Review of Income and Wealth
Series 48, Number 3, September 2002

THE ROLE OF MOBILITY IN OFFSETTING INEQUALITY: A NONPARAMETRIC EXPLORATION OF THE CPS

BY CATHERINE DE FONTENAY

University of Melbourne

TUE GØRGENS*

Australian National University

AND

HAOMING LIU

National University of Singapore

This paper explores how annual earnings mobility offsets annual earnings inequality, using matched CPS data. Mobility in the economy is estimated using nonparametric quantile regression, for which we adapt state-of-the-art smoothing techniques. Mobility is measured through the churning process (changes in earnings given initial earnings) in order to identify different mobility patterns for different earnings groups. For instance, upward mobility in high earners is far weaker than its converse, downward mobility for low earners. We assess the (positive or negative) contribution to offsetting of each pattern in mobility. Innovations in our approach also allow us to identify trends and minute changes in mobility, and to pinpoint which changes in mobility have offset the increases in inequality observed over the decades.

1. INTRODUCTION

Annual earnings inequality in the U.S. has rapidly increased in the past few decades.[1] Much research has been devoted to investigating whether the increase in annual earnings inequality has been accompanied by an off-setting increase in the mobility of individuals within the earnings distribution.[2] Mobility or churning is defined as changes in earnings (or alternatively, in each person's relative position in the earnings distribution) over time. In effect, given an observed level of yearly inequality, greater mobility of all individuals would imply less inequality over longer periods. However, increases in mobility for specific earnings groups may not have an offsetting effect (may not imply lower long-term inequality).

Note: The authors would like to acknowledge the support of the Australian Research Council and the research assistance of Thomas Walker. We thank David Ruppert for making his MATLAB programs available on the Internet and for clarifying details. We also thank Garry Barrett and Lisa Magnani for useful comments.

*Correspondence to: Tue Gørgens, Research School of Social Sciences, Australian National University, Canberra ACT 0200, Australia (Tue.Gorgens@anu.edu.au).

[1]See for instance Levy and Murnane (1992), Juhn, Murphy, and Pierce (1993), Brauer and Hickok (1995), Bishop, Formby, and Thistle (1997), Gottschalk and Smeeding (1997), and Johnson (1997).

[2]See for instance Gottschalk and Moffitt (1994), Gittleman and Joyce (1996), and Buchinsky and Hunt (1999).

This paper introduces nonparametric techniques for assessing the level and evolution of mobility for the U.S. population and for determining the extent to which mobility has offset inequality.[3]

To examine mobility, this paper looks at the churning process for the entire population, rather than considering individual-specific effects. We estimate nonparametrically the conditional distribution of earnings changes given initial earnings. The analysis is based on matched cross-sectional data from the Current Population Survey (CPS), which has important advantages and disadvantages. The size and representativeness of the CPS sample allows us to derive reliable nonparametric estimates. The major drawback is that only two years of observations can be matched for each individual in the CPS. Thus we can only relate outcomes across two years.

Mobility offsets inequality to the extent that earnings changes reverse the previous earnings distribution. Put simply, upward mobility for low earners and downward mobility for high earners (positive and negative earnings changes, respectively) are offsetting effects, even at a constant level of yearly inequality. Some of these offsetting effects take the form of nonlinearities and asymmetries in the distribution of earnings changes and are therefore difficult to capture in a parametric framework. Using nonparametric techniques allows us to identify which types of mobility offset inequality and which aggravate inequality. For instance, downward mobility was much stronger for the highest earners than upward mobility (more high earners had negative changes in earnings, and the negative changes were larger in size), implying that mobility tended to offset inequality at the upper end of the distribution. However, in addition to upward mobility, the lowest earners experienced significant downward mobility, which reduced total offsetting.

Moreover we identify *changes* in mobility for different earnings groups that have offset the trend of increasing inequality over the decades. For example, we find that the improved upward mobility of low earners in the 1970s helped offset their worsening earnings. Such offsetting changes are disentangled from the changes in mobility that account for increasing inequality.

We also identify changes in mobility from one year to the next, and over the course of the business cycle. One conclusion is that during economic downturns, the structure of mobility for the lowest earners worsens far more than for other groups. This phenomenon explains most of the increase in inequality during downturns. To reach these conclusions, we isolate the changes in inequality from year to year, by calculating the "steady-state" earnings distribution that would result if the current year's level of mobility were to persist indefinitely. This steady-state captures how changes in earnings from the previous year (the churning process) affect inequality.

The literature on the whole has taken a very different approach to evaluating mobility, focusing instead on a variance decomposition. In any period, each individual's earnings are viewed as the realization of a random variable, the mean

[3]Our methodology is related to that of Burkhauser *et al.* (1999), who compared the (nonparametrically estimated) densities of yearly household income in the U.S. and the U.K. across phases of business cycles. The nonparametric approach was informative, in that the increase in inequality was traced to a decline in the mass in the middle of the distribution, and an increase of mass in the upper tail of the distribution. However they did not study mobility.

and variance of which are specific to the individual and the period. Such a perspective involves no loss of generality. However, to estimate the mean and variance, substantial structure must be imposed across time or across individuals, and these assumptions may not be justified. Generally the effect of time is parameterized simply, and an individual's mean and variance are assumed otherwise constant.[4] Based on these estimates, the variance of population earnings can be decomposed into the variance of individuals' means plus the mean of individuals' variances. If individuals' means and variances are indeed constant, the terms above are measures of inequality and mobility, respectively (Gottschalk and Moffitt, 1994).

This approach is attractive in that it allows one to measure individual fixed effects, but it has drawbacks for studying the evolution of mobility. The variance decomposition framework is not amenable to studying year-to-year changes in mobility, because time is already assumed to enter in a specific form. Conversely, our approach of examining population dynamics from year to year allows us to examine even minute changes in mobility. The cost is that individual fixed effects cannot be considered.

The paper is organized as follows. Sections 2 presents the nonparametric methodology. Section 3 introduces the data and addresses the issues of possible bias arising from the CPS matching procedure. Section 4 presents results on the role of mobility in offsetting inequality, and points out which effects are captured through the use of nonparametrics. Section 5 examines changes in year-to-year mobility over the business cycle. Section 6 concludes.

2. Methodology

2.1. *Relationship between Inequality and Mobility*

To assess the extent of churning in the economy, we look at how individuals' earnings have changed depending on their initial earnings. The analysis centers on the conditional distribution of *changes* in log-earnings in year t, given log-earnings in year $t-1$; this is our measure of mobility from each earnings level in year $t-1$. Individual log-earnings in year t are denoted y_{it}, and changes in log-earnings are denoted $\Delta y_{it} = y_{it} - y_{i,t-1}$. We examine Δy_{it} rather than y_{it} because there is strong correlation between y_{it} and $y_{i,t-1}$, over the entire sample (the correlation coefficients range from 0.54 to 0.72). Consequently the relationship between y_{it} and $y_{i,t-1}$ is dominated by the 45-degree line, and other patterns are less clearly distinguishable. We consider log-earnings rather than earnings in order to maintain comparability with the literature (Gottschalk and Moffitt, 1994, among others). Considering log-earnings focuses attention on relative changes in earnings.

In examining the log-earnings distribution, we abandon the variance decomposition approach taken in previous studies and instead employ a

[4]For instance, strong autocorrelation in earnings is likely across periods where an individual holds the same job. Gottschalk and Moffitt (1994) noted that an individual's earnings tend to vary less during a job spell than between job spells. These findings can be partially accounted for by relaxing specific assumptions (Baker and Solon (forthcoming), for instance, allow for a specific form of autocorrelation in earnings), but the structure necessary for informative measures is still restrictive.

decomposition which focuses on the relationship between changes in log-earnings and initial log-earnings. Although we primarily discuss distributions in terms of quantiles, it is easier to understand the decomposition in terms of densities. Hence, let g_t denote the density of log-earnings in year t, and let f_t denote the conditional density of the change in year t log-earnings given year $t-1$ log-earnings. Then

$$g_t(y) = \int_{-\infty}^{\infty} f_t(y - v|v) g_{t-1}(v)\, dv. \tag{1}$$

Considering population distributions rather than variances and means of an individual's earnings will prove to be a very useful perspective. In particular, it is useful for inquiring to what extent mobility has offset earnings inequality: imagine a situation in which g_t and f_t are stable over time. Any such observed g_t distribution could have been generated by an infinite number of f_t distributions, each with different implications for longer-term inequality. If we restrict ourselves to the two-year perspective allowed by the data, each possible f_t would imply a different distribution of two-year inequality. In other words, each possible f_t has different offsetting effects.

2.2. *Nonparametric Quantile Estimation*

This section outlines the methodology for obtaining nonparametric estimates of conditional quantiles of Δy_{it} given $y_{i,t-1}$. Let the conditional distribution function of Δy_{it} at α given $y_{i,t-1} = \delta$ be

$$F_{\Delta y_{it}}(\alpha|\delta) = \Pr(\Delta y_{it} \leq \alpha | y_{i,t-1} = \delta). \tag{2}$$

The γ-th conditional quantile is defined by

$$Q_{\Delta y_{it}}(\gamma|\delta) = \inf\{\alpha\colon F_{\Delta y_{it}}(\alpha|\delta) \geq \gamma\}, \tag{3}$$

and unconditional quantiles are defined equivalently.

There are two main approaches to estimating conditional quantiles. The approach adopted here is to invert an estimate of the conditional distribution function. What limited evidence is available suggests that the choice of estimator is not crucial (Yu and Jones, 1998). The choice of approximation technique is more significant, however. There is evidence that local linear or higher order polynomial fits perform better at the boundary of the support of the independent variable than do local constant fits (Yu and Jones, 1997). The quantiles presented below are estimated by fitting local quadratic polynomials to the data.

Local polynomial regression was developed as a method for estimating conditional means, but it can be adapted to estimating the conditional distribution function. The distribution function can then be inverted to get estimates of the conditional quantiles, as in equation (3). The idea of local polynomial regression is as follows. Suppose that $(X_1, Y_1), \ldots, (X_n, Y_n)$ is a set of independent observations and that the objective is to estimate $E(Y_i|X_i = \delta)$ for a given δ. Let K be a kernel function, let b_δ be a (local) bandwidth, and define $K_\delta(u) = b_\delta K(u/b_\delta)$. To estimate the conditional mean consider fitting a p-th order polynomial to the data using weighted least squares with weights $K_\delta(\delta - X_i)^{1/2}$. Specifically, estimate $\mathrm{E}(Y_i|X_i = \delta)$ by $\hat{\beta}_{\delta,0}$, where $\hat{\beta}_\delta = (\hat{\beta}_{\delta,0}, \hat{\beta}_{\delta,p}, \ldots, \hat{\beta}_{\delta,p})'$ is the solution to

$$\min_{\beta_\delta} \sum_{i=1}^{n} (Y_i - \beta_{\delta,0} - \beta_{\delta,1}(X_i - \delta) - \cdots - \beta_{\delta,p}(X_i - \delta)^p)^2 K_\delta(\delta - X_i). \tag{4}$$

The well-known solution to this problem is $\hat{\boldsymbol{\beta}}_\delta = (\mathbf{X}'_\delta \mathbf{W}_\delta \mathbf{X}_\delta)^{-1} \mathbf{X}'_\delta \mathbf{W}_\delta \mathbf{Y}$, assuming the inverse exists, where

$$\mathbf{Y} = \begin{bmatrix} Y_1 \\ \vdots \\ Y_n \end{bmatrix}, \quad \mathbf{X}_\delta = \begin{bmatrix} 1 & (X_1 - \delta) & \cdots & (X_1 - \delta)^p \\ \vdots & \vdots & \ddots & \vdots \\ 1 & (X_n - \delta) & \cdots & (X_n - \delta)^p \end{bmatrix}, \quad \mathbf{W}_\delta = \mathrm{diag} \begin{bmatrix} K_\delta(\delta - X_1) \\ \vdots \\ K_\delta(\delta - X_n) \end{bmatrix}.$$

If $p = 0$ (local constant fit), the estimator reduces to the usual Nadaraya–Watson kernel estimator. Using $p > 0$ has been shown to have several advantages, including better behavior of the estimator when δ is near the boundary of the support of X_i.[5]

To estimate the conditional distribution function of Δy_{it} given $y_{i,t-1} = \delta$, let $1(\cdot)$ denote the usual indicator function, that is, $1(A) = 1$ if the event A is true and $1(A) = 0$ otherwise. Exploiting the fact that

$$\Pr(\Delta y_{it} \leq \alpha | y_{i,t-1} = \delta) = \mathrm{E}(1(\Delta y_{it} \leq \alpha) | y_{i,t-1} = \delta), \tag{5}$$

we estimate $F_{\Delta y_{it}}(\alpha|\delta)$ using local polynomial regression by substituting $1(\Delta y_{it} \leq \alpha)$ for Y_i and $y_{i,t-1}$ for X_i in the above derivation. Having obtained an estimate, $\hat{F}_{\Delta y_{it}}$, we estimate the conditional quantiles by[6]

$$\hat{Q}_{\Delta y_{it}}(\gamma|\delta) = \inf\{\alpha : \hat{F}_{\Delta y_{it}}(\alpha|\delta) \geq \gamma\}. \tag{6}$$

To implement the estimator it is necessary to choose K, b_δ, and p. The choice of the kernel function has been shown to be inconsequential. We use the quartic kernel function $K(u) = 1(|u| \leq 1)(15/16)(1 - u^2)^2$. The remaining problem is to choose p and b_δ. No optimal data-driven bandwidth selector has been developed for the conditional quantile estimation problem. The only guidance is provided by Yu and Jones (1998), who suggested appropriate rule-of-thumb modifications for adapting a state-of-the-art bandwidth selector developed for the problem of estimating the conditional mean. We follow their suggestion and use Ruppert's (1997) empirical-bias bandwidth selector (EBBS) for local polynomial mean regression. The EBBS is quite complicated and requires the selection of many "tuning" parameters, and we refer to Ruppert (1997) for details of the procedure. To study the performance of the estimator for various choices of p, tuning parameters of EBBS, and the rule-of-thumb modifications, we carried out a set of Monte Carlo experiments using a data generating process with properties similar to our log-earnings data. We found that $p = 2$ yielded good results for all quantiles and that the rule-of-thumb modifications yielded higher mean squared errors than using the unmodified bandwidths. Accordingly, the estimates presented here are

[5] See for example the articles by Fan (1992) and Ruppert and Wand (1994) or the monograph by Wand and Jones (1995).

[6] As pointed out by Yu and Jones (1998), the fact that $\hat{F}_{\Delta y_{it}}$ is not necessarily bounded between zero and one does not cause problems for computing $\hat{Q}_{\Delta y_{it}}$.

local quadratic fits using the same EBBS bandwidth b_δ at the point δ for estimating conditional quantiles as for estimating the conditional mean.[7]

Our analysis is mainly graphical. In order to indicate the degree of accuracy we have calculated pointwise confidence bands for all the estimates of inequality and changes in earnings. Confidence bands are shown in Figures 5, 6, 11 and 13, but have been omitted in other graphs where they would obscure features in main curves. Omitted confidence bands are of similar magnitude to those shown. Note that all of the conclusions drawn in the paper are sufficiently broad that they are not dependent on the exact size of the confidence bands. The confidence bands are computed using a simple bootstrap with 500 draws and undersmoothing to reduce the asymptotic bias of the bands.[8]

3. Data and Data Problems

3.1. *Data*

The data examined in this paper are drawn from the Current Population Survey (CPS) March Annual Demographic Supplement files. Once an address is selected for the CPS, the household at the address is in the survey for four months, exits for eight months, and then returns for an additional four months. This sample design allows matching of data from a household in its first March interview with data from the household at the same address in its second March interview the following year, so long as the household members have not moved. Thus our matched data constitute a sequence of overlapping two-year panels from the survey years 1968–99, corresponding to the work years 1967–98. Due to changes in the household identifier the survey years 1971–72, 1972–73, 1976–77, and 1985–86 could not be matched.

Our earnings variable is annual wage earnings of the individual (wage and salary income) deflated by the Consumer Price Index.[9] Our sample includes all prime-age males with positive annual earnings in both years. Individuals with zero earnings in either year tend to be either self-employed or out of the labor force (not looking for work): in the 1987 CPS, 60 percent of those reporting zero earnings were self-employed, and an additional 30 percent did not look for work.

[7] The full EBBS algorithm has three steps: (1) Estimate the mean function and compute squared residuals using an *ad hoc* small bandwidth. (2) Estimate the kurtosis coefficient from standardized squared residuals; compute EBBS-bandwidths and estimate the variance function from standardized squared residuals. (3) Reestimate the mean using EBBS-bandwidths with the variance estimate from step 2.

[8] To achieve undersmoothing, the bandwidths were scaled by a factor $1.2n^{-0.05}$ (approximately 0.77). Since the bands contain less bias than the main curves, comparing the two provides an indication of the bias in the curves. When evaluating curve estimates and confidence bands, note that estimates computed on the same sample are correlated across quantiles.

[9] The CPI has been criticized for overstating inflation, among other shortcomings. The CPI is used as a deflator here since no commonly accepted alternative is available, although some alternatives with desirable properties do exist, such as the PCE (Personal Consumption Expenditures) price index or the CPI-X1, which treats housing costs more consistently. Real earnings measured under these alternative deflators exhibit more growth over time. Note however that the choice of deflator has no effect on earnings spreads in terms of log-earnings. Also, the choice of deflator has a negligible effect on the estimates of mobility for any given year: the annual inflation rates differ between estimators by less than 0.025, which is not visible on the scale of the graphs presented.

Given that our aim is to analyze changes in people's labor market outcomes (as opposed to changes in total income), and to make comparisons with the literature on that subject, we and others exclude groups with a marginal attachment to the labor market, such as these.[10] On this basis we also exclude the young and the old, who are likely to be students or semi-retired: our sample is restricted to respondents between 25 and 54 years of age (inclusive) in the first year.[11]

Earnings data are subject to topcoding in almost every year, and various components of earnings are topcoded separately. Therefore there is a risk of underestimating the change in log-earnings for the very highest earners. However, the degree of topcoding is small and varies from year to year, ranging from no topcoded observations in (earnings year) 1984 to just under 4 percent in 1994.[12] Given that the quantiles of log-earnings changes for high earners are not appreciably different in shape in 1984 versus other years (as we will show), we conjecture that the bias is not severe.[13]

The resulting matched samples include 4,532 to 9,077 observations in every two-year period. The size of the samples allows us to derive accurate nonparametric estimates of the conditional distribution of earnings, given past earnings. Furthermore, the estimated conditional distribution is representative for the U.S. population, given that the CPS covers the entire civilian population. Finally, the CPS has a significant advantage in that it covers a long period, making it possible to track mobility since the 1960s.

In contrast, most empirical research on earnings instability is based on two longitudinal datasets. Of these, the subsample of the Michigan Panel Study on Income Dynamics (PSID), used by Gottschalk and Moffitt (1994), has a comparable sample of only 2,730. The National Longitudinal Survey of Youth (NLSY), used by Buchinsky and Hunt (1999), is restricted to individuals aged 14 to 22 years old in 1979 and is therefore not representative of the entire population. The obvious advantage of these longitudinal datasets is that earnings patterns can be studied at the individual level over an extended period of time. Thus datasets are better suited to the variance decomposition analysis prevalent in the literature.

[10]Apart from the exclusion of women, this sample is similar to the "positive earnings sample" studied by Gittleman and Joyce (1996), which consisted of everybody (both sexes) who had "positive earnings in both years, who were not self-employed in either year, and who were 25–55 years of age in the first year. These authors also studied a "full time year round" sample. They found less mobility in the latter sample, but otherwise results were qualitatively similar to the positive earnings sample. Our sample is also comparable to the PSID subsample studied by Gottschalk and Moffitt (1994) which consisted of white males who had positive earnings, who were 20–59 years of age sometime during 1968–88, and who were not students.

[11]We estimated mobility from a larger sample, of age range 25–59, to check the sensitivity of results to our choice of age range. There was little difference in results.

[12]Due to inflation, the degree of topcoding tends to decrease over time. The topcoding level was adjusted upwards in survey years 1982, 1985 and 1996, corresponding to earnings years 1981, 1984 and 1995.

[13]The question arises of whether the increases in mobility identified in this paper are driven by increases in measurement error. Fortunately there is no evidence to suggest that measurement error in the CPS data increased over time; if anything, measurement error should have decreased, as the survey was computerized in 1994. Moreover, we did not find a significant "structural break" in 1994, which suggests that measurement error has not strongly affected (or biased) the results.

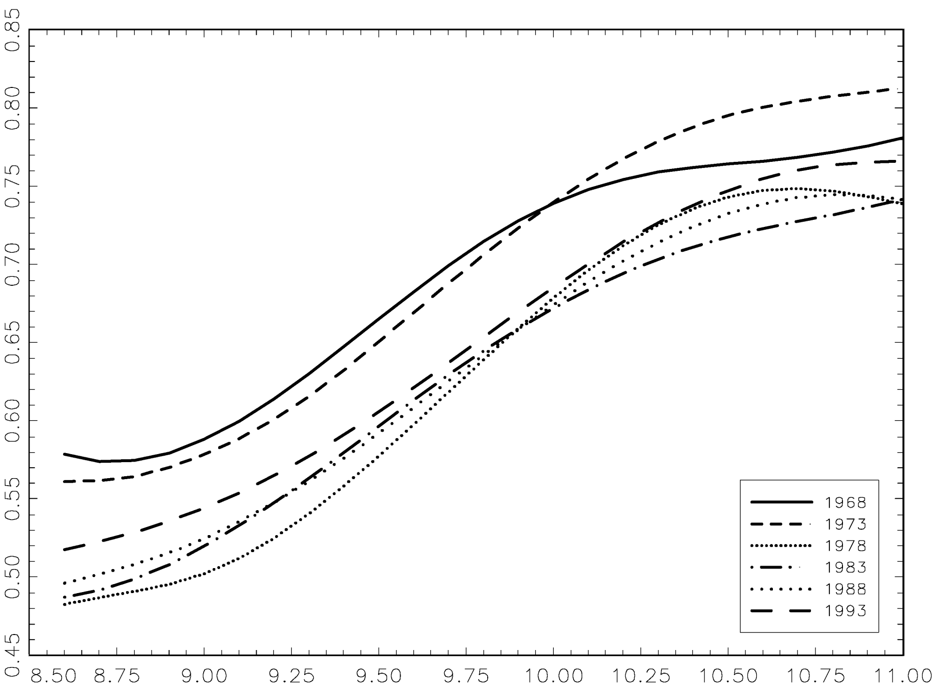

Figure 1. Conditional Probability of Matching with Following Year Given y_{it} Plotted Against y_{it}

3.2. *Evaluating the Effect of Matching Bias in the Data*

The CPS data do not contain consistent unique personal identifiers across years. Following the suggestion of the CPS documentation, we adopt a two-step matching procedure. We first match household units over the two year period. Next, persons within matched households are matched based on personal line number (only available between 1964–75 and 1979–99), race, sex, age, relation to household heads, and education level. Individuals who moved out of sampled households were not tracked, and as a result the matching rate is low. The matching rate is 65 percent on average, ranging from 47 percent (those with a first interview in 1975) to 77 percent (those with a first interview in 1970). The computer program is based on the Stata program written by Finis Welch.[14] Peracchi and Welch (1995), in their study of the matched versus unmatched sample, similarly found an average match rate of 2/3 for their 1979–90 sample, and large annual fluctuations.

Unfortunately, matching is not independent of earnings, which raises the question of bias. Figure 1 plots for selected years the estimated percentage of individuals who completed a first interview and who were matched with a second interview, conditional on their first year earnings.[15] (We obtained similar shapes

[14]The Stata matching program is described in Stata Bulletin 12, and can be downloaded from http://www.stata.com/stb/stbl2/dmll.

[15]The estimates were computed using local quadratic fits with the EBBS bandwidths, as described in Section 2.2.

in plots of the estimated percentage of individuals who completed a second interview and who were matched with a first interview in the previous year, conditional on their second year earnings. These plots are omitted for brevity.) The figure confirms that matching is not independent of earnings. Individuals experiencing low earnings are far more likely to have changed address (or refused to be interviewed) in the previous period or to change address in the following period.

The relationship between earnings and matching means that the earnings distributions and the inequality estimated from the matched subsample are biased. The question is whether the estimates of mobility are biased as well. The fact that the matched subsample does not have the same earnings distribution as the population does not necessarily imply that the mobility estimates will be biased, as we are estimating mobility conditional on initial earnings.

The population analyzed in this paper consists of individuals with close attachment to the labor force, in particular individuals who have positive earnings in two consecutive years. Information about this population is available only in the matched sample, and we have therefore no direct way of assessing whether our mobility estimates are biased. However, the marginal earnings distributions for the "total" population, including individuals with zero earnings in one or two years, can be estimated from both the corresponding "total" sample and from the matched subsample. To estimate the bias resulting from imperfect matching, we compare the actual distribution of year t earnings with the predicted distribution generated as in equation (1), from the marginal earnings distribution in year $t-1$ for the "total" sample and the conditional distribution of earnings changes given year $t-1$ earnings estimated for the matched subsample. If the mobility estimates are nearly unbiased, the actual and the predicted distributions should be very close.

To implement this comparison, our estimates of the conditional distribution of log-earnings changes for individuals with positive consecutive earnings must be augmented by estimates for those respondents with zero earnings in year $t-1$ as well as by estimates of the conditional probability of earnings in year t being zero. Mathematically, let γ_t^{all} denote the probability of zero earnings among all respondents in year t and let g_t^{all} denote the subdensity of log-earnings (so g_t^{all} integrates to $1-\gamma_t^{\text{all}}$). Further, let $\hat{f}_t$ denote the conditional distribution of changes estimated on the matched subsample with $\hat{f}_t(-\infty|v)$ indicating the probability of transition to zero earnings conditional on initial log-earnings being v, and let $\hat{\phi}_t^{\text{all}}$ denote the conditional density of year t log-earnings for respondents with zero earnings in year $t-1$ with $\hat{\phi}_t(-\infty)$ indicating the probability of transition to zero earnings. The predicted earnings distribution is then calculated as

$$\begin{aligned} \hat{\gamma}_t^{\text{all}} &= \int_{-\infty}^{\infty} \hat{f}_t(-\infty|v) g_{t-1}^{\text{all}}(v)\, dv + \hat{\phi}_t(-\infty)\gamma_{t-1}^{\text{all}} \\ \hat{g}_t^{\text{all}} &= \int_{-\infty}^{\infty} \hat{f}_t(y-v|v) g_{t-1}^{\text{all}}(v)\, dv + \hat{\phi}_t(y)\gamma_{t-1}^{\text{all}}. \end{aligned} \tag{7}$$

The actual and predicted distributions are remarkably close together, as illustrated in Figure 2. The results suggest that the conditional distributions estimated

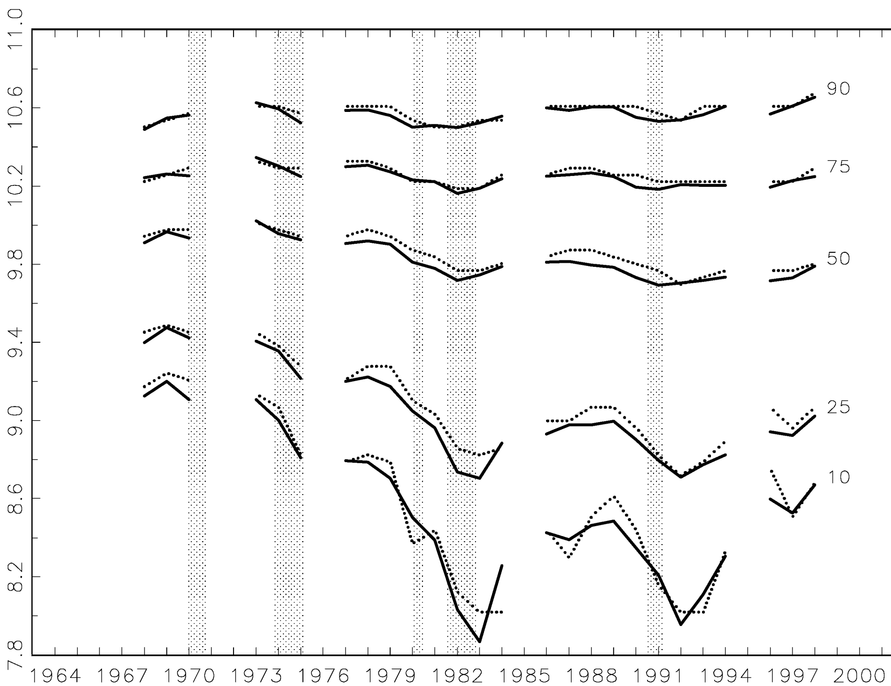

Figure 2. Percentiles of Actual (Solid) and Predicted (Dotted) y_{it} for the "Total" Sample Plotted Against t

on the matched sample are very similar to the distribution for the "total" sample. Thus there seems to be little bias in our results due to the matching procedure.[16]

4. Results: Mobility Offsetting Inequality

4.1. *Inequality*

While inequality in a year can be measured through a number of summary statistics, it is most informative to consider the full cross-sectional distribution of log-earnings. Figure 3 shows selected quantiles of the log-earnings distributions over the period 1963–98. The solid lines represent estimates for the main population of interest, males aged 25 to 54, restricted to individuals with positive earnings in two consecutive years (the "matched++" sample). Recall that low earners are less likely to be matched and the number of low earners is therefore underestimated for this population. The dotted lines in Figure 3 indicate the quantiles for the larger population of all those with positive earnings in the given year (the "all + " sample). No matching procedure is needed to estimate quantiles for this

[16]Our findings agree with those of Peracchi and Welch (1995). They found household heads of unmatched households were younger, less educated, and less likely to be employed, and that the weekly wage was lower in the unmatched sample. However, they found no systematic bias in the estimates of the transitions between labor force states after controlling for sex, age and initial labor force status.

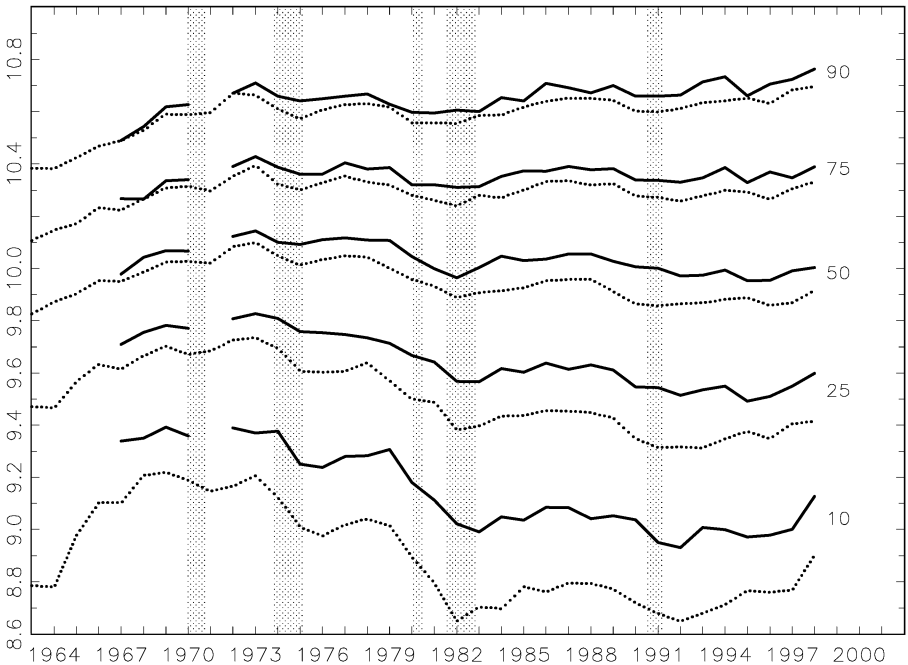

Figure 3. Percentiles of y_{it} for the "Matched++" (Solid) and "All+" (Dotted) Samples Plotted Against t

population, so these estimates are not biased, but this larger population includes individuals who had zero earnings in the other year, and consequently includes too many low earners. Quantiles for the population of interest lie in between the quantiles of these two distributions. Fortunately, the evolution of the quantiles in the two samples is very similar, suggesting a similar pattern for the population of interest.

Estimates for the "all+" sample are available for a longer period than for the "matched++" sample, thus allowing us to see more of a cyclical pattern in the earnings distribution. The estimates for the longer period highlight that the late 1960s and early 1970s was an unusual period, with high earnings and little inequality. In fact all quantiles but the 90th reached their maximum in 1973. In the "all+" sample, the median was about 10 percent higher in 1998 than in 1963, but 18 percent lower than in 1973, and approximately the same as in 1967.[17] Comparison with the timing of recessions (which are shaded in this and subsequent figures), as defined by the NBER, suggests that the quantiles fluctuated from year to year roughly in line with the business cycle.

The spread of the earnings distribution increased substantially over the period, implying a large increase in inequality, as found by other researchers. Figure

[17]Real earnings show more growth over the period under alternative deflators such as the CPI-XI or the PCE price index. However the period around 1973 remains a period of unusually high earnings, and a maximum for the lower quantiles.

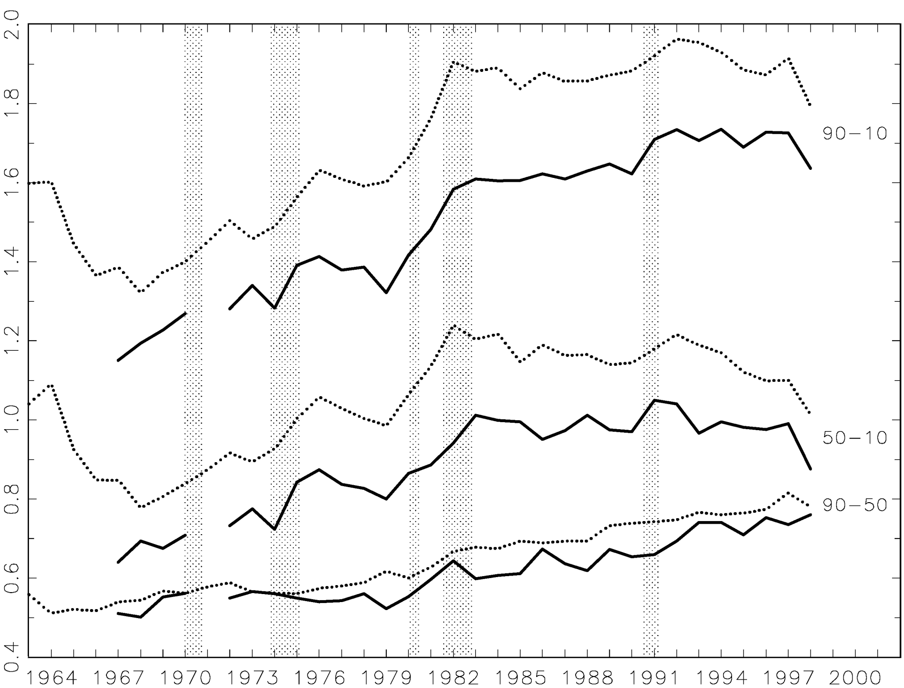

Figure 4. Percentile Spreads of y_{it} for the "Matched++" (Solid) and "All + " (Dotted) Samples Plotted Against t

4 shows the 90–10, 50–10, and 90–50 percentile spreads of the log-earnings distribution for the two samples. The 90–10 percentile spread of earnings fell from 1963 to 1968, then showed a strong increase from 1968 to 1982, and stabilized thereafter. The composition of this spread increase is clear from the other two curves: the 90–50 spread increased steadily over the entire period, although the trend was perhaps stronger after 1982. The 50–10 spread followed a pattern resembling a periodic cycle, returning by 1997 to nearly its 1963 level. There was a strong downswing in the 50–10 spread to its low in 1968 and increase until about 1982, and then the second half of the cycle is flatter. Note also that the fluctuations in the 50–10 spread, which loosely followed the business cycle, accounted for most of the year-to-year variation in the 90–10 spread, as the 90–50 spread was much more steady. Earnings of the lowest quantile seem to be most sensitive to the business cycle.

Based on these figures, the evolution of inequality seems to fall into several well-defined periods: a period of lessening inequality, from 1963 to 1967, and two periods of increasing inequality, from 1967 to 1982, and from 1982 onwards. Increasing inequality from 1967 to 1982 was primarily marked by a widening of the 50–10 spread, which corresponds in Figure 3 to a fall in earnings for the lowest earners (both in absolute terms, and relative to other quantiles). Increasing inequality after 1982 was primarily accounted for by the widening 90–50 spread, which corresponds in Figure 3 to an increase in earnings of the highest earners (both absolute and relative).

4.2. *Introduction to Mobility*

Figure 5 plots the quantiles of one-year individual changes in log-earnings against time. (Recall that these estimates suffer from "attrition" bias due to underrepresentation of low earners in the matched data; but estimates conditioned on initial earnings are not biased.) The figure indicates a remarkable widening in the distribution over time: there was a slight widening between the 25th and 75th percentile, and a substantial widening between the highest and lowest percentiles. The 90–10 spread almost doubled over the period; thus at least a fifth of the changes were twice as large in the later years. In other words, mobility greatly increased over the period, as more people saw larger changes in their earnings in the later years.

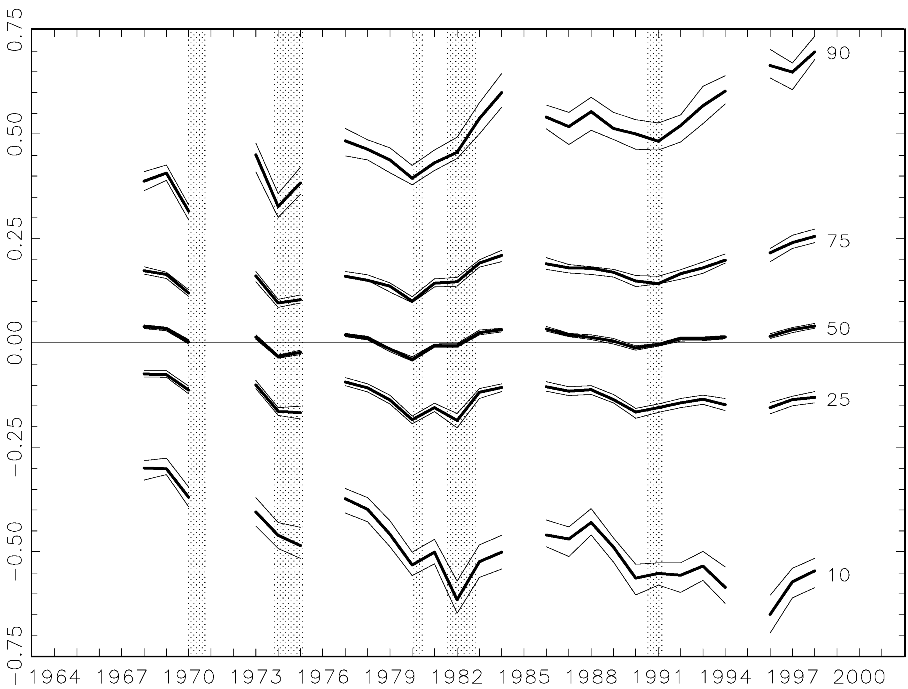

Figure 5. Percentiles of Δy_{it} Plotted Against t, with 95% Confidence Bands

Interestingly, both upward and downward mobility increased. Over the entire period the upper and lower quantiles were highly symmetric, and the median was steady near zero; thus the 90–50 and the 50–10 spreads were nearly identical throughout the period. The median annual earnings increases were between –4 and 5 percent.

The consistent symmetry of the distribution would seem to suggest that the distribution of changes was symmetric for most subsets of the population, and that mobility could be measured as a variance around a mean (as in the variance decomposition literature). However, this impression is belied by the distribution of changes conditional on initial earnings, which varies greatly by initial earnings, and is skewed at many levels.

4.3. *Characteristics of Mobility and Their Offsetting Role*

The conditional quantiles of changes in earnings, given earnings in the previous year, are plotted in Figure 6 for 1984. No observations were topcoded in 1984, and therefore it seems an appropriate year in which to consider the complete distribution of changes, and describe mobility's overall characteristics. However, we will show that the shape of the distribution was virtually identical in all years, including in years in which topcoding was more pronounced; therefore it is sufficient to examine 1984.

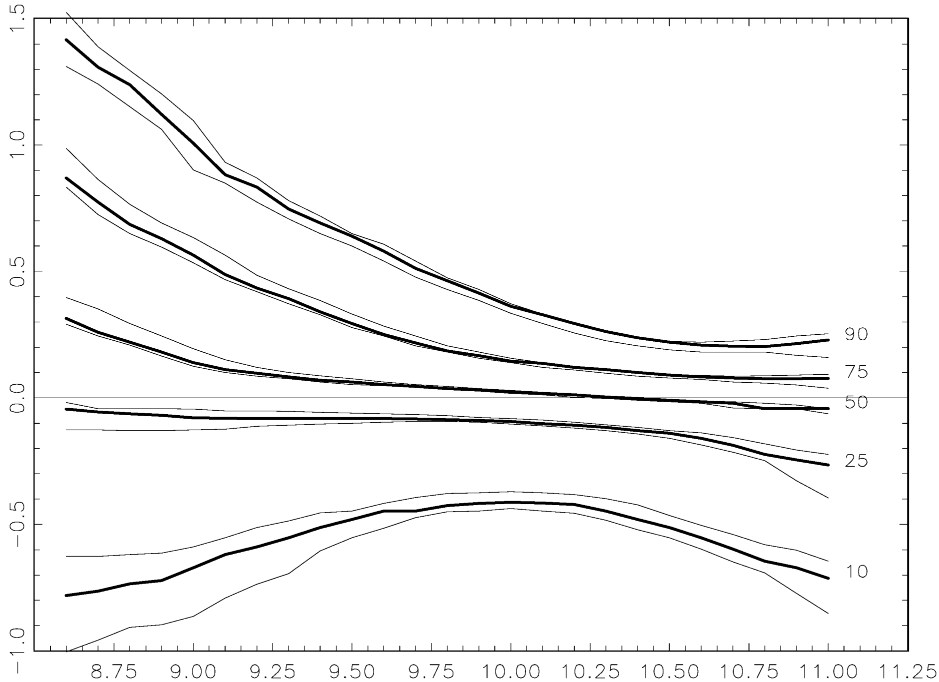

Figure 6. Conditional Percentiles of Δy_{it} Given $y_{i,t-1}$ for 1984 Plotted Against $y_{i,t-1}$, with 95% Confidence Bands

Broadly speaking, mobility offsets inequality to the extent that earnings changes move low earners upwards and high earners downwards. The term we adopt to describe this process is "reversion," by which we mean the extent to which the current year's inequality is reversed by the subsequent change in earnings. Figure 6 exhibits significant reversion. The median sloped downwards and reached negative levels for high values of initial log-earnings. (This pattern is confirmed by Figure 7, which plots medians at ten-year intervals.) Thus the majority of individuals with low earnings in any given year experienced substantial increases in the following year, and the majority of individuals with high earnings experienced decreases, ranging from slight to large. Consider those with earnings at or below 9.0 in 1983, roughly the tenth percentile. The 10 percent with highest changes saw increases of more than 100 percent in their earnings. The median change was large and positive, ranging from about 15 to 30 percent.

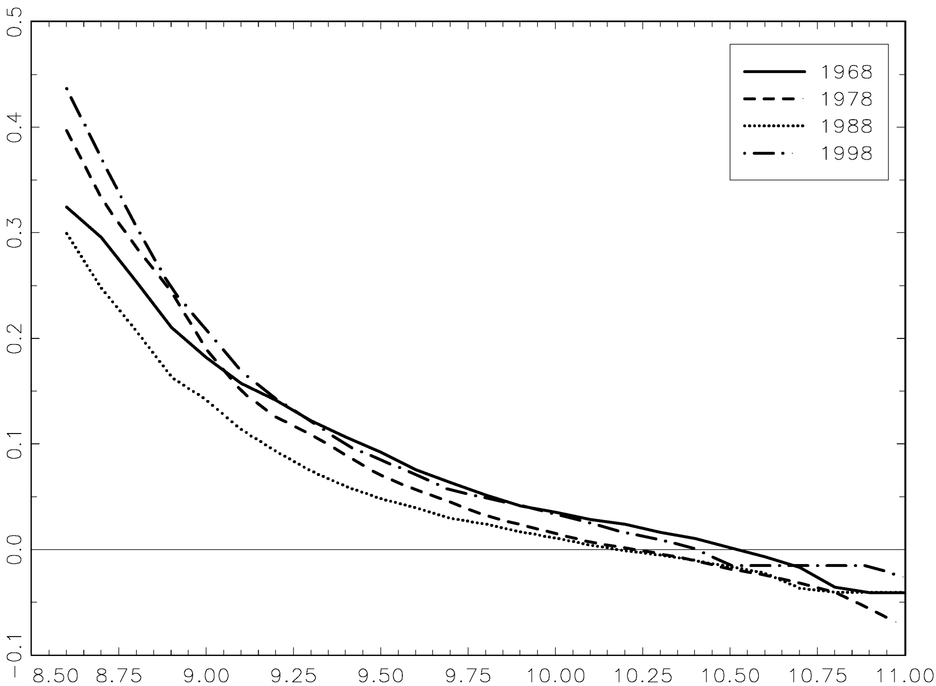

Figure 7. Conditional Medians of Δy_{it} Given $y_{i,t-1}$ Plotted Against $y_{i,t-1}$

Note however that although the amount of reversion was substantial, earnings in adjoining years were positively correlated. It is clear that a plot of earnings (rather than changes) against the previous year's earnings follows the 45-degree line, for those with midrange initial earnings. If earnings across years are highly correlated, then longer-term inequality will not be much lower than yearly inequality.

The spread of earnings changes varied considerably depending on initial earnings. Relative to earners around the median, low and high earners were far more mobile. For example, in 1984 the 90–10 spread of changes for those with initial earnings below the 10th percentile was more than twice as large as for those with median initial earnings. Low earners experienced both strong upward mobility and significant downward mobility: over 25 percent of low earners experienced negative changes in earnings and 10 percent of low earners experienced a fall of more than 70 percent. (Gottschalk and Moffitt (1994) also found that low earners had the highest variance in earnings.) Note that the downward mobility of low earners works against reversion.

While the spread was symmetric around the median for low earners, it was highly asymmetric for high earners. Measured relative to the median change, negative changes for high earners tended to be much larger than positive changes. Of the earners whose initial log-earnings were near 10.6 in 1983 (which corresponds to the 90th percentile), 10 percent saw increases of more than 20 percent in their earnings, while 10 percent experienced decreases of more than 55 percent. The median change was between 0 and –5 percent.

The approximate pointwise confidence bands are wider for the 10th percentile than for other percentiles, and wider for low and high earners than for middle earners. The estimated curve is not always at the center of the confidence band; this reflects the smaller bias of the confidence bands (which comes at the cost of higher variance). The confidence bands suggest that there is nonnegligible uncertainty about the estimated curves, but it is also apparent that the general conclusions are robust.

The general patterns identified in 1984 prove to be quite consistent over time, as mentioned above. Figures 7, 8 and 9 graph the medians and quantile spreads of changes conditional on initial earnings, at ten-year intervals, and their shape is very similar to 1984. (Figures 11, 12 and 13 plot variations from year to year in more detail, showing that 1984 was not an unusual year.) The 50–10 spread was U-shaped with a minimum around 10, which was slightly above median initial earnings in most years. The spread of changes for the lowest earners and highest earners was twice to three times as large as the minimum value, and larger for the lowest earners than the highest earners. In contrast the 90–50 spread was generally downward-sloping relative to earnings, confirming the asymmetry in changes for high earners due to their high downward mobility.

How does reversion in the churning process translate into offsetting inequality in the longer run? Offsetting effects are most clearly understood in a

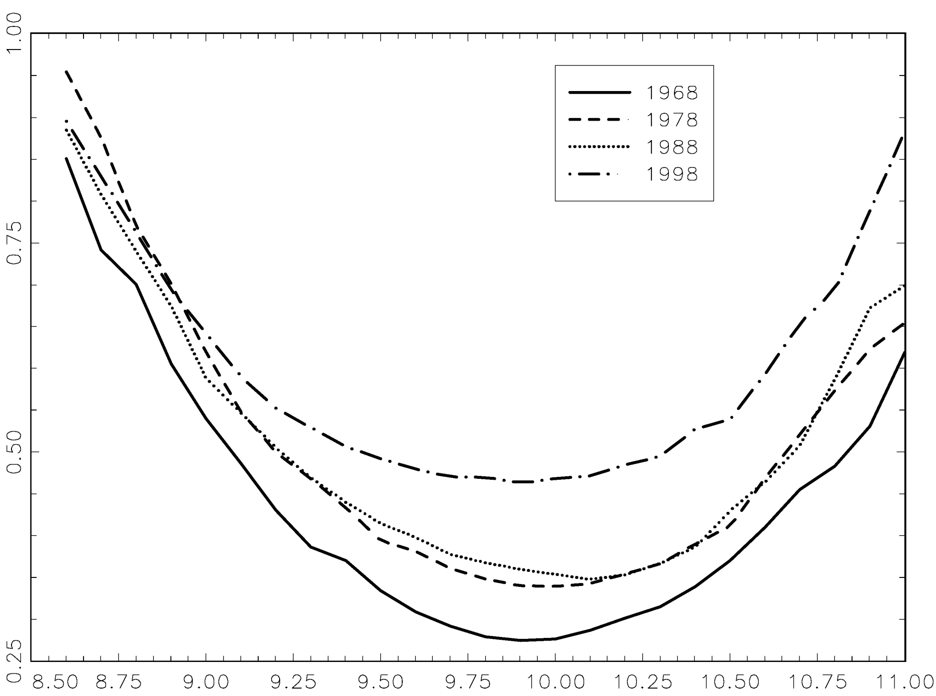

Figure 8. Conditional 50–10 Percentile Spreads of Δy_{it} Given $y_{i,t-1}$ Plotted Against $y_{i,t-1}$

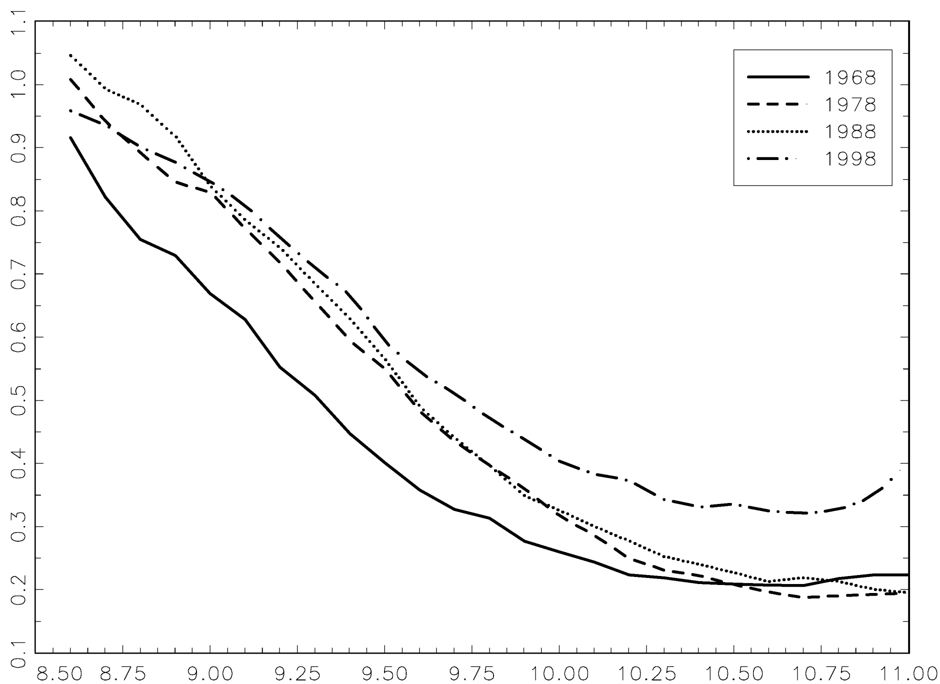

Figure 9. Conditional 90–50 Percentile Spreads of Δy_{it} Given $y_{i,t-1}$ Plotted Against $y_{i,t-1}$

context where the distributions of yearly earnings and yearly mobility are constant over time. The spread (and inequality) of earnings over longer periods will necessarily be smaller than over one year, but the structure of mobility determines how much smaller. Holding the yearly earnings distribution constant, observing a higher degree of upward mobility of high earners must imply less upward mobility for a lower earnings group, and hence more unequal earnings over long periods. Observing more downward mobility for the highest earners implies more upward mobility or less downward mobility for lower earners. Similarly, more upward (downward) mobility of low earners implies less (more) unequal earnings over long periods. The degree of reversion is therefore indicative of the extent of offsetting.

Throughout this period the upward mobility of low earners acted as an offtetting effect that was partially undone by their downward mobility. Contrarily the offsetting effect of high earners' downward mobility was not counteracted by significant upward mobility. Consequently we should see more offsetting at the upper end of the distribution than at the lower end. This is in fact what we observe in Figure 10, which plots the distribution of two-year average log-earnings (corresponding to the geometric average of earnings) as well as yearly log-earnings for the "matched++" sample. Notice that much of the effect of reversion is not visible when inequality is measured over a mere two years: inequality over two years was less than yearly inequality, but not much less. The two curves were only noticeably different around the 90th percentile, as the 90th percentile over

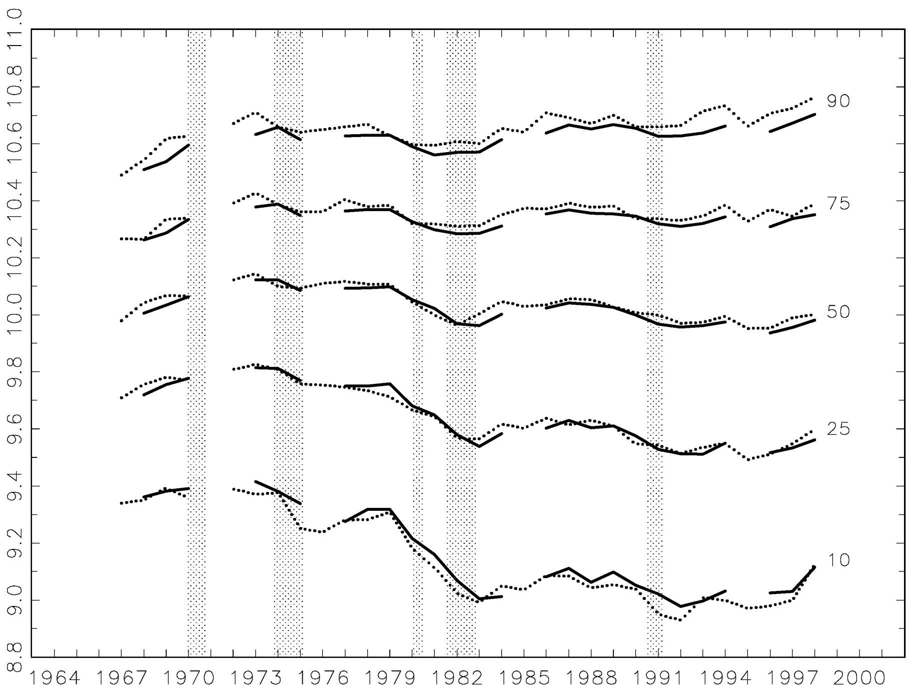

Figure 10. Percentiles of $0.5y_{i,t-1} + 0.5y_{it}$ (Solid) and of y_{it} for the "Matched++" Sample (Dotted) Plotted Against t

two years was well below the yearly 90th. At least part of this difference is due to greater offsetting at the upper end of the distribution.

4.4. *Changes in Mobility Offsetting Changes in Inequality*

Here we concentrate on changes in mobility over the three decades, in order to assess whether increasing inequality over the decades has been offset by increases in reversion. There is a complication: given that mobility is defined as earnings changes, changes in the earnings distribution must be reflected in mobility. Thus increases in inequality will be reflected in changes in mobility that "aggravate inequality" by reducing reversion. Fortunately, disentangling those changes from offsetting changes is relatively straightforward in these data.

The evolution of the distribution over time is studied in more detail by plotting the conditional median and spreads of changes for selected levels of initial earnings against time. (Refer to Figures 7, 8 and 9 for the general pattern of these spreads by initial earnings.) Figure 11 plots the conditional median for individuals with initial log-earnings of 9.0, 9.5, 10.0, and 10.5; Figures 12 and 13 plot the conditional 50–10 and 90–50 percentile spreads, respectively.

Beginning with the cyclical variation in the figures, note that median changes were roughly cyclical, relative to recessions as indicated by shaded areas in the figure. This is particularly true for the lowest earners considered. The variation

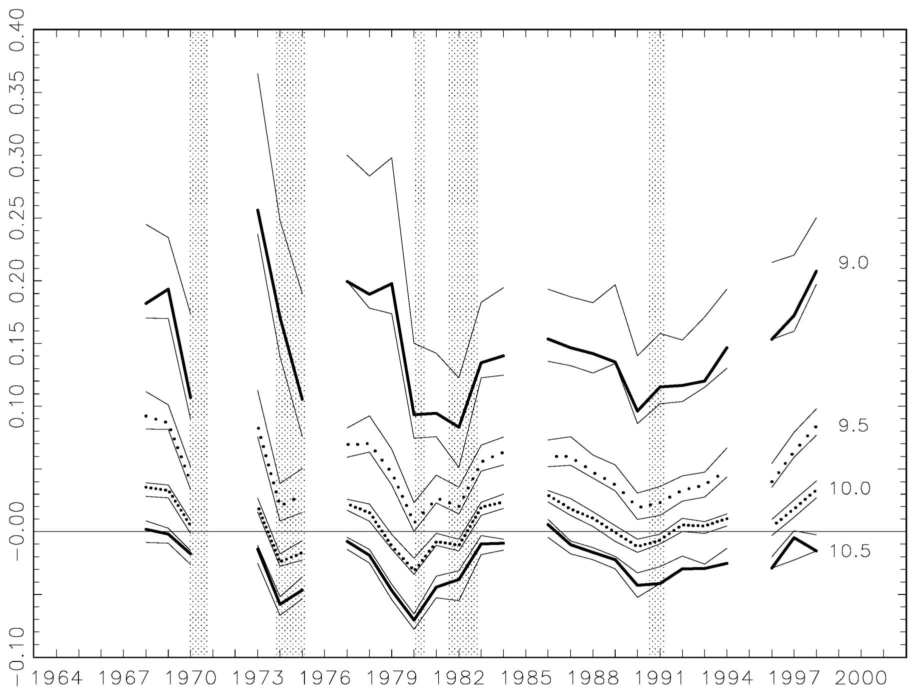

Figure 11. Conditional Medians of Δy_{it} Given $y_{i,t-1} = *$ Plotted Against t

in the 50–10 spread was also strong, and coincided with recessions, while there was much less variation in the 90–50 spreads.

Given that we are attempting to relate changes in mobility to changes in inequality, it seems natural to examine mobility in the two "phases" identified in the evolution of inequality, 1967 to 1982, and 1982 onwards. We focus our attention on the 50–10 and 90–50 spreads, as median changes showed no strong trend over the period, beyond a slight U-shape with its low around 1982.

Recall that rising inequality in the 1967–82 period took the form of lower earnings for the bottom of the distribution; this is *reflected* in a large increase of the 50–10 spread for low earners, implying greater downward mobility. However, their worsening situation was offset to some extent by an increase in their upward mobility (an increase in the 90–50 spread for low earners) and by more downward mobility for higher earners. We conclude that offsetting mobility increased from 1967 to 1982, at the same time as increasing inequality was reflected in aggravating changes in mobility.

From 1982 onwards, inequality rose more slowly, and took the form of increased earnings for high earners. That increase is reflected in the slow increase in the 90–50 spread for high earners, implying greater upward mobility for high earners. Downward mobility for low earners dropped to a lower level after 1986, implying a slight offsetting effect, but otherwise there are no strong changes in mobility. We conclude that there was little increase in offsetting during the later period, as opposed to the 1967–82 period.

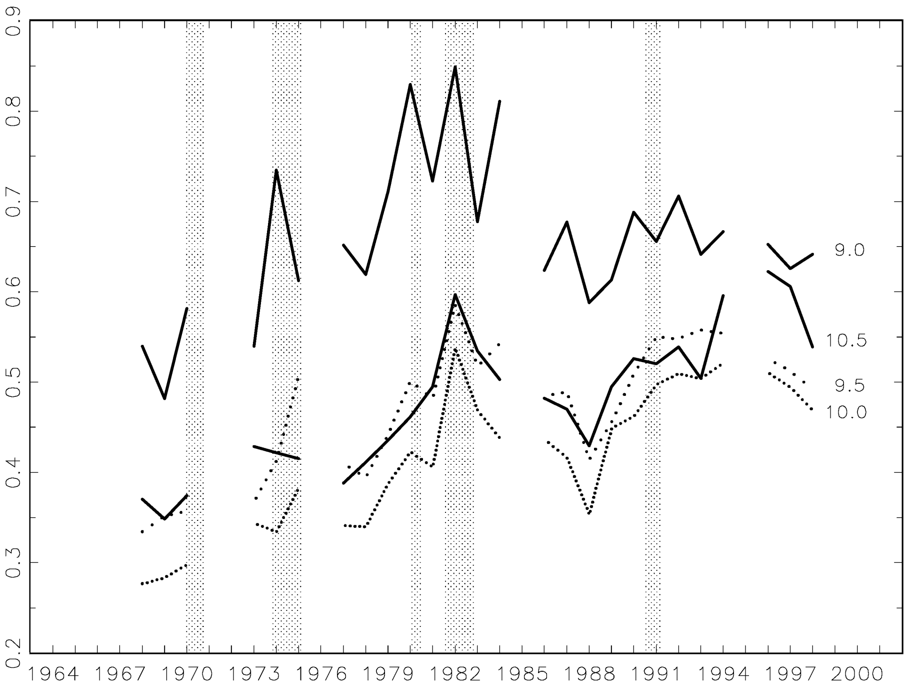

Figure 12. Conditional 50–10 Percentile Spreads of Δy_{it} Given $y_{i,t-1} = *$ Plotted Against t

These findings are roughly similar to those of Gittleman and Joyce (1996), who regressed the correlation between earnings across two years against a quadratic time trend. They found a U-shaped pattern, declining in the 1970s, increasing in the 1980s, with a turning point around 1982; but their regression method did not identify the more stable structure of mobility in later years. Likewise Gottschalk and Moffitt, 1994 found an increase in the "variability" of earnings when comparing the 1970–78 period to 1979–87. In contrast, the small literature on transitions between quintiles (or deciles) found that the number of transitions declined over the 1980s (Arkes, 1998; Buchinsky and Hunt, 1999). This approach measures changes in relative position, rather than changes in earnings, and therefore identifies slightly different patterns. For instance, Arkes (1998) found no effect of the business cycle on mobility, while we find a strong increase in downward mobility for low earners; but this is unlikely to impact relative positions. Moreover, relative position is measured through fairly broad categories: therefore as inequality increases and the distribution widens, there will be less movement between quintiles.

5. Steady-state Distributions

In the discussion above, we identified two aspects of inequality's relationship to mobility: first, that mobility can offset inequality, and second, that mobility reflects changes in inequality (that is, a change in inequality must necessarily come

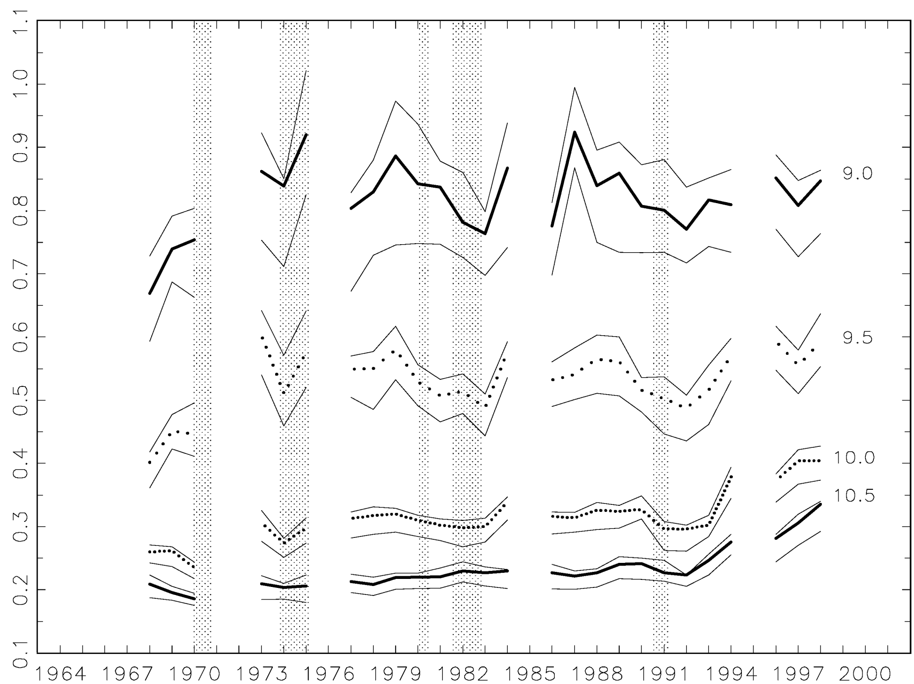

Figure 13. Conditional 90–50 Percentile Spreads of Δy_{it} Given $y_{i,t-1} = *$ Plotted Against t, with 95% Confidence Bands

about through changes in earnings). This section attempts to isolate the latter interdependence more clearly. By distilling the implications for inequality, through a measure we call the "steady state" distribution, we can more clearly identify when the change in mobility from one year to another has worsened or improved inequality.

Recall from equation (1) that changes in g_t over time come about partly through changes in f_t and partly through the convolution process itself. That is, even if f_t does not change over time g_t may not be in a steady-state equilibrium. To isolate the effect of the churning process in a given year from the distribution of earnings in the previous year, we compute and compare the steady state distributions corresponding to f_t for each year in the sample. The steady-state is the distribution of earnings that would hold if mobility f_t were forever at its current level. That is, given our estimated density $\hat{f}_t$ we calculate the density $\hat{g}_t^*$ such that

$$\hat{g}_t^*(y) = \int_{-\infty}^{\infty} \hat{f}_t(y - v|v)\hat{g}_t^*(v)\, dv. \tag{8}$$

The subscript t indicates that $\hat{g}_t^*$ is computed using the churning process in year t. The estimated steady-state distribution g_t^* serves to isolate the effect of churning in a given year. In effect, the degree of churning has been transformed into an earnings distribution. As such, it serves to measure the yearly inequality implied

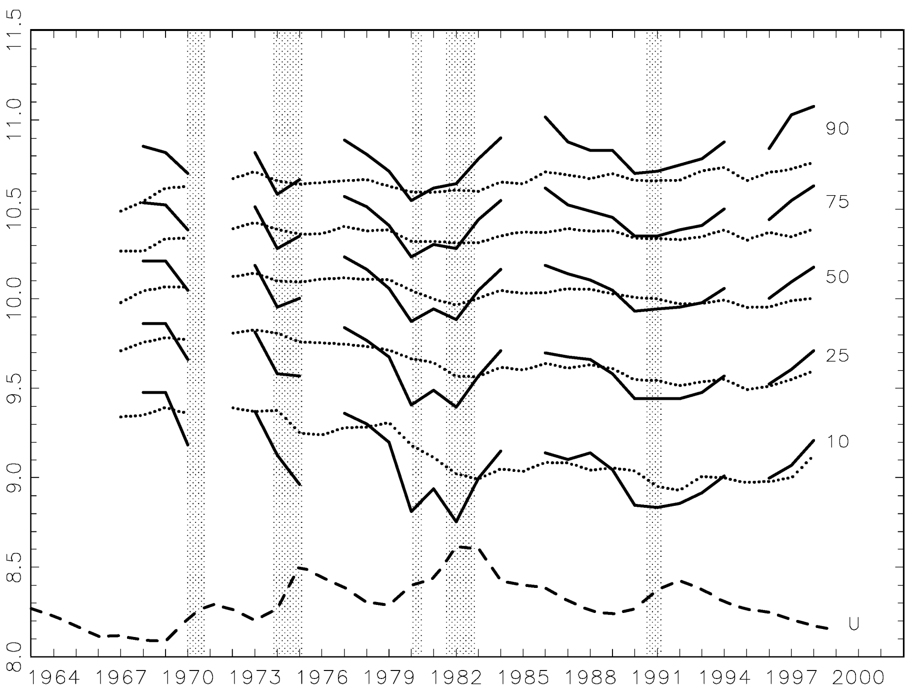

Figure 14. Percentiles of y_{it} for the "Steady State" Distribution (Solid) and the "Matched++" Sample (Dotted) Plotted Against t with Unemployment Rate (Dashed)

by the distribution of earnings changes. The effect of previous years has been purged out.

Figure 14 graphs the quantiles of the estimated steady-state distributions over time, as well as the quantiles of the actual distributions for the "matched++" subsample. The annual average unemployment rate is also graphed below the curves (although not to scale), highlighting the cyclical variation in employment conditions caused by the business cycle. Note that changes in inequality are much more pronounced in the steady-state distributions, because the steadying influence of past effects has been eliminated. As a result the implied effects of changes in mobility are much easier to identify. The effect of the business cycle on mobility is much clearer, as the quantiles move down with recessions and up with the recovery period. It also appears that economic upturns tend to reduce inequality, in that the upturn earnings for the lowest quantiles rise further than other quantiles during the booms. Inequality also seems to be increasing in the downturns, although not identically in every year of the recession (as the precision of the measure allows us to determine). The correlation between the unemployment rate and the 90th percentile is –0.36, whereas it is much stronger for the 10th percentile, –0.64. Similarly, the correlation between the unemployment rate and the 90–50 spread of earnings is only –0.06, whereas it is 0.66 for the 50–10 spread.

The changes in inequality over the business cycle can be related to our analysis of cyclical variation. Recall that while individuals at all quantiles experienced year-to-year variations in earnings, particularly over the business cycle, the effects were strongest for low earners. Both the median and the 90–50 spread of changes were much more cyclical for low earners. Thus the source of the worsened situation of low earners during recessions was worsened mobility for those who were low earners in the previous year. This structure of mobility has more serious welfare implications than if recessions were marked by more downward mobility for another group, such as median earners, because offsetting is less. In that sense, mobility in the cyclical changes is "aggravating" rather than "offsetting."

6. Conclusion

This paper has introduced several powerful techniques for the analysis of mobility. Nonparametric quantile estimation allowed us to examine the full distribution of conditional changes in earnings given initial earnings. Based on the conditional distribution we identified differences in upward and downward mobility for different earnings groups, both their spread and their variation over time. The analysis pinpointed which changes in mobility were acting to offset inequality. Our estimates of the conditional distribution were used to derive steady-state distributions, through a convolution technique, and thereby confirm that greater variation in mobility for low earners was driving the cyclical pattern in inequality.

References

Arkes, J., "Trends in Long-run versus Cross-section Earnings Inequality in the 1970s and 1980s," *Review of Income and Wealth*, 44, 199–213, 1998.

Baker, M. and G. Solon, "Earnings Dynamics and Inequality among Canadian Men, 1976–1992: Evidence from Longitudinal Income Tax Records," *Journal of Labor Economics*, forthcoming.

Bishop, J. A., J. P. Formby, and P. D. Thistle, "Changing American Earnings Distributions: One-half Century of Experience," *Empirical Economics*, 22, 501–14, 1997.

Brauer, D. A. and S. Hickok, "Explaining the Growing Inequality in Wages Across Skill Levels," *Federal Reserve Bank of New York Economic Policy Review*, 1, 61–75, 1995.

Buchinsky, M. and J. Hunt, "Wage Mobility in the United States," *The Review of Economics and Statistics*, 81, 351–68, 1999.

Burkhauser, R. V., A. Crews Cutts, M. C. Daly, and S. P. Jekins, "Testing the Significance of Income Distribution Changes over the 1980s Business Cycle: A Cross-national Comparison," *Journal of Applied Econometrics*, 14, 253–72, 1999.

Fan, J., "Design-adaptive Nonparametric Regression," *Journal of the American Statistical Association*, 87, 998–1004, 1992.

Gittleman, M. and M. Joyce, "Earnings Mobility and Long-run Inequality: An Analysis using Matched CPS data," *Industrial Relations*, 35, 180–96, 1996.

Gottschalk, P. and R. Moffitt, "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity*, 2, 217–54, 1994.

Gottschalk, P. and T. Smeeding, "Cross-national Comparisons of Earnings and Income Inequality," *Journal of Economic Literature*, XXXV, 633–87, 1997.

Johnson, G. E., "Changes in Earnings Inequality: The Role of Demand Shifts," *Journal of Economic Perspectives*, 11, 41–54, 1997.

Juhn, C., K. M. Murphy, and B. Pierce, "Wage Inequality and the Rise in Returns to Skill," *Journal of Political Economy*, 101, 411–41, 1993.

Levy, F. and R. J. Murnane, "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations," *Journal of Economic Literature*, 30, 1333–81, 1992.

Peracchi, F,. and F. Welch, "How Representative are Matched Cross-sections? Evidence from the Current Population Survey," *Journal of Econometrics*, 68, 153–79, 1995.

Ruppert, D., "Empirical-bias Bandwidths for Local Polynomial Nonparametric Regression and Density Estimation," *Journal of the American Statistical Association*, 92, 1049–62, 1997.

Ruppert, D. and M. P. Wand, "Multivariate Locally Weighted Least Squares Regression," *Annals of Statistics*, 22, 1346–70, 1994.
Wand, M. P. and M. C. Jones, *Kernel Smoothing*, Chapman and Hall, London, 1995.
Yu, K. and M. C. Jones, "A Comparison of Local Constant and Local Linear Regression Quantile Estimation," *Computational Statistics and Data Analysis*, 25, 159–66, 1997.
———, "Local Linear Quantile Regression," *Journal of the American Stastical Association*, 93, 228–37, 1998.