Session 2B: Economic Mobility and Vulnerability
Time: Monday, August 25, afternoon
Session Organizer (s): Timothy Smeeding and Stephan Klasen
Discussant: Gordon Anderson

Paper Prepared for the 30th General Conference of
The International Association for Research in Income and Wealth

Portoroz, Slovenia, August 24–30, 2008

Has income growth in Britain become more pro-poor?

by

Stephen P. Jenkins

and

Philippe Van Kerm

For additional information please contact:

Stephen P. Jenkins
Institute for Social and Economic Research
University of Essex, Colchester CO4 3SQ, United Kingdom
Email: stephenj@essex.ac.uk

Philippe Van Kerm
CEPS/INSTEAD
B.P. 48, L-4501 Differdange, Luxembourg
Email: philippe.vankerm@ceps.lu

Has income growth in Britain become more pro-poor?*

Stephen P. Jenkins

Institute for Social and Economic Research
University of Essex, Colchester CO4 3SQ, United Kingdom
Email: stephenj@essex.ac.uk

Philippe Van Kerm

CEPS/INSTEAD
B.P. 48, L-4501 Differdange, Luxembourg
Email: philippe.vankerm@ceps.lu

20 July 2008

*This research was supported by core funding from the University of Essex and the Economic and Social Research Council for the Research Centre on Micro-Social Change at ISER. The research is also part of the MeDIM project (*Advances in the Measurement of Discrimination, Inequality and Mobility*) supported by the Luxembourg 'Fonds National de la Recherche' (contract FNR/06/15/08) and by core funding for CEPS/INSTEAD from the Ministry of Culture, Higher Education and Research of Luxembourg.

Abstract

Assessments of who is getting better off over time typically summarize changes in the incomes of particular groups, e.g. the poor or the rich, or for subgroups such as lone parent families and other families with children. These calculations ignore the fact that these groups change composition over time: the same individuals are not being compared. To assess whether this year's poor (or rich) are gainers or losers, one has to track the fortunes of individuals using longitudinal data. Using data from the British Household Panel Survey, we compare the patterns of individual income growth over the period 1992–1996 with those of the period 1999–2003. We develop methods providing a longitudinal perspective, and show that the pattern of income growth became more pro-poor between periods, and in a different manner than is revealed by conventional analysis. The results are consistent with the aims of the New Labour government to reduce pensioner and child poverty.

Keywords: pro-poor growth; income mobility; redistribution; British Household Panel Survey

JEL Classification: D31; D63; I32

1 Introduction

The income distribution in each year can be characterized as a Parade of Dwarfs and a few Giants (Pen, 1971).[1] In each year, every individual in the population is represented by a person who has a height proportional to the individual's income, and these representatives are lined up in order of height with the shortest at the front. Income growth over time corresponds to changes in the heights of the Parade participants. How to assess whether income growth between two years *A* and *B* has been pro-poor or pro-rich is the subject of this paper, together with empirical evidence on the topic for Britain.

The conventional approach to assessing who has got better off is to compare incomes at a series of common points in the Parades for *A* and *B*, e.g. the incomes one-tenth, one-fifth, half-way, or three-quarters of the way along each Parade. Growth incidence curves, commonly used in development economics (Ravallion & Chen, 2003), are an example of this approach. Alternatively, cross-year comparisons are made of the shares of total income held by the poorest tenth, the poorest fifth, the richest tenth, and so on. The UK's official statistics on the personal income distribution, *Households Below Average Income*, report both selected percentiles and income group shares, and Figure 1 illustrates the picture of income growth derived from the first of these.[2]

Over the four years between fiscal years 1994/95 and 1998/99, income growth was relatively homogeneous across the middle half of the distribution, at a rate of around 10%. At the tails, the experience was quite different: the 5^{th} percentile grew by 3.1% but the 95^{th} percentile grew by 11.6%. By contrast, over the four years between fiscal years 1998/99 and 2002/03, income growth was greater the lower the percentile (with the exception of the 5^{th} percentile). The growth rate was almost 14% at the median and 17% at the 15^{th} percentile, but nearer 10% at the 75^{th} percentile and above. As a result, overall inequality, as measured by e.g. the Gini coefficient, declined somewhat. *Households Below Average Income* also provides information each year about the incomes of different population subgroups, with breakdowns by family type and work status. For example, Department for Work and Pensions (2008), Table 2.4ts, reports that the median income of a couple with two children aged 5 and 14 increased by 7.8% between 1994/95 and 1998/99, and 13.5% between 1998/99 and 2002/03. For a lone parent with two children aged 5 and 14, the corresponding increases were 8.1% and 13.6%.

As a means of describing *who* has got better off, income change statistics of the type just described have a fundamental flaw. Looking at the change in the income of the

[1] See also Jenkins & Cowell (1994).

[2] See also Brewer *et al.* (2008). The data are derived from the annual cross-sectional survey, the *Family Resources Survey*. The current HBAI series began in 1994/95, and the latest year for which statistics are available is 2006/07. The periods used in Figure 1 were defined to broadly correspond with the periods used in our own analysis reported later.

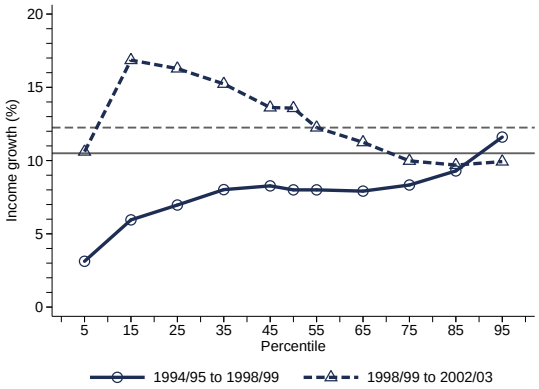

Figure 1. Income growth (%) of selected percentiles, 1994/95–1998/99 and 1998/99–2002/03. The horizontal lines show the growth rate of mean income in each period. Source: authors' calculations from Department for Work and Pensions (2008), Table 2.1ts.

person at the middle of Parade *A* and the person at the middle of Parade *B* ignores the fact that the persons concerned are not the same individual. Over time, people change their position in the income Parade. Mobility means that the people who are poor this year are not the same people who were poor last year: some leave poverty and others enter. More generally, income groups of every kind change their composition over time. Similarly, the people who are lone parents in one year are not the same as those who are lone parents in another year: some remarry or their children become non-dependent, and others become lone parents via divorce, widowhood, or an extra-partnership birth.

To assess whether this year's poor (or rich) are gainers or losers, one has to track the fortunes of these individuals, not of income groups such as 'the poor' whose composition may change from one year to the next. Longitudinal data are therefore required, as it is only these that enable one to link the income of a specific individual in Parade *A* with her income in Parade *B*, and hence calculate income growth at the individual level.

In this paper, we examine the extent to which income growth in Britain since the early 1990s benefited the poor rather than the rich, or vice versa, and examine how patterns of income growth changed between 1992–1996 and 1999–2003. We develop a properly longitudinal approach to assessing the progressivity of income growth, applying techniques developed by Jenkins & Van Kerm (2006) and Van Kerm (2006), and contrast them with the conventional approaches to assessment described above. By contrast with much of the related literature, all our statistics are accompanied by calculations of associated sampling variances, so that the impact on inference of sampling variability may be considered.

Our substantive results contribute to the understanding of the redistributive effects of the policies introduced by the Labour government that came to power in May 1997, and how these redistributive effects compare with those of the Conservative government that

it replaced. Labour's principal motivation was to reduce poverty rather than to reduce inequality. When asked in 2001, whether it was acceptable for the gap between rich and poor to get bigger, the Prime Minister stated that 'the issue isn't in fact whether the very richest person ends up becoming richer. The issue is whether the poorest person is given the chance that they don't otherwise have', and 'the most important thing is to level up, not level down'.[3] With respect to our distinction between cross-sectional and longitudinal perspectives, we note that there is ambiguity in the Prime Minister's statement about which he was referring to.

The Labour government has aimed to raise employment rates ('work first') by increasing the rewards to work relative to not working, and to ensure that working families were not poor ('making work pay'), and thereby also reduce child poverty rates. The principal programmes directed at these ends were Working Families Tax Credit (WFTC), introduced in October 1999, providing means-tested support to low income working families that was more generous than its predecessor (Family Credit). (WFTC was itself replaced by the Working Tax Credit and Child Tax Credit programmes from April 2003.) There were other changes during this period that raised support for families with children, notably increases in universal Child Benefit, and increases in the child allowances in Income Support, the social assistance benefit for those not in paid work. A national minimum wage was introduced in April 1999. To reduce pensioner poverty, a means-tested minimum income guarantee for poor pensioners was introduced in April 1999, and became the Pensioner Credit in October 2003.[4]

To be clear, our analysis sheds light on who gained and who lost in an aggregate sense – we do not examine the causal redistributive impact of particular policy measures (or the business cycle). We describe individual income growth for individuals at all positions in the income Parade, contrasting the patterns for a period when the Conservative government was in power with a period when the Labour government was in power, also adding some details about the changing circumstances of subgroups of particular interest such as families with children and pensioners.

We show that, from a longitudinal perspective, the pattern of individual income growth has been pro-poor under both Conservative and Labour governments, but it was more so after Labour came to power, at least when absolute growth is considered rather than proportionate growth. This contrasts with the conventional cross-sectional approach (illustrated by Figure 1) which suggests that growth was pro-rich in the mid-1990s but pro-poor after the 1997 change of government.

Section 2 reviews the longitudinal methods we use to summarize the 'pro-poorness'

[3]Interview with the Prime Minister on BBC Newsnight, 5 June 2001: transcript at `http://news.bbc.co.uk/1/hi/events/newsnight/1372220.stm`.

[4]See inter alia Brewer & Shephard (2004) and Bennett & Millar (2005) for further details of the policies introduced by the Labour government.

of income growth, and compares them with the conventional methods. Our data, from the British Household Panel Survey (BHPS), and the measure of income, are described in Section 3. The estimates of income growth are presented in Section 4. Section 5 contains a summary and conclusions. Appendix A explains the statistical techniques used to derive our estimates and their bootstrapped sampling variances. Appendix B compares income distributions trends derived from the BHPS with those derived from the *Households Below Average Income* which are based on the annual *Family Resources Survey*.

2 Describing individual income growth

To measure the income growth between two years t and $t+\tau$ for a set of individuals, one requires information about income at t and $t+\tau$ for every individual. That is, one requires data about the joint distribution of income at t and $t+\tau$, and not simply the two marginal distributions for each year. Denote the bivariate distribution of income as

$$H_{t,t+\tau}(x,y) = \Pr[X \leq x, Y \leq y]$$

where X and Y are jointly distributed random variables that describe incomes at year t and $t+\tau$ respectively. Because marginal distributions can be derived from the joint distribution ($F_t(x) = \int H_{t,t+\tau}(x,s)ds$), any summary statistic based on a marginal distribution for a given year (e.g. the Gini coefficient, or percentiles), and changes between t and $t+\tau$ in these statistics, can be expressed as a function of the joint distribution $H_{t,t+\tau}$. Panel data on incomes also allow us to estimate $H_{t,t+\tau}$. With repeated cross-sectional data on incomes, one can estimate marginal distributions but cannot observe individual-level income changes. In this section, we develop graphical and numerical methods that provide a longitudinal perspective on individual-income growth based on $H_{t,t+\tau}$.

Income *mobility profiles* are graphs that show the income growth experienced by people at every position along the base-year income Parade (Van Kerm, 2006). To construct these profiles, one needs a definition of 'income growth', i.e. a choice of functional form for an individual income growth function $\delta\left(x,y;H_{t,t+\tau}\right)$ which summarizes income change for a person with income x at time t and y at time $t+\tau$. We consider the cases of *absolute* income growth in which case $\delta\left(x,y;H_{t,t+\tau}\right) := y-x$, and *proportionate* income growth in which case $\delta\left(x,y;H_{t,t+\tau}\right) := \log(y) - \log(x)$. Equiproportionate income increases correspond to larger absolute income increases for poor people than for rich people.[5] The mobility profile shows the expected value (average) of $\delta\left(x,y;H_{t,t+\tau}\right)$ at each Parade position.

[5] A variety of individual income growth functions, each capturing different distinct aspects of income changes, have been used in the income mobility literature and many income mobility indices can be expressed as population averages of them. See Fields *et al.* (2002) and Cecchi & Dardanoni (2003) for examples.

More formally, define the normalized rank $p \in [0,1]$ in an income Parade, where $x(p) = F_t^{-1}(p)$ and $F_{t+\tau}(s|x(p))$ is the conditional distribution of period $t+\tau$ incomes given initial income $x = x(p)$. The mobility profile is a plot of pairs

$$\left(\int \delta\left(x(p), s; H_{t,t+\tau}\right) dF_{t+\tau}(s|x(p)); p \right)$$

over $p \in [0,1]$. Put differently, the mobility profile is a plot of individual income growth function values against individuals' normalized ranks in the base year income distribution.[6]

The mobility profile thus reveals how income growth is distributed according to position in the base year income distribution. If changes in fortunes are the same for everyone, the profile is horizontal. If income growth is pro-poor (or 'progressive'), then the profile has a negative slope; if income growth is pro-rich ('regressive'), then the profile has a positive slope.[7]

Comparisons of income mobility profiles for different time periods facilitate discussion of whether income growth has become more pro-poor or more pro-rich over time. In this connection, Van Kerm (2006) showed that, if a pair of profiles (or cumulated profiles) does not cross, conclusions robust to the relative weight that one attaches to people at different ranks in the distribution can be drawn about which period has the greater degree of pro-poor growth.

The mobility profile differs fundamentally from the growth incidence curve that is widely used in the development economics literature to summarize 'pro-poor growth' (Ravallion & Chen, 2003), and also from closely related devices such as the poverty growth curve (Son, 2004). A growth incidence curve is, in our notation, a plot of pairs

$$\left(\delta\left(x(p), y(p); H_{t,t+\tau}\right); p\right)$$

over $p \in [0,1]$, with $\delta\left(x, y; H_{t,t+\tau}\right) := 100\left(\frac{y}{x} - 1\right)$, $x(p) = F_t^{-1}(p)$ and $y(p) = F_{t+\tau}^{-1}(p)$.[8] Thus a growth incidence curve is based on the two marginal distributions, and records changes in quantile values. Each income pair may refer to the incomes of two different individuals. By contrast, the income mobility profile tracks the income changes of each

[6] See Van Kerm (2006) for further discussion and illustrations. The mobility profile is an example of the technique known as fractile graphical analysis introduced by Mahalanobis (1960). See also Sen (2005). Independently of our earlier work, Grimm (2007) advocated the use of similar devices.

[7] Measurement error and transitory variation impart a negative slope to mobility profiles separately from 'pure' income growth effects. We return to this regression to the mean issue in the empirical application below.

[8] In our empirical work, we summarize proportionate growth using $\delta\left(x, y; H_{t,t+\tau}\right) := \log(y) - \log(x)$ rather than $\delta\left(x, y; H_{t,t+\tau}\right) := 100\left(\frac{y}{x} - 1\right)$ in order to provide a direct link with the mobility profile for proportionate growth. Unsurprisingly the two definitions result in growth incidence curves of virtually identical shape.

individual as summarized by the individual income growth function (δ), which depends on the joint bivariate distribution of income. In our empirical application, we show how the two approaches lead to different conclusions about patterns of income growth.

We supplement our graphical analysis with numerical indices. These allow us not only to summarize the amount of individual income growth overall, but also to provide a link between changes in marginal distributions and individual income changes. Jenkins & Van Kerm (2006) showed that the change in inequality between two years can be written in terms of a component representing the degree of progression of individual income growth and a component representing the amount of reranking in the income Parade associated with differential income growth. In short, Inequality Change equals Reranking minus Progressivity. More formally, the decomposition states that:

$$\mathrm{GINI}(F_{t+\tau};\upsilon) - \mathrm{GINI}(F_t;\upsilon) = R(H_{t,t+\tau};\upsilon) - P(H_{t,t+\tau};\upsilon)$$

where $\mathrm{GINI}(F_t;\upsilon)$ is the generalized Gini coefficient for the marginal income distribution F_t.[9] The standard Gini coefficient arises when $\upsilon = 2$, and this is the case on which we focus in this paper.[10]

The change in the Gini coefficient represents the 'net' effect on inequality in marginal distribution of all the individual income changes. The decomposition identifies a 'gross' effect, $P(H_{t,t+\tau};\upsilon)$, summarizing the growth between time t and $t+\tau$ of individuals' incomes, which is combined with a pure mobility effect – the effect of the reranking associated with differential income growth. The reranking component, $R(H_{t,t+\tau};\upsilon)$, corresponds to the changes in income group composition that we referred to in the Introduction. Put another way, if every one retained the same rank in the income Parade over time, then the gross and net effects of income change would be the same, and mobility profiles and growth incidence curves would simply be two ways of describing the same phenomena rather than being fundamentally different.

$P(H_{t,t+\tau};\upsilon)$ measures the degree of progressivity in income growth. It summarizes how much income growth favours people with low income in the base year and so is a measure of how pro-poor income growth is. It is defined as

$$P(H_{t,t+\tau};\upsilon) = \int w\left(F_t(x);\upsilon\right)\left[\frac{y}{\mu_{t+\tau}} - \frac{x}{\mu_t}\right] dH_{t,t+\tau}(x,y)$$

[9]The generalized Gini coefficient was introduced by Donaldson & Weymark (1980, 1983) and Yitzhaki (1983).

[10]In this case, the decomposition can also be interpreted in terms of the difference between the Lorenz curves for the two marginal distributions in question (Jenkins & Van Kerm, 2006). $P(H_{t,t+\tau};2)$ is twice the area between the Lorenz curve for base year incomes and the concentration curve for current year incomes based on base year income ranks, and $R(H_{t,t+\tau};2)$ is twice the area between the latter curve and the current year Lorenz curve.

where $w(p;\upsilon) = \upsilon(1-p)^{\upsilon-1}$, $\upsilon > 1$, is a weight function also used to define the generalized Gini coefficient, with greater weight being given to persons with lower normalized rank in the base year distribution (i.e. poorer). In particular, for the standard Gini coefficient, $w\left(F_t(x);\upsilon\right)$ decreases linearly from 2 to 0 as p increases from 0 to 1. $P(H_{t,t+\tau};\upsilon)$ is therefore equal to a weighted integral of the income mobility profile for $\delta\left(x,y;H_{t,t+\tau}\right) := \frac{y}{\mu_{t+\tau}} - \frac{x}{\mu_t}$. (The individual income growth function summarizes changes in income shares in this case, since incomes in each year are normalized by the mean income for that year.) The greater that $P(H_{t,t+\tau};\upsilon)$ is, the more pro-poor is income growth.

The second component of the decomposition, $R(H_{t,t+\tau};\upsilon)$, summarizes the reranking associated with individual income growth. If income growth were to favour people with low income at time t, but did not lead to any changes of position in the income Parade, then inequality would fall. One would conclude that income growth is pro-poor. However, growth could be even more pro-poor if the income gains of those with lower incomes were such that they overtook people who had initially higher incomes. In this case, the effect on inequality in the marginal distribution is not known a priori. Inequality may fall but can also rise despite pro-poor growth. This was the case in the USA in the 1980s when inequality grew substantially (Jenkins & Van Kerm, 2006). By contrast, observe that regressive income growth – that favouring the rich in the base year – is necessarily associated with an increase in inequality.[11] The reranking effect that may offset the equalizing effect of pro-poor growth is summarized by $R(H_{t,t+\tau};\upsilon)$:

$$R(H_{t,t+\tau};\upsilon) = \int \left[w\left(F_t(x);\upsilon\right) - w\left(F_{t+\tau}(y);\upsilon\right)\right] \frac{y}{\mu_{t+\tau}} dH_{t,t+\tau}(x,y).$$

The decomposition also reminds us that changes over time in the progressivity of income growth cannot be inferred from trends in inequality index changes: the degree to which income growth becomes more pro-poor or more pro-rich depends also on changes over time in the extent to which individuals change their positions in the income Parade.

Details of the estimators used to obtain point estimates and standard errors of the statistics referred to in this section are provided in Appendix A.

3 Data

We use panel data from the first 14 waves of the British Household Panel Survey (BHPS), corresponding to survey years 1991–2004. The BHPS's first wave is a nationally representative sample of the population of Great Britain living in private households in 1991.

[11] See also O'Neill & Van Kerm (2008) for illustrations of heterogeneous patterns of progressivity, reranking and inequality change.

Original sample respondents (including each partner from a dissolved Wave 1 partnership) have been followed and they, and their co-residents, were interviewed annually subsequently. Children in original sample households are also interviewed when they reach the age of 16 years. Thus the sample remains broadly representative of the population of Britain over time.[12]

To assess whether patterns of income growth changed over time, we contrast estimates from two distinct time periods. The first covers 1992 to 1996, when there was a Conservative government, headed by John Major, and a second period covering 1999 to 2003, when there was a Labour government (from May 1997), headed by Tony Blair. Our choice of years and the gap between the time periods was constrained by the years for which income data were available (see below) and by the fact that we wished to clearly differentiate between periods that were distinctly 'Conservative' and distinctly 'Labour'. Initial analyses using a moving four-year window over the entire 1991–2004 period suggest that our results are not substantially affected by the specific definitions of the 'Conservative' and 'Labour' periods: see Appendix C.

Our prior expectation, largely formed by the cross-sectional analyses of the type summarized in Figure 1 and the nature of the policies introduced, was that income growth would be more pro-poor under Labour than under the Conservatives. And yet, at the same time, it was unclear to us *ab initio* whether the mild decline in inequality of the (marginal) income distribution from the late 1990s to the early 2000s was also associated by significant reranking. As demonstrated by our decomposition earlier, the degree of progressivity of income growth depends on both factors.

Our 1992–1996 sample is composed of respondents for whom income data are available for survey years 1991–1993 and 1994–1996 (because our income measure for each year is a three-year longitudinal average – see below), and 1998–2003 sample is composed of respondents with income data for 1997–1999 and 2002–2004. The two sub-samples have an overlapping membership, which has implications for the calculations of standard errors for our estimators (see Appendix A). After dropping cases with missing data on income, the 1992–1996 sample is composed of 6,537 individuals and the 1999–2003 sample is composed of 6,093 individuals.

Our measure of the living standard of individuals is conventional – the income of the household to which they belong, adjusted for differences in household size and composition using an equivalence scale, and expressed in constant prices. Specifically, income is 'current net household income', which is the sum across all household members of cash income from work, capital income, private and public transfers, minus direct taxes and occupational pension contributions. The income reference period is the month prior to the

[12]Data from the extension samples for Scotland, Wales, and Northern Ireland that began in the late 1990s are not used. For a detailed discussion of the BHPS and documentation, see Lynn (2006) and `http://www.iser.essex.ac.uk/ulsc/bhps/doc`.

interview or the most recent relevant period, converted to a weekly equivalent pro rata.[13] Total money income was equivalized using the modified OECD equivalence scale, and expressed in January 2006 pounds using a 'before housing costs' price index series from the Department for Work and Pensions. The BHPS net income data were constructed to adhere as closely as possible to the 'before housing costs' income definition that is used in the UK's official income statistics, the *Households Below Average Income* series, cited in the Introduction.[14] In Appendix B, we show that estimates of cross-sectional trends in (marginal) income distributions derived from the BHPS data are broadly consistent with the corresponding trends derived from HBAI data.

All income distribution statistics are sensitive to measurement errors and transitory variations in income, but the issue is particularly relevant when estimators are based on individual-level measures of income change. To mitigate the impact of these factors, we smoothed respondents' incomes over three consecutive years: our measure of income for a person in year t is the average of her income in years $t-1$, t and $t+1$.[15] The use of three-year averaged income follows Gottschalk & Danziger (2001) and Jenkins & Van Kerm (2006).[16]

In some analyses, we compare the experiences of individuals classified according to their family type. The definition of family type is as in the HBAI reports, namely: 'pensioner' (single or couple), 'couple with children', 'couple without children', 'single with children', 'single without children'.[17]

4 Patterns of income growth in Britain, and their changes over time

In this section we contrast patterns of income growth over the period 1992–1996 with the patterns for 1999–2003. The conventional cross-sectional approach, illustrated by Figure 1 which is based on published HBAI data, suggests that income growth was distinctly

[13]The use of a current rather than annual measure of income is standard in Britain. Böheim & Jenkins (2006) argue that the BHPS current and annual income measures provide very similar estimates of most distributional summary statistics.

[14]The BHPS net income data are available as an unofficial supplement to the official BHPS release (Bardasi *et al.*, 1999, Levy *et al.*, 2006). Fourteen waves of data are currently available. For more details of the HBAI definition of income, see Department for Work and Pensions (2008).

[15]In addition, mobility profiles were estimated using a robust estimator that limits the leverage of the largest income changes: see Appendix A.

[16] In current work, we are also considering applications of the exploratory bounding methods developed by Fields *et al.* (2003) which were applied to single-year incomes.

[17]See Bardasi *et al.* (1999) and Levy *et al.* (2006) for details of the definitions.

more pro-poor in the second period compared to the first. We reexamine this finding using our BHPS data, using the three-year-averaged income variable to control for potential measurement error and transitory variation, and exploiting access to the unit record data to provide a much more detailed picture. The picture provided by growth incidence curves is followed by a longitudinal analysis employing mobility profiles and decompositions of changes in Gini coefficients.

4.1 Income growth: a cross-sectional perspective

Figure 2 shows growth incidence curves for 1992–1996 (left) and 1999–2003 (right) of two types. The graphs in the top panel show income growth in absolute terms; the graphs in the bottom panel show income growth in proportionate terms, as measured by changes in log income. (The graphs in the bottom panel correspond to growth incidence curves as conventionally defined – see above.) We used both definitions of growth so that there was a direct correspondence with the two definitions used to construct mobility profiles.

Over the period 1992–1996, income growth was approximately £10 on average. In absolute terms, the pattern was regressive as the incomes at the bottom of the distribution increased by less than the average, while incomes at the top of the distribution increased by substantially more than the average. In the later period, the pattern of growth was more pro-poor because the lowest incomes also rose. The result was that income change was distinctly more homogeneous right across the distribution (the apparently higher growth at the highest percentiles may simply reflect sampling variability). When we look instead at proportionate growth (bottom panel), the contrast between the two periods is somewhat different. In the earlier period, growth was broadly equiproportionate (the common rate across percentiles is about 5%) but, in the later period, the pattern was more pro-poor, with income growth distinctly higher at the lowest percentiles.

So, according to the conventional cross-sectional approach, income growth became more pro-poor whether defined in absolute or proportional terms.[18] How does this picture change if we take a longitudinal perspective, and take account of the fact that individuals move between income groups over time?

[18] Analysis based on microsimulation models finds that redistribution can be attributed to Labour's policies. Sefton & Sutherland (2005) compare estimated incomes in 2004/05 under the 2004/05 tax-benefit rules and estimated incomes in 2004/05 had policies remained as they were in 1997. Their overall conclusion is that 'the impact of the tax and benefit policies introduced since 1997 imply a significant redistribution towards those on low incomes, compared with what would have happened if the old system has simply been adjusted for inflation' (p. 244).

4.2 Income growth: a longitudinal perspective

We now assess the distribution of income growth from an individual-level perspective. Figure 3 shows income mobility profiles for 1992–1996 (left) and 1999–2003 (right). Although there is some superficial similarities with a growth incidence curve – the horizontal axis of each graph has the same scale – mobility profiles are fundamentally different because they exploit the longitudinal nature of the data. They show the growth in the income of individuals (rather than the growth of percentiles). The top two curves of Figure 3 show the income growth in absolute terms (the units are January 2006 pounds) and the bottom two curves show income growth in proportionate terms (the units are log January 2006 pounds).

From the longitudinal perspective, the pattern of income growth is quite different from that suggested by the growth incidence curves. In both periods, growth is clearly pro-poor, characterized by greater growth for those who were initially at the bottom. This is regardless of whether an absolute or proportionate definition of growth is used. The key difference now is that an increase in pro-poorness between the two periods is only apparent when income growth is defined in absolute terms. Those who were initially in the bottom half of the distribution or, more particularly, initially in the poorest fifth, increased their income by more than those in the top half of the distribution. However, although income growth was pro-poor in proportionate terms in both periods, there was no apparent increase in its degree between the two periods.

The negative slopes of the mobility profiles describe a form of regression to the mean. There is an issue of how to gauge the extent to which this is genuine – attributable to genuine economic phenomena – or reflecting the effects of measurement errors and transitory variations in income. If there is measurement error of the 'classical' form (uncorrelated with the true value and over time), then the expected income increase for someone with a below-average income is positive and is negative for someone with above average income.[19] So, some of the observed progressivity in income growth may be spurious. Our use of a three-year income average aims to reduce the impact of this problem by smoothing out transitory variations and measurement error while still keeping track of the more substantial variations.[20] Moreover we have no reason to expect the effects of measurement error and transitory variation to have increased over time. We are therefore inclined to attribute the *changes* in pro-poorness that we observe to genuine economic phenomena.

[19] An analogy may help. If one rolls a standard die, the expected number of spots at any roll is 3.5 (the sum of the possible scores divided by six). If the first roll in fact produces a 1, then the expected increase in the score when the die is rolled again is positive (+2.5). By contrast, if a 6 comes up first, the expected gain at the second roll is negative (–2.5). So, despite there being no association between the first and the second rolls (the die is fair), there is a correlation between the initial outcome and the change in outcome.

[20] In current work, we are also exploring the application of the bounding methods used by Fields *et al.* (2003).

The fact that our qualitative conclusions about the pro-poor pattern of income growth and its changes over time depend on whether individual income growth is measured in absolute terms or proportionate terms is a reminder that assessments of the degree of progressivity of income growth depend on the way in which the individual income growth function, $\delta\left(x, y; H_{t,t+\tau}\right)$, is defined. Also relevant is the weight given to income growth in different parts of the income distribution. Observe, for example, that if one were to ignore how growth was distributed along the income range, one would prefer the 1999–2003 situation because the overall average growth rate was slightly higher for that period. But the more weight that assessments give to growth among the poorest, the more likely it is that the later period's pattern is preferred.

It is therefore of interest to consider the robustness of the conclusion that income growth was more progressive in the second period than the first. The tools relevant for these checks are the mobility profile dominance results of Van Kerm (2006), which link particular configurations of profiles (dominance comparisons) with unambiguous rankings according to social evaluation functions respecting a small number of desirable properties. Implementation of these checks indicates that the 1999–2003 pattern of income growth dominates the 1992–1996 pattern at the 'second order.' This means that if one weights the income growth of people according to the inverse of their rank in the base year distribution, i.e. giving more importance to growth for people with lower income positions, the 1999–2003 pattern is preferable for any weight function that is positive, decreasing, and decays to zero at the top normalized income rank (Van Kerm, 2006). The Gini weight function we use is a particular example of such a function.

How can one reconcile the fact that the growth incidence curve for 1992–1996 suggested that there was income growth at the top of the income distribution, whereas the mobility profile for the same period suggests the opposite? The answer is that most of the increase in the top percentiles was due to income gains by individuals who did not belong to the top income group in 1992 but did in the later year. There was reranking over time in the income Parade. We provide additional information about changes in group membership shortly.

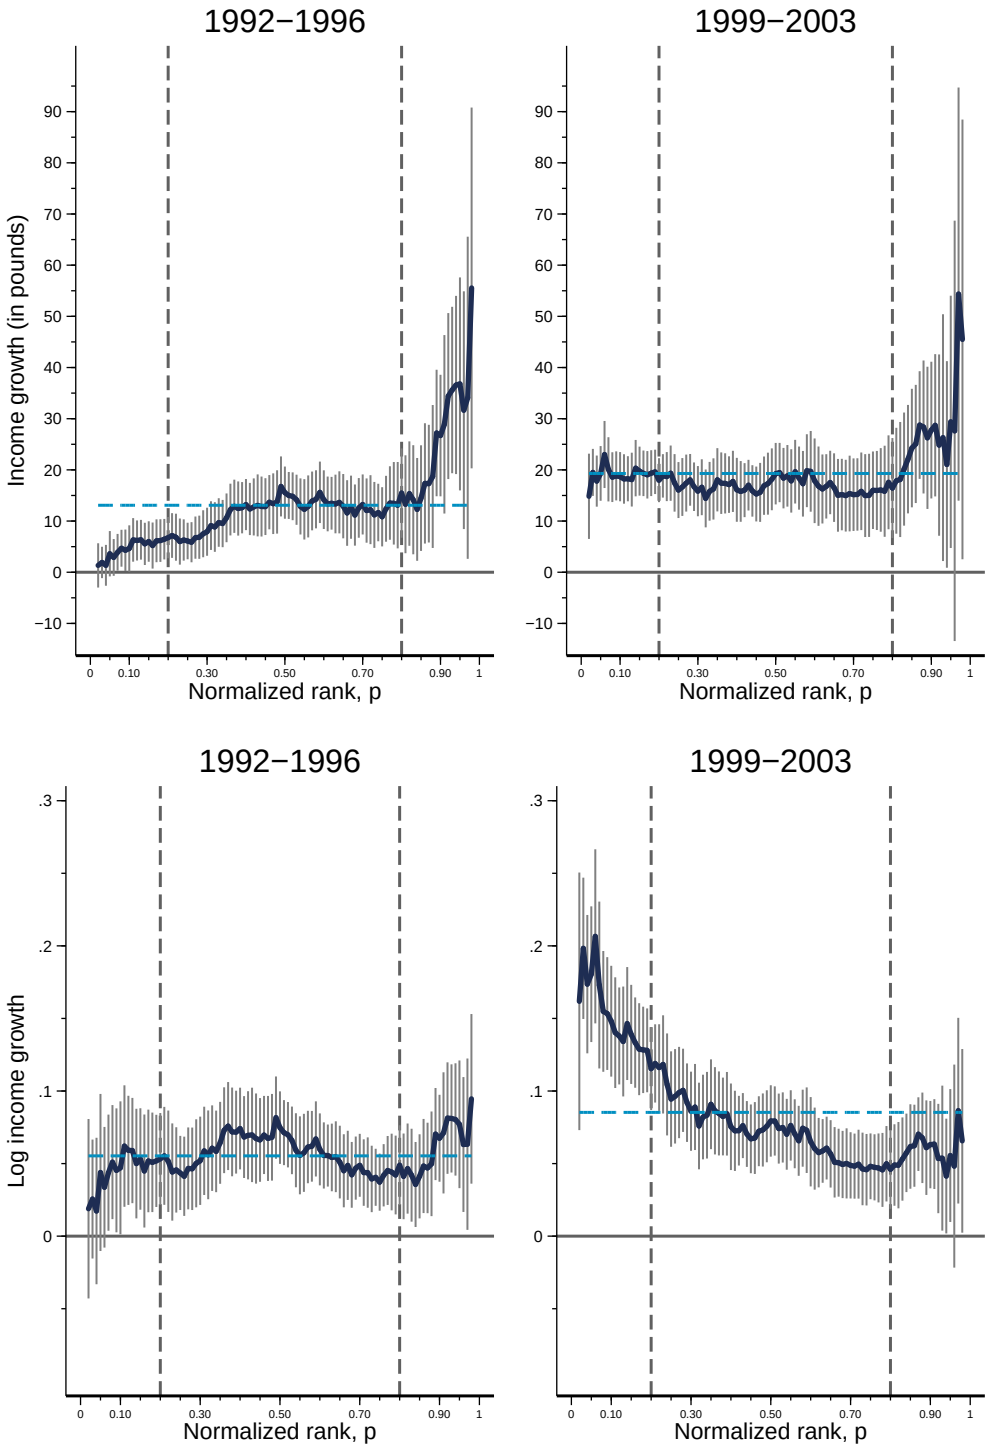

Figure 2. Growth incidence curves: income growth by percentile, 1992–1996 (left) and 1999–2003 (right). Top panel shows absolute growth (change in income); bottom panel shows proportionate growth (change in log income). Horizontal dashed lines mark the percentage growth of mean income in each period. Vertical bars show pointwise two-standard-error variability bands.

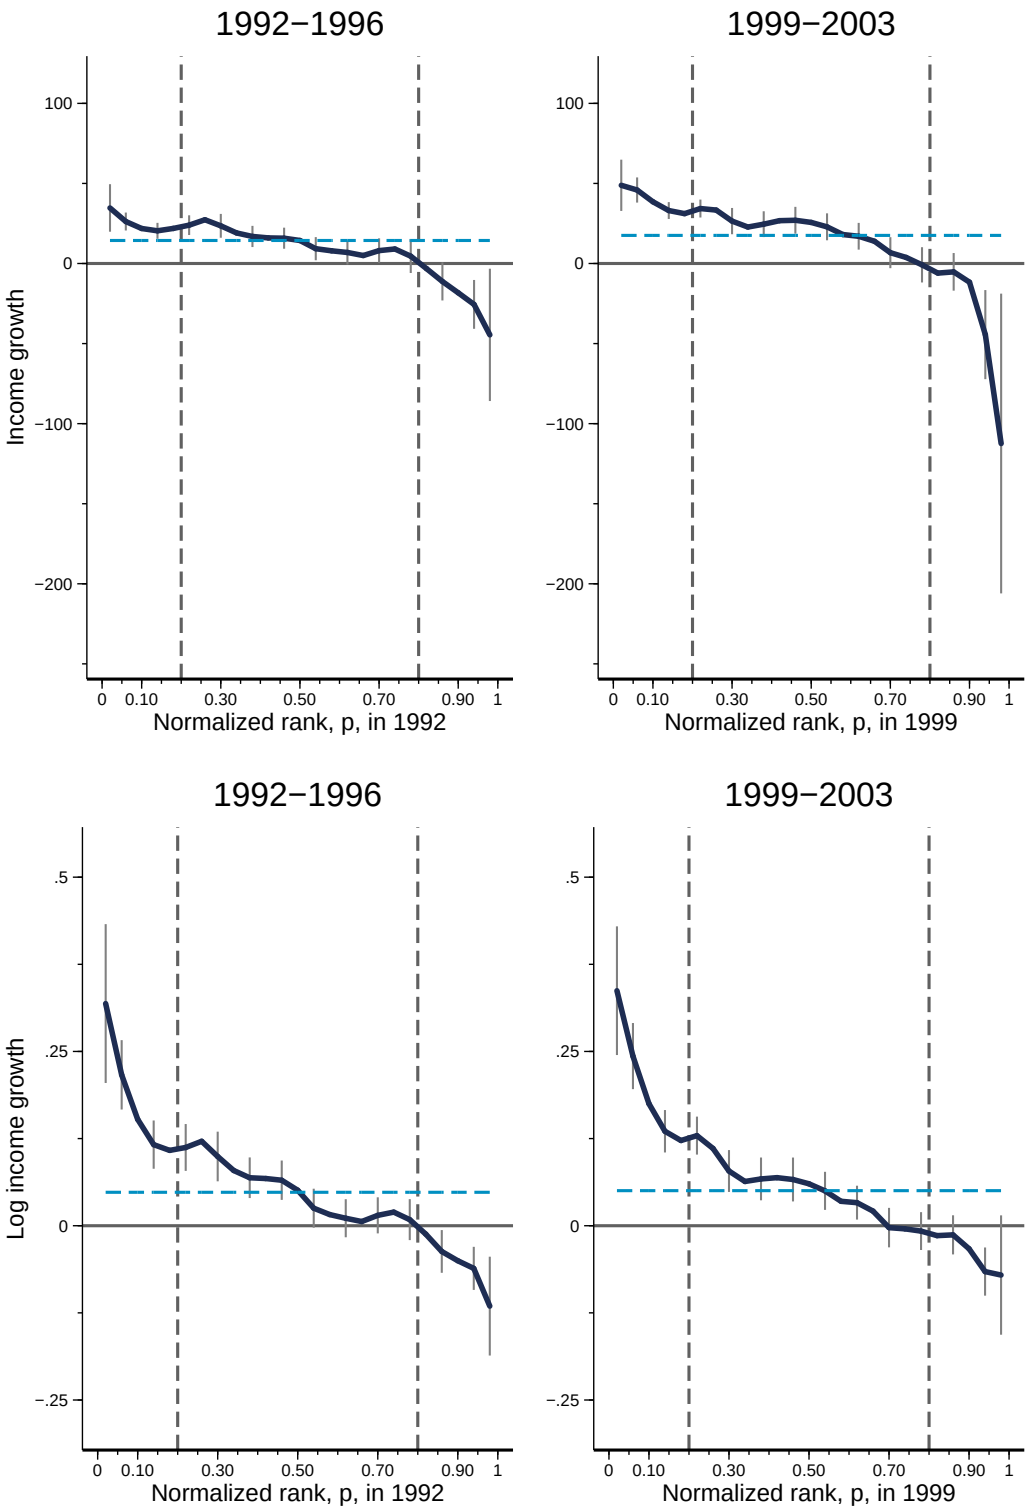

Figure 3. Income mobility profiles: average individual income growth by base year (normalized) rank in the income distribution, 1992–1996 (left) and 1999–2003 (right). Top panel shows absolute growth (change in income); bottom panel shows proportionate growth (change in log income). Horizontal dashed lines mark population averages (irrespective of base year rank). Vertical bars show pointwise two-standard-error variability bands.

We now put numerical flesh on the patterns revealed so far. Table 1 reports, for both periods, average income in the initial and final years, and average income growth over the period for the population as a whole, as well as for individuals who were in the poorest fifth or the richest fifth in the relevant base year. (These two subgroups are composed of individuals to the left and to the right of the vertical bars on Figure 3 respectively.)

Although income growth for the population as a whole increased on average in absolute terms between periods (but not in proportionate terms) slightly, this was well within the bounds of sampling variability. In both periods, those who started in the poorest fifth had larger income gains than those who started in the richest fifth (both in absolute terms and a fortiori in proportionate terms), but the difference was bigger for the 1999–2003 period than for 1992–1996. Note however that is only the absolute increase between periods, from £31 to £47, that is statistically significant. The change in log income growth from +19 to +22 log points may represent sampling variability. Unsurprisingly, the greater gains of those initially poorer were not strong enough to offset the initial income differences between the two groups. The average income among the richest fifth in the final year remained substantially higher (despite their losses) than the average income of the poorest fifth (despite their gains). The last column of Table 1 reveals an interesting symmetry with respect to changes in income group composition over time. In both periods, the fraction of individuals that left the bottom quintile group was almost the same as the fraction of people that left to top quintile group, 35%.

Table 1. Income growth by income groups, 1992–1996 (A) vs. 1999–2003 (B)

	Mean income				Mean growth				Proportion	
	initial		final		income		log income		stayers	
	A	B	A	B	A	B	A	B	A	B
All population	226	264	240	282	14	18	0.05	0.05	1.00	1.00
	(2)	(4)	(3)	(3)	(2)	(3)	(0.01)	(0.01)	(0.00)	(0.00)
Poorest 20%	93	112	123	159	31	47	0.19	0.22	0.67	0.65
	(1)	(2)	(4)	(3)	(3)	(3)	(0.02)	(0.02)	(0.02)	(0.02)
Richest 20%	421	490	404	459	–18	–31	–0.08	–0.08	0.69	0.65
	(6)	(14)	(7)	(10)	(6)	(11)	(0.01)	(0.02)	(0.02)	(0.02)

Notes: A refers to 1992–1996. B refers to 1999–2003. Bootstrap SEs in parentheses. Individual income at year t is the average over years $t-1$, t and $t+1$. Income group membership is based on income in the initial year.

4.3 Inequality trends related to pro-poor growth and reranking

Decompositions, for each of the two periods, of the change in the Gini coefficient into progressivity and reranking components considered are reported in Table 2. The decomposition quantifies the extent of pro-poor growth in terms of individual income progression (the 'gross' effect of income changes) and the associated change in the marginal

distribution (the 'net' effect). The latter is summarized by the change in the Gini coefficient, whereas the former is summarized by P. The difference between the two is the reranking component, R.

The first three columns of Table 2 report estimates of the initial and final Gini coefficients obtained from both the 1992–1996 and 1999–2003 samples, and the change in the Gini derived from these. The next two columns report the estimates of the decomposition components, P and R, and the last two columns show the components expressed as a percentage of initial Gini. As expected from the configurations of the mobility profiles, P is positive in both time periods: growth was pro-poor. But the measure is larger for the second period, 0.075 rather than 0.060, i.e. some 25% higher. Expressed as percentage of the initial year Gini, the increase in P is about the same magnitude. Thus, over the first period, the progressivity of income growth was associated with a small increase in the Gini coefficient. Whereas reranking almost fully offset the progressivity of income growth between 1992 and 1996, this was not the case between 1999 and 2003. Progressivity increased and reranking decreased, and the result was a (small) decline in inequality measured by the Gini coefficient.

Table 2. Gini change, progressivity and reranking

	Gini			Components			
	Initial	Final	Change	R	P	R (as % of initial Gini)	P (as % of initial Gini)
1992–1996	0.290	0.294	0.005	0.065	0.060	22.3	20.6
	(0.004)	(0.005)	(0.004)	(0.004)	(0.005)	(1.4)	(1.6)
1999–2003	0.288	0.270	–0.019	0.056	0.075	19.6	26.1
	(0.007)	(0.005)	(0.006)	(0.003)	(0.007)	(1.2)	(1.8)

Notes: Bootstrap SEs in parentheses. See text for the definition of the decomposition components.

4.4 Income growth for population subgroups

Table 3 reports average income at both periods and average income growth measures for population subgroups defined in terms of family type in the initial year. The top panel summarizes the experience of all subgroup members; the bottom panel focuses on the subgroup members in the poorest fifth of the base year distribution. The latter breakdowns allow us to focus in more detail on groups of particular interest, though at the cost of small cell sizes and less precision to the estimates.

Looking at the subgroups as a whole (top panel), the calculations reveal that for each family type, proportionate income growth was much the same in both periods. It is only when growth is defined in absolute terms that subgroup differences appear – being most apparent for couples with children and singles with children. We note that families with

children were a specific target of government policy. On the other hand, we also observe that absolute income growth declined between periods for pensioners, another target group. Part of this apparent contradiction is resolved when we examine the bottom panel of Table 3. When the focus is on the individuals in the poorest fifth, the calculations reveal a large increase in income growth between period in both absolute terms (from £12 to £28) and in proportionate terms (about 6% to about 14%). The increase in income growth for the poorest lone parents is also substantial between the periods, rising from £12 to £28). Both these results are consistent with policies directed at poverty reduction rather than inequality reduction.

It is important to emphasize that subgroup status refers to status in the initial time period. (Strictly speaking, it is status in the middle year of the three years used to form the three-year income average.) Between periods, the degree of fluidity was fairly stable: at least 20% of each non-pensioner group belonged to a different family type four years later.

To help link these changes in fortunes to the configurations of mobility profiles shown in Figure 2 for the population as a whole, Figures 4 and 5 show mobility profiles for each subgroup separately. These subgroup profiles are shown in the top section of each plot. In the bottom section of each plot, we show the subgroup's size relative to the total population ('pop. proportion'; right-hand scale) at each base year normalized rank of the population income distribution. The graphs show that pensioners are concentrated in the bottom third of the distribution. The poorest fifth is dominated by representatives from families with children, both lone parent families and couple families. Those most likely to be found in the higher income ranges are childless couples. The population mobility profile value at each normalized rank is the weighted sum over the subgroup values at that rank, where the weights are the subgroup proportions.

Figures 4 and 5 tell a similar story to Table 3. Within each period 1992–1996 and 1999–2003, most groups benefited from pro-poor income growth, especially groups relatively concentrated towards the lower income groups. However there was a distinct increase in pro-poorness, at least if one focuses on the bottom of the distribution and on growth defined in absolute terms. There was an increase in pro-poor growth that benefited those family types concentrated at the bottom, notably families with children and pensioners. These patterns are consistent with the aims of the Labour government to reduce pensioner and child poverty.

Table 3. Income growth by family type, 1992–1996 (A) and 1999–2003 (B)

	Sample size		Mean income initial		Mean income final		Mean growth income		Mean growth log income		Proportion stayers	
	A	B	A	B	A	B	A	B	A	B	A	B
Among all individuals:												
Pensioners	1002	978	174	209	182	211	8	2	0.01	0.02	1.00	1.00
	(17)	(18)	(4)	(5)	(5)	(4)	(2)	(3)	(0.01)	(0.01)	(0.00)	(0.00)
Couples with children	2845	2594	213	252	233	282	19	30	0.08	0.07	0.79	0.77
	(50)	(48)	(4)	(5)	(4)	(5)	(3)	(4)	(0.01)	(0.01)	(0.01)	(0.01)
Couples without children	1426	1450	302	341	304	342	2	0	–0.01	–0.01	0.69	0.68
	(24)	(27)	(5)	(8)	(6)	(7)	(5)	(6)	(0.01)	(0.01)	(0.02)	(0.02)
Single with children	380	332	138	167	161	215	23	47	0.17	0.16	0.68	0.69
	(15)	(15)	(7)	(8)	(8)	(8)	(5)	(7)	(0.03)	(0.03)	(0.04)	(0.04)
Single without children	884	739	251	301	277	324	27	23	0.07	0.07	0.75	0.80
	(17)	(14)	(6)	(17)	(9)	(9)	(5)	(11)	(0.02)	(0.02)	(0.02)	(0.01)
Among poorest 20% in base year:												
Pensioners	336	336	97	115	110	143	12	28	0.06	0.14	0.81	0.76
	(11)	(11)	(1)	(2)	(2)	(3)	(2)	(2)	(0.01)	(0.02)	(0.03)	(0.03)
Couples with children	468	357	90	112	134	164	43	52	0.27	0.25	0.52	0.45
	(20)	(20)	(2)	(3)	(9)	(6)	(8)	(7)	(0.04)	(0.04)	(0.04)	(0.05)
Couples without children	63	75	90	111	152	148	62	37	0.31	0.23	0.12	0.32
	(6)	(6)	(5)	(4)	(14)	(8)	(14)	(8)	(0.08)	(0.05)	(0.07)	(0.08)
Single with children	200	166	88	106	114	171	26	64	0.20	0.24	0.55	0.52
	(11)	(11)	(2)	(4)	(6)	(8)	(5)	(9)	(0.04)	(0.05)	(0.06)	(0.06)
Single without children	88	75	91	109	140	201	49	92	0.30	0.40	0.39	0.30
	(5)	(5)	(3)	(4)	(8)	(19)	(8)	(19)	(0.05)	(0.07)	(0.06)	(0.06)

Notes: A refers to 1992–1996. B refers to 1999–2003. Bootstrap SEs in parentheses. Individual income at year t is the average over years $t-1$, t and $t+1$. Group membership is based on status in the initial year.

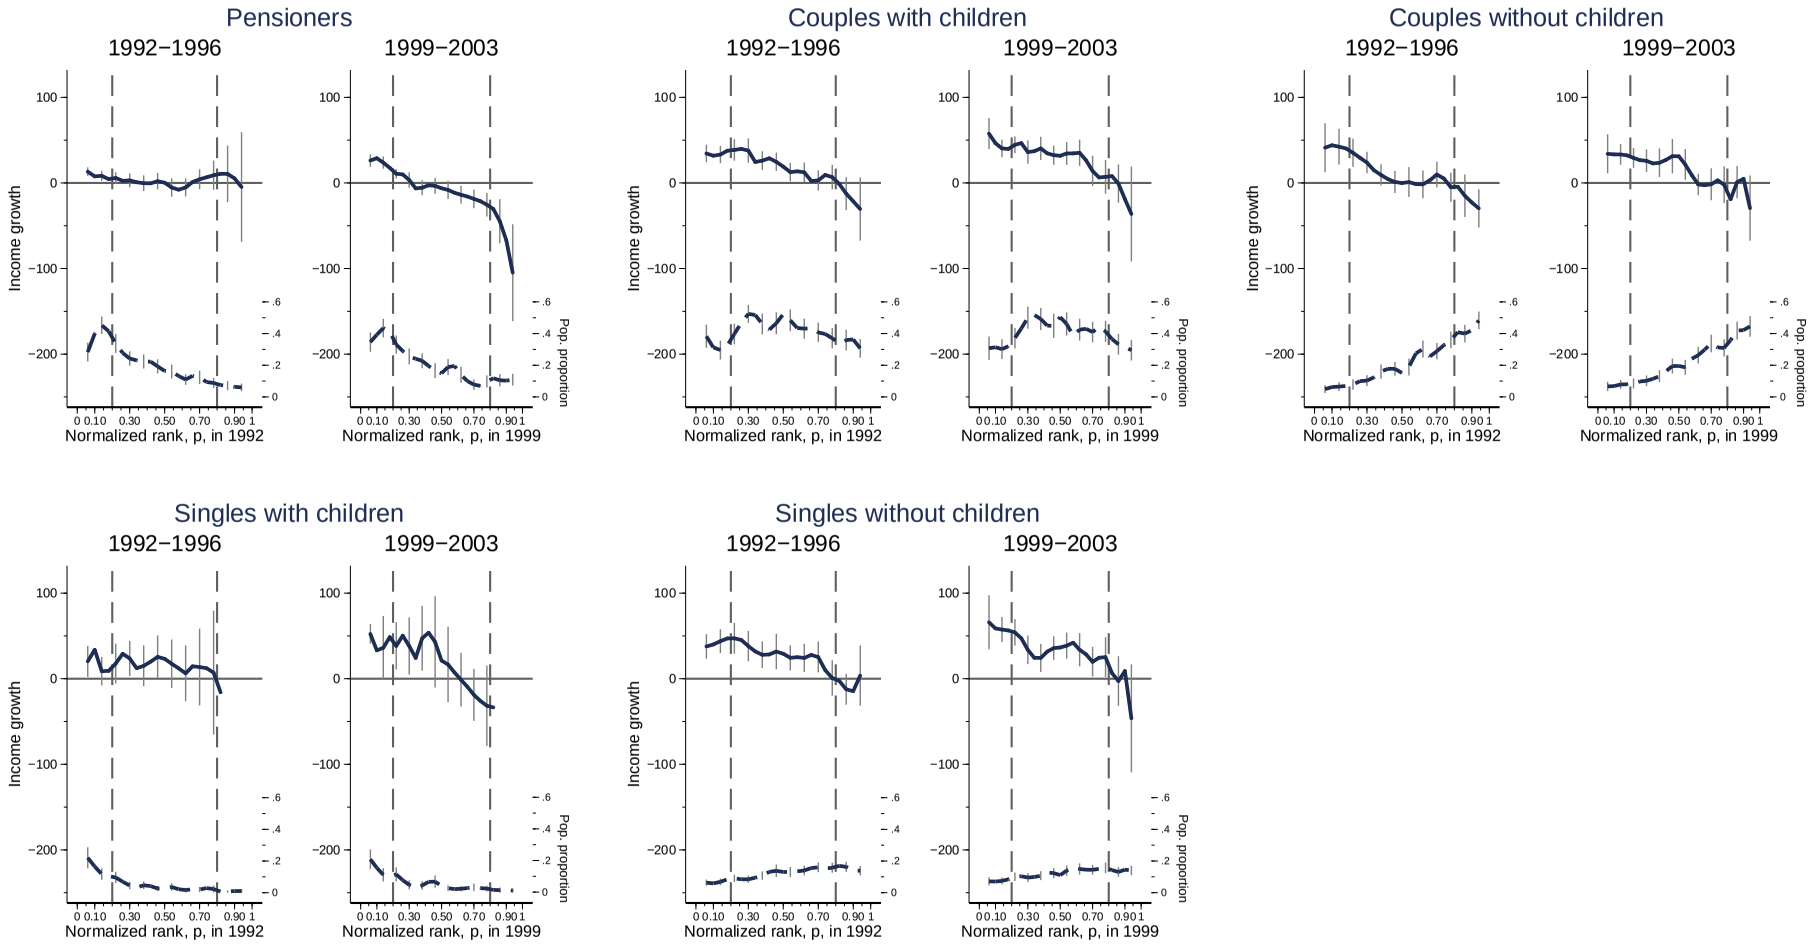

Figure 4. Income mobility profiles by family type (absolute income growth), and subgroup population proportions. The top section of each plot shows average individual income growth among each subgroup by base year position in the population income distribution. The dotted line in the bottom section of each plot shows the population proportion of the family type at each normalized rank of the overall income distribution in the initial year. Vertical bars show two-standard-error variability bands.

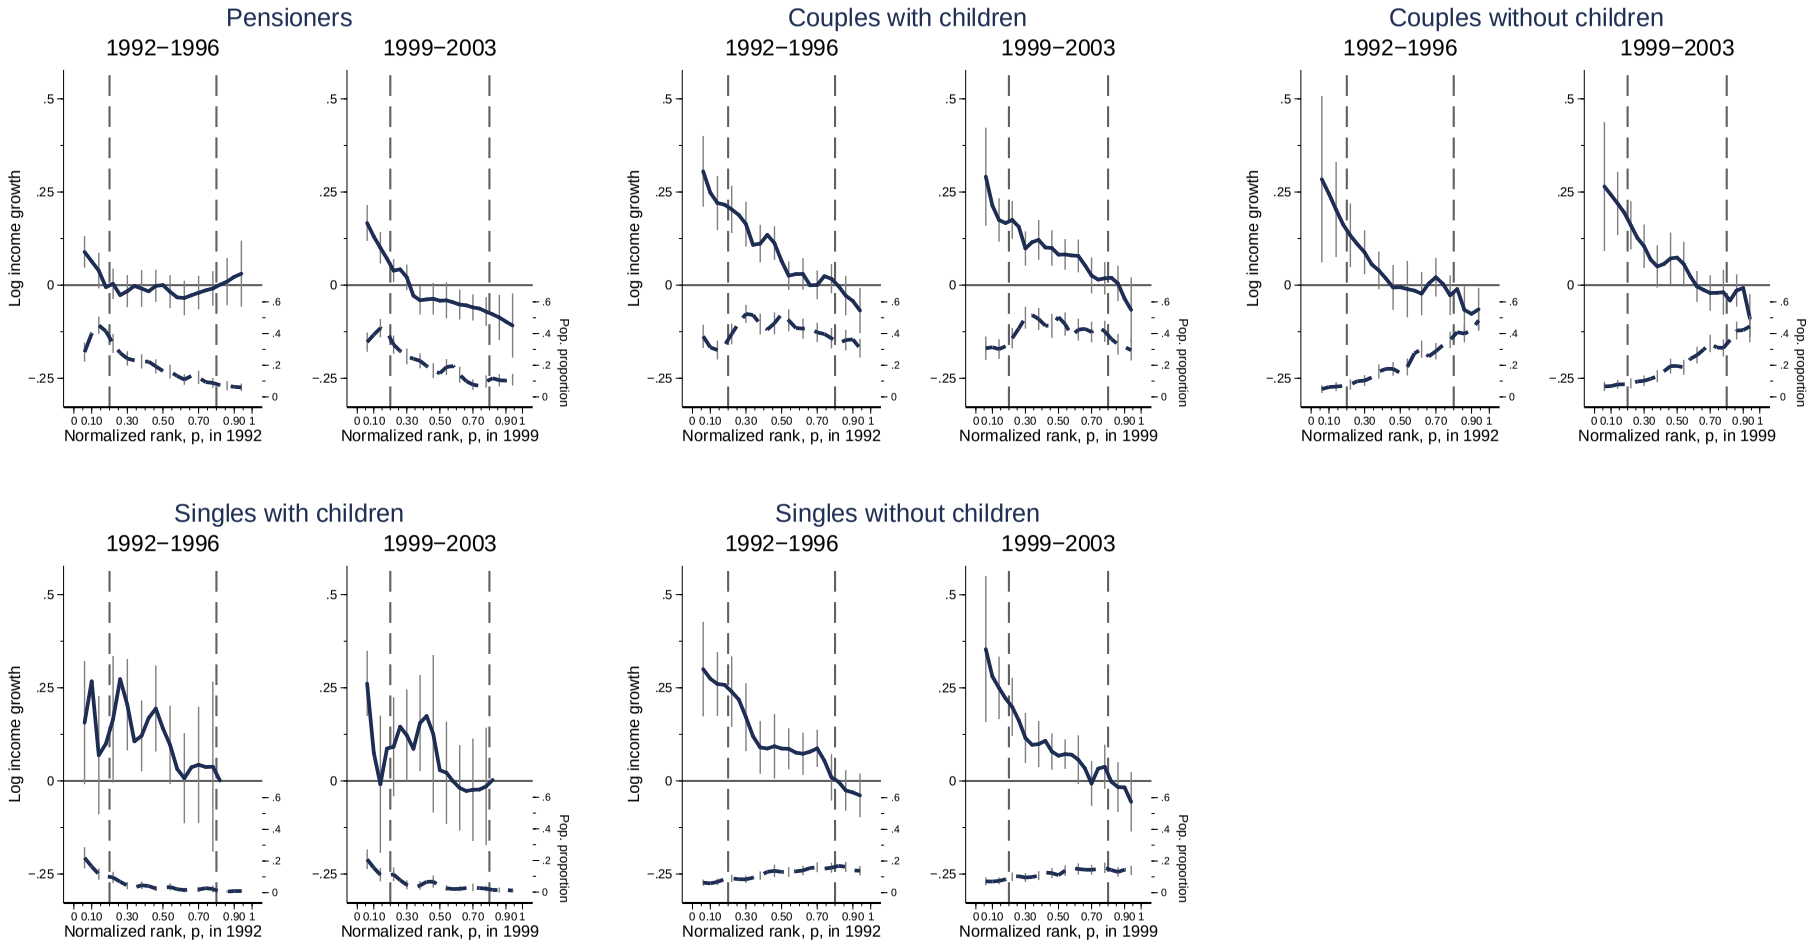

Figure 5. Income mobility profiles by family type (proportionate income growth), and subgroup population proportions. The top section of each plot shows average individual log income growth among each subgroup by base year position in the population income distribution. The dotted line in the bottom section of each plot shows the population proportion of the family type at each normalized rank of the overall income distribution in the initial year. Vertical bars show two-standard-error variability bands.

5 Summary and Conclusions

We have argued that assessments of who has gained most from income growth over time should use longitudinal perspectives to complement conventional cross-sectional approaches. Assessments of who is getting better off over time that summarize changes in the incomes of particular groups, e.g. changes in income for the poor or the rich, or for subgroups such as lone parent families and other families with children, ignore the fact that these groups change composition over time: the same individuals are not being compared. To be sure, there is interest in whether the income of the poorest person (whoever that person is and regardless of their previous position), increases over time. But to assess whether this year's poor (or rich) are gainers or losers, one has to track the fortunes of individuals (not groups) using longitudinal data.

We have developed methods for describing and summarizing patterns of income growth from a longitudinal perspective, together with estimates of sampling variability that account for the dependent nature of the sampling that is inherent in this analysis.

Using data from the British Household Panel Survey, we have compared patterns of individual income growth over the period 1992–1996 with patterns over the period 1999–2003. We have shown that conclusions about whether patterns of income growth became more pro-poor depend on the perspective used and the definition of income growth. According to conventional analysis using a cross-sectional perspective, income growth became more pro-poor using both absolute and proportionate growth definitions. However, from a longitudinal perspective, an increase in progression of income growth is most clearly apparent using absolute income growth definition. The picture of greater progressivity is more muted when viewed in terms of proportionate changes. Focusing on particular groups shows that there has been increase in income growth among groups specifically targeted by the Labour government's anti-poverty programmes, namely families with children and pensioners but, again, this is apparent in terms of absolute rather than proportionate income growth. The sensitivity of the conclusions to how income growth is defined raises intriguing questions about how changes in income distribution over time should be assessed.

References

Bardasi, E., Jenkins, S. P. & Rigg, J. A. (1999), 'Documentation for derived current and annual net household income variables, BHPS waves 1–7', ISER Working Paper 99-25, Institute for Social and Economic Research, University of Essex, UK, downloadable from `http://www.iser.essex.ac.uk/pubs/workpaps/pdf/1999-25.pdf`.

Bennett, F. & Millar, J. (2005), 'Making work pay?', *Benefits*, **13**:28–33.

Böheim, R. & Jenkins, S. P. (2006), 'A comparison of current and annual measures of income in the British Household Panel Survey', *Journal of Official Statistics*, **22**:733–758.

Brewer, M., Muriel, A., Phillips, D. & Sibieta, L. (2008), 'Poverty and inequality in the UK: 2008', IFS Commentary 104, Institute for Fiscal Studies, London, UK, downloadable from `http://www.ifs.org.uk/comms/comm105.pdf`.

Brewer, M. & Shephard, A. (2004), *Has Labour made work pay?*, Joseph Rowntree Foundation, downloadable from `http://www.jrf.org.uk/bookshop/eBooks/1859352626.pdf`.

Cecchi, D. & Dardanoni, V. (2003), 'Mobility comparisons: Does using different measures matter?', in *Research on Economic Inequality*, vol. 9, JAI Press, 113–145.

Cleveland, W. S. (1979), 'Robust locally weighted regression and smoothing scatterplots', *Journal of the American Statistical Association*, **74**(368):829–836.

Cleveland, W. S. & Grosse, E. (1991), 'Computational methods for local regression', *Statistics and Computing*, **1**(1):47–62.

Department for Work and Pensions (2008), *Households Below Average Income. An analysis of the income distribution 1994/95-2006/07*, Department for Work and Pensions, London, downloadable from `http://www.dwp.gov.uk/asd/hbai/hbai2007/contents.asp`.

Donaldson, D. & Weymark, J. A. (1980), 'A single-parameter generalization of the Gini indices of inequality', *Journal of Economic Theory*, **22**:67–86.

— (1983), 'Ethically flexible Gini indices for income distributions in the continuum', *Journal of Economic Theory*, **29**:353–358.

Fields, G. S., Cichello, P. L., Freije, S., Menéndez, M. & Newhouse, D. (2003), 'For richer or for poorer? Evidence from Indonesia, South Africa, Spain, and Venezuela', *Journal of Economic Inequality*, **1**(1):67–99.

Fields, G. S., Leary, J. B. & Ok, E. A. (2002), 'Stochastic dominance in mobility analysis', *Economics Letters*, **75**:333–339.

Gottschalk, P. & Danziger, S. (2001), 'Income mobility and exits from poverty of American children', in B. Bradbury, S. P. Jenkins & J. Micklewright (eds.), *The Dynamics of Child Poverty in Industrialised Countries*, chap. 5, Cambridge University Press, Cambridge, USA, 135–153.

Grimm, M. (2007), 'Removing the anonymity axiom in assessing pro-poor growth', *Journal of Economic Inequality*, **5**(2):179–197.

Hastie, T. & Loader, C. (1993), 'Local regression: Automatic kernel carpentry', *Statistical Science*, **8**(2):120–129.

Jenkins, S. P. & Cowell, F. A. (1994), 'Dwarfs and giants in the 1980s: trends in the UK income distribution', *Fiscal Studies*, **15**:99–118.

Jenkins, S. P. & Van Kerm, P. (2006), 'Trends in income inequality, pro-poor income growth and income mobility', *Oxford Economic Papers*, **58**(3):531–548.

Levy, H., Zantomio, F., Sutherland, H. & Jenkins, S. P. (2006), *Documentation for Derived Current and Annual Net Household Income Variables, BHPS waves 1-14*, Institute for Social and Economic Research, University of Essex, UK, downloadable from `http://www.data-archive.ac.uk/findingData/snDescription.asp?sn=3909`.

Lynn, P. (ed.) (2006), *Quality Profile: British Household Panel Survey. Version 2.0: Waves 1 to 13: 1991-2003.*, Institute for Social and Economic Research, University of Essex, Colchester, UK, downloadable from `http://www.iser.essex.ac.uk/ulsc/bhps/quality-profiles/BHPS-QP-01-03-06-v2.pdf`.

Mahalanobis, P. C. (1960), 'A method of fractile graphical analysis', *Econometrica*, **28**(2):325–351.

O'Neill, D. & Van Kerm, P. (2008), 'An integrated framework for analysing income convergence', *The Manchester School*, **76**(1):1–20.

Pen, J. (1971), *Income Distribution*, Penguin Books.

Ravallion, M. & Chen, S. (2003), 'Measuring pro-poor growth', *Economics Letters*, **78**(1):93–99.

Sefton, T. & Sutherland, H. (2005), 'Inequality and poverty under New Labour', in J. Hills & K. Stewart (eds.), *A More Equal Society*, chap. 11, Policy Press, Bristol, UK, 231–249.

Sen, B. (2005), 'Estimation and comparison of fractile graphs using kernel smoothing techniques', *Sankhyã, The Indian Journal of Statistics*, **67**(2):305–334.

Son, H. H. (2004), 'A note on pro-poor growth', *Economics Letters*, **82**:307–314.

Van Kerm, P. (2006), 'Comparisons of income mobility profiles', IRISS Working Paper 2006-03, CEPS/INSTEAD, Differdange, Luxembourg, downloadable from `http://www.ceps.lu/iriss/documents/irisswp58.pdf`.

Yitzhaki, S. (1983), 'On an extension of the Gini inequality index', *International Economic Review*, **24**(3):617–628.

Appendix A: Estimation

All quantiles and measures of inequality, progressivity, and reranking, were estimated using sample analogues for measures defined with continuous distributions.

Growth incidence curves require estimation of percentiles only. In contrast, income mobility profiles are conditional expectations and require 'smooth' estimators of these. The statistics reported in the paper were estimated using local linear regression methods as these have good behaviour near the boundaries of the support of the data. Specifically, we used the locally weighted regression (LOESS) introduced by Cleveland (1979) (see also Cleveland & Grosse, 1991, Hastie & Loader, 1993). The technique involves determining local neighbourhoods around each of a series of p values and using the sample observations falling in these neighbourhoods to estimate the mobility profile on the grid points using (locally) weighted least squares regression. We estimated the profiles at 25 equally spaced points 0.02, 0.06, ..., 0.98. Local neighbourhoods were defined as $p \pm 0.075$ so that approximately 15 percent of the sample fell in each neighbourhood. To protect the estimates of the mobility profiles from being unduly influenced by outliers in the distribution of income changes conditional on p), we applied the 'robust' LOESS estimator, i.e. down-weighting outlier observations in the manner described by Cleveland (1979).

The sampling variability of all estimates was assessed using standard errors derived from bootstrap resampling methods. A block bootstrap procedure was implemented to reflect the dependence of the sample over time resulting from the longitudinal nature of the data. Resampling consisted of sampling with replacement from the sample of households interviewed in wave 1 of the BHPS and their descendants and split-off households. The full response history over waves 1–14 of all members of the selected households, as well as all respondents that later joined these households, was then selected to form a bootstrap replicate of the panel across waves 1 to 14. Sub-samples for the periods 1992–1996 and 1999–2003 were then drawn from these bootstrap replicates according to the selection rules described in Section 3. This procedure ensures that the sampling dependence of the two sub-samples was preserved. Resampling in the first wave of the survey was done independently within sampling strata. All standard error estimates are based on 500 bootstrap replications. The vertical bars in all graphs show pointwise two-standard-error variability bands.

Sample weights were used to compute all estimates. Our two longitudinal samples were weighted with the BHPS cross-sectional enumerated individual weights of year $t+\tau$. We did not use the BHPS longitudinal weights because our analysis samples include sample joiners ('temporary sample members'), for whom there are no longitudinal weights available: for wave $t + \tau$, the weights exist only for sample members (and their children) who are respondents at every wave from wave 1 to wave $t + \tau$.

Because our longitudinal samples t–$t+\tau$ include only a subsample of the year $t+\tau$ cross-section sample, there is no guarantee that the cross-section weights that we use fully adjust the longitudinal sub-samples to the British population characteristics. Reassuringly, however, estimates of key income distribution statistics from 'cross-section' versions of three-year-average incomes samples are close to estimates from our sub-samples: see Table 4.

Table 4. Estimates of summary statistics from cross-section samples and longitudinal sub-samples

	1992		1996		1999		2003	
	XS	LS1	XS	LS1	XS	LS2	XS	LS2
Sample size	8537	6537	8646	6537	8460	6093	7622	6093
	(65)	(61)	(72)	(61)	(76)	(63)	(76)	(63)
Mean income	225	226	240	240	260	264	283	282
	(2)	(2)	(2)	(3)	(3)	(4)	(3)	(3)
10th percentile	94	96	98	101	113	117	136	136
	(2)	(2)	(2)	(2)	(1)	(2)	(2)	(2)
Median income	199	200	215	216	231	233	253	252
	(3)	(3)	(3)	(3)	(3)	(3)	(3)	(3)
90th percentile	383	383	411	409	422	427	459	454
	(5)	(6)	(6)	(6)	(6)	(6)	(7)	(7)
Gini coefficient	0.295	0.290	0.300	0.294	0.288	0.288	0.272	0.270
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.007)	(0.005)	(0.005)
P90/P10 ratio	4.07	3.98	4.18	4.06	3.74	3.65	3.38	3.35
	(0.09)	(0.10)	(0.10)	(0.11)	(0.07)	(0.08)	(0.07)	(0.07)

Notes: XS refers to the sample of respondents with valid three-year-smoothed individual income at year t. LS1 is our 1992–1996 longitudinal sub-sample and LS2 is our 1999–2003 longitudinal sub-sample. See text for details on sample construction. Bootstrap SEs reported in parentheses.

Appendix B: BHPS estimates of trends in marginal income distributions

This Appendix presents summary statistics about trends in the marginal income distributions derived from the BHPS, and compares them with trends derived from the *Households Below Average Income* series used in the official income distribution statistics, and derived from the annual cross-sectional *Family Resources Survey*. The BHPS samples used for this exercise differ from those used in the main text: the sample for each year consists of all individuals with a non-missing income for that year. The definition of 'income' is the same but there is no three-year averaging.

Figures 6 and 7 summarize BHPS-based income distribution trends in Britain in the period 1992–2004. Our estimates based on BHPS data reveal trends that are consistent

with those reported by *Households Below Average Income* and non-official analyses such as Brewer *et al.* (2008). There was positive income growth for each income group over the period as a whole, with a slow-down in the rate particularly for the richest percentiles at the beginning of the 2000s. The proportion of the population with an income below 60% of median income (the 'proportion poor') declined slightly over the period as a whole after a rise at the start of the 1990s, and changed little in the 2000s. Inequality, whether measured using the Gini coefficient, the ratio of the 90$^{\text{th}}$ percentile to the 10$^{\text{th}}$ percentile ('P90/P10'), or the ratio of the income share of the richest fifth to the income share of the poorest fifth ('S80/S20'), changed little over most of the 1990s except over the period 1998–2001. The main difference between the trends derived from BHPS data and *Households Below Average Income* data is that the former indicate a steeper fall in poverty and inequality in the period 1997–2001 than do the latter (cf. Brewer *et al.* (2008), Chapter 3).

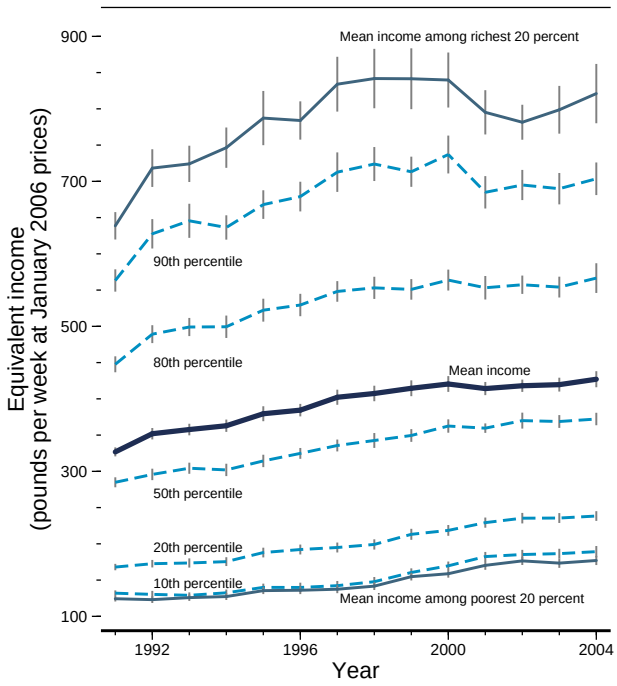

Figure 6. Growth in mean income, selected percentiles, and mean income among the poorest 20% and the richest 20%. Source: authors' calculations from BHPS data (see main text). Vertical bars show two-standard-error variability bands.

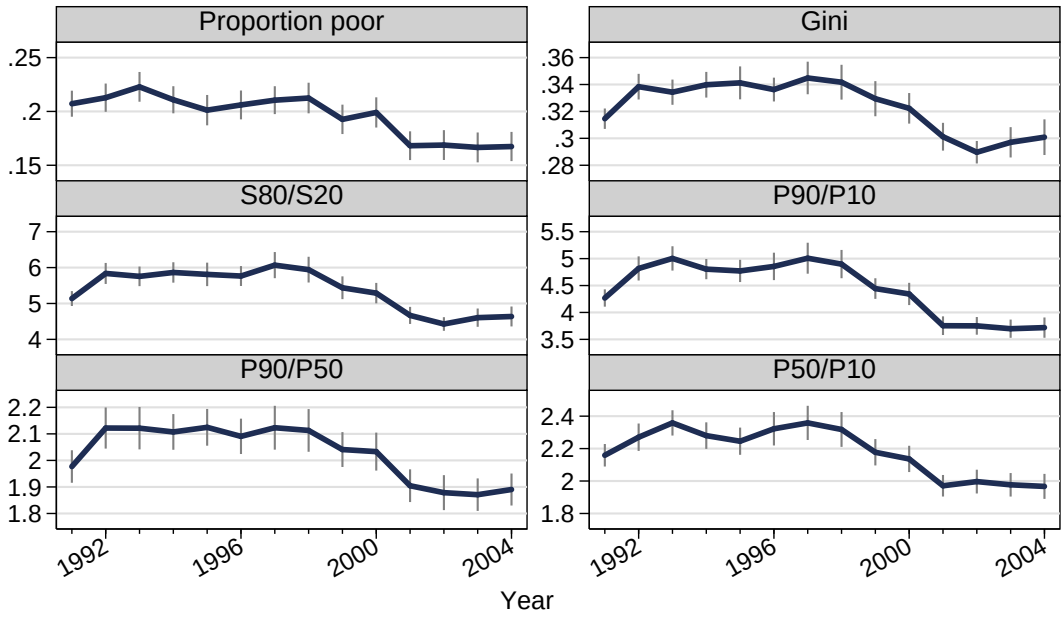

Figure 7. Trends in poverty and inequality. Source: authors' calculations from BHPS data (see main text). Vertical bars show two-standard-error variability bands.

Appendix C: Growth incidence curves, income mobility profiles and Gini change decompositions for all t–$t+4$ periods covered by the BHPS

This Appendix presents growth incidence curves, income mobility profiles and Gini change decompositions for every t–$t+4$ period covered by our BHPS data for survey years 1991–2004. By presenting estimates based on a moving four-year window, we aim to illuminate the extent to which the results presented in the main text for two particular 'Conservative' and 'Labour' periods are sensitive to use of differently-defined periods. In addition to providing estimates based on three-year-averaged income (as in the main text), we also present estimates based on single-year incomes (i.e. without longitudinal averaging).

Figure 8 shows growth incidence curves for the case with three-year averaged income, and Figure 9 shows the corresponding plots for the case with single-year income. The results may be compared with Figure 2.

Income mobility profiles are shown in Figure 10 (three-year-averaged income) and Figure 11 (single-year income). Compare the estimates with Figure 3.

Estimates of the decomposition of changes in the Gini coefficient are shown in Figure 12 for the three-year-averaged income case, and in Figure 13 for the single-year income case. Compare the estimates with Table 2.

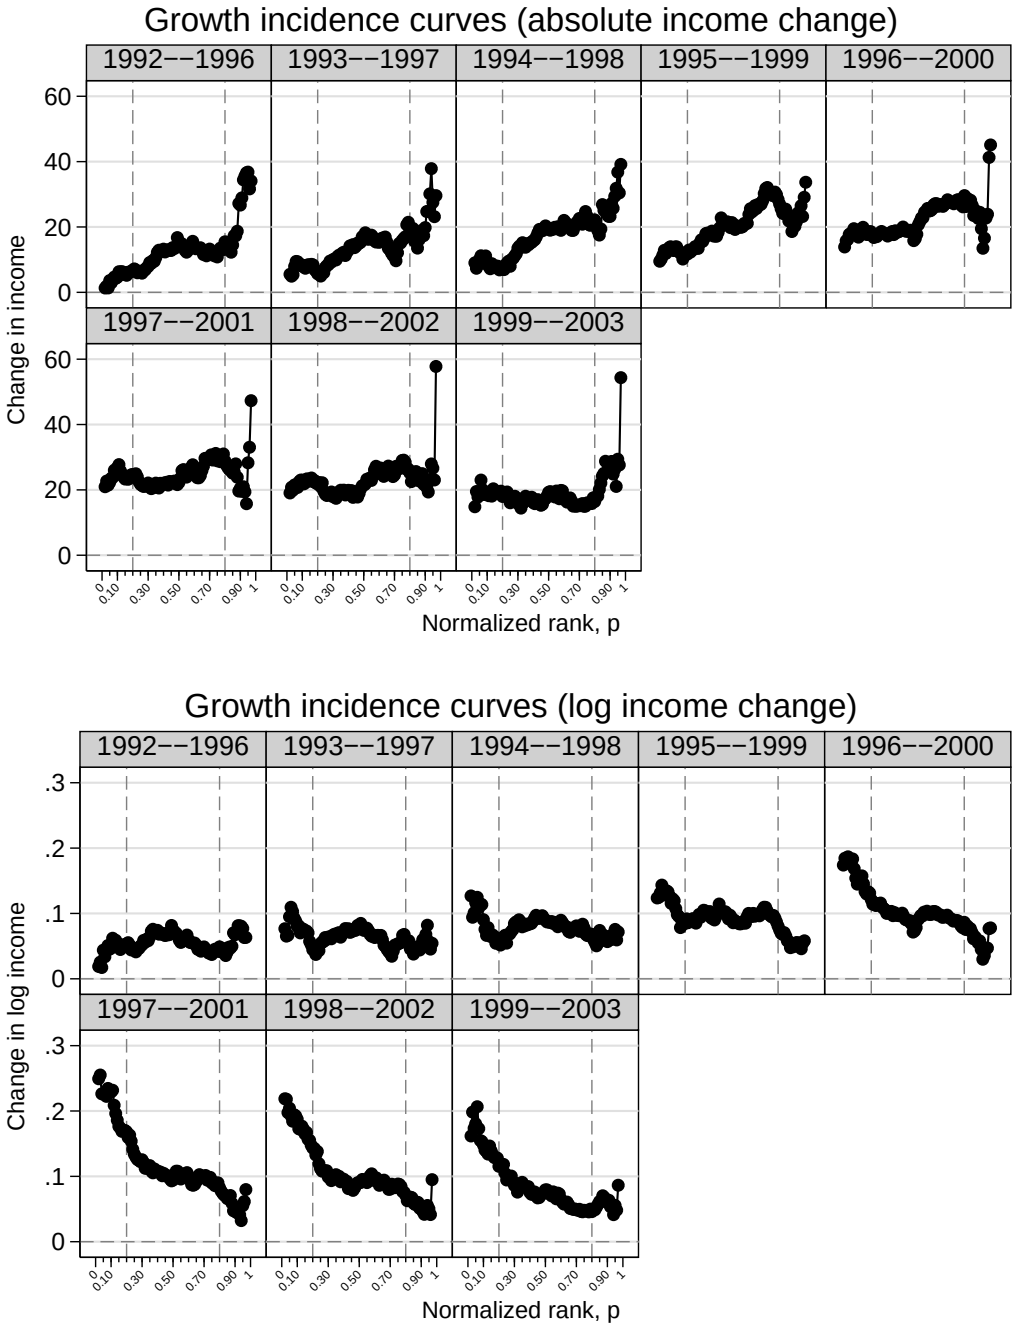

Figure 8. Growth incidence curves: top panel shows absolute growth (change in income); bottom panel shows proportionate growth (change in log income). Based on three-year-average incomes.

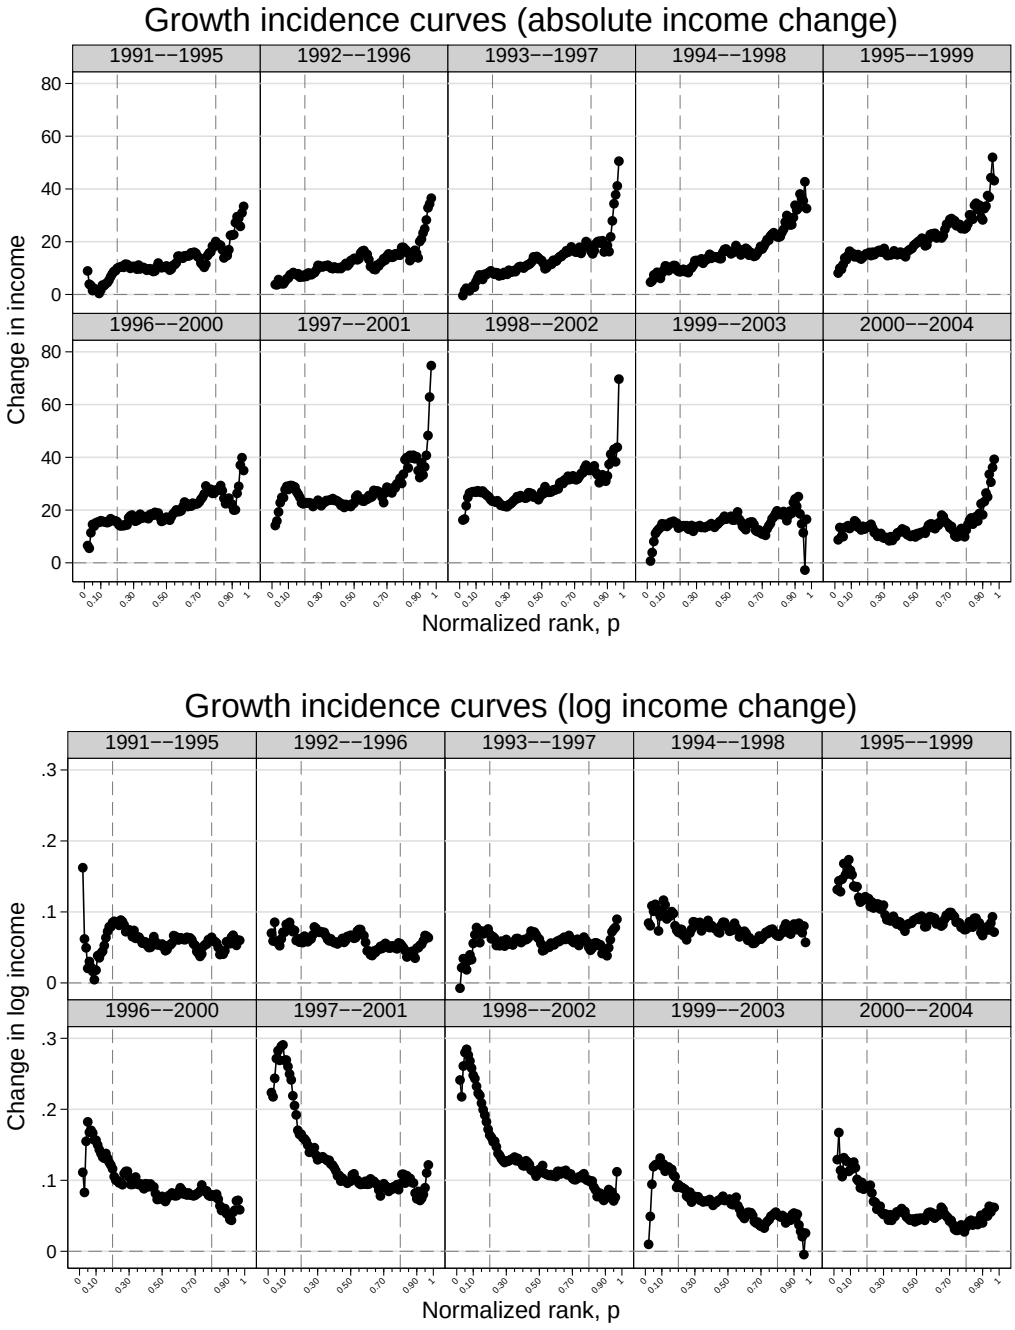

Figure 9. Growth incidence curves: top panel shows absolute growth (change in income); bottom panel shows proportionate growth (change in log income). Based on single-year incomes.

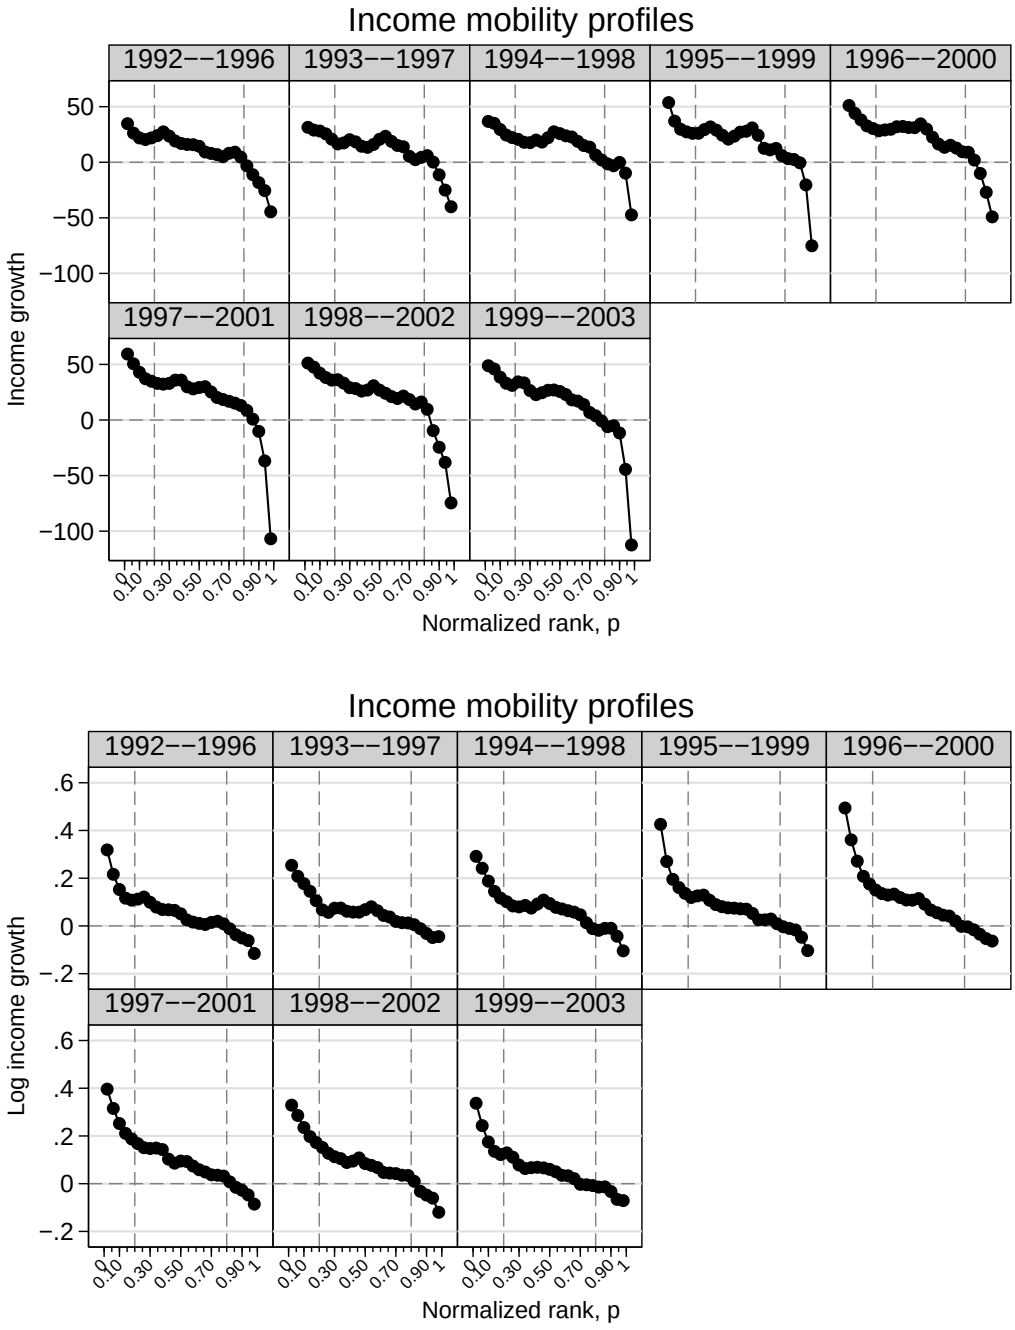

Figure 10. Income mobility profiles: top panel shows absolute growth (change in income); bottom panel shows proportionate growth (change in log income). Based on three-year-average incomes.

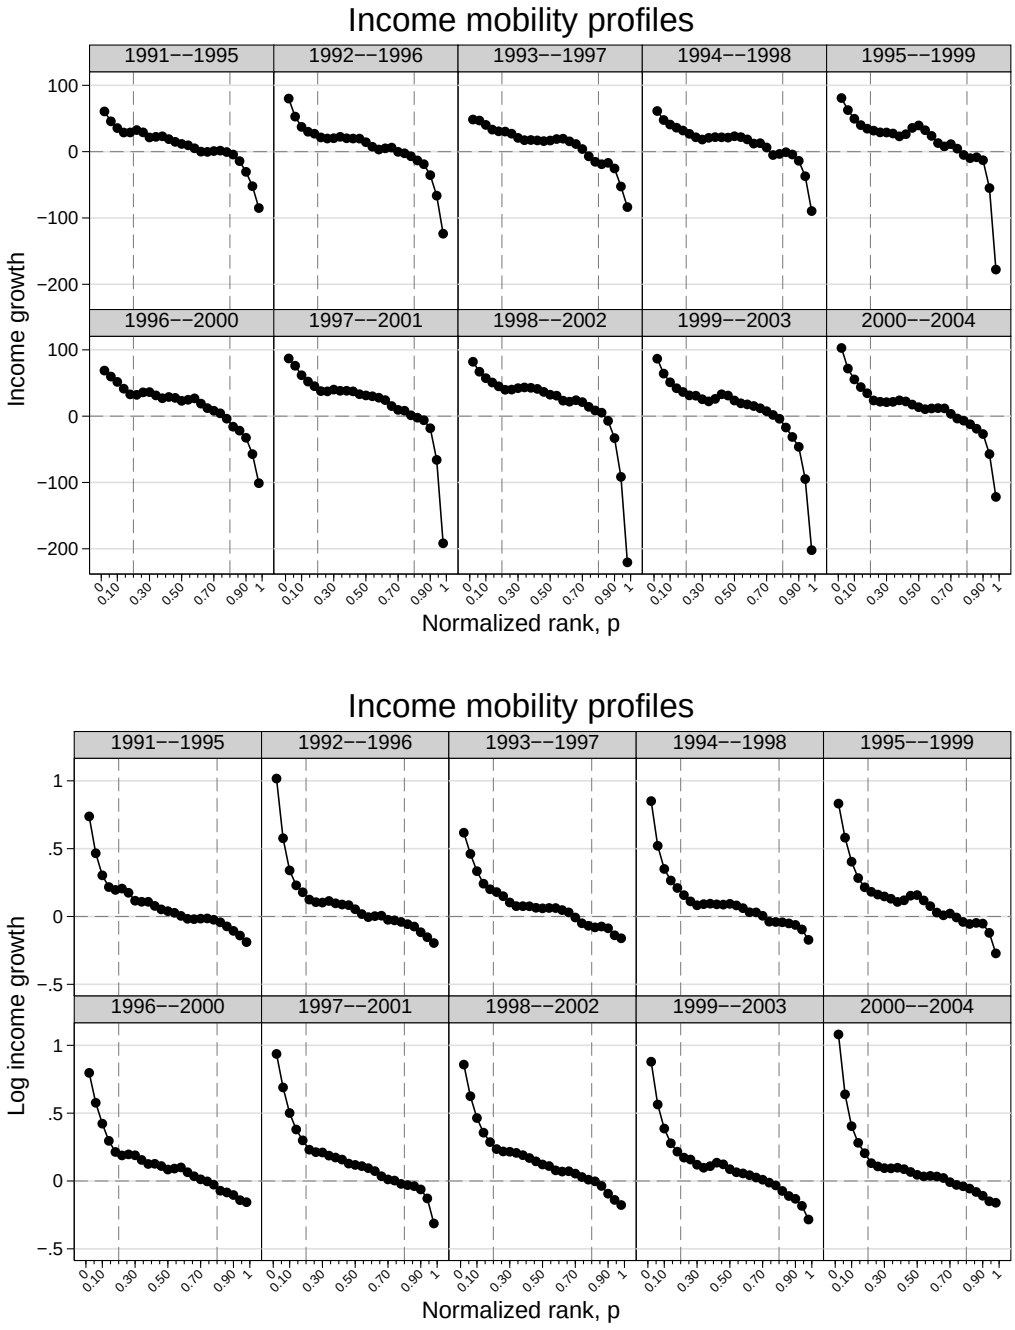

Figure 11. Income mobility profiles: top panel shows absolute growth (change in income); bottom panel shows proportionate growth (change in log income). Based on single-year incomes.

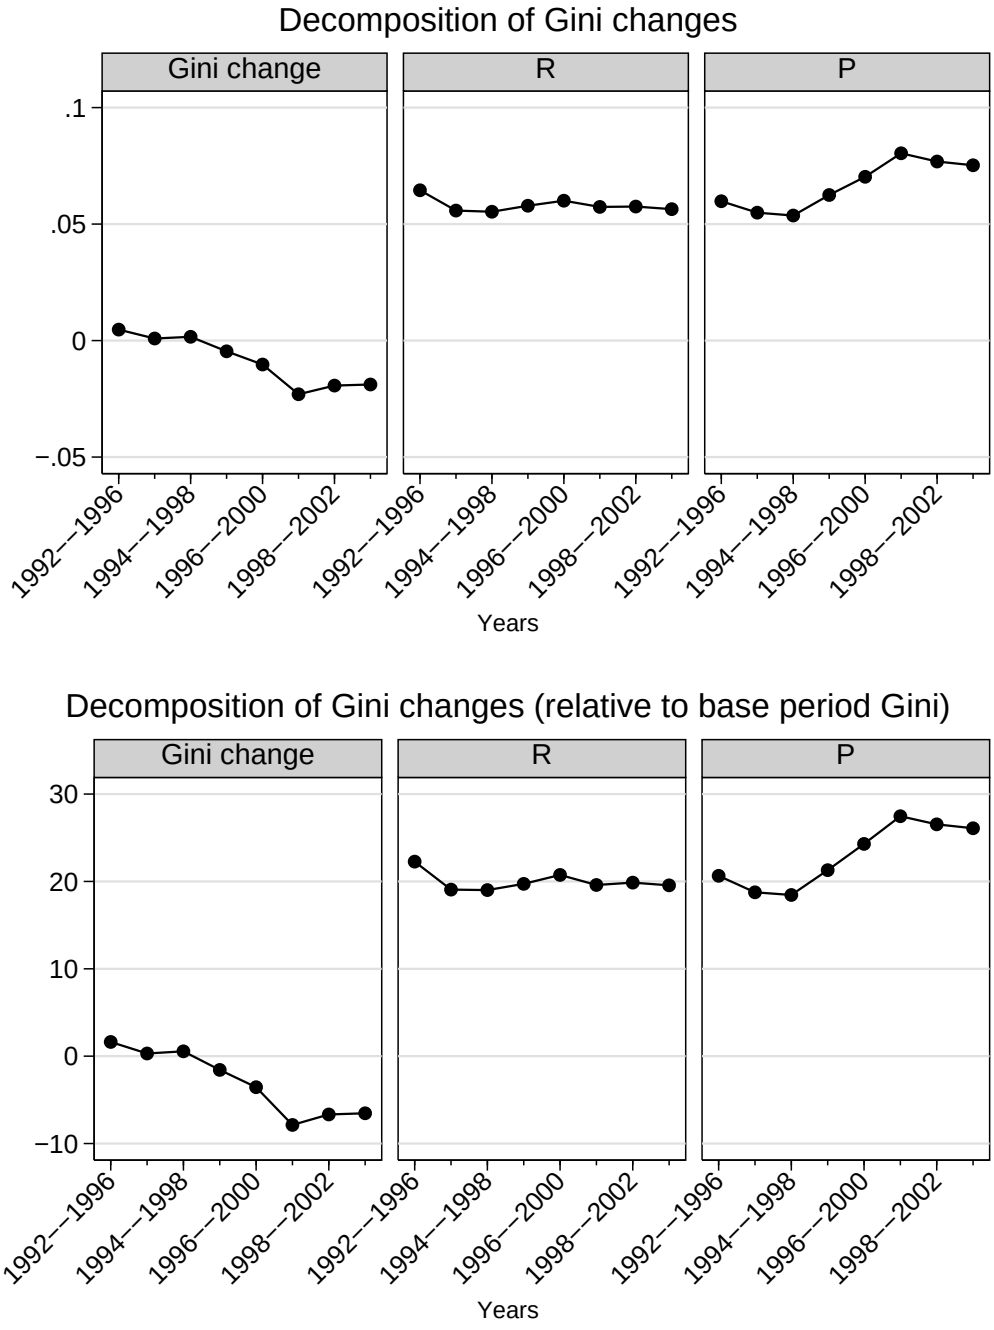

Figure 12. Decompositions of changes in Gini coefficients: top panel shows estimates of the change in the Gini and of the R and P components; bottom panel shows estimates expressed as a percentage of base period Gini. Based on three-year-average incomes.

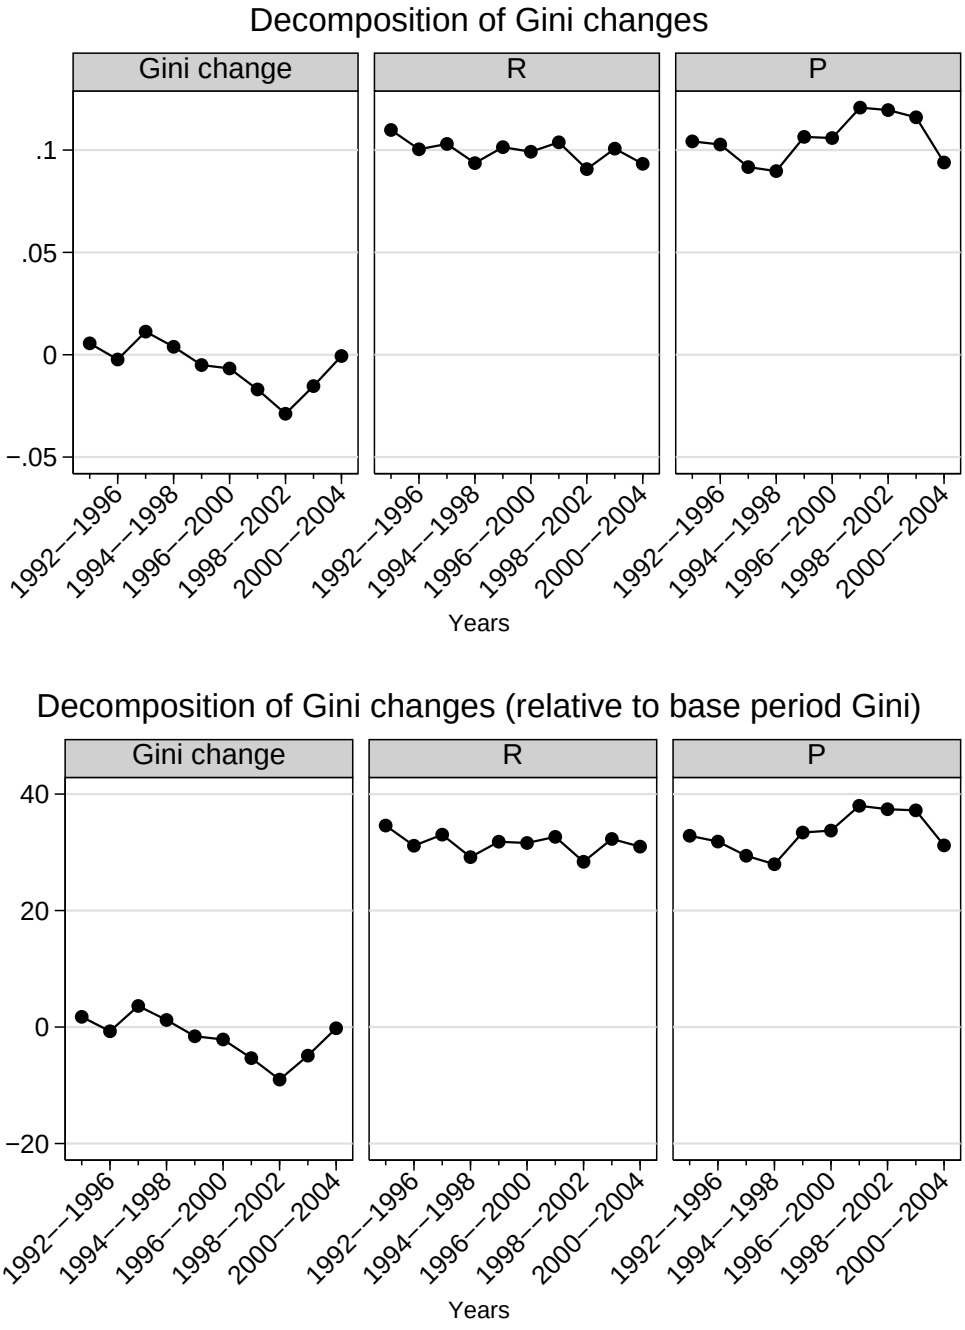

Figure 13. Decompositions of changes in Gini coefficients: top panel shows estimates of the change in the Gini and of the R and P components; bottom panel shows estimates expressed as a percentage of base period Gini. Based on single-year incomes.