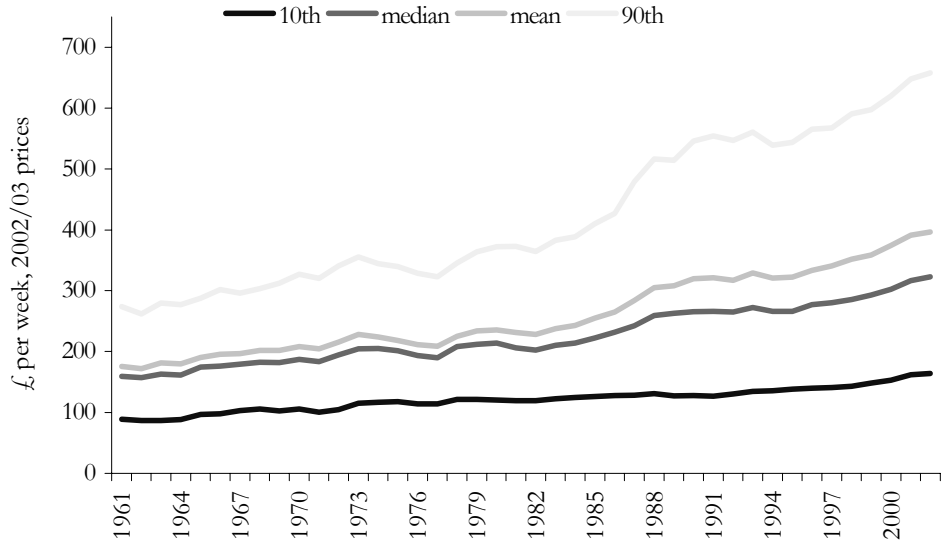

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Sources: Family Expenditure Survey, 1961 to 1993/94; Family Resources Survey, 1994/95 onwards.

The pattern of changes in the median income has been quite similar to the pattern for the mean income, although it is interesting to note that from the early 1980s onwards, the mean income has tended to grow faster than the median. This is because, unlike the median, the mean is a measure of the average that is affected by the size of incomes at the top of the income scale, and it is here that income growth has been strongest over this period.

This point can also be seen from the 90th percentile line on Figure 3.2, which shows that over the 40-year period, the incomes of the richest tenth have risen considerably faster than those at points lower down the income scale, and almost twice as fast as those of the poorest tenth (the income of the richest tenth grew 140 per cent in real terms between 1961 and 2002/03, whilst the income of the poorest tenth grew 86 per cent over this period).

These patterns of income growth at different points in the income distribution suggest that there have been significant changes in income inequality over time. These are the subject of the remainder of this chapter.

3.3 What Has Happened to Income Inequality?

As discussed in Chapter 2, there are a number of different ways of measuring inequality. Each provides a different sort of summary of 'the gap between rich and poor'. The most commonly used summary measure of inequality is the Gini coefficient, and it is with this that we begin our description of changes in inequality over time. We then go on to show how some other measures of inequality compare.

The Gini coefficient

The rapid growth in income inequality over the 1980s is the starkest feature of the changes in the Gini coefficient over the last 40 years. Our previous work has shown that over the 1960s and 1970s, the Gini coefficient followed a fluctuating pattern, before rising sharply over the 1980s (see, for example, Goodman, Johnson and Webb (1997)). Figure 3.3 confirms this, showing that the Gini fluctuated closely around an average 'long-run' level of around 0.25, with total fluctuations of no more than 1.5 percentage points in either direction over the 1960s and 1970s. The most notable change during this time was the downward trend in the Gini coefficient over much of the 1970s, during a period of explicitly egalitarian pay policies.

The subsequent path of the Gini coefficient involves change of a different order of magnitude, which makes the earlier fluctuations appear small in contrast. The rise in inequality, which started in the late 1970s, occurred at a steady pace through the early part of the 1980s and accelerated over the later part of the decade. The total increase in the Gini coefficient for income was 8.9 percentage points between 1979 and 1990 (8.4 percentage points between 1980 and 1990). This increase has been documented elsewhere as amongst the fastest income inequality growth anywhere in the industrialised world over the period (see, for example, Atkinson (1999)).

Over the 1990s, the Gini coefficient appears to have taken a more fluctuating path once again, highlighting the fact that the experience of inequality in the 1990s and beyond has been quite different from that of the 1980s. This is an important point, which appears to have been little remarked upon hitherto (although see Brewer et al. (2004) for a discussion of the late 1990s). Although there was an overall growth in the Gini coefficient between 1990 and 2000 (of around 1.2 percentage points), this growth occurred mostly in the second half of the 1990s, and the magnitude of change is very small compared with the magnitude of change seen in the previous decade.

Figure 3.4 brings this out more clearly, showing the path of the Gini coefficient for income over each decade from the 1960s through to the early 2000s, with the start of each decade marked by the year '0'. Whilst the Gini coefficient ended the 1980s almost 30 per cent higher than its level in 1980, it ended the 1990s just 2 per cent higher than its level at the start of that decade.

FIGURE 3.3

The income Gini coefficient

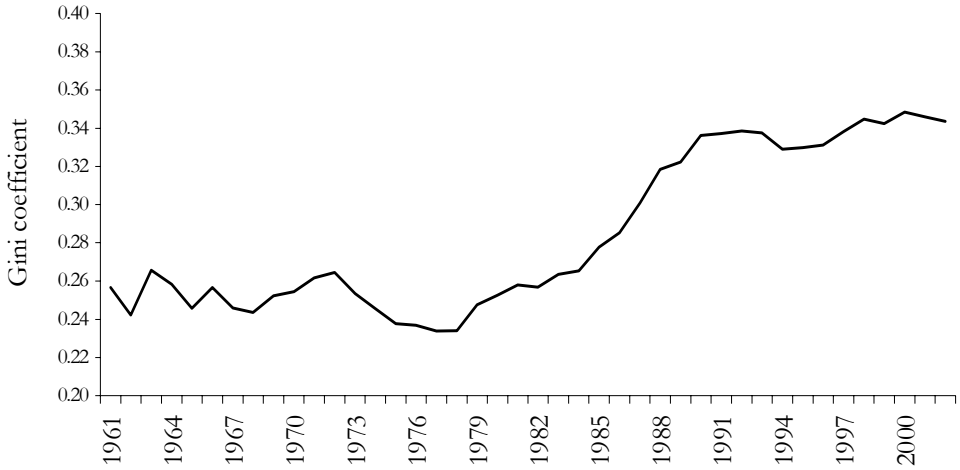

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Sources: Family Expenditure Survey, 1961 to 1993/94; Family Resources Survey, 1994/95 onwards.

FIGURE 3.4

Path of the Gini coefficient over different decades

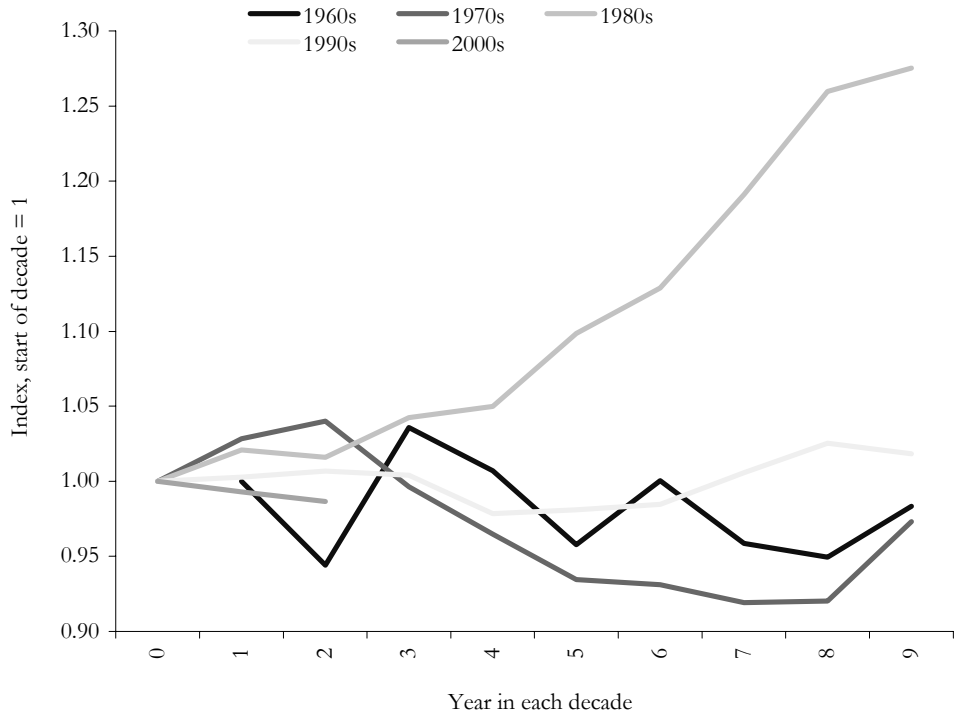

Notes: Points are drawn for calendar years up to and including 1992, and financial years thereafter. The 1960s series is indexed at 1.0 in the year 1961 as this is when our data series begins.

Sources: Family Expenditure Survey, 1961 to 1993/94; Family Resources Survey, 1994/95 onwards.

Over the first two years of the 2000s, income inequality as measured by the Gini coefficient has declined slightly, though not yet by a statistically significant amount.[5] In spite of these downward movements in inequality, however, the Gini coefficient is still at historically high levels — approaching 10 percentage points higher than it was in 1979. So although inequality has stopped rising, it has certainly not returned to the levels we saw before the big increase of the 1980s.

BOX 3.1

The switch from the FES to the FRS: what difference did it make?

It is important to bear in mind that assessing the trends over the 1990s is complicated by a possible measurement discontinuity during the middle part of the decade. The introduction of the Family Resources Survey (FRS) in 1994/95 meant some changes in the measurement of income both in the new FRS and in the ongoing Family Expenditure Survey (FES). Our series uses FES data up until 1993/94, and then switches to using the FRS from 1994/95 onwards.

Although the official income series based on the FES has not been continued, Figure 3.5 shows our calculations of the path of the Gini coefficient based on the FES series for BHC income, had it been continued for five years after the introduction of the FRS. It can be seen that the drop in income inequality between 1993/94 and 1994/95 is slightly more pronounced on the FES series, and the rise in inequality between 1996/97 and 1998/99 is slightly sharper: overall, the FES series shows a slightly bigger increase in inequality over the mid to late 1990s than our series shows.

However, it is important to remember that the FES series itself may not be fully consistent over this time, since a number of questions relating to employment and self-employment were changed to align them more closely to those in the FRS.

FIGURE 3.5

Comparison of the Gini coefficient using the Family Resources Survey and the Family Expenditure Survey

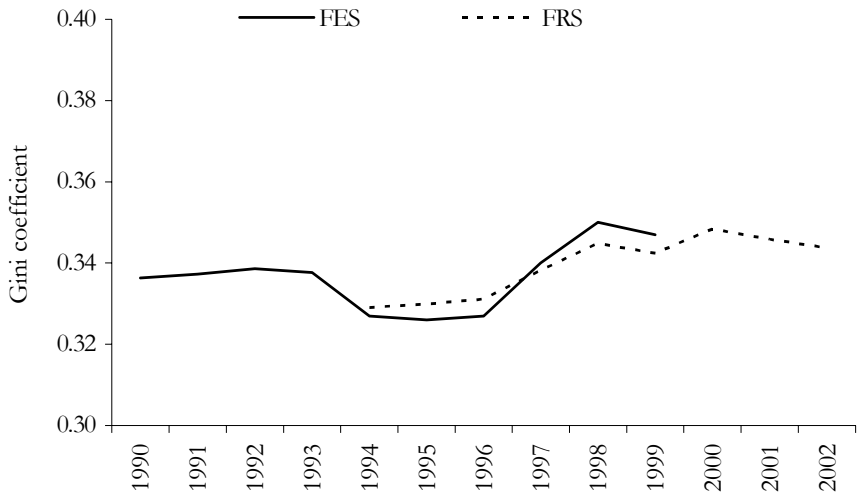

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Sources: Family Expenditure Survey and Family Resources Survey.

[5]In the Appendix, we set out confidence intervals around the Gini coefficient, and present tests of the statistical significance of changes in the Gini coefficient between pairs of years, in order to explore how robust our findings are to possible sampling error in our data.

We now go on to look more closely at these changes in inequality, by considering how different measures of inequality compare, and how income changes at different parts of the income distribution have driven the movements in the Gini coefficient. Given that the most striking developments in inequality occur after 1979, we will use this date as our base period for much of what follows in this Report.

The 90/10 ratio

As we have shown, the growth in inequality measured by the Gini coefficient was much slower over the 1990s than in the preceding decade. In this section, we show that what growth in inequality there was has been driven by changes at the very top and the very bottom of the income distribution.

First, we can illustrate this by showing the path of the 90/10 ratio over time. The 90/10 ratio is another very commonly used measure of inequality, which quantifies the gap between the income of an individual near the top of the distribution and the income of an individual near the bottom. When we compare inequality indices, there is no reason to think that their levels will be comparable, as they are in different units. Thus comparison is assisted if we index both inequality measures back to have a common value (1.0) in a particular year (1979), as we have done in Figure 3.6. This enables us to focus only on the relative change in the paths of the two measures since that point.

The growth of the 90/10 ratio is in some respects quite similar to that of the Gini coefficient, showing rising inequality over the 1980s and some of the 1990s, and slightly falling inequality over the 2000s. However, it is interesting to note that the 90/10 ratio grew quite a bit more rapidly than the Gini in the late 1980s, fell more sharply in the early 1990s and was flatter during the later part of the 1990s and early 2000s.

These differences between the paths of the two measures become particularly relevant when we attempt to make an overall assessment of what happened to income inequality over the 1990s. Although the Gini coefficient suggests that inequality

FIGURE 3.6

The Gini coefficient and the 90/10 ratio for income

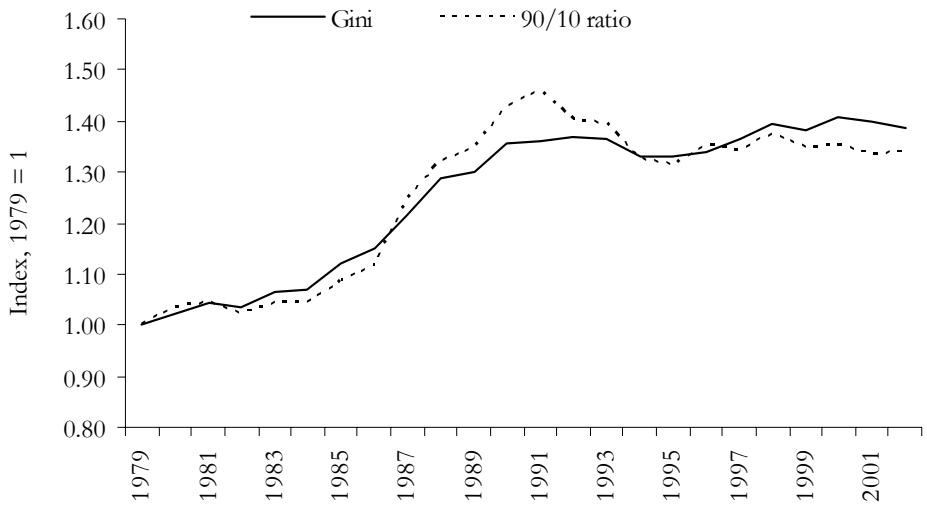

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Sources: Family Expenditure Survey, 1979 to 1993/94; Family Resources Survey, 1994/95 onwards.

FIGURE 3.7

Path of the 90/10 ratio over different decades

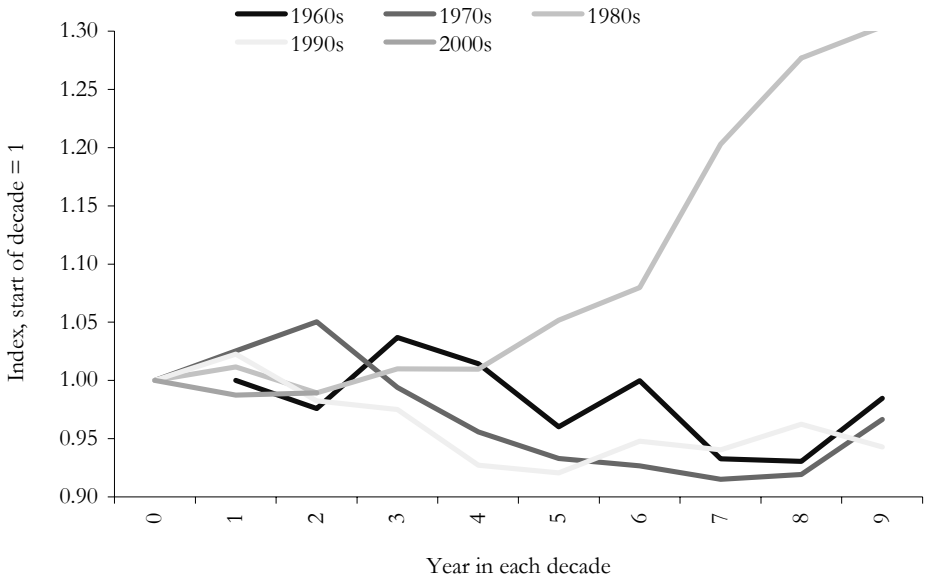

Notes: Points are drawn for calendar years up to and including 1992, and financial years thereafter. The 1960s series is indexed at 1.0 in the year 1961 as this is when our data series begins.

Sources: Family Expenditure Survey, 1961 to 1993/94; Family Resources Survey, 1994/95 onwards.

continued to rise over the course of the decade, the 90/10 ratio by contrast suggests somewhat falling inequality over the same time, ending the decade at around 95 per cent of its level at the start. As with the Gini coefficient, however, the pattern of change over the 1990s much more closely resembles the changes in earlier decades, with the 1980s standing out as the exception (see Figure 3.7).

What drives the differences between these two measures is that the Gini coefficient is a measure of inequality that takes into account the incomes of every single household at all points in the income scale, whereas the 90/10 ratio compares just two individual points — one near the top and one near the bottom — of the distribution. In fact, as we will show, the main drivers of the apparent increase in inequality over the 1990s captured by the Gini coefficient are the income changes at the very top of the income distribution and changes at the very bottom, which are not captured by the 90/10 ratio.

Before we consider changes at the very top and very bottom of the income distribution, it is worth understanding what has happened to relative incomes within the 90/10 interval. Figure 3.8 compares the paths of the 10^{th}, 25^{th}, 50^{th}, 75^{th} and 90^{th} percentiles of income, over the 1980s and the 1990s. The left-hand graph covers the 1980s and is indexed at 1980=1, whilst the right-hand graph covers the 1990s and early 2000s and is indexed at 1990=1. The differences between the two panels are quite striking. Over the 1980s, the growth in incomes was much stronger at the top end of the income distribution than in the middle, and again much stronger in the middle than at the bottom. Over the 1990s, the strongest income growth overall amongst the percentile points shown was at the 10^{th} and 25^{th} percentiles, although the income growths in different parts of the distribution are also far more similar than in the previous decade.

FIGURE 3.8

Real income at percentile points of the income distribution: the 1980s and 1990s compared

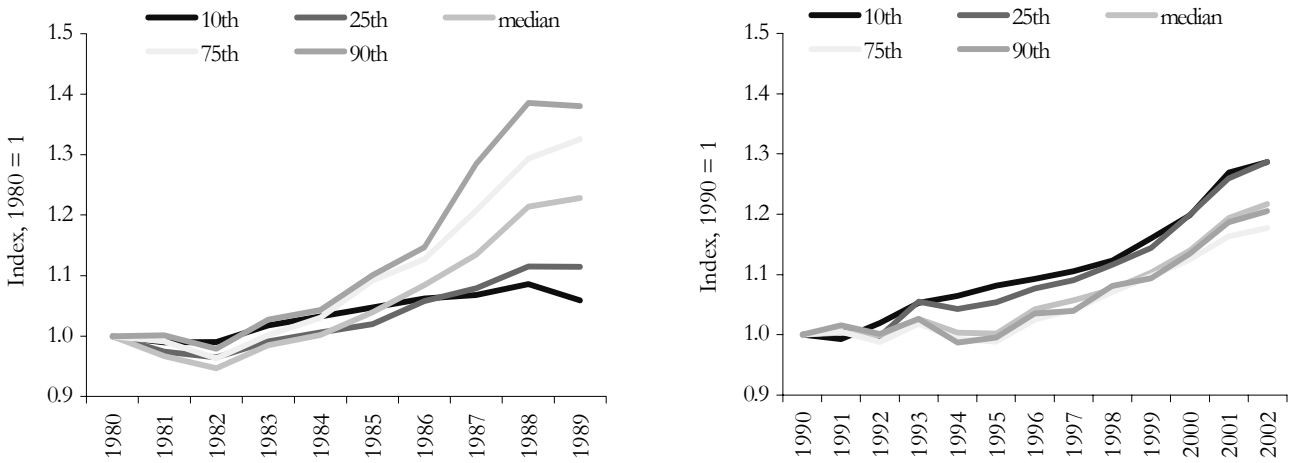

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Sources: Family Expenditure Survey, 1980 to 1993/94; Family Resources Survey, 1994/95 onwards.

Income shares

Given that it seems that it is income changes outside of the 90/10 interval that appear to explain most of the rise in inequality during the 1990s, it is especially interesting to hone in on changes at the extremes of the distribution. We can do this in a number of ways, but here we look at the *share* of total income held by different groups of the population, ranked according to their income, focusing particularly on groups at the upper and lower ends of the distribution.[6]

Looking first at the income shares of the population divided into tenths (or around 5.7 million people) according to their income: if incomes were equally distributed, then each decile group would account for 10 per cent of the total income. Figure 3.9 shows that this is far from the case. The share of the top 10 per cent — that is, all individuals in the top tenth, including the very richest — continued to grow over much of the 1990s. By 2002/03, the richest tenth of the population held about 28 per cent of the total income. By contrast, the share of the bottom tenth — that is, all individuals in the bottom tenth, including the very poorest — declined over much of the decade, being at about 2.8 per cent of total income by 2002/03. The income shares of the decile groups between these richest and poorest groups have been fairly constant over the 1990s by comparison.

We can look more closely at the top and bottom ends of the distribution by dividing the population more finely, into percentile groups, or hundredths of the population (around 570,000 people). In this case, if incomes were completely equally distributed, then each percentile group would hold precisely 1 per cent of the total income. As the left-hand panel of Figure 3.10 shows, the top 1 per cent in our sample has taken an increasing share of total income since 1979, with the sharp rise

[6]Figure 2.6 in Brewer et al. (2004) instead compared the percentage change in real incomes at different percentile points of the income distribution between 1996/97 and 2002/03. In this Report, we consider summary statistics that allow us to see changes across a number of years, and are therefore not as sensitive to choice of start and end years as that analysis.

continuing over the 1990s. By the start of the 2000s, the top 1 per cent held around 8 per cent of total income, compared with income shares of around 5.8 per cent in 1990 and 3.5 per cent in 1980. Compared with these 'very rich' household incomes, the income shares of percentile groups within the rest of the top 10 per cent have stayed relatively constant over the 1990s. The right-hand panel of Figure 3.10 shows, by contrast, what has happened to income shares within the poorest tenth of the population. Here, we can see that amongst the bottom 5 per cent in particular, income shares continued to fall over the 1990s.

FIGURE 3.9

Income shares by decile group

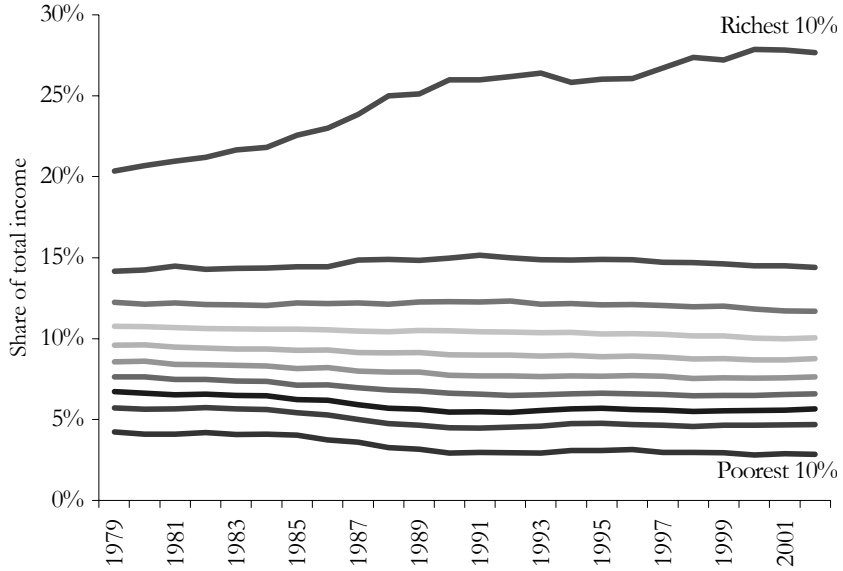

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Sources: Family Expenditure Survey, 1979 to 1993/94; Family Resources Survey, 1994/95 onwards.

FIGURE 3.10

Income shares by percentile group amongst the top and bottom 10 per cent

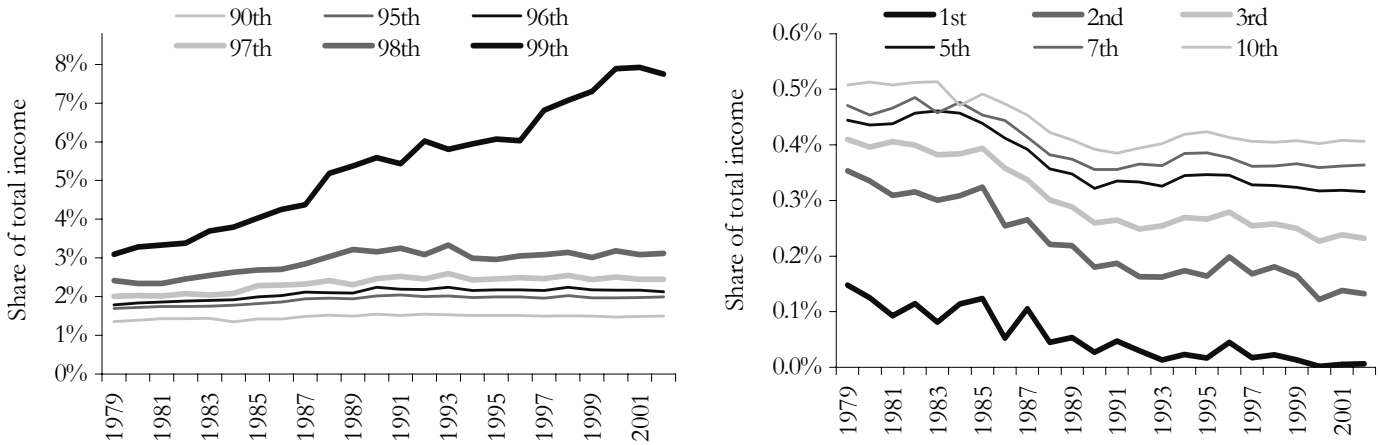

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Sources: Family Expenditure Survey, 1979 to 1993/94; Family Resources Survey, 1994/95 onwards.

BOX 3.2

Who are the 'very rich' and where does their income come from?

The top 1 per cent — around the richest half-a-million people — in Britain receive around 8 per cent of the total income. The increasing share of the top 1 per cent in total income has been a notable feature of the changes in income inequality in Britain over the last two decades, and has also been a feature of the changes in other countries, including the USA and Canada (Piketty and Saez, 2003; Saez and Veall, 2003).

According to our analysis, the top 1 per cent of individuals have incomes *after tax* of over £82,000 per year, expressed as the equivalent for a couple with no children. Almost a third of the top 1 per cent are in families with someone who is full-time self-employed (compared with less than a tenth across the population as a whole). Nearly 60 per cent of the top 1 per cent live in London and the South-East (whilst 33 per cent of the British population as a whole live in London and the South-East). Very few of the top 1 per cent are either lone parents or single pensioners; other family types are all fairly well represented at the very top.

As can be seen from Figure 3.11, the top 1 per cent receive a considerably higher proportion of their income from self-employment and from investments than do those lower down the income distribution. Earnings are also a sizeable source of income for this group, though less important than for other middle- and high-income individuals.

What explains the growing income share of the top 1 per cent? We do not have conclusive answers to this question, but possible explanations that have been put forward include changes in the norms surrounding executive pay; it is also possible that cuts in the top rates of income tax over the 1980s may have led to more rapid capital accumulation at the very top of the income scale; increasing reliance on more variable income sources, such as self-employment, may also have had a role to play. (For a fuller discussion of some of these issues, see Atkinson (2003) for the UK and Piketty and Saez (2003) for the USA.)

FIGURE 3.11

Sources of income across income decile groups and in the top 1 per cent, 2002/03

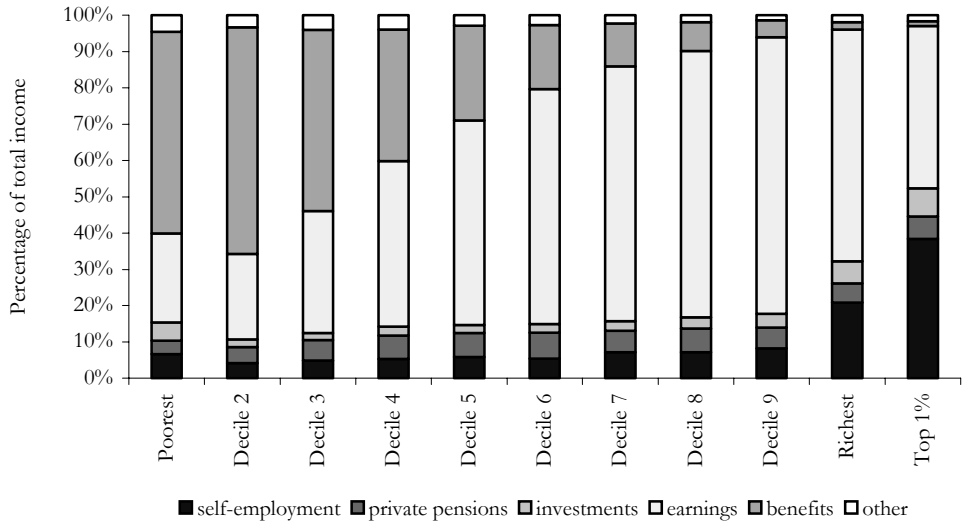

Note: Negative sources of income are set to zero for the calculation of income shares.

Source: Family Resources Survey.

BOX 3.3

Explaining the trends in income inequality: what changed over the 1990s?

Although the main purpose of this Report is a descriptive one, it is also important to set out some explanations for the trends we have described, and in particular why the major growth in income inequality appears to have halted. Here, we briefly set out a number of contributory factors.

What has happened to earnings inequality?

Earnings make up the most important source of income for households, and so changes in the distribution of earnings are very important in explaining overall patterns of income inequality. Earnings are determined both by the amount individuals are paid per hour and by the number of hours they work.

One clear picture to emerge when looking at the distribution of earnings is that earnings inequality, measured both in terms of hourly wages and in terms of weekly earnings, did not grow as rapidly over the 1990s as over the 1980s. Although men's earnings continued to become more unequal, inequality in women's earnings did not increase over the 1990s. Looking more closely within the distribution of men's earnings, it appears that the recent growth in inequality in men's earnings has been driven largely by a widening gap between the top and the middle of the earnings distribution (see Machin (2003)).[a]

What changed to make inequality in wages stop growing as fast? Much of the rapid increase in earnings inequality over the 1980s has been attributed to increasing wage returns to education, as increased demand for skilled workers outpaced supply: in part, this is due to skill-biased technical change, favouring the more highly skilled (see Card and DiNardo (2002) for a critical evaluation of the evidence on this). Over the 1990s, the evidence that we have suggests that the returns to education have stopped growing (see Sianesi (2003)), as skill-biased technical innovation has continued but the supply of high-skilled workers has also increased further.

Employment changes

A major part of the story of why inequality grew over the 1980s was the increases in both unemployment and inactivity rates, with a growing gap between the fortunes of those who had jobs and those who did not. In the 1990s, by contrast, the effect of employment and participation changes seems to have been somewhat more equalising, rather than polarising.

Unlike the recession of the early 1980s, the recession during the early 1990s and the subsequent early recovery seem to have had an equalising effect on the distribution of income (and also of expenditure: see Chapter 4). In part, this is explained by the composition of the job loss and subsequent job gain over this time (with the recession of the early 1990s sometimes referred to as the 'middle-class recession'[b]).

Employment growth in the latter part of the 1990s has also contributed to a stabilising in the distribution of income: in particular, it has halted the growth in the number of workless households, particularly amongst families with children (see Gregg and Wadsworth (2003a)). The evidence on whether more low-skilled workers have found jobs over this time is more mixed, and seems to be more geographically varied (see Gregg and Wadsworth (2003b)).

Continues overleaf.

BOX 3.3 continued

Taxes and benefits

Tax and benefit policy has also had a very important role to play, both in increasing inequality over the 1980s and in halting inequality growth over the latter part of the 1990s. Clark and Leicester (2004) show that fiscal policy changes over the 1980s — particularly reductions in top rates of income tax, increases in indirect taxes, and real and nominal freezes on a number of different benefits — all contributed to rising inequality over this time. The major redistributive policies of the Labour government since the late 1990s have had the opposite effect, slowing what would otherwise have been a much bigger growth in income inequality (see Brewer et al. (2004) and Clark and Leicester (2004)).

[a]Differences between the wages of men and women have also continued to narrow over the 1990s and early 2000s, though the impact of this on household income inequality is not clear.

[b]For example, see Wilkinson (1999).

One important question to consider is the extent to which the patterns shown in these graphs are genuine reflections of the changes in the incomes of the very richest and poorest, or whether they could simply be the result of measurement error or sampling variation in the data. As we noted in Chapter 2, our income series contains an adjustment to the incomes of the very rich, based on data from the Survey of Personal Incomes (SPI), which is designed to improve the robustness of the series at the very top end. The incomes of many of those in the top 1 per cent of our data have been 'SPI-adjusted' in this way. Correcting incomes at the very bottom for misreporting or for over- or under-sampling is much more problematic and has not been attempted here.

Our findings on the changes at the very top are also consistent with data from other, potentially more reliable, sources. For example, data from the Inland Revenue based on individual-level tax returns rather than household survey data show that the 99th percentile of individual incomes amongst taxpayers rose much faster than incomes at other points in the distribution of taxpayer incomes over the 1990s. These data suggest a similar pattern for changes in inequality to the one presented here.

Taking a wider perspective, recent work by Atkinson (2003) looking at the changes in top incomes over the whole of the twentieth century in the UK based on super-tax, surtax and income tax returns data also confirms that rising inequality at the very top of the income distribution has been an important phenomenon which has continued through the 1990s. This work shows that the share of the top 1 per cent — and, even higher than this, the top 0.5 per cent — in total income rose sharply in the second half of the 1990s.

3.4 Historical Context: Are We More Unequal than Ever?

The work by Professor Atkinson cited above enables us to learn about income changes at the very top of the income distribution over nearly the whole of the twentieth century. Although we can learn much less from this work about changes at the middle and bottom of the income distribution — and hence measures of inequality that are based on all points across the income distribution — the information it reveals on the share of total income held by the very top puts the changes in inequality set out in this chapter into very interesting context.

FIGURE 3.12

Top income shares

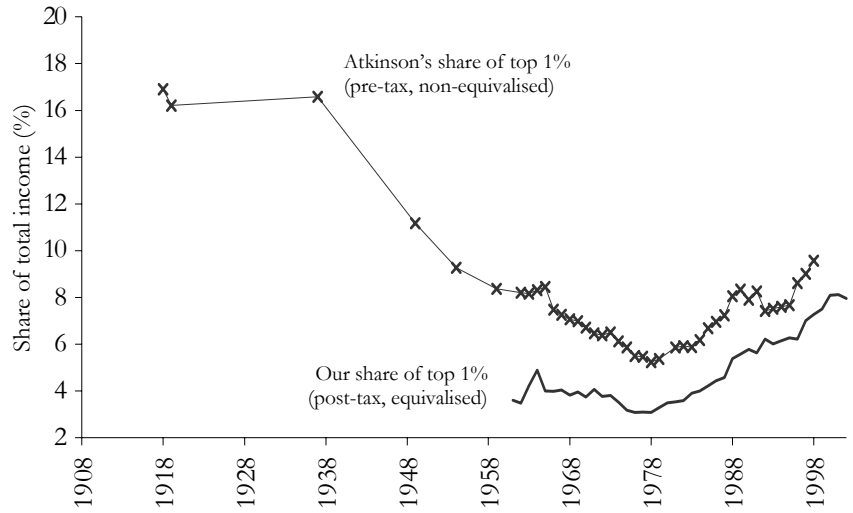

Sources: 'Atkinson's share of top 1%' comes from Atkinson (2003). 'Our share of top 1%' is based on the Family Expenditure Survey up to and including 1993/94 and on the Family Resources Survey from 1994/95 onwards.

Figure 3.12 compares the results set out in Atkinson (2003) and our top income shares shown in Figure 3.10. Atkinson's series is based on pre-tax, unequivalised incomes, which are more unequally distributed than the post-tax equivalised incomes we consider in this Report, with the top 1 per cent taking a bigger income share. Although the levels are different, the trends in top income shares on these measures since the early 1960s have been quite similar.[7] What Atkinson's series adds to our knowledge is that the rise in the share of the incomes of the 'very rich' over the 1980s and 1990s in fact served to reverse a protracted period of equalisation, returning the income shares of the richest back to levels not seen since the 1950s. But, interestingly, it also shows that the level of inequality captured by the income shares of the very rich still remains considerably below the levels we saw before the Second World War, and is of a different magnitude from that of the large income shares of the very rich seen prior to the First World War.

Atkinson's work also shows that the nature of the changes in the income distribution over the 1990s described in this chapter echoes changes in previous eras. In particular, the combination of a fairly stable distribution of income together with localised changes at the very top was also a feature of the changes during much of the inter-war period, as well as after the Second World War. Atkinson's work also highlights the fact that rather than trending steadily either up or down over much of the last century, a typical pattern has tended more to be one of fluctuations and plateaux around a broader trend. Our investigation of how income inequality changed over the 1990s also suggests a return to this sort of pattern.

3.5 Conclusions

The rapid income inequality growth experienced in Britain over the 1980s clearly dominates any assessment of the changes in income inequality over the last four decades. In comparison, the changes in inequality that took place over the 1960s and

[7]This similarity in trends may arise because the series are based in part on the same data sources: Atkinson's series over this period is based on data from the SPI, whilst our series also uses income from the SPI to adjust incomes at the very top of the scale.

1970s are much smaller, whilst the changes over the 1990s and early part of the 2000s are also comparatively small.

Our judgement about whether inequality rose or fell during the 1990s depends on the exact measure of inequality chosen. Measures of inequality that focus on differences in incomes across the whole population, including the very richest and the very poorest, tend to suggest a widening in the income distribution, particularly over the second half of the 1990s. Such measures include the Gini coefficient and the income shares presented in this chapter.

Changes in inequality measured in this way have been driven by rapidly rising incomes at the top of the income scale, particularly amongst the top half-a-million people, as well as by incomes at the very bottom of the scale that have not kept pace with general income growth across the population.

By contrast, measures of inequality that do not consider the very top and bottom ends of the income scale, such as the 90/10 ratio, suggest a pattern of falling income inequality in the first part of the 1990s, and then stable inequality over the later part of the 1990s.

Since 2000, it appears that income inequality has stabilised, though any falls in the headline measures of inequality are not large enough to be statistically significant. Irrespective of the measure of inequality chosen, income inequality remains near a 40-year high.

These trends can be placed in the context of movements in income inequality over the whole of the twentieth century, which suggests that the rise in inequality over the 1980s returns the level of inequality back to that last seen in Britain in the 1950s. The pattern of localised changes in the structure of incomes against a backdrop of relative stability amongst the bulk of the population also echoes changes taking place in earlier decades.

CHAPTER 4
Expenditure Inequality

We found in Chapter 3 that although inequality in incomes increased dramatically over the 1980s, the growth over the 1990s and early 2000s was much slower, and was driven largely by changes at the top and bottom of the income distribution. As we discussed in Chapter 2, measuring living standards by household expenditure — a useful proxy for consumption — gives us an alternative view of people's living standards, and one that is more likely to reflect the longer-run, or more 'permanent', circumstances of individuals. Previous work[8] has shown that in the 1980s, there was a steep rise in expenditure inequality, though not as steep as the rise in income inequality. In this chapter, we look at what has happened to inequality in expenditure over time, focusing particularly on the 1990s.

4.1 The Distribution of Expenditure

A picture of the distribution of expenditure

We saw in Chapter 3 that the distribution of income is skewed. Figure 4.1 shows a similar picture to that found in Figure 3.1: it shows the number of people in 1999/2000[9] living in households with different equivalised expenditure levels, grouped into £10 bands. Expenditure is defined as total household expenditure, including that on durables and housing, and is uprated to 2002/03 prices.[10] The distribution shown on the graph has been truncated at £1,000 per week. As for income, we find that the distribution of expenditure is also skewed, with 66 per cent of individuals living in households with equivalised expenditure that is below the mean of £330 per week.

Figure 4.2 shows the paths of the real values of expenditure from 1974 to 1999/2000 at the mean, the median and the 10th and 90th percentile points. What emerges from this picture is a very similar story to what we found for income in Chapter 3: while expenditure at the top of the distribution (the 90th percentile) has risen by around 40 per cent since 1974, expenditure growth at the bottom of the distribution (the 10th percentile) has been much less, at around 17 per cent.

However, while it is true that growth at the 90th percentile point in both distributions has outstripped growth at the 10th percentile point, growth at the 90th percentile point in the income distribution over this period was 73 per cent, which is much greater than the growth rate of 40 per cent at the 90th percentile point in the expenditure distribution. Over the same period, growth at the 10th percentile point of the income distribution (27 per cent) also exceeded growth at the 10th percentile point of the expenditure distribution (17 per cent), but not to the same extent as at the 90th percentile point. This suggests that while inequality in the expenditure distribution has increased over time, it has done so to a lesser extent than the increase in income inequality that we found in Chapter 3.

---

[8] See, for example, Goodman, Johnson and Webb (1997).

[9] This is the latest year for which we have constructed a consistent measure of housing expenditure.

[10] See Chapter 2 for more details on this measure.

FIGURE 4.1

The expenditure distribution in 1999/2000

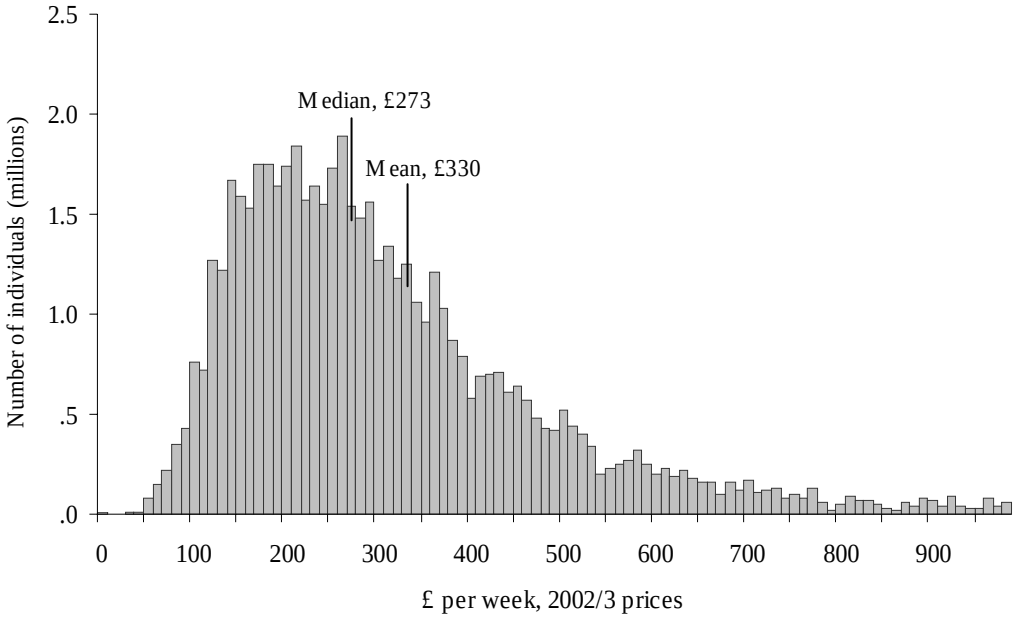

Source: Authors' calculations using Family Expenditure Survey, 1999/2000.

FIGURE 4.2

The real value of expenditure

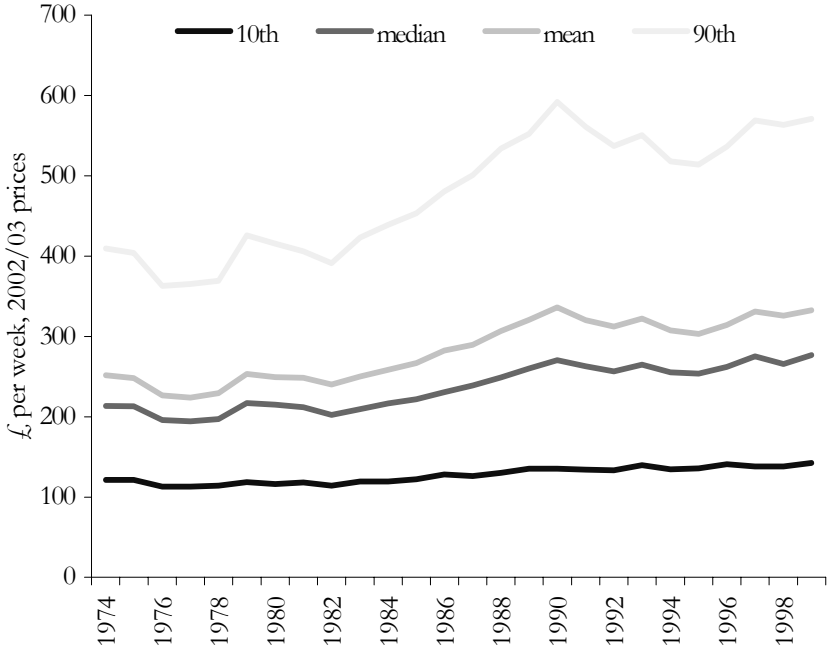

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

Expenditure inequality over time

Figure 4.3 shows the evolution of the Gini coefficient for expenditure in Great Britain between 1974 and 1999/2000.[11] The Gini remained below 0.3 from 1974 to 1984, but with different trends at the beginning and end of that period. From 1974 to 1977, the Gini was on a downward trend, falling by 2 percentage points in that period. However, this decline was eliminated by a sharp rise of more than 2 percentage points between 1977 and 1979. The Gini then remained quite stable at around 0.29 until 1984. Between 1984 and 1990, a quite rapid increase took place, seeing the Gini rise above 0.3 in 1985 and continue to grow, hitting a peak of 0.34 in 1990, almost a quarter above its 1977 trough.

Between 1990 and 1995, however, there was a sustained reduction in the expenditure Gini coefficient which saw it fall back to a low of 0.30 in 1995/96, returning inequality to its mid-1980s level. This decline was followed by three years of further increases, to 0.32 in 1998 — still some way below the 1990 level, however. These rises were checked in 1999/2000, with the Gini coefficient falling back just over half a percentage point, leaving it 2.5 percentage points (or 7 per cent) below its 1990 level. The decline in expenditure inequality over the 1990s did little to reverse the much bigger rises seen in the 1980s. At the end of the 1990s, expenditure inequality was still at the level experienced in 1987 and 3 percentage points (10 per cent) above the level at the beginning of the 1980s.

Over the period since 1974 as a whole, the expenditure Gini increased but only fairly marginally — in the order of 8 per cent or 2.5 percentage points. This disguises perhaps four distinct phases of movement — a decline of almost 2 percentage points in the mid-1970s, a sustained rise totalling over 7 percentage points over the late 1970s and 1980s, a decline of around 3.5 percentage points in the early 1990s and a rise of just over 1 percentage point in the late 1990s. It remains to be seen whether the decline from 1998 to 1999 marked the beginning of a sustained fall once again or was merely a one-off.

FIGURE 4.3

The expenditure Gini coefficient

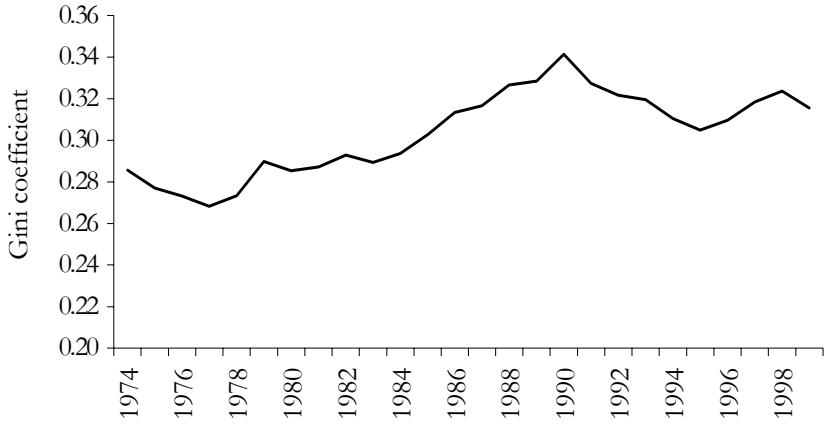

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

[11]The Appendix reports confidence intervals around the Gini and tests whether between-year changes in inequality are statistically significant.

FIGURE 4.4

The expenditure Gini coefficient and the 90/10 ratio

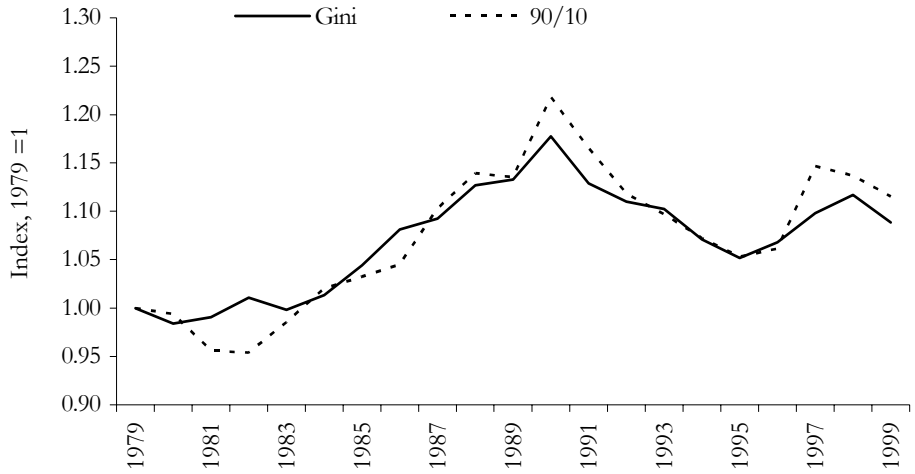

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

Figure 4.4 compares the expenditure Gini coefficient and the 90/10 ratio. As in Chapter 3, we focus on the period since 1979 and base both series at 1 in that year. Despite the fact that the 90/10 ratio does not take into account expenditure other than at the 10^{th} and 90^{th} percentile points, the picture of expenditure inequality that emerges does not differ too dramatically from what emerges with the Gini coefficient which takes into account the *whole* distribution. For example, although the growth in inequality between 1979 and 1990 is greater on the 90/10 ratio measure of inequality, both measures fall between 1990 and 1995. This contrasts with measures of income inequality, where the 90/10 ratio fell but the Gini coefficient was broadly flat over the first half of the 1990s (see Figure 3.6).

It is interesting to note that Blow, Leicester and Oldfield (2004) find that the 90/10 ratio for a measure of expenditure which excludes housing follows a slightly different trend over the 1990s from the trend followed by the measure of expenditure used in this chapter (which includes housing). Although the 90/10 ratio for expenditure excluding housing costs moves broadly in the same direction over the first and second halves of the 1990s, this measure of inequality is higher at the end of the decade than it was at the beginning. This differs from the path for the 90/10 ratio shown in Figure 4.4 (for expenditure including housing), which is lower at the end of the 1990s than at the beginning. The reasons why we include housing costs in our measure of expenditure were discussed in Chapter 2, as were issues regarding the measurement of housing consumption.

Figure 4.5 directly compares the Gini coefficients for expenditure and income. The picture can broadly be divided into three periods. During the 1980s, the growth in income inequality far outstripped the growth in expenditure inequality. Nevertheless, both measures of inequality were generally moving in the same direction. This contrasts with the first half of the 1990s (1990–1995/96), when income inequality was broadly flat but expenditure inequality declined by around 11 per cent (or 4 percentage points). Since 1995/96, both income and expenditure inequality have moved in parallel, rising by around 4 per cent (or 1 percentage point). So, in

FIGURE 4.5

The income and expenditure Gini coefficients

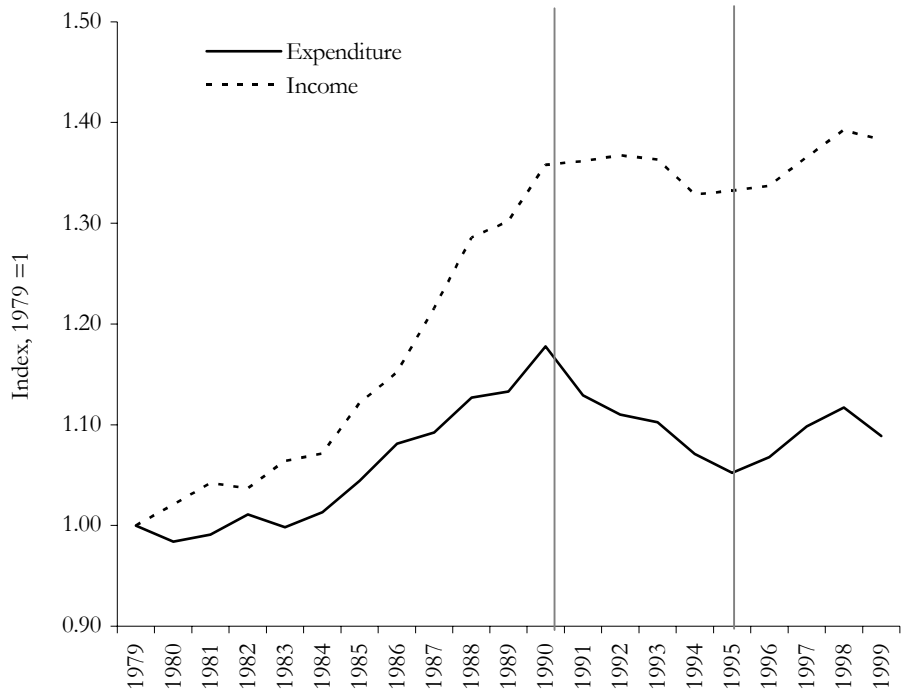

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

summary, whichever measure of material well-being we use to measure inequality, it is clear that inequality increased dramatically over the 1980s. However, the paths of inequality of income and expenditure look rather different from each other in the 1990s.

The conclusion that we reach about the picture of inequality today depends very much on whether we use income or expenditure as our measure of inequality. Looking at expenditure, we would conclude that, despite the rise in inequality that has occurred since 1995/96, inequality is still lower than was seen at the peak in 1990. However, turning to income inequality, we would conclude that we are more unequal than at any time during the last four decades.

Changes in percentile points of income and expenditure

An observed increase in inequality can be driven by changes in incomes or expenditures at different parts of the distribution. For example, an observed increase in inequality can be driven by an increase in incomes or expenditures at the top of the distribution while those in the rest of the distribution remain flat. However, a similar increase in inequality could also be driven by a fall in incomes or expenditures at the very bottom of the distribution while the remainder of the distribution remains flat. These two situations might have different implications for policies designed to reduce inequality or poverty.

Figures 4.6 and 4.7 show income and expenditure at percentile points from different parts of the income and expenditure distributions over the 1980s and 1990s. Taking the 1980s first (Figure 4.6): this was a period when both income and

FIGURE 4.6

Real income and expenditure at percentile points of their distributions: the 1980s

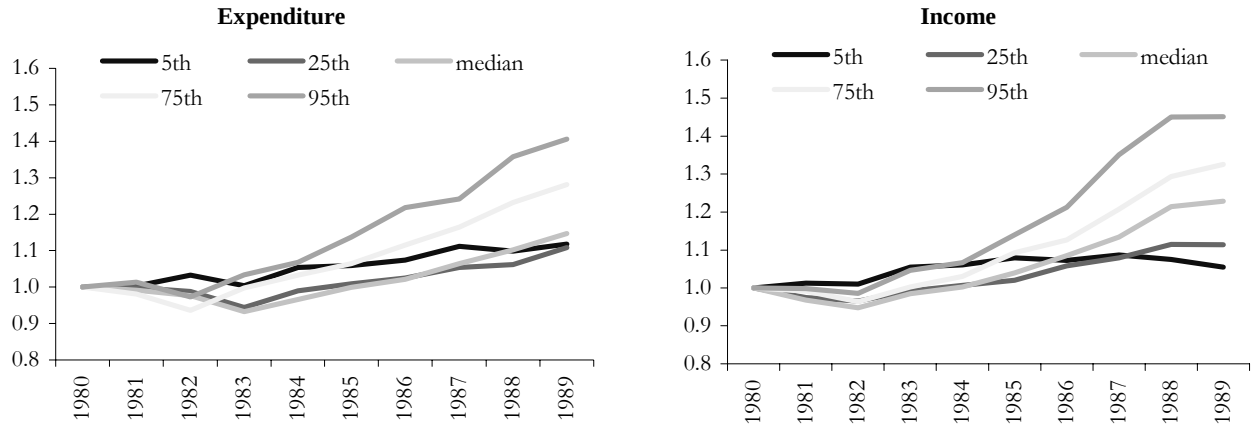

Note: Points are drawn for calendar years.

Source: Authors' calculations using the Family Expenditure Survey.

FIGURE 4.7

Real income and expenditure at percentile points of their distributions: the 1990s

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

expenditure inequality rose, but income inequality rose far more than expenditure inequality. After 1982, there was a general widening of both the income and expenditure distributions but the extent of this widening was less in the expenditure distribution.

Figure 4.7 shows income and expenditure at percentile points of the expenditure and income distribution in the 1990s. We know that inequality in expenditure fell in the first half of the 1990s whereas inequality in income was broadly flat. Figure 4.7 highlights exactly which parts of the distribution were driving these changes. Expenditure at the bottom of the distribution grew in the first half of the 1990s while expenditure at the top fell. This contrasts with the income distribution, where incomes across the board did not diverge very much. In the second half of the 1990s, the situation was quite different. Expenditure in the upper part of the distribution grew at a faster rate than expenditure in the lower part. This was also true of income, leading

to the very similar increases in inequality measured by income and by expenditure seen in Figure 4.5.

4.2 Explanations

The divergence in income and expenditure inequality trends that occurred over the 1980s is well documented. This Report has shown that the divergence continued in the first half of the 1990s, to such an extent that the two measures moved in opposite directions, but they have moved in parallel since 1995. One important explanation that has been put forward to explain the divergence in expenditure and income inequality over the 1980s is that income became more transitory, meaning that it became subject to more short-term volatility.[12] Chapter 2 discussed the differences between consumption and income and outlined the argument that consumption may be a better measure of household welfare because it reflects 'permanent' differences in resources (or lifetime resources) rather than current resources. If incomes become more volatile but people are able to save and borrow in order to smooth their consumption over their lifetime, this would be reflected in higher income inequality but constant consumption inequality (other things being equal). Using data covering the period 1968–92, Blundell and Preston (1998) find that towards the end of this period, there was a strong growth in transitory (i.e. non-permanent) income inequality. It could be the case that this increase in income volatility continued into the first half of the 1990s, which would explain the continued divergence of income and expenditure inequality. However, what is interesting is that since 1995, the two measures of inequality have moved in parallel. This would imply an end to the growth in income volatility that has so far been used to explain the divergence before this time. If this is the case, the question as to why we have seen an end to the growth is an interesting one and should be the focus of future research. A further issue for future research would be to examine why we saw a fall in 'permanent' income inequality (as reflected in the reduction in inequality in expenditure) in the first half of the 1990s.

Increased income volatility is not the only explanation for a divergence in income and expenditure inequality. The extent to which expenditure measures 'permanent' income depends on the ability and willingness of households to smooth their consumption over time. Two important mechanisms by which households can finance consumption in times of low income are to run down savings and to borrow. One explanation for the divergence of expenditure and income inequality could be that households' ability to borrow has increased over time. There is very little evidence on whether or not there has been a relaxation of credit constraints allowing different types of households to smooth their consumption more easily, not least because of the lack of long time series of micro-data that record credit and debt in the UK.[13,14] However, since 1988, the FES allows us to look at the percentage of spending that is obtained by means of some kind of credit.[15]

[12]See Blundell and Preston (1998).

[13]Some recent studies have looked at access to credit arrangements (for example, Bridges and Disney (2004)) but there is little evidence on how this has changed over time.

[14]Even if such data did exist, it would be hard to identify whether households are in fact credit constrained, because those who do not have credit or debt are not necessarily credit constrained as they may choose not to use credit arrangements. Even data that record applications for credit that are refused do not allow us to identify all households that are credit constrained because those who are most constrained may not apply for credit.

[15]The types of credit that are recorded are hire-purchase agreement, formal loan, mail-order credit and credit card. Note also that charge card, while not strictly a form of credit, is also included.

FIGURE 4.8

Percentage of spending that is obtained by means of credit

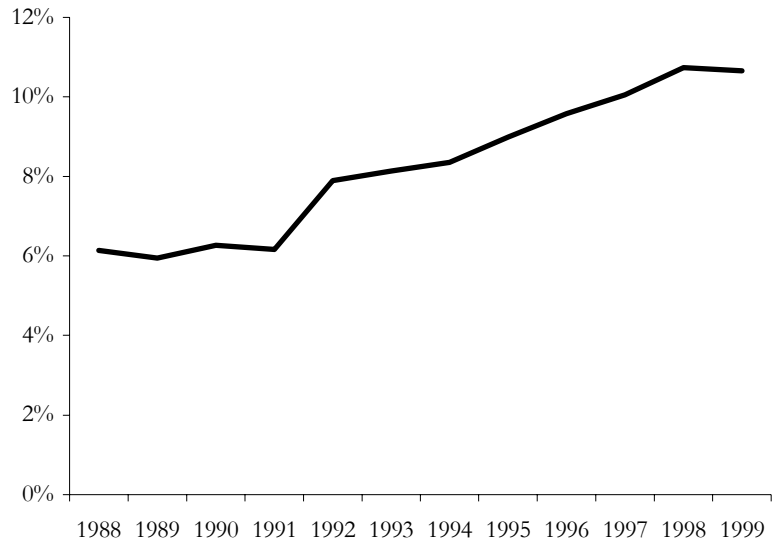

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

Figure 4.8 shows that since 1988, and particularly between 1992 and 1999, there has been an increase in the share of spending bought using credit.[16] These data are by no means perfect for examining whether household borrowing has increased, because transactions carried out via a credit card may be paid off immediately out of current income. Nevertheless, Figure 4.8 may be an indication that households are borrowing more. This could tell us that households have found credit easier to obtain or it could simply reflect the fact that if incomes have become more volatile, households will have a greater need to borrow in order to smooth consumption.

Another explanation for divergence in trends in income and expenditure inequality relates to measurement error. Errors in the reporting or recording of any measure of resources will tend to increase measured inequality but do not reflect any true increase in inequality of well-being. Measurement error can occur for a number of reasons. One is simple misreporting of resources. However, if this were to explain the divergence in income and expenditure inequality, misreporting of incomes would have to have risen over time relative to misreporting of expenditure. There is no reason to believe that this has happened, so while measurement error of this kind cannot be discounted altogether, it is unlikely to account fully for the divergence between income and expenditure inequality.

In addition to measurement error caused by misreporting of income or expenditure, there is another issue relating to measurement that could potentially explain, at least in part, the divergence between income and expenditure inequality. As discussed in Chapter 2, expenditure is a proxy for consumption but there are important differences between the two. One of these is the way in which expenditure on durables is measured. In the FES, information about spending on durables is collected in two

[16]Although in the graph there appears to be a discontinuity in the series between 1991 and 1992, there are no details of any change in the way in which spending on credit was recorded between these two years in the documentation that accompanies the FES. Even if there is a discontinuity in the series between these two years, the overall conclusion — that the share of credit in total spending has increased over the period — does not change.

alternative ways. Data on some items (such as cars) are collected using a method known as retrospective recall. This involves asking respondents to recall how much they spent on a particular durable during a specified time period (three months or a year). Data on other items are collected using the normal diary method. However, neither of these methods measures the true object of interest, which is the flow of services that the durable provides. Both methods, but particularly the diary method, mean that spending on durables is recorded in a very 'lumpy' way: those who did purchase a durable during the two-week diary period (or three-month recall period) will have a level of spending that is higher than the true flow of services provided by the durable during that period, while those who did not purchase the durable will have zero expenditure. In practice, anyone who owns a certain durable will have received some flow of services from that durable during the time period.

The lumpy nature of recorded expenditure on durables is a source of variability in the data which does not reflect true inequality of welfare. However, if this source of variation does not change over time, it should not affect the conclusions we reach about the *changes* in expenditure inequality over time. While we have taken care to ensure that our measure of spending is consistent over time,[17] changes in spending patterns could mean that this source of variability in the data increases or decreases with time. For example, if, over time, households on average devoted a greater share of their budget to durables at the expense of a good that exhibits lower variability, this would lead to an increase in inequality over time, other things being equal. However, simply looking at the share of spending on durables is not enough to tell us the contribution that durable spending has made to inequality over time because there could have been changes in the amount of variability in spending on durables across different people. For example, if the frequency of durable purchases increased, this would lead to a lower number of zero expenditures in the data, which (other things being equal) would reduce this source of variability and lead to an observed decrease in expenditure inequality. The effect on inequality would need to take account of both of these factors. In order to shed some light on this issue, we examine in the next section the contribution that expenditure on durables and other types of goods and services makes to inequality and how this has changed over time.

4.3 Components of Spending and Their Contribution to Inequality

Total expenditure is the sum of spending on many different items, and household spending patterns have changed over time. For example, the average share of spending devoted to food declined from 24.6 per cent to 14.8 per cent between 1975 and 1999, and the average share devoted to holidays increased from 2.2 per cent to 5.1 per cent over the same period (see Blow, Leicester and Oldfield (2004)). In this section, we examine the contribution of broad types of goods to expenditure inequality since 1979 to see how changing spending patterns may help explain the changes in expenditure inequality over time. We pay particular attention to the contribution to inequality that spending on durables has made, in order to determine whether this is a likely explanation for the divergence between expenditure and income inequality.

[17] In particular, the collection method for expenditure on some durable items in the measure of total spending in the FES changed from being diary-based to retrospective-recall-based in the late 1980s. However, both diary and retrospective recall versions exist throughout, which enables us to correct for this inconsistency; diary expenditure on these items is used throughout.

FIGURE 4.9

The Gini coefficient and half the squared coefficient of variation

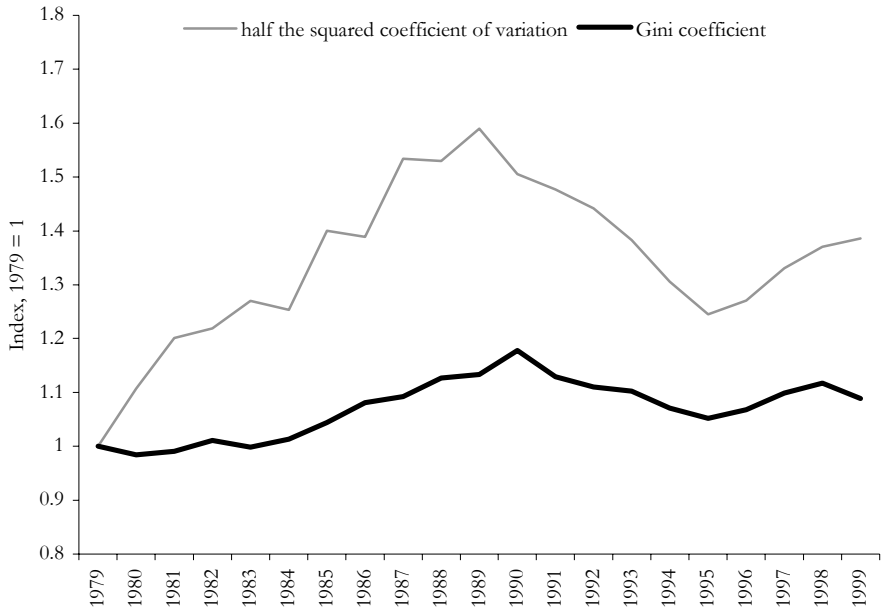

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

In order to examine the contribution that different goods have made to inequality, we need to decompose the measure of inequality. As discussed in Chapter 2, not all inequality measures decompose neatly, and, unfortunately, the Gini coefficient is one of them. For this reason, we will rely on an alternative measure, half the squared coefficient of variation (COV from now on). We start by comparing the COV[18] to the Gini coefficient in Figure 4.9 (both series are indexed at 1 in 1979). The obvious point to note from this graph is that the magnitude of growth and subsequent decline in the COV far outweighs that measured by the Gini coefficient. This is partly because, unlike the Gini coefficient, the COV is not bounded between 0 and 1, so changes in this measure are also not bounded. It is also because the COV is more sensitive to outliers than the Gini coefficient. However, for the purposes of this section, the fact that the direction of change is broadly the same for both measures means that we can apply any results that we find in this section more generally.

In looking at how various goods and services have contributed to overall inequality since 1979, we consider six very broad types of goods and services. These are: basics (food, domestic fuel and light, and clothing), other non-durable goods, leisure services (holidays, entertainment and catering), non-leisure services, durable goods, and housing.

The contribution of different goods and services to inequality will depend on the inequality of expenditure within that particular good, the share of that good in total expenditure, and its correlation with total expenditure. Table 4.1 shows how these factors vary across the six broad groups of goods in 1980, 1990 and 1999/2000. The first column shows the amount of within-source inequality, which measures how

[18]Because half the squared coefficient of variation is sensitive to outliers, this measure of inequality is very noisy across time. For this reason, in all measures of inequality in this section that use half the squared coefficient of variation, we use a three-year moving average of the series.

unequally spending within each good is distributed. In 1980, the good that displays the highest amount of within-source inequality is non-leisure services. However, in 1990 and 1999/2000, it is durables that display the highest amount of within-source inequality. The second column shows the factor share, which is the average proportion of the household budget allocated to a particular good. Of the groups of goods that we define here, basics and other non-durables take up the largest share of household budgets, although the proportion spent on basics declined between 1980 and 1999/2000. The third column shows the relationship between spending on each type of good and total spending. This can be a number between –1 and +1, where a negative number tells us that low-spending households tend to spend a higher amount on that good than high-spending households and a positive number tells us that low-spending households tend to spend a lower amount on that good than high-spending households. The overall contribution of each good to inequality (the last column) is the product of the other three columns multiplied by the square root of the COV.[19] Because all six types of goods are positively correlated with total spending (meaning that expenditure-rich households spend more than the expenditure-poor on each good), they all contribute positively to inequality.

Figure 4.10 shows the amount that each good contributes to total inequality and Figure 4.11 shows that contribution as a percentage of the total. In all the years except 1984 and 1985, the good that contributed the most to inequality was durables. Since the mid-1980s, leisure services have also been a major contributor, and the extent of this contribution increased from around 16 per cent to 28 per cent over the period from 1979 to 1999/2000. By the end of the period, durables and leisure services contributed almost equal amounts to inequality. Although at the beginning of the period, both non-leisure services and other non-durables were amongst the highest contributors, by the end of the period, their contributions had fallen somewhat, by around 5 percentage points. Basics and housing contribute relatively little to

TABLE 4.1

Detailed decomposition of expenditure inequality by broad components of spending

	Within-source inequality	*Factor share*	*Correlation with total spending*	*Contribution to inequality*
1980				
Basics	0.337	0.337	0.487	*0.025*
Other non-durables	0.592	0.237	0.623	*0.039*
Leisure services	1.346	0.110	0.504	*0.033*
Non-leisure services	1.647	0.110	0.552	*0.046*
Durables	1.542	0.124	0.595	*0.051*
Housing	0.676	0.083	0.292	*0.007*
1990				
Basics	0.378	0.253	0.469	*0.023*
Other non-durables	0.552	0.222	0.591	*0.038*
Leisure services	1.548	0.133	0.573	*0.062*
Non-leisure services	1.299	0.117	0.527	*0.042*
Durables	1.981	0.126	0.617	*0.081*
Housing	0.795	0.149	0.445	*0.028*
1999/2000				
Basics	0.431	0.225	0.512	*0.025*
Other non-durables	0.555	0.232	0.597	*0.038*
Leisure services	1.357	0.159	0.642	*0.070*
Non-leisure services	1.051	0.124	0.543	*0.036*
Durables	1.583	0.142	0.645	*0.073*
Housing	0.743	0.118	0.229	*0.010*

Source: Authors' calculations using the Family Expenditure Survey.

[19]See Goodman, Johnson and Webb (1997) for more details of this decomposition.

FIGURE 4.10

Contribution of components of spending to overall inequality

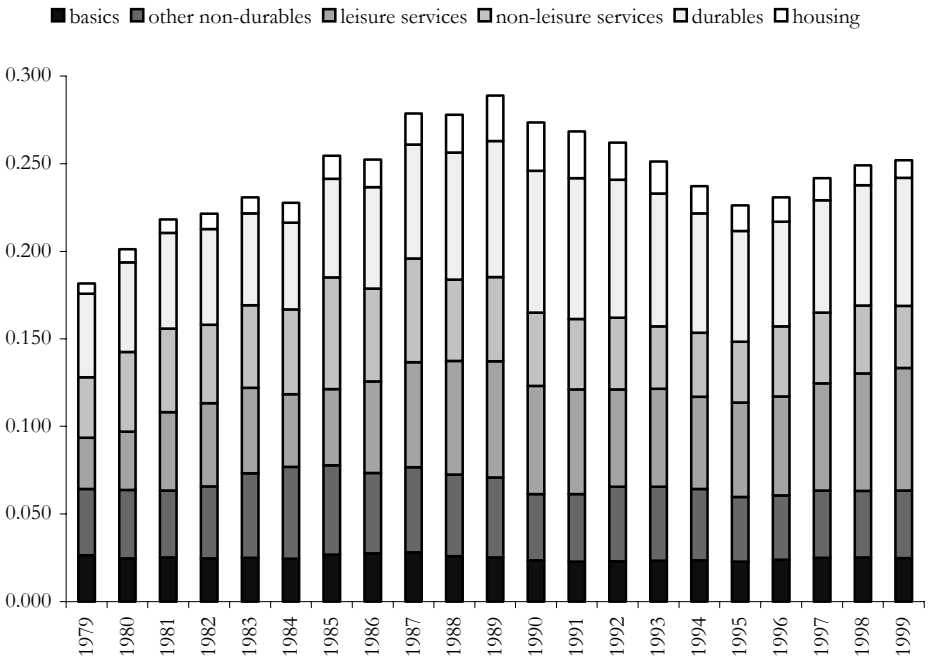

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

FIGURE 4.11

Percentage contributions of components of spending to overall inequality

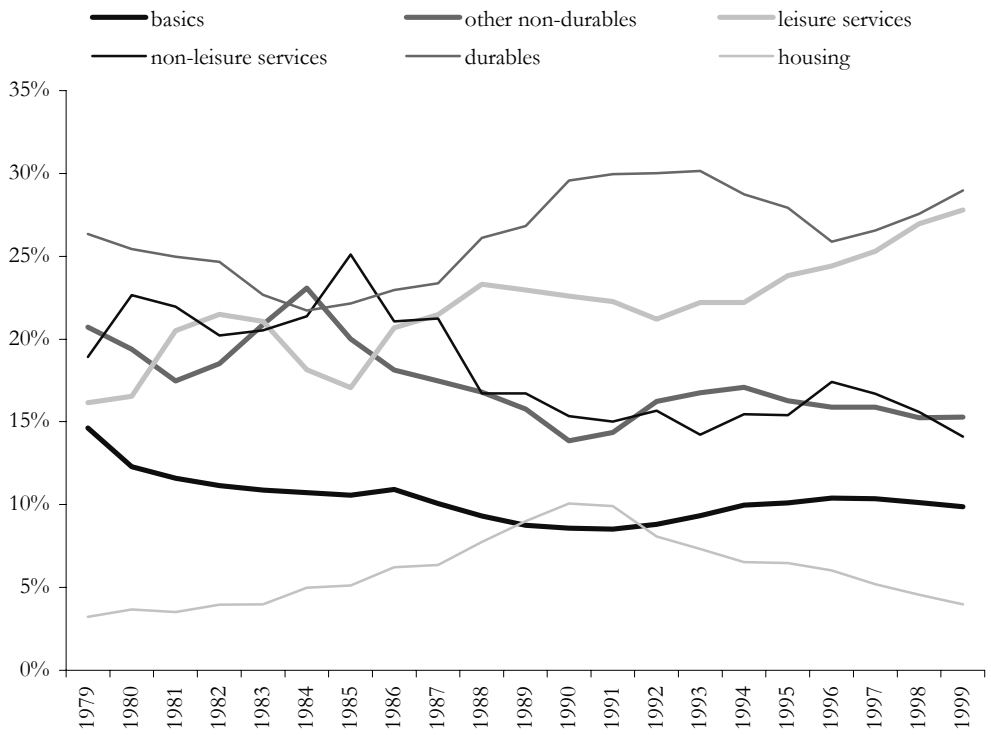

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

inequality. Looking at Table 4.1, this is because of a combination of the low within-source inequality that these goods exhibit and a lower correlation with total spending compared with the other goods we have defined. However, although the amount that housing contributes to overall inequality is small, some large proportional changes in its contribution were seen over the period. Over the 1980s, the percentage contribution increased from around 3 per cent to over 10 per cent in 1990. Over the 1990s, the amount that housing contributed to inequality fell — almost back to the level seen at the beginning of the 1980s.

In the previous section, we discussed how changing patterns of spending on durable items might explain some of the divergence between income and expenditure inequality. This argument relies on a reduction in the extent to which durables contributed to inequality over the period in which the divergence occurred (the 1980s and the first half of the 1990s). Although there was a small reduction in the extent to which durables contributed to inequality over the first half of the 1990s, this is not enough, on its own, to explain fully the divergence in income and expenditure inequality.

4.4 Conclusions

This chapter has highlighted the divergence in income and expenditure inequality in the 1980s and found that this divergence continued into the first half of the 1990s. In the first half of the 1990s, income inequality was fairly constant whereas there was a general compression of the expenditure distribution, with expenditures in the top half of the distribution growing at a slower rate than those in the lower half.

Since 1995/96 (until 1999/2000), the two measures of inequality have moved in parallel, with expenditure and income in different parts of the distributions growing at similar rates. Our latest expenditure data are for 1999/2000, whereas the income data continue up to 2002/03. It will be interesting to see whether the two measures continue to move in parallel between these years or whether they once again diverge.

We discussed some explanations for the divergence. One explanation — that income has become more volatile, or transitory — has been studied in some detail in previous research for the 1980s, but more work is needed to see if this explanation continues to hold into the mid-1990s. Similarly, future research might also focus on whether and why there was a reduction in ‘permanent’ income inequality in the first half of the 1990s and why income and expenditure inequality have moved in parallel since then.

CHAPTER 5
Conclusions

This Report has set out what has happened to income and expenditure inequality in the 1990s and early 2000s, comparing the changes with those in previous decades. We have shown that the 1990s were quite different from the 1980s in terms of both income and expenditure inequality.

Both income and expenditure inequality returned to a fluctuating path in the 1990s, in contrast to the sustained increases in inequality that took place over the 1980s on both measures of material well-being. Despite these fluctuations, although expenditure inequality has fallen slightly from its peak in 1990, income inequality remains near a 40-year high.

Although the fall in expenditure inequality over the early part of the 1990s suggests that 'permanent' income differences may have narrowed over this time, inequality in both income and expenditure continued to rise in the second half of the decade. In the early 2000s, it appears that inequality has fallen slightly again, but not yet by an amount that is statistically significant. The fact that there has been no big reversal in income inequality, despite very large redistribution by the government, suggests that the era of high inequality ushered in during the 1980s may be here for some time to come.

APPENDIX
Sampling Error and Changes in Inequality over Time

All our results are based on samples of the population and we are using these samples to make inferences about the entire population. Because of this, we have to be careful in interpreting results because any change in measures of inequality, particularly from one year to the next, could, at least in part, be due to sampling variation. We can see how precise our estimates are by attaching confidence intervals to the results. A 95 per cent confidence interval implies that there is a 95 per cent probability that the true population estimate lies within this interval. Figures A.1 and A.2 show 95 per cent confidence intervals, which have been analytically derived, around the income Gini and the expenditure Gini coefficients.

As well as placing bounds around the point estimate of the Gini coefficient, we can also test whether any change in the Gini coefficient between any two years is statistically significantly different from zero. The significance of the changes in the Gini coefficient between pairs of years is shown in Tables A.1 and A.2. A '+' indicates that there was an increase in inequality between year *t* (the year displayed in the rows) and year *s* (the year displayed in the columns) that is unlikely to be just due to sampling variation. A '–' indicates that there was a decrease in inequality between year *t* and year *s* that is unlikely to be just due to sampling variation. A '0' indicates that any change is not significantly different from zero.

FIGURE A.1

95% confidence intervals around the income Gini coefficient

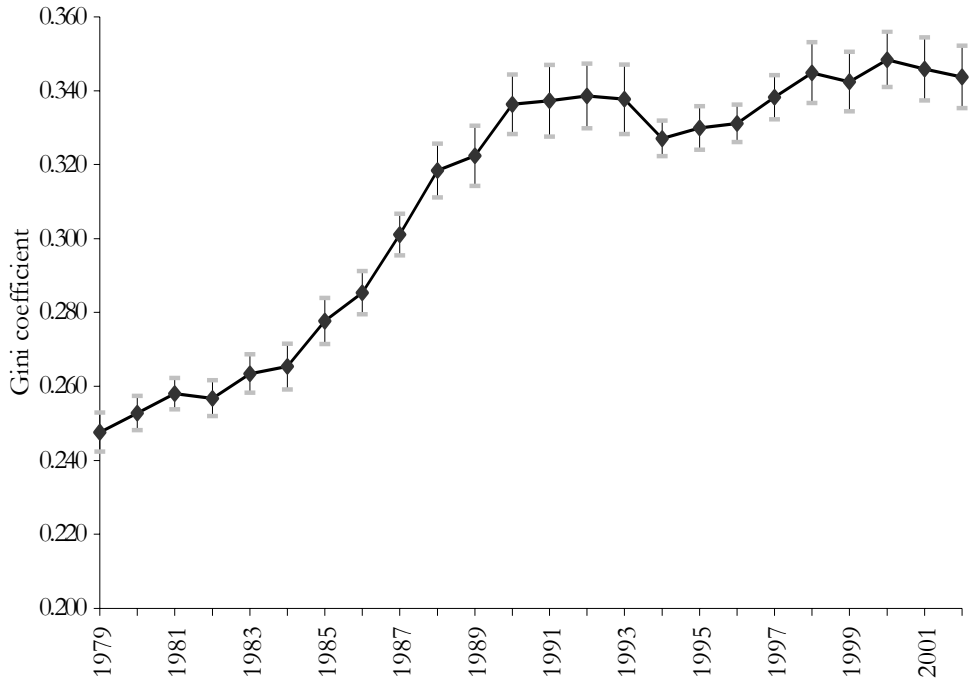

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Sources: Income data from the Family Expenditure Survey from 1979 to 1993/94 and the Family Resources Survey thereafter.

FIGURE A.2

95% confidence intervals around the expenditure Gini coefficient

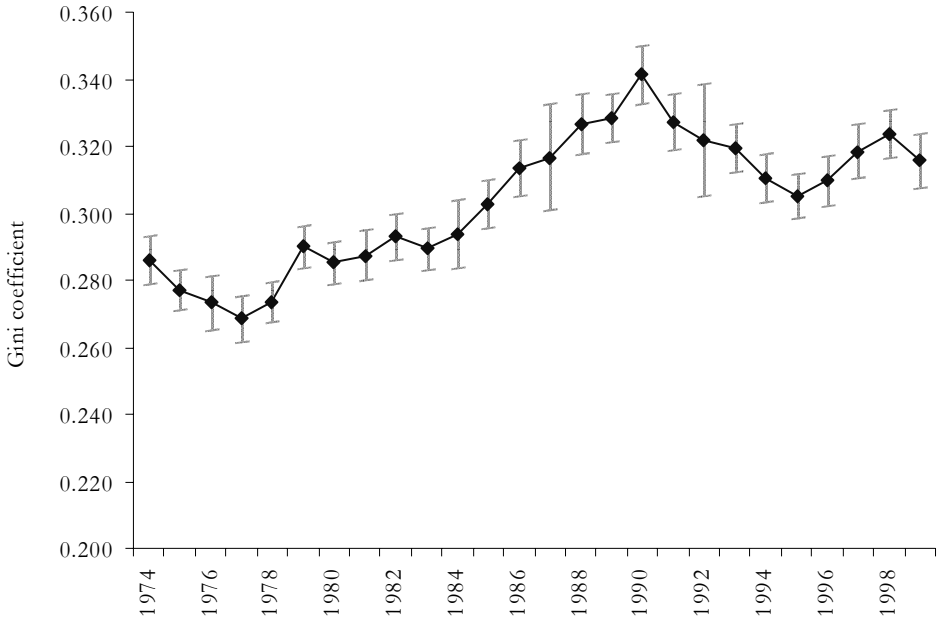

Note: Points are drawn for calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

TABLE A.1

Test for significance of differences in income Gini coefficients between pairs of years

	Year *s*																							
	79	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	00	01	02	
		0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	79
			0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	80
				0	0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	81
					0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	82
						0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	83
							+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	84
								0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	85
									+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	86
										+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	87
											0	+	+	+	+	+	+	+	+	+	+	+	+	+	88
												+	+	+	+	0	0	0	+	+	+	+	+	+	89
													0	0	0	0	0	0	0	0	0	+	0	0	90
														0	0	0	0	0	0	0	0	0	0	0	91
															0	0	0	0	0	0	0	0	0	0	92
																0	0	0	0	0	0	0	0	0	93
																	0	0	+	+	+	+	+	+	94
																		0	0	+	+	+	+	+	95
																			0	+	+	+	+	+	96
																				0	0	+	0	0	97
																					0	0	0	0	98
																						0	0	0	99
																							0	0	00
																								0	01
																									02

Year *t*

Note: Years refer to calendar years up to and including 1992, and financial years thereafter.

Sources: Authors' calculations using the Family Expenditure Survey and the Family Resources Survey.

TABLE A.2

Test for significance of differences in expenditure Gini coefficients between pairs of years

Year *s*																									
74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	
	0	–	–	–	0	0	0	0	0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	74
		0	0	0	+	0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	75
			0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	76
				0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	77
					+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	78
						0	0	0	0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	79
							0	0	0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	80
								0	0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	81
									0	0	0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	82
										0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	83
											0	+	+	+	+	+	+	+	+	+	+	+	+	+	+	84
												0	0	+	+	+	+	+	+	0	0	0	+	+	+	85
													0	+	+	+	+	0	0	0	0	0	0	+	0	86
														0	+	+	0	0	0	0	–	0	0	0	0	87
															0	+	0	0	0	–	–	–	0	0	–	88
																+	0	0	0	–	–	–	0	0	+	89
																	–	–	–	–	–	–	–	–	–	90
																		0	0	–	–	–	0	0	–	91
																			0	–	–	–	0	0	0	92
																				0	–	0	0	0	0	93
																					0	0	0	+	0	94
																						0	+	+	0	95
																							0	+	0	96
																								0	0	97
																									0	98
																										99

Year *t*

Note: Years refer to calendar years up to and including 1992, and financial years thereafter.

Source: Authors' calculations using the Family Expenditure Survey.

Looking first at the changes in the income Gini coefficient, Table A.1 shows that the great majority of years since 1981 have seen inequality that is statistically significantly higher than it was two years earlier, though for a number of years in the mid-1980s inequality rose significantly year on year. It is also interesting to notice that income inequality in 2002/03 remained significantly higher than its 1996/97 level. Table A.2 shows that generally, year-on-year changes in expenditure inequality are not statistically significantly different from zero. However, over longer periods of time, we find many instances of statistically significant increases or decreases in inequality.

References

Atkinson, A. (1999), 'The distribution of income in the UK and OECD countries in the twentieth century', *Oxford Review of Economic Policy*, vol. 15, no. 4, pp. 56–75.

— (2003), 'Top incomes in the United Kingdom over the twentieth century', University of Oxford, Nuffield College, mimeo (www.nuff.ox.ac.uk/users/atkinson/TopIncomes20033.pdf).

Blow, L., Leicester, A. and Oldfield, Z. (2004), *Consumption Trends in the UK, 1975–99*, London: Institute for Fiscal Studies.

Blundell, R. and Preston, I. (1998), 'Consumption inequality and income uncertainty', *Quarterly Journal of Economics*, vol. 113, pp. 603–40.

Brewer, M., Goodman, A., Myck, M., Shaw, J. and Shephard, A. (2004), *Poverty and Inequality in Britain: 2004*, Commentary no. 96, London: Institute for Fiscal Studies.

Bridges, S. and Disney, R. (2004), 'Use of credit and arrears on debt among low-income families in the United Kingdom', *Fiscal Studies*, vol. 25, pp. 1–25.

Card, D. and DiNardo, J. (2002), 'Skill biased technological change and rising wage inequality: some problems and puzzles', National Bureau of Economic Research, Working Paper no. 8769.

Clark, T. and Leicester, A. (2004), 'Inequality and two decades of British tax and benefit reforms', *Fiscal Studies*, vol. 25, pp. 129–58.

Cowell, F. (1995), *Measuring Inequality*, LSE Handbooks in Economics Series, London: London School of Economics.

Department for Work and Pensions (2004), Households Below Average Income 1994/95–2002/03, Leeds: Corporate Document Services.

Goodman, A., Johnson, P. and Webb, S. (1997), *Inequality in the UK*, Oxford: Oxford University Press.

Gregg, P. and Wadsworth, J. (2003a), 'Workless households and the recovery', in R. Dickens, P. Gregg and J. Wadsworth (eds), *The Labour Market under New Labour*, Basingstoke and New York: Palgrave Macmillan.

— and — (2003b), 'Labour market prospects of less skilled workers over the recovery', in R. Dickens, P. Gregg and J. Wadsworth (eds), *The Labour Market under New Labour*, Basingstoke and New York: Palgrave Macmillan.

Hills, J. (ed.) (1996), *New Inequalities: The Changing Distribution of Income and Wealth in the United Kingdom*, Cambridge: Cambridge University Press.

Jenkins, S. (1991), 'Income inequality and living standards: changes in the 1970s and 1980s', *Fiscal Studies*, vol. 12, no. 1, pp. 1–28.

— (1995), 'Accounting for inequality trends: decomposition analyses for the UK 1971–86', *Economica*, vol. 62, pp. 29–63.

— and Cowell, F. (1994), 'Dwarfs and giants in the 1980s: trends in the UK income distribution', *Fiscal Studies*, vol. 15, no. 1, pp. 99–118.

Johnson, P. and Stark, G. (1989), 'Ten years of Mrs Thatcher: the distributional consequences', *Fiscal Studies*, vol. 10, no. 2, pp. 29–37.

— and Webb, S. (1993), 'Explaining the growth in UK income inequality: 1979–1988', *Economic Journal*, vol. 103, pp. 429–35.

Lakin, C. (2004), 'The effects of taxes and benefits on household income, 2002–03', London: Office for National Statistics (www.statistics.gov.uk/articles/economic_trends/effects_taxes_benefits_household_income_02-03/Taxesbenefits2002-03.pdf).

Machin, S. (2003), 'Wage inequality since 1975', in R. Dickens, P. Gregg and J. Wadsworth (eds), *The Labour Market under New Labour*, Basingstoke and New York: Palgrave Macmillan.

McClements, L. (1977), 'Equivalence scales for children', *Journal of Public Economics*, vol. 8, pp. 191–210.

Piketty, T. and Saez, E. (2003), 'Income inequality in the United States 1913–1998', *Quarterly Journal of Economics*, vol. 118, pp. 1–39.

Saez, E. and Veall, M. (2003), 'The evolution of high incomes in Canada, 1920–2000', National Bureau of Economic Research, Working Paper no. 9607.

Sefton, T. (2002), 'Recent changes in the distribution of the social wage', London School of Economics, Centre for the Analysis of Social Exclusion, CASE Paper no. 62.

Shorrocks, A. F. (1982), 'Inequality decomposition by factor components', *Econometrica*, vol. 50, pp. 193–212.
Sianesi, B. (2003), 'Returns to education: a non-technical summary of CEE work and policy discussion', Institute for Fiscal Studies, mimeo.
Wilkinson, H. (1999), 'Balancing work life and personal life: new workers', *Guardian*, 6 January.
Women and Equality Unit (2003), 'Individual income 1996/97 – 2001/02', Crown Copyright, www.womenandequalityunit.gov.uk/research/full_ind_inc_02.pdf.