Stephen P. Jenkins, Philippe Van Kerm

Assessing individual income growth

Article (Accepted version)
(Refereed)

Original citation:
Jenkins, Stephen P. and van Kerm, Philippe (2016) *Assessing individual income growth.* Economica, 83 (332). pp. 679-703. ISSN 1468-0335

DOI: 10.1111/ecca.12205

Assessing individual income growth

Stephen P. Jenkins

Department of Social Policy

London School of Economics and Political Science, London WC2A 2AE, U.K.

Email: s.jenkins@lse.ac.uk

Philippe Van Kerm

Luxembourg Institute of Socio-Economic Research

11, Porte des Sciences, Campus Belval, L-4366 Esch-sur-Alzette, Luxembourg

Email: philippe.vankerm@liser.lu

April 6, 2016

Abstract

We develop methods for describing distributions of income growth across individuals and for comparing changes in growth distributions over time. The methods include graphical devices ('income growth profiles') and dominance conditions, and also summary indices, together with associated methods of estimation and inference. Taking an explicitly longitudinal perspective, our approach illuminates clearly who are the gainers and the losers, and also provides distributionally-sensitive assessments – ones that allow the income growth for different individuals to be weighted differently. Our empirical application shows that the pattern of income growth in Britain over the period 1992–1996 was less pro-poor than that for 1998–2002 and not significantly different from the pattern for 2001–2005.

Keywords: individual income growth; pro-poor growth; progressive income growth; income mobility; income growth profile

JEL Classification: D31; D63; I32

Introduction

The income distribution in each year can be characterized as a Parade of Dwarfs and a few Giants (Pen, 1971). Each 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 Parade participants. There is typically a distribution of income changes – there are winners and losers – a feature that is missed by looking at the average or aggregate income growth rate, the focus of most discussion. This point is highlighted by Recommendation 4 of the Stiglitz-Sen-Fitoussi Commission ('Give more prominence to the distributional aspects when assessing social progress'), supported by their statement that '[a]verage income, consumption and wealth are meaningful statistics, but they do not tell the whole story about living standards. For example, a rise in average income could be unequally shared across groups, leaving some households relatively worse-off than others.' (Stiglitz et al., 2009, 13). In this paper, we address the issue of how to summarize income growth that is unequally shared. We develop methods for distributionally-sensitive assessments of income growth and illustrate them using data for Britain.

Researchers often assess changes in the personal distribution of income growth by comparing incomes at a series of common points in the income Parades for different years, for example at the deciles or vingtiles of each Parade. Growth incidence curves, graphs that summarize these calculations, are used in development economics (Ravallion and Chen, 2003) and have also been used to summarize income growth in the UK (see e.g. Joyce et al., 2010, Chapter 3). There is also a related literature developing indices of the pro-poorness of income growth: see the review by Essama-Nssah and Lambert (2009).

As a means of describing *who* has got better off or worse off, growth incidence curves and pro-poor growth indices miss some important aspects. Looking at the change in the income of the person at a specific quantile of Parade A and the person at the same quantile of Parade B ignores the fact that the persons concerned are not the same individuals. Over time, people change their position in the income Parade. As a recent US assessment of distributional impact of the Great Recession put it: 'cross-sectional data do not necessarily tell us how *individual* households are faring over time, since the group of households in the bottom 20 percent changes each year: some previously high-earnings households move into the bottom 20 percent, and some households that were previously in the bottom 20 percent move out of it' (Perri and Steinberg, 2012, 12; emphasis in original). Similarly, people move into and out of top income groups (Auten et al., 2013). More generally, every income group throughout the income range changes its composition over time (as we confirm later). To assess whether the individuals who are poor (or rich) this year are gainers or losers, one has to track the fortunes of individuals not the fortunes of income groups such as 'the poor' or 'the rich' whose composition may change from one year to the next. Longitudinal data are required since it is only these data that enable one to link the income of a specific individual in Parade A with her income in Parade B, and hence calculate the income growth that each individual experiences.

The need for distributionally-sensitive measures of income growth was recognized before Stiglitz et al. (2009). For example, Ahluwalia and Chenery (1974) criticized the weighting scheme underlying assessments based on the change in GDP per capita or average income (the growth rate for each quantile group is weighted by its share in total income). They proposed a more general approach in which 'the rate of increase in the welfare of society as a whole can [be] defined as a weighted sum of the growth of all groups', where each group's weight could be 'set according to the degree of distributional emphasis required' (Ahluwalia and Chenery, 1974, 39). They discuss in particular weighting schemes in which the weights are either 'in proportion to their numbers ("one man, one vote") or inversely proportional to their initial income levels ("poverty weights")' (Chenery, 1974, xvi). The former scheme provides equal weighting of each person's income growth. 'Poverty' weighting schemes include ones in which the weights are 'a declining function of the rank in the income distribution' (Klasen, 1994, 259).[1] This is the scheme underlying the approach we propose, and we discuss it in more detail later.

In this article, we develop a comprehensive framework for describing and evaluating the personal distribution of income growth, and changes in these distributions over time. We motivate our approach with discussion of previous literature on the topic in Section 1. The approach itself is set out in Section 2. We develop graphical devices (income growth profiles and cumulative income growth profiles) for describing distributions of individual-level income growth and show that non-intersections of profiles correspond to dominance according to general classes of social evaluation functions. We also develop distributionally-sensitive summary indices, and these include the average growth rate as a special case. In Section 3, we show that issues of statistical inference raised by our approach can be addressed using bootstrap methods for dependent clustered data. The following two sections apply the measurement and estimation methods to analyze patterns of income growth in Britain over subperiods in the two decades since the early 1990s. The data, derived from the British Household Panel Survey (BHPS) are described in Section 4, and our findings are presented in Section 5. Section 6 contains a summary and conclusions. Various extensions and sensitivity analyses are reported in Appendices.

1 Previous literature

There are three components to the social evaluation of income growth: a measure of income growth for each individual (call it δ_i), the individual-specific weight applied to that measure, and the way in which the personal distribution of weighted growth is aggregated across individuals.

If an equal weighting scheme is adopted (and subject to an appropriate definition of δ, about which more later), social evaluations of income growth reduce to comparisons of the (univariate) personal distributions of δ. See the applications of first-order stochastic dominance checks by Fields et al. (2002) and Chen (2009), for example. For the Fields and Ok (1999) index of directional income mobility, δ is the change over time in log income, there are equal weights, and the aggregate index is the population average of the δ.

More general systems of weighting of the δ_i are characterized axiomatically by Demuynck and Van de gaer (2012), for example, who posit an axiom of Priority for Lower Growth that states that aggregate growth is increased more by an increment to the income growth of someone with a low growth rate than someone with a high growth rate. In the aggregate growth index that they derive, different 'priorities' are summarized parametrically, and equal weights is a special case. However, these indices do not reflect the principal concerns of Ahluwalia and Chenery (1974), Klasen (1994), Stiglitz et al. (2009), and most subsequent literature, because the weights depend only on each individual's income growth rate. An income growth rate of 5% for person i and for person j is treated the same regardless of whether i was rich and j was poor initially.

It is more plausible to suppose that aggregate social evaluations of growth processes should depend not only on individual growth but also on information about base-year (or final-year) positions. In this case, evaluations based on utilitarian social welfare functions (SWFs) are one obvious approach. By these, we mean SWFs that are the population average of individuals' utility functions in which base-period and final-period incomes are arguments, as in Atkinson and Bourguignon (1982). For instance,

$$W(H) = \int\int v(x, y) dH(x, y)$$

where H is the joint distribution of income at base- and final-period income and $v(x, y)$ is the individual utility associated with base-period income x and final-period income y. (The form of $v(x, y)$ is common to all individuals.) The definitive analysis using this framework is by Bourguignon (2011), who derives dominance conditions characterizing when one personal distribution of income growth ('growth process') is socially preferred to another distribution.

For the purposes at hand, Bourguignon (2011) rewrites the Atkinson-Bourguignon utility function to be defined over base-period personal income x and the *change* in personal income between base-period and final-period c (so, $\delta_i = c_i = y_i - x_i$):

$$W(H) = \int\int v(x, x + c) dH(x, x + c) = \int\int u(x, c) dH(x, x + c)$$

where $u(x, c)$ is the individual utility re-expressed as a function of base-period income and income change (i.e. $u(x, c) = v(x, x + c)$). Bourguignon shows that *if the marginal distributions of base-year income are identical*, a social preference for distribution of income growth C over distribution D for all SWFs increasing in c (class V_1) requires the *conditional* CDF of income growth to be everywhere no greater in C than in D at every point along the base-period income parade (which also implies that the average income change in C is greater than in D at each point). In other words, the conditional distribution of income change in C must first-order dominate that in D for any base-period income level. Drawing on the results of Atkinson and Bourguignon (1987), Bourguignon (2011) derives a second dominance condition for a narrower class of SWFs, those that are not only increasing in individual income growth but also concave: an increase in income is valued more if it is received by individuals who are initially poorer (class V_2). For this class, social preference for C over D requires that the conditional CDF of income changes given base-period income *lower or equal to x* in C first-order dominates that in D for all base year income x. Class V_2 includes SWFs defined over final-year income only and CES-like functions in which base- and final-year income are substitutes but not necessarily perfectly so, as in the 'permanent' (longitudinally-averaged) income case.

Bourguignon's results highlight the challenges facing researchers wishing to take a distributionally-sensitive approach to assessments that takes both base-year and final-year income positions into account. On the one hand, the results usefully delineate conditions under which personal distributions of income growth can be unanimously ordered, while using a relatively conventional SWF-based approach. On the other hand, the conditions are remarkably restrictive. They are applicable only if the marginal distributions in the base year are identical, which is never likely to be the case when comparing growth distributions for the same country for different periods or between different countries.[2] And, in any case, the dominance conditions are demanding, and may well not be satisfied.[3] Thus, summary indices are almost certainly likely to be required to order distributions of individual income growth, but Bourguignon's (2011) paper does not consider these – or methods of statistical inference. We do so in this paper.

In sum, there is a need for a different approach to the measurement of distributionally-sensitive income growth, one that is both practically feasible and has consistent welfare foundations. The aim of this paper is to provide this approach and to demonstrate its empirical implementation, including methods for estimation and inference.

2 A framework for evaluating distributions of individual income growth

This section develops methods for comparing distributions of individual income growth corresponding to two time periods. (The methods could also be used for comparisons between growth for two regions or countries.) We propose graphical devices ('income growth profiles') for describing distributions of income growth, and derive relationships between configurations of income growth profiles, orderings according to social evaluation functions defined over bivariate income distributions, and classes of scalar indices of individual income growth.

2.1 Distributionally-sensitive evaluations

A social evaluation function for patterns of individual income growth can be written as $\Upsilon(H)$ where H is the joint distribution function of two positive random variables X and Y describing income distributions for years t and $t+\tau$ respectively: $H(x,y) = \Pr[X \leq x, Y \leq y]$. Loosely speaking, individual income growth is the change from X to Y, a form of directional income movement to use the language of Fields and Ok (1999). We propose that the evaluation function $\Upsilon(H)$ satisfy seven properties, discussed in turn.

Property 1. Individualistic and additive evaluations $\Upsilon(H)$ is the sum of individual-level income growth evaluations, $U(x,y;H)$, each of which summarizes income growth for a person with income x in year t and y in year $t+\tau$. $U(x,y;H)$ may be positive, zero, or negative, since an individual's income may rise, stay the same, or fall over time. In other words, social evaluations are individualistic and additive, as commonly assumed in the income inequality and poverty measurement literatures. We allow U to depend on the overall bivariate income distribution function H, so that in general the evaluation of a person's income growth may depend on the incomes of other individuals in the population.

Property 2. Replication invariance Comparisons of two distributions are independent of the size of the populations (as in the inequality and poverty measurement literatures).

Properties 1 and 2 imply that the overall social evaluation is a per-capita average:

$$\Upsilon(H) = \int_{z_-}^{z_+} \int_{z_-}^{z_+} U(r, s; H) dH(r, s) \tag{1}$$

where $[z_-, z_+]$ is the support of X and Y.

Property 3. Socially-weighted individual income growth We assume that the individual-level growth evaluation function for each person is the product of a component summarizing how much growth the person experiences (assessed in a manner common to all individuals) and a social weight that depends on the person's income rank in the base year. (This way of specifying an evaluation function follows Chakravarty (1984).) The aggregate evaluation function (1) can therefore be written as

$$\Upsilon(H) = \int_{z_-}^{z_+} \int_{z_-}^{z_+} w(F_X(r))\delta(r, s) dH(r, s) \tag{2}$$

where the social weight is $w(F_X(r)) > 0$, F_X is the marginal cumulative distribution of base-year incomes, and $\delta(x, y)$ is the 'growth distance function' common to all individuals. Social weights depend on base-year ranks, and hence are independent of the particular distribution of base-year incomes. With our specification for $\Upsilon(H)$, our approach departs from Bourguignon's (2011). So far, $\Upsilon(H)$ could be interpreted as an aggregate Atkinson-Bourguignon bivariate SWF. However, our aim is not to derive an aggregate welfare measure for the joint bivariate distribution of income but, rather, to make an evaluation of the patterns of *growth* that is sensitive to how individual growth is distributed along the initial income ordering.[4]

Property 4. Normalisation We also assume that w is scaled such that the average social weight is one ($\int w(F_X(r))\,dF_X = 1$). This makes $\Upsilon(H)$ a weighted average of individual income growth.

Property 5. Social preference for progressive income growth. That is, we suppose $w'(p) \leq 0$, which means that the impact on $\Upsilon(H)$ of an increment to $\delta(\cdot)$ is greater (or no smaller), the lower the individual's income in the base year (t). Giving greater weight to initially-poorer individuals builds in a social preference for greater equality in final-year ($t+\tau$) incomes than in base-year (t) incomes, other things being equal. An important reference point is $w(p) = 1$, for all p. This is the boundary case of neutrality towards differential income growth: $\Upsilon(H)$ is then simply the population average of the individual-level growth statistics, which we write for brevity as $\overline{\delta}$.

Property 6. Income growth is directional Since individuals' income changes may be positive or negative, we suppose that the growth distance function is directional. We follow Fields and Ok (1999, p. 460) who define this to mean that $\delta(x,y) = -\delta(y,x)$ and also $\delta(x,\rho x) > \delta(x,x)$ for all $\rho > 1$, where $\delta(\cdot)$ is a continuous function. We also introduce the normalization $\delta(x,x) = 0$. These conditions ensure that positive income growth for an individual corresponds to a social improvement ($\Upsilon(H)$ increases) and negative individual income growth corresponds to a reduction in $\Upsilon(H)$, ceteris paribus.

Property 7. The growth distance function is either (a) scale invariant, or (b) translation invariant. Case (a) refers to a requirement that $\delta(\lambda x, \lambda y) = \delta(x,y)$. Case (b) requires $\delta(x+\eta, y+\eta) = \delta(x,y)$. In our empirical application, we focus on two specifications for $\delta(\cdot)$. The first is $\log(y) - \log(x)$ which is the case that we shall refer to as proportionate income growth (satisfying case a); the second is $y - x$ which we refer to as absolute income growth (case b).[5] Remember that because income growth is directional (Property 6), $\delta(\cdot)$ may be positive or negative. Also, in conjunction with the other properties, scale invariance of $\delta(\cdot)$ implies scale invariance of $\Upsilon(H)$, and translation invariance of $\delta(\cdot)$ implies translation invariance of $\Upsilon(H)$.[6]

2.2 Income growth profiles and dominance relations

An alternative but equivalent expression for $\Upsilon(H)$ is useful for subsequent discussion. Since $H(r,s) = F_{Y|X=r}(s) \times F_X(r)$ where $F_{Y|X=x}$ is the cumulative distribution of final-year incomes conditional on base-year income x, we can rewrite (2) as

$$\begin{aligned} \Upsilon(H) &= \int_{z_-}^{z_+} \left(\int_{z_-}^{z_+} w(F_X(r))\delta(r,s)dF_{Y|X=r}(s) \right) dF_X(r) && (3) \\ &\equiv \int_{z_-}^{z_+} w(F_X(r))\, \mathrm{E}\left(\delta\left(r,y\right)|X=r\right)\, dF_X(r) && (4) \\ &= \int_0^1 w(p)\, \mathrm{E}\left(\delta\left(x(p),y\right)|X=x(p)\right)\, dp && (5) \end{aligned}$$

where $\mathrm{E}\left(\delta\left(r,y;H\right)|X=r\right)$ is the expected (average) income growth for individuals with a base-year income r. Equation (5) uses the change of variable $p = F_X(r)$, so $p \in [0,1]$ is the normalized rank in the base-year income distribution corresponding to income r, and $x(p) = F_X^{-1}(p)$ is the income corresponding to rank p in the base-year distribution. For brevity, we refer to the conditional expectation as $m(p)$ and rewrite (5) as:

$$\Upsilon(H) = \int_0^1 w(p)\, m(p)\, dp. \qquad (6)$$

Thus, instead of writing the aggregate evaluation in terms of a weighted average of the individual evaluations, we can express it equivalently in terms of weighted averages of expectations of individual-level income growth conditional on position in the base-year income parade, $m(p)$. We use this alternative representation when discussing income growth profiles below.[7]

Unambiguous orderings of pairs of distributions of individual income growth according to social evaluation functions with the properties just discussed correspond to dominance defined in terms of income growth profiles and cumulative income growth profiles.

An *income growth profile* is the plot of $m(p)$ against p, given a definition of $\delta(\cdot)$.[8] The income growth profile reveals how income growth is distributed according to position in the base year income distribution. If income growth – assessed using $m(p)$ – is the same for everyone, the profile is horizontal. The income growth profile has negatively-sloped sections over the ranges of p where individual income growth decreases as p increases (progressive income growth); and the profile has positively-sloped sections over the ranges of p where individual income growth increases as p increases (regressive income growth).[9]

Cumulative income growth profiles are plots of the average income growth for people with an initial income *at or below* a given percentile $x(p)$ in the base-year distribution, i.e. among the poorest $100p$ percent. That is, one plots

$$\frac{1}{p}\int_0^p m(q)dq \tag{7}$$

against $p \in [0, 1]$. The resulting graph plots areas below the income growth profile – analogous to the way that a generalized Lorenz curve shows areas below a quantile function. The slope of the cumulative income growth profile may be positive or negative at different p values. The cumulative profile's height at $p = 1$ is $\overline{\delta}$.

We may now state two results concerning comparisons of distributions of income growth. For proofs (using integration by parts) and further discussion, see Van Kerm (2006, 2009). We consider comparisons between period $\mathcal{A}$ which refers to income growth between year t and year $t + \tau$, and period $\mathcal{B}$ which refers to income growth between year s and year $s + \tau$. Assessments are made using the evaluation functions described earlier: the comparison is of $\Upsilon(H^{\mathcal{A}})$ and $\Upsilon(H^{\mathcal{B}})$ for periods $\mathcal{A}$ and $\mathcal{B}$.

Proposition 1 (Income growth profile (first order) dominance) *Let* $m^{\mathcal{A}}(p)$ *and* $m^{\mathcal{B}}(p)$ *denote the income growth profiles for periods* $\mathcal{A}$ *and* $\mathcal{B}$ *respectively.* $m^{\mathcal{A}}(p) \geq m^{\mathcal{B}}(p)$ $\forall p \in [0, 1]$ *if and only if* $\Upsilon(H^{\mathcal{A}}) \geq \Upsilon(H^{\mathcal{B}})$ *for any* Υ *with* $w(p) > 0$*.*

Proposition 2 (Cumulative income growth profile (second order) dominance) *Let* $C(p) = \frac{1}{p}\int_0^p m(q)dq$*.* $C^{\mathcal{A}}(p) \geq C^{\mathcal{B}}(p)$ $\forall p \in [0, 1]$ *if and only if* $\Upsilon(H^{\mathcal{A}}) \geq \Upsilon(H^{\mathcal{B}})$ *for any* Υ *with* $w(p) > 0$ *and* $w'(p) \leq 0$*.*

Income growth profile dominance states that finding the profile for period $\mathcal{A}$ is nowhere below and somewhere above the profile for period $\mathcal{B}$ is equivalent to the distribution of individual income growth for $\mathcal{A}$ being preferred to the distribution for $\mathcal{B}$ for any nonnegative social weight function $w(p)$. Cumulative income growth profile dominance states that finding the cumulative profile for period $\mathcal{A}$ is nowhere below and somewhere above the profile for period $\mathcal{B}$ is equivalent to the distribution of individual income growth for $\mathcal{A}$ being preferred to the distribution for $\mathcal{B}$ for any positive non-increasing weight function $w(p)$.

2.3 Indices of progressivity-adjusted growth and return-to-progressivity

Income growth profile dominance provides robust but only partial orderings of distributions of income growth with respect to the shape of the social weight function $w(\cdot)$ in the evaluation function Υ. Complete orderings by scalar indices that incorporate a social preference for progressive growth can be derived with further assumptions about the $w(\cdot)$. We refer to these as indices of 'progressivity-adjusted growth'.

In our analysis, we employ a class of single-parameter indices Υ^v in which the social weight function is defined using the rank-dependent scheme that is implicit in the generalized Gini inequality index (Donaldson and Weymark, 1980, Yitzhaki, 1983):

$$w(p) = v\,(1-p)^{v-1} \tag{8}$$

where $v \geq 1$. For all $0 < p < 1, w(p) > 0$, and $w'(p) \leq 0$ as long as $v > 1$, so the dominance results provide orderings for these indices. The larger that v is, the faster the decrease in the social weight as p increases, and hence the greater the preference for progressive income growth. If $v = 2$, weights decrease linearly with p from 2 to 0; if $1 < v < 2$, $w(p)$ is concave; and if $v > 2$, $w(p)$ is convex. If $v = 1$, $\Upsilon^v = \bar{\delta}$.

Υ^v reflects both levels of income growth and differences between individuals. There is also interest in summarizing the social return to the progressivity of individual income growth *per se*, i.e. the increase in social welfare that arises over and above the average growth experienced by the population as a whole (which varies with the business cycle). A natural index of the return to progressive growth is $G^v = \Upsilon^v - \bar{\delta}$. This measure can be interpreted as the difference between observed average growth and equally-distributed-equivalent growth – the growth which, if received uniformly by each individual, would yield the same overall evaluation as the observed average growth were it also received uniformly. The more positive that G^v is, the more progressive is income growth; negative values correspond to regressive income growth.[10]

3 Estimation and inference

3.1 Estimation of income growth profiles

Estimation of progressivity-adjusted growth indices, growth incidence curves and poverty growth curves is relatively straightforward because they involve standard estimators for concentration indices, quantiles, and incomplete means; see inter alia Chotikapanich and Griffiths (2001), Barrett and Donald (2009), and Verma and Betti (2011). We therefore focus our discussion on the estimation of income growth profiles.

Income growth profiles are examples of fractile graphs (Mahalanobis, 1960), functions that capture the relationship between an outcome (here income growth) and the *fractiles* of a covariate (here base-year income). Fractile graphs were developed to facilitate comparisons of regression functions over alternative conditioning variables when the conditioning variables have different distributions. Their estimation involves two stages. The first involves calculating $p_i = \hat{F}_t(x_i)$, the (fractional) rank of observation i with first period income x_i.[11]

The second stage is estimation of the conditional mean function of the outcome of interest non-parametrically conditional on the estimated fractional ranks, a problem for which a number of alternative estimators are available (Sen, 2005). We use Cleveland's (1979) locally weighted regression (LOESS). LOESS involves determination of overlapping local neighbourhoods around each of a series of points p spanning the range $[0, 1]$, and then using data for the sample observations that fall in each neighbourhood to estimate by (weighted) regression the expected value of $\delta\left(x(p), y\right)$ at each p, denoted $\hat{m}(p)$. Local regression methods are appropriate in this setting because of their well-known good behaviour near the boundaries of the support of the data, unlike running means or standard Nadaraya-Watson kernel regression estimators. The consistency of local regression estimators of fractile graphs is demonstrated by Sen (2005) who also analyzes their asymptotic properties.[12]

We estimate $\hat{m}(p)$ at 19 equally spaced points over the range of p (0.05, 0.10, ..., 0.95). Local neighbourhoods are defined as $p \pm 0.085$ so that approximately 17 percent of the sample falls in each neighbourhood.[13] We estimate cumulative income growth profiles by numerical integration of the LOESS estimates of the income growth profiles using a trapezoidal rule (Press et al., 2007). A bandwidth $p \pm 0.085$ was chosen for estimation of $\hat{m}(p)$ after experimentation. This choice is a compromise between not under-smoothing the profiles and not over-smoothing the cumulative profiles (estimates of cumulative income growth profiles are smoother than income growth profiles because of the cumulation process). We use the same bandwidth for all time periods. Estimates based on other bandwidths are available on request.

3.2 Resampling-based inference

The sampling variability of our estimates is computed using a non-parametric block (panel) bootstrap procedure. This accounts for sample dependence that arises from the longitudinal nature of the data – we use 16 waves of data from the BHPS corresponding to survey years 1991–2006 (see the next section). Such procedures are outlined by Cameron and Trivedi (2005, Chapter 11) but, as far as we are aware, they have not been applied systematically in previous research on individual income growth issues.

Resampling is from the sample of households interviewed in wave 1 of the BHPS. The full response history (including periods of non-participation if any) over BHPS waves 1–16 of all members of the selected households, plus their descendant split-off households and all respondents that later joined these households, is then selected to form a bootstrap replicate of the complete panel from waves 1 through 16. To deal with survey design features (potentially small stratum sizes in particular), resampling from BHPS wave 1 households is done by cluster (primary sampling unit) within each sample stratum using the repeated half-sample bootstrap algorithm of Saigo et al. (2001).[14] Let $\mathbf{X}^b$ denote one bootstrap replication b of the full BHPS waves 1–16 sample, so constructed. Our analysis examines changes over time in the distribution of income growth that we can observe in the BHPS; we contrast income growth between four periods 1992–1996, 1995–1999, 1998–2002 and 2001–2005. To preserve the dependence and overlapping membership of the four subsamples associated with each period, it is only in a second step that subsamples for each of the sub-periods compared are formed from the bootstrap replicates $\mathbf{X}^b$ of the full panel (according to selection rules described in more detail below). Denote the replication b subsample for period $(t, t+\tau)$ by $S_t(\mathbf{X}^b)$.

All of the statistics of interest described in Section 2 (including coordinates of the (cumulative) income growth profiles at any p) are then estimated on each replicate subsample $S_t(\mathbf{X}^b)$ with $b \in \{1, 2, \ldots, B\}$, where $B = 999$. We denote any such estimate $\hat{\theta}_t^b \equiv \theta(S_t(\mathbf{X}^b))$. We also calculate the difference between the $\hat{\theta}_t^b$ for each of the three later periods $(t, t+\tau)$ and that for the baseline period 1992–1996, $\hat{\Delta}_t^b \equiv \theta(S_t(\mathbf{X}^b)) - \theta(S_{1992}(\mathbf{X}^b))$. By computing all estimates for the four sub-periods and the between-period difference statistics using the same sets of bootstrap replication, we capture the covariances between different measures across time.

Our B replicates of $\hat{\theta}_t^b$ and $\hat{\Delta}_t^b$ are used to estimate the sampling error of our point estimates $\hat{\theta}_t$ and $\hat{\Delta}_t$, with the errors calculated as:

$$\hat{s}(\hat{\theta}_t) = \sqrt{\frac{1}{B-1}\sum_{b=1}^{B}(\hat{\theta}_t^b - \bar{\theta}_t)^2} \tag{9}$$

where $\bar{\theta}_t$ is the average of $\hat{\theta}_t^b$ over the B replications. Pointwise 95% variability bands are calculated using the bias-corrected percentile method (Efron, 1981).

We also use our bootstrap replications $\hat{\Delta}_t^b$ to evaluate the statistical significance of the *sign* of the observed changes over time. We do so by computing $\mathcal{P}(\hat{\Delta}_t) = \Pr[\mathrm{sgn}(\hat{\Delta}_t^b) \neq \mathrm{sgn}(\hat{\Delta}_t)]$ which we evaluate as

$$\mathcal{P}(\hat{\Delta}_t) = \frac{1}{B}\sum_{b=1}^{B}\mathbf{1}(\mathrm{sgn}(\hat{\Delta}_t^b) \neq \mathrm{sgn}(\hat{\Delta}_t)) \tag{10}$$

where $\mathbf{1}(\cdot)$ evaluates to 1 if the expression in parentheses is true and 0 otherwise. $\mathcal{P}(\hat{\Delta}_t)$ gives the proportion of bootstrap replications for which the sign of $\hat{\Delta}_t^b$ is different from the sign observed in the point estimates $\hat{\Delta}_t$. $\mathcal{P}(\hat{\Delta}_t)$ can be understood as the smallest α such that $(1-2\alpha)$ variability bands for $\hat{\Delta}_t$ do not include zero. We interpret these as giving the probability of incorrectly inferring the sign of the difference from the sign of the point estimate of the difference.

We adopt similar procedures for assessing the statistical significance of the dominance relations between (cumulative) income growth profiles. We base our tests of income growth profile dominance of period $\mathcal{A}$ (years $t, t+\tau$) over the period $\mathcal{B}$ (years $s, s+\tau; s > t$), on the statistic

$$\mathrm{MP}(t,s) = \min_p \left[\; \hat{m}^{\mathcal{A}}(p) - \hat{m}^{\mathcal{B}}(p) \; \right] \tag{11}$$

and tests of cumulative income growth profile dominance on the statistic

$$\mathrm{CMP}(t,s) = \min_p \left[\; \frac{1}{p} \left(\int_0^p \hat{m}^{\mathcal{A}}(q)dq - \int_0^p \hat{m}^{\mathcal{B}}(q)dq \right) \; \right] \tag{12}$$

where p takes the values $\{0.05, 0.10, \ldots, 0.95\}$. Income growth profile dominance for period $\mathcal{A}$ over period $\mathcal{B}$ is established if $\mathrm{MP}(t,s) \geq 0$, and cumulative income growth profile dominance if $\mathrm{CMP}(t,s) \geq 0$.

The $\mathrm{MP}(t,s)$ and $\mathrm{CMP}(t,s)$ statistics are similar to the Kolmogorov-Smirnov statistics used in some tests of stochastic dominance. Adapting the unrestricted bootstrap approach for stochastic dominance of Maasoumi and Heshmati (2000) to the analysis of income growth profile dominance, we claim observed dominance relations to be statistically significant at the $100(1-\alpha)$ percent confidence level if

$$\left(\frac{1}{B} \sum_{b=1}^{B} \mathbf{1}(\mathrm{MP}^b(t,s) < 0) \right) < \alpha, \tag{13}$$

that is, if the proportion of bootstrap replications in which dominance is *not* observed (where it is observed in the original sample) is smaller than α.[15] Cumulative income growth profile dominance is checked similarly.

4 BHPS data: definition of income and subsamples

We use data from waves 1–16 of the British Household Panel Survey (BHPS), corresponding to survey years 1991–2006.[16] BHPS wave 1 is a nationally representative sample of the population of Great Britain living in private households in 1991. Original sample respondents (including each partner from a dissolved wave 1 partnership) have been followed over time and they, and their co-residents, were interviewed annually. Children in original sample households are also interviewed as adults when they reach the age of 16 years. The BHPS following rules ensure that the sample remains broadly representative of the population of Britain over time.[17]

To assess whether patterns of income growth have changed over time, we report results for four periods, i.e. for four values of t for pairs of years t and $t + \tau$. The first period is 1992–1996. We treat this period as a baseline against which we contrast three subsequent periods: 1995–1999, 1998–2002, and 2001–2005.[18] Our analysis sheds descriptive light on the distribution of individual income growth and its progressivity, and how these have been changing over time. We do not seek to identify causal effects associated with particular policy reforms or particular governments in power at the time (Conservative in the first period, Labour in the third and fourth periods).

Our measure of an individual's income adheres as closely as possible to the 'before housing costs' net household income definition that is used in the UK's official income statistics, the *Households Below Average Income* series.[19]

Income is the income of the household to which a person belongs, 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 interview or the most recent relevant period, converted to a weekly equivalent pro rata.[20] Total money income is equivalized using the modified-OECD equivalence scale, and expressed in January 2008 pounds using a 'before housing costs' price index series supplied by the Department for Work and Pensions.

In order to reduce the potential impact on our estimates of measurement error and transitory income fluctuations, each person's income is measured using a three-year longitudinal average: an individual's income measure for year t is the arithmetic average of observed income for years $t-1$, t, and $t+1$.[21] Longitudinal averaging of an individual's income smooths out errors and other transitory variation if measurement errors are classical (see Appendix B). Also, to prevent outlier income values exerting undue influence, analysis throughout is based on samples from which incomes in the top 1% and the bottom 1% have been excluded (prior to taking three-year averages). Trimming of this kind is common in analysis of income dynamics: see e.g. Gottschalk and Moffitt (2009).

Given these definitions, our sample for estimation of growth over the period 1992–1996 is composed of respondents for whom income data are available for six survey years (1991, 1992, 1993, and 1995, 1996, 1997), and similarly for the other three time periods. The number of individuals in each subsample is: 6,088 (1992–1996), 6,130 (1995–1999), 5,789 (1998–2002), and 5,451 (2001–2005). Since the BHPS is an on-going panel survey, each of the four subsamples corresponding to the four periods has an overlapping membership. If there were no missing data on income, no attrition (or death) and no panel joiners (by coresidence or birth), each of the four sub-samples would comprise the same set of respondents observed at different points in time. Although sample overlap is incomplete in practice, there is dependence between the four subsamples and this needs to be taken account when calculating the standard errors of estimates. We use the bootstrap resampling algorithm discussed in Section 3.

Sample weights are used to compute all estimates. Our four longitudinal samples are weighted using the BHPS cross-sectional enumerated individual weights of year $t+\tau$. We do not use the BHPS longitudinal weights because our analysis samples include sample joiners, and no longitudinal weights are provided for them.

5 Patterns of individual income growth in Britain and their changes over time

In this section, we assess patterns of individual income growth in Britain over the 1990s and 2000s using the methods developed earlier. We have also undertaken extensive sensitivity analyses that confirm that our results are not driven by regression-to-the-mean and measurement error, or by sample ageing and lifecycle income changes. For brevity, these checks are reported in Appendix B and Appendix C.

5.1 Patterns of income growth: the longitudinal perspective

Figure 1 shows income growth profiles (left) and cumulative income growth profiles (right) for the four time periods considered. The top two panels show income growth in absolute terms (the units are January 2008 pounds) and the bottom two panels show income growth in proportionate terms (the units are log January 2008 pounds). In all graphs, vertical dashed lines are used to demarcate the poorest and richest fifths of each base-year income distribution.

[Insert Figure 1 about here]

Income growth profiles are negatively-sloped, broadly speaking, regardless of whether an absolute or proportionate definition of growth is used, and for each of the four periods (albeit with a positive slope over some small ranges of p). That is, from a longitudinal perspective, the pattern of individual income growth is progressive: the lower the rank in the base-year distribution, the greater the expected income growth. Expected income growth, absolute or proportionate, is positive for the majority of individuals, but negative for individuals in the richest fifth in the base year.

The income growth profiles do not differ markedly across the four time periods. No curve lies completely above another at all values of p, so there is no first order dominance result. Taking all periods together, the estimates indicate that individuals starting in the poorest fifth experience an income increase of approximately 15%–25% over the subsequent four years, i.e. around £30–£50 per week. For those starting in the richest fifth, the corresponding change is around –5%.

Differences between periods in patterns of income growth are more clear cut when we consider cumulative income growth profiles, especially when individual growth is measured in proportionate terms. The non-intersection of cumulative profiles suggests second-order dominance results, and unambiguous rankings by all members of our class of progressivity-adjusted growth measures. In the proportionate growth case, the ranking of the distributions of income growth for the periods is, from highest to lowest, 1998–2002, 1995–1999, 1992–1996, and 2001–2005. In the absolute growth case, the same ranking applies except that there is no dominance result concerning comparisons of 2001–2005 with 1995–1999 and 1992–1996: the cumulative profiles intersect in these cases.

These conclusions ignore sampling variability. Figures 2 and 3 show *differences* between the profile for 1992–1996 and each of the other three profiles, together with pointwise bootstrapped 95% variability bands. In each figure, the top panel refers to income growth profiles and the bottom panel to cumulative income growth profiles. Figure 2 refers to the absolute income growth case; Figure 3 refers to the proportionate income growth case. The other feature that is summarized in each of the graphs is the difference between periods in overall average individual income growth (the mean of the δ_i values, $\bar{\delta}$), shown as a horizontal dashed line. This provides an additional reference point for assessing change over time (see the discussion in Section 2).

The relatively large variability bands for the pointwise differences in profiles in many cases, combined with their overlap with the reference point of zero, might suggest that we are unlikely to find MP and CMP dominance, and we check this shortly using the methods explained earlier. (Overall profile dominance does not require that variability bands lie wholly above zero at every value of p.) Notwithstanding these remarks, cumulative profile differences and proportionate growth income growth generally provide more clear cut results generally. The difference curves and associated variability bands lie above zero at all values of p in Figure 3(e). This pattern implies that individual income growth was distinctively more progressive – and more pro-poor, specifically – in the 1998–2002 period than in 1992–1996. The comparisons between each of the other two periods and 1992–1996 do not stand out in this way.

[Insert Figure 2 about here]

[Insert Figure 3 about here]

The patterns just described are confirmed by the test statistics MP and CMP for assessing overall dominance. The statistics estimated for all 12 possible dominance comparisons across the four sub-periods examined are reported in Table 1. Remember MP(t, s) and CMP(t, s) give the smallest difference between the income growth profiles (or cumulative profile) for periods t and s over 19 equally spaced percentile ranks at which the profiles are estimated. Dominance of period t profile over period s profile therefore corresponds to a positive value for MP(t, s) or CMP(t, s). Income growth profile (first order) dominance is never observed in our samples: all statistics are negative. Cumulative income growth profile dominance (second order) is observed in 10 pairwise comparisons. These correspond to the cells in Table 1 with positive entries (emboldened to highlight them). These refer to dominance of 1998–2002 over all three other periods (for both absolute and proportionate changes), for 1995–1999 over 1992–1996 (for both absolute and proportionate changes) and over 2001–2005 (for proportionate changes) and for 1992–1996 over 2001–2005 (for proportionate changes). However, for dominance to be *statistically significant* there has to be dominance across virtually all bootstrap replications (see earlier). Only three of the ten dominance results satisfy this criterion: dominance of 1998–2002 over 1992–1996 for both measures of income change and dominance of 1998–2002 over 2001–2005 for proportionate change.

[Insert Table 1 here]

We now put numerical flesh on the patterns that we have graphed. The upper panel of Table 2 reports, for each period, average income in the initial and final years, and average income growth over the period for the population as a whole. (The average growth estimates are estimates of the progressivity-adjusted growth index Υ^v with $v = 1$, i.e. $\overline{\delta}$.) The lower panel reports, for each statistic, the difference between the value for each of the three later periods and the 1992–1996 reference period. Corresponding estimates are shown in the middle and right-hand panels for individuals who were in the poorest fifth or the richest fifth in the relevant base year.

Table 2 confirms that average income growth for the population as a whole increased in each period, and in both absolute terms (from £15 to £25) and proportionate terms (from 5% to 8%) until 1998–2002. Income growth then fell to a level similar to that for the initial reference period; indeed, to a lower level in the proportionate growth case (£14 or 3%).

In every period, those who started in the poorest fifth had larger income gains than average (e.g. £35 compared to £15 in 1992–1996), and substantially larger gains than those who started in the richest fifth who experienced negative income growth in all periods (e.g. –£19 in 1992–1996). The contrasts between the income change of the poorest and richest fifths were largest for the 1998–2002 period. This is a manifestation of the greater progressivity of income growth that we highlighted earlier. We summarize this progressivity using the progressivity-adjusted growth indices Υ^v shortly.

The greater income gains for those starting in the bottom fifth are not sufficiently large to enable them to catch up with the richest fifth but the gap between them is reduced substantially. The ratio of the mean income of the richest fifth to the mean income of the poorest fifth in the base year ranges from 3.9 (441/113 in 1992) to 3.4 (508/149 in 2001) whereas the final-year ratios of mean income for the same two groups range from 2.9 (422/147 in 1996) to 2.6 (488/189 in 2005). Put differently, initial mean income differences between the richest fifth and the poorest fifth are reduced by approximately a quarter because income growth is progressive.

These calculations summarize income changes for groups defined by base-year income position, but income groups change their composition over time. Information about the prevalence of changes in income group composition is provided in the two columns labelled '% move' in Table 2. The statistics show the proportion of individuals in a base-year income group (poorest fifth; richest fifth) that are not in that group in the final year.

There is an interesting symmetry: in each period, the fraction of individuals moving from the poorest fifth, or moving from the richest fifth, is the same – around one third. This is a relatively large proportion, and its size underscores the potential importance of taking a longitudinal perspective.

[Insert Table 2 here]

Income growth patterns are summarized further in Table 3 using our progressivity-adjusted growth indices (Υ^v) and return-to-progressivity indices (G^v). Estimates for each period are shown in the top panel of the table, with the absolute growth measures to the left and proportionate growth measures to the right. We vary the progressivity-sensitivity parameter v from 2 (linear Gini-like weights) to 4 (placing greater weight on gains for initially-poorer individuals). It turns out that all the estimates of Υ^v suggest the same trend over time. Progressivity-adjusted growth in every period was higher in each period than the preceding one, except in the final period when it fell. This was also the trend shown by the 'mean growth' estimates reported in Table 2 – these correspond to the case Υ^1.

Summaries of the differences in indices (relative to 1992–1996 values) are shown in the bottom panel of the table. The middle set of differences estimates, for 1998–2002, correspond to the case for which we found statistically significant second-order dominance. For this period, the difference in each index is large compared to the corresponding difference for other periods, and the associated 'p-value' is very small. Varying v makes little difference to the estimates. For each v, income growth is about £11, or 4%, and larger in 1998–2002 than in 1992–1996. The lack of variation in the difference estimates with changes in v is also a feature of the comparisons for the other two periods. Observe, however, that the estimated differences involving 2001–2005 are near zero and lack statistical significance, whereas the corresponding estimates for 1995–1999 are also positive.

The larger estimates of Υ^v for 1998–2002 partly reflect the greater average growth in this period reported in Table 2 and partly reflect greater growth progressivity. The estimates of G^v isolate the latter component. The return-to-progressivity is positive for all periods, and significantly so. But the *differences* in G^v between periods are less marked and never significant.

[Insert Table 3 here]

6 Concluding remarks

We have argued that analysis of the pattern of income growth, and its changes over time, should employ longitudinal perspectives to complement conventional repeated cross-section approaches. It is of interest to know how the incomes of specific individuals change over time, not only how the incomes associated with Parade positions (income 'slots') change over time. We have therefore developed methods that relate configurations of income growth profiles to unanimous orderings by aggregate measures of progressivity-adjusted income growth. For statistical inference, we have shown how to apply bootstrap resampling methods that take account of the dependent clustered samples that are inherent in this type of analysis.

Our estimates for four four-year periods in recent British history suggest that, from a longitudinal perspective, income growth is generally progressive. Over any particular period, income growth is greater for those with lower incomes in the base-year distribution. This is not an artefact of measurement error or sample ageing. This pattern is quite different from the picture provided by growth incidence curves (a repeated cross-section perspective): broadly speaking, they suggest that income growth over a four-year period has been typically been regressive in terms of absolute income changes though more distributionally neutral in terms of proportionate income changes. For further discussion, see Jenkins and Van Kerm (2011).

References

Ahluwalia, M. S. and Chenery, H. (1974). The economic framework,. In Chenery, H., Ahluwalia, M. S., Bell, C. L. G., Duloy, J. H., and Jolly, R., editors, *Redistribution* *with Growth*, chapter 2, pages 38–51. Oxford University Press.

Atkinson, A. B. and Bourguignon, F. (1982). The comparison of multi-dimensioned distributions of economic status. *Review of Economic Studies*, 49(2):183–201.

Atkinson, A. B. and Bourguignon, F. (1987). Income distribution and differences in needs. In Feiwel, G., editor, *Arrow and the Foundations of the Theory of Economic Policy*, pages 350–370. MacMillan, London.

Auten, G., Gee, G., and Turner, N. (2013). Income inequality, mobility, and turnover at the top in the US, 1987–2010. *American Economic Review*, 103(3):168–172.

Barrett, G. F. and Donald, S. G. (2009). Statistical inference with generalized Gini indices of inequality, poverty, and welfare. *Journal of Business and Economic Statistics*, 27(1):1–7.

Böheim, R. and Jenkins, S. P. (2006). A comparison of current and annual measures of income in the British Household Panel Survey. *Journal of Official Statistics*, 22(4):733–758.

Bourguignon, F. (2011). Non-anonymous growth incidence curves, income mobility and social welfare dominance. *Journal of Economic Inequality*, 9(4):605–627.

Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics: Methods and Applications*. Cambridge University Press, Cambridge.

Chakravarty, S. R. (1984). Normative indices for measuring social mobility. *Economics Letters*, 15(1–2):175–180.

Chen, W.-H. (2009). Cross-national differences in income mobility: evidence from Canada, the United States, Great Britain and Germany. *Review of Income and Wealth*, 55(1):75–100.

Chenery, H. (1974). Introduction. In Chenery, H., Ahluwalia, M. S., Bell, C. L. G., Duloy, J. H., and Jolly, R., editors, *Redistribution With Growth*, pages xiii–xx. Oxford University Press.

Chotikapanich, D. C. and Griffiths, W. (2001). On calculation of the extended Gini coefficient. *Review of Income and Wealth*, 47(4):541–547.

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

Demuynck, T. and Van de gaer, D. (2012). Inequality adjusted income growth. *Economica*, 79:747–765.

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. and Weymark, J. A. (1980). A single-parameter generalization of the Gini indices of inequality. *Journal of Economic Theory*, 22(1):67–86.

Efron, B. (1981). Nonparametric standard errors and confidence intervals. *The Canadian Journal of Statistics / La Revue Canadienne de Statistique*, 9(2):139–158.

Essama-Nssah, B. and Lambert, P. J. (2009). Measuring pro-poorness: a unifying approach with new results. *Review of Income and Wealth*, 55(3):752–778.

Fan, J. and Gijbels, I. (1996). *Local Polynomial Modelling and its Applications.* Chapman and Hall, London.

Fields, G. S., Leary, J. B., and Ok, E. A. (2002). Stochastic dominance in mobility analysis. *Economics Letters*, 75(3):333–339.

Fields, G. S. and Ok, E. A. (1999). Measuring movement of incomes. *Economica*, 66(264):455–471.

Gottschalk, P. and Danziger, S. (2001). Income mobility and exits from poverty of American children. In Bradbury, B., Jenkins, S. P., and Micklewright, J., editors, *The Dynamics of Child Poverty in Industrialised Countries*, chapter 5, pages 135–153. Cambridge University Press, Cambridge MA.

Gottschalk, P. and Moffitt, R. (2009). The rising instability of U.S. earnings. *Journal of Economic Perspectives*, 23(4):3–24.

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

Jenkins, S. P. (2010). The British Household Panel Survey and its income data. ISER Working Paper 2010-33, Institute for Social and Economic Research, University of Essex, Colchester, UK. Downloadable from `http://www.iser.essex.ac.uk/pubs/workpaps/pdf/2010-33.pdf`.

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

Jenkins, S. P. and Van Kerm, P. (2009). The measurement of economic inequality. In Salverda, W., Nolan, B., and Smeeding, T. M., editors, *The Oxford Handbook on Economic Inequality*, chapter 3, pages 40–67. Oxford University Press, Oxford.

Jenkins, S. P. and Van Kerm, P. (2011). Trends in individual income growth: measurement methods and British evidence. STICERD Discussion Paper PEP08, STICERD, London School of Economics. Downloadable from `http://sticerd.lse.ac.uk/dps/pep/pep08.pdf`.

Joyce, R., Muriel, A., Phillips, D., and Sibieta, L. (2010). Poverty and inequality in the UK: 2010. IFS Commentary 116, Institute for Fiscal Studies, London. Downloadable from `http://www.ifs.org.uk/comms/comm116.pdf`.

Klasen, S. (1994). Growth and well-being: introducing distribution-weighted growth rates to reevaluate U.S. post-war economic performance. *Review of Income and Wealth*, 40(3):251–272.

Levy, H. and Jenkins, S. P. (2012). *Documentation for Derived Current and Annual Net Household Income Variables, BHPS waves 1–18*. Institute for Social and Economic Research, University of Essex, UK. Downloadable from `http://www.esds.ac.uk/search/indexSearch.asp?ct=xmlSn&q1=3909`.

Maasoumi, E. and Heshmati, A. (2000). Stochastic dominance amongst Swedish income distributions. *Econometric Reviews*, 19(3):287–320.

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

Palmisano, F. and Van de gaer, D. (2013). History dependent growth incidence: A characterization and an application to the economic crisis in Italy. Working paper 43, Dipartimento di Scienze Economiche e Metodi Matematici, Università di Bari.

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

Perri, G. and Steinberg, J. (2012). Inequality & redistribution during the Great Recession. Economic Policy Paper 2012-1, Federal Reserve Bank of Minneapolis, Minneapolis, MN. Downloadable from `http://www.minneapolisfed.org/pubs/eppapers/12-1/epp_12-1_inequality.pdf`.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical Recipes: the Art of Scientific Computing*. Cambridge University Press, New York, 3rd edition.

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

Saigo, H., Shao, J., and Sitter, R. R. (2001). A repeated half-sample bootstrap and balanced repeated replications for randomly imputed data. *Survey Methodology*, 27(2):189–196.

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

Stiglitz, J. E., Sen, A. K., and Fitoussi, J. P. (2009). Report by the Commission on the Measurement of Economic Performance and Social Progress. Downloadable from `http://www.insee.fr/fr/publications-et-services/default.asp?page=dossiers_web/stiglitz/documents-commission.htm`.

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`.

Van Kerm, P. (2009). Income mobility profiles. *Economics Letters*, 102(2):93–95.

Van Kerm, P. (2013). Repeated half-sample bootstrap resampling. United Kingdom Stata Users' Group Meeting 2013 10, Stata Users Group. Downloadable from `http://www.stata.com/meeting/uk13/abstracts/materials/uk13_vankerm.pdf`.

Verma, V. and Betti, G. (2011). Taylor linearization sampling errors and design effects for poverty measures and other complex statistics. *Journal of Applied Statistics*, 38(8):1549–1576.

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

Notes

[1]Whereas Ahluwalia and Chenery (1974) referred only to weights that differed by (income quantile) group, Klasen (1994) is clear that the welfare weights could be individual-specific if one has unit-record data available.

[2]Bourguignon's (2011) empirical illustration compares a counterfactual growth process with the observed growth process for a set of countries in the global income distribution, and so the base-year distributions are identical, by construction.

[3]Weaker conditions are available for a narrower class of SWFs (class V_3), but the properties are harder to interpret, and may not command wide agreement: Bourguignon's (2011) third result, providing second-order dominance conditions, requires restrictions on a particular third-order derivative of the utility function.

[4]We discuss the welfare content of Υ and our distributionally-sensitive measures further in Appendix A.

[5]These are commonly-used ways of summarizing income growth, but not the only possibilities, e.g. a parametric class of scale-invariant growth distance functions is $(2^\beta/\beta)(y^\beta - x^\beta)/(x+y)^\beta, \beta \neq 0$, which is the 'arc percentage' divided by 100 in the case $\beta = 1$. A translation-invariant class is $(2^{1-\gamma}/\gamma)[\exp(\gamma y) - \exp(\gamma x)]/[\exp(\gamma x) + \exp(\gamma y)], \gamma \neq 0$. The limiting cases of $\beta = 0$ and $\gamma = 0$ are the proportionate and absolute growth distance functions cited in the text. We could also have considered the case of 'true' proportionate growth, defined as $(y-x)/x$.

[6]If $\delta\,(\lambda x, \lambda y) = \delta(x,y)$, $\overline{\delta}$ is the Fields and Ok (1999) index of directional income movement with $c = 1$.

[7]Most income mobility indices, regardless of mobility concept, can be written in the form shown by equations (3)–(6) (Jenkins and Van Kerm, 2009, Van Kerm, 2009).

[8]See Van Kerm (2006, 2009) for further discussion and illustrations. The income growth profile is labelled *mobility profile* by Van Kerm (2009). Grimm (2007) indepen-

dently proposed a similar device.

[9]Classical measurement error and transitory variability in incomes impart a negative slope to income growth profiles. We address this regression-to-the-mean issue in substantial detail in Appendix B.

[10]Since the first version of our paper (Jenkins and Van Kerm, 2011), Palmisano and Van de gaer (2013) have provided an axiomatic characterisation of these indices.

[11]Since individual income data typically contain ties (e.g. because members of the same household are generally assigned identical disposable income), we recommend use of estimators for the cumulative distribution function that assign equal rank for equal income. See e.g. Chotikapanich and Griffiths (2001).

[12]We have also undertaken Monte-Carlo simulation analysis using a data generation process that mimicks the joint distribution of household income and size and year-on-year dependence in incomes that appears in the panel data we use below. The analysis confirms that the estimators adopted are consistent and that the bootstrap methods for inference (described below) lead to confidence intervals with good coverage probability. Our simulation results prompted us to adopt a second order polynomial in the LOESS procedure which have better behaviour near the boundary of the support of p than a standard first order polynomial. Details are available from the authors.

[13]LOESS differs from other local polynomial regression estimators (Fan and Gijbels, 1996) in the definition of the neighbourhood. LOESS uses a nearest neighbour approach that selects a fixed proportion of the sample around each grid point. Local polynomial regression typically uses a kernel-based neighbourhood of fixed width around grid points. In our application the two approaches are equivalent since the δ_i are regressed on ranks p_i that are, by definition, uniformly distributed, and so a fixed fraction of the data falls into neighbourhoods of fixed width. In Jenkins and Van Kerm (2011), we report results based on a 'robust' LOESS approach (Cleveland, 1979) that aims to protect against the effects of outlying observations on conditional mean estimates. The findings are very similar to

those reported here.

[14]The repeated half-sample bootstrap algorithm is also versatile as it has been demonstrated to lead to valid bootstrap inference with both smooth and non-smooth statistics (Saigo et al., 2001). We use the Stata package *rhsbsample* for generating the replication weights (Van Kerm, 2013).

[15]We are in effect testing a null hypothesis of dominance against an alternative of non-dominance. Alternative approaches such as permutation tests or restricted bootstrap tests used in stochastic dominance analysis which allow tests of the null hypothesis of non-dominance are difficult to apply here given the complex dependence of the subsamples compared.

[16] When we began our research, BHPS net income data were available for 16 waves only. The BHPS ended in 2008. We have also undertaken analysis using the full 18-wave data set, but it does not change the conclusions reported below.

[17]The BHPS is less representative of immigrant groups arriving in Britain after 1991. We do not use data from the extension samples for Scotland, Wales, and Northern Ireland that began in the late 1990s because it would require use of complicated probability weighting schemes and the temporal coverage of the data is relatively short. For detailed discussion of the BHPS and its income data, see Jenkins (2010).

[18]Our conclusions are not substantially affected by the specific choices of the periods for which results are reported. This is confirmed by calculations based on rolling four-year windows for the entire period (results not reported for brevity).

[19]BHPS net income data files are an unofficial supplement to the official BHPS release, and documented by Levy and Jenkins (2012) and Jenkins (2010). Jenkins (2010) shows that BHPS cross-sectional income distributions closely match HBAI ones. For more details of the HBAI definition of income, see Department for Work and Pensions (2008).

[20]The use of a current rather than annual measure of income is standard in Britain.

Böheim and Jenkins (2006) show that the BHPS current and annual income measures provide very similar estimates of distributional summary statistics.

[21]The use of three-year averaged income follows Gottschalk and Danziger (2001) and Jenkins and Van Kerm (2006). Use of single-year non-averaged incomes provides qualitatively similar results to those reported below.

Acknowledgements

This research was partly supported by core funding from the University of Essex and the UK Economic and Social Research Council for the Research Centre on Micro-Social Change at the Institute for Social and Economic Research (awards RES-518-28-001 and ES/L0009153). 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 Higher Education and Research of Luxembourg. We gratefully acknowledge comments from two anonymous referees, Gordon Anderson, Francesco Andreoli, Tony Atkinson, John Hills, Markus Jäntti, Matthew Lindquist, Bodhi Sen, participants in seminars at ECRU (Univ. Catholique de Louvain), KU Leuven, University of Palermo, Oslo Fiscal Studies, University of Luxembourg, SOFI (Univ. Stockholm), STICERD (LSE), the workshop on 'Measuring discrimination, inequality and deprivation' at CEPS/INSTEAD, and IARIW, BHPS, and ECINEQ conferences.

(a) Absolute income growth profiles

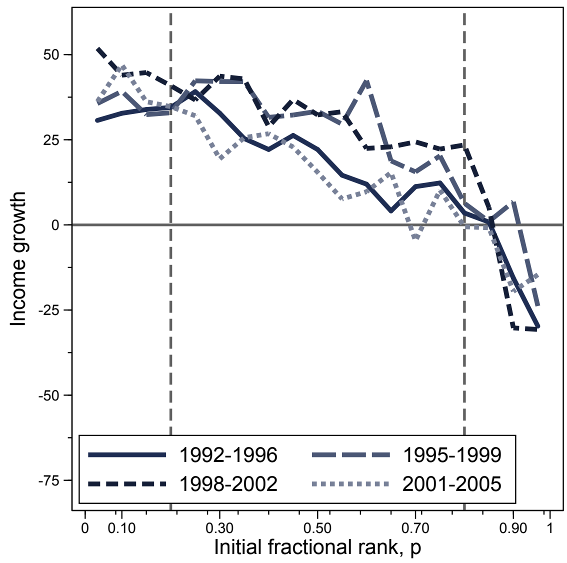

(b) Cumulative absolute income growth profiles

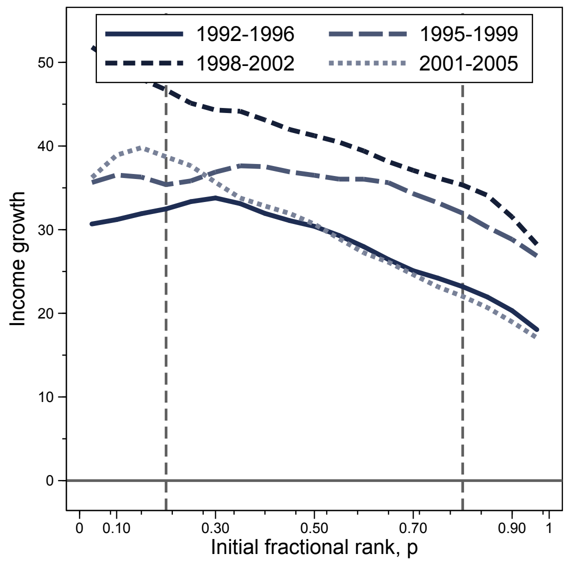

(c) Proportionate income growth profiles

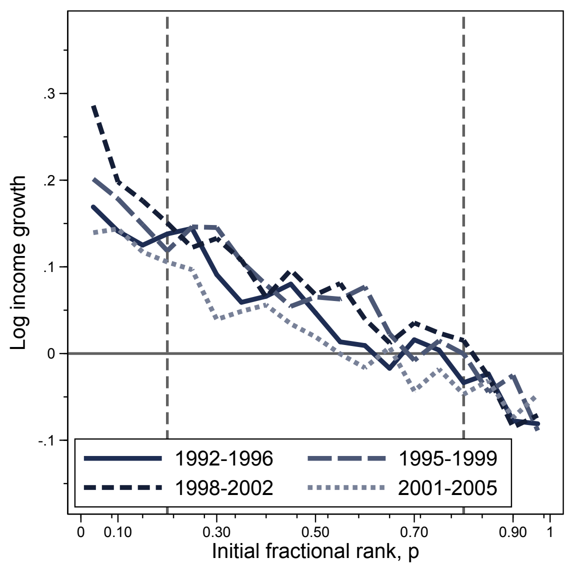

(d) Cumulative proportionate income growth profiles

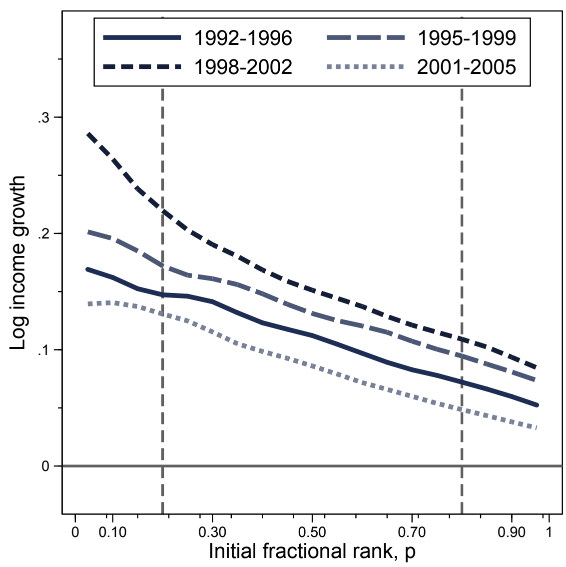

Figure 1. Income growth profiles and cumulative income growth profiles

Note. Top panel shows absolute growth (change in real income, pounds); bottom panel shows proportionate growth (change in log real income).

Income growth profiles

(a) 1995–1999 vs. 1992–1996 (b) 1998–2002 vs. 1992–1996 (c) 2001–2005 vs. 1992–1996

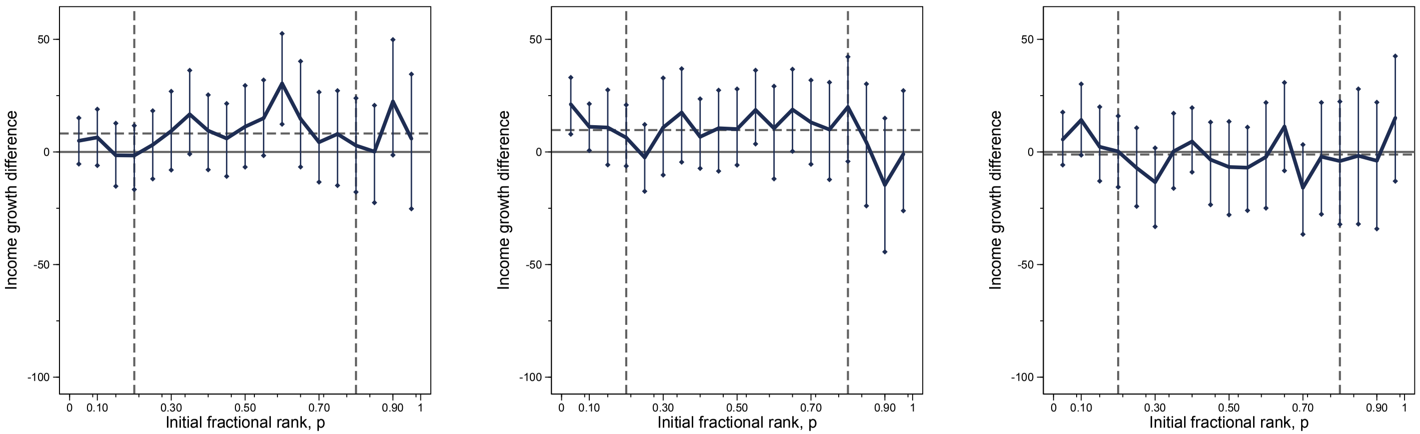

Cumulative income growth profiles

(d) 1995–1999 vs. 1992–1996 (e) 1998–2002 vs. 1992–1996 (f) 2001–2005 vs. 1992–1996

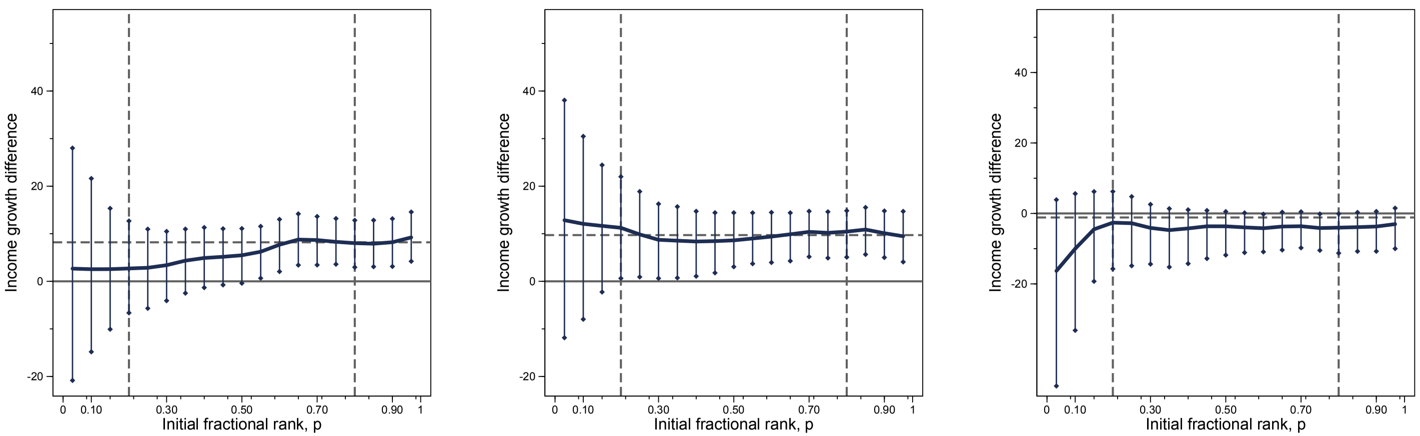

Figure 2. Differences in income growth profiles (relative to 1992–1996), with 95% pointwise variability bands: absolute income growth (change in real income, pounds)

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

Income growth profiles

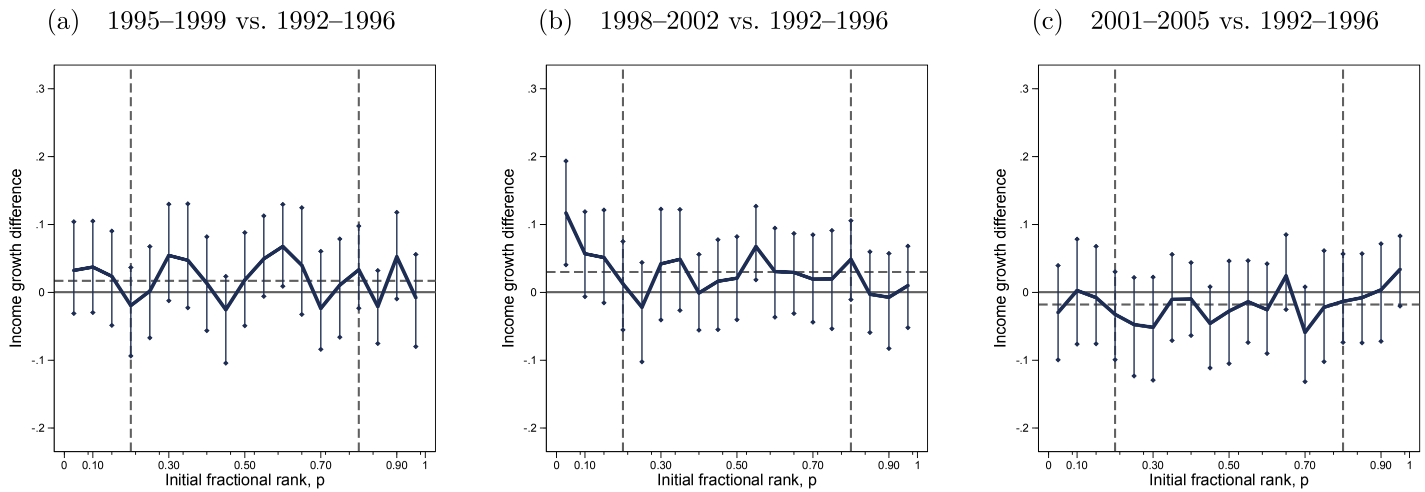

Cumulative income growth profiles

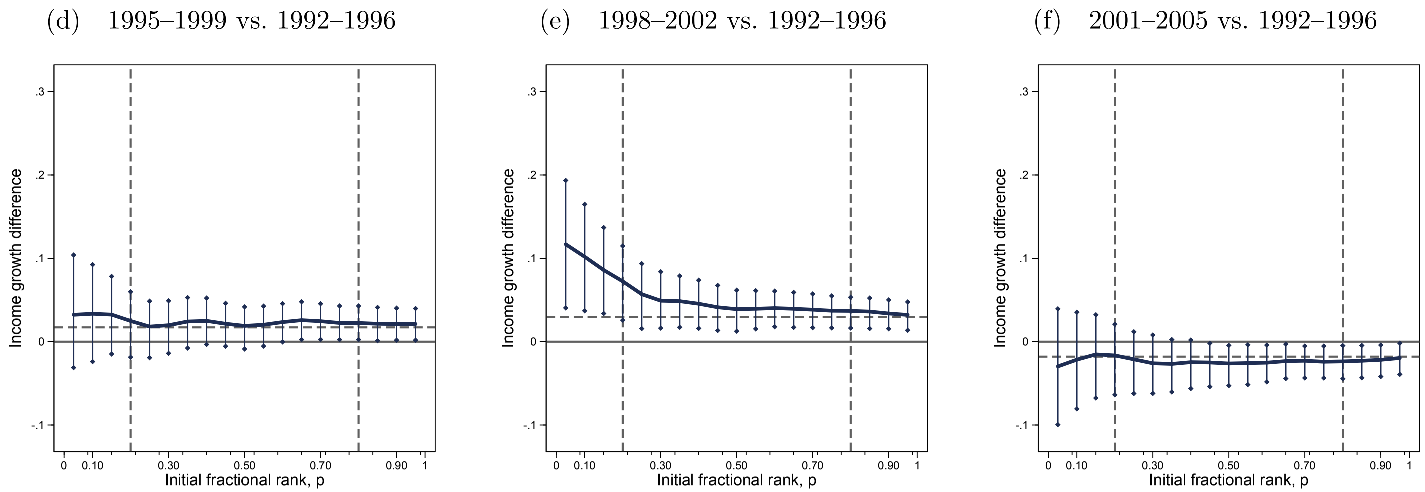

Figure 3. Differences in income growth profiles (relative to 1992–1996), with 95% pointwise variability bands: proportionate income growth (change in log real income)

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

Table 1: MP and CMP statistics for dominance checks across sub-periods with bootstrap confidence intervals and p-values

	Income growth		Log income growth	
	First order	Second order	First order	Second order
	MP(t, s)	CMP(t, s)	MP(t, s)	CMP(t, s)
Dominance statistics for 2001-2005 (period t) over ...				
1992-1996 (period s)	–16	–1	–0.06	–0.03
	(–19 ; 9)	(–6 ; 5)	(–0.06 ; –0.04)	(–0.07 ; 0.01)
	[1.000]	[0.785]	[1.000]	[0.997]
1995-1999 (period s)	–33	–10	–0.11	–0.06
	(–41 ; 16)	(–15 ; –4)	(–0.13 ; –0.07)	(–0.12 ; –0.03)
	[1.000]	[1.000]	[1.000]	[1.000]
1998-2002 (period s)	–29	–16	–0.15	–0.15
	(–33 ; 22)	(–26 ; –8)	(–0.20 ; –0.10)	(–0.21 ; –0.08)
	[1.000]	[1.000]	[1.000]	[1.000]
Dominance statistics for 1998-2002 (period t) over ...				
1992-1996 (period s)	–15	**10***	–0.02	**0.03***
	(–37 ; 0)	(7 ; 15)	(–0.04 ; 0.01)	(0.01 ; 0.05)
	[0.992]	[0.000]	[0.997]	[0.003]
1995-1999 (period s)	–37	**1**	–0.06	**0.01**
	(–63 ; 17)	(–3 ; 6)	(–0.09 ; –0.01)	(–0.01 ; 0.03)
	[1.000]	[0.257]	[1.000]	[0.171]
2001-2005 (period s)	–16	**7**	–0.02	**0.05***
	(–33 ; 0)	(1 ; 16)	(–0.04 ; 0.01)	(0.03 ; 0.07)
	[0.999]	[0.070]	[0.987]	[0.000]
Dominance statistics for 1995-1999 (period t) over ...				
1992-1996 (period s)	–2	**2**	–0.03	**0.02**
	(1 ; 2)	(–3 ; 8)	(–0.03 ; –0.01)	(–0.01 ; 0.04)
	[0.996]	[0.329]	[1.000]	[0.245]
1998-2002 (period s)	–17	–16	–0.08	–0.08
	(–20 ; 11)	(–30 ; –6)	(–0.13 ; –0.04)	(–0.17 ; –0.02)
	[1.000]	[1.000]	[1.000]	[0.998]
2001-2005 (period s)	–9	–4	–0.04	**0.04**
	(–14 ; 1)	(–13 ; 6)	(–0.08 ; 0.01)	(0.01 ; 0.06)
	[1.000]	[0.851]	[0.993]	[0.047]
Dominance statistics for 1992-1996 (period t) over ...				
1995-1999 (period s)	–30	–9	–0.07	–0.03
	(–42 ; 14)	(–13 ; –3)	(–0.07 ; –0.04)	(–0.07 ; 0.00)
	[1.000]	[1.000]	[1.000]	[0.998]
1998-2002 (period s)	–21	–21	–0.12	–0.12
	(–21 ; 15)	(–33 ; –12)	(–0.18 ; –0.07)	(–0.19 ; –0.05)
	[1.000]	[1.000]	[1.000]	[1.000]
2001-2005 (period s)	–15	–8	–0.03	**0.02**
	(–18 ; 7)	(–16 ; 1)	(–0.05 ; –0.01)	(–0.02 ; 0.04)
	[1.000]	[0.981]	[1.000]	[0.362]

Notes: For each comparison, the table reports point estimates for the MP and CMP statistics; a *positive* value indicates dominance in the sample (emboldened to highlight them). Bootstrap 95 percent confidence intervals for MP and CMP are reported in parentheses. Figures in square brackets give the proportion p of bootstrap replications in which we do *not* observe dominance (asterisks mark instances with $p < 0.01$).

Table 2: Income growth by income group (total population, poorest and richest 20%)

	All				Bottom 20% (at base period)					Top 20% (at base period)				
	Mean income		Mean growth		Mean income		Mean growth		% move	Mean income		Mean growth		% move
	Initial	Final	Income	Log income	Initial	Final	Income	Log income		Initial	Final	Income	Log income	
Estimates														
1992–1996	249	264	+15	+0.05	113	147	+35	+0.16	35.1	441	422	−19	−0.07	33.1
	(3)	(3)	(2)	(0.01)	(2)	(3)	(2)	(0.01)	(2.1)	(5)	(7)	(5)	(0.01)	(2.0)
1995–1999	259	283	+23	+0.06	122	158	+36	+0.18	33.4	453	442	−11	−0.05	32.3
	(3)	(3)	(2)	(0.01)	(2)	(3)	(2)	(0.01)	(2.4)	(6)	(8)	(6)	(0.01)	(2.2)
1998–2002	278	303	+25	+0.08	128	176	+48	+0.23	32.0	480	461	−19	−0.05	36.0
	(3)	(3)	(2)	(0.01)	(2)	(3)	(3)	(0.01)	(2.3)	(6)	(7)	(6)	(0.01)	(2.0)
2001–2005	302	316	+14	+0.03	149	189	+40	+0.14	34.5	508	488	−20	−0.06	33.8
	(3)	(4)	(2)	(0.01)	(2)	(4)	(3)	(0.02)	(2.4)	(7)	(8)	(6)	(0.01)	(2.0)
Difference from 1992–1996														
1995–1999	10	19	8	0.02	9	11	1	0.02	−1.7	12	20	7	0.01	−0.8
	[0.000]	[0.000]	[0.000]	[0.021]	[0.000]	[0.000]	[0.250]	[0.126]	[0.252]	[0.012]	[0.008]	[0.193]	[0.278]	[0.451]
1998–2002	29	39	10	0.03	16	29	13	0.06	−3.1	39	39	0	0.01	2.9
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.241]	[0.000]	[0.000]	[0.475]	[0.259]	[0.117]
2001–2005	53	52	−1	−0.02	36	42	6	−0.02	−0.6	67	66	−2	0.01	0.7
	[0.000]	[0.000]	[0.384]	[0.026]	[0.000]	[0.000]	[0.061]	[0.185]	[0.634]	[0.000]	[0.000]	[0.408]	[0.354]	[0.343]

Notes: The top panel reports point estimates (with bootstrap standard errors reported in parentheses; see text for description of the bootstrap resampling algorithm). The bottom panel reports differences between estimates for the 1992–1996 period and each of the subsequent periods. Figures in square brackets give the proportion of bootstrap replications for which the sign of the difference in point estimates is reversed (see text for details). 'Mean growth' estimates for all individuals show $\bar{\delta} = \Upsilon^1$.

Table 3: Progressivity-adjusted growth indices (Υ^v) and return-to-progressivity indices (G^v)

	Progressivity-adjusted growth, Υ^v						Return-to-progressivity, $G^v = W^v - \bar{\delta}$					
	Income growth			Log income growth			Income growth			Log income growth		
	$v = 2$	$v = 3$	$v = 4$	$v = 2$	$v = 3$	$v = 4$	$v = 2$	$v = 3$	$v = 4$	$v = 2$	$v = 3$	$v = 4$
Estimates												
1992–1996	25	27	28	0.08	0.10	0.10	8	10	10	0.04	0.05	0.06
	(2)	(2)	(2)	(0.01)	(0.01)	(0.01)	(1)	(1)	(2)	(0.00)	(0.01)	(0.01)
1995–1999	33	34	33	0.10	0.12	0.12	7	7	6	0.04	0.05	0.06
	(2)	(2)	(2)	(0.01)	(0.01)	(0.01)	(1)	(1)	(1)	(0.00)	(0.00)	(0.01)
1998–2002	36	37	37	0.11	0.13	0.14	9	11	10	0.04	0.06	0.07
	(2)	(2)	(2)	(0.01)	(0.01)	(0.01)	(1)	(2)	(2)	(0.00)	(0.00)	(0.01)
2001–2005	24	28	29	0.06	0.08	0.08	8	12	13	0.03	0.05	0.06
	(2)	(2)	(2)	(0.01)	(0.01)	(0.01)	(1)	(1)	(2)	(0.00)	(0.00)	(0.01)
Difference from 1992–1996												
1995–1999	8	6	5	0.02	0.02	0.02	−1	−3	−4	0.00	0.00	0.00
	[0.000]	[0.000]	[0.007]	[0.014]	[0.023]	[0.034]	[0.175]	[0.072]	[0.034]	[0.366]	[0.421]	[0.537]
1998–2002	11	10	10	0.03	0.03	0.04	1	0	0	0.01	0.01	0.01
	[0.000]	[0.000]	[0.000]	[0.000]	[0.002]	[0.003]	[0.315]	[0.453]	[0.464]	[0.126]	[0.162]	[0.166]
2001–2005	−1	0	1	−0.02	−0.02	−0.02	1	2	3	0.00	0.00	0.00
	[0.463]	[0.384]	[0.279]	[0.019]	[0.037]	[0.057]	[0.313]	[0.210]	[0.156]	[0.232]	[0.383]	[0.463]

Notes: The top panel reports point estimates (with bootstrap standard errors reported in parentheses). See text for description of the bootstrap resampling algorithm. The bottom panel report differences between estimates for the 1992–1996 period and each of the subsequent periods. Figures in square brackets give the proportion of bootstrap replications for which the sign of the difference in point estimates is reversed (see text for details). $\bar{\delta} = \Upsilon^1$ is the overall average growth rate.

Online supplementary material to

Assessing individual income growth

Stephen P. Jenkins

Department of Social Policy

London School of Economics and Political Science, London WC2A 2AE, U.K.

Email: `s.jenkins@lse.ac.uk`

Philippe Van Kerm

Luxembourg Institute of Socio-Economic Research

11, Porte des Sciences, Campus Belval, L-4366 Esch-sur-Alzette, Luxembourg

Email: `philippe.vankerm@liser.lu`

January 15, 2016

Appendix A The welfare content of distributionally-sensitive growth measures

Appendix A.1 Υ as an Atkinson-Bourguignon SWF

As explained in Section 2, our income growth evaluation function differs from an Atkinson-Bourguignon bivariate SWF in that we focus on the evaluation of income *growth*. This is distinct from the overall evaluation of the bivariate welfare of an Atkinson-Bourguignon SWF: comparisons of growth distributions according to Υ need not necessarily be consistent with comparisons according to W. However, as demonstrated by Bourguignon (2011), the distinction is unimportant if one is able to restrict comparisons to bivariate distributions that have identical base-period incomes. In this case, evaluation according to W is essentially the same as an evaluation of the income change distribution (with initial incomes given).

The evaluation Υ combined with a directional δ function and positive w can then be interpreted as a special case of the generic families of welfare functions V_1 considered in Bourguignon (2011). His first-order dominance conditions therefore apply to Υ. Considering $\delta(x, y) = y - x$ and further restricting w to be non-increasing, Υ defines a special case of welfare function belonging to V_2 and Bourguignon's second-order dominance conditions apply. It is also a member of V_3, except that it is a limiting case that restricts the second derivative of the individual utility function v_{yy} with respect to second-period income to be equal to zero. This means that inequality in second-period incomes *conditional on first-period income* does not affect overall welfare according to the resulting W.

The assumption that inequality in second-period incomes conditional on first-period income does not influence social evaluation is not exceptional in the context of mobility measurement. For example, Bénabou and Ok (2001) state that "(W)e are interested in mobility as an equalizer of ex-ante opportunities (or welfare), not of ex-post outcomes. Thus (...) our concern is not whether future realized income distributions will be more or less equal than the current one. They could be much more unequal, but if this primarily reflects shocks which were unpredictable on the basis of initial conditions there is little disparity of opportunity." (p. 2). One can also interpret this standpoint in the context of an Atkinson-Bourguignon SWF where second-period income is replaced by the expectation of second-period income conditional on base-period income in the individual utility functions, $v(x, \bar{y}_x) = v(x, x + \bar{\delta}_x) = u(x, \bar{\delta}_x)$. This idea is similar to the modification of the Atkinson-Bourguignon SWF proposed by Gottschalk and Spolaore (2002).

Appendix A.2 Υ and welfare change from an ex ante perspective

The formulation just described of social welfare measurement from an *ex ante* perspective as advocated in Bénabou and Ok (2001) leads to further social welfare content for the Υ measure.

The normative content of the index of $\Upsilon(H)$ can be linked in a straightforward way to linear, rank-dependent (or 'Yaari') social welfare functions (Yaari, 1987, 1988, Lambert, 2001). According to a Yaari SWF, the level of social welfare associated with a random variable with distribution function F is given by the functional

$$V(F) = \int_0^1 w(p) F^{-1}(p) dp. \tag{14}$$

Accordingly, $V_X \equiv V(F_X)$ and $V_Y \equiv V(F_Y)$ denote the level of (cross-section) social welfare in the initial and final period respectively. The welfare growth between the two periods is

$$\Delta V = V_Y - V_X \tag{15}$$

Let us now define $\bar{y}_X(x) = \mathrm{E}(y|X = x)$, the expected second-period income given initial income x. Define now $\bar{F}_Y$ as the distribution function of this *expected* second period income. Accordingly, $\bar{V}_Y \equiv V(\bar{F}_Y)$ is the social welfare associated with the distribution of expected second-period incomes,

$$V(\bar{F}_Y) = \int_0^1 w(p) \bar{F}_Y^{-1}(p) dp. \tag{16}$$

Observe that we focus here on $V(\bar{F}_Y)$ which differs from the actual welfare level in the second period $V(F_Y)$ by the degree to which income is dispersed around its expected value. The difference between base-period welfare and expected final-period welfare is given by

$$\begin{aligned}
\bar{\Delta}V &= \bar{V}_Y - V_X && (17)\\
&= \int_0^1 w(p)\bar{F}_Y^{-1}(p)dp - V_X && (18)\\
&= \int_0^1 w(p)\bar{F}_Y^{-1}(p)dp - \int_0^1 w(p)F_X^{-1}(p)dp && (19)
\end{aligned}$$

$\bar{\Delta}V$ is the statistic of interest for Bénabou and Ok (2001) who interpret it as an indication of the progressivity of the growth process and a measure of equalization of opportunities.

The connection between $\bar{\Delta}V$ and $\Upsilon(H)$ is immediate:

Proposition A1 *For $\delta(x, y) = y - x$, $\bar{\Delta}V = \Upsilon(H)$ if $\bar{y}_X(x)$ is monotone increasing in x.*

Proposition A1 states that if expected second-period income is increasing in initial income—which means that an individual's rank remains the same in the distribution of initial-period incomes and in the distribution of expected second-period incomes—then $\Upsilon(H)$ captures the welfare gain in the distribution of expected second period incomes compared to welfare in the initial distribution. Under this condition, an increase in $\Upsilon(H)$ unambiguously increases expected future welfare. (The proof of Proposition A1 derives trivially from the invariance properties of quantiles under monotonic transformations.) The monotonicity property is described in Shaked and Shanthikumar (2007, chapter 1) and is simply assumed by Bénabou and Ok (2001) who claim that '[e]mpirical plausibility requires that future income prospects increase smoothly with the current level, in the sense of first order stochastic dominance' (p.3).

Appendix A.3 Υ and intertemporal welfare from an ex ante perspective

The arguments can be extended in a straightforward manner to intertemporal welfare.

Consider now the distribution of the sum of initial and final period incomes ($z = x+y$), $F_{XY}(y)$, and the social welfare associated with it, $V_{XY} \equiv V(F_{XY})$.[1]

Define $\bar{F}_{XY}$ as the distribution function of *expected* total (or permanent) income: first period income plus expected second period income, $\bar{z}_X(x) = x + \bar{y}_X(x)$. $\bar{V}_{XY} \equiv V(\bar{F}_{XY})$ is the social welfare associated to the distribution of total income as expected from period 1.

We can now express the welfare gain induced by the pattern of expected income growth in comparison with a replication of the status quo (i.e. if everyone were to retain the same income income: $x = y$ and $z = 2x$):

$$\bar{\Gamma} V = \bar{V}_{XY} - 2V_X \tag{20}$$

$$= \int_0^1 w(p) \left(\bar{F}_Y^{-1}(p) + x(p) \right) \, dp - \int_0^1 w(p) 2x(p) \, dp \tag{21}$$

$$= \int_0^1 w(p) \left(\bar{F}_Y^{-1}(p) - x(p) \right) \, dp \tag{22}$$

The relationship between $\bar{\Gamma} V$ and Υ is as follows:

Proposition A2 *For $\delta(x, y) = y - x$, $\bar{\Gamma} V = \Upsilon(H)$ if $\bar{z}_X(x)$ is monotone increasing in x.*

Proposition A2 states that if $\bar{z}_X(x)$ (or equivalently $x(p)+m(p)$) is monotonically increasing in x—which means that an individual's rank remains the same in the distribution of initial-period incomes and in the distribution of expected total income—then $\Upsilon(H)$ captures the welfare gain of expected income growth in the distribution of expected inter-temporal welfare compared to welfare in the status quo distribution.[2] An increase in $\Upsilon(H)$ therefore unambiguously enhances welfare in expected permanent income. As a corollary,

Corollary A1 *For $\delta(x, y) = y - x$, $-\frac{m'(p;\delta)}{x'(p)} < 2$ for all $p \in [0, 1]$ implies $\bar{\Gamma} V = \Upsilon(H)$.*

This corollary relates the monotonicity condition to a condition on the relative slopes of the income growth profile and the initial period quantile function.

Proposition A2 and Corollary A1 derive straightforwardly from the monotonicity assumption and the invariance of quantiles under monotonic transformations. We know that

$$\bar{\Gamma}V = \int_0^1 w(p)\bar{F}_X Y^{-1}(p)dp - 2V_X$$

If $\delta(x, y) = y - x$, $m(F_X(x); \delta) = \bar{y}_X(x) - x$. After a change of variable $p = F_X(x)$: $m(p; \delta) = \bar{y}_X(x(p)) - x(p) = \bar{z}_X(x(p)) - 2x(p)$ where $x(p) = F_X^{-1}(p)$ is the quantile function of initial period incomes. Therefore

$$\begin{aligned}
\Upsilon(H) &= \int_0^1 w(p)\, m(p; \delta)\, dp \\
&= \int_0^1 w(p)\, (\bar{z}_X(x(p)) - 2x(p))\, dp \\
&= \int_0^1 w(p)\, \bar{z}_X(x(p))\, dp - 2\int_0^1 w(p)x(p)\, dp \\
&= \int_0^1 w(p)\, \bar{z}_X(x(p))\, dp - 2V_X.
\end{aligned}$$

The invariance property of quantiles under monotonic transformations implies $\bar{F}_{XY}^{-1}(p) = \bar{z}_X(x(p))$ and therefore increasing monotonicity of $\bar{z}_X(x)$ implies increasing monotonicity in $\bar{z}_X(x(p))$ and therefore

$$\begin{aligned}
\Upsilon(H) &= \int_0^1 w(p)\, \bar{F}_{XY}^{-1}(p)\, dp - 2V_X \\
&= \bar{V}_{XY} - 2V_X \\
&= \bar{\Gamma}V
\end{aligned}$$

Increasing monotonicity in $\bar{z}_X(x(p))$ can be empirically assessed from the sign of its first derivative for any p, which can be expressed in terms of the derivatives of both the base-period quantile function and of the income growth profile:

$$\begin{aligned}
&\bar{z}'_X(x(p)) > 0 \\
\Leftrightarrow\ & (m(p; \delta) + 2x(p))' > 0 \\
\Leftrightarrow\ & m'(p; \delta) > -2x'(p) \\
\Leftrightarrow\ & -\frac{m'(p; \delta)}{x'(p)} < 2
\end{aligned}$$

So $-\frac{m'(p;\delta)}{x'(p)} < 2$ for any p implies increasing monotonicity of $\bar{z}_X(x(p))$ which by Proposition A2 leads to Corollary A1 that $\bar{\Gamma}V = \Upsilon(H)$.

Appendix B Regression to the mean and measurement error

It might be argued that the negative slopes of our income growth profiles are simply due to regression to the mean. The reasoning is that if there is classical measurement error (errors 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.[3] Hence, some of the observed progressivity in income growth evidenced in the income growth profiles may be spurious. Our use of a three-year income average aims to reduce the impact of this problem by smoothing out measurement errors (and transitory variability) so that the economically substantive variations are better monitored. (As it happens, use of single-year incomes led to broadly similar findings but with steeper income growth curves at extreme ranks.) One response to the regression to the mean argument is that it is not so relevant to comparisons of patterns over time. There is no change e.g. in the BHPS design that leads us to expect the nature of measurement error to have varied over time. A bias in the profile slope may be consistent with no bias in the estimated difference between profiles. Another response to the classical measurement error argument is that, in reality, measurement error may not be classical. Gottschalk and Huynh (2010) argue persuasively that factors such as mean reversion in errors offset biases arising from the variance inflation aspect of measurement error, so that bias in measures of mobility may be negligible.[4]

Validation data are not available to us to examine this issue, and so our investigations of the impact of measurement error employed two other strategies. The first is a procedure analogous to an instrumental variables (IV) approach to reduce the impact of measurement errors. The second focuses on samples with more reliably-measured incomes. We discuss these in turn.

Our 'IV' approach uses definitions of income and income growth that break the link between measurement error in base-year income level and in income growth, thereby offsetting any regression-to-the-mean driven by classical measurement error. We approximate individual base-year ranks, $F_t(y_{it})$, by $\tilde{F}_t(\tilde{y}_{it})$ where $\tilde{F}_t$ is the distribution function of $\tilde{y}_{it} = y_{it-1} + y_{it+1}$. The idea is to use income lags and leads to approximate current period income and to derive base-year income rankings from this. We consider income change of the form $\delta\left(y_{it}, y_{it+\tau}; H\right)$ where y_{is} is income for year s (*not* a three-year smoothed average). These definitions imply that no data are used simultaneously to determine base-year rank and income change; it is as if people's ranks in the distribution of lag and lead incomes are used as instruments for their ranks in the distribution of current-year income.

Estimates employing the instrumented base-year ranks are presented in Appendix Figures B1, B2, and B3. The results are similar to those reported earlier. There is clear progressivity of income growth: income growth profiles remain negatively sloped.[5] The estimated trends over time are also reassuringly similar to those reported earlier, including the distinctive pro-poor nature of the 1998–2002 period. Table B1 shows all dominance results. All qualitative results are maintained, with similar levels of statistical significance.

Our second strategy is to recalculate our estimates using subsamples for which we believe incomes are better measured. The first subsample is our main sample but excluding individuals belonging to a household in which at least one person is in full-time self-employment. The second subsample drops instead individuals belonging to households with income components imputed by BHPS staff because of item non-response. (We also employed a third subsample excluded both individuals in households with self-employment and also those with imputed incomes, but subsample numbers were prohibitively small.) Self-employment income is notoriously difficult to capture reliably in surveys. Imputations lead to measurement error because of prediction error: to the extent that income is imputed in the final or initial year of a period, imputation is likely to lead to error in measurement of income change over time. One issue with the subsample strategy is that it can lead to substantial reductions in sample size. For example, dropping individuals from households with self-employed workers (in either the base or final year of each period) leads to reduction in sample size of about 20%–25%. Dropping

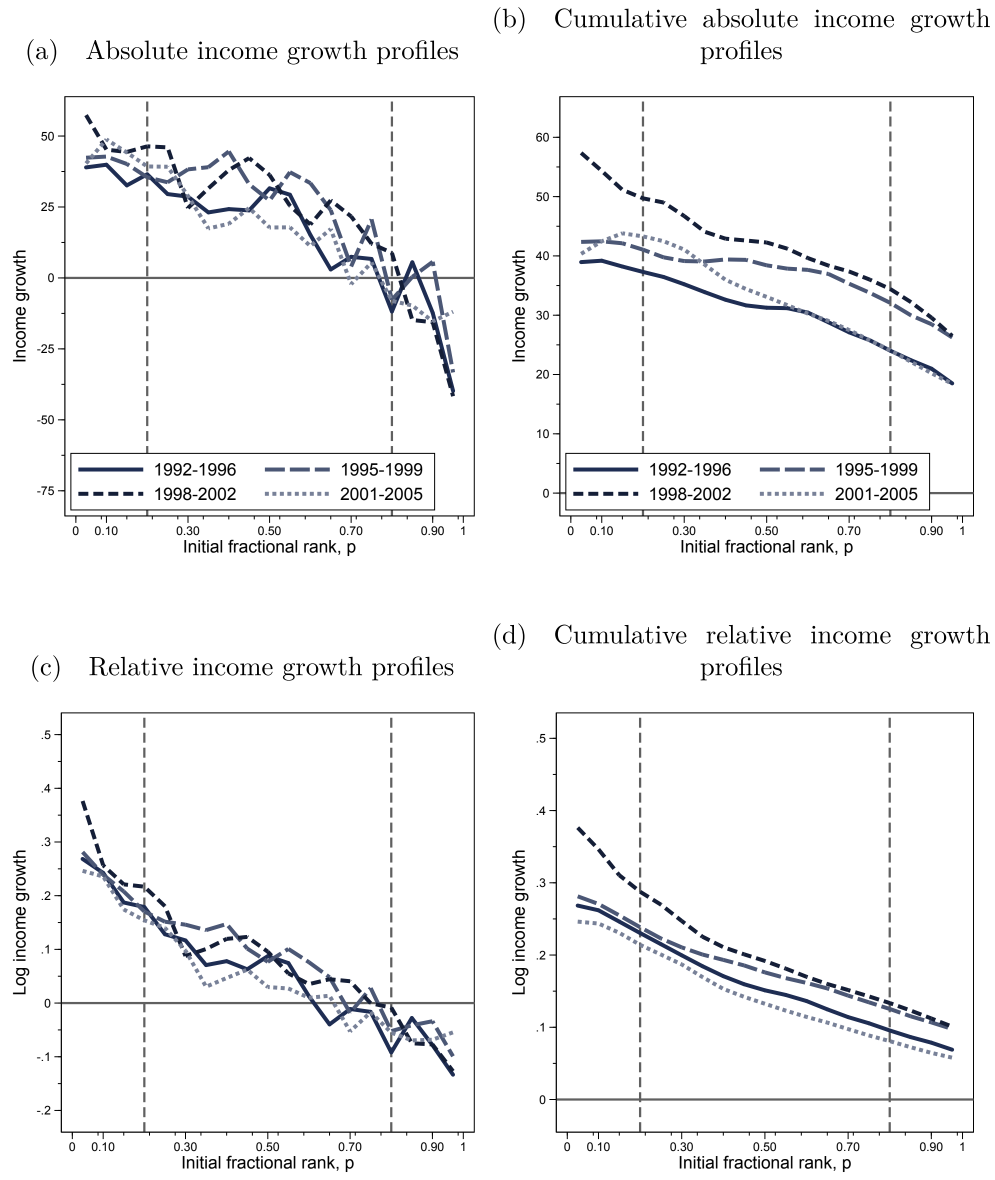

Figure B1. Income income growth profiles with proxied base year rank and no three-year-average smoothing

Note. Top panel shows absolute growth (change in real income in pounds); bottom panel shows proportionate growth (change in log real income).

Income growth profiles

(a) 1995–1999 vs. 1992–1996

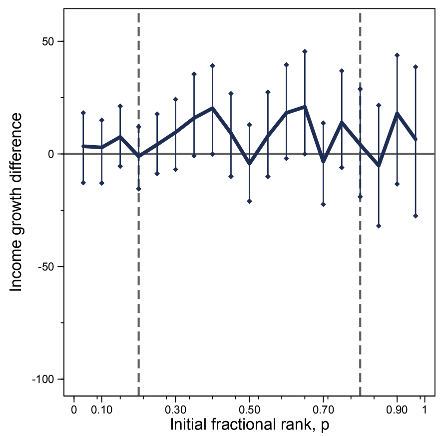

(b) 1998–2002 vs. 1992–1996

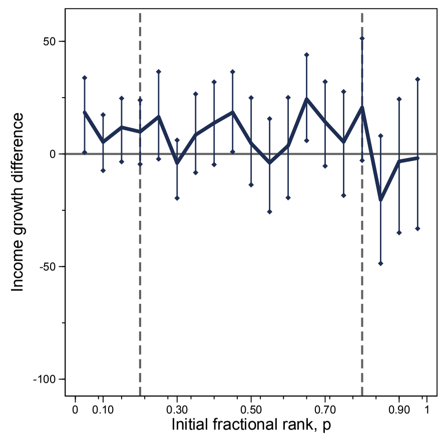

(c) 2001–2005 vs. 1992–1996

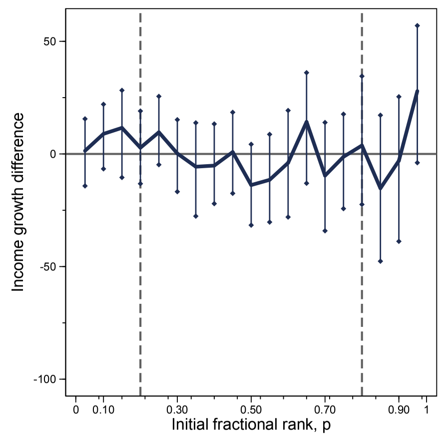

Cumulative income growth profiles

(d) 1995–1999 vs. 1992–1996

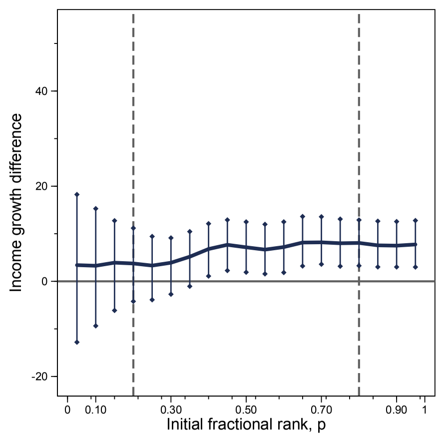

(e) 1998–2002 vs. 1992–1996

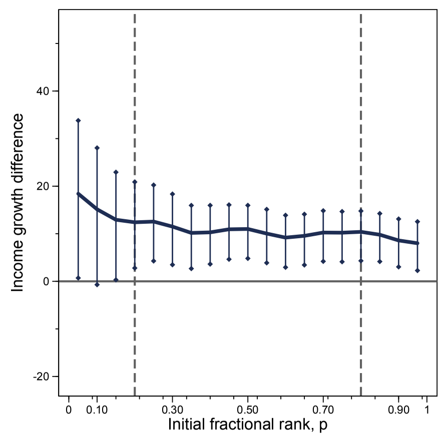

(f) 2001–2005 vs. 1992–1996

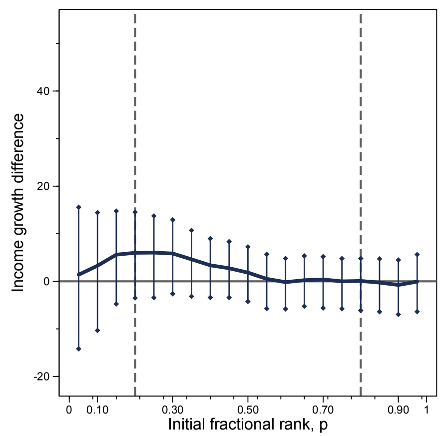

Figure B2. Differences between absolute income growth profiles with proxied base year rank and no three-year-average smoothing

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

Income growth profiles

(a) 1995–1999 vs. 1992–1996

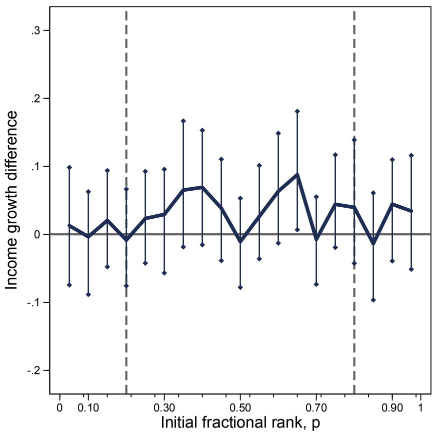

(b) 1998–2002 vs. 1992–1996

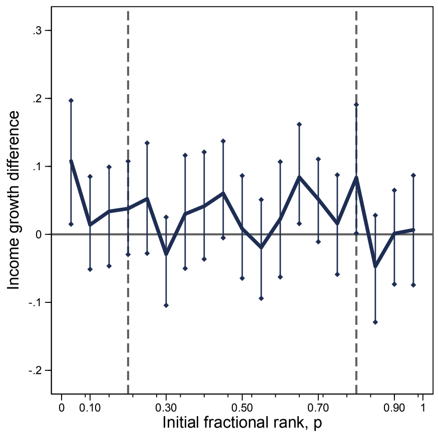

(c) 2001–2005 vs. 1992–1996

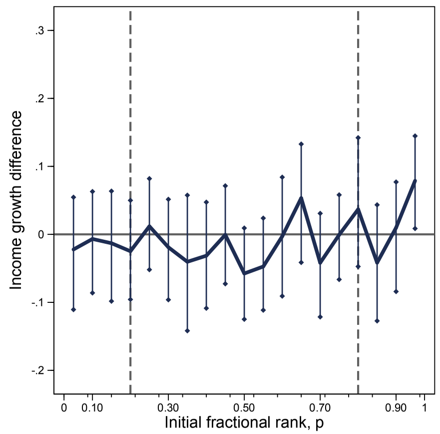

Cumulative income growth profiles

(d) 1995–1999 vs. 1992–1996

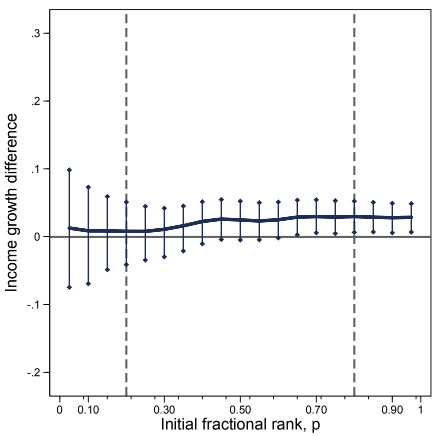

(e) 1998–2002 vs. 1992–1996

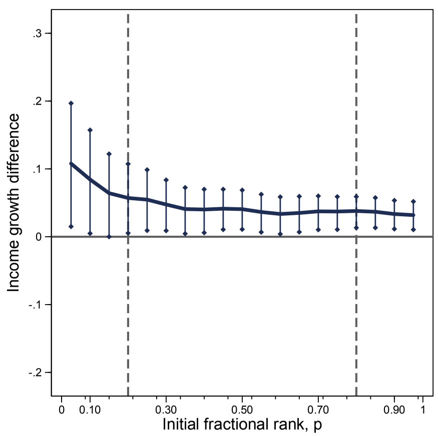

(f) 2001–2005 vs. 1992–1996

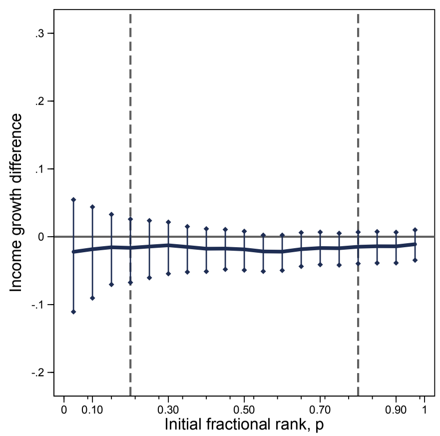

Figure B3. Differences between proportionate income growth profiles with proxied base year rank and no three-year-average smoothing

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

Table B1: MP and CMP statistics for dominance checks across sub-periods with bootstrap confidence intervals and p-values (with proxied base year rank and no three-year smoothing)

	Income growth		Log income growth	
	First order	Second order	First order	Second order
	$MP(t, s)$	$CMP(t, s)$	$MP(t, s)$	$CMP(t, s)$
Dominance statistics for 2001-2005 (period t) over ...				
1992-1996 (period s)	–15	**1**	–0.06	–0.02
	(–18 ; 7)	(–14 ; 15)	(–0.06 ; –0.03)	(–0.11 ; 0.05)
	[1.000]	[0.393]	[1.000]	[0.682]
1995-1999 (period s)	–25	–2	–0.11	–0.04
	(–28 ; 15)	(–17 ; 11)	(–0.13 ; –0.05)	(–0.12 ; 0.04)
	[1.000]	[0.611]	[1.000]	[0.827]
1998-2002 (period s)	–24	–17	–0.13	–0.13
	(–27 ; 15)	(–32 ; –3)	(–0.18 ; –0.09)	(–0.21 ; –0.05)
	[1.000]	[0.988]	[1.000]	[0.999]
Dominance statistics for 1998-2002 (period t) over ...				
1992-1996 (period s)	–20	**18**	–0.05	**0.11***
	(–43 ; 4)	(1 ; 33)	(–0.08 ; 0.00)	(0.01 ; 0.19)
	[0.999]	[0.013]	[0.999]	[0.008]
1995-1999 (period s)	–21	**15**	–0.06	**0.09**
	(–28 ; 10)	(–2 ; 29)	(–0.07 ; –0.03)	(0.00 ; 0.18)
	[1.000]	[0.039]	[1.000]	[0.020]
2001-2005 (period s)	–30	**17**	–0.07	**0.13***
	(–59 ; 10)	(3 ; 32)	(–0.13 ; –0.02)	(0.05 ; 0.21)
	[0.999]	[0.012]	[0.999]	[0.001]
Dominance statistics for 1995-1999 (period t) over ...				
1992-1996 (period s)	–5	**3**	–0.01	**0.01**
	(–6 ; 0)	(–13 ; 18)	(–0.01 ; 0.00)	(–0.07 ; 0.10)
	[1.000]	[0.301]	[0.998]	[0.363]
1998-2002 (period s)	–18	–15	–0.10	–0.10
	(–20 ; 12)	(–30 ; 1)	(–0.15 ; –0.05)	(–0.18 ; 0.00)
	[1.000]	[0.961]	[1.000]	[0.981]
2001-2005 (period s)	–21	**2**	–0.04	**0.03**
	(–48 ; 5)	(–11 ; 16)	(–0.09 ; 0.01)	(–0.04 ; 0.11)
	[1.000]	[0.401]	[0.991]	[0.178]
Dominance statistics for 1992-1996 (period t) over ...				
1995-1999 (period s)	–21	–4	–0.09	–0.01
	(–21 ; 15)	(–18 ; 12)	(–0.10 ; –0.07)	(–0.10 ; 0.08)
	[1.000]	[0.708]	[1.000]	[0.640]
1998-2002 (period s)	–24	–18	–0.11	–0.11
	(–28 ; 14)	(–34 ; –1)	(–0.13 ; –0.06)	(–0.20 ; –0.02)
	[1.000]	[0.988]	[1.000]	[0.993]
2001-2005 (period s)	–28	–1	–0.08	**0.02**
	(–48 ; 12)	(–16 ; 14)	(–0.12 ; –0.02)	(–0.05 ; 0.11)
	[1.000]	[0.613]	[1.000]	[0.320]

Notes: For each comparison, the table reports point estimates for the MP and CMP statistics; a *positive* value indicates dominance in the sample (emboldened to highlight them). Bootstrap 95 percent confidence intervals for MP and CMP are reported in parentheses. Figures in square brackets give the proportion p of bootstrap replications in which we do *not* observe dominance (asterisks mark instances with $p < 0.01$).

individuals in households with imputation of at least one major household income component reduces the sample by up to 75%–80%. As a result, we increased the bandwidth used to derive the income growth profile to $h = 0.2$ for the 'no self-employed' subsample and to $h = 0.3$ for subsamples excluding imputed data.

Excluding self-employed households leads to no substantial change in results. Profiles are not markedly flatter (except for very high p). When observations with imputed incomes are excluded, income growth profiles remain negatively-sloped but they become much flatter. Broad conclusions about general trends remain valid, although the much reduced sample sizes makes it difficult to draw robust conclusions with confidence and statistical significance is weaker.

(a) Absolute income growth profiles

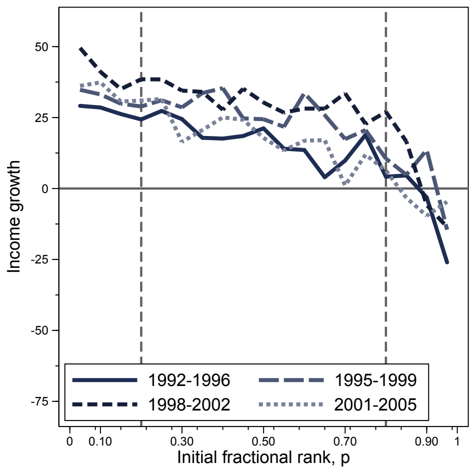

(b) Cumulative absolute income growth profiles

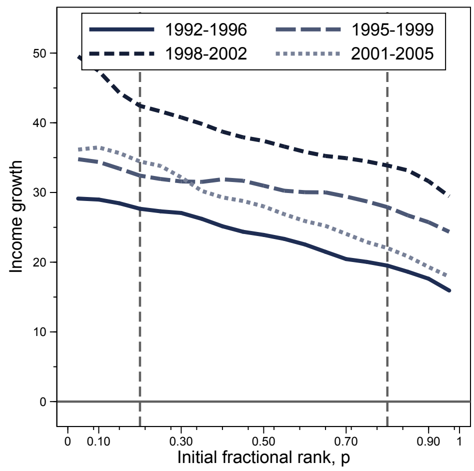

(c) Relative income growth profiles

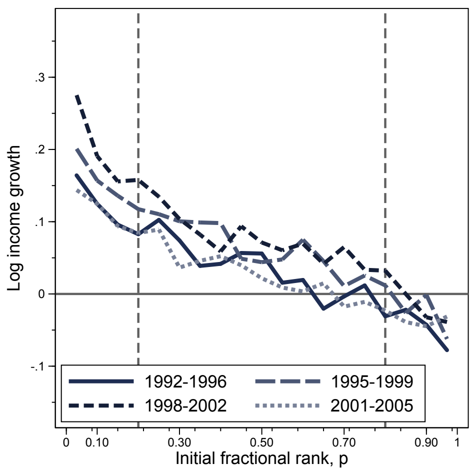

(d) Cumulative relative income growth profiles

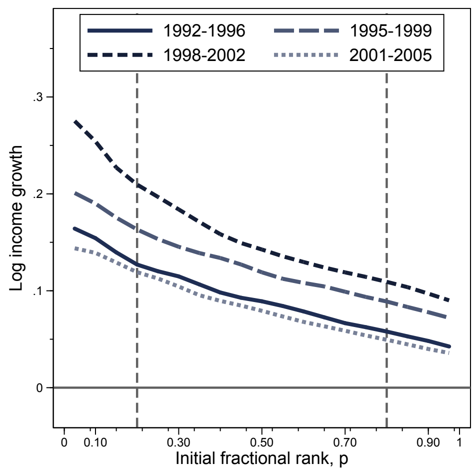

Figure B4. Income growth profiles (no self-employed subsample)

Note. Top panel shows absolute growth (change in real income in pounds); bottom panel shows proportionate growth (change in log real income).

Income growth profiles

(a) 1995–1999 vs. 1992–1996

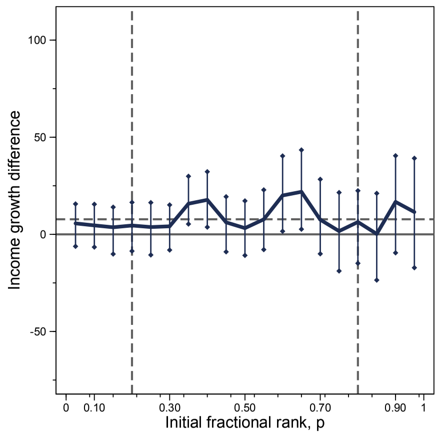

(b) 1998–2002 vs. 1992–1996

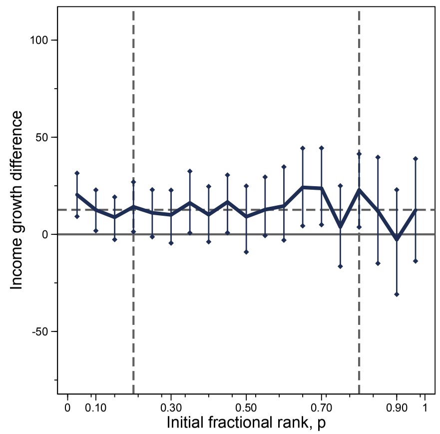

(c) 2001–2005 vs. 1992–1996

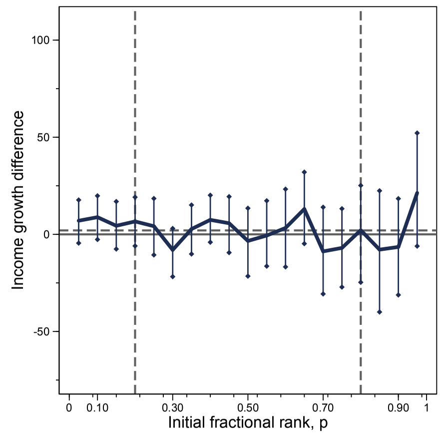

Cumulative mobility profiles

(d) 1995–1999 vs. 1992–1996

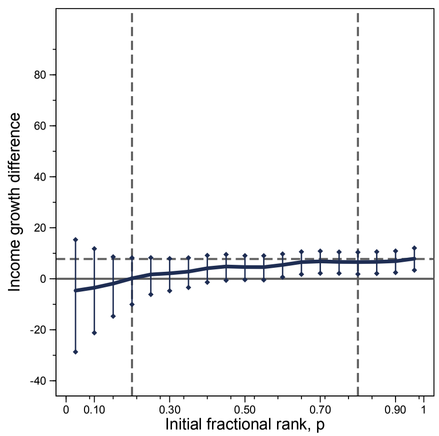

(e) 1998–2002 vs. 1992–1996

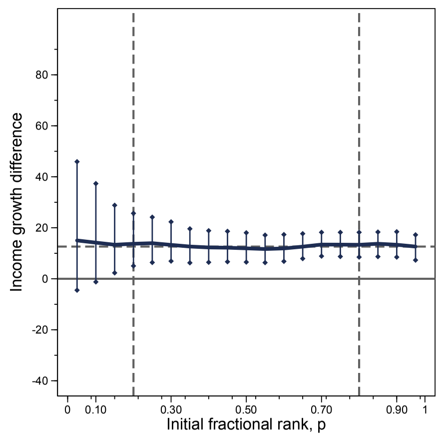

(f) 2001–2005 vs. 1992–1996

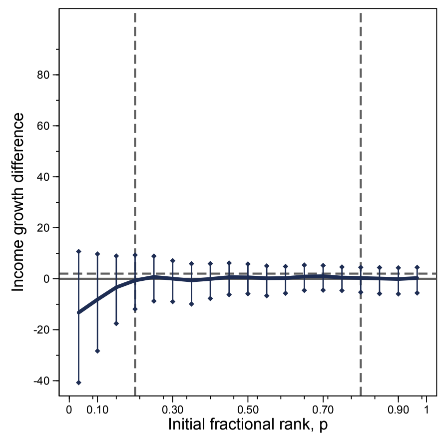

Figure B5. Absolute income growth profile difference (relative to 1992–1996) with pointwise variability bands (no self-employed subsample)

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

Income growth profiles

(a) 1995–1999 vs. 1992–1996

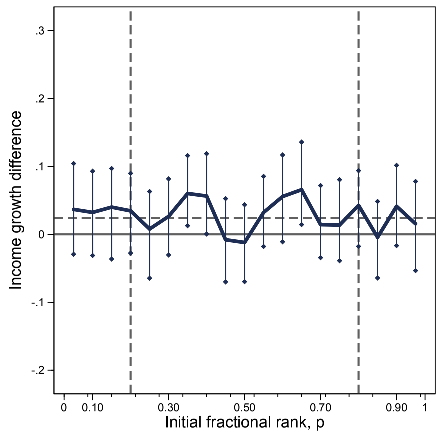

(b) 1998–2002 vs. 1992–1996

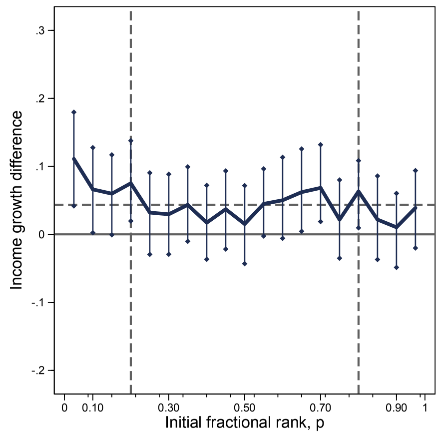

(c) 2001–2005 vs. 1992–1996

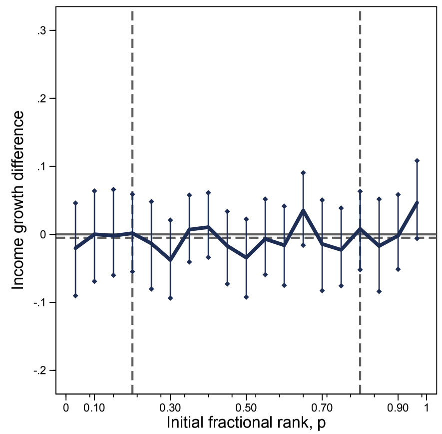

Cumulative mobility profiles

(d) 1995–1999 vs. 1992–1996

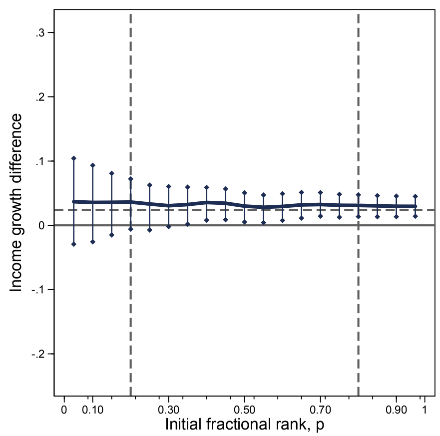

(e) 1998–2002 vs. 1992–1996

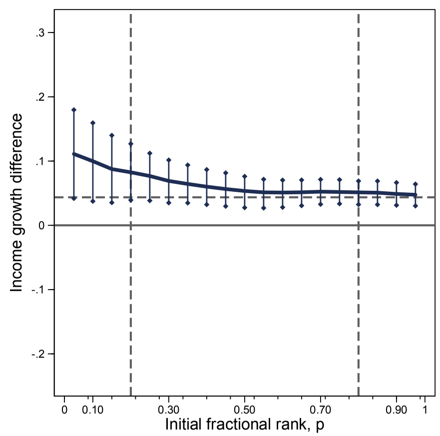

(f) 2001–2005 vs. 1992–1996

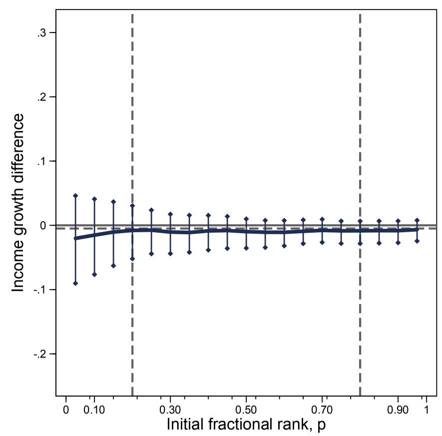

Figure B6. Proportionate income growth profile difference (relative to 1992–1996) with pointwise variability bands (no self-employed subsample)

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

(a) Absolute income growth profiles

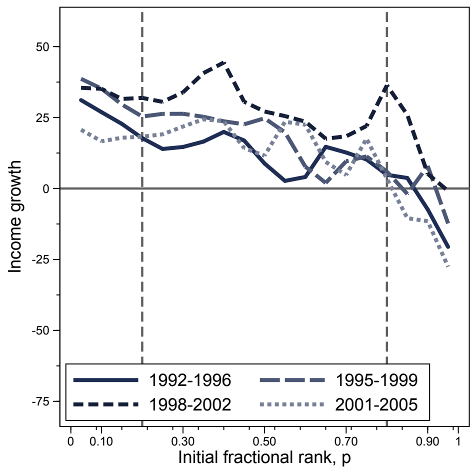

(b) Cumulative absolute income growth profiles

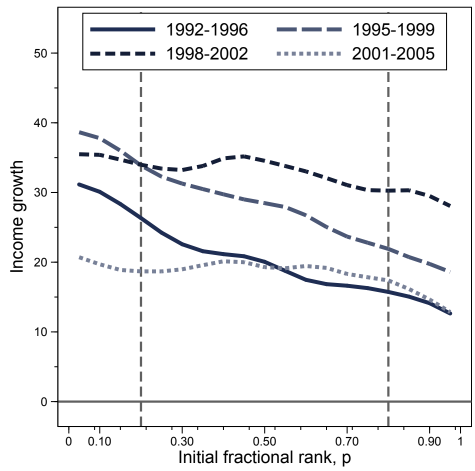

(c) Relative income growth profiles

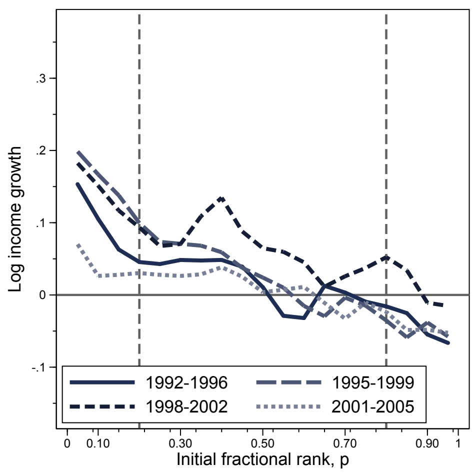

(d) Cumulative relative income growth profiles

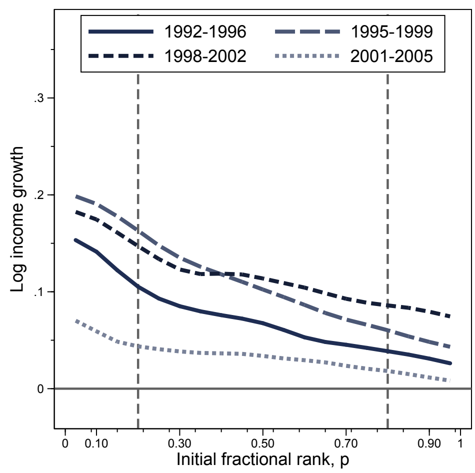

Figure B7. Income growth profiles (no imputed income subsample)

Note. Top panel shows absolute growth (change in real income in pounds); bottom panel shows proportionate growth (change in log real income).

Income growth profiles

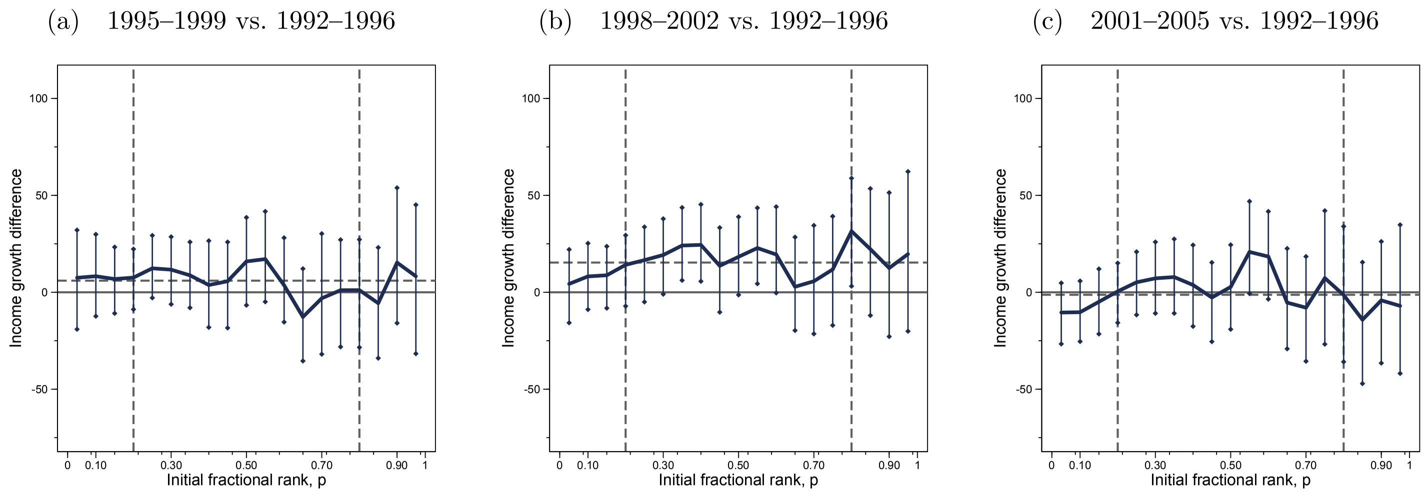

Cumulative mobility profiles

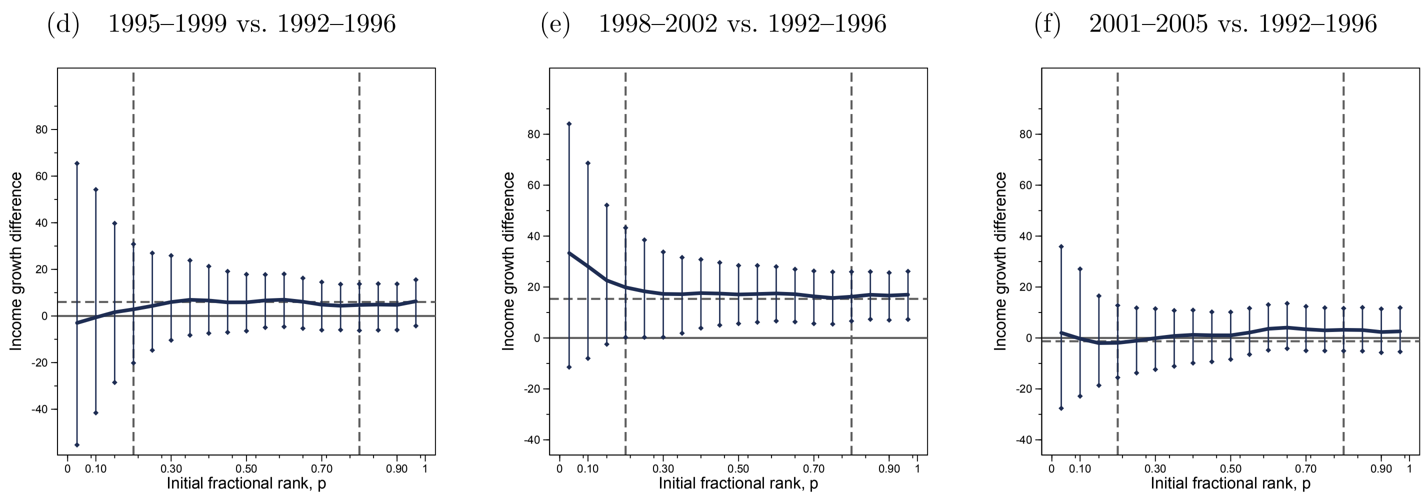

Figure B8. Absolute income growth profile difference (relative to 1992–1996) with pointwise variability bands (no imputed income subsample)

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

Income growth profiles

(a) 1995–1999 vs. 1992–1996

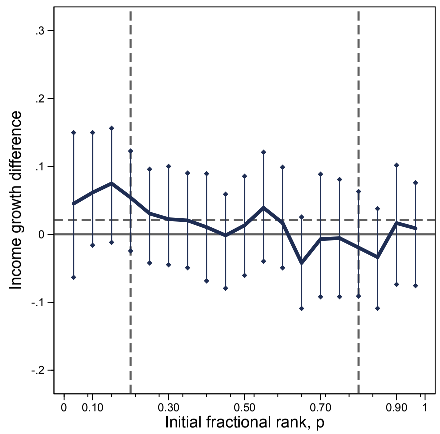

(b) 1998–2002 vs. 1992–1996

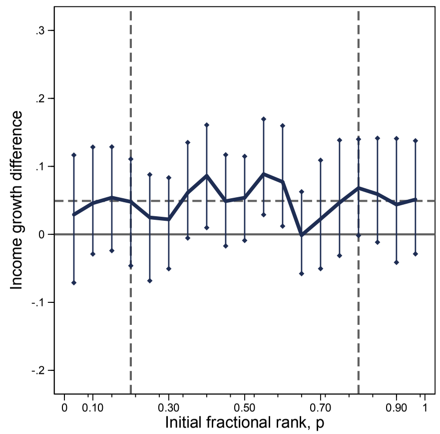

(c) 2001–2005 vs. 1992–1996

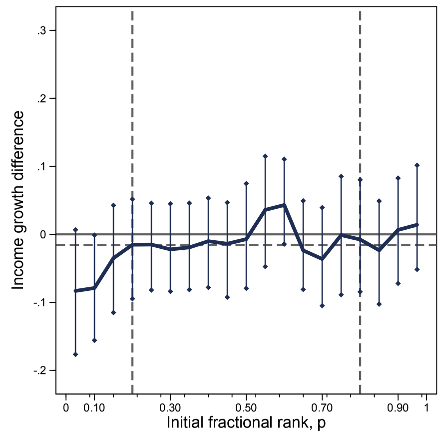

Cumulative mobility profiles

(d) 1995–1999 vs. 1992–1996

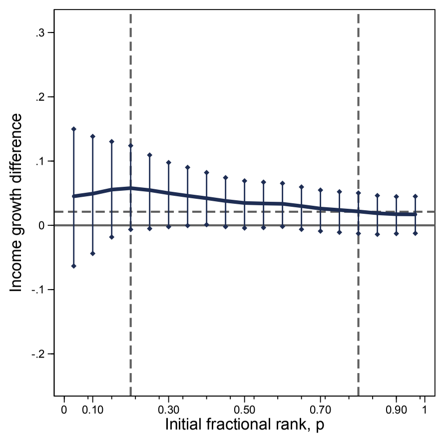

(e) 1998–2002 vs. 1992–1996

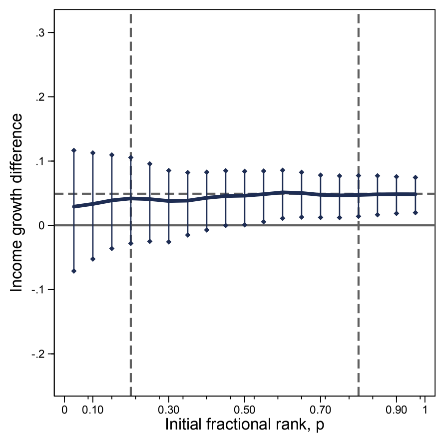

(f) 2001–2005 vs. 1992–1996

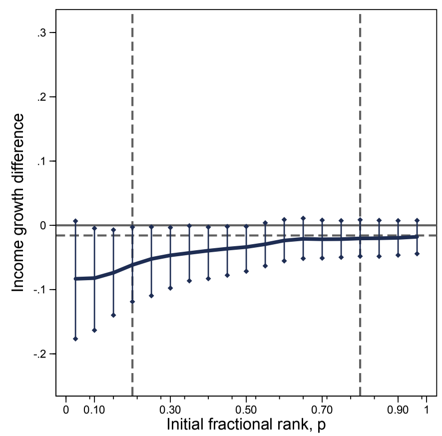

Figure B9. Proportionate income growth profile difference (relative to 1992–1996) with pointwise variability bands (no imputed income subsample)

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

Appendix C Accounting for age: sample ageing and life-cycle dynamics

Consideration of simple life-cycle earnings dynamics suggests a systematic association between income, income growth and age. While the association in terms of equivalized disposable income is much more complex than what simple earnings dynamics may suggest (Jenkins, 2009, 2011), it is a valid concern that income growth patterns, and differences in them between periods, partly reflect (i) the ageing of the panel sample and (ii) straightforward life-cycle patterns. We consider these two issues in this section.

Sample ageing

Incomes vary systematically with age (on average) and so, if the age composition of the sample changes over time, part of the *change* in the distribution of income growth across the four time periods considered may reflect this phenomenon. The BHPS design is intended to preserve representativeness and minimize the possibility of the sample ageing (through its following rules, and post-hoc using weights). However, we observe a gradual ageing of our samples over time: mean age increases from 38.26 in the 1992-93-94 sample to 42.72 in the 2006-07-08 sample. To reassure ourselves that sample ageing is not driving our results, we have also computed our estimates using 'standardized' samples. These are samples in which the age structure of the 1992–1996 subsample is imposed on all three later subsamples.[6]

Standardisation is implemented by reweighting subsample observations so that the age distributions for each subsample are the same as the distribution in the first of the four periods. The reweighting factor for observation i with age a_i in the subsample for the later period beginning in year s is defined as

$$\pi_i(s) = \frac{\Pr(\, a = a_i \mid t = 1992\,)}{\Pr(\, a = a_i \mid t = s\,)} \tag{23}$$

where $\Pr(\, a = a_i \,|\, t = s\,)$ is the relative frequency of observations with age a_i in the subsample for period $(s, s+\tau)$. Using Bayes' rule, the reweighting factor can be expressed equivalently as

$$\pi_i(s) = \frac{\Pr(\,t = 1992 \mid a = a_i\,)}{\Pr(\,t = s \mid a = a_i\,)} \frac{\Pr(\,t = s\,)}{\Pr(\,t = 1992\,)} \tag{24}$$

where $\Pr(\,t = s\,)$ is the proportion of observations from period $(s, s+\tau)$ in a pooled sample of 1992–1996 and period $(s, s+\tau)$ observations, and $\Pr(\,t = s \mid a = a_i\,)$ is the proportion of period $(s, s+\tau)$ observations in this pooled sample given age equal a_i. The reweighting factors are then multiplied by the sample weights and applied in all computations for periods other than 1992–1996. We estimated reweighting factors using expression (24) with the conditional probabilities computed from a probit model in which the covariates represent a linear spline in age with knots placed at ages 5, 12, 25, 35, 45, 60, and 70.

The estimates are displayed in Figures C1 and C2. Differences over time for age-standardized estimates are virtually identical to the non-standardized estimates that we presented earlier.

Adjusting for life-cycle income changes

In an elementary process of earnings dynamics over the life-cycle, incomes follow an inverted U-shape. This suggests that earnings grow faster when earnings are low over a person's lifetime profile and decrease when earnings are high. In this elementary model, a downward sloping 'progressive' income growth profile therefore arises because of simple life-cycle earnings dynamics. In practice, the pattern of lifetime income is not so elementary because incomes include many other sources than one's own earnings and because the individual earnings process is not characterized by smooth changes over time. Also, the income growth profile averages over a population of heterogenous individuals mixing heterogenous earnings processes of different levels and slopes. Nevertheless we investigate in this sub-section the effect of netting out individual life-cycle effects.

We consider three ways of controlling for life-cycle effects. First and most simply, we ran preliminary regressions of individual incomes on a polynomial in age and derive an 'age-corrected income' as the sum of income predicted at age 40 and the regression residual for each observation. This is done separately at each wave of the survey. Income

Income growth profiles

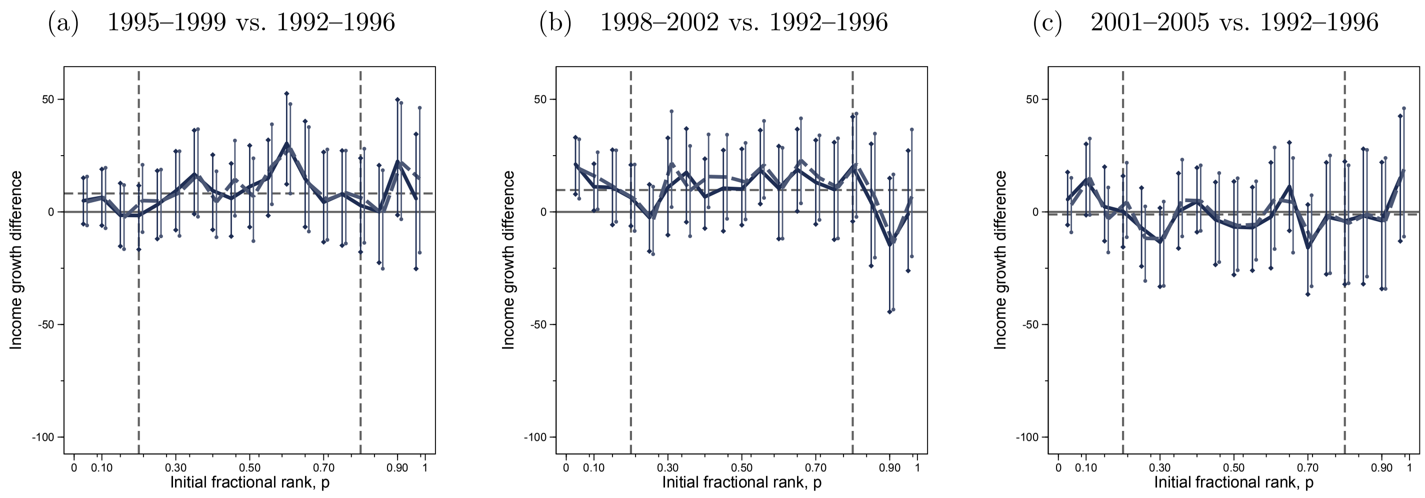

Cumulative income growth profiles

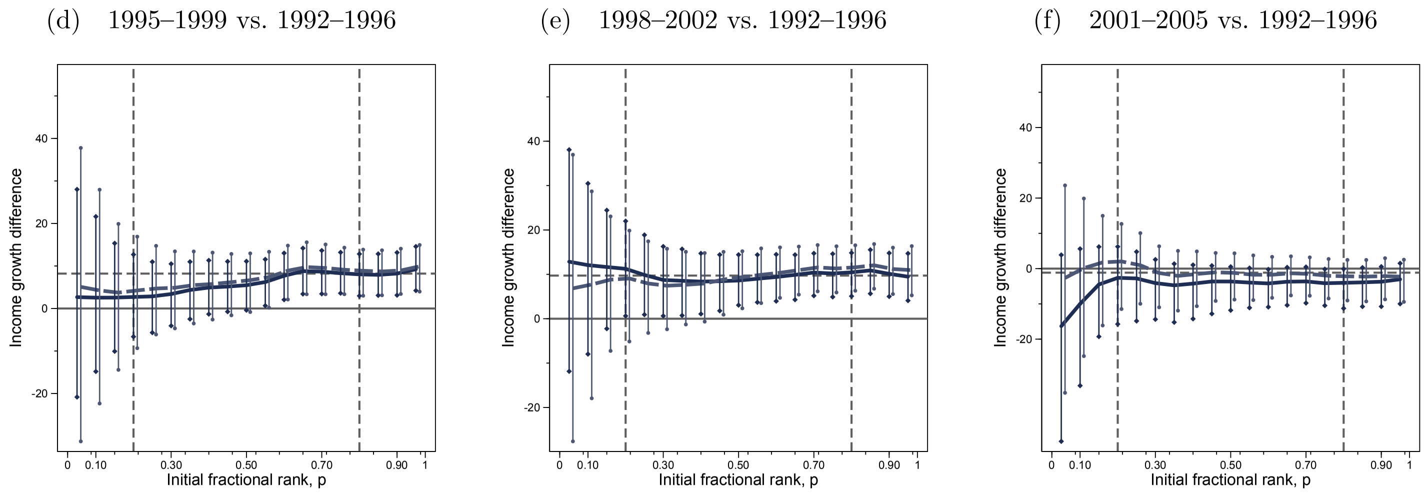

Figure C1. Differences in mobility profiles (relative to 1992–1996) with (dashed line) and without (solid line) age-standardization of subsamples, with 95% pointwise variability bands: absolute income growth in pounds

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

Income growth profiles

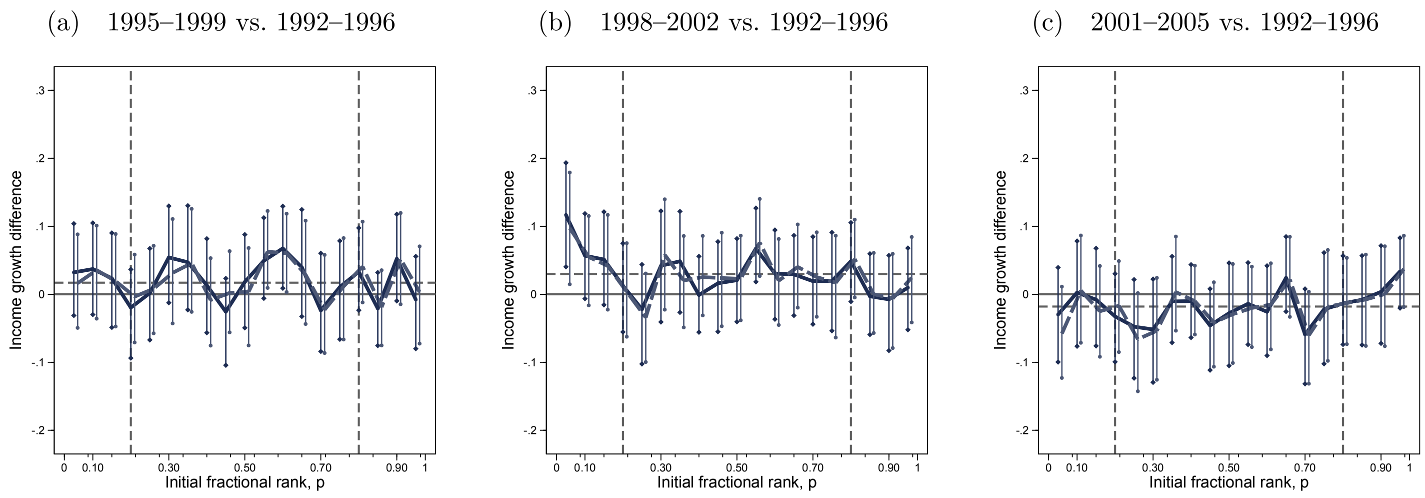

Cumulative mobility profiles

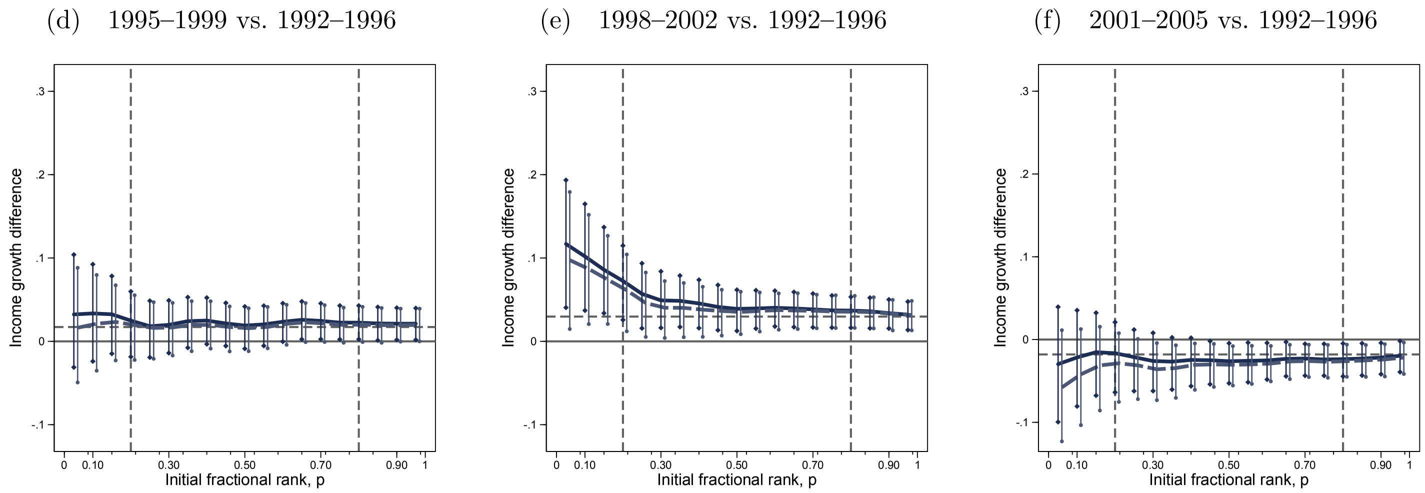

Figure C2. Differences in mobility profiles (relative to 1992–1996) with (dashed line) and without (solid line) age-standardization of subsamples, with 95% pointwise variability bands: proportionate income growth

Note. The horizontal dashed line in each graph shows the between-period difference in the overall average growth rate ($\bar{\delta}$).

growth profiles are then derived on the basis of age-adjusted income – that is, both initial ranks and income growth are 'age corrected' as per individual incomes at age 40. The second approach involves regression of individual income *growth* on a polynomial in age and construction 'age corrected income growth' as the sum of income growth predicted at age 40 and the regression residual. Income growth profiles then use unadjusted income to rank individuals but conditional income growth uses age-adjusted income growth where income growth is as per predicted at age 40. Finally, the third approach consists in ranking individuals within respondents of the same age group when deriving income growth profiles.[7] This procedure uniformly distributes individuals of different age over the initial income rank ordering, thereby avoiding any systematic age-related relationship between initial income rank and income growth. In this scenario, social weight is now assumed to be determined by the relative position of individuals *within* their age group. While the first approach is standard in the literature on earnings dynamics (Gottschalk and Moffitt, 2009, see, e.g.,), the third avoids additional regression assumptions but makes a normative adjustment to the measurement approach.

Estimates of income growth profiles using the three age-adjustments are provided in Figures C3, C4 and C5. None of the adjustments modify the shape of the income growth profiles significantly. The pro-poor nature of the patterns still hold and the contrast of the four time-periods is largely unaffected.

(a) Absolute income growth profiles

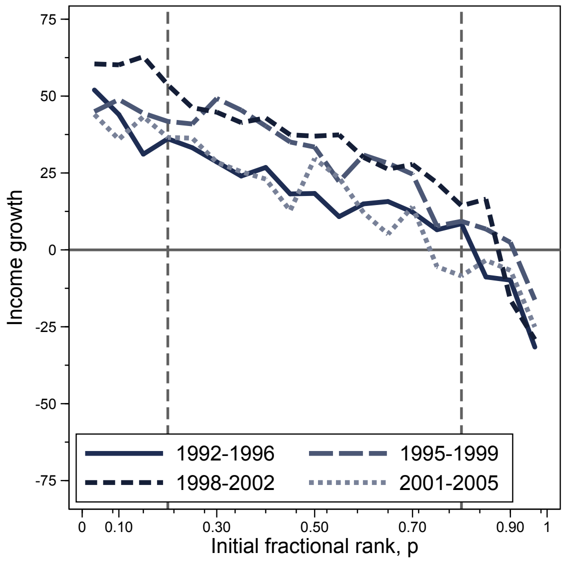

(b) Cumulative absolute income growth profiles

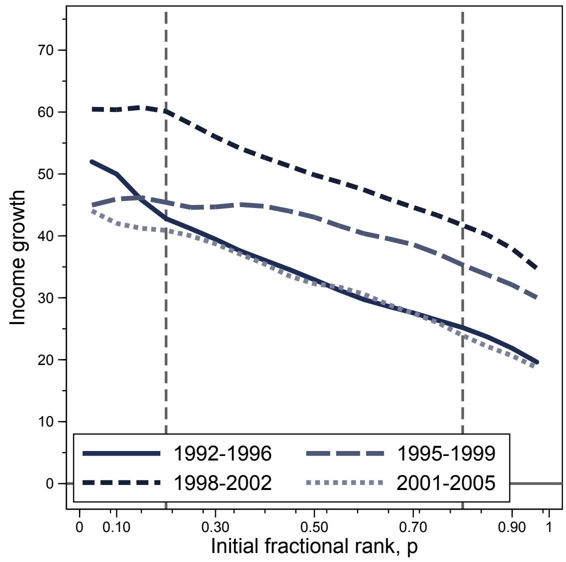

(c) Proportionate income growth profiles

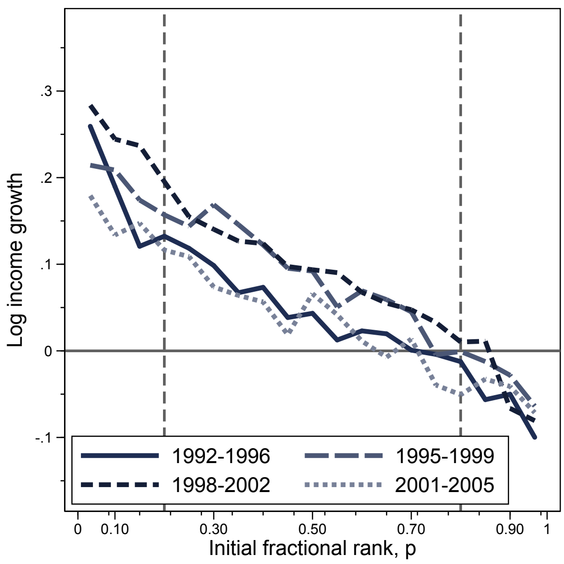

(d) Cumulative proportionate income growth profiles

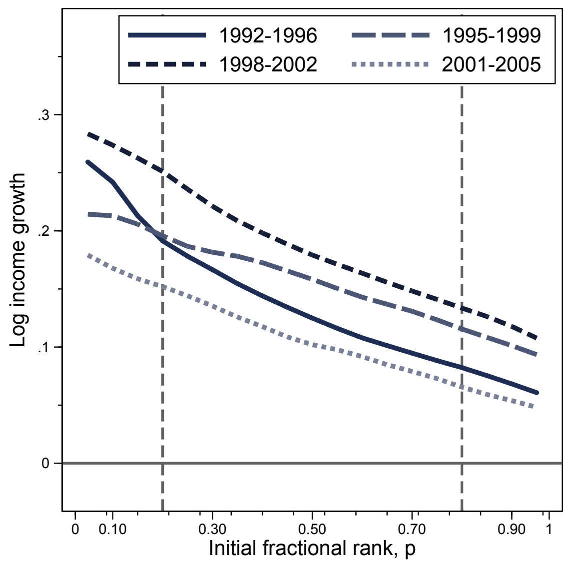

Figure C3. Income growth profiles adjusted for life-cycle effects using regression-based age-corrected individual incomes (predicted to age 40)

(a) Absolute income growth profiles

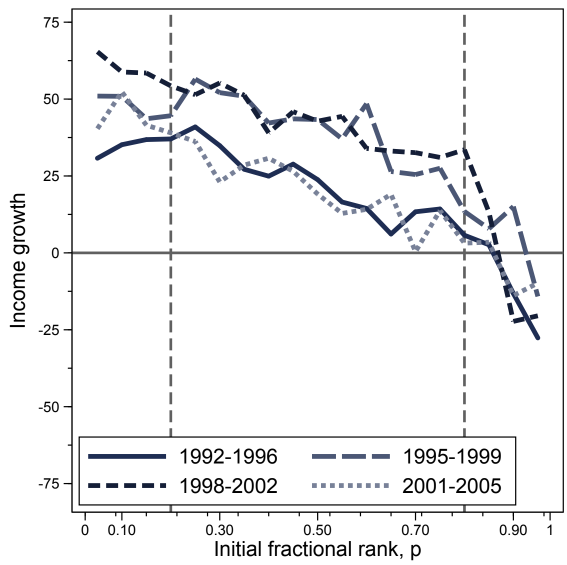

(b) Cumulative absolute income growth profiles

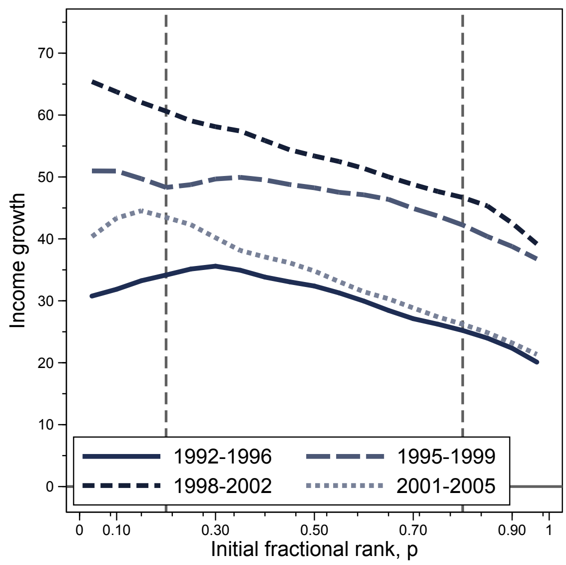

(c) Proportionate income growth profiles

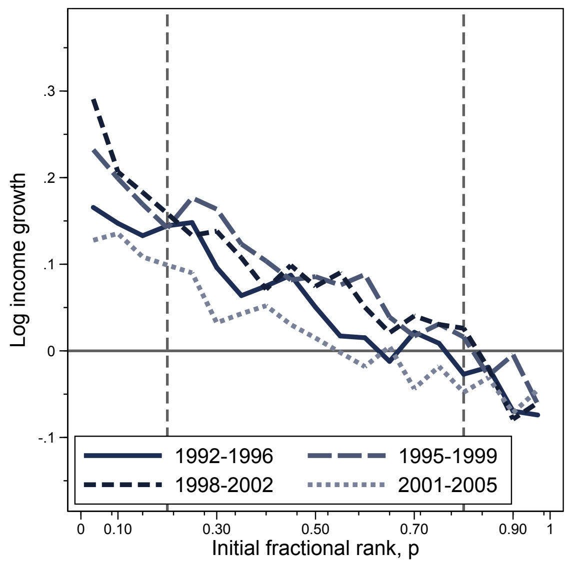

(d) Cumulative proportionate income growth profiles

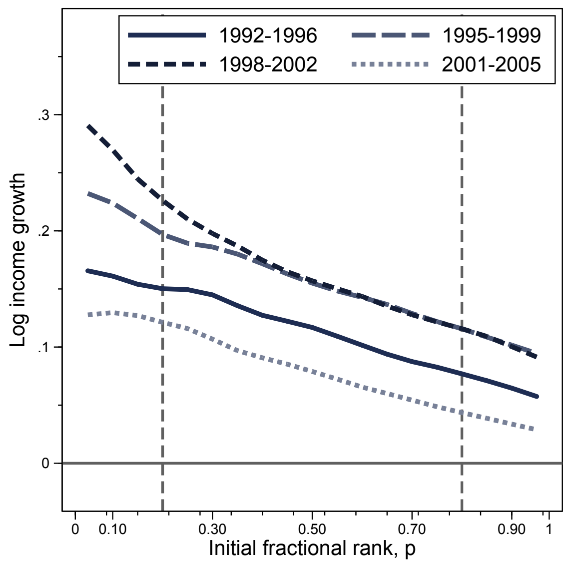

Figure C4. Income growth profiles adjusted for life-cycle effects using regression-based age-corrected income *growth* (predicted to age 40)

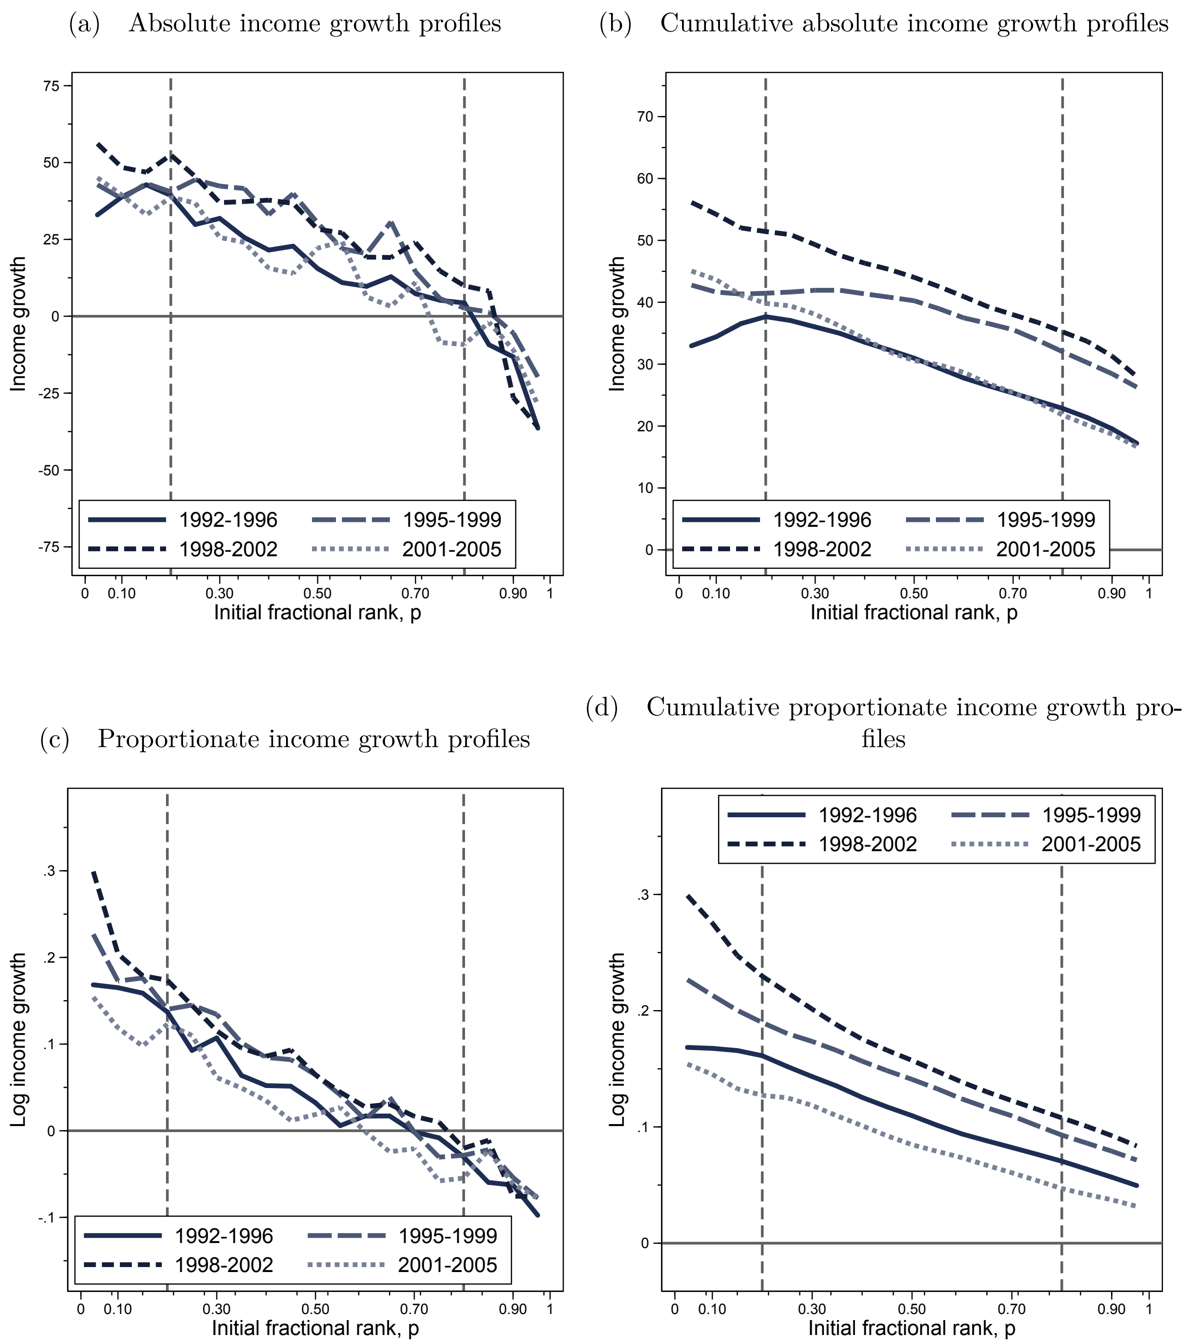

(a) Absolute income growth profiles

(b) Cumulative absolute income growth profiles

(c) Proportionate income growth profiles

(d) Cumulative proportionate income growth profiles

Figure C5. Income growth profiles adjusted for life-cycle effects using base-year rank computed *within* age groups

References

Bénabou, R. and Ok, E. A. (2001). Mobility as progressivity: Ranking income processes according to equality of opportunity. Working Paper 8431, National Bureau of Economic Research, Cambridge MA, USA.

Bourguignon, F. (2011). Non-anonymous growth incidence curves, income mobility and social welfare dominance. *Journal of Economic Inequality*, 9(4):605–627.

Chakravarty, S., Dutta, J., and Weymark, J. (1985). Ethical indices of income mobility. *Social Choice and Welfare*, 2:1–21.

Dardanoni, V. (1993). Measuring social mobility. *Journal of Economic Theory*, 61:372–394.

Dragoset, L. M. and Fields, G. S. (2006). U.S. earnings mobility: comparing survey-based and administrative-based estimates. ECINEQ Working Paper 2006–55, Society for the Study of Economic Inequality, Milan. Downloadable from `http://www.ecineq.org/milano/WP/ECINEQ2006-55.pdf`.

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

Gottschalk, P. and Huynh, M. (2010). Are earnings inequality and mobility overstated? The impact of non-classical measurement error. *Review of Economics and Statistics*, 92(2):302–315.

Gottschalk, P. and Moffitt, R. (2009). The rising instability of U.S. earnings. *Journal of Economic Perspectives*, 23(4):3–24.

Gottschalk, P. and Spolaore, E. (2002). On the evaluation of economic mobility. *Review of Economic Studies*, 69(1):191–208.

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

Jenkins, S. P. (2009). Spaghetti unravelled: a model-based description of differences in income-age trajectories. ISER Working Paper 2009-30, Institute for Social and Economic Research.

Jenkins, S. P. (2011). *Changing Fortunes: Income Mobility and Poverty Dynamics in Britain*. Oxford University Press, Oxford.

Lambert, P. J. (2001). *The Distribution and Redistribution of Income, a Mathematical Analysis*. Manchester University Press, Manchester, UK, 3rd edition.

Shaked, M. and Shanthikumar, J. G. (2007). *Stochastic Orders*. Springer-Springer+Business Media, New York.

Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica*, 55(1):99–115.

Yaari, M. E. (1988). A controversial proposal concerning inequality measurement. *Journal of Economic Theory*, 44(2):617–628.

Notes

[1]The distribution of the sum of initial and final period incomes is a function of the joint distribution H: $F_{XY}(y) = \int_0^y \int_0^t dH(t-s,s)ds\,dt$.

[2]This result is similar to that derived by Dardanoni (1993) for distributions summarised in terms of transition matrices. Assessment based on comparisons of the distribution of the sum of base- and final-period income against a replication of base-period income is also the starting point of Chakravarty et al. (1985).

[3]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.

[4]Gottschalk and Huynh (2010) fit a measurement error model to linked survey and administrative record data on the earnings of US men. See also Dragoset and Fields (2006) who use the same data, but a larger collection of mobility measures, and find that qualitative findings are similar for administrative and survey data. Quantitative findings often differ more substantially but no systematic pattern is found. Fields et al. (2003) and Grimm (2007) use a calibrated version of a model similar to that of Gottschalk and Huynh (2010) to place bounds on the impact of measurement error on estimates of the correlation between incomes in two years.

[5]The curves are somewhat steeper than in our baseline estimates, but note that the income change reported here is not based on 3-year smoothing. Comparable estimates not based on 3-year smoothing and without proxying base yeear rank are substantially steeper, as expected.

[6]This is similar to the technique of 'direct standardisation' that is commonly used in demography.

[7]To estimate conditional ('within group') ranks, data are first binned into age-intervals of 6 years and conditional ranks are estimated for observations within each bin using local kernel smoothing over the two adjacent bins.