LOUSY AND LOVELY JOBS: THE RISING POLARIZATION OF WORK IN BRITAIN

Maarten Goos and Alan Manning*

Abstract—This paper shows that the United Kingdom since 1975 has exhibited a pattern of job polarization with rises in employment shares in the highest- and lowest-wage occupations. This is not entirely consistent with the idea of skill-biased technical change as a hypothesis about the impact of technology on the labor market. We argue that the "routinization" hypothesis recently proposed by Autor, Levy, and Murnane (2003) is a better explanation of job polarization, though other factors may also be important. We show that job polarization can explain one-third of the rise in the log(50/10) wage differential and one-half of the rise in the log(90/50).

I. Introduction

ECONOMISTS writing about the impact of technology on the labor market in recent years have tended to emphasize the role played by skill-biased technical change (SBTC), the idea that technology is biased in favor of skilled workers and against unskilled workers. The idea of SBTC has primarily been used to explain rising wage inequality (see Katz and Autor (1999) for a survey of a very large literature). But a recent paper by Autor, Levy, and Murnane (ALM) (2003) has argued for a more nuanced way of understanding the impact of technology in general (and computers in particular) on the labor market.[1] They argue persuasively that technology can replace human labor in routine tasks—tasks that can be expressed in step-by-step procedures or rules—but (as yet) cannot replace human labor in nonroutine tasks.[2]

The ALM hypothesis is intuitively plausible and they provide evidence that industries in which routine skills were heavily used have seen the most adoption of computers, and this has reduced the usage of routine skills in those industries (see Spitz (2006) for similar evidence for Germany). But, if the ALM routinization hypothesis is correct, then we might expect to see evidence for it in other areas: this is the aim of this paper.

The basic idea is the following. The SBTC hypothesis predicts that demand for "skilled" jobs is rising relative to that for "unskilled" jobs, while the ALM hypothesis suggests a more subtle impact of technology on the demand for labor of different skills. The routine tasks in which technology can substitute for human labor include jobs like craft manual jobs and bookkeeping jobs that require precision and, hence, were never the least-paid jobs in the labor market. The nonroutine tasks which are complementary to technology include skilled professional and managerial jobs that tend to be in the upper part of the wage distribution. The nonroutine manual tasks that make up many of the most unskilled jobs such as cleaning are not directly affected by technology, but the impact of technology in other parts of the economy is likely to lead to a rise in employment in these unskilled jobs. If this is true, then the impact of technology will be to lead to rising relative demand in well-paid skilled jobs (that typically require nonroutine cognitive skills) and in low-paid least-skilled jobs (that typically require nonroutine manual skills) and falling relative demand in the "middling" jobs that have typically required routine manual and cognitive skills—a process we call job polarization. This paper documents that the pattern of employment changes in Britain over the period 1975–1999 does show job polarization.

A literature related to the idea of job polarization already exists—the "job-quality" debate in the United States. Some of the early papers on the rise in U.S. wage inequality (for example, Bluestone & Harrison, 1988) argued that there was an increasing number of low-wage jobs and a shrinking number of middling jobs. This argument was controversial even at the time (for example, see Kosters & Ross, 1988) and most labor economists came to the conclusion that the problem for low-skill workers was a declining number of jobs for them rather than an increasing number (see Burtless, 1990). But, one can still find a number of papers from the 1990s continuing to address the major themes of the job-quality debate (see, for example, Costrell, 1990; Howell & Wolff, 1991; Levy & Murnane, 1992; Juhn, Murphy, & Pierce, 1993; Murphy & Welch, 1993; Gittleman & Howell, 1995; Ilg, 1996; Farber, 1997; Acemoglu, 1999, 2001; Juhn, 1999; Ilg & Haugen, 2000; Wright & Dwyer, 2003). Although these studies do differ slightly in their conclusions, common themes do emerge, most notably that, in the last 30 years, there has been a very big increase in the number of high-paid jobs and (probably) an increase in the number of low-paid service jobs—this is broadly consistent with the job polarization prediction of the ALM hypothesis, although few of these papers offer this interpretation.

Received for publication January 22, 2004. Revision accepted for publication January 23, 2006.

* Department of Economics, Catholic University Leuven and Erasmus University, and Centre for Economic Performance, London School of Economics, respectively.

We would like to thank David Autor, Francis Green, Frank Levy, Coen Teulings, and two referees for comments on this paper, and David Autor and Francis Green for making their data available.

[1] See also Card and DiNardo (2002) for the argument that SBTC is not as successful in explaining wage inequality as commonly thought.

[2] The idea of jobs as a set of tasks that differ in how easily machines can displace human labor is not new and goes back to at least Herbert Simon (1960). Unlike more expansive computer scientists of that time, Simon had a clear sense of what computers would and would not be able to do, and his predictions are broadly consistent with both the ALM hypothesis and the ideas advanced in this paper. Simon starts from the prediction that complex information-processing programs will supplant labor in routine jobs intense in many relatively simple and repetitive eye-brain-hand sequences. Consequently, workers will sort into nonroutine jobs requiring the flexible use of the brain, eyes, hands, and legs.

The Review of Economics and Statistics, February 2007, 89(1): 118–133

The plan of the paper is as follows. In section II, we use the U.S. data from Autor, Levy, and Murnane (2003) to show that the jobs that require nonroutine tasks tend to be at the top and bottom of the wage distribution, while the jobs that require routine tasks tend to be in the middle, thus leading to the job polarization prediction. Section III describes the data used for the United Kingdom. Section IV then documents how job polarization can be observed in the United Kingdom between 1975 and 1999 when the quality of jobs is defined by their median wage. There has been a growth in lousy jobs (mainly in low-paying service occupations) together with a (much larger) growth in lovely jobs (mainly in professional and managerial occupations in finance and business services) and a decline in the number of middling jobs (mainly clerical jobs and skilled manual jobs in manufacturing). We document that one sees these trends using all measures of employment, for men and women together or separately and for all definitions of "jobs" that we use. We also show that a method used by Juhn, Murphy, and Pierce (1993) and Juhn (1999) to predict employment growth at each percentile of the wage distribution also supports the hypothesis of job polarization. And although the pattern of changes in the occupational structure of employment is broadly consistent with the ALM hypothesis, other factors may be important and section V considers some of them. We discuss the potential importance of changes in the composition of the labor force (e.g., from the rising labor market participation of women, the changing age and education structure), the structure of consumer demand, and trade. It is likely that all of these factors are important for employment changes in at least some occupations, but none of these hypotheses seem able to explain the broad sweep of job polarization.

As an increase in the relative demand for low-wage workers (relative to middling workers) is not in line with the predictions of the SBTC hypothesis, sections VI and VII consider the evidence most commonly cited in favor of that hypothesis. Section VI considers the rise in the employment of nonmanual workers. We argue that the pattern of within- and between-industry changes in employment observed at the one-digit occupation level is consistent with the ALM hypothesis that technical progress has displaced the labor of clerical and manual workers in all sectors of the economy, but that differential productivity growth between manufacturing and service sectors has led to the growth in low-wage service employment (as originally proposed by Baumol, 1967). Section VII documents that the well-known shift toward more educated labor has largely occurred within jobs and that there has been a rapid rise in educational attainment of workers even in the worst jobs. There are two possible interpretations of this. First, that there has been SBTC within jobs as we define them so that the consensus view on the importance of SBTC is correct. Secondly, that as the educational attainment of all groups in the population has risen but the job distribution has become more polarized, some educated workers are forced into the low-skill jobs at the bottom end of the distribution. The attraction of this view is that it can explain why there has been a simultaneous rise in the returns to education (the demand for educated workers has increased as the number of good jobs has increased) and in the level of overeducation as some have claimed. Distinguishing between these hypotheses requires evidence on changing skill requirements within jobs that is hard to find. We review two pieces of evidence that might shed light on these questions, although they are somewhat contradictory in their implications.

Section VIII considers the extent to which the observed job polarization can explain the rise in wage inequality between the 1970s and 1990s. We find that a modest part of the rise in wage inequality can be explained by the polarization of jobs alone, but that once one includes the fact that wage growth seems to be monotonically positively related to the quality of jobs, one can explain most of the evolution of wage inequality. The implication is that the rise in "within-group" wage inequality that others have emphasized is more a product of a restricted definition of a "group," and that if one includes job controls then it largely disappears. However, the finding that the wages in the lousy jobs are falling relative to those in the middling jobs presents something of a problem for the ALM hypothesis, as one might expect the opposite if relative demand is rising in the lousy relative to the middling jobs. The final section concludes.

II. Routine Jobs, Nonroutine Jobs, and Technical Change

This section shows how the Autor, Levy, and Murnane (2003) view of the impact of technology on the demand for different skills predicts job polarization. ALM use the U.S. Dictionary of Occupational Titles (DOT) to associate particular occupations with the intensity of use of five particular types of tasks. The types of tasks included in the analysis are chosen to represent those that are affected in different ways by technology—they label them nonroutine cognitive, nonroutine interactive, routine cognitive, routine manual, and nonroutine manual (see Autor, Levy, and Murnane (2003) for a more-detailed description of the tasks given these labels).

ALM then show that industries that were relatively intensive users of occupations that use routine tasks had more computerization and that the extent of the use of routine skills has fallen in these industries. Here we pursue an angle of the ALM hypothesis that ALM do not develop—namely, that jobs that can be routinized are not distributed uniformly across the wage distribution. The central idea is that nonroutine manual jobs are concentrated in the lower percentiles of the wage distribution, whereas nonroutine cognitive and interactive jobs are concentrated in the top end of the wage range, with routine jobs concentrated in the middle.

TABLE 1.—DOT TASK DENSITY BY WAGE PERCENTILES

DOT Task Measure	Mean DOT Task Measure	Fraction of Workers above Mean DOT Task Measure by Wage Percentiles		
		≤ 33	33–66	≥ 66
Nonroutine cognitive	3.755	0.17	0.48	0.88
Nonroutine interactive	2.417	0.03	0.14	0.59
Routine cognitive	4.582	0.37	0.63	0.43
Routine manual	3.901	0.28	0.58	0.35
Nonroutine manual	1.198	0.49	0.33	0.31

Notes: Task inputs are measured as in ALM (2003) and are between zero and ten. The mean DOT task measure is the 1977 mean across three-digit occupations. Wage percentiles are taken from the 1983 CPR MORG (Current Population Survey Merged Outgoing Rotation Group) file.

ALM argue that the nonroutine cognitive and interactive tasks are complementary to technology, the routine tasks are substitutes, and the nonroutine manual tasks are not directly affected. However this should not be taken to mean there will be no effects of technology on employment in occupations that primarily consist of nonroutine manual tasks. The reason is the general equilibrium effect first identified by Baumol (1967)—employment will shift toward jobs in which productivity growth is low (because technology is not applied there) in order to keep the balance of output in different products. Baumol applied his argument to the shift in employment from manufacturing to services, but it is relevant in the current context as well. As a result, technological progress can be expected to result in job polarization, with employment growth in lovely and lousy jobs and employment falls in middling jobs.

Table 1 presents a simple way of showing that the nonroutine jobs are concentrated at the top and bottom of the wage distribution. We use wage information from the 1983 CPS MORG (Current Population Survey Merged Outgoing Rotation Group) file and assign to each individual the five task measures in 1977 used by ALM based on their occupations.[3] All skills are measured on a ten-point scale, although these should not be taken to be comparable across tasks. Table 1 tabulates the fraction of workers that have DOT scores above the overall mean DOT score for the five different tasks as a percentage of total employment within the three terciles of the wage distribution. For example, only 17% of all workers in the lowest-paid occupations are in jobs that require above-average nonroutine cognitive skills. But 88% percent of workers in the highest-paid occupations are in jobs that require above-average nonroutine cognitive skills. A similar picture holds for the nonroutine interactive skills: occupations intensive in nonroutine interactive skills are concentrated in the upper part of the wage distribution. In contrast, routine-intensive occupations are concentrated in the middle. Of workers in occupations earning between the 33rd and 66th wage percentiles, 63% require above-average routine cognitive and 58% above-average routine manual skills. These numbers are higher than for any other specified wage range. Finally, the lowest-paid occupations require a higher fraction of nonroutine manual skills and its fraction is higher than for any other occupation paying higher wages.

This section has shown some direct evidence that workers in the middling jobs used to do routine tasks, while workers in lousy and lovely jobs did nonroutine tasks. Since nonroutine jobs are concentrated in both tails of the wage distribution, the ALM hypothesis predicts an increasing polarization of the workforce into lousy and lovely jobs. This predicted process of polarization provides an explanation for the empirical "facts" in an ongoing debate about the quality of jobs mentioned in the introduction.

III. The Data

The main data in this paper comes from Britain, but we would expect the task composition of occupations and the impact of technology to be very similar to that observed in the United States. The data used in this paper come from two sources, the New Earnings Survey (NES) and the Labor Force Survey (LFS). The New Earnings Survey is an annual panel data set that started in 1968, though the first year for which computerized records are available is 1975, the sample being all individuals whose National Insurance number ends in 14. In April of each year the tax records are used to contact the employer of each of these workers; the employer reports information on pay, hours, and, importantly for this paper, occupation and industry. Although the NES is in theory a random sample, it is known to undersample certain groups in practice, notably part-time workers (if weekly earnings fall below the threshold for paying National Insurance, then they are unlikely to appear in the tax records) and those who have changed jobs recently (as the sampling frame is drawn up early in the year and the survey is likely to be sent to the wrong employer in April).

For this reason we supplement the NES with data from the Labor Force Survey. The LFS was first conducted in 1975, then every two years until 1983, then annually until 1992, and quarterly since then (when a panel component was also introduced). The LFS has a much smaller sample than the NES (and until 1993 it did not contain any wage data) but does have the advantage that it is closer to a random sample.

In this paper we define a job as a particular occupation or as a particular occupation in a particular industry. The

[3] We are grateful to David Autor for making the DOT data available to us. The year 1983 is the earliest year for which the DOT occupations can be merged into CPS data.

FIGURE 1.—PERCENTAGE CHANGE IN EMPLOYMENT SHARE BY JOB-QUALITY DECILE

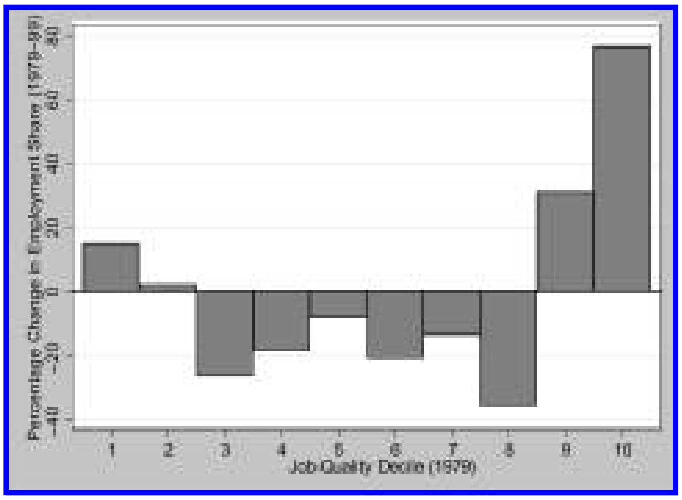

Notes: Employment data are taken from the LFS using three-digit SOC90 codes. Employment changes are taken between 1979 and 1999. Quality deciles are based on three-digit SOC90 median wages in 1979 taken from the NES.

FIGURE 2.—EMPLOYMENT GROWTH BY JOB MEDIAN WAGE

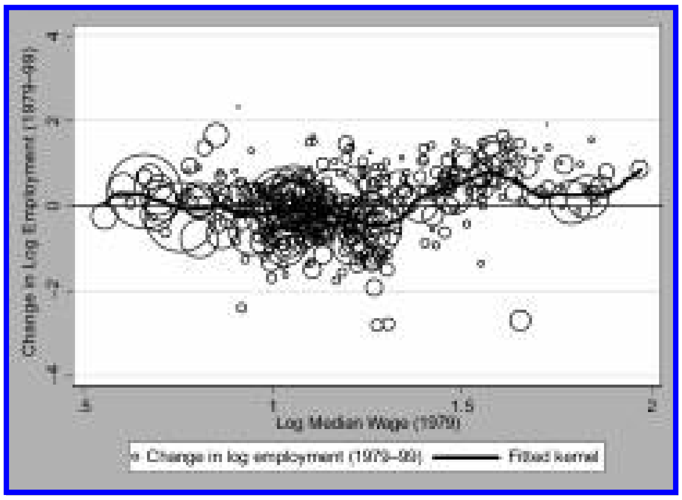

Notes: Employment data are taken from the LFS using three-digit SOC90 codes. Employment changes are taken between 1979 and 1999. Wages are three-digit SOC90 median wages in 1979 taken from the NES.

occupation part corresponds to the main usage of the term *job*—the question in the LFS used to obtain the information on occupation is, "What was your main job in the week ending Sunday?" The industry part of the definition of a job is more problematic, but other papers in this area have used a similar definition and there are significant industry effects on wages even once one has controlled for occupation. However, it is important to realize that the occupation part of our definition is much more important than the industry part, as one gets very similar results whether a job is defined by occupation alone or by an occupation-industry interaction.

We have explored using different levels of disaggregation by occupation and industry, and the results seem robust to the level chosen. We have restricted the results reported in this paper to using three-digit occupation codes only (allowing for approximately 370 jobs) as well as the interaction of a three-digit occupation and one-digit industry classification (allowing for a maximum of 3,700 jobs, although in practice only about 1,600 exist because not all occupations are represented in all industries). One might wonder about whether some jobs that are disappearing and new ones are popping up. In practice this does not seem to be a problem: of the occupations that existed in the 1970s, all still have workers in them in the late 1990s, and there are essentially no new occupations that cannot be put into the 1970s classification.

IV. Trends in the Quality of Jobs

We start by looking at long-term trends in the quality of jobs. To do this obviously requires a measure of the quality of a job. We first do this in a very simple way by using the median wage in the job at the beginning of the period (see OECD (2001) and Meisenheimer (1998) for a discussion of other ways of discussing the quality of jobs). One can think of it as a single-index model of skill—see Card and Lemieux (1996). However, we then also take a slightly different approach based on the analysis of Juhn, Murphy, and Pierce (1993), with very similar results.

First consider how the proportional change in employment from the late 1970s to the late 1990s is related to the initial level of wages. If the SBTC hypothesis is correct, then one would expect to see a monotonic positive relationship between employment growth and initial wages. Figure 1 groups occupations into the lowest 10%, the second-lowest 10%, up to the top 10% based on their median wage and cell size in 1979. For example, the worst job-quality decile captures 10% of all workers employed in the lowest-paid occupations. Figure 1 shows large growth in the share of employment in the top two deciles, but also growth, albeit smaller, in the share of jobs in the bottom decile. Also, there has been a significant decline in middling jobs. Though the increase in the number of workers with bad jobs has been lower than the increase in the number of workers with good jobs, employment polarization into low-paid and high-paid work is clear from figure 1. It is this process of job polarization that is the central theme of this paper.

Figure 2 presents the LFS data for the period 1979–1999, where the size of the circles denotes the initial employment level in each occupation. On figure 2, we also include a kernel regression estimate of the mean of employment growth conditional on job quality.[4] There is certainly no striking evidence of a positive monotonic relationship between employment growth and initial log median wages as the literature on SBTC might have led one to expect. Moreover, one can discern the J-shaped relationship that is going to appear in the regression results.

[4] These are Nadaraya-Watson estimates, using a bandwidth of 0.1 and an Epanechnikov kernel.

TABLE 2A.—THE RELATIONSHIP BETWEEN EMPLOYMENT GROWTH AND INITIAL MEDIAN WAGE: MEN AND WOMEN TOGETHER

Sample	Sample Period	Data	Employment Measure	β_1	β_2	Fraction in Declining Section
Men + Women	1979–99	LFS (occ)	Employment	−4.541 (0.700)	2.107 (0.297)	52.93
Men + Women	1976–95	NES (occ)	Employment	−3.412 (0.664)	1.373 (0.267)	72.57
Men + Women	1979–99	LFS (occXind)	Employment	−4.804 (0.472)	2.109 (0.198)	62.80
Men + Women	1976–95	NES (occXind)	Employment	−3.957 (0.378)	1.581 (0.151)	74.69
Men + Women	1979–99	LFS (occ)	Hours	−4.218 (0.785)	2.047 (0.327)	28.42
Men + Women	1976–95	NES (occ)	Hours	−3.603 (0.775)	1.576 (0.319)	56.85
Men + Women	1979–99	LFS (occXind)	Hours	−4.331 (0.514)	1.969 (0.213)	49.67
Men + Women	1976–95	NES (occXind)	Hours	−4.145 (0.435)	1.748 (0.178)	62.22

Notes: Regressions are weighted by job cell size in the initial period. Occupation uses three-digit SOC90 codes. Industry uses one-digit SIC80 codes.

Figures 1 and 2 relate to one measure of employment, one definition of a job, and one survey (the LFS). One would like to know whether the results are robust or not. Because it is tedious to present graphs for every possible outcome, we turn to a simple regression to summarize our results.

A. *Regression Estimates*

The models we estimate are of the quadratic form:

$$\Delta n_j = \beta_0 + \beta_1 w_{j0} + \beta_2 w_{j0}^2, \tag{1}$$

where Δn_j is the change in log employment in job *j,* and w_{j0} is the initial log median wage in the job.

We experiment with a number of different measures of employment and jobs. Table 2A presents estimates combining employment for men and women. The top half of table 2A measures employment in terms of bodies, using different definitions for a job and different surveys: we report results from the LFS and the NES using either three-digit occupation codes only or the interaction of a three-digit occupation code with a one-digit industry code. But the results tell a similar story. The linear term in equation (1) is negative and the quadratic term positive, implying a U-shaped relationship between employment growth and the initial level of wages. One might be concerned that the downward-sloping part of this relationship contains no data points but, as the final column in table 2A makes clear, this is not the case: substantial numbers of workers are in the downward-sloping part of the relationship.[5] These regressions support the view that there has been polarization in the quality of jobs, with the employment growth being at the extreme ends of the distribution. It should also be noted that the parameter estimates for the LFS and NES are very similar, which suggests that the nonrandom sampling in the NES is not too serious a problem. We have also experimented with further aggregation or disaggregation in the jobs classification, but this does not seem to make a great deal of difference to the qualitative results.

One might think that these results are misleading because much of the growth in employment has been in part-time jobs and these tend to be low paid. Hence, the estimates in the top half of table 2A might be thought to overstate the employment growth in low-paid occupations. However, when we measure employment in terms of total hours, the results are very similar, so this does not explain away the observed job polarization. One might also think that the feminization of employment can explain this job polarization, with women accounting for the growth in relatively low-paid occupations. But, as tables 2B and 2C show, one observes similar patterns for male and female employment considered separately although, in fact, the trends are more marked for men.[6]

One might also be concerned that the quality ranking of jobs changes a lot over time so that the patterns of employment growth are sensitive to the point in time at which job quality is measured. Table 3 shows this is not the case: one gets similar results if one uses median wages at the end of the 1990s to rank jobs. This reflects the well-known fact that there is very considerable stability in the occupational wage structure.

B. *An Alternative Approach: Juhn, Murphy, and Pierce*

So far we have defined the quality of a job by the median wage in that job. Although this approach has the virtue of simplicity in that it enables us to label specific jobs as good or bad, it does ignore the fact that there is substantial wage dispersion within jobs. One approach to dealing with this issue is taken by Juhn, Murphy, and Pierce (1993)—al-

[5] Inspection of the kernel regression line in Figure 2 should make it clear that this estimate of the proportion in the downward-sloping section is not an artifact of the quadratic specification adopted.

[6] Note that the last column in Table 2C shows a missing if the quadratic term is not significantly different from zero.

TABLE 2B.—THE RELATIONSHIP BETWEEN EMPLOYMENT GROWTH AND INITIAL MEDIAN WAGE: MEN

Sample	Sample Period	Data	Employment Measure	β_1	β_2	Fraction in Declining Section
Men	1979–99	LFS (occ)	Employment	−5.807 (1.317)	2.447 (0.482)	39.66
Men	1976–95	NES (occ)	Employment	−3.080 (1.097)	1.267 (0.389)	43.33
Men	1979–99	LFS (occXind)	Employment	−6.039 (0.719)	2.413 (0.265)	55.84
Men	1976–95	NES (occXind)	Employment	−4.697 (0.535)	1.783 (0.191)	68.91
Men	1979–99	LFS (occ)	Hours	−5.022 (1.361)	2.246 (0.502)	27.98
Men	1976–95	NES (occ)	Hours	−4.732 (1.266)	1.981 (0.463)	39.10
Men	1979–99	LFS (occXind)	Hours	−5.622 (0.755)	2.337 (0.281)	45.48
Men	1976–95	NES (occXind)	Hours	−5.906 (0.618)	2.309 (0.226)	64.32

Notes: Regressions are weighted by job cell size in the initial period. Occupation uses three-digit SOC90 codes. Industry uses one-digit SIC80 codes.

TABLE 2C.—THE RELATIONSHIP BETWEEN EMPLOYMENT GROWTH AND INITIAL MEDIAN WAGE: WOMEN

Sample	Sample Period	Data	Employment Measure	β_1	β_2	Fraction in Declining Section
Women	1979–99	LFS (occ)	Employment	−1.580 (1.025)	1.222 (0.505)	—
Women	1976–95	NES (occ)	Employment	−0.657 (0.686)	0.584 (0.310)	—
Women	1979–99	LFS (occXind)	Employment	−3.363 (0.840)	1.942 (0.411)	54.69
Women	1976–95	NES (occXind)	Employment	−2.227 (0.517)	1.256 (0.239)	50.95
Women	1979–99	LFS (occ)	Hours	−1.441 (1.177)	1.415 (0.597)	—
Women	1976–95	NES (occ)	Hours	−0.776 (0.815)	0.887 (0.401)	—
Women	1979–99	LFS (occXind)	Hours	−3.199 (0.934)	2.034 (0.466)	34.17
Women	1976–95	NES (occXind)	Hours	−2.650 (0.618)	1.659 (0.306)	29.58

Notes: Regressions are weighted by job cell size in the initial period. Occupation uses three-digit SOC90 codes. Industry uses one-digit SIC80 codes.

TABLE 3.—THE RELATIONSHIP BETWEEN EMPLOYMENT GROWTH AND TERMINAL MEDIAN WAGE

Sample	Sample Period	Data	Employment Measure	β_1	β_2	Fraction in Declining Section
Men + Women	1979–99	LFS (occ)	Employment	−1.915 (0.491)	0.839 (0.166)	29.59
Men + Women	1976–95	NES (occ)	Employment	−2.920 (0.387)	1.090 (0.127)	54.96
Men + Women	1979–99	LFS (occXind)	Employment	−0.581 (0.289)	0.403 (0.097)	—
Men + Women	1976–95	NES (occXind)	Employment	−2.416 (0.231)	0.915 (0.076)	50.36
Men + Women	1979–99	LFS (occ)	Hours	−1.651 (0.519)	0.806 (0.171)	17.22
Men + Women	1976–95	NES (occ)	Hours	−2.770 (0.433)	1.101 (0.140)	38.97
Men + Women	1979–99	LFS (occXind)	Hours	−0.271 (0.286)	0.356 (0.094)	—
Men + Women	1976–95	NES (occXind)	Hours	−2.506 (0.264)	1.003 (0.085)	37.21

Notes: Regressions are weighted by job cell size in the terminal period. Occupation uses three-digit SOC90 codes. Industry uses one-digit SIC80 codes.

FIGURE 3.—THE IMPACT OF JOB POLARIZATION ON EMPLOYMENT GROWTH BY WAGE PERCENTILE

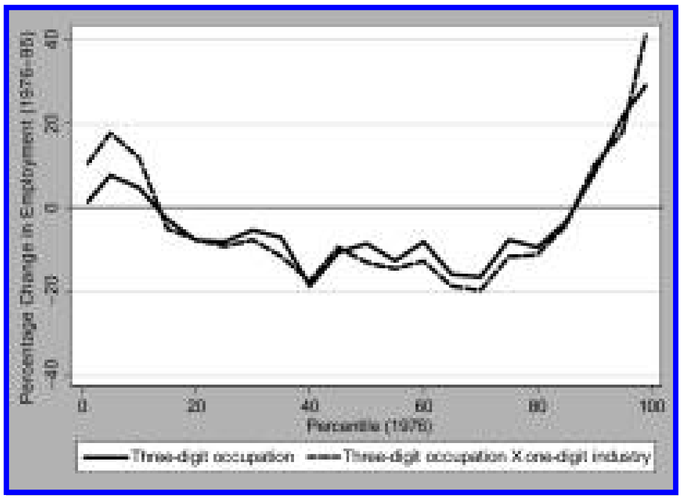

Notes: Data are taken from the NES using three-digit SOC90 codes. Employment changes are taken between 1976 and 1995. Percentiles are the 1976 wage density percentiles.

though that paper is better known for other contributions—and Juhn (1999).

They assume that each job (defined here as an occupation) potentially uses labor from each percentile of the wage distribution. They compute the share of labor from each percentile used in each job in a base year. Assuming that these "factor shares" remain constant, one can then predict changes in employment for each percentile of the wage distribution by allowing changes in the total levels of employment in each occupation. Note that now the same job potentially contributes to the predicted change in employment for each percentile rather than contributing only once as in the single-index approach taken above.

Figure 3 plots these predicted employment changes between 1976 and 1995 for different percentiles of the 1976 wage distribution. As figure 3 shows, employment growth is positive for all workers earning less than the 11th percentile and more than the 86th percentile. Predicted growth at the top end is strongest, between 35% and 45%. Growth at the 5th percentile is between 8% and almost 20% whereas employment in the middling jobs is in decline. The conclusions derived are therefore the same as those derived from our more simplistic approach in the previous section that there has been increased polarization in the quality of jobs. It is noteworthy that only in the top three deciles does one see evidence of the positive relationship between skill and employment change as predicted by SBTC.

C. *Employment Growth by Occupation*

What are the sorts of jobs that are growing and declining? Table 4 presents a "top ten" by job growth for occupations that have cells of a respectable size using the LFS data.[7] The first column specifies the occupation. Since growth in the best jobs has been stronger than employment growth in the bad jobs, most of the jobs reported in table 4 pay above median hourly wages, as can be seen from the second column in the table. Columns (3) and (4) report estimated employment levels by occupation, and the final column calculates the percentage change in employment between 1979 and 1999.

Most of the top ten rapidly growing jobs are specialized occupations mainly in finance and business service indus-

[7] There are some dangers in doing this because the occupations at the extremes of the employment change distribution are quite likely to be ones for which a number of factors reinforce each other.

TABLE 4.—TOP TEN OCCUPATIONS BY JOB GROWTH

Occupation	Median Wage in 1979	Employment in 1979	Employment in 1999	Percent Change in Employment
All	3.052	24,332,613	27,343,467	12.373
Care assistants and attendants	2.345	103,837	539,407	419.474
Software engineers	5.008	34,009	171,769	405.065
Management consultants and business analysts	4.745	18,811	81,803	334.868
Computer systems and data processing managers	5.065	43,239	178,701	313.286
Computer analysts and programmers	4.842	76,083	302,617	297.745
Educational assistants	2.272	45,040	173,763	285.793
Hospital ward assistants	2.572	7,460	26,986	261.705
Actors, entertainers, stage managers, and producers	4.719	22,549	73,030	223.870
Treasurers and company financial managers	5.105	37,794	119,812	217.015
Financial institution and office managers	4.511	107,138	322,608	201.114

Notes: Employment data are taken from the LFS using three-digit SOC90 codes. Wages are 1979 median hourly wages taken from the NES using three-digit SOC90 codes.

TABLE 5.—BOTTOM TEN OCCUPATIONS BY MEDIAN WAGE

Occupation	Median Wage in 1979	Employment in 1979	Employment in 1999	Percent Change in Employment
All	3.052	24,332,613	27,343,467	12.373
Hairdressers and barbers	1.745	123,986	96,073	−22.513
Bar staff	1.832	119,455	188,319	57.647
Shelf fillers	1.938	49,699	97,144	95.462
Sales assistants	1.939	954,200	1,321,251	38.466
Retail cash desk and checkout operators	1.969	112,816	218,581	93.749
Petrol pump forecourt attendants	1.979	13,304	9,935	−25.321
Kitchen porters	2.003	178,758	143,092	−19.952
Waiters and waitresses	2.020	124,780	187,391	50.177
Cleaners	2.132	854,535	649,362	−24.009
Beauticians	2.145	24,536	28,946	17.972

Notes: Employment data are taken from the LFS using three-digit SOC90 codes. Wages are 1979 median hourly wages taken from the NES using three-digit SOC90 codes.

tries located at the top end of the wage distribution. But positions 1, 6, and 7 in the top ten, however, are taken by low-paid jobs—care, education, and hospital assistants. And, just outside the top ten one finds large increases in the number of hotel porters, merchandisers, window dressers, and travel and flight attendants, among other low-paid service occupations, that are intense in nonroutine manual tasks.

To document this, table 5 lists the ten lowest-paying jobs given they are of considerable size, their median wage, and employment in 1979 and 1999. The biggest absolute increase in those jobs listed has been for sales assistant and checkout operators. Given the emphasis in the literature on SBTC, the presence of the good jobs in table 4 is probably no surprise, but strong growth in many bad jobs in table 5 might be more surprising. However, this pattern is exactly what we would expect to see according to the ALM hypothesis, as the rapidly growing lousy jobs are all ones where it has proved difficult to substitute machines or computers for human labor. To see further evidence supportive of the ALM hypothesis, table 6 lists the bottom ten jobs by job growth. A comparison of the median job wages with the overall median suggests the decline in jobs has been largest for middling jobs in manufacturing occupations.

So far we have presented evidence for job polarization as an important phenomenon in the United Kingdom over the past 25 years and suggested that the pattern of employment changes is broadly consistent with the ALM view of the

TABLE 6.—BOTTOM TEN OCCUPATIONS BY JOB GROWTH

Occupation	Median Wage in 1979	Employment in 1979	Employment in 1999	Percent Change in Employment
All	3.052	24,332,613	27,343,467	12.373
Boring and drilling machine setters and setter operators	3.584	29,276	1,731	−94.086
Coal mine laborers	3.696	29,782	1,818	−93.892
Face trained coal-mining workers, shotfirers, and deputies	5.237	76,301	5,095	−93.322
Grinding machine setters and operators	3.557	56,426	8,164	−85.531
Laborers in foundries	3.219	14,801	2,505	−83.070
Laborers in engineering and allied trades	3.025	58,243	12,758	−78.095
Electrical, energy, boiler, and related plant operatives and attendants	3.684	36,352	8,009	−77.968
Spinners, doublers, and twisters (in textiles and tannery process operatives)	2.802	16,941	4,173	−75.363
Originators, compositors, and print preparers (in printing and related trades)	3.404	48,878	12,162	−75.116
Rail signal operatives and crossing keepers	3.010	13,761	3,571	−74.045

Notes: Employment data are taken from the LFS using three-digit SOC90 codes. Wages are 1979 median hourly wages taken from the NES using three-digit SOC90 codes.

FIGURE 4.—THE IMPACT OF CHANGES IN THE COMPOSITION OF THE LABOR FORCE

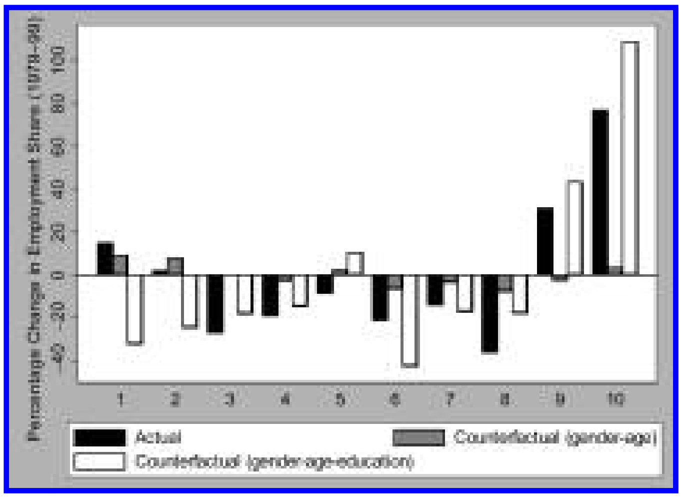

Notes: Employment data are taken from the LFS using sing three-digit SOC90 codes. Employment changes are taken between 1979 and 1999. Quality deciles are based on three-digit SOC90 median wages in 1979 taken from the NES. The first bar gives actual employment changes as in figure 1. The second bar gives counterfactual employment changes keeping the three-digit occupational composition within 24 gender-age cells constant over time. The third bar gives counterfactual employment changes keeping the three-digit occupational composition within 96 gender-age-education cells constant over time.

impact of technology on the demand for labor rather than the simple SBTC hypothesis. But, job polarization could be driven by factors other than technology—either changes in the supply of labor or changes in demand for reasons other than technology. The next section provides a discussion of these issues.

V. Alternative Hypotheses for Job Polarization

A. *Changes in Labor Supply*

It is possible that changes in the structure of the labor force can explain some of the changes in the occupational structure of employment. The most important such changes are the increased feminization of the labor force and the increase in educational attainment, though changes in the age structure and the proportion of immigrants could also conceivably be important. To assess the importance of these changes we did the following counterfactual exercise. First, we divided the labor force into cells (described in more detail below). Then keeping the initial occupational structure of employment within cells constant, we computed what the change in occupational employment would have been if the only change was the changing relative size of the cells in the overall labor force.

Figure 4 shows the results for the changes divided by deciles. The first column in each decile shows the actual change in employment in each decile (this is the same as figure 1). The second column then shows the predicted change in employment in each decile when the labor force is divided into two gender and twelve age cells.[8] The predicted changes in the occupational structure are small compared to the actual. There is a small predicted rise in employment in the lousy jobs that is primarily caused by the increasing proportion of women in the labor force who are concentrated in the lowest-wage jobs. However, this counterfactual takes no account of the substantial occupational upgrading of women over this period so is likely to overstate the true contribution of the feminization of employment to job polarization.

The obvious omission from the above counterfactual is education: there has been very substantial educational upgrading over this period (for example, the fraction of workers that have education at "A" level or above—at least twelve years of education—increased from 0.25 in 1979 to 0.55 in 1999). The final column for each decile in figure 4 shows what happens if one divides the labor force into four educational categories as well as the gender and age categories used above. The counterfactual predicts rapid growth in employment in the lovely jobs but also predicts large falls in the lousy jobs for the simple reason that these jobs used to employ large numbers of low-educated workers that have been a rapidly declining share of the labor force. Hence, while dividing the labor force by education can help to explain the growth in lovely jobs, it is unhelpful in explaining the growth in lousy jobs. The conclusion must be that changes in the structure of labor supply are unable to explain the broad pattern of job polarization.

B. *Changes in Labor Demand Other than Technology*

Changes in the occupational structure of employment may also be caused by changes in the demand for different sorts of labor that are not caused by technology. Here we briefly discuss two of these: trade and the structure of product demand.

The role of international trade and outsourcing has been a perennial alternative hypothesis to technology as a potential explanation of changes in wage inequality. It undoubtedly has been important for some occupations (for example, the large decline in spinners, doublers, and twisters seen in table 6 is the result of a continuing shift of textiles to countries where labor is cheaper). And trade may be more important in the future, such as with the outsourcing of more skilled jobs. But the overall assessment of Freeman (2003) is that trade has had much smaller impacts on labor markets than commonly believed.[9] We are not going to investigate this in detail here, but one would have to oppose this

[8] We do not include immigrants as a separate category, as the fraction of foreign-born workers in the U.K. labor force only rose from 7.3% in 1979 to 8.9% in 1999, making it a rather unimportant factor.

[9] This assessment is consistent with many other studies. For example, Feenstra and Hanson (1998) estimate that expenditures on high-technology capital such as computers are about twice as important as outsourcing in explaining variation in relative wages of nonproduction workers in the United States between 1979 and 1990. And, Borjas, Freeman, and Katz (1997) find that immigration has had a larger impact

TABLE 7.—SHIFT-SHARE ANALYSIS OF EMPLOYMENT SHARES BY OCCUPATION

Occupation	Wage	NES Total	NES Within	NES Between	LFS Total	LFS Within	LFS Between
Professional occupations *(NM)*	5.914	1.709	1.127	0.582	3.733	2.838	0.895
Managers and administrators *(NM)*	4.117	5.204	4.588	0.616	5.606	5.271	0.335
Associate professional and technical occupations *(NM)*	3.823	2.579	1.700	0.879	4.466	3.446	1.020
Craft and related occupations *(M)*	3.277	−8.158	−3.738	−4.420	−7.883	−3.461	−4.422
Plant and machine operatives *(M)*	3.055	−5.579	−1.809	−3.770	−5.195	−1.362	−3.833
Clerical and secretarial occupations *(NM)*	2.841	1.291	−1.879	3.171	−2.105	−5.388	3.283
Personal and protective service occupations *(NM/M)*	2.668	3.516	1.969	1.547	3.502	1.732	1.770
Other occupations *(M)*	2.558	−2.527	−2.775	0.248	−3.398	−3.564	0.166
Sales occupations *(NM)*	2.132	1.964	0.817	1.147	1.272	0.487	0.785

Notes: Employment changes are taken between 1979 and 1999 for the LFS and between 1976 and 1995 for the NES. Reported wages are 1979 median hourly wages taken from the NES using one-digit SOC90 occupations. The decomposition is done using one-digit SIC80 industry codes.

conclusion to argue that trade was the most important factor behind job polarization. And trade and technology may not be competing explanations—it is quite plausible that the jobs that can be routinized are the ones that are most likely to be shifted abroad.

Potentially more important are changes in the structure of the demand for different products that then have consequences for the demand for different occupations. For example, the rise in the number of care and hospital assistants seen in table 4 is partly the result of more old people, and more old people being cared for outside the family. But it is important to realize that technology also plays a very important role here, as the increase in the demand for care has not been met by any great improvements in the productivity of caring because of the difficulty in applying technology to nonroutine tasks like caring.[10] Also, excluding care and hospital assistants from the regressions presented in table 2 does not explain away the increasing importance of low-wage work. This reflects that the dramatic decline in many jobs in manufacturing seen in table 6 is mostly the result of relatively inelastic consumer demand together with the rapid productivity growth in those sectors where products are produced in ways that have proved relatively easy to routinize.

In sum, none of the other hypotheses considered here seem to have the ability to explain the basic feature of job polarization, though they are undoubtedly important for some specific occupations. In contrast, the ALM hypothesis does seem to have this broad explanatory power. But, before we uncritically accept the ALM hypothesis, we need to understand the evidence that is often quoted in support of SBTC. The next two sections consider two of these—the growth in nonmanual employment and the rise in the educational attainment of the workforce.

on the skill composition than trade in the United States between 1980 and 1990.

[10] For example the *Economist* of March 13, 2004, quoted the inventor of the first industrial robot (Unimate, employed by General Motors in 1961), Joe Engelberger, as saying that care of the elderly is the opportunity the robotics industry should be pursuing as "every highly industrialized nation has a paucity of help for vast, fast-growing ageing populations."

VI. Understanding the Growth in Nonmanual Employment

A number of papers (for example, Berman, Bound, & Griliches, 1994; Berman, Bound, & Machin, 1998; Machin & Van Reenen, 1998) have presented evidence that employment has shifted toward nonmanual jobs and that this shift has been much more important within than between manufacturing industries. The fact that nonmanual jobs tend to be better paid than manual jobs is interpreted as evidence that technical change is biased toward more skilled workers. And the fact that most of the shifts are within industries suggests that this trend is related to technical change that has very pervasive effects on all sectors of the economy. As a statement about the average quality of jobs, this conclusion is undoubtedly right: our data also suggest that the "average" job quality is increasing. But the binary distinction between manual and nonmanual is simply not able to capture the increased polarization we have also argued is important.[11]

If the shift-share analysis is done for broader occupation groups and for the whole economy, not just manufacturing, we get the results presented in table 7. For each occupation, table 7 reports a manual/nonmanual indicator taken from the LFS (*M* or *NM,* respectively). Occupations are ranked by their median wage. Then, for each of the two data sets, the first column reflects the total percentage point change in the share of each occupation group between 1979 and 1999. The second column measures the percentage point change due to changes within industries, whereas the final column reports the change due to workers moving between industries.

The results are rather more nuanced than earlier studies would suggest and in line with the ALM hypothesis. There is a large increase in the employment shares of managerial and professional workers, an increase in lovely jobs that is mostly within industries. Both craft workers and machine

[11] Indeed, any binary distinction between low- and high-skill workers (whether in theoretical or empirical work) cannot have separate mean and variance effects.

FIGURE 5.—CHANGE IN FRACTION OF WORKERS BY EDUCATION AND JOB MEDIAN WAGE

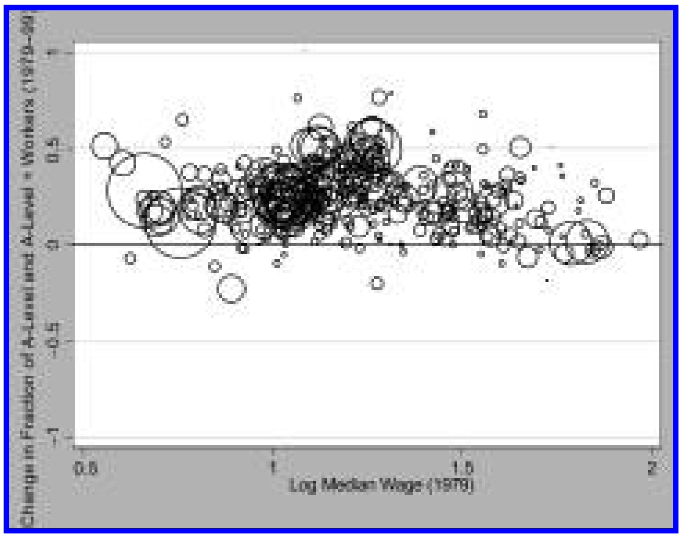

Notes: Employment data are taken from the LFS using three-digit SOC90 codes.

operatives have large negative within and between components, reflecting both the impact of technical change and the shift toward services. Routine clerical occupations have large negative employment effects within industries with a sizeable positive between component reflecting the shift to services. The increase in the employment share of low-paid personal and protective services and sales occupations has a large within and between component, reflecting the fact that technology has not managed to do these jobs and reflecting the shift toward services.

Therefore studies that use a simple manual/nonmanual split (usually out of necessity rather than choice) and concentrate on manufacturing miss important features of the way the structure of employment is evolving. If one broadens one's view, then one does see evidence for the ALM hypothesis.

VII. Education and Occupation

Another piece of widely cited evidence in favor of the SBTC hypothesis is that there has been a rapid increase in the level of educational attainment together with a rise in the returns to education. It is true that there is a lot of evidence that the average educational attainment of workers within jobs has changed. The evidence in figure 4 that the assumption of a fixed occupational structure within education groups predicts a fall in lousy jobs when there has been a rise implies a rise in the educational attainment of workers in low-wage occupations. A more direct way of seeing this is figure 5, which shows the change in the fraction of workers that have education at "A" level or above—at least twelve years of education—for each occupation. Almost all occupations show an increase, evidence of educational upgrading within occupations as the changes are above the horizontal axis.[12]

There are two interpretations of these findings. First, that what we have defined as a job is not constant over time and the educational and/or skill requirements within jobs have risen, possibly because of SBTC within jobs. Secondly, that as the educational attainment of the labor force has increased and middling jobs become relatively scarcer, some educated workers have been forced to take lousier jobs than previously. This is the idea of the literature on overqualification (see, for example, Sicherman, (1991); Hartog, (2000); and, for the United Kingdom, Green, McIntosh, & Vignoles, (1999); Chevalier, (2000); and Green & McIntosh, (2002)), which typically finds that high proportions of people report that they are employed in jobs for which their educational qualifications are unnecessary. Employers may also respond by raising the minimum educational standards to get certain jobs—what is known as credentialism. To distinguish these two hypotheses requires some information on changes in skill use within occupations. This information is not so easy to find, but we present two pieces of disparate information relevant to the question.

First, consider the data on the use of the five DOT measures used by Autor, Levy, and Murnane (2003). Table 8 presents data on the average level of skill use in 1977, the change from 1977 to 1991, and the decomposition of this change into a within-occupation and a between-occupation component. Panel A of Table 8 pools all occupations together and shows an overall increase in nonroutine cognitive and interactive tasks, a decrease in routine tasks (especially cognitive ones), and a smaller decrease in nonroutine manual tasks. But, the decomposition suggests that, within occupations, there is a rise only in the nonroutine interactive tasks and all other skills show declines, the decline being particularly large for the routine cognitive tasks. But from the point of view of educational upgrading, it is what is happening in the lousy jobs that is perhaps of more interest. Panel B therefore does the same exercise for jobs in the bottom half of the wage distribution. Again one sees a big rise in the nonroutine interactive tasks and large declines in routine tasks. But, most of the increase in skill requirements is between occupations: within-occupation task requirements are generally falling. There is little evidence here that there is substantial SBTC within occupations (Spenner (1983) reaches similar conclusions).

Our second piece of evidence on changing skill requirements within occupations comes from the U.K. Social Change and Economic Life Initiative (SCELI) survey conducted in 1986 and the 2001 Skills Surveys.[13] Both of these surveys asked workers about the educational qualifications

[12] The rise in this proportion is smallest in some of the high-wage jobs, but this is because the proportion of educated workers in these jobs was already close to one in the 1970s, leaving little scope for educational upgrading using "A" levels as the cutoff.

[13] We are grateful to Francis Green for doing these computations for us.

TABLE 8.—DOT TASK SHIFTS WITHIN AND BETWEEN OCCUPATIONS

DOT Task Measure	Mean 1977	Change 1977–1991		
		Total	Within	Between
Panel A: All Occupations				
Nonroutine cognitive	3.755	0.084	−0.047	0.131
Nonroutine interactive	2.417	0.504	0.137	0.367
Routine cognitive	4.582	−0.854	−0.564	−0.290
Routine manual	3.901	−0.146	−0.025	−0.121
Nonroutine manual	1.198	−0.132	−0.094	−0.038
Panel B: Low-wage Occupations				
Nonroutine cognitive	3.338	−0.027	−0.106	0.079
Nonroutine interactive	2.169	0.367	0.019	0.348
Routine cognitive	3.929	−1.116	−0.871	−0.245
Routine manual	3.879	−0.224	−0.065	−0.159
Nonroutine manual	0.847	−0.037	−0.032	−0.005

Notes: Task inputs are measured as in ALM (2003) and are between zero and ten. For panel A, the reported means are weighted using 463 three-digit COC occupations. Panel B uses 208 occupations with hourly earnings below overall average wages using 1984 CPS data. Changes between 1977 and 1991 are measured using three-digit COC occupations and employment changes between 1984 and 1997.

necessary to get the job they do and whether these qualifications were necessary to do the job. Only data at the one-digit occupation level are comparable in the two data sets. Table 9 presents some relevant information. The second column gives the change from 1986 to 2001 in the educational qualifications needed to get a job where qualifications are measured on a five-point scale with one representing no qualifications and five a college degree. In all occupations there is a rise in the level of qualifications required, with a very large rise in sales occupations and elementary occupations. This could reflect greater skill requirements within occupations or a greater use of credentialism. There is evidence (shown in the third column) that more workers report in 2001 that the education required to get the job is not necessary to do the job but, in the absence of any information on the extent to which education is underutilized, one cannot know whether this effect is large enough to outweigh the positive effect on skill levels of an increase in the level of education required.

These two pieces of disparate evidence are not entirely consistent. The DOT data do not suggest any significant skill upgrading within occupations, while the SCELI/SS data suggest an increase in the level of education required by employers although also an increasing proportion of workers reporting that this education is unnecessary to do the job. But, it does seem that the supply of skills may be increasing faster than the demand in the bottom half of the distribution because the extent of overqualification does not seem to be falling over time and, according to some estimates (for example, Felstead, Gallie, & Green, 2002), is actually increasing.[14] Consistent with this, Felstead, Gallie, and Green report that there is excess demand for workers with no qualifications, an excess supply of people with low-level qualifications, and a rising use of credentialism among the lowest-level occupations.

TABLE 9.—CHANGES IN SKILL REQUIREMENTS WITHIN JOBS, 1986–2001

Occupation	Change in Education Level Required to Get Job, 1986–2001	Change in Fraction Reporting Required Education Not Necessary to Do Job, 1986–2001
Managerial	0.25	0.014
Professional	0.12	0.021
Associate professional	0.31	0.072
Clerical	0.10	0.046
Craft	0.25	0.043
Personal services	0.50	0.110
Sales	0.54	0.093
Operatives	0.08	0.063
Elementary	0.24	0.076

Notes: Data come from 1986 SCELI data and 2001 Skills Survey. Education required is measured on a five-point scale with one being no qualifications and five a college degree.

VIII. Job Polarization and the Rise in Wage Inequality

All the analysis so far has been about the quantity side of the labor market—what is happening to the employment of different types of workers. But, the polarization of employment could also be expected to have led to increased wage inequality. Of some interest is what fraction of the rise in wage inequality can be explained by this polarization of employment. This is the subject of this section.

In figure 6 we present the evolution of two measures of actual wage inequality over the period 1976–1995, the log(90/50) and the log(50/10) differentials, as well as a prediction of what would have happened if the only change in the wage distribution taking place is the change in the distribution of jobs in the economy. To this end we assign everyone in the base year (here, 1976) a weight that is equal to the total number of workers in a job in a given year divided by the job cell size in the initial period. We then

[14] It has always been something of a puzzle to reconcile these findings of widespread overqualification with rising employment and relative wages of educated workers. Our finding of increased job polarization can explain why both phenomena may coexist. The increased supply of skills that has been necessary to meet the increased number of lovely jobs poses a problem for the increased number of lousy jobs. Because there has been an increase in the mean but no increase in the variance of educational qualifications, those in lousy jobs are increasingly likely to have higher levels of education than necessary for doing the job.

FIGURE 6.—HOW MUCH OF ACTUAL WAGE DISPERSION CAN BE EXPLAINED BY JOB POLARIZATION?

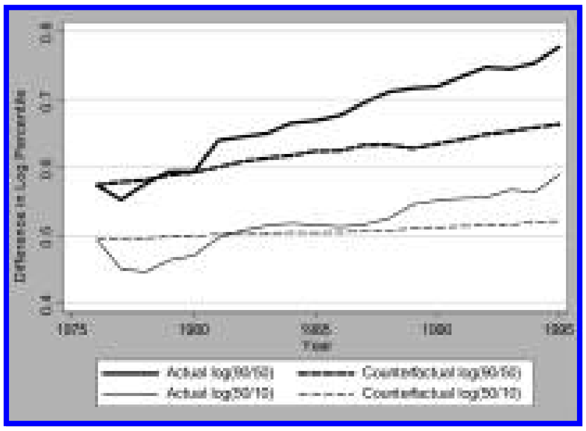

Notes: Data are taken from the NES. The figure uses three-digit SOC90 codes as the definition of a job. The counterfactual keeps constant median wage and wage dispersion within occupations.

FIGURE 7.—THE IMPACT OF JOB POLARIZATION AND CHANGING RELATIVE WAGES ACROSS JOBS ON WAGE INEQUALITY

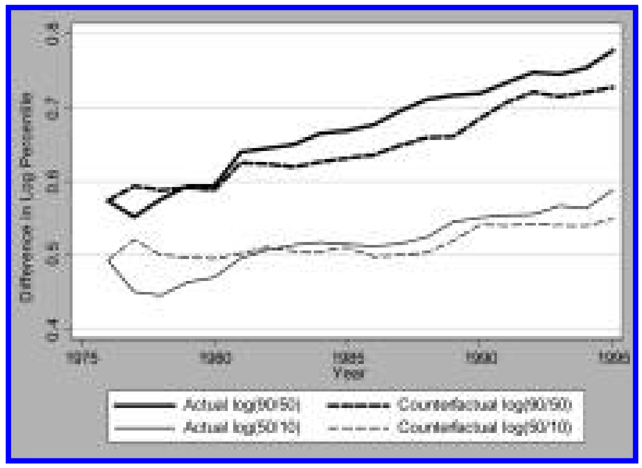

Notes: Data are taken from the NES. The figure uses three-digit SOC90 codes as the definition of a job. The counterfactual keeps constant wage dispersion within occupations but allows the actual median wage to vary in line with the data.

compute counterfactual percentiles of the reweighted wage distribution. As is well known, actual wage inequality rose very strongly in this period following a fall in 1977 (the result of the Social Contract incomes policy then in place). The rise in inequality is somewhat larger at the top of the distribution than at the bottom. Since the counterfactual log median increases only very little, the rises in the counterfactual log(90/50) and log(50/10) reflect large polarization. In comparison with the actual changes, increased job polarization can explain 33% of the increase in the log(50/10) differential between 1976 and 1995 and 54% of the increase in the log(90/50) wage differential.

The remaining rise in wage inequality can be thought of as coming from one of two sources: differential changes in median wages across jobs and within-job wage inequality. For example, wage inequality will rise if median wages have risen faster in good jobs than in bad jobs. The other potential source of increased wage inequality is an increase in within-job pay dispersion. To look for evidence of this, table 10 reports regression estimates of median wage growth onto the log of the initial median wage.[15] Since estimates will be biased downward when using the initial wage on both sides of the regression equation, we use the NES and run regressions using wages in 1977 (rather than 1976) as a covariate. All point estimates are positive and all are statistically significant. These results suggest that allowing median wages to change over time while keeping the variance of pay within each job constant could close the actual-counterfactual gap further. The implications for wage inequality are presented in figure 7. Here, we do the reweighting described earlier and also adjust wages in every job cell by the change in log median wage in that cell. Now 51% of the increase in the log(50/10) differential between 1976 and 1995 and 79% of the increase in the log(90/50) wage differential can be explained.

[15] We experimented with the inclusion of a quadratic term but this was never significant.

TABLE 10.—RELATIONSHIP BETWEEN WAGE GROWTH AND INITIAL MEDIAN WAGE

Sample	Sample Period	Data	Relative to 1976	Relative to 1977
Men + Women	1976–95	NES	0.239	0.267
		(occ)	(0.029)	(0.029)
Men + Women	1976–95	NES	0.109	0.156
		(occXind)	(0.013)	(0.014)
Men	1976–95	NES	0.374	0.403
		(occ)	(0.033)	(0.033)
Men	1976–95	NES	0.178	0.237
		(occXind)	(0.017)	(0.017)
Women	1976–95	NES	0.284	0.371
		(occ)	(0.052)	(0.052)
Women	1976–95	NES	0.086	0.188
		(occXind)	(0.027)	(0.029)

Notes: Regressions are weighted by initial job cell size in terms of employment. The dependent variable is the change in log wages between 1976 and 1995, the regressors the log wage in 1976 or 1977.

FIGURE 8.—THE CHANGING R^2 IN THE U.K. EARNINGS FUNCTION

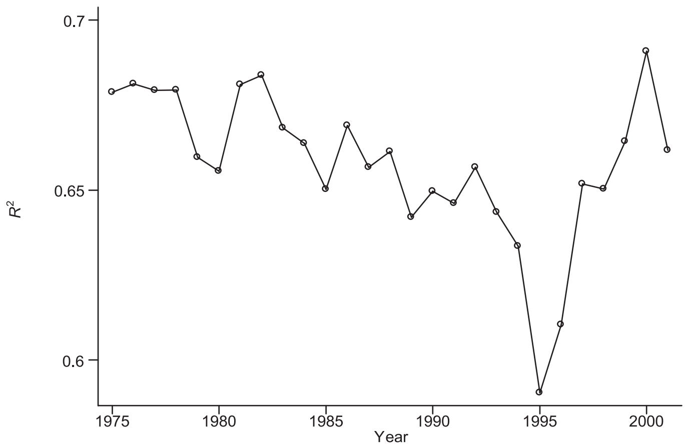

Notes: Data are taken from the NES. The dependent variable is log hourly earnings and the covariates included are age, industry, and occupation dummies, all interacted with gender. The dip in 1995 and 1996 is the result of considerably more workers reporting implausibly low wages of below £0.5 per hour. If the data is trimmed the dip is much less pronounced.

A. Within and Between Group Wage Inequality

One implication of this is that the rise in within-job wage inequality has a relatively minor part to play in explaining the overall rise in wage inequality. This is in contrast with some studies that try to explain wage inequality in terms of age and education that typically find that most of the rise in inequality is due to rising within-group wage inequality (see Levy and Murnane (1992) and Katz and Autor (1999) for a survey of the U.S. literature and Machin (2002) for the U.K.). The studies are correct given the variables they use to try to explain the rise in wage inequality, but the evidence here suggests that this conclusion is sensitive to how the groups are defined. Unfortunately, a small industry has been established based on the premise that wage inequality has risen very markedly among "identical" workers and has been building theoretical explanations of this "fact."

One particularly simple way to understand this is to consider what is happening to the R^2 in earnings functions. Figure 8 graphs the R^2 from an earnings function estimated for each year on the NES in which the dependent variable is log hourly earnings, and the covariates include a complete set of dummies for age, industry, and occupation, all interacted with gender. There are two things to note: first the R^2 is high—averaging almost two-thirds—compared to the one-third found in a standard specification using the U.S. Current Population Survey (U.S. CPS). Secondly, there is no marked trend in the R^2 over time if one also includes the most recent years for which data are available (we believe this is also true for the U.S. CPS, see for example Lemieux (2002)). The consequence is that the rise in the residual variance can explain only one-third of the total rise of the variance in log wages.

The conclusion that the importance of within-group wage inequality depends on the controls one includes in an earnings function seems also consistent with U.S. studies that have more detailed controls than is usual in earnings functions. For example, Dunne et al. (2000) have controls for establishment fixed effects (which are obviously better than industry) and find they can explain much of the rise in wage inequality by widening between-plant wage gaps. It seems likely that much of these wage gaps between plants can be explained in terms of the characteristics (in gender, age, education, and occupation) of the workers within them—for example, Hellerstein, Neumark, and Troske (1999) find that a fairly rudimentary set of controls (fewer than thirty) can explain 40% of the variation in average wages across establishments.

B. Explaining Relative Wage Changes

It is relatively easy to explain why wages at the top of the distribution have been increasing relative to the median as an increase in the demand in good jobs that has not been matched by an increase in supply. But, it is not clear why the increase in the demand for bad jobs has not resulted in a rise in wages at the bottom relative to the median. The rise in the number of bad jobs has coincided with a decline in their pay, not just relative to the good jobs that are increasing in number but also relative to the middling jobs that are decreasing in number. If the labor market is competitive, this does not seem consistent with a view in which technology causes a shift in the demand for different types of labor but the supply curve is stable, and the observed changes in wages and employment are simply movements along this supply curve.[16]

However, it may be that the average skill level of workers in different jobs is changing in a different way. For example it may be the less skilled who are displaced from the middling jobs, so the average skill of those who remain rises. In line with this hypothesis, figure 5 suggests somewhat greater educational upgrading in middling than in lousy jobs. If this is the case then the average level of human capital may have risen in middling relative to lousy jobs, and this could account for the relative wage movements. This is also in line with the literature on displaced workers (for example, Kletzer, 1998; Farber, 1999) showing that it was blue-collar workers in manufacturing who disproportionately suffered displacement and who also suffered large earnings losses. It would therefore be interesting to know whether routinization has caused displacement from middling to lousy jobs.

Another possible explanation for why wages have been falling in lousy jobs relative to those in middling jobs is to think of the labor market as being noncompetitive in some way. There are a number of ways in which this could be

[16] The discussion-paper version of this paper, Goos and Manning (2003), presents a simple three-skill competitive model of the labor market that can be used as a more formal justification for the discussion that follows. Juhn (1994) presents a model in which she claims that a fall in the demand for middling jobs reduces wages more at the bottom, but hers is really only a model with two types of skills but middling people with some combination of both skills. In this example there is no well-defined sense of a fall in demand for middling jobs.

done. For example, Acemoglu (2001) presents a model of a labor market with frictions in which an increase in the supply of skilled workers encourages employers to create more lovely and lousy jobs and fewer middling jobs. In this type of model "supply creates its own demand" and there is no need to resort to demand shocks caused by technological change to explain job polarization. But it is a little bit hard to see how supply shocks of this type can explain the pattern of changes in occupational employment documented above—technology seems much more plausible as an explanation for these changes. But, the Acemoglu story may have some relevance for explaining what is happening within occupations when employers often have a decision about what level of skill to require of workers doing these jobs.

Another noncompetitive explanation is that institutions have changed in such a way as to lead to a fall in wages at the bottom end of the wage distribution. There is now a small literature in the United States (DiNardo, Fortin, & Lemieux (1996); Lee, 1999; Teulings, 2000) that suggests that the evolution of unionization and the minimum wage can do a very good job in explaining what is happening to the bottom half of the wage distribution. The United Kingdom has also seen a marked decline in unionization, a decline in minimum wages (though they were never very strong), and the indexation of welfare benefits to prices, not wages. Perhaps these changes can account for the rise in wage inequality in the bottom half of the distribution in the 1980s.

IX. Conclusions

There is little doubt that technology has a powerful impact on the labor market. But, the dominant current view about the nature of its impact, the hypothesis of skill-biased technical change, is only a partial truth and cannot explain all of the important changes in the labor market such as job polarization (see Card and DiNardo (2002) for an additional list of puzzles and problems). Crudely, the SBTC hypothesis can explain what is happening in the top half of the wage distribution but not the bottom half. However, the more nuanced view about the impact of technology proposed by Autor, Levy, and Murnane (2003)—the routinization hypothesis—provides a plausible explanation for why the demand for middling jobs has fallen and why we see the process of job polarization.

Job polarization may also be able to explain the increase in U.K. and U.S. lower-tail wage inequality during the 1980s as workers are displaced from middling jobs toward lousy jobs. However, as polarization continues (possibly reinforced by an expansion of routine task offshoring), lower-tail wage inequality could decrease as the scarcity of displaced workers drives up the relative wages of workers in lousy jobs. It would therefore be interesting to know whether polarization can explain the recent slowdown in rising lower-tail wage inequality in the United Kingdom and the United States and whether similar labor market changes can be observed in the continental European countries that have not had the rises in wage inequality seen in the United Kingdom and the United States.

REFERENCES

Acemoglu, Daron, "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence," *American Economic Review* 89 (1999), 1259–1278.

——— "Good Jobs versus Bad Jobs," *Journal of Labor Economics* XIX (2001), 1–21.

Autor, David H., Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics* CXVIII (2003), 1279–1333.

Baumol, William J., "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review* LVII (1967), 415–426.

Berman, Eli, John Bound, and Zvi Griliches, "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures," *Quarterly Journal of Economics* CIX (1994), 367–397.

Berman, Eli, John Bound, and Stephen Machin, "Implications of Skill-Biased Technological Change: International Evidence," *Quarterly Journal of Economics* CXII (1998), 1245–1279.

Bluestone, Barry, and Bennet Harrison, "The Growth of Low-Wage Employment: 1963–1986," *American Economic Review* LXXVIII (1988), 124–128.

Borjas, George J., Richard B. Freeman, and Lawrence F. Katz, "Searching for the Effect of Immigration on the Labor Market," *American Economic Review,* 86:2 (May 1996), 246–251.

Burtless, Gary, *A Future of Lousy Jobs? The Changing Structure of U.S. Wages* (Washington, DC: Brookings Institution Press, 1990).

Card, David, and John DiNardo, "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics* XX (2002), 733–783.

Card, David, and Thomas Lemieux, "Wage Dispersion, Returns to Skill, and Black-White Wage Differentials," *Journal of Econometrics* LXXIV (1996), 319–361.

Chevalier, Arnaud, "Graduate Over-education in the UK," LSE, Centre for the Economics of Education discussion paper no. 7 (2000).

Costrell, Robert M., "Methodology in the 'Job Quality' Debate," *Industrial Relations* XXIX (1990), 94–110.

DiNardo, John., Nicole Fortin, and Thomas Lemieux, "Labor Market Institutions and the Distribution of Wages, 1973–1992," *Econometrica* LXIV (1996), 1001–1044.

Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth Troske, "Wage and Productivity Dispersion in US Manufacturing: The Role of Computer Investment," NBER working paper 7465 (2000).

Farber, Henry S., "Job Creation in the United States: Good Jobs or Bad?" Princeton Industrial Relations Section working paper no. 385 (July 1997).

——— "Mobility and Stability: The Dynamics of Job Change in Labor Markets," in Orley Ashenfelter and David Card (Eds.), *Handbook of Labor Economics,* vol. 3 (San Diego, CA: Elsevier Science, 1999).

Feenstra, Robert C, and Gordon H. Hanson, "The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979–1990," *Quarterly Journal of Economics* CXIV (1999), 907–940.

Felstead, Alan, Duncan Gallie, and Francis Green, "Work Skills in Britain, 1986–2001," Department for Education and Skills (2002).

Richard B. Freeman, "Trade Wars: The Exaggerated Impact of Trade in Economic Debate," NBER working paper 10,000 (September 2003).

Gittleman, Maury B, and David R. Howell, "Changes in the Structure and Quality of Jobs in the United States: Effects by Race and Gender, 1973–1990," *Industrial and Labor Relations Review* XLVIII (1995), 420–440.

Goos, Maarten, and Alan Manning, "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," LSE CEP discussion paper no. 604 (Dec. 2003).

Green, Francis, and Steven McIntosh, “Is there a Genuine Underutilisation of Skills Amongst the Over-Qualified,” LSE CEP working paper (2002).

Green, Francis, Steven McIntosh, and Anna Vignoles, “Overeducation and Skills—Clarifying the Concepts,” LSE CEP DP no. 435 (1999).

Hartog, J., “On Returns to Education: Wandering Among the Hills of ORU Land,” in Hans Heijke and Joan Muysken (Eds.), *Education and Training in a Knowledge-Based Economy* (London: Macmillan, 2000).

Hellerstein, Judith K, David Neumark, and Kenneth R. Troske, “Wages, Productivity and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations,” *Journal of Labor Economics* VII (1999), 409–446.

Howell, David R., and Edward N. Wolff, “Trends in the Growth and Distribution of Skills in the U.S. Workplace, 1960–1985,” *Industrial and Labor Relations Review* XLIV (1991), 486–502.

Ilg, Randy E., “The Nature of Employment Growth, 1989–1995,” *Monthly Labor Review* CXIX (1996), 29–36.

Ilg, Randy E., and Steven E. Haugen, “Earnings and Employment Trends in the 1990s,” *Monthly Labor Review* (March 2000), 21–33.

Juhn, Chinhui, “Wage Inequality and Industrial Change: Evidence from Five Decades,” NBER working paper no. 4684 (1994).

——— “Wage Inequality and Demand for Skill: Evidence from Five Decades,” *Industrial and Labor Relations Review* LII (1999), 424–443.

Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce, “Wage Inequality and the Rise in Returns to Skill, *Journal of Policital Economy* CI (1993), 410–442.

Katz, Lawrence F., and David H. Autor, “Changes in the Wage Structure and Earnings Inequality” in Orley Ashenfelter and David Card (Eds.), *Handbook of Labor Economics,* vol. 3A (pp. 1463–1555), (Amsterdam: North-Holland, 1999).

Kletzer, Lori G., “Job Displacement,” *Journal of Economic Perspectives* 12:1 (Winter 1998), 115–136.

Kosters, Marvin H., and Murray N. Ross, “A Shrinking Middle Class?” *Public Interest* 1988, 3–27.

Lee, David S., “Wage Inequality in the United States During the 1980s: Rising Dispersion or Falling Minimum Wage?” *Quarterly Journal of Economics* CXIV (1999), 941–1023.

Lemieux, Thomas, “Decomposing Wage Distributions: A Unified Approach,” *Canadian Journal of Economics* 35:4 (November 2002), 646–688.

Levy, Frank, and Richard J. Murnane, “U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanation,” *Journal of Economic Literature* 30:3 (September 1992), 1333–1381.

Machin, Stephen, “Labor Market Inequality and Changes in the Relative Demand for Skills,” LSE mimeograph (2002).

Machin, Stephen, and John van Reenen, “Technology and Changes in Skill Structure: Evidence from Seven OECD Countries,” *Quarterly Journal of Economics* CXIII (1998), 1215–1244.

Meisenheimer, Joseph R. II, “The Services Industry in the ‘Good’ versus ‘Bad’ Jobs Debate,” *Monthly Labor Review* (February 1998), 22–47.

Murphy, Kevin M., and Finis Welch, “Occupational Change and the Demand for Skill, 1940–1990,” *American Economic Review* LXXXIII (1993), 122–136.

OECD, “The Characteristics and Quality of Service Sector Jobs,” *OECD Employment Outlook* (2001), 89–128.

Sicherman, Nachum, “Overeducation in the Labor Market,” *Journal of Labor Economics* IX (1991), 101–122.

Simon, Herbert A., “The Corporation: Will It Be Managed by Machines?” in M. L. Anshen and G. L. Bach (Eds.), *Management and Corporations 1985* (New York: McGraw-Hill, 1960).

Spenner, Kenneth I., “Deciphering Prometheus: Temporal Change in the Skill Level of Work,” *American Sociological Review* XLVIII (1983), 824–837.

Spitz, Alexandra, “Technical Change, Job Tasks, and Rising Educational Demands: Looking Outside the Wage Structure,” *Journal of Labor Economics* 24:2 (April 2006), pp. 235–270.

Teulings, Coen N., “Aggregation Bias in Elasticities of Substitution and the Minimum Wage Paradox,” *International Economic Review* XLI (2000), 359–398.

Wright, Erik Olin, and Rachel Dwyer, “The Patterns of Job Expansions in the United States: A Comparison of the 1960s and 1990s,” *Socio-Economic Review* 1 (2003), 289–325.