A Distributional Approach to Measuring Changes in Economic Inequality

Mark S. Handcock
Department of Statistics & Operations Research
New York University
New York, NY 10012

Martina Morris
Department of Sociology
Pennsylvania State University
University Park, PA 16802

Annette Bernhardt
Institute on Education and the Economy
Columbia University
New York, NY 10027

Social scientists and economists are increasingly interested in techniques for comparing changes in distributional shape as well as changes in mean-levels. We present a statistical approach to interdistributional comparison that relies on the relative distribution: a graphical and statistical tool for representing location and shape differences, applicable to both continuous and discrete data. We also develop two application-oriented measures based on the relative distribution. The first is a set of relative polarization indices that identify growing density in the tails of the distribution. The second is a decomposition technique for isolating the marginal impact of changes in population composition.

We illustrate the flexibility and utility of this approach in an economic application, examining the widely publicized recent rise in wage inequality. We use the relative distribution of men's hourly wages to graphically display and summarize the growing wage polarization between 1975 and 1993. During the same years we find a similar polarizing trend in the weekly hours worked, which suggests that the polarization in wages may be due in part to the change in the mix of part-time and full-time workers. Decomposing the relative distribution, however, we find that this compositional change makes only a modest contribution to the rise in wage inequality.

KEY WORDS: Comparison change analysis; Grade distribution; Interdistributional dissimilarity; Relative distribution; Polarization.

Mark S. Handcock is Associate Professor of Statistics, Department of Statistics & Operations Research, Leonard N. Stern School of Business, New York University, New York, NY 10012. Martina Morris is Associate Professor of Sociology, Department of Sociology, The Pennsylvania State University, 614 Oswald Tower, University Park, PA 16802. Annette D. Bernhardt is Senior Research Associate, Institute on Education and the Economy, Teachers College, Columbia University, New York, NY 10027. The authors thank Jeffrey S. Simonoff for numerous useful comments that have greatly improved this paper.

1. INTRODUCTION

In an increasing number of social science applications, the comparison of an attribute across groups requires consideration of more than the usual summary measures of location and variation. A good example can be found in recent analyses of earnings inequality (Butler and McDonald 1987, Grove and Hannum 1986, Picot, Myles and Wannel 1990). Traditionally, research on inequality has relied heavily on measures that capture differences in average earnings between groups or rough measures of dispersion over time. These summary measures leave untapped much of the information inherent in a distribution. With the recent growth of wage inequality within groups and simultaneous changing patterns of inequality between groups (Bernhardt, Morris and Handcock 1995, Blau and Kahn 1994) has come a renewed interest in statistical approaches that can capture such changes (DiNardo, Fortin and Lemieux 1995, Ginther 1995, Hildenbrand and Hildenbrand 1986, Marron and Schmitz 1992)

In this spirit, we develop the relative distribution as a statistical tool for fully representing differences between distributions. The relative distribution provides a general integrated framework for analysis: a graphical component that simplifies exploratory data analysis and display, a statistically valid basis for the development of hypothesis-driven summary measures, and the potential for decomposition that enables one to examine hypotheses regarding the origins of distributional changes within and between groups. In this paper we demonstrate the use of the relative distribution for each of these analytic tasks. The integration of the different components of analysis, the capability of representing full distributional information, and the flexibility for representing the factors that are of substantive interest, make this approach well suited to emerging research questions in many fields.

We use these methods to analyze the magnitude and origins of the recent growth in earnings inequality. Changes in the distribution of white men's hourly wages between 1975 and 1993 can be expressed as a relative wage distribution. The growing polarization in earnings is clearly visible in the graph of the relative density, and the set of polarization indices quantify the impacts of wage–upgrading and –downgrading during this period. Many factors have been suggested as potential explanations for this rising inequality. Initial hypotheses regarding the impact of changing demographics (the baby boom and the greater labor force participation of women) and business cycles suggested that the trend might be temporary. The stability of the trends, however, has increasingly led researchers to focus on more permanent macro–economic changes in trade and technology and their impacts on the restructuring of work. A central hypothesis has centered on the increasing use of part-time and contingent workers. As businesses substitute lower–paid temporary workers for permanent staff, one might expect the aggregate wage distribution to display a growing polarization. We examine this hypothesis here, identifying the contribution of changes in the distribution of hours worked to the growth in wage inequality.

Section 2 provides the majority of the technical development and a description of the data used in the application to wage inequality. We define the relative distribution, identify its history and connections to previous methods, and discuss inference issues. We also introduce a set of summary measures, the median, upper and lower polarization indices, designed to capture the changes in wages and hours worked that are of particular substantive interest. In Section 3 we use these methods to examine the univariate changes in the distributions of hourly wages and weekly hours worked from 1975 to 1993. In Section 4 we pursue the connection between the observed polarization in both wages and hours, stratifying the change in the relative wage distribution by part–time, full–time, and over–time workers. We conclude by presenting a technique for bivariate decomposition, and apply this to the relative distribution of wages to isolate the effect of changes in the distribution of hours worked.

2. The relative distribution

Let Y_0 be a random variable representing a baseline measurement for a population (e.g., hourly wages). It is assumed that Y_0 is nonnegative with cumulative distribution function (CDF) $F_0(y)$ and density $f_0(y)$ (when the latter is defined). The restriction of the support of the random variables is motivated by the nature of the applications and can be removed with only a minor change in notation.

Suppose we observe another measurement Y. It is assumed that Y is nonnegative with distribution function $F(y)$ and density $f(y)$ (when the latter is defined). Typically Y is the measurement for a separate group or the same group during a later time period. The objective is to measure the relative prevalence of Y relative to the baseline distribution Y_0.

For continuous distributions, we define the *relative distribution* of Y to Y_0 to be the random variable:

$$X = F_0(Y) \tag{2.1}$$

The cumulative distribution function (CDF) of X is

$$G(p) \;=\; F(F_0^{-1}(p)) \qquad 0 \le p \le 1. \tag{2.2}$$

where $F_0^{-1}(p) = \inf_y\{y \mid F_0(y) \;\ge\; p\;\}$ is the left continuous inverse CDF of Y_0. The corresponding density, when it exists, is

$$g(p) = \frac{f(F_0^{-1}(p))}{f_0(F_0^{-1}(p))} \qquad 0 \le p \le 1$$

where p represents the proportion of values.

The relative CDF $G(p)$ has a particularly interesting interpretation: a proportion $G(p)$ of the Y are below the values of a proportion p of Y_0. Note that the *dollar* amount is constant in both cases ($\$x \;=\; F_0^{-1}(p) \;=\; F^{-1}(G(p))$ representing the cut–point). The relative density $g(p)$ has the direct interpretation as the ratio of the target population density to the baseline population density at a given level ($F_0^{-1}(p)$). In general, the relative distribution is invariant to the scale of the distributions (up to a monotone transformation). For example, the same relative distribution would result from comparison of log–hourly wages or hourly wages itself.

If the two distributions are identical then the relative distribution is uniform in the following sense: If F is continuous then X has the uniform distribution on $[0, 1]$ so that the relative density is constant on $[0, 1]$. For F not necessarily continuous, $G(p) \le p \quad 0 \le p \le 1$, with equality failing if and only if p is not in the closure of the range of F (Shorack and Wellner 1986).

2.1 History and Background

The relative CDF $G(r)$ is implicitly a theoretical P–P plot of F against F_0, an empirical version of which was considered by Wilk & Gnanadesikan (1968). It is the plot $\{F(r), F_0(r) : r \ge 0\}$. See also Chambers, Cleveland, Kleiner & Tukey (1983). In addition the $G(r)$ is essentially a receiver operating characteristic (ROC) curve used in the evaluation of the performance of medical tests for separating two populations (Cambell 1994).

A closely related quantity to the relative density $g(r)$ is that of the *density ratio*: $h(x) \;=\; f(x)/f_0(x),\;\; x \in \mathbb{R}$ considered by Silverman (1978). It is a key element of discriminant analysis (Hand 1982) and likelihood–ratio methods. Note that $h(x) = g(F_0(x))$ and $g(r) = h(F_0^{-1}(r))$. Cwik and Mielniczuk (1989) developed a kernel estimate for $h(x)$ and show uniform a.s. convergence. Gijbels and Mielniczuk (1995) generalized these results (to the Radon–Nikodym derivative) and determined the rates of uniform a.s. convergence.

The relative density itself has been explicitly studied in at least two contexts. Parzen (1992) studies the relative distribution as part of "comparison change analysis," and refers to it as the "comparison density". Separately, Cwik and Mielniczuk (1993) consider kernel density estimation for the relative density. They refer to it as the "grade density". They also consider band–width selection for their estimator. The work presented here has developed independently from these two contexts.

2.2 Estimating the Relative Distribution

In this section we describe the methodology used in the discussion of the subsequent sections. Let $Y_1, Y_2, \ldots, Y_m$ be independently and identically distributed from the distribution F. Similarly, let $Y_{01}, Y_{02}, \ldots, Y_{0n}$ be independently and identically distributed from the distribution F_0. We consider the situation where the $\{Y_j\}_{j=1}^m$ are independent of $\{Y_{0i}\}_{i=1}^n$. This is the usual situation when the sample is a result of a cross–sectional sample survey. If the sample is a result of a longitudinal survey the samples will tend to be correlated. While the development is similar, this case will be considered elsewhere. Let $F_m(y) = \frac{1}{m}\sum_{j=1}^m \mathcal{I}(Y_j \leq y)$ be the empirical distribution function of Y and $F_{n0}(y) = \frac{1}{n}\sum_{i=1}^n \mathcal{I}(Y_{0i} \leq y)$ be the empirical distribution function of Y_0. Here $\mathcal{I}(\cdot)$ is the indicator function.

The natural estimator of the relative CDF $G(p) = F(F^{-1}(p))$ in (2.2) is:

$$G_{n,m}(p) = F_m(F_{n0}^{-1}(p)) \qquad 0 < p < 1$$

This estimator can be viewed as a U–statistic with an estimated parameter ($\lambda = F_0^{-1}(p)$ is estimated by $\hat{\lambda} = F_{n0}^{-1}(p)$). The relative density is estimated by using a local–polynomial smoother (e.g., Hastie & Loader 1993) to a binned version of the relative data (see Section 2.6). The confidence intervals can then be derived from within the framework of the local–polynomial smoother. These results can be used to calculate (simultaneous) confidence bands for $G(p)$ and $g(p)$ based on $G_{n,m}(p)$. This methodology is developed in Handcock and Janssen (1995a).

2.3 Measuring Polarization

In this section we develop a measure of interdistributional inequality that isolates shape changes from level changes, and summarizes the extent of polarization in the relative distribution. Polarization is of particular interest in the study of wage inequality, because it captures a discrepancy in living standards that is hidden when only trends in average wages are examined. While visual examination of the (relative) distributions shows whether polarization is occurring, we would also like to be able to quantify the observed patterns. The polarization index defined here and its decomposition provide a flexible and sensitive method for measuring the comparative density in the center or tails of the distribution. It plays the same role as the difference in Gini coefficients, coefficients of variation, or variances of log–values in measuring interdistributional inequality (c.f., Grove and Hannum 1986). The polarization index however is based on the relative distribution, provides a more intuitive link between graphical display and numerical summary, and can be decomposed to compare the growth in the upper and lower tails.

Isolating differences in distributional shape requires that differences in location be removed through deflation. While it is common to inflate prior year earnings by a Consumer Price Index–based measure, the resulting time series may still display location shifts. As we are working in a quantile framework, we inflate by median–matching the distributions. That is, we inflate the reference values by the ratio of the median of the target distribution to the median of the reference distribution. (that is, $Y_0 \times \text{median}(Y)/\text{median}(Y_0)$). The measures and the relative density would be unchanged if we chose instead to deflate to target values to the reference values. Since the typical difference between these distributions, as measured by the median difference, is zero, measures based on the median–matched relative distribution isolate other aspects of interdistributional inequality (i.e., shape changes). For example, if one distribution is more polarized than another we would expect a U–shaped median–matched relative distribution. If the distributions differ only in their level, then the median–matched relative distribution would be approximately uniform.

2.4 The Median Relative Polarization Index

Ideally, we would like a statistic that measures the deviations of the relative distribution from the uniform distribution (i.e., no divergence) and that emphasizes the deviations in both the upper and lower tails. We consider the following measure:

Let ρY_0 denote the random variable Y_0 multiplied by a factor ρ, so that the CDF of ρY_0 is $F_0(y/\rho)$. In particular to median–match Y_0 with Y, we will consider $\rho = \text{median(F)}/\text{median}(F_0)$. Define the *median relative polarization index* (MRP) of Y relative to Y_0 as:

$$RP(F;\ F_0) \quad = \quad 4\int_0^1 \mid x - \frac{1}{2} \mid g_m(x)\ dx \quad - \quad 1,$$

where g_m is the relative density of Y to ρY_0. This measure weighs the mass in the upper and lower tails more heavily than the mass in the center. It is a scaled version of mean absolute deviation about the median of the median–matched relative distribution. The scaling produces an index that varies between -1 and 1, positive values representing more polarization, i.e., increases in the tails of the distribution, and negative values representing less polarization, i.e., convergence towards the center of the distribution. If the median–matched relative distribution is uniform the index will be zero. The MRP has a number of useful characteristics. For example, it has a natural symmetry property: the MRP of Y_0 to Y is equal to the negative of the MRP of Y to Y_0. It is also invariant to monotone transformations of the distributions: Let $h(\cdot)$ be a monotone function on the support of Y_0. Then the MRP of $h(Y)$ to $h(Y_0)$ is equal to the MRP of Y to Y_0.

2.5 Upper and Lower Relative Polarization Indices

We would like also to compare the contributions made by the upper and lower tails to the MRP. For example, is wage upgrading more prevalent than wage downgrading? The MRP is decomposable into the contributions made by components above and below the median of the relative distribution. We define the lower polarization index (LRP) by:

$$LRP(F;\ F_0) \quad = \quad 8\int_0^{\frac{1}{2}} \mid x - \frac{1}{2} \mid g_m(x)\ dx \quad - \quad 1,$$

The upper polarization index (URP) is defined in a similar way:

$$URP(F;\ F_0) \quad = \quad 8\int_{\frac{1}{2}}^1 \mid x - \frac{1}{2} \mid g_m(x)\ dx \quad - \quad 1$$

These indices decompose the MRP in the sense that:

$$RP(F;\ F_0) \quad = \quad \frac{1}{2}LRP(F;\ F_0) \quad + \quad \frac{1}{2}URP(F;\ F_0)$$

The upper and lower indices have properties similar to the MRP. The scaling produces upper and lower indices that vary between -1 and 1, positive values representing more polarization, i.e., increases in the tail of the distribution, and negative values representing less polarization, i.e., convergence towards the center of the distribution. If the median–matched relative distribution is uniform below (above) the median then the lower (upper) index will be zero.

2.6 Inference for the Relative Polarization Indices

As in Section 2.2, we will investigate the properties of the natural estimator of $RP(F;\ F_0)$

$$\widehat{RP(F;\ F_0)} \equiv RP(F_m;\ F_{n0})$$

Define the *relative data* $\{X_1, X_2, \ldots, X_m\}$ by:

$$X_j = F_{n0}(Y_j) \qquad j = 1, \ldots, m$$

then the estimate has the following appealing expression:

$$\widehat{RP(F;\ F_0)} \quad = \quad \frac{4}{m}\sum_{j=1}^{m}\left|X_j - \frac{1}{2}\right| \quad - \quad 1$$

This expression follows from the definitions of the empirical distribution function and some algebra. Note that the $\{X_j\}_{j=1}^{m}$ are not independent as they depend on the $\{Y_{0i}\}_{i=1}^{n}$. However they will be close to uncorrelated (their pairwise correlation is $O(n^{-1})$).

The natural estimators of LRP and URP are $\widehat{LRP(F; F_0)} \equiv LRP(F_m;\ F_{n0})$ and $\widehat{URP(F;\ F_0)} \equiv URP(F_m;\ F_{n0})$, respectively. They can be also be easily expressed in terms of the quasi–data:

$$\widehat{LRP(F;\ F_0)} \quad = \quad \frac{8}{m}\sum_{j=1}^{m}\left|X_j - \frac{1}{2}\right|\mathcal{I}(X_j \leq \frac{1}{2}) \quad - \quad 1$$

$$\widehat{URP(F; F_0)} = \frac{8}{m}\sum_{j=1}^{m}\left|X_j - \frac{1}{2}\right|\mathcal{I}(X_j > \frac{1}{2}) - 1$$

Inference for the polarization indices for continuous data is developed in Handcock and Janssen (1995b). The statistical properties of these estimators for grouped data are summarized in the Appendix. These results are used to calculate (simultaneous) confidence bands for $RP(F;\ F_0), LRP(F;\ F_0)$ and $LRP(F;\ F_0)$. Often we would like to test if the median, upper or lower relative polarization indices in a given situation is statistically significantly different from zero. These tests can be constructed directly from the results given in the Appendix. If the sample size is not small (i.e., $n, m > 30$), we can use the Gaussian approximation to the exact distributions of the estimate as the basis for a test. If the sample sizes are small the bootstrap can be used to determine the sampling distribution of the estimate and the corresponding critical values.

2.7 Data

The data are drawn from the U.S. Current Population Survey (CPS) in its annual March Supplement (1976 and 1994). The selected sample consists of white males, aged 16–66 and excludes the self–employed, full–time students, and those in the military and in farming. Estimates of the hourly wage are obtained from two questions. In each survey, respondents were asked, “In the weeks that ... worked, how many hours did ... usually work per week?” and “During 19XX in how many weeks did ... work even for a few hours? Include paid vacation and sick leave as work.” They were also asked asked how much income they received in wage and salary before deductions during 19XX. The reported earnings were top–coded by \$50,000 in 1975 and \$99,999 in 1993. A small proportion of the reported annual earnings fall in the top–coded category (0.5%) . We derive hours worked last year by multiplying the reported hours worked per week last year by the reported weeks worked last year. The hourly wage for the given year is then derived by dividing the total wage and salary income by the total number of hours worked. In order to measure the wages from the two surveys on a common dollar scale, the 1975 hourly wages have been inflated to 1993 hourly wages using the median ratio as a relative earnings inflator (that is, $Y_0 \times \text{median}(Y)/\text{median}(Y_0)$.)

Hourly wages from the March CPS have been used by researchers (Juhn, Murphy and Pierce 1993, Murphy and Welch 1992) to study changes in the structure of wages. Other reasonable measures include weekly wages, yearly earnings, and direct measurement of hourly wages via the May CPS outgoing rotation groups. The different questions addresses by these measures has been discussed in DiNardo, et al. 1995.

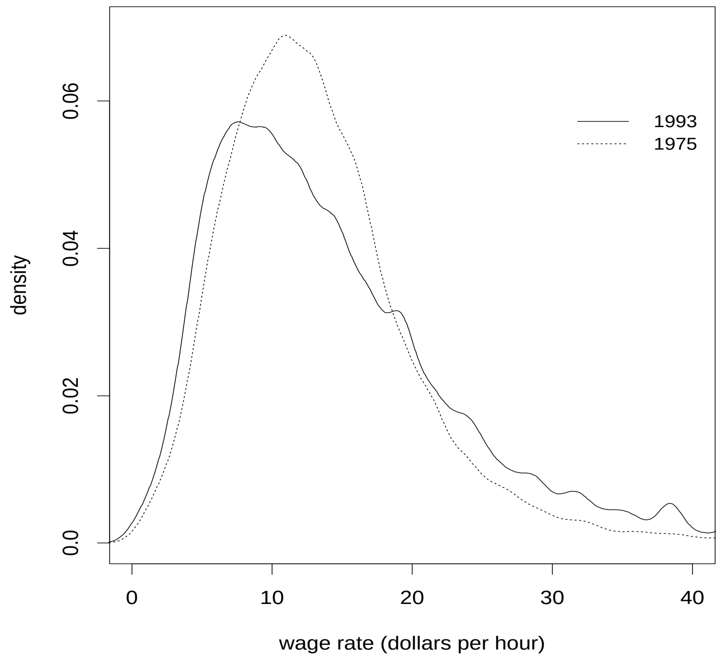

Figure 1. The distributions of hourly wages in 1975 and 1993 expressed in 1993 dollars.

3. APPLICATION TO THE GROWTH IN INEQUALITY

Previous research has shown that the earnings of American workers have undergone a series of dramatic changes since the early 1970s. During the 1970s, real wages began to stagnate, especially for workers with low skill and education levels. Poverty rates rose after decades of steady contraction, and the convergence of black on white earnings slowed noticeably. Most important, by the mid–80s analyses were documenting an unprecedented growth in wage inequality, on the order of 20–30% (Levy and Murnane 1992). Subsequent research has focused on the role played by demographic changes and industrial shifts (Danziger and Gottschalk 1993, Ginther 1995, and the references therein), attempted to distinguish the effects of supply and demand (Bound and Johnson 1992, DiNardo, et al. 1995, Katz and Murphy 1992 and the references therein), and begun to examine the role of technology and the reorganization of work (Cappelli 1996, DiNardo and Pischke 1996, Harrison 1994).

In this context, we use the relative distribution in order to take a closer look at how white men's wages have changed over the last two decades.

3.1 Changes in the Distribution of Hourly Wages

Figure 1 provides a straight–forward graphical comparison of the density of white men's hourly wages in 1993 to that in 1975. As the samples sizes here are $n = 27,273$ for the 1993 data and $m = 20,776$ for the 1975 data, the sampling variability will have only a modest effect on the form of the kernel density estimates. The spread of the 1993 wages is clearly larger than the spread of the 1975 wages. This observation can be verified using summary measures of location, scale, and skewness. In particular, the mean of the 1993 values is larger than that of the 1975 values, while the median is smaller. This suggests the importance of changes in distributional shape.

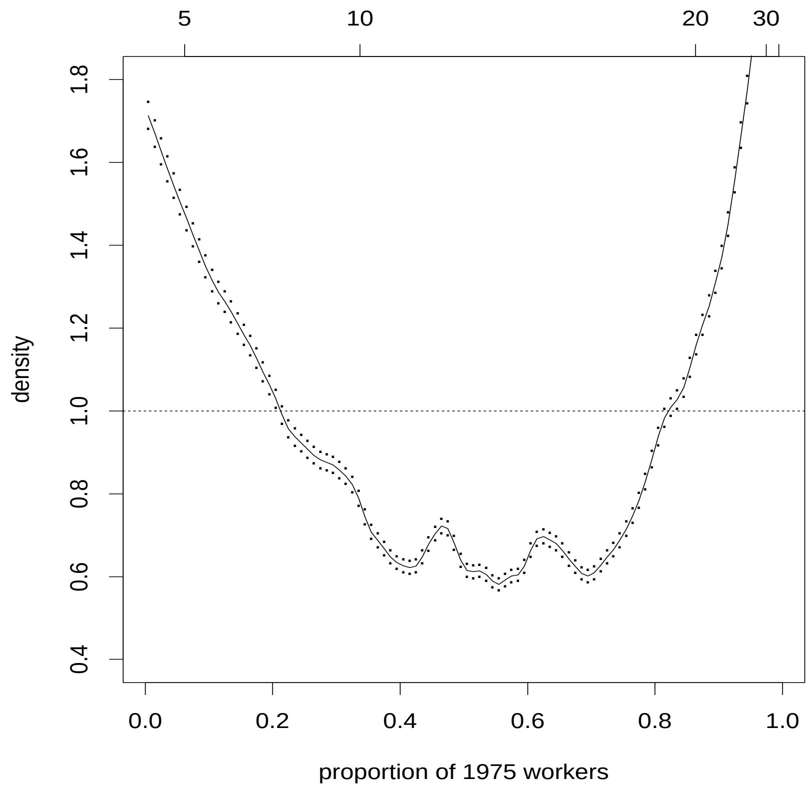

Figure 2. The relative density of hourly wages in 1993 to 1975 expressed in 1993 dollars. The upper axis is labeled in thousands of 1993 dollars. The dotted lines are 95% pointwise confidence bounds.

Figure 2 shows the relative density of the 1993 to the 1975 hourly wages. If the two distributions were identical, the relative density would be uniform. We can see, however, that there is a substantial difference between the shapes of the two distributions. The hourly wages of 1993 workers are over–represented in the lower and upper quantiles of the 1975 hourly wages. They are correspondingly under–represented in the middle 60% of the 1975 distribution. The frequency of 1975 workers does not match that of 1993 workers until about the 20% quantile and again at the 80% quantile of the 1975 hourly wages distribution. These observations are not readily apparent from the direct comparison of the density estimates.

The relative density enhances comparison of the two densities in several ways. First, it directly compares their relative frequency in terms of a ratio, which is easier to understand both visually and numerically (Cleveland and McGill 1984). In contrast to the direct graphical overlay in Figure 1, the relative density brings into sharp focus differences between the two individual distributions at each point on the scale. In this sense, it is similar to a difference graph (Chambers et al. 1983) – it provides a visual signal for something that is visible but not easy to process in the original overlay graph. For example, we see that the 1993 workers were over 1.8 times more likely to earn at the highest wage levels than the 1975 workers, but also about 1.7 times more likely to earn at the lowest wage levels. We also obtain a direct visual signal of the "declining middle class:" workers in 1993 were only 70% as likely to fall in the interquartile range of the 1975 earnings distribution.

The relative distribution provides a powerful tool for identifying differences in shape between two distributions, while still remaining close to the original data. It allows the researcher to isolate key characteristics of the individual distributions, and (like any good graphical method) the result is a much more accessible, intuitively meaningful and informative analysis than that afforded by summary statistics such as the mean, interquartile range, or the Gini index (Kakwani 1980). The relative distribution does not replace the direct graphical overlay; rather, it complements the overlay, by focusing on those characteristics of the individual distributions important for comparing the two. Taken together, the two figures provide both absolute and relative comparisons.

Figure 2 clearly shows that the hourly wage distribution of white men has polarized during the last two decades. The estimate of the MRP of the relative distribution of hourly wages in Figure 2 is 0.164. A 95% confidence interval for the MRP is 0.154 to 0.174, indicating that the MRP is statistically significantly greater than zero. The size and sign of the estimate confirms the impression left by the graphical display.

For comparison, two Gaussian distributions with the same MRP would have a ratio of standard deviations of 1.30. The full set of polarization indices is shown in Table 1.

Table 1. Relative Polarization Indices for Earnings from 1993 to 1975

Polarization Index	Estimate	95% Confidence Interval	p–value
Median Index	0.164	0.154 – 0.174	0.000
Lower Index	0.135	0.095 – 0.175	0.000
Upper Index	0.193	0.152 – 0.233	0.000

The estimated lower and upper polarization indices for Figure 2 are 0.135 and 0.193, respectively. This indicates that both the upper and lower parts of the hourly wages distribution are significantly positively polarized, but that upgrading in wages was the stronger trend.

3.2 Changes in the Distribution of Weekly Hours Worked

To explain the rise in wage inequality, researchers have started to look at the restructuring and reorganization of the American firm and its effects on workers (Cappelli 1994, Harrison 1994). In particular, studies have documented a dramatic increase in the use of part–time and temporary workers (so called "contingent" workers) as substitutes for full–time workers (so called "core" workers) during the period examined here. This change in business practices might be expected to generate a bifurcation of work status in the labor force, and may in part explain the growth in the dispersion of hourly wages: hourly wages for part–time workers are on average 70% of those of full–time workers (Belous 1989).

One measure of this change in the mix of part and full–time workers is the distribution of weekly hours worked. Substitution of contingent for full–time workers should result in an increased dispersion in weekly hours worked, and stylistically take the form of polarization. If such a pattern obtains, it would suggest that research pursue the connection between firm–level restructuring and reorganization on the one hand, and increased wage inequality on the other.

The relative distribution can easily be used to test whether the distribution of weekly hours worked has changed over time. This example displays the applicability of these methods to discrete data. Figure 3 shows the relative distribution of weekly hours worked per year in 1993 to that in 1975 (we will use the term work schedule below). Conceptually, this relative distribution is similar to the one constructed for wages, though there is no need for inflation here as the scale is the same in both time periods. The graph is not nearly as smooth because of natural discreteness in reported hours around standard work week schedules (e.g., 35, 37.5 or 40 hours per week). Results concerning the estimation and inference issues for discrete data are given in the Appendix.

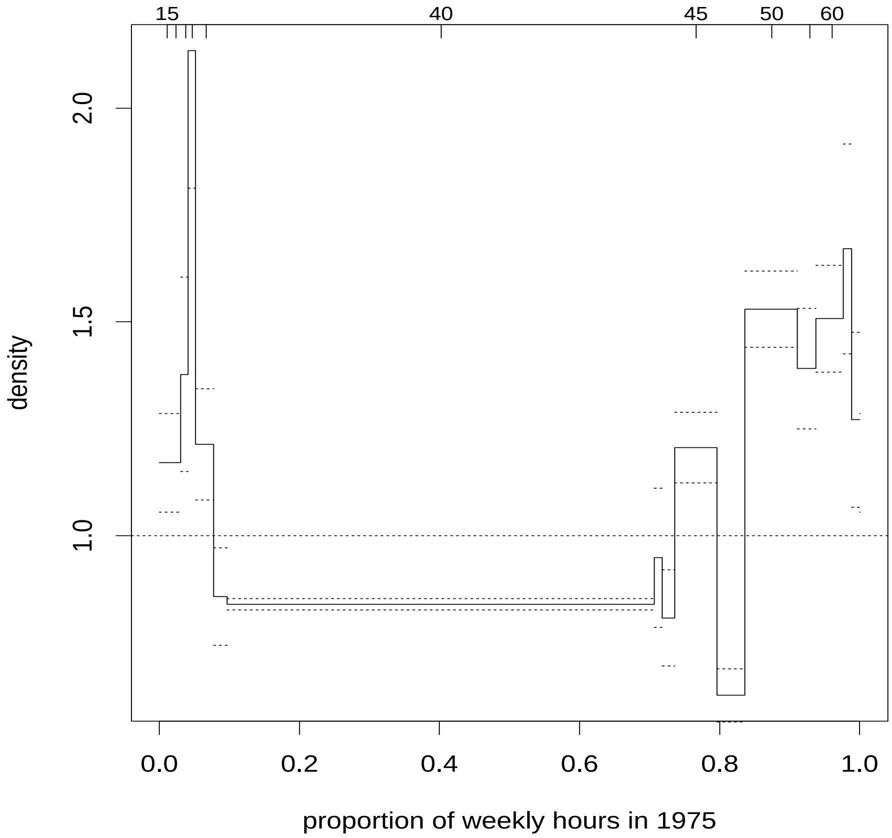

Figure 3. The relative distribution of usual weekly hours worked in 1993 to that in 1975. The upper axis is labeled in 1993 weekly hours worked. The dotted lines are 95% pointwise confidence bounds.

A greater proportion of workers in 1993 are reporting lower and (especially) higher weekly hours, giving the figure the U–shape of a polarized relative distribution. The bar from the 10% through the 70% quantiles in the figure represents individuals working 40 hours per week, indicating that $70\% - 10\% = 60\%$ of workers in 1975 were working the equivalent of a standard 40 hour week. The relative density for this group is 0.84, indicating that such workers were only 84% as common in 1993 (that is, $84\% \times 60\% = 50\%$ were working the equivalent of a standard 40 hour week). Note also the thin peak of 2.1 at 30 hours indicating that this group was relatively rare in

1975 and over twice as frequent in 1993. Another interesting feature is the relative infrequency (62%) of the 48 hour week in 1993, breaking the overall pattern of over–time hours worked. These workers are likely to be working six standard eight hour days per week rather than the usual five days. The fraction of workers reporting less than the standard full–time work week has grown by about 30% from 1975–1993. This is consistent with information from standard sources on the rise of the temporary help industry and the increasing number of part–time jobs (see, for example, Belous 1989). Just as interesting is that substantially more of the 1993 workers are putting in *longer* working hours than their counterparts in 1975. There are proportionately 50–60% more individuals working the longest hours in 1993 compared to 1975. Thus there has been a clear divergence in weekly hours worked over this time frame: a hollowing out of the center of the distribution reflecting a decline in the standard work week and a concomitant rise in the fraction of part–time and "over–time" workers.

The estimate of the MRP of the relative distribution of the work schedule in Figure 3 is 0.104. A 95% confidence interval for the MRP is 0.093 to 0.114, again indicating that the MRP is statistically significantly greater than zero. This is consistent with the graphical display, and corresponds to a ratio of standard deviations of 1.18 for Gaussian distributions. The full set of polarization indices is shown in Table 2. The estimated lower and upper relative polarization indices are 0.015 and 0.223, indicating relatively more growth in the upper tail.

Table 2. Relative Polarization Indices for Hours Worked (1993 to 1975)

Polarization Index	Estimate	95% Confidence Interval	p–value
Median Index	0.104	0.093 – 0.114	0.000
Lower Index	0.015	-0.005 – 0.036	0.138
Upper Index	0.223	0.181 – 0.264	0.000

4. Linking changes in wages and changes in hours worked

The fact that both hourly wages and work schedules have polarized since the mid–70s suggests that there may be a link between the two sets of changes. In this section, we begin to explore this possibility. We disaggregate the overall wage distribution and polarization indices for subgroups of workers defined by the number of weekly hours worked. This approach will allow us to characterize within– and between–group changes in the wage distribution over time.

Figure 4 shows the 1993 to 1975 relative distributions of hourly wages separately within three groups defined by weekly hours worked: those working less than the standard work week ($<$ 35 hours per week), those working a standard work week (35–40 hours per week), and those working more than the standard work week ($>$ 40 hours per week). The graphs were constructed as follows. First, the population in each work group is identified, separately for 1975 and 1993. We then form the relative distribution (within each subgroup) of hourly wages in 1993 to hourly wages in 1975. Conceptually, we are comparing the distributional shape of the hourly wages of a given group of workers to their counterparts at an earlier time.

The wage distributions graphed in Figure 4 have not been median–matched within each group (though the overall distributions were median matched). As a result, the relative density represents both within–group shape shifts and relative group median shifts (see Bernhardt, Morris and Handcock, 1995 for a discussion of decomposing the relative distribution into median and shape shifts). A relative median shift would occur if the subgroup made relatively greater (lesser) median gains than the workforce as a whole. In this case, only the over–time group had a substantial median shift – their median earnings rose 17% more than the median earnings of the entire workforce – hence the upward sloping relative distribution in their graph. Full–time workers experienced a modest relative median loss of 4%, and part–time workers held level. Thus their two panels in Figure 4 essentially represent the shape changes of interest here. To isolate the pure shape changes, the polarization indices – which net out the median shifts by median matching within group – are presented in Figure 5.

The disaggregated findings can be interpreted as follows. If each of the group–specific polarization indices is close to 0, then holding changes in work schedule constant, there is no residual polarization in wages. This would suggest that between group differences – that is, the polarization in work schedules – may have generated the polarization in hourly wages. If instead all of the group–specific polarization indices are essentially the same as the overall workforce indices, then holding the changes in work schedule constant does nothing to reduce the observed polarization in wages. This would suggest that the polarization in work schedules has contributed little to the polarization in wages. In the first case, a median–matched graph of the relative density would show a uniform distribution, in the second it would show the same polarization as the overall relative distribution.

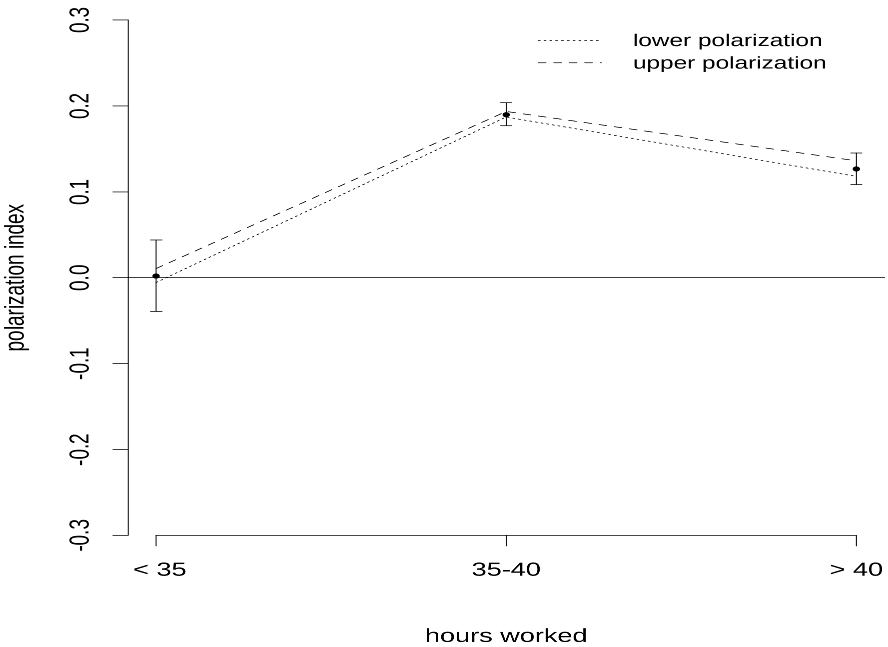

Figure 5. Relative polarization indices for the relative distribution of hourly wages for workers in 1993 to workers 1975 for three different weekly hours worked groups. The error bars represent 95% pointwise confidence bounds for the MRP.

In fact, we see a mix of these two scenarios. Workers on the standard work week show a wage distribution that is similar in shape and slightly more polarized than that of the population as a whole (compare Figures 4 and 5 to Figure 2 and Table 1). Those who worked less than the standard work week show a remarkably uniform distribution, indicating that the distribution of hourly wages for this group has little changed over this time period. In contrast, those who worked more than the standard work week show both an upshifting and substantial polarization. As their average hourly wage is about 9% higher than the overall workforce in 1993, we might expect these distributional shifts to amplify the upper polarization in the overall wage distribution beyond that expected by the growth in their numbers.

At this point, we can make several initial conclusions. The shifts in work schedules that we observed do not *completely* account for the polarization in wages, because there is evidence of residual polarization within the main group of full–time workers. At the same time, however, the shifts in work schedules likely had *some* effect. This effect is difficult to establish from the three within–group graphs because the scales on the horizontal axis are standardized to within–group quantities, so the group position on the overall scale (visible on the top axis) is not easy to decode. The question, then is how large the composition effect is.

5. ADJUSTING THE RELATIVE WAGE DISTRIBUTION FOR CHANGES IN WEEKLY HOURS WORKED

Here we present a method for decomposing the relative distribution of wages, in order to isolate the effect of changes in the distribution of hours worked. Our goal is to arrive at a relative distribution of hourly wages that is "adjusted for" the changing distribution of work schedules. That is, it represents what the relative wage distribution would have looked like if there had been no change in the distribution of hours worked. In the process, we also isolate a relative distribution that measures the compositional effect. The details of this approach are provided in Handcock & Janssen (1995a).

5.1 Decomposing the Relative Distribution of Hourly Wages

Let Y_A be the random variable that describes the (hypothetical) population of work schedules in 1975 with a marginal distribution equal to that in 1993 and with the conditional wage distributions given the weekly hours worked as in 1975.

We can now determine the effects of the differences in the covariate (work schedule) by comparing Y_A to Y_0. Let $X_{A,0} = F_0(Y_A)$ be the relative distribution of Y_A to Y_0. We can interpret $X_{A,0}$ as the random variable describing the effect on the 1975 wage distribution of the distributional difference in work schedules between 1993 and 1975. It reflects the direct compositional effect of weekly hours worked. Note that $X_{A,0}$ will have a uniform distribution when the relative distribution of work schedules in 1975 and 1993 (Figure 3) is uniform.

Let $X_{1,A} = F_A(Y)$ be the relative distribution of Y to Y_A. This relative distribution compares Y to Y_0 with the same work schedule distribution. Note that $X_{1,A}$ will have a uniform distribution when the only difference between the two populations is the distribution of work schedules. We can interpret $X_{1,A}$ as the random variable describing the differences between Y and Y_0 not directly due to differences in the work schedules. It is the relative distribution had the distribution of hours worked been the same in the two populations.

These two effects form a decomposition of the relative distribution of Y to Y_0

in the sense that $X_{1,A}$ is the relative distribution of $X_{1,0}$ to $X_{A,0}$. The decomposition can be graphically represented in terms of the corresponding densities. Denote the densities of $X_{A,0}$ and $X_{1,A}$ by g_0^A, and g_A^1, respectively. Mathematically the relationship between the densities is:

$$g(r) \;=\; g_0^A(r) \times g_A^1(p) \qquad\qquad \text{where} \quad p = F_{A,0}(r), \quad 0 \le r \le 1$$

Note that if r is the percentile in 1975 for a given wage then p is the percentile in the hypothetical population of wages of that same value. Heuristically, we can represent the decomposition of the relative densities by:

$$\begin{array}{c}\text{overall relative}\\ \text{density}\end{array} = \begin{array}{c}\text{relative density representing}\\ \text{compositional effect of hours worked}\end{array} \times \begin{array}{c}\text{hours–adjusted}\\ \text{relative density}\end{array} \quad (5.1)$$

By comparing plots of g, g_0^A, and g_A^1 side–by–side we can gauge the relative size and nature of the components.

Figure 6 graphically represents the decomposition (5.1) of the relative wage density in terms of the impact of changes in weekly hours worked. Figure 6(a) is simply the (unadjusted) relative density of Figure 2 that is being decomposed. Figure 6(b) represents the component of (a) that is attributable to the effect of relative changes in the distribution of weekly hours worked. Figure 6(c) then represents the *adjusted* relative density of 1993 to 1975 wages.

The shift in weekly hours worked had a modest polarizing effect on the distribution of hourly wages, an effect due to the positive relationship between hours worked and hourly wages. In terms of the magnitude of the effect, Figure 6 makes clear that the shift in work schedules did not drive the majority of the rising inequality in wages (since Figure 6(c) shows a considerable amount of residual polarization). At the same time, the decomposition has clearly isolated one of the factors contributing to the rise in wage inequality, as the graph of the effect (Figure 6(b)) shows noticeable polarization.

6. Discussion

We have presented a general framework for the comparative analysis of distributional difference and change. The framework is based on the relative distribution, a concept that builds on earlier techniques such as P–P plots and comparison densities, and expands the central insight of these techniques into a comprehensive and flexible framework for distributional analysis. The relative distribution can be used as the basis for exploratory, descriptive and analytical techniques. With the increase in information and complexity that accompanies the shift from simple means–based to fully distributional analyses, visualizing the data is an important aid to understanding. We show that the graph of the relative density provides a powerful tool for identifying differences between distributions, remaining close to the original data and retaining a simple intuitive meaning. A key question of interest in many applications is the comparative density in the tails of a distribution. We therefore develop a set of summary measures based on the relative density that can be used to test hypotheses about distributional polarization, and the comparative changes in the upper and lower tail. These polarization indices play the same role as the Gini coefficient or other indices of relative dispersion used to represent inequality in economic applications. Finally, in most application settings description and summarization of the patterns is followed by multivariate analysis. We present here a bivariate decomposition technique for isolating the marginal impact of changes in population composition. This decomposition enables one to distinguish the impact of changes in population mix (a demographic process) from changes in attribute allocation (a social or economic process). We have developed a complementary decomposition technique elsewhere that can be used to identify the contribution of within–group median and shape changes to overall between–group changes in relative distribution (Bernhardt, Morris and Handcock, 1995). Taken together, the integration of these different components into a single analytic framework provides a powerful approach to the analysis of full distributional information.

In our application of these methods, we analyze the growing inequality in white men's hourly wages from 1975 to 1993. The U–shaped graph of the relative wage density and the polarization indices document the marked polarization of the wage distribution that has occurred during this period. The "declining middle class" is clearly evident in the graph: the fraction of middle–wage earners drops by about 30%, while the fraction in the highest and lowest wage levels each rises by nearly 80%, We find a similar polarization has occurred in weekly hours worked during this period. The relative work schedule density graph is also U–shaped: the fraction working the standard 35–40 hour work week falls by 16%, and there are corresponding increases in both part–time and over–time workers. The polarization indices here indicate that the relative increase in over–time workers is substantially larger than the increase in part–

time workers. This similarity in marginal wage and work schedule changes suggests that the change in the mix of part, full and over–time workers could be driving a substantial part of the growth in wage inequality. Using stratified analyses and the decomposition technique, however, we find this is not the case. The exploratory graphics demonstrate substantial residual polarization in the main full–time work group (which continues to comprise about 50% of the workforce), and in the over–time group. The findings from the decomposition analysis do show a modest polarizing effect of the compositional change in work schedules, but the residual polarization is about twice as large as the composition effect all the way across the earnings scale. Unexpectedly, the composition effects were slightly weaker in the highest earnings level. Thus the greater increase in upper tail of the work schedule distribution did not account for much of the increase in the upper tail of the wage distribution.

A potentially more complicated story may lie behind these findings, one that underlines the importance of developing full multivariate techniques for analyzing relative distributions. From exploratory work not shown here, we know that young workers in 1993 were more likely to be in the lower quantiles of the wage distribution than in 1975. The older workers, by contrast, were more likely to be in the upper wage quantiles in 1993 than in 1975. A similar pattern holds in the relation between age and hours worked. Compared to their 1975 counterparts, young workers in 1993 were relatively more likely to be employed in part–time jobs. For mature workers, this trend is much less pronounced and is more than offset by a move into longer numbers of weekly hours worked. Thus, a portion of the age differences in the rise in wage inequality may be a function of age differences in the changing distribution of weekly hours worked. The findings above, however, indicate that a similarity in marginal changes does not necessarily translate into a connection at the joint distributional level. This is a compelling task for future research, both substantively and methodologically.

Appendix: technical results

In this appendix we summarize results about the estimates of the relative density and polarization indices for discrete data used to construct the graphical displays and numerical summaries used throughout the paper. The results for continuous data and further development can be found in Handcock and Janssen (1995a,b).

Measurement of the relative density based on group–level information

If the distributions are discrete a development similar to that give for continuous data is quite natural. Let the common support of the random variables Y and Y_0 be $\{x_i\}_{i=1}^{Q}$ where Q need not be finite. The discrete relative density can be defined by:

$$g(i) = \frac{F(x_i) \; - \; F(x_{i-1})}{F_0(x_i) \; - \; F_0(x_{i-1})} \qquad i \; = \; 1, \ldots, Q$$

where $F(x_0) = F_0(x_0) = 0$. A graphical display similar to the plot of the relative density in the continuous case can be created by plotting $\{g(i), F_0(x_i)\}_{i=1}^{Q}$. The estimation of, and inference for, the relative density is straight forward where F and F_0 are estimated by F_{n0} and F_m, respectively.

Property A.1:
The estimate of the relative density is

$$\widehat{g}(i) = \frac{F_m(x_i) \; - \; F_m(x_{i-1})}{F_{n0}(x_i) \; - \; F_{n0}(x_{i-1})} \qquad i \; = \; 1, \ldots, Q.$$

The distribution of $\widehat{\mathbf{g}} = \{\widehat{g}(1), \ldots, \widehat{g}(Q)\}$ is

$$\widehat{\mathbf{g}} \quad \sim \quad \mathrm{AN}\left\{\mathbf{g}, \frac{1}{m}\Omega_p \; + \; \frac{1}{n}\Omega_q\right\}$$

where

$$\Omega_p \quad = \quad \begin{cases} -g(i)g(j) & i \neq j \\ g(i)(\frac{1}{\pi(i)} - g(i)) & i = j \end{cases}$$

$$\Omega_q \quad = \quad \begin{cases} -g(i)g(j) & i \neq j \\ g^2(i)(\frac{1-\pi(i)}{\pi(i)}) & i = j \end{cases}$$

where

$$\pi(i) \; = \; F_0(x_i) \; - \; F_0(x_{i-1}), \qquad i \; = \; 1, \ldots, Q$$

as $m \to \infty, m/n \to \kappa^2 < \infty$.

The MRP for the group–level data can be defined by:

$$RP(F;\; F_0) \quad = \quad 1 + 8\int_0^{\frac{1}{2}} \widehat{G}(x)\, dx \quad - \quad 4\int_0^1 \widehat{G}(x)\, dx$$

where $\widehat{G}(x) = \sum_{i: x_i \leq x} g(i)$. This expression is analogous to the definition for continuous data (See Handcock and Janssen 1995b)

Estimation of the group–level median relative polarization index

Property A.2:

$$\begin{aligned}
\mathbb{E}\left[RP(\widehat{F};\ F_0)\right] &= RP(F;\ F_0)\ +\ \frac{(Q-1)}{n}RP(F;\ F_0) + o(\frac{1}{n}) + o(\frac{1}{m}) \\
\mathbb{V}\left[RP(\widehat{F};\ F_0)\right] &= \frac{16}{(Q-2)^2}\left[(Q+1)^2 p^0_{Q_2} - 2(Q+1)p^1_{Q_2} + p^2_Q - \mu_r^2\right]\frac{1}{m} \\
&+ \frac{16}{(Q-2)^2}\left[(Q+1)^2 q^0_{Q_2} - 2(Q+1)q^1_{Q_2} + q^2_Q - \mu_r^2\right]\frac{1}{n} \\
&+ o(\frac{1}{n}) + o(\frac{1}{m})
\end{aligned}$$

where

$$p^l_k = \frac{1}{Q}\sum_{i=1}^{k} i^k g(i) \qquad q^l_k = \frac{1}{Q}\sum_{i=1}^{k} i^k g^2(i) \qquad l = Q_2, Q, \quad k = 0, 1, 2$$

$$\mu_r \ = \ (Q+1)p^0_{Q_2} - 2p^1_{Q_2} + p^1_Q$$

$Q_2 = \lceil(Q/2)$.

In addition, $RP(\widehat{F};\ F_0)$ is asymptotically Gaussian as $m \to \infty, m/n \to \kappa^2 < \infty$.

Estimation of the upper and lower indices of polarization

Property A.3:

Under the hypothesis $H_0 : F = F_0$,

$$\mathbb{E}\left[LRP(\widehat{F};\ F_0)\right] \ = \ \mathbb{E}\left[URP(\widehat{F};\ F_0)\right] \ = \ \begin{cases} \frac{1}{n+1} & n \text{ even} \\ \frac{1}{n} & n \text{ odd} \end{cases}$$

$$\mathbb{V}\left[LRP(\widehat{F};\ F_0)\right] \ = \ \mathbb{V}\left[URP(\widehat{F};\ F_0)\right] \ = \ \frac{\sigma^2_{\text{LRP}}}{nm}(m+n+1),$$

where

$$\sigma^2_{\text{LRP}} \ = \ \begin{cases} \frac{5}{3} \ - \ \frac{2}{3(n+1)^2} & n \text{ even} \\ \frac{5}{3} \ - \ \frac{2}{3n(n+1)} & n \text{ odd} \end{cases}.$$

In addition, $LRP(\widehat{F};\ F_0)$ and $URP(\widehat{F};\ F_0)$ are asymptotically Gaussian as $n \to \infty$ or $m \to \infty$ or both.

REFERENCES

BARTEL, A. & DINARDO, J. (1995), "Technological Change and the Skill Acquisition of Young Workers," National Bureau of Economic Research, Working Paper No. 5107.

BELOUS, R. S. (1989), *The Contingent Economy: The Growth of the Temporary, Part–Time, and Sub–Contracted Workforce.* NPA Report #239, Washington, D.C.: National Planning Association.

BLAU, F. D. & KAHN, L. M. (1994), "Rising Wage Inequality and the U.S. Gender Gap," *American Economic Review* **84**, 23-28.

BUTLER, R. J. & MCDONALD, J. B. (1987), "Interdistributional Income Inequality," *Journal of Business and Econometrics Statistics* **5**, 13-18.

BOUND, J. & JOHNSON, G. (1992), "Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations," *American Economic Review* **82**, 371-391.

CAMBELL, G. (1994), "Advances in Statistical Methodology for Evaluation of Diagnostic and Laboratory Tests," *Statistics in Medicine* **13**, 499-508.

CAPPELLI, P. (1994), "The Effect of Restructuring on Employees," in *Looking Ahead: The Restructuring of Employment,* Washington, DC: National Planning Association.

CAPPELLI, P. (1996), "Rethinking Employment," Forthcoming *British Journal of Industrial Relations.*

CHAMBERS, J. M., CLEVELAND, W. S., KLEINER, B. & TUKEY, P. A. (1983), *Graphical methods for data analysis.* CA: Wadsworth.

CLEVELAND, W. S. & MCGILL, R. (1984), "Graphical perception: Theory, experimentation, and application to the development of graphical methods," *Journal of the American Statistical Association* **79**, 531-554.

CWIK, J. & MIELNICZUK, J. (1989), "Estimating density ratio with application to discriminant analysis," *Communications in Statistics* **A**, 3057-3069.

CWIK, J. & MIELNICZUK, J. (1993), "Data–dependent bandwidth choice for a grade density kernel estimate," *Statistics and Probability Letters,* 397-405.

DANZIGER, S. & GOTTSCHALK, J. (1993), Eds. *Uneven Tides: Rising Inequality in America* New York: Russell Sage Foundation.

DINARDO, J., FORTIN, N. M. & LEMIEUX, T. (1995), "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," Forthcoming

Econometrica.

DiNardo, J. & Pischke, J. (1996), "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?," National Bureau of Economic Research, Working Paper No. 5606.

Ginther, D. K. (1995), "A Nonparametric Analysis of the U.S. Earnings Distribution," Discussion paper #1067-95, Institute for Research on Poverty, University of Wisconsin–Madison.

Gijbels, I. & Mielniczuk, J. (1995), "Asymptotic Properties of Kernel Estimators of the Radon–Nikodym Derivative with Applications to Discriminant Analysis," *Statistica Sinica* 261-278.

Grove, D. J. & Hannum, R. (1986), "On Measuring Intergroup Inequality," *Sociological Methods and Research* **15**, 142-159.

Hand, D. J. (1982), *Kernel discriminant analysis.* New York: Wiley.

Handcock, M. S. & Janssen, P. (1995a), "Statistical Inference for the Relative Distribution and Relative Density," Technical Report, Department of Statistics and OR, New York University.

Handcock, M. S. & Janssen, P. (1995b), Statistical "Properties of the Relative Distribution and Relative Polarization Indices," Technical Report, Department of Statistics and OR, New York University.

Harrison, T. (1994), *Lean and Mean: The Changing Landscape of Corporate Power in the Age of Flexibility.* New York: Basic Books.

Hastie, T. & Loader, C. (1993), "Local regression: Automatic kernel carpentry," (Disc: p129-143), *Statistical Sciences* **8**, 120-129.

Hildenbrand, K. & Hildenbrand, W. (1992), "On the Mean Income Effect: A Data Analysis of the U.K. Family Expenditure Survey," in *Contributions to Mathematical Economics,* ed. Andreu Mas–Colell and Werner Hildenbrand. New York: North Holland.

Kakwani, N. (1980), *Income Inequality and Poverty.* New York: Oxford Press.

Katz, L. F. & Murphy, K. M. (1992), "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics* **107**, 35-78.

Levy, F. & Murname, R. (1992), "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations," *Journal of Economic Literature* **30**, 1333-81.

Juhn, C., Murphy, K. M. & Pierce, B. (1993), "U.S. Wage Inequality and the Rise in Returns to Skill," *The Journal of political economy* **101**, 410-422.

MORRIS, M., BERNHARDT, A. D. & HANDCOCK, M. S. (1995), "Women's Gains or Men's Losses? A closer look at the Shrinking Gender Gap in Earnings," *American Journal of Sociological,* **101**, 2, 302-328.

MARRON, J. S. & SCHMITZ, H. P. (1992), "Simultaneous Density Estimation of Several Income Distributions," *Econometric theory* **8**, 476-488.

MURPHY, K. M. & WELCH, F. (1992), "The Structure of Wages," *Quarterly Journal of Economics* **107**, 285-326.

PICOT, G., MYLES, J. & WANNEL, T. (1990), "Good Jobs/Bad Job and the Declining Middle 1967-86," Research paper, No. 28, Statistics Canada, Business and Labor Market Analysis Group, Ottawa, Ontario.

PARZEN, E. (1992), "Comparison change analysis," *Nonparametric Statistics and Related Topics,* 3-15.

SILVERMAN, B. W. (1978), "Density ratios, empirical likelihood and cot death," *Applied Statistics,* 26-33.

SHORACK, G. R. & WELLNER, J. A. (1986), *Empirical Processes with Applications to Statistics.* New York: Wiley.

WILK, M. B. & GNANADESIKAN, R. (1968), "Probability Plotting Methods for the Analysis of Data," *Biometrika* **55**, 1-17.

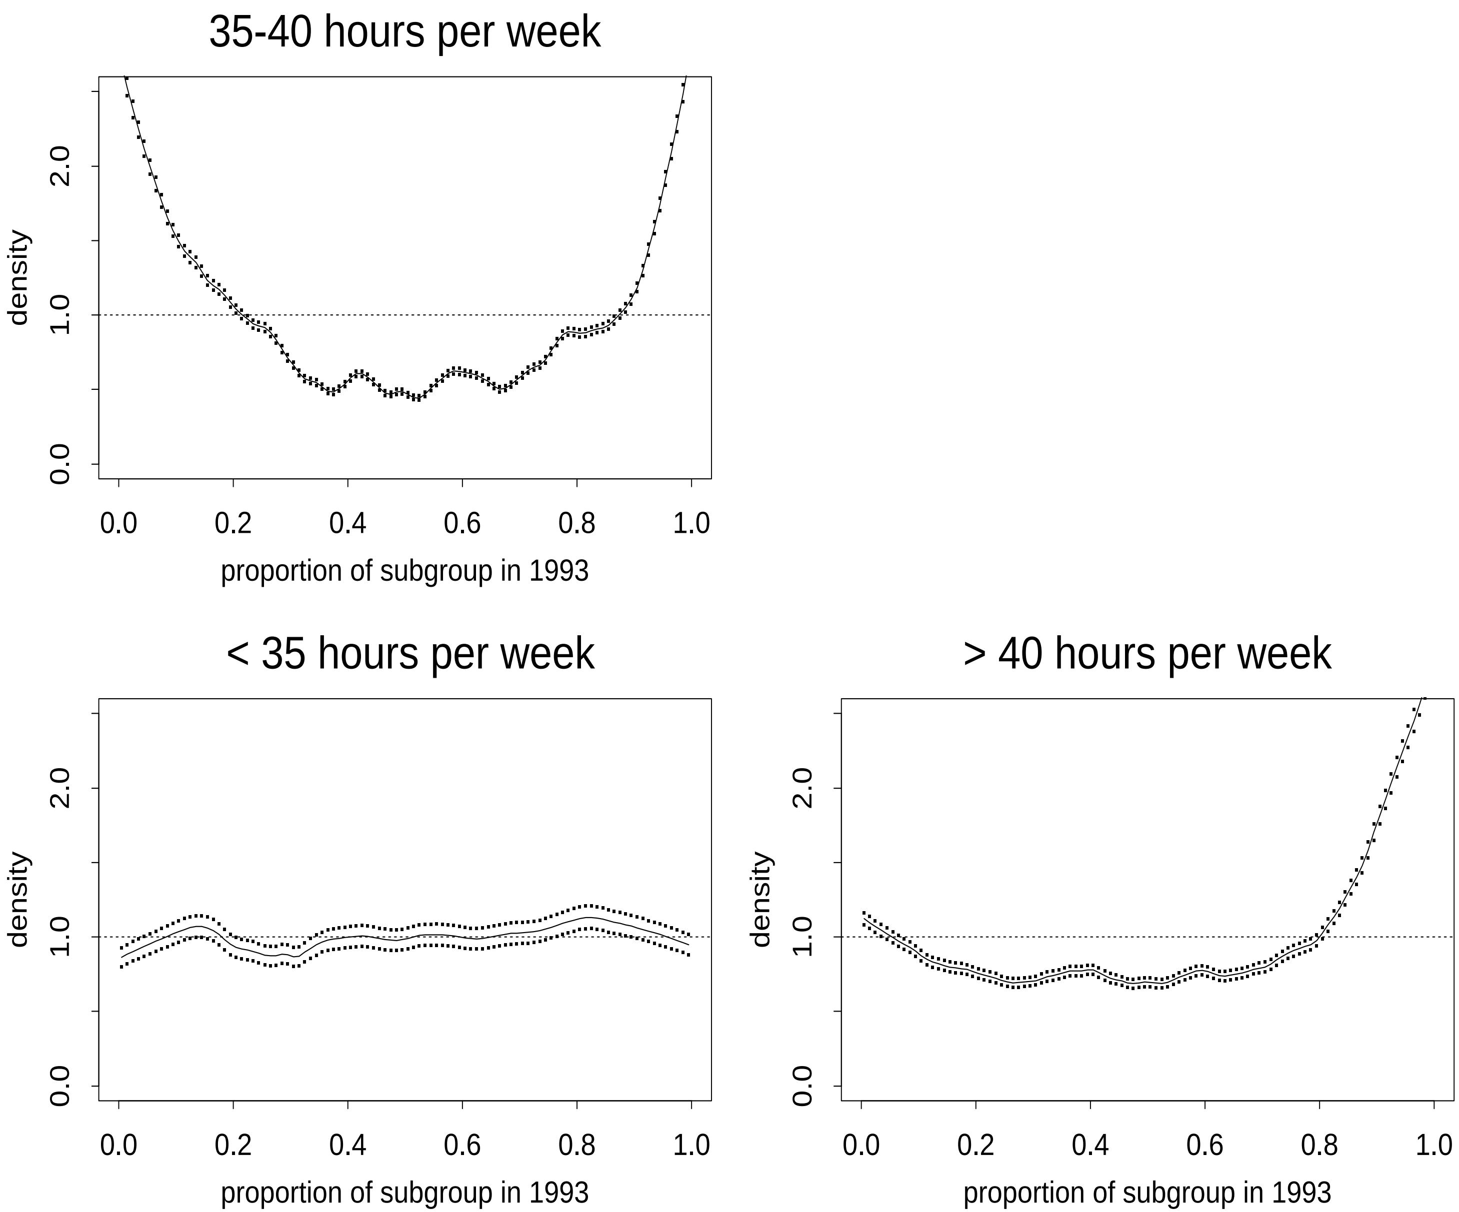

Figure 4. Relative densities of hourly wages for workers in 1993 to workers in 1975 for three different work groups: part–time (< 35 hours per week), full–time (35–40 hours per week), and over–time (> 40 hours per week). The doted lines are 95% pointwise confidence bounds.

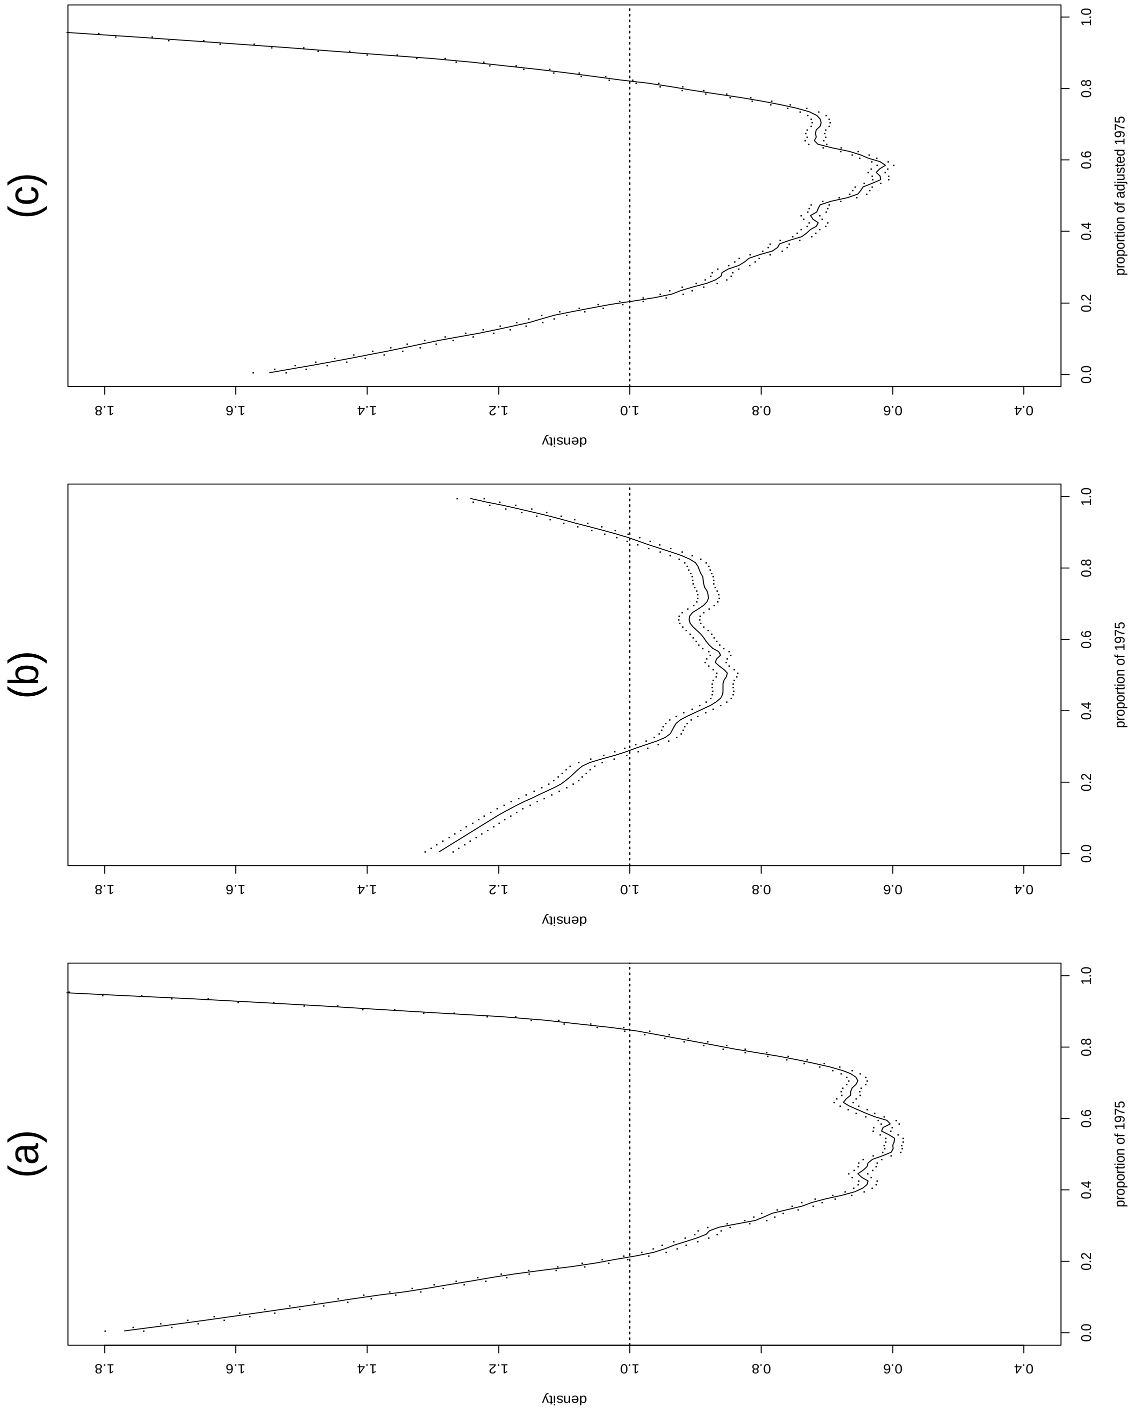

Figure 6. Decomposition of the difference between the distribution of 1975 and 1993 hourly wages. (a) The unadjusted relative density of 1993 to 1975 hourly wages; (b) The relative compositional effect of weekly hours worked from 1993 to 1975 on (c); (c) The relative density of 1993 to 1975 hourly wages adjusted for the relative compositional effect of weekly hours worked.