The Effects of Industry Employment Shifts on the U.S. Wage Structure, 1979–1995

Robert G. Valletta

Economist, Banking and Regional Section. John DiNardo provided very helpful discussions of the statistical approach used. I am grateful to Fred Furlong, Ken Kasa, and especially Mary Daly for their careful review of this paper. I also thank Nan Maxwell and session participants at the 1996 Western Economic Association Meetings for comments on an early version. Randy O'Toole provided useful research assistance.

The trend toward increasing U.S. wage inequality during the 1980s is well documented. I investigate the role of employment shifts from goods-producing to service-producing industries in contributing to increased inequality during the period 1979–1995. Earlier analyses revealed that average earnings are lower, and earnings inequality is higher, for service-producing workers than for goods-producing workers. For both reasons, an increasing share of service employment may increase earnings inequality.

I analyze the effect of broad industry employment shifts by using a recently developed statistical technique, which I term "conditionally weighted density estimation." This technique enables investigation of the effects of changing industry employment shares on the complete distribution of earnings, conditional on changes in other earnings-related characteristics. The results show at most a small effect of industry employment shifts on growing inequality in male hourly earnings.

The trend toward increasing U.S. wage inequality during the 1980s is well documented and extensively analyzed (for example, Bound and Johnson 1992; DiNardo, Fortin, and Lemieux 1996; Juhn, Murphy, and Pierce 1993; Karoly 1992; Katz and Murphy 1992). During this decade, earnings inequality increased both across and within industry sectors and worker groups, and the return to measurable skills (particularly formal education) increased substantially. Research in this area has focused largely on assessing the contribution to rising earnings inequality of factors such as the declining real minimum wage, declining unionism, changing supply and demand across worker groups, increased international trade, and skill-biased technological change. Each of these factors appears to have played a role in increased U.S. earnings inequality during the decade.

An additional factor that may have contributed to increased earnings inequality during the 1980s and earlier, however, is the substantial employment shift in the U.S. from goods-producing to service-producing industries. A common stereotype associated with service-producing jobs is that they pay less than goods-producing jobs. Consistent with this belief, studies such as Blackburn (1990) report that average earnings are lower, and earnings inequality higher, in service-producing jobs than in goods-producing jobs. For both reasons, an increasing share of service employment may increase earnings inequality. Thus, in popular and academic discussion, the shift from goods to services has been cited as a reason for increased inequality and a declining middle class (for example, see Bluestone and Harrison 1982, 1988).

In this paper, I examine the contribution of such industry employment shifts to changing earnings inequality from 1979 to 1995. As described in more detail in Section I, several papers have examined this issue. For example, Maxwell (1989, 1990) and Bluestone and Harrison (1988) both included measures of relative manufacturing employment in their analyses of changing inequality and low-wage employment over the periods 1947–85 and 1963–86, respectively. Each found that employment shifts out of manufacturing have played an important role. However, the use of aggregate time-series data may obscure the role of underlying forces such as changing skill attributes. Blackburn (1990) examined the impact of industry employment shifts on earnings inequality using individual level data from various March Current Population Survey files and found them to have a noticeable but limited influence. In contrast, Murphy and Welch (1993) and Juhn, Murphy, and Pierce (1993) found no effect of industry shifts on average wages and the variance of the wage distribution, respectively.

From an academic perspective, then, the exact contribution of industry employment shifts to rising earnings inequality remains an open question. I attempt to resolve this debate by applying a recently developed methodology that is particularly well suited to analyzing the contribution of broad economic changes to earnings inequality. The technique—which I call "conditionally weighted density estimation"—was recently developed by DiNardo, Fortin, and Lemieux (1996) and applied to the analysis of increased earnings inequality. Their technique enables estimation of the effects of broad economic changes on the entire distribution of earnings. Most studies of rising earnings inequality have focused on explaining changes in the mean or variance of the distribution, or changes in expected wage differentials across labor market groups. In contrast, the conditionally weighted density approach is far less restrictive and applies particularly well when there is no strong *a priori* knowledge about what portions of the earnings distribution are most affected by the factor being examined. For example, conditional weighted density estimation enables examination of whether wages have become more disperse due to widening of the tails or movement from the middle to the tails, a distinction that is important for distinguishing among different explanations of increased inequality (one of which is the "deindustrialization" hypothesis of Bluestone and Harrison 1982).

In general, the technique of DiNardo, et al., enables estimation of a distribution under counterfactual assumptions about the state of the world, which in turn reveals the distributional impact of the state of the world as it actually evolved. My focus is on the effect of changing industry employment shares. In particular, the technique enables me to answer the question, "How would the distribution of earnings look in 1995 if industry employment shares had remained as they were in 1979?" Furthermore, it produces two depictions of how the earnings distribution has been altered by the modeled changes: (1) a visual depiction obtained through comparison of kernel density estimates of the earnings distribution; (2) quantitative comparison based on calculation of parametric inequality measures (standard deviation, quantile dispersion measures, the Gini coefficient, etc.). Both depictions are based on a comparison of calculations that use the original data and survey sampling weights with calculations for which the sampling weights are modified by estimated conditioning weights. This procedure is described heuristically in Section II, with analytic details provided in the Appendix.

To estimate the role of changing industry employment shares, I use data from the 1979 and 1995 Current Population Surveys, as described in Section III. Much of the literature focuses on widening earnings inequality during the 1980s. However, a recent paper by Karoly (1996) finds that increasing inequality continued during the early 1990s. Despite this continued increase during the period covered by my analysis, and despite finding that the service sector exhibits lower average earnings and higher earnings variation, I find at most a small independent impact of industry employment shifts on dispersion in the lower half of the male earnings distribution. These results are described in detail in Section IV of the paper, with conclusions provided in Section V.

I. Earnings Inequality and Changing Industry Employment

A large number of papers in recent years have attempted to attribute increasing earnings inequality during the 1980s to a variety of observable factors (e.g., Bound and Johnson 1992, Katz and Murphy 1992, Blackburn, Bloom, and Freeman 1990, Juhn, Murphy, and Pierce 1993). These authors typically focused on regression-based decompositions or similar analysis based on worker groups defined by earnings-related characteristics, using either aggregated time-series data or yearly individual data.

One recent methodological advance in this literature is the application of kernel density estimation, which provides visual depiction of the entire distribution of earnings. The use of kernel density estimation as an exploratory data analysis tool has long been recognized (see Silverman 1986). In the analysis of changing earnings inequality, kernel density estimates provide a useful visual depiction of how the distribution of earnings has changed over time and where in the distribution the largest changes have been concentrated. Given the lack of strong prior knowledge on where in the distribution the largest changes have occurred, and the focus in the literature on parametric measures such as the variance in earnings, this is an important advance. For example, Levy and Murnane (1992) noted that standard scalar measures of inequality may not distinguish among alternative sources of increasing inequality that have differing economic and social implications, since these measures do not identify the portion of the earnings distribution on which changes have occurred.

Burkhauser, et al., (1996) recently applied kernel density estimation to the analysis of changing inequality. They examined changes in the distribution of family earnings in the U.S., U.K., and Germany during the 1980s. In this form, kernel density estimation serves essentially as a smoothed histogram, thereby providing visual insight into changing inequality. Burkhauser, et al., found that rising inequality in family earnings in the U.S. was characterized primarily by large but unequal income gains in the middle of the family income distribution.

Although kernel density estimation is useful for such exploratory analysis and visual characterization of distributions, its direct use as an analytical tool is limited. In contrast, conditional density estimation enables a full range of analytical applications. Conditional density estimation methods proceed by reestimating the entire distribution of earnings after accounting for various earnings determinants, or by reweighting the distribution according to conditional probabilities. For example, Juhn, Murphy, and Pierce (1993) applied a regression-based conditioning approach. They used the cumulative distribution function of residuals obtained from wage equations to decompose changes in inequality measures into portions due to changes in observable personal characteristics, changes in the returns to those characteristics, and changes in the distribution of unobservables. They found an increasing contribution of unobservables to rising earnings inequality in the 1980s.

DiNardo, Fortin, and Lemieux (1996; henceforth DFL) and DiNardo and Lemieux (1994) improved on previous methods by conditioning through the use of estimated weights. They combined the estimated conditioning weights with sample survey weights to produce an adjusted earnings distribution. This is a flexible procedure that provides semiparametric estimates of the entire distribution of earnings under various counterfactual assumptions. The adjusted distribution can be compared with the original distribution both visually, using appropriately reweighted kernel density estimates, and quantitatively, by comparing dispersion measures from the adjusted and unadjusted distributions.

DFL used their technique to estimate how much earnings inequality would have risen between 1979 and 1992 if the real minimum wage and union membership density in the U.S. had remained at their 1979 levels. Comparison to the actual amount by which earnings inequality rose revealed the impact of the declining minimum wage and declining union membership, conditional on changes in other important variables (such as individual skill attributes). Because the minimum wage affects only the lower portion of the earnings distribution, the technique's ability to reveal features of the entire distribution is particularly salutary. Both papers reported important contributions of a declining real minimum wage and declining unionism to increasing U.S. earnings inequality during the period 1979–88.

These authors, however, did not examine the role of changing industry employment patterns. During most of the postwar period, the share of service-producing jobs in the U.S. has increased substantially. These shifts will alter the distribution of earnings if either the level or dispersion in earnings is different across the goods-producing and service-producing sectors.

Previous work that analyzed the effect of industry employment shifts on earnings inequality typically used aggregated data. Using aggregate time series data, Maxwell (1989) found that the increasing share of service sector employment relative to manufacturing employment explains a substantial portion of increasing inequality over the period 1947–1985; she attributed much of this to declining unionization (Maxwell 1990). Also using aggregate data, Bluestone and Harrison (1988) found a corresponding effect on low-wage employment for the period 1963–86.

In contrast, Blackburn (1990) examined the influence of changing industry structure and other factors on earnings inequality using individual level data from various March Current Population Survey files and found only a limited impact of industry employment shifts. Similarly, Murphy and Welch (1993) and Juhn, Murphy, and Pierce (1993) found no effect of industry shifts on average wages and the variance of the wage distribution, respectively. Furthermore, Schweitzer and Dupuy (1995) used kernel density techniques and found substantial convergence in the goods-producing sector and service-producing sector wage distributions through 1993, which suggests a limited impact of industry employment shifts on inequality. Thus, evidence on the role of industry employment shifts in increased earnings inequality is mixed.

I build on previous work by using weighted density estimation to assess the contribution of changing industry employment shares to increasing earnings inequality. As noted, this enables more flexible and detailed assessment of the impact of industry shifts on the structure of earnings than do other approaches.

II. Methods

Kernel Density Estimation

Kernel density estimation is a flexible, largely nonparametric means of estimating the underlying distribution from which an empirical distribution is sampled.[1] The estimated densities essentially serve as "smoothed histogram" representations of a distribution, and as such are useful for exploratory data analysis. This subsection describes the basics of kernel density estimation, and the next

1. Silverman (1986) discusses non-parametric density estimation in detail, and Delgado and Robinson (1992) provide a useful summary of econometric applications.

two subsections describe the estimation and incorporation of conditioning weights into density estimation.

The kernel density estimate of a univariate distribution based on a random sample $(W_1, \ldots, W_n)$ of size n with sampling weights $\theta_1, \ldots, \theta_n$ (normalized so that $\Sigma_i \theta_i = n$) is:

$$(1) \quad f_h(w_j) = \frac{1}{n} \sum_{i=1}^{n} \frac{\theta_i}{h} K\left(\frac{w_j - W_i}{h}\right) \quad \text{for } j = 1,2,\ldots, m.$$

In this expression, K is the kernel function, h is the bandwidth, and m is the number of points at which the density function is evaluated.[2] Several alternatives are available for the function K, although they typically are probability density functions (and therefore are symmetric and integrate to 1 over the range of W). For each evaluation point w_j, these functions assign to the W's estimation weights that decline (smoothly or abruptly) as the W's move farther from w_j. The subscript j denotes evenly spaced values of w, with the choice of m depending largely on computing resources and the data. The full estimation essentially involves sliding a window (of width $2h$) across the range of W_i, with m density estimates computed at equal intervals.

The choice of h has been subject to substantial discussion in the literature and is generally acknowledged to be more important than the choice of kernel function. Various "optimal bandwidth selection" rules are available. Rather than investigating this issue in detail, I follow DiNardo and Lemieux (1994) in setting the bandwidth equal to .075 for all ln(hourly earnings) estimates provided below. This falls within the range of bandwidths selected by the optimal method of Sheather and Jones (1991) for similar data in DFL. This bandwidth also does a good job of capturing important visual features of the distribution of hourly earnings, such as the spike at the minimum wage. I use the Epanechnikov kernel function, which yielded results identical to a Gaussian kernel in comparison tests.

Conditional Weighted Density Estimation

In this section, I describe how simple estimated reweighting functions can be obtained and applied to the estimation of earnings distributions that embody counterfactual assumptions. In the text I describe these procedures heuristically; the exact derivation—which is conceptually simple but notationally complex—is provided in the Appendix.

Consider the distribution of wages w in year t, conditional on individual characteristics X and a measure of industry employment patterns E:

$$(2) \quad f_t(w) \equiv f(w; t_w = t, t_{E|X} = t, t_X = t).$$

This identity is notational; it shows that the distribution of w is defined in year t, conditional on the distribution of X and E (conditional on X) in the same year. In the empirical work, I focus on $t_w = 1995$, and I measure industry employment patterns by a dummy variable indicating whether each worker is in the broad goods-producing or service-producing sector.

The essence of the test is to investigate the effect of holding $t_{E|X}$ at earlier year (1979) levels—i.e., to estimate what the distribution of earnings would be if the distribution of goods-producing versus service-producing jobs had remained the same as in 1979. The simplest way to do this is to upweight individuals in the goods sector by a factor that is proportional to the decrease in the share of goods-producing jobs in the economy (and similarly downweight service sector workers). However, this simple test ignores any changes in the relationship between earnings-related characteristics and the probability of being in different broad industry sectors. For example, if the movement to services was exclusively by low-skilled workers, then the shift toward services did not have a substantial independent effect on the earnings distribution. Thus, we need to estimate the 1995 distribution of earnings with the industry employment distribution, *and its relationship to X,* held to its 1979 level.

In terms of the notation in (2), we are interested in:

$$(3) \quad f(w; t_w = 95, t_{E|X} = 79, t_X = 95).$$

This expression represents the density that would be observed if the probability (conditional on individual characteristics X) of being employed in goods-producing industries retained its 1979 level and structure, but workers were otherwise paid according to the earnings schedule prevailing in 1995. As shown in the Appendix, this distribution can be expressed as the original unconditional distribution of earnings in 1995, with observations reweighted by a function $\psi_{E|X}$. This function represents the change in the probability between 1979 and 1995 that an observation defined by characteristics X is observed in the goods-producing or service-producing sector.

Intuitively, to obtain the density of earnings that would prevail if the structure of conditional industry affiliation remained as it was in 1979, we downweight individuals in the 1995 sample whose characteristics would have made them less likely to work in the same sector in 1979 as they worked in 1995. These conditional probabilities can be estimated as the fitted values obtained from standard binary variable models; in the empirical work below, I use the probit model to estimate these conditional probabilities.

2. See Silverman (1986) for a detailed discussion of kernel density techniques.

As long as the unconditional probability of being in the goods sector or the relationship of the X's to that probability have changed, then the estimated weighting function $\hat{\psi}_{E \div X}$ will differ from one, and the counterfactual density will differ from the observed density. In general, because the probability of working in the goods-producing sector declined between 1979 and 1995, compared to the unconditional density the reestimated density will attach more weight to individuals currently working in the goods-producing sector and less weight to those in the service-producing sector.

In addition to accounting for the impact of changes in industry employment shares, the technique enables us to account explicitly for the impact of changes in the X vector of earnings-related characteristics. This serves as a useful basis for comparison and also enables us to account for interactions between the X's and industry structure (as described in the next subsection).

The distributional effect of changes in the X's can be modeled by again estimating weights and applying them to the 1995 earnings distribution (see DFL for details). This estimated weighting function—$\psi_X(X)$—is equal to the relative probability of observing an individual with characteristics X in the 1979 versus the 1995 sample, normalized by the unconditional probabilities of being in either sample. As long as the distribution of X's changed between the two years (for example, through higher average educational attainment), the weights ψ_X will alter the estimated distribution. In the empirical work, the function ψ_X is calculated based on fitted values from probit equations that estimate the probability of observing an individual with characteristics X in the 1979 versus the 1995 data set. Workers in the 1995 sample with characteristics that make them relatively more likely to be observed in the 1979 sample will receive more weight in the conditional density estimation than they do in the unconditional density estimation.[3]

Reweighted Estimates and Comparisons

I now describe how the conditioning weights are used in the estimation. Briefly, the estimated conditioning weights are used to modify the sampling weights; I term this process "conditional weighted kernel density estimation." Comparison of the original and adjusted distributions reveals the effects of interest.

For each of the estimated weighting functions ($\psi_{E \div X}$ and ψ_X), the conditional weighted kernel density estimates are obtained by multiplying the sampling weights for each observation (θ) by the estimated conditioning weights ($\hat{\psi}$). The combination of sampling weights and conditioning weights produces three distributions of earnings:

(1) Population weighted distribution: $f(w,\theta)$
(2) Distribution adjusted for industry employment structure: $f_e(w,\theta\,\hat{\psi}_{E \div X})$
(3) Distribution adjusted for individual characteristics: $f_x(w,\theta\,\hat{\psi}_{E \div X}\,\hat{\psi}_X)$

where

θ = survey sampling weight
$\hat{\psi}_{E \div X}$ = estimated conditioning weight for industry employment structure
$\hat{\psi}_X$ = estimated conditioning weight for individual characteristics.

These new weights can be incorporated directly into the estimation of the kernel densities, which requires only slight modification of equation 1:

$$(4) \quad f_h(w_j) = \frac{1}{n}\sum_{i=1}^{n}\frac{\theta_i\psi_i}{h}K\left(\frac{w_j - W_i}{h}\right) \quad for\ j = 1,2,\ldots,m.$$

The result is a different kernel density estimate for each weighting scheme. Graphical depiction and comparison of the sample weighted and conditionally weighted kernel estimates provide a visual representation of the impact of changing industry distribution and individual characteristics.

Furthermore, the reweighting procedure enables calculation of the effect of the modeled change on any distributional statistic: moments (such as the mean and variance), quantile differences (the difference in earnings measured at specific cumulative points on the distribution), and parametric inequality indices (for example, the Gini and Theil indices). This procedure is particularly simple. Distributional statistics for the adjusted distribution are obtained by replacing the population weights by their product with the estimated conditioning weights when calculating the distributional statistics, a procedure easily handled by software that allows weighted tabulations.

One objection to the procedure outlined above is that it gives industry employment shifts precedence over changing individual characteristics in assessing the contribution of each factor. Given this ordering, any interactions between the two factors—for example, due to increasing concentration of unskilled workers in service industries—will be attributed to industry employment shifts. A useful check on the results, then, is to reverse the ordering of the estimation, which entails accounting for the impact of the X's

3. In attempting to control for changes in educational attainment over their sample frame, Schweitzer and Dupuy (1995, p. 20) apply a restricted version of conditional weighted kernel density estimation: for each observation, they scale the sampling weight up or down to reflect a larger or smaller number of individuals with similar educational attainment in the base year.

first and then assessing the impact of changing industry employment shares. I report results from this procedure below; it requires reformulation of the conditioning weights, as described in DFL.

In terms of the treatment of the conditioning variables (X), the conditional weighting procedure is closely related to standard regression-based decompositions of variance. Regressions typically are used, however, to estimate the mean of a distribution. The advantage of the weighted kernel density procedure is that it estimates the entire conditional distribution, as opposed to analyzing distributional characteristics one-by-one, and therefore provides a more flexible method than regression techniques for investigating distributional changes. Regression techniques would require a potentially lengthy search for the exact effect of industry employment shifts on the wage structure; conditionally weighted density estimation provides an immediate visual representation, and it enables estimation of any desired dispersion measure.

III. DATA

The data used in this study are the merged outgoing rotation group files, or Annual Earnings Files, from the Current Population Survey for the years 1979 and 1995 (CPS–AEF). Each month, members of the outgoing rotation group of CPS sample households (about one quarter of the sample) are asked questions concerning earnings on their current job. Pooled over the 12 months in a year, these files provided me with approximately 150,000 observations per year, after sample restrictions. I dropped observations with allocated values for earnings or hours and limited the analysis to individuals aged 16–64. To focus clearly on the goods/services distinction, I eliminated agricultural workers from the sample. I inflated 1979 earnings to 1995 levels using the GDP deflator for personal consumption expenditures,[4] and I dropped earnings observations with values below $1/hour and above $200/hour (in 1995 dollars).[5]

I focus on hourly earnings data from the CPS–AEF for several reasons. First, this provides a large, representative data set for a period characterized by substantial changes in earnings inequality. An alternative is to use data from the March CPS Annual Demographic Surveys, which collect information on labor market experience and earnings in the entire previous year. Although these data are extensively used in the study of inequality (for example, in Blackburn 1990, Burtless 1990, Juhn, Murphy and Pierce 1993, and Katz and Murphy 1992), the yearly earnings data are affected by job changes and labor supply factors. Also, formation of point-in-time earnings measures in the March CPS requires dividing by weeks worked and hours worked, which may introduce measurement error. The primary alternative—the use of yearly earnings for full-time, full-year workers—would narrow the sample undesirably for my experiment.

I begin my analysis in 1979 rather than earlier in the 1970s because the rate at which inequality increased was faster in the 1980s than in the 1970s, particularly for low-skilled workers (Bound and Johnson 1992). The 1995 data are the most recent available, and they produce the added benefit of enabling comparison across similar points in the business cycle: the unemployment rate was 5.8% in 1979 and 5.6% in 1995. Furthermore, Burtless (1990) found that cyclical effects on inequality were small to nonexistent in the 1980s. Given that he focused on yearly wage and salary earnings, this concern is mitigated further by my use of hourly earnings data, which are relatively insensitive to variation in hours and weeks worked over the year.

I use the CPS sample earnings weights for all estimates reported in this paper. Unlike DFL, however, I do not weight by hours worked. This enables greater flexibility in isolating job composition shifts associated with the shift from goods-producing to service-producing industries. For example, if service sector jobs are more likely to be part-time, and if part-time jobs pay less than full-time jobs, weighting by hours worked would undesirably downweight the wage inequality created by such shifts. Thus, my focus is on the distribution of earnings by job, rather than by hour.

IV. RESULTS

Summary Statistics and Densities

Table 1 shows summary statistics for ln(hourly earnings) for the 1979 and 1995 samples, with separate panels stratified by sex. I list mean ln(earnings), and the standard deviation as a simple dispersion measure. I provide a major industry breakdown in addition to the overall goods/services distinction, and I show employment shares by industry. These figures show substantially higher hourly earnings dispersion in 1995 than in 1979, a large reduction in mean real earnings for men, and a small increase in mean real earnings for women. There was a substantial decline between those years in the share of goods-producing jobs in

4. The deflator does not affect the dispersion measures, but it is useful for comparisons of means over time.

5. I did not directly account for top-coding of weekly earnings. Although this may affect the dispersion measures, the top-code is roughly at the same level in real terms in 1979 and 1995. To the extent that the share of very high wage workers increased over the period, the estimates in this paper may understate increasing dispersion due to increased mass in the upper tail.

TABLE 1

Mean and Standard Deviation of ln(Hourly Earnings), Industry Employment Shares, by Sex, Year, and Industry

A. Men

	1979			1995		
Industry	Mean	S.D.	Share	Mean	S.D.	Share
Total	2.55	.497	1.0	2.47	.641	1.0
Goods-Producing	2.62	.434	.430	2.52	.550	.339
Mining	2.78	.398	.016	2.65	.576	.009
Construction	2.62	.471	.097	2.47	.537	.091
Durable Manufacturing	2.64	.406	.208	2.56	.540	.152
Nondurable Manufacturing	2.57	.448	.108	2.49	.568	.087
Service-Producing	2.50	.534	.570	2.45	.681	.661
Trans., Comm., & Public Utilities	2.73	.446	.095	2.63	.602	.102
Wholesale Trade	2.58	.477	.048	2.49	.619	.053
Retail Trade	2.20	.459	.140	2.09	.621	.160
Finance, Insurance, & Real Estate	2.71	.559	.038	2.71	.679	.046
Services	2.46	.555	.178	2.49	.693	.240
Government	2.72	.433	.071	2.73	.558	.059
Total Observations		74,671			83,931	

B. Women

	1979			1995		
Industry	Mean	S.D.	Share	Mean	S.D.	Share
Total	2.16	0.44	1.0	2.23	0.61	1.0
Goods-Producing	2.21	.351	.197	2.25	.519	.132
Mining	2.49	.413	.003	2.52	.542	.002
Construction	2.28	.389	.010	2.34	.546	.012
Durable Manufacturing	2.27	.342	.091	2.30	.497	.057
Nondurable Manufacturing	2.14	.337	.094	2.19	.524	.063
Service-Producing	2.14	.459	.803	2.18	.540	.868
Trans., Comm., & Public Utilities	2.43	.408	.041	2.45	.543	.046
Wholesale Trade	2.22	.363	.024	2.28	.513	.024
Retail Trade	1.89	.349	.197	1.84	.504	.191
Finance, Insurance, & Real Estate	2.22	.354	.081	2.39	.538	.083
Services	2.18	.436	.414	2.31	.625	.474
Government	2.37	.397	.047	2.47	.502	.051
Total Observations		62,681			82,153	

Note: All tabulations are weighted by the CPS earnings weight, and 1979 earnings were inflated to 1995 levels using the GDP deflator for personal consumption expenditures.

the economy, particularly for men. Most of this change arose from a decline in manufacturing jobs (particularly durable manufacturing) and a rise in jobs in the services (narrowly defined) sector.

A key revelation from Table 1 is little or no convergence in the goods-producing and service-producing earnings distributions between 1979 and 1995. For men, there was a substantial decrease in the mean and substantial increase in the standard deviation in both broad sectors, and mean earnings are much higher in the goods-producing sector than in the service-producing sector. In contrast, for women mean earnings are very similar across the two broad sectors, and became more so over the period. Like men's jobs, however, for women earnings dispersion increased for all sectors, and the service sector as a whole (and for most subcategories) exhibits higher dispersion than does the goods sector. In general, this table is consistent with the view that the shift from goods-producing to service-producing jobs has increased inequality in hourly earnings, although this effect is likely to be much more pronounced for men.

Figures 1 and 2 show kernel density estimates of several unadjusted earnings distributions. Figure 1 shows the 1979 and 1995 distributions of hourly earnings, for men in Panel A and women in Panel B. These figures confirm the pattern identified in Table 1 of increasing dispersion for both men and women, and also the declining mean for men, between 1979 and 1995. For men, Figure 1 reveals that much of the increased dispersion is due to a shift from the middle of the distribution to the lower part, although there also is some added mass in the right tail of the 1995 distribution. For women, there appears to be a more uniform increase in dispersion across the upper and lower portions of the distribution. Furthermore, although not explicitly labeled in these figures, each distribution exhibits a pronounced spike at the real minimum wage, which declined substantially between 1979 and 1995. In their formal analysis, DFL attribute much of the increased inequality during 1979–1992 to the declining real minimum wage; Figure 1 also illustrates this effect, extended out by three years to 1995.

Figure 2 shows the earnings distributions for the goods-producing and service-producing sectors, by year and sex. For both men and women, the distributions in the two sectors have become more alike over time. However, each of these distributions became more disperse between 1979 and 1995, with the degree of dispersion remaining higher in services than in goods in all cases.[6]

FIGURE 1

EARNINGS DISTRIBUTIONS, 1979 AND 1995

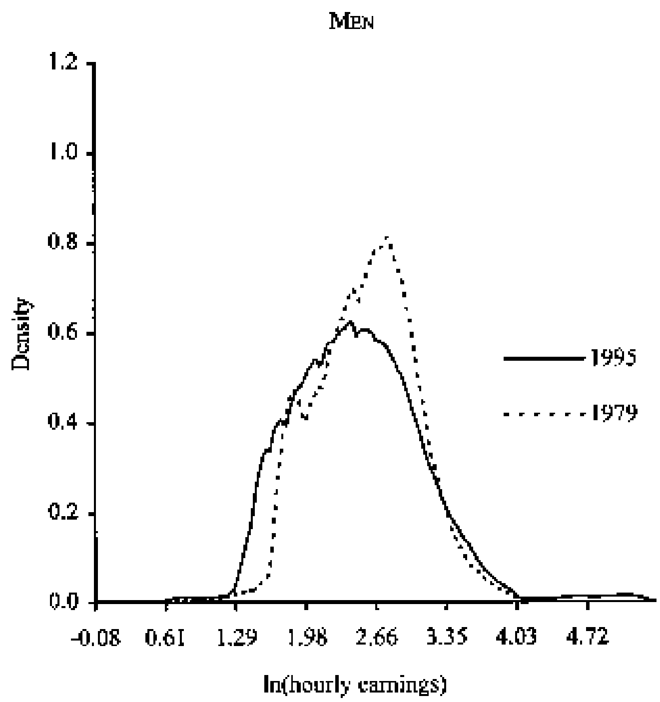

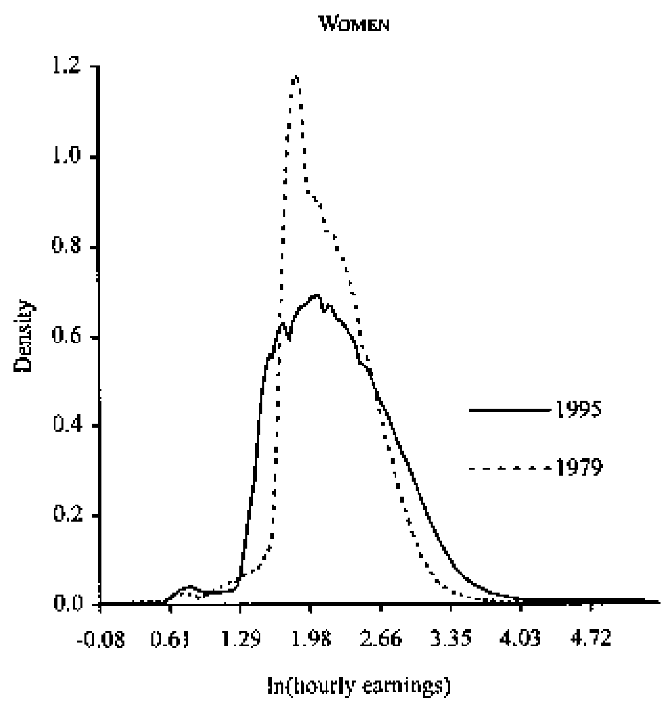

6. These tabulations and figures conflict somewhat with the results of Schweitzer and Dupuy (1995), who reported a substantial increase in overlap of the goods and service sector earnings distributions from 1979 to 1993. To the extent that my results differ from theirs, it is probably due to different sample definition: they used March CPS data on full-time workers who worked at least 39 weeks in the previous year, and they pooled men and women in their sample.

FIGURE 2

EARNINGS DISTRIBUTIONS, 1979 AND 1995, BY SEX

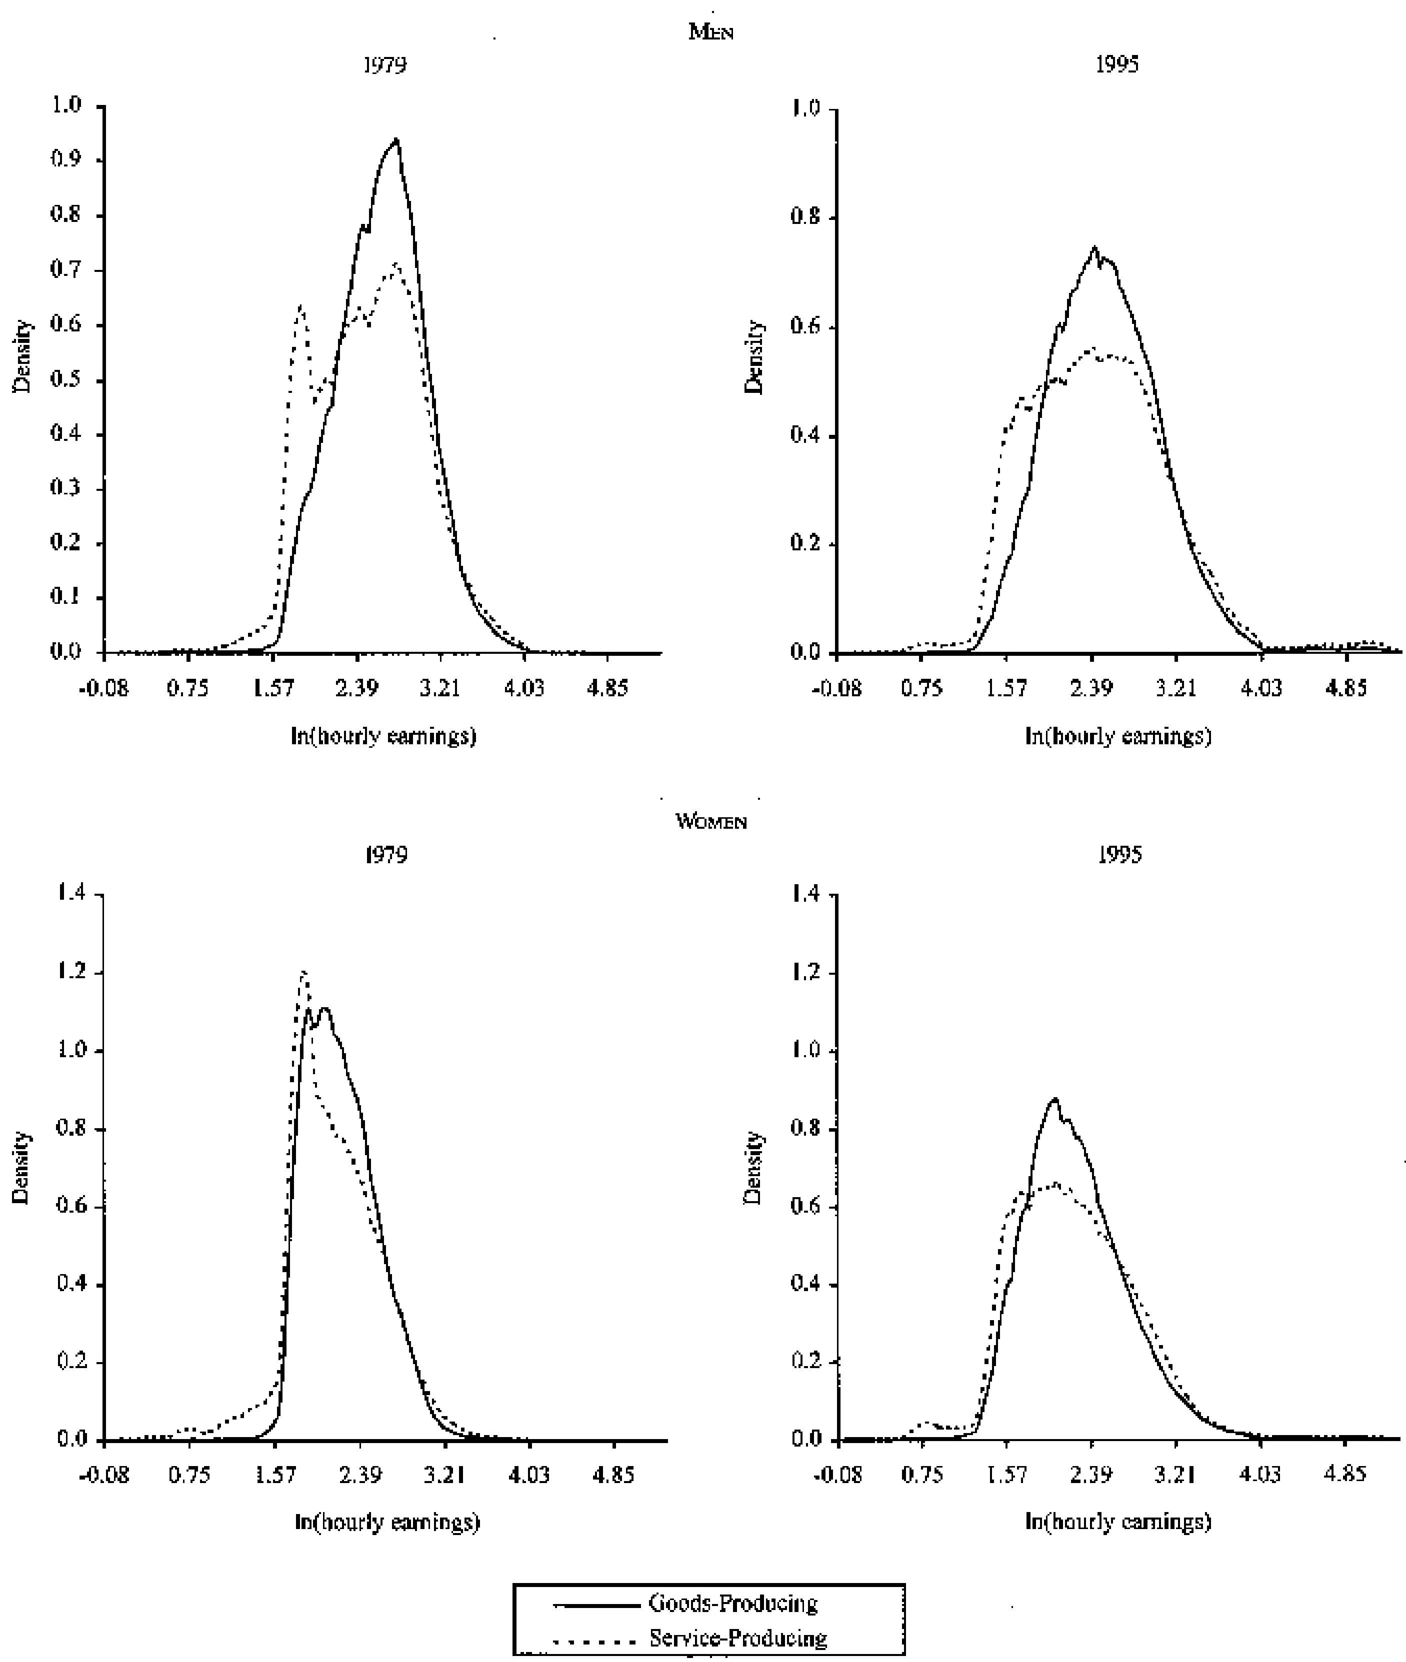

One difference between the two broad sectors that may help to explain the different earnings distributions is in the share of part-time jobs. Table 2 lists the mean and variance of earnings by broad sector and part-time status (and by sex). For men and women, mean earnings are lower, and the variance in earnings is higher, in part-time than in full-time jobs. Several changes occurred between 1979 and 1995 for both men and women. The most noticeable change is a large increase in the variance of earnings within all part-time categories listed; this increase dwarfs the increased

TABLE 2

MEAN AND STANDARD DEVIATION OF LN(HOURLY EARNINGS), BY PART TIME AND GOODS/SERVICES STATUS

A. MEN

	1979			1995		
INDUSTRY	MEAN	S.D.	SHARE[a]	MEAN	S.D.	SHARE[a]
TOTAL	2.55	.497	1.0	2.47	.641	1.0
Full Time	2.61	.464	.912	2.52	.547	.851
Part Time	2.01	.498	.088	2.23	.990	.149
GOODS-PRODUCING	2.62	.434	.430	2.52	.550	.339
Full Time	2.64	.421	.969	2.53	.498	.917
Part Time	2.12	.542	.031	2.46	.954	.083
SERVICE-PRODUCING	2.50	.534	.570	2.45	.681	.661
Full Time	2.58	.496	.870	2.51	.573	.818
Part Time	1.99	.487	.130	2.18	.990	.182
TOTAL OBSERVATIONS		74,671			83,931	

B. WOMEN

	1979			1995		
INDUSTRY	MEAN	S.D.	SHARE[a]	MEAN	S.D.	SHARE[a]
TOTAL	2.16	.441	1.0	2.23	.605	1.0
Full Time	2.23	.411	.724	2.31	.520	.700
Part Time	1.96	.457	.276	2.05	.738	.300
GOODS-PRODUCING	2.21	.351	.197	2.25	.519	.132
Full Time	2.22	.343	.915	2.26	.472	.877
Part Time	2.08	.415	.085	2.19	.780	.123
SERVICE-PRODUCING	2.14	.459	.803	2.18	.540	.868
Full Time	2.23	.431	.679	2.32	.528	.673
Part Time	1.95	.458	.321	2.05	.735	.327
TOTAL OBSERVATIONS		62,681			82,153	

Note: All tabulations are weighted by the CPS earnings weight, and 1979 earnings were inflated to 1995 levels using the GDP deflator for personal consumption expenditures.

[a] The full-time/part-time employment shares sum to 1 within each industry category (total, goods, services).

variance for full-time jobs. Mean earnings in part-time jobs also increased, particularly for men. Similarly, the part-time job share increased by 4 to 6 percentage points in most categories; the exception is service-producing women, for whom the share of part-time jobs remained essentially constant between 1979 and 1995. This latter fact suggests that any industry shift effects on inequality associated with increased part-time work will be greater for men than for women.

The following section presents conditionally weighted results. However, it is illustrative first to investigate the unconditional effect of the goods/services shift. This unconditional effect is obtained by upweighting the 1995 goods-producing sector observations by the relative goods share in 1979 versus 1995 (and downweighting the service sector observations by a similarly formed ratio for that sector). Figure 3 shows the impact on the male earnings distribution of this reweighting scheme, which does not account for any changes in the distribution of or returns to other earnings related characteristics. Relative to the actual distribution, the adjusted distribution has more weight around the median and less in the lower portion. This first pass at depicting the impact of the goods/services shift is consistent with the stereotypical view that the growing services share (as embodied in the solid "actual" line) is partially responsible for the erosion of the middle-class job base.

FIGURE 3

UNCONDITIONAL EFFECT OF GOODS/SERVICES SHIFT, MEN

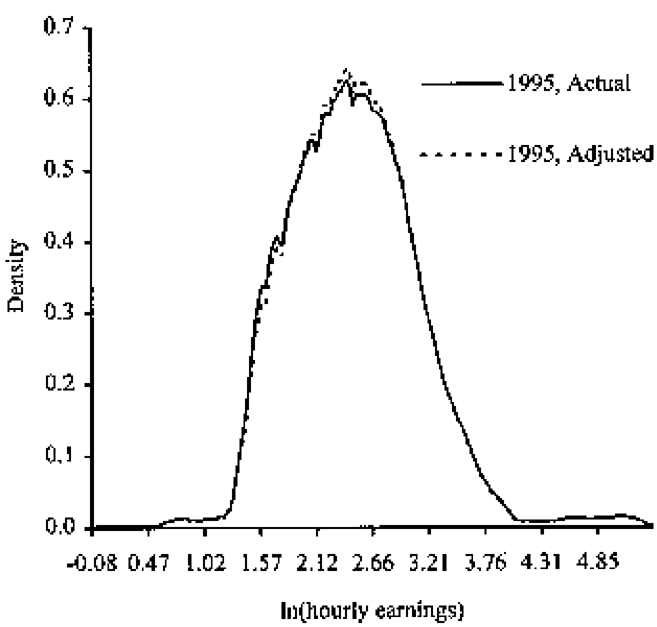

However, the corresponding figure for women (not shown) exhibits a much smaller unconditional impact of the goods/services shift.

Conditionally Weighted Density Estimates

The tabulations and densities in the previous section show the unconditional difference in the earnings distribution over time and across the goods-producing and service-producing sectors. It is likely, however, that the distribution of earnings-related characteristics differs across the two sectors, and that the earnings distributions conditional on these characteristics will differ less than the unconditional distributions. It is therefore important to condition on observables. To this end, I use a basic vector of X variables that includes a linear measure of educational attainment, potential experience and its square, two race dummies, three region dummies, and dummies for SMSA residence and marital status. Also, because the results in Table 2 suggest the potential importance of shifts between full-time and part-time jobs, I add a dummy for part-time work in additional analyses.

Figures 4 (men) and 5 (women) present the key results. For each panel, comparison of the solid line to the counterfactual dotted line shows the impact of the modeled change (industry structure or individual characteristics) as it actually evolved. Panel A in both figures shows the effect of accounting for the net shift from goods-producing to service-producing jobs between 1979 and 1995. This effect is estimated by reweighting the distribution through use of the conditioning weight $\psi_{E\div X}$. Panel B for each of these figures shows the effect of changing individual characteristics, which is estimated by use of the conditioning weight ψ_X.

For men, the adjusted distribution in Figure 4A reveals a small but discernible impact of industry employment shifts on the distribution of earnings. The adjusted distribution has slightly less mass in the lower portion and slightly more at or just above the middle; other portions of the adjusted and unadjusted distributions are nearly identical. This mass shift in the lower and middle portions is consistent with but smaller than the unconditional effect of the goods/services shift depicted in Figure 3. For women (Figure 5A), the conditional effect is barely discernible.

Figures 4B and 5B illustrate the impact of changing individual characteristics on the male and female earnings distributions. Their main impact for both men and women, as revealed by the comparison of the solid (actual) line to the dotted (adjusted) line, was to shift the distribution to the right. Also, the change in female characteristics and returns to them stretched the distribution somewhat from the median to the right.

FIGURE 4A

EFFECT OF GOODS/SERVICES SHIFT, MEN

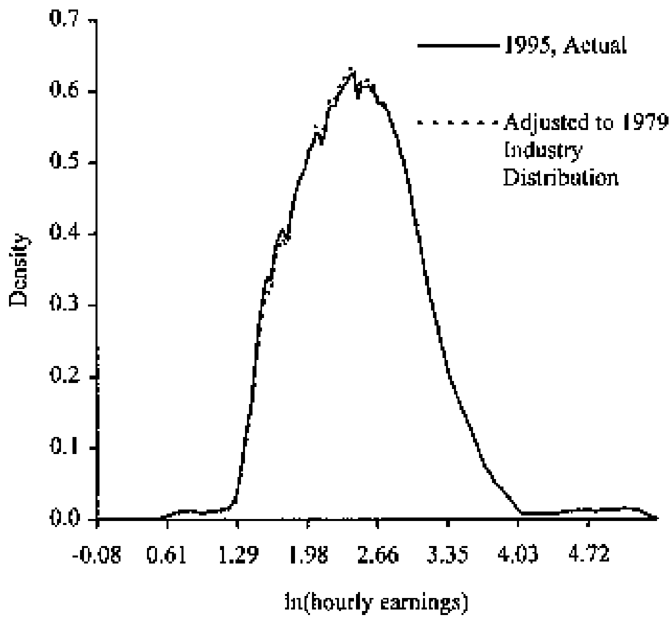

FIGURE 4B

EFFECT OF CHANGING *X*'S, MEN

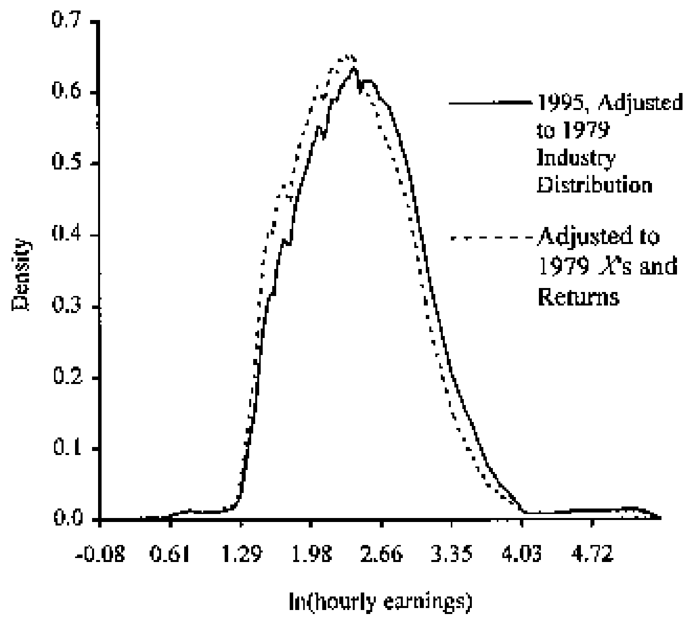

FIGURE 5A

EFFECT OF GOODS/SERVICES SHIFT, WOMEN

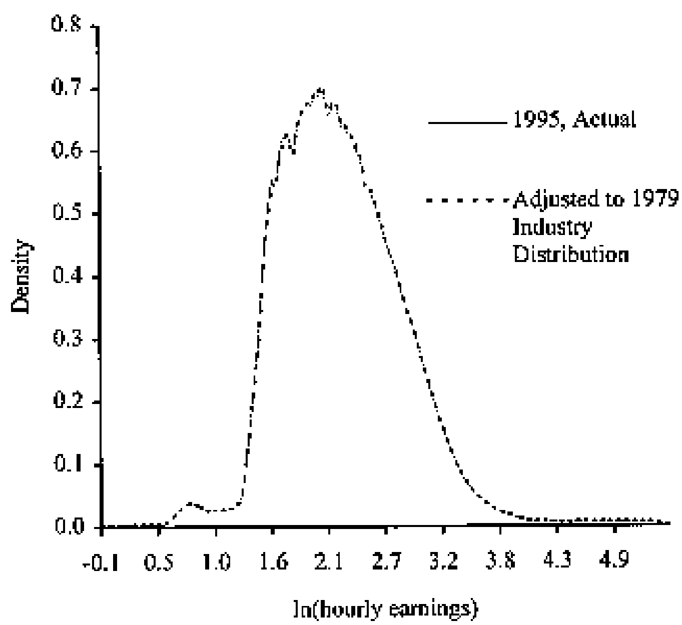

FIGURE 5B

EFFECT OF CHANGING *X*'S, WOMEN

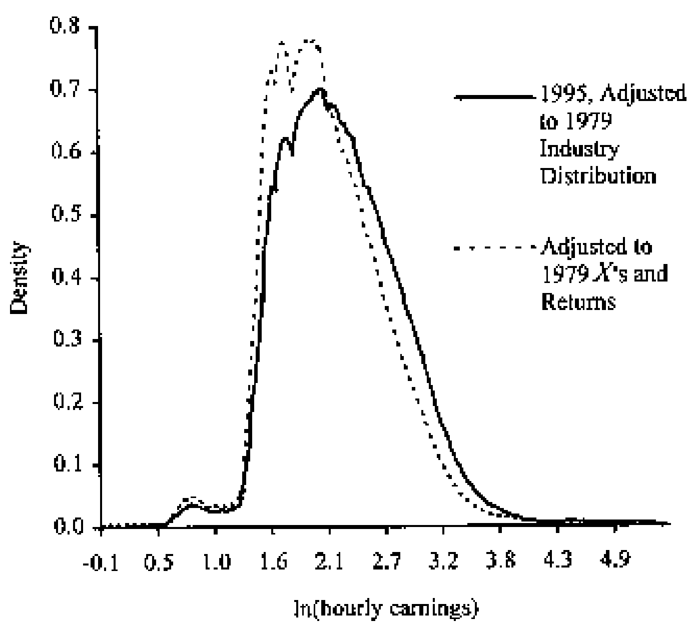

These figures provide a useful visual representation of the measured effects. As noted in Section I, however, the conditional weighted densities also can be used to assess the quantitative contribution of changing industry structure to changes in mean earnings and various dispersion measures. These results are reported in Table 3, for men and women separately. I analyze changes in the mean, standard deviation, and several quantile differences. I also analyze changes in two commonly used parametric inequality measures, the Gini coefficient and Theil's entropy measure.[7] For each measure, I list the total change in the measure between 1979 and 1995 and the amount explained by changing industry employment shares and changing individual characteristics.[8]

The results reported under the column "Goods vs. Services" in Table 3 show that broad industry employment shifts explain a small to moderate amount of the change in several earnings distribution measures for men. Broad industry shifts explain about 10% of the declining mean and 5% of the rising standard deviation. The largest impact on the changing quantile differences is for the 10–50 differential: the goods/services shift explains 43% of the increased dispersion in that range of the distribution. However, the lower portion of the male distribution changed far less than the upper portion; for example, the widening in the 10–50 differential is less than a quarter of the widening in the 50–90 differential. Among other measures, the goods/services shift also explains approximately 14% of the increase in the 10–90 differential. For women, the goods/services shift offset the increase in mean earnings somewhat, but had very little effect on the dispersion measures.

The final column of Table 3 shows the contribution of changing individual characteristics to the mean and dispersion measures. The impact of changing characteristics on mean male earnings was counterfactual. Also, although changing characteristics explain a substantial amount of the change in the 10–50 and 25–75 differentials, they explain very little of the increase in the other dispersion measures.

In contrast, for women the increase in mean earnings is more than fully explained by changing individual characteristics. These characteristics also explain a substantial portion of the change in various dispersion measures, including nearly 20% of the change in the standard deviation and the 10–50 differential and approximately 30% of the changes in the 5–95 and 25–75 differentials.

7. The 10–90 differential, for example, is defined as (ln(earnings at the 90th percentile) – ln(earnings at the 10th percentile)); the other quantile measures are defined similarly. See DFL for a definition of the Gini and Theil indices.

8. Unlike DFL, I do not provide a full decomposition of the change in each measure.

TABLE 3

Contribution of Changing Industry Shares and Individual Attributes to the Changing Earnings Distribution, 1979–1995

Statistic	Total Change[a]	Portion explained by: Goods vs. Services	Portion explained by: Individual Characteristics[b]
A. Men			
Mean	–.079	–.008 (.102)	.090 (–1.14)
Standard Deviation	.143	.008 (.052)	.014 (.095)
10–90[c]	.259	.037 (.143)	.010 (.038)
10–50	.047	.020 (.426)	.017 (.368)
50–90	.213	.017 (.081)	–.007 (–.034)
25–75	.178	–.001 (–.005)	.037 (.209)
5–95	.379	.026 (.069)	.032 (.083)
Gini	.109	.004 (.042)	.002 (.021)
Theil	.178	.007 (.041)	–.004 (–.024)
B. Women			
Mean	.076	–.005 (–.068)	.128 (1.68)
Standard Deviation	.164	.004 (.024)	.030 (.184)
10–90	.466	.030 (.065)	.057 (.123)
10–50	.288	.021 (.075)	.053 (.185)
50–90	.179	.009 (.050)	.004 (.022)
25–75	.212	–.002 (–.012)	.072 (.339)
5–95	.515	0 (0)	.151 (.292)
Gini	.121	.002 (.016)	.016 (.129)
Theil	.180	.003 (.015)	.017 (.094)

Note: Percentage of total change explained is shown in parentheses.

[a] Difference between statistic in 1979 and 1995.

[b] The individual attributes include a linear measure of educational attainment, potential experience and its square, two race dummies, three region dummies, and dummies for SMSA residence and marital status.

[c] This is defined as the change between 1979 and 1995 in (ln(earnings at the 90th percentile) – ln(earnings at the 10th percentile)). The other quantile measures are defined similarly.

The conditional results presented thus far are based on specifications that exclude any control for hours worked. However, as suggested by the tabulations presented in Table 2, there may be potentially important interactions between industry structure, earnings, and the share of part-time jobs. I therefore estimated additional models with a dummy variable for part-time work added to the list of individual characteristics. The results from this model are presented in Table 4.[9] The results in the second column reveal that inclusion of the part-time dummy in the conditioning equation reduces the estimated impact of the goods/services shift for men. Although the goods/services shift still explains about 26% of the increase in the 10–50 differential, the other estimated impacts are close to zero in percentage terms.

In contrast, inclusion of the part-time dummy substantially increases the share of the change in dispersion accounted for by individual characteristics. For men in Panel A, individual characteristics explain approximately 45% of the increase in the standard deviation and the Gini and Theil indices, and from 15% to more than 100% of the increase in the quantile dispersion measures. For women in Panel B, individual characteristics explain approximately 30% of the increase in the standard deviation and the Gini and Theil indices, and from 10 to 55% of the increase in the quantile dispersion measures.

I performed two primary checks of the robustness of these results. First, the results may be sensitive to the ordering of attribution imposed. Above, I assessed the contribution of the goods/services shift first, and the contribution of the *X*'s second; with this ordering, any interaction effects between the two are attributed to the goods/services shift. Therefore, I also performed the analysis in reverse order, letting the *X*'s affect the distribution first. This did not noticeably change the results for women. For men, however, this order reversal largely eliminated the impact of the goods/services shift on the 10–50 differential. Also, although the order reversal substantially increased the effect on the 25–75 differential in the model that excludes the part-time dummy, it did not do so in the model that includes the part-time dummy. Thus, it appears that in regard to their impact on male earnings inequality, there are important interaction effects between individual characteristics—particularly working part time—and the probability of working in goods versus services.

Another objection to the basic framework is that it does not account for changes in the general structure of the economy between 1979 and 1995. One way to assess how im-

9. I do not report corresponding kernel density estimates in additional figures, because in this model the actual density and density adjusted for industry shifts are indistinguishable.

TABLE 4

CONTRIBUTION OF CHANGING INDUSTRY SHARES AND INDIVIDUAL ATTRIBUTES TO THE CHANGING EARNINGS DISTRIBUTION, 1979–1995, PART TIME DUMMY ADDED

		Portion explained by:	
Statistic	Total Change[a]	Goods vs. Services	Individual Characteristics[b]
A. Men			
Mean	–.079	–.002 (.031)	.094 (–1.18)
Standard Deviation	.143	.002 (.015)	.066 (.460)
10–90[c]	.259	.019 (.072)	.086 (.331)
10–50	.047	.012 (.257)	.055 (1.17)
50–90	.213	.007 (.032)	.031 (.147)
25–75	.178	–.004 (–.021)	.040 (.224)
5–95	.379	.002 (.005)	.177 (.466)
Gini	.109	.001 (.011)	.046 (.423)
Theil	.178	.002 (.012)	.083 (.467)
B. Women			
Mean	.076	–.001 (–.020)	.118 (1.55)
Standard Deviation	.164	.002 (.013)	.053 (.322)
10–90	.466	.010 (.022)	.099 (.212)
10–50	.288	.007 (.024)	.079 (.274)
50–90	.179	.003 (.019)	.020 (.111)
25–75	.212	.000 (.002)	.115 (.540)
5–95	.515	0 (0)	.174 (.338)
Gini	.121	.001 (.009)	.034 (.285)
Theil	.180	.002 (.010)	.051 (.284)

Note: Percentage of total change explained is shown in parentheses.

[a] Difference between statistic in 1979 and 1995.

[b] The individual attributes include a linear measure of educational attainment, potential experience and its square, two race dummies, three region dummies, and dummies for SMSA residence, marital status, and whether worked part-time.

[c] This is defined as the change between 1979 and 1995 in (ln(earnings at the 90th percentile) – ln(earnings at the 10th percentile)). The other quantile measures are defined similarly.

portant such changes might be is to reverse the temporal ordering of the analysis: i.e., rather than imposing the 1979 industry and characteristics structure on the 1995 distribution of earnings, impose the 1995 structure on the 1979 distribution of earnings. Again, the results differ across the models that include or exclude the part-time dummy. In the model that excludes it, the results are very similar to those using the original temporal ordering. In the model that includes the part-time dummy, however, the estimated goods/services shift impact on the 10–50 differential is largely eliminated but replaced by a comparable impact on the 25–75 differential.

Overall, the estimated small effect of the goods/services shift on earnings dispersion in the lower half of the male distribution seems sensitive to the treatment of part-time work in the model. This finding, combined with the absence of an effect for women, suggests that the increase in part-time work by men in the services industry, as identified in Table 2, played a key role in any existing industry shift effects on earnings inequality. Furthermore, the most important measured characteristic in these models is part-time work. Inclusion of the part-time dummy in the model increases the share of increased dispersion explained by individual characteristics substantially, to nearly one-half for men and nearly one-third for women.

V. Conclusions

In this paper, I investigated the extent to which a substantial net shift from goods-producing to service-producing jobs altered the U.S. distribution of individual earnings between 1979 and 1995. Relative to previous work in this area, my paper's primary contribution is to apply recently developed semi-parametric estimation techniques to the problem. The analyses revealed four key empirical findings:

(1) Average earnings are lower and the dispersion of earnings is higher in service-producing than in goods-producing jobs.

(2) Consistent with (1), the unconditional effect of the shift from goods-producing to service-producing jobs has been to increase dispersion in the lower half of the earnings distribution.

(3) When individual characteristics are introduced into the model, a smaller but detectable impact on the lower half of the male earnings distribution remains. The quantitative impact was to increase the 10–50 earnings differential by several percentage points, nearly half the total change. However, this change is small relative to the large changes that occurred in upper half of the male earnings distribution. No similar effect was found for women.

(4) Result (3) is sensitive to controlling for part-time work. Although the estimated impact of the goods/services shift on the 10–50 differential largely withstands inclusion of a part-time dummy, additional checks revealed that this result is not fully robust to reversing the ordering of attribution or temporal ordering in the model.

The results from this analysis provide at best weak evidence in support of the view that the shift from goods-producing to service-producing jobs made an independent contribution to the erosion of middle-class earnings in the U.S. To the extent that an effect was isolated, its largest contribution was in the lower portion of the male distribution, which is consistent with the stereotype that shrinkage of the manufacturing sector has helped to erode the middle-class job base. However, this effect appears largely due to increased incidence of part-time work by men, particularly in the service-producing sector, which exhibited a sharp increase in earnings variance in part-time jobs. To the extent that increased part-time work by men was voluntary, this trend has limited adverse implications. However, if this trend reflects demand-side constraints, it may bode poorly for men stuck in part-time jobs. Furthermore, the increased incidence of part-time work appears to have made a large contribution to growing inequality for both men and women. The exact contribution of part-time work to growing inequality merits further investigation.

Overall, my results are much closer to those of authors (such as Juhn, Murphy, and Pierce 1990) who find no industry shift effects on earnings inequality than to those of authors (such as Maxwell 1989, 1990) who find large industry shift effects on earnings inequality. However, one key drawback of my approach is its broad measure of industry sectors (goods-producing vs. service-producing). It might be interesting to incorporate a finer industry breakdown into the analysis, although this extension may be problematic for the conditional weighted density estimation framework. In the meantime, additional applications of the conditionally weighted approach, as developed in DFL, appear warranted. This procedure is relatively easy to implement, and it is very powerful in regard to uncovering distributional changes and in its ability to perform additional tests on the altered distributions.

APPENDIX

Derivation of the Conditioning Weights

This appendix provides the derivation of the conditioning weights, $\psi_{E|X}$ and ψ_X, described heuristically in Section II. This discussion largely follows that in DiNardo, Fortin, and Lemieux (1996; DFL), although they provide a more complete and therefore more complex decomposition of changing earnings inequality.

Using the notation in the text, consider the distribution of wages w in year t, conditional on a vector of individual characteristics X and a dummy variable (E) indicating whether the worker is in a goods-producing sector job:

$$\text{(A1)} \qquad f_t(w) \equiv f(w; t_w = t, t_{E|X} = t, t_X = t).$$

DFL show that a distribution such as (A1) can be expressed as:

$$\text{(A2)} \qquad f_t(w) = \iint f(w|E, X, t_w = t)\, dF(E|X, t_{E|X} = t)\, dF(X|t_X = t).$$

In this equation, $f(.)$ is the conditional distribution of w and $F(.)$ is the joint distribution of w, E, and X. The right-hand side of (A2) indicates that the distribution of earnings in a given year can be expressed as the conditional distribution multiplied by the marginals (the first of which also is conditional) and integrated over E and X.

We are interested in the distribution of w if the distribution of E conditional on X is held to its 1979 structure:

$$\text{(A3)} \qquad f(w; t_w = 95, t_{E|X} = 79, t_X = 95).$$

Using (A2), this distribution can be expressed as:

$$\text{(A4)} \qquad f_t(w; t_w = 95, t_{E|X} = 79, t_X = 95)$$

$$= \iint f(w|E, X, t_w = 95)\, dF(E|X, t_{E|X} = 79)\, dF(X|t_X = 95)$$

$$= \iint f(w|E, X, t_w = 95)\, \psi_{E|X}(E, X)\, dF(E|X, t_{E|X} = 95)\, dF(X|t_X = 95),$$

where $\psi_{E|X}(E,X)$ is a reweighting function to be defined momentarily. It is very important to note that except for $\psi_{E|X}$, (A4) is identical to (A2) with $t = 95$— i.e., the distribution that we are interested in is equal to the unconditional distribution of earnings in 1995, with observations reweighted by the function $\psi_{E|X}$. If we can estimate $\psi_{E|X}$, it is straightforward to incorporate it and obtain the counterfactual distribution expressed in (A4) by using the observed univariate, unconditional distribution of wages in 1995.

The reweighting function is defined (identically) as:

$$\text{(A5)} \quad \psi_{E|X}(E, X) \equiv \frac{dF(E|X, t_{E|X} = 79)}{dF(E|X, t_{E|X} = 95)}$$

$$= E \left(\frac{\Pr(E = 1|X, t_{E|X} = 79)}{\Pr(E = 1|X, t_{E|X} = 95)} \right)$$

$$+ (1 - E) \left(\frac{\Pr(E = 0|X, t_{E|X} = 79)}{\Pr(E = 0|X, t_{E|X} = 95)} \right).$$

The first line identity in (A5) is obtained by substituting the expression on the right side into (A4) and canceling out the denominator. The second line is derived by noting that E only takes the values 0 or 1, so that:

$$\text{(A6)} \qquad dF(E|X, t_{E|X} = t) \equiv E \Pr(E = 1|X, t_{E|X} = t) + (1 - E) \Pr(E = 0|X, t_{E|X} = t).$$

The second equality in (A5) follows from the recognition that one term in this expression will always equal zero.

This weight represents the change in the probability between 1979 and 1995 that an observation defined by characteristics X is in the goods-producing or service-producing sector. The probabilities in (A6) are easily recognized as expressions from standard binary dependent variable models. These conditional probabilities can be obtained by estimating a model such as a probit or logit and then using the fitted values. I use the probit equation

$$\text{(A7)} \qquad \Pr(E = 1|X, t_{E|X} = t) = \Pr(\varepsilon > -H(X)\beta)$$
$$= 1 - \Phi(-H(X)\beta)$$

to obtain the structure of $E|X$ at time t, where ε is a normally distributed random variable. In (A7), $H(X)$ is a vector function of X designed to capture the conditional relationship being modeled, and β is a vector of estimated coefficients. This equation is estimated for both the 1979 and 1995 samples, and the coefficients are retained. We use these results to fit the probabilities in (A5) using the 1995 sample X's, combined with the 1979 coefficients for the numerator and the 1995 coefficients for the denominator. The resulting estimated weights, $\hat{\psi}_{E|X}$, are incorporated into the kernel density estimation or into the tabulation of distributional statistics, as described in Section II.

A modification of this procedure enables us to account for the impact of changes in the X vector of earnings-related characteristics. In this case, the weighting function is obtained through a simple application of Bayes' Law:

(A8) $$\psi_X(X) = \frac{\Pr(t_X = 95)}{\Pr(t_X = 79)} \frac{\Pr(t_X = 79|X)}{\Pr(t_X = 95|X)}.$$

This function represents the relative probability of observing an individual with characteristics X in the 1979 versus the 1995 sample, normalized by the unconditional probabilities of being in either sample. As long as the distribution of X's changed between the two years (for example, through higher average educational attainment), the weights ψ_X will alter the estimated distribution.

The function ψ_x is estimated by pooling the 1979 and 1995 data sets, and then estimating a binary dependent variable model for a dummy variable indicating the sample from which the observation is obtained. The conditional probabilities $\Pr(t_X = t|X)$ are obtained by forming fitted probabilities for workers in the 1995 sample, based on their X values. The unconditional probabilities, $\Pr(t_X = t)$, are simply the weighted shares of the 1979 and 1995 samples in the pooled sample. Estimation is then performed on the 1995 sample, with the estimated weights $\hat{\psi}_X$ modifying the sampling weights (as described in Section II).

Two points should be noted. First, the conditional probability estimating equation (A7) has no behavioral interpretation; it simply permits conditioning out the effect of covariates (X) which may be related to the factor (industry employment shifts) whose effect we are attempting to estimate. Second, due to potentially important interactions between the effects of industry shifts and changing individual characteristics, I also estimated models that reverse the order of attribution, by first assessing the contribution of the X's, and then assessing the contribution of E. The exact procedure is described in DFL.

REFERENCES

Blackburn, McKinley L. 1990. "What Can Explain the Increase in Earnings Inequality Among Males?" *Industrial Relations* 29 (3) pp. 441–456.

__________, David E. Bloom, and Richard B. Freeman. 1990. "The Declining Economic Position of Less Skilled American Men." In *A Future of Lousy Jobs,* ed. Gary Burtless, pp. 31–67. Washington, D.C.: Brookings.

Bluestone, Barry, and Bennett Harrison. 1988. "The Growth of Low-Wage Employment: 1963–86." *American Economic Review* 78 (2) pp. 124–128.

__________, and __________. 1982. *The Deindustrialization of America.* New York: Basic Books.

Bound, John, and George Johnson. 1992. "Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations." *American Economic Review* 82 (3) pp. 371–392.

Burkhauser, Richard V., Amy D. Crews, Mary C. Daly, and Stephen P. Jenkins. 1996. *Income Mobility and the Middle Class.* Washington, D.C.: American Enterprise Institute.

Burtless, Gary. 1990. "Earnings Inequality over the Business and Demographic Cycles." In *A Future of Lousy Jobs,* ed. Gary Burtless, pp. 77–116. Washington, D.C.: Brookings.

Delgado, Miguel A., and Peter M. Robinson. 1992. "Nonparametric and Semiparametric Methods for Economic Research." *Journal of Economic Surveys* 6 (3) pp. 201–250.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach." *Econometrica* 64 (5) pp. 1001–1044.

__________, and Thomas Lemieux. 1994. "Diverging Male Wage Inequality in the United States and Canada, 1981–88: Do Unions Explain the Difference?" Irvine Economic Paper No. 93–94-16 (June). University of California, Irvine.

Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill." *Journal of Political Economy* 101 (3) pp. 410–442.

Karoly, Lynn. 1996. "Anatomy of the U.S. Income Distribution: Two Decades of Change." *Oxford Review of Economic Policy* 12 (1) pp. 77–96.

__________. 1992. "Changes in the Distribution of Individual Earnings in the United States: 1967–86." *Review of Economics and Statistics* 74 (1) pp. 107–115.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–87: Supply and Demand Factors." *Quarterly Journal of Economics* 107 (1) pp. 35–78.

Levy, Frank, and Richard J. Murnane. 1992. "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations." *Journal of Economic Literature* 30 (September) pp. 1333–1381.

Maxwell, Nan. 1990. *Income Inequality in the United States,1947–85.* Westport, CT: Greenwood Press.

__________. 1989. "Demographic and Economic Determinants of United States Income Inequality." *Social Science Quarterly* 70 (2) pp. 245–264.

Murphy, Kevin M., and Finis Welch. 1993. "Industrial Change and the Rising Importance of Skill." In *Uneven Tides:Rising Inequality in America,* eds. Sheldon Danziger and Peter Gottschalk. New York: Russell Sage Foundation.

Schweitzer, Mark E., and Max Dupuy. 1995. "Sectoral Wage Convergence: A Nonparametric Distributional Analysis." Working Paper 95-20 (December). Federal Reserve Bank of Cleveland.

Sheather, S.J., and M.C. Jones. 1991. "A Reliable Data-based Bandwidth Selection Method for Kernel Density Estimation." *Journal of the Royal Statistical Society* B 53 (3) pp. 683–690.

Silverman, B.W. 1986. *Density Estimation for Statistics and Data Analysis.* London: Chapman and Hall.