

The Polarization of the U.S. Labor Market*

David H. Autor
MIT and NBER

Lawrence F. Katz
Harvard University and NBER

Melissa S. Kearney
The Brookings Institution and NBER

January 2006

* We are grateful to Daron Acemoglu, Eli Berman, Frank Levy, and Sendhil Mullainathan for superb suggestions and to Michael Anderson and Tal Gross for excellent research assistance. Autor acknowledges support from the National Science Foundation (CAREER SES-0239538) and the Sloan Foundation. Katz acknowledges support from the Mellon Foundation and the Radcliffe Institute for Advanced Study.

The Polarization of the U.S. Labor Market
David H. Autor, Lawrence F. Katz, and Melissa S. Kearney
NBER Working Paper No.
January 2006
JEL No. J3, D3, O3

ABSTRACT

This paper analyzes a marked change in the evolution of the U.S. wage structure over the past fifteen years: divergent trends in upper-tail (90/50) and lower-tail (50/10) wage inequality. We document that wage inequality in the top half of the distribution has displayed an unchecked and rather smooth secular rise for the last 25 years (since 1980). Wage inequality in the bottom half of the distribution also grew rapidly from 1979 to 1987, but it has ceased growing (and for some measures actually narrowed) since the late 1980s. Furthermore we find that occupational employment growth shifted from monotonically increasing in wages (education) in the 1980s to a pattern of more rapid growth in jobs at the top and bottom relative to the middles of the wage (education) distribution in the 1990s. We characterize these patterns as the “polarization” of the U.S. labor market, with employment polarizing into high-wage and low-wage jobs at the expense of middle-wage work. We show how a model of computerization in which computers most strongly complement the non-routine (abstract) cognitive tasks of high-wage jobs, directly substitute for the routine tasks found in many traditional middle-wage jobs, and may have little direct impact on non-routine manual tasks in relatively low-wage jobs can help explain the observed polarization of the U.S. labor market.

David H. Autor
Department of Economics
MIT
Cambridge, MA 02142
and NBER
dautor@mit.edu

Lawrence F. Katz
Department of Economics
Harvard University
Cambridge, MA 02138
and NBER
lkatz@harvard.edu

Melissa S. Kearney
The Brookings Institution
1775 Massachusetts Ave. NW
Washington, DC 2003
and NBER
mkearney@brookings.edu

Much research (surveyed in Lawrence F. Katz and David H. Autor [1999] and Daron Acemoglu [2002]) documents a substantial widening of the U.S. wage structure since the late 1970s, largely driven by increases in educational wage differentials and residual (within-group) wage inequality. The growth in overall wage inequality was most rapid during the 1980s, and involved a spreading out of the entire wage distribution. Rapid secular growth in the demand for skills, partly driven by skill-biased technical change (SBTC), combined with a sharp slowdown in the growth of the relative supply of college workers (starting in the early 1980s) help to explain these wage changes. The erosion of labor market institutions protecting low- and middle-wage workers—the minimum wage and unions—further contributed to rising wage inequality, especially during the 1980s.

Recent work challenges these conclusions emphasizing a slowdown in the growth of U.S. wage inequality over the last fifteen years (David Card and John DiNardo 2002; Thomas Lemieux 2006). This “revisionist” literature views the surge in wage inequality of the 1980s as an “episodic” event driven by institutional forces and argues that the more “modest” inequality growth in the 1990s is inconsistent with a key role for SBTC and other market forces.

In this paper, we reconsider this revisionist view, focusing on a marked change in the evolution of the U.S. wage structure over the past 15 years: divergent trends in upper- and lower-tail wage inequality. In Section I, we document that wage inequality in the top half of the distribution has exhibited an unchecked secular rise for the last 25 years (since 1980), but it has ceased growing since the late 1980s (and for some measures narrowed) in the bottom half of the distribution. In Section II, we show that the pattern of employment growth differed sharply in the 1990s versus the 1980s with more rapid growth of employment in jobs (occupations) at the bottom and top relative to the middle of the wage (and skill) distribution. Borrowing the

terminology of Maarten Goos and Alan Manning (2003) from work on Britain, we characterize this pattern as the “polarization” of the U.S. labor market, with employment polarizing into high-wage and low-wage jobs at the expense of traditional middle-skill jobs. We document that quantity and price changes covary positively throughout the distribution of earnings—both in the 1980s, when the wage structure was spreading monotonically, and in the 1990s, when the wage structure was polarizing. These patterns suggest a central role for labor demand shifts in explaining wage structure changes of the last twenty-five years. We show in Section III how a simple model of computerization (based on Autor, Frank Levy, and Richard Murnane [2003]) in which computers complement non-routine (abstract) cognitive tasks, substitute for routine tasks, and may have little impact on non-routine manual tasks, can help explain the observed polarization of the U.S. labor market. Section IV concludes.

I. Divergent Upper- and Lower-Tail Wage Inequality

Figure 1 displays the evolution of the 90-50 and 50-10 log hourly wage differentials for all workers (males and females combined) from 1973 to 2004 using (hours-weighted) wage data from the Current Population Survey (CPS) May samples (for 1973-78) and Merged Outgoing Rotation Group (MORG) samples (for 1979 to 2004).¹ The plot shows similar substantial increases in upper-half (90-50) and lower-half (50-10) wage inequality from 1979 to 1987, expanding the 90-10 log wage differential by 21 log points. But the trends in top-half and bottom-half inequality diverge after 1987 with upper-half wage inequality continuing to rise steadily over the last 15 years while the growth in lower-half inequality ceased in the late 1980s

¹ All of our sample selection and data processing steps for CPS wage data are the same as those described in the Data Appendix of Autor, Katz, and Melissa S. Kearney (2005a) extended to cover wage data for 2004.

(with an actual contraction in the 50-10 log wage differential of 4 log points from 1987 to 2004).² The net effect has been a slower growth of only 4 log points in the 90-10 log wage differential since 1987. These trends are robust across numerous samples and data sources. Monotonic rising wage inequality in the 1980s and diverging upper- and lower-tail wage inequality since the late 1980s are observed for males and females separately, for hourly wages (based on annual earnings divided by hours worked in the previous calendar year) from the CPS March Samples, and for the weekly wages of full-time workers in both the March and MORG CPS samples (Autor, Katz, and Kearney 2005a).³

Figure 2 illustrates more precisely where in the wage structure the divergence of upper- and lower-tail wage inequality trends has occurred. Figure 2 plots cumulative log hourly real earnings growth (indexed using the Personal Consumption Expenditures implicit price deflator) by wage percentile for 1973 to 1988 and for 1988 to 2004. The figure shows an almost linear spreading out of the entire wage distribution (from the 3rd to 97th percentile) in the first half of the sample (driven by changes in the 1980s). In contrast, wage growth has polarized since 1988, with faster wage growth in the bottom quartile than in middle two quartiles and with the most rapid rise and a continued spreading out of the wage distribution in the top quartile.

The divergence of upper- and lower-tail wage inequality since the late 1980s also holds for residual wage inequality and is robust to (and even reinforced by) adjustments for labor force composition (demographic) changes (Autor, Katz, and Kearney 2005b).⁴ Similar patterns are

² This divergence in trends in the 90-50 and 50-10 log hourly wage differentials using CPS data has also been emphasized by Lawrence Mishel, Jared Bernstein, and Sylvia Allegretto (2005).

³ The magnitude of wage inequality growth is larger for both males and females when analyzed separately than for the overall sample (given the narrowing of the gender gap), and wage inequality grew by more over the last fifteen years using March rather than the MORG CPS samples.

⁴ Between-group wage inequality similarly shows a faster growth in the top-half than the bottom-half of the distribution since the 1980s. For example, Lemieux (2005) documents the growing “convexification”

found for total employee compensation in the micro data from Employment Cost Index (Brooks Pierce 2001) and for annual wage and salary earnings based on Internal Revenue Service (IRS) micro data (Ian Dew-Becker and Robert J. Gordon 2005). Thomas Piketty and Emmanuel Saez (2003) further document the persistent divergence of annual earnings in the very upper-end of distribution (the top 1%) relative to other workers since the 1970s using IRS tax-return data.

Some researchers speculate that the growing skewness of the upper part of the wage distribution reflects changes in social norms (Piketty and Saez 2003), and others (Card and DiNardo 2002; Lemieux 2006) argue that the slowdown of 90-10 wage inequality growth in the 1990s speaks in favor of institutional factors (especially the minimum wage) and against SBTC and other market mechanisms for changes in the wage structure. But the time pattern of changes in the minimum wage does not do much to explain the evolution of wage inequality even in the bottom half of the distribution outside the 1979 to 1987 period (Autor, Katz and Kearney 2005a). The ‘social norms’ explanation—while potentially promising—needs further refinement to be directly testable.

We examine how a more nuanced view of technological change than the usual SBTC model (based on a single-index view of skill) may contribute to an explanation for observed wage inequality trends. Building on Autor, Levy and Murnane (2003), our conceptual framework observes that the first-order impact of computerization is to displace (substitute for) a set of ‘middle skilled’ routine cognitive and manual tasks, such as bookkeeping, clerical work and repetitive production tasks.⁵ If these routine tasks are more complementary to high-skilled, abstract tasks (problem-solving, management and coordination) than to ‘non-routine manual’

of U.S. returns to schooling over recent decades with the returns to post-secondary schooling (especially post-college training) rising sharply and returns to lower levels of education remaining largely unchanged.
⁵ Levy and Murnane (2004) provide a rich elaboration and extension of these ideas. In earlier work, Chinhui Juhn (1994) discusses evidence and explanations for a decline in demand for ‘middle-skill’ workers.

tasks (such as those performed by truck drivers, waiters and security guards), the computerization of routine work can lead to a polarization of the structure of employment and earnings. The model predicts that a period of wage structure polarization should be accompanied by employment polarization, a pattern that is strongly supported by our analysis below.

II. Job Polarization Trends

We examine trends in the “quality,” skill content, and task content of U.S. jobs since 1980. We follow Goos and Manning’s (2003) analysis of the U.K. in exploring how U.S. employment growth by occupation has been related to skill (or job quality) proxied by initial wages or educational levels.

We first sort occupations into percentiles by their median hourly wage in 1980 (actually based on 1979 earnings) using data from the 1980 Census Integrated Public Use Microsample (IPUMS). Figure 3 plots employment growth, measured as a share of total hours worked in the economy, from 1980-1990 and 1990-2000, against 1980 occupational wage percentile using employment data from the 1980, 1990, and 2000 IPUMS.⁶ For the 1980s, the figure shows declining employment at the bottom end of the occupational distribution and rather monotonic increases in employment shares as one moves up the wage structure (with a little dip from the 55th to 80th percentile). In sharp contrast, employment growth in the 1990s appears to have polarized. We observe rapid employment growth in low-wage and high-wage jobs in the 1990s and slow growth in the middle of the wage distribution.

⁶ We employ a consistent set of occupation codes for Census years 1980, 1990 and 2000 developed by Peter B. Meyer and Anastasiya M. Osborne (2005). We use a locally weighted smoothing regression (bandwidth 0.8 with 100 observations) to fit the relationship between decadal growth in occupational employment share and occupations’ initial skill percentile in the 1980 skill distribution, measured by median occupational wages or average years of schooling. Figures 3 and 4 plot the fitted values from these regressions.

This pattern of job growth (and possible labor demand shifts) corresponds well with the observed pattern of wage structure changes shown in Figure 2. In a comparison of the 1980s to the 1990s, we find that employment growth is more rapid in the 1990s below the 25th and above the 65th percentiles of the occupational wage distribution and declined in the middle (from the 25th to the 65th percentiles).

We test the robustness of these patterns by utilizing an alternative skill definition: the mean years of schooling in an occupation in 1980. Figure 4 displays the analogous plot of U.S. occupational employment growth in the 1980s and 1990s by 1980 occupational education percentiles. Using this skill measure, we find near-monotonic (almost linear) increases in employment growth moving up the education distribution in the 1980s. We also find that employment growth has polarized in the 1990s, with more rapid employment growth in the bottom end of the education distribution than in the middle. The most rapid (and diverging) employment gains are concentrated in the top part of the education distribution.

Acemoglu (1999) also finds U.S. job polarization from the mid-1980s to early 1990s using more aggregated industry-occupation cells, and Goos and Manning (2003) document job growth polarization by detailed occupations from 1979 to 1999 for Britain—another country with a large increase in wage inequality. Furthermore, Spitz-Oener (2006) reports job polarization in West Germany, a labor market where wage inequality has been rather quiescent until recently. These similar patterns of recent employment growth polarization for several rich countries are suggestive of similar market demand shifts favoring jobs at the top and bottom relative to the middle of the skill (wage) distribution. But the much larger growth in wage inequality in the United States and Britain than in Germany in recent decades also suggests that differences in

labor market and educational institutions may greatly affect how such demand shifts translate into wage structure changes.

A third methodology for measuring U.S. employment structure trends is to directly examine changes in job task content. Autor, Levy, and Murnane (2003) take this approach using data on detailed tasks from the Dictionary of Occupational Titles aggregated to 3-digit Census occupations. Though not depicted, we have extended their data analysis of employment growth in industry-gender-education cells using CPS data through 2002. We find that employment growth was most rapid in the 1990s for cells intensive in non-routine cognitive tasks (those most complementary with computerization), was declining and decelerating in the 1990s for cells most intensive in routine cognitive and manual tasks (those most substitutable for computers), and ceased declining in the 1990s for (typically low-wage) jobs intensive in non-routine manual tasks.

These patterns of employment growth by wages, education, and task intensity suggest that labor demand shifts have favored low- and high-wage workers relative to middle-wage workers over the last fifteen years. This pattern stands in contrast to the shifts in labor demand evident during the decade of the 1980s, which appear to have been monotonically rising in skill. We now offer a conceptual framework for interpreting these trends.

III. A Model and Interpretation

Our framework, building on Autor, Levy, Murnane (2003), considers how a decline in the real price of computing power may lead to a polarization of work, as measured by traditional skill measures. The model builds on four observations that we believe are well-supported by the case-study and representative evidence (see Levy and Murnane [2004]). First, computer

capital—denoted by K in our model and measured in efficiency units—is a close substitute for human labor in routine cognitive and manual tasks, such as bookkeeping, clerical work and repetitive production tasks. Second, routine task input—embodied in either computer capital or human labor—is a complement to workers engaged in abstract reasoning tasks such as problem solving, coordination, and other high-level management activities. Third, there exists a panoply of non-routine manual tasks for which computers currently neither directly substitute nor strongly complement, such as the job tasks performed by truck drivers, security guards, waiters and janitors. For brevity, we refer to these three categories of tasks as Abstract (A), Routine (R) and Manual (M), and we consider them to roughly correspond to high, intermediate and low skilled occupations as measured by earnings or education.

We observe that workers' ability to engage in specific tasks is partially contingent on their education. For simplicity, we assume that there are two types of workers: College workers, who can perform Abstract tasks, and High School workers who can freely substitute between Routine and Manual tasks.

The exogenous driving force in our model is the well-documented precipitous decline in the price of computing power of recent decades. Computerization—by which we mean the falling price of computer power—lowers the price of Routine task input and increases demand for Routine tasks. We sketch the formal model here and refer the reader to the Theory Appendix for full details.

Output in this economy, priced at unity, is produced using the aggregate Cobb-Douglas production function $Y = A^\alpha R^\beta M^\gamma$ with $\alpha, \beta, \gamma \in (0, 1)$, $\alpha + \beta + \gamma = 1$, and where A, R and M denote the three tasks above. Abstract and manual tasks are performed by workers who supply labor inputs, L_A and L_M . Routine tasks can be performed either by workers who supply L_R or by

computer capital, K , measured in efficiency units, which is a perfect substitute for L_R .

Computer capital is supplied perfectly elastically to Routine tasks at price ρ , which is falling over time at an exogenous rate.

There are a large number of workers with educational supplies assumed exogenous. A fraction $\theta \in (0,1)$ are High School workers and the remaining $1-\theta$ are College workers. Each College worker is endowed with one efficiency unit of Abstract skill, which she supplies inelastically to Abstract tasks. Each High School worker i is endowed with one efficiency unit of skill in Manual tasks and η_i efficiency units of skill in Routine tasks, where η is continuously distributed on the unit interval with positive probability mass at all points $\eta \in [0,1]$. High school workers do not possess Abstract skills.⁷

The supply of High School labor to Routine and Manual tasks is determined by self-selection. Let w_m and w_r each the wages paid to Manual and Routine tasks. Each High School worker i chooses to supply either one efficiency unit of labor to Manual tasks if $\eta_i < w_m / w_r$ or supplies η_i efficiency units of labor to Routine tasks otherwise. Thus labor supply to Manual (Routine) tasks is weakly upward (downward) sloping in w_m / w_r .

Equilibrium in the model occurs when: (1) the economy operates on the demand curve of the aggregate production function; (2) each factor is paid its marginal product ($w_A = \partial Y / \partial A$, $w_R = \partial Y / \partial R$ and $w_M = \partial Y / \partial M$); and (3) the labor market clears so that no worker wishes to reallocate labor input among tasks.

⁷ A more complete version of this model would allow for endogenous longer-run educational supply responses to college versus high school earnings opportunities, endow college workers with routine as well as abstract skills, endow high school workers with some abstract skills, and allow for heterogeneity in abstract skills among college workers. Heterogeneity in abstract skills among college workers will tend to reinforce the spreading out of the top end of the wage distribution in response to declining real computer prices. Endogenous education responses could eventually work to offset some of the impacts of computer price declines on wage inequality.

Because computer capital is a perfect substitute for routine labor input, it is immediate that $w_r = \rho$. Hence, a decline in ρ reduces w_r one for one. We are interested in the effect of a decline in ρ on: (1) the equilibrium quantity of Routine task input; (2) the allocation of labor between Routine and Manual tasks; (3) the wage paid to each task; and (4) the observed wage in each job type—which, in the case of Routine tasks, may differ from the wage paid per efficiency unit of Routine task input.

Since own-factor demand curves are downward sloping, a decline in ρ raises demand for Routine tasks. This demand can be supplied by either additional computer capital or additional Routine labor input. Due to worker self-selection, the additional demand will be supplied by computer capital. The reason, as we demonstrate formally in the Theory Appendix, is that Manual and Routine tasks are q-complements—a rise in Routine input (spurred by the fall in ρ) raises the marginal productivity of Manual task input.

Note, however, that when ρ (and thereby w_r) falls, a subset of workers—those with relatively low η —self-select from Routine to Manual tasks. This additional labor supply works against the beneficial effects of computerization on the Manual wage. In fact, it is possible for both w_m and w_r to fall when ρ declines. Nevertheless, the wage of Manual relative to Routine tasks (w_m / w_r) unambiguously rises.

We distinguish, however, between the wage of Routine tasks measured in efficiency units and the observed wages of workers in routine tasks, which are affected by composition. As workers self-select out of Routine tasks, the remaining Routine workers have above average Routine skills, meaning that the observed Routine wage can either rise or fall (both overall and relative to w_m) as ρ declines.

As with Manual workers, workers in Abstract tasks benefit from a rise in Routine task input due to q-complementarity. In the case of Abstract tasks, however, there is no countervailing labor supply response to buffer the positive impact of computerization on the Abstract wage. Consequently, computerization unambiguously raises w_a , both absolutely and relative to w_r and w_m . A more realistic production function would reinforce this pattern by allowing for greater complementarity between Routine and Abstract tasks than between Routine and Manual tasks.

In what sense does this stylized framework represent a polarization of work? The key observation of the model is that computers do not appear to directly substitute for the lowest skilled workers; rather they appear to displace a set of ‘middle skilled’ routine tasks. Displacing this ‘middle’ generates polarization through three mechanisms. First, it directly lowers the wage of middle-skill tasks. Second, it raises the wage of high-skilled (Abstract) tasks through q-complementarity. Finally, it has ambiguous effects on the wages of low-skilled (Manual) tasks due to offsetting impacts of q-complementarity and additional labor supply emanating from workers displaced from Routine tasks. Thus, computerization always raises ‘upper-tail’ inequality in our model—that is the, wage gap between Abstract and Routine tasks. But it can either expand or compress ‘lower-tail’ inequality—that is, the wage gap between Routine and Manual tasks—depending on whether the q-complementarity or labor supply effect dominates.

It is potentially useful to think of the model as moving successively between two equilibria as the price of computer capital declines. In the initial equilibrium, computerization displaces workers from Routine and into Manual tasks, thus depressing wages in both the ‘middle’ and (potentially) the ‘bottom’ of the distribution even as the Abstract wage increases. When ρ falls sufficiently far, however, the model exhibits a second equilibrium in which Routine tasks are performed exclusively by K . In this equilibrium, only the q-complementarity effect is operative;

further declines in ρ raise wages in both Manual and Abstract tasks. Thus, the model is heuristically consistent with an initially ‘monotone’ wage impact of computerization followed by a polarization of the earnings distribution as job tasks performed by workers ‘hollow out’ to encompass only Abstract and Manual tasks.⁸ A similar (complementary) impact of international outsourcing could arise from the declining international communication and coordination costs associated with improvements in information technology (see Levy and Murnane [2004] and Pol Antras, Luis Garicano, and Esteban Rossi-Hansberg [2006]).

IV. Conclusions

In our view, the key new wage structure ‘facts’ documented by recent analyses are the smooth, secular rise in upper-tail inequality over the last twenty five years coupled with an expansion and then compression of lower-tail inequality. These facts are not easily explained by standard institutional stories. For example, the minimum wage explanation for trends in lower-tail inequality only fits the data well to 1987 (Autor, Katz, and Kearney 2005a) and cannot explain why employment in low-wage occupations fell precipitously as the minimum wage dropped (nor why upper-tail inequality has risen steadily). The changing ‘social norms’ explanation for the remarkable skewing of the upper-reaches of the income distribution since the early 1970s, while intriguing, needs to be formulated in more precise manner with clear testable predictions.

⁸ The computerization of relatively low-wage, routine jobs (such as those in lower-wage manufacturing industries and some clerical work) might be easier to initially envision and implement than for other work. And it may be that only with further improvements in software and computer-based technologies does the re-organization of more middle-wage, middle-skill tasks (high-paid precision production work and middle management work) become possible. This process could reinforce a pattern of the substitution for workers by computers moving up the wage (skill) structure over time.

Our analysis offers surprisingly unambiguous evidence that demand shifts are likely to be a key component of any cogent explanation for the changing US wage structure. In particular, we find that quantity and price changes covary positively throughout the entire distribution of earnings—and this is true in the 1980s, when the wage structure was spreading monotonically, and in the 1990s, when the wage structure was polarizing. We believe that a conceptual framework focused on the changing distribution of job task demands, spurred directly by advancing information technology and indirectly by its impact on international outsourcing, goes some distance towards understanding the recent polarization of the U.S. wage structure. The similarity across several countries (the United States, the United Kingdom, and Germany) in employment growth polarization combined with some differences in wage structure changes suggests the need to seriously consider the roles of market forces, institutional factors, and their interactions in explaining the evolution of national wage structures.

References

- Acemoglu, Daron. "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence." *American Economic Review*, December 1999, 89(5), pp. 1259-78.
- _____. "Technical Change, Inequality and the Labor Market." *Journal of Economic Literature*, June 2002, 40(1), pp. 7-72.
- Antras, Pol; Luis Garicano and Esteban Ross-Hansberg. "Offshoring in a Knowledge Economy." *Quarterly Journal of Economics*, February 2006, 121(1), forthcoming.
- Autor, David H.; Lawrence F. Katz, and Melissa S. Kearney. "Trends in U.S. Wage Inequality: Re-Assessing the Revisionists." National Bureau of Economic Research (Cambridge, MA) Working Paper No.11627, September 2005a.
- _____. "Rising Wage Inequality: The Role of Composition and Prices." National Bureau of Economic Research (Cambridge, MA) Working Paper No.11628, September 2005b.
- Autor, David H.; Frank Levy, and Richard J. Murnane. "The Skill Content of Recent Technological Change: An Empirical Investigation." *Quarterly Journal of Economics*, November 2003, 118(4), pp. 1279-1333
- Card, David and John DiNardo. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics*, October 2002, 20(4), pp. 733-83.
- Dew-Becker, Ian and Robert J. Gordon. "Where Did the Productivity Go? Inflation Dynamics and the Distribution of Income." National Bureau of Economic Research (Cambridge, MA) Working Paper No.11842, December 2005.
- Goos, Maarten and Alan Manning. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." Mimeo, London School of Economics, September 2003.
- Juhn, Chinhui. "Wage Inequality and Industrial Change: Evidence from Five Decades." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 4684, March 1994.
- Katz, Lawrence F. and David H. Autor. "Changes in the Wage Structure and Earnings Inequality" in O. Ashenfelter and D. Card, eds. *Handbook of Labor Economics*, volume 3A Amsterdam: North Holland, pp. 1463-1555.
- Lemieux, Thomas. "Post-Secondary Education and Increasing Wage Inequality." University of British Columbia, December 2005.
- _____. "Increased Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *American Economic Review*, 2006, 95, forthcoming.

Levy, Frank and Richard J. Murnane. *The New Division of Labor*. New York: Russell Sage, 2004.

Meyers, Peter B. and Anastasiya M. Osborne. "Proposed Category System for 1960–2000 Census Occupations." U.S. Bureau of Labor Statistics Working Paper No. 383, September 2005.

Mishel, Lawrence; Jared Bernstein and Sylvia Allegretto. *The State of Working America, 2004-05*. Ithaca, NY: ILR Press, 2005.

Pierce, Brooks. "Compensation Inequality." *Quarterly Journal of Economics*, November 2001, 116(4), pp. 1493-1525.

Piketty, Thomas and Emmanuel Saez. "Income Inequality in the United States, 1913-1998." *Quarterly Journal of Economics*, February 2003, 118(1), pp. 1-39.

Spitz-Oener, Alexandra. "Technical Change, Job Tasks and Rising Educational Demands: Looking Outside the Wage Structure." *Journal of Labor Economics*, 2006, 24, forthcoming.

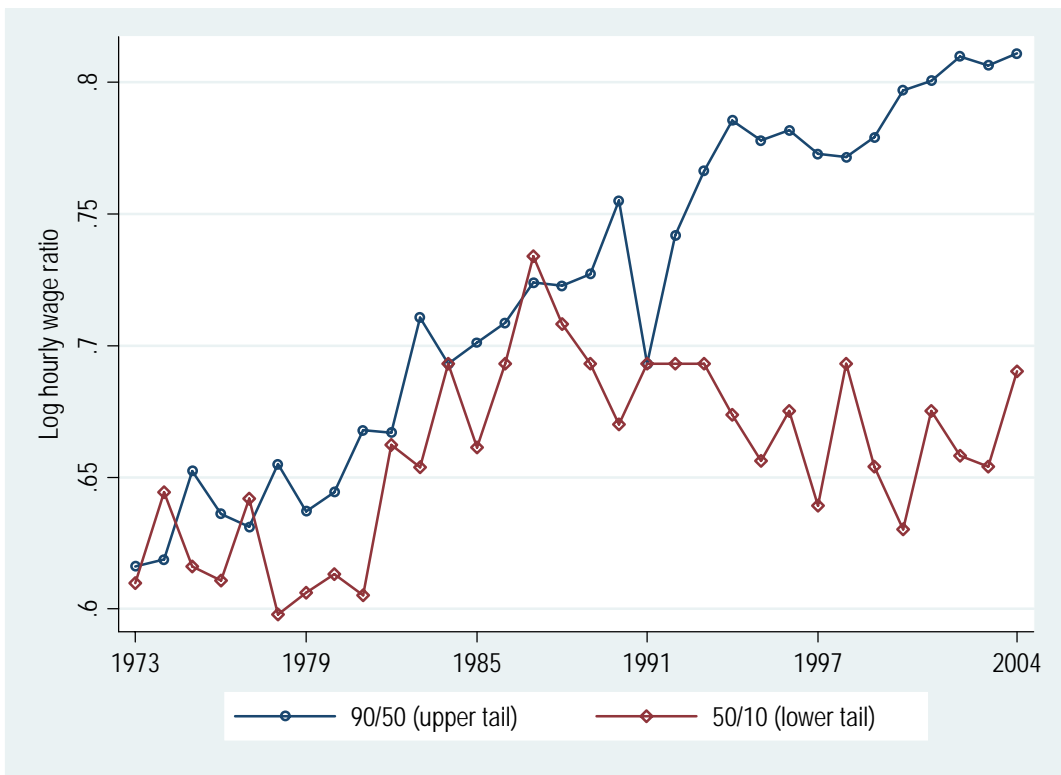


Figure 1. Male and Female Log Hourly 90/50 and 50/10 Earnings Ratios.
 Source: Current Population Survey May and Monthly Files, 1973 - 2004.

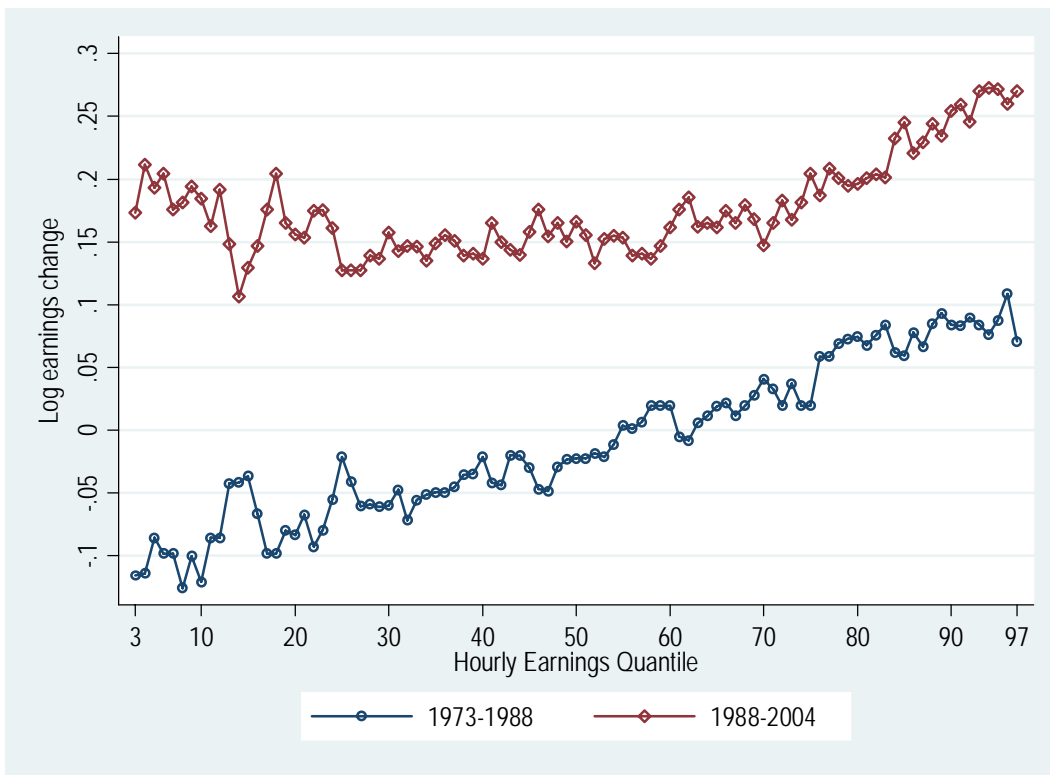


Figure 2. Changes in Male and Female Log Real Hourly Earnings by Percentile 1973 - 1988 and 1988 - 2004. Source: see Figure 1.

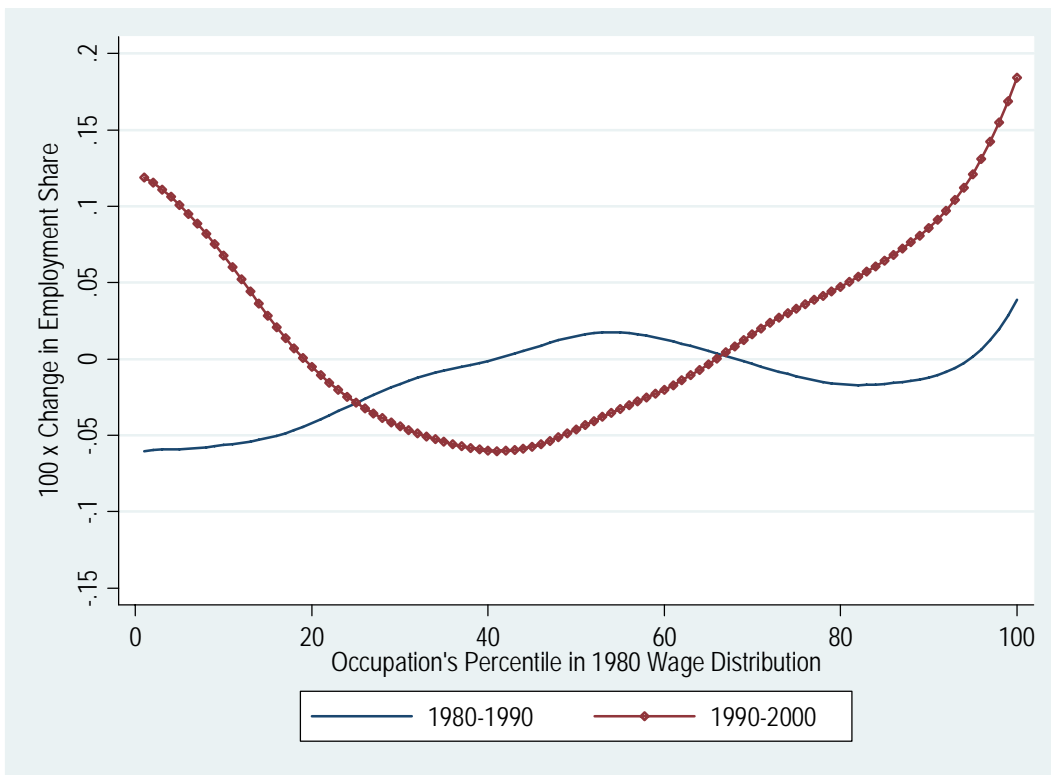


Figure 3. Smoothed Changes in Occupational Employment Shares 1980 - 2000, with occupations ranked by their 1980 Median Wage. Source: Census Integrated Public Use Microsamples, 1980, 1990 and 2000.

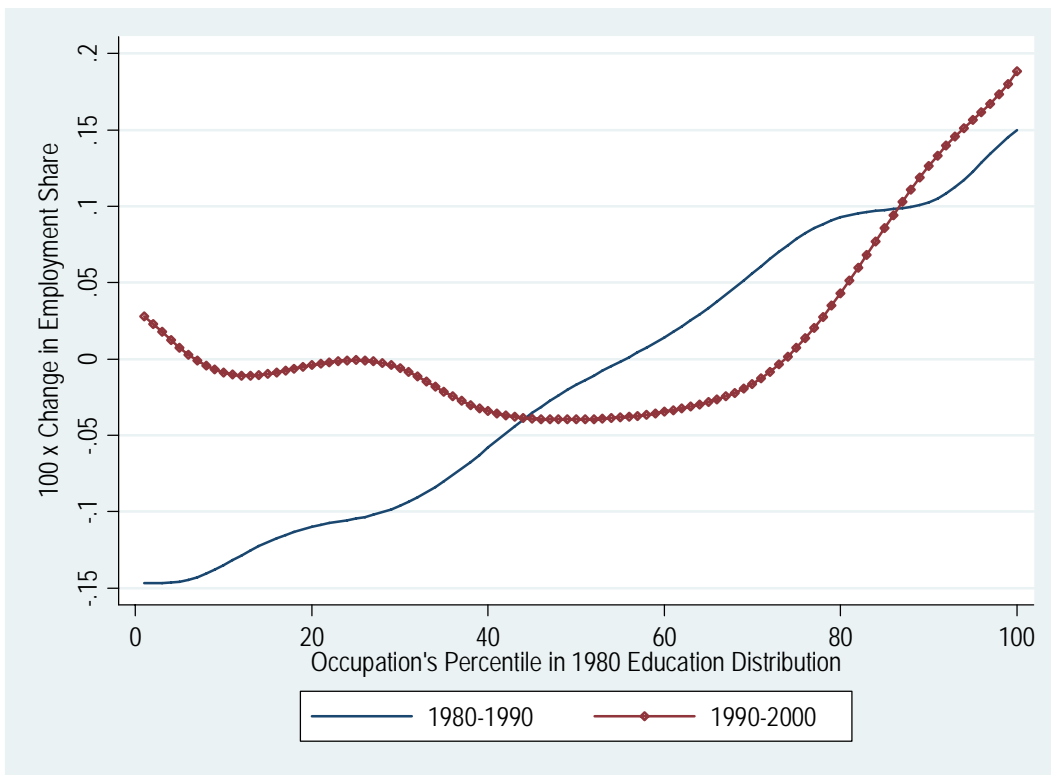


Figure 4. Smoothed Changes in Occupational Employment Shares 1980 - 2000, with occupations ranked by their 1980 Average Years of Schooling.
 Source: see Figure 3.

The Polarization of the U.S. Labor Market: Theory Appendix

David H. Autor
MIT and NBER

Lawrence F. Katz
Harvard University and NBER

Melissa S. Kearney
The Brookings Institution and NBER

December 30, 2005

1 Motivation

Motivating ideas:

1. Workplace tasks may be usually subdivided into three broad groups:
 - (a) Abstract problem-solving and managerial tasks. These tasks are not well structured and require non-routine cognitive skills.
 - (b) Routine tasks. These may either be cognitive or physical tasks that they follow closely prescribed sets of rules and procedures and are executed in a well-controlled environment.
 - (c) Manual tasks. These tasks (which ALM call non-routine manual) do not require abstract problem solving or managerial skills but are nevertheless difficult to automate because they require some flexibility in a less than fully predictable environment. Examples include: truck drivers, security guards, unskilled medical personnel, janitors and house-cleaners, high speed check-keyers, construction workers, many in-person servers.
2. Routine tasks are complementary to Abstract tasks and perhaps also to Manual tasks (though probably less so).
3. Routine tasks are readily substituted with computer capital. The continually falling price of computer capital drives this substitution.
4. The falling price of computer capital generates an incentive for workers engaged in Routine tasks to switch to other tasks.
5. It is more difficult for workers displaced from Routine tasks to shift ‘up’ to Abstract tasks than it is for them to shift ‘down’ to Manual tasks.

2 Model

2.1 Production

Aggregate output in this economy (priced at unity) is given by the Cobb-Douglas production function

$$Y = A^\alpha R^\beta M^\gamma,$$

where A , R and M are Abstract, Routine and Manual tasks, with exponents $\alpha, \beta, \gamma \in (0, 1)$ respectively, and $\alpha + \beta + \gamma = 1$. Abstract and manual tasks can *only* be performed by workers who supply labor inputs, L_A and L_M . Routine tasks can be performed *either* by workers who supply L_R or by computer capital, K , measured in efficiency units, which is a perfect substitute for L_R .

2.2 Factor supplies

Computer capital is supplied perfectly elastically to Routine tasks at price ρ per efficiency unit. The secularly declining price of computer capital is the exogenous driving force in this model.

There is a large number of income-maximizing workers in this economy, each endowed with a vector of three skills, $S_i = (a_i, r_i, m_i)$, where lower-case letters denote an individual's skill endowment for the three production tasks.

Workers are of two types. A fraction $\theta \in (0, 1)$ are High School (H) workers and the remaining $1 - \theta$ are College (C) workers. For simplicity, all college workers are assumed identical. Each is endowed with one efficiency unit of Abstract skill: $S^C(a, r, m) = (1, 0, 0)$.¹ The labor supply of each college worker is $L^C = (1, 0, 0)$.

All high school workers are equally skilled in Manual tasks but differ in their ability in Routine tasks. We write the skill endowment of high school worker i as $S_i^H(a, r, m_i) = (0, \eta_i, 1)$, where η is a continuous variable distributed on the unit interval with positive probability mass at all points $\eta \in (0, 1)$. The labor supply of High School worker i is $L_i^H(a, r, m) = (0, \lambda_i \eta_i, (1 - \lambda_i))$ where $\lambda_i \in [0, 1]$. Each High School worker chooses λ_i to maximize earnings.

2.3 Equilibrium concept

Equilibrium in this model occurs when:

1. Productive efficiency is achieved—that is, the economy operates on the demand curve of the aggregate production function for each factor.
2. All factors are paid their marginal products.
3. The labor market clears; no worker wishes to reallocate labor input among tasks.

2.4 Productive efficiency

The wage of each factor is given by:

$$\begin{aligned} w_a &= \frac{\partial Y}{\partial A} = \alpha A^{\alpha-1} R^\beta M^\gamma, \\ w_r &= \frac{\partial Y}{\partial R} = \beta A^\alpha R^{\beta-1} M^\gamma, \\ w_m &= \frac{\partial Y}{\partial M} = \gamma A^\alpha R^\beta M^{\gamma-1}. \end{aligned}$$

¹It would be a very minor matter to instead assume that $L_C(a, r, m) = (1, 1, 1)$ with $w_a > w_r, w_m$ for all relevant cases. This would ensure that college-workers always supply Abstract labor—and of course this tendency would only be reinforced by a falling price of computer capital.

2.5 Self-selection of workers to tasks

The supply of College labor to Abstract tasks is inelastic.

The supply of High School labor to Routine and Manual tasks is determined by self-selection. Each High School worker i chooses to supply one efficiency unit of labor to Manual tasks if $\eta_i < w_m/w_r$, and supplies η_i efficiency units of labor to Routine tasks otherwise. We can write the labor supply functions to Manual and Routine tasks as $L_M(w_m/w_r) = \theta \sum_i 1[\eta_i < w_m/w_r]$ and $L_R(w_m/w_r) = \theta \sum_i \eta_i \cdot 1[\eta_i \geq w_m/w_r]$, where $1[\cdot]$ is the indicator function. Observe that $L'_M(\cdot) \geq 0$ and $L'_R(\cdot) \leq 0$.

2.6 Equilibrium and comparative statics

Since computer capital is a perfect substitute for routine labor input, it is immediate that $w_r = \rho$ and hence a decline in ρ reduces w_r one for one.²

We are interested in the effect of a decline in ρ on:

1. The equilibrium quantity of Routine task input
2. The allocation of labor between Routine and Manual tasks
3. The wage paid to each task
4. The observed wage in each job type (which in may differ from the wage per efficiency unit in Routine tasks)

A decline in ρ raises demand for Routine tasks, since own-factor demand curves are downward sloping ($R'(\rho) < 0$). This demand can be supplied by either additional computer capital or Routine labor input. Due to worker self-selection, the additional demand will be supplied by computer capital.

To see this, let η^* equal the Manual skill level of the marginal worker, such that $\eta^* = w_m/w_r$. Rewriting η^* using the marginal productivity conditions:

$$\eta^* = \frac{w_m}{w_r} = \frac{\gamma R}{\beta L_M(\eta^*)}.$$

Differentiating with respect to $-\rho$ (a decline in the price of K) gives

$$-\frac{\partial \eta^*}{\partial \rho} = \frac{\gamma}{\beta} \left[\frac{\partial \eta^*}{\partial \rho} \cdot \frac{RL'_M(\eta^*)}{L_M(\eta^*)^2} - \frac{\partial R/\partial \rho}{L_M(\eta^*)} \right] = -\frac{\gamma L_M(\eta^*) \cdot \partial R/\partial \rho}{\beta L_M(\eta^*)^2 + \gamma RL'_M(\eta^*)} > 0. \quad (1)$$

A decline in ρ raises the relative Manual/Routine wage.

²Technically, ρ only binds w_r from above. In an earlier period, when computer capital was far more expensive, it's plausible that $w_r < \rho$. For the period under study, we assume that this constraint binds.

Summing up these wage implications:

$$\begin{aligned}
-\frac{\partial w_r}{\partial \rho} &= -1, \\
-\frac{\partial w_m}{\partial \rho} &= -\gamma A^\alpha \left[\underbrace{\beta L_M^{\gamma-1} R^{\beta-1} \frac{\partial R}{\partial \rho}}_{(-)} + \underbrace{(\gamma-1) R^\beta L_M (\eta^*)^{\gamma-2} L'_M (\eta^*) \frac{\partial \eta^*}{\partial \rho}}_{(+/0)} \right] \leq 0, \\
-\frac{\partial w_a}{\partial \rho} &= -\alpha A^{\alpha-1} \left[\underbrace{\beta R^{\beta-1} \frac{\partial R}{\partial \rho}}_{(-)} + \underbrace{\gamma L_M (\eta^*)^{\gamma-1} L'_M (\eta^*) \frac{\partial \eta^*}{\partial \rho}}_{(-)} \right] > 0.
\end{aligned}$$

A decline in the price of computer capital lowers the wage of Routine labor input, raises the wage of Abstract labor input through two channels of q-complementarity—increased use of Routine task input and increased labor supply to Manual task input—and has ambiguous implications for the wage of Manual task input (due to the countervailing effects of q-complementarity between Routine and Manual tasks and increased labor supply to Manual tasks).

Though, as established above, a decline in ρ yields a larger proportionate fall in the wage of Routine than Manual tasks ($-\partial(w_m/w_r)/\partial\rho > 0$), the *observed* log wage differential between workers in Routine and Manual jobs may *rise* despite the fall in w_r/w_m . The reason is that a decline in w_r leads to marginal workers with lower values of η to exit Routine jobs, inducing a *positive* compositional shift in the pool of workers in Routine occupations ($-\partial E[\eta|\eta > \eta^*]/\partial\rho > 0$).

Summarizing:

1. A decline in the price of computer capital causes an increase in demand for Routine task input.
2. This increase is entirely supplied by computer capital as the price decline causes a corresponding reduction in labor supply to Routine tasks and an increase in labor supply to Manual tasks.
3. The reduction in the price of computer capital has the following implications for wages levels measured in *efficiency units*:

$$-\frac{\partial w_r}{\partial \rho} < 0, -\frac{\partial w_m}{\partial \rho} \leq 0, -\frac{\partial w_r/w_m}{\partial \rho} < 0, -\frac{\partial w_a}{\partial \rho} > 0.$$

4. The reduction in the price of computer capital has the following implications for *observed* wage levels:

$$-\frac{\partial \hat{w}_r}{\partial \rho} \leq 0, -\frac{\partial \hat{w}_m}{\partial \rho} \leq 0, -\frac{\partial \hat{w}_r/w_m}{\partial \rho} \leq 0, -\frac{\partial \hat{w}_a}{\partial \rho} > 0,$$

where ‘hats’ over wage variables denote observed values that do not adjust for changes in occupational skill composition (e.g., $\partial E[\eta|\eta > \eta^*]/\partial\rho$).³

³For Manual and Abstract tasks, compositional shifts are nil by assumption.

Remark 1: The model does not pin down the ranking of wages in Abstract, Routine and Manual tasks; these levels depend on labor supplies and ρ . For workers who switch from Routine to Manual tasks as ρ falls, the manual wage must be higher than the Routine wage ($w_m > \eta_i w_r$). But for inframarginal Routine workers, it must be the case that the Routine wage is higher than the Manual wage. This observation has an important empirical implication: if there are *any* workers remaining in the Routine job, their *observed* wage (i.e., not accounting for composition) must be higher than the wage in the Manual job since there is no skill heterogeneity in Manual tasks. Hence, even in cases where $w_r < w_m$, it will be true that $\hat{w}_r > \hat{w}_m$ provided that $L_R(w_m/w_r) > 0$.

Remark 2: In an equilibrium in which H workers supply both Routine tasks and Manual tasks, K is a direct substitute for some H workers and a complement to others. In this setting, a decline in ρ causes a ‘widening’ of wage inequality by lowering w_m relative to w_a (moreover, w_m may fall in absolute terms). In an equilibrium with ρ sufficiently low such that no workers remain in Routine tasks ($w_m/w_r > 1$), further declines in ρ unambiguously benefit both High School and College workers and hence do not augment inequality. (Given the Cobb-Douglas form, both groups benefit equally, so this has no effect on w_a/w_m .)