

The Changing Nature of Wage Inequality*

Thomas Lemieux
Department of Economics
University of British Columbia
#997-1873 East Mall
Vancouver, BC, V6T 1Z1
Canada
Email: tlemieux2interchange.ubc.ca
Fax: (1) 604-822-5915

ABSTRACT

The paper reviews recent developments in the literature on wage inequality, with a particular focus on why inequality growth has been particularly concentrated in the top end of the wage distribution over the last 15 years. Several possible institutional and demand-side explanations are discussed for the secular growth in wage inequality in the United States and other advanced industrialized countries.

JEL Classification: J24, J31, J51

Keywords: Inequality, institutions, technological change

*: A preliminary version of this paper was presented at the 20th Annual Conference of the European Society for Population Economics held in Verona in June 2006. I would like to thank Christian Dustmann and three anonymous referees for their comments, and the Social Sciences and Humanities Research Council of Canada (through TARGET) for research support.

1. Introduction

This paper reviews recent changes in the nature of wage inequality, and discusses proposed explanations for these changes. Most of the discussion focuses on the case of the United States that has been studied in much more detail than other countries. I first discuss the consensus that emerged on the sources of changes in inequality in the 1980s in Section 2, but raise several difficulties that were already apparent at the time. In Section 3, I discuss in detail the more recent (since 1990) changes in inequality and explain how they pose further challenges for the explanations that were suggested for the 1980s. I also present some evidence from the Current Population Survey (CPS) to complement the discussion.

The main message of Section 3 is that the nature of the changes in inequality has been dramatically altered over the last fifteen years. While the growth in inequality in the 1980s was pervasive, it has been concentrated at the top end of the distribution since then. Sections 4 and 5 review possible explanations for these changes. Section 4 discusses the role of institutional factors, while Section 5 focuses on explanations based on demand factors and, in particular, technological change. Finally, I conclude in Section 6.

2. The 1990s consensus

A representative view of inequality in 1980 is Blinder (1980, p. 416) who remarked that “...when we turn to consider the distribution of economic welfare, --economic equality, as it is commonly called, the stylized fact is one of constancy.” Blinder’s view was based on the latest data available at the time he wrote his paper that showed that “income inequality was just about the same in 1977... as in 1947.” In hindsight, we now realize that as Blinder and other observers at the time correctly saw a remarkably constant income distribution, the U.S. economy had entered a period of sharp increase in income and earnings inequality.

Throughout the 1980s, a number of researchers started noticing a definite trend in increasing inequality. For instance, Bluestone and Harrison (1988) pointed out that the share of “low-wage” jobs had increased sharply in the first half of the 1980s. This generated some controversy, in part because it was hard to tell at the time whether this

change was simply a consequence of the deep recession of 1981-1982. Even in the second half of the 1980s, a number of papers concluded that recent changes in inequality were, at best, modest. For example, using data up to 1985, Blackburn and Bloom (1987) concluded that “The time profile of earnings inequality, measured across individual workers, has been quite flat since the late 1960s”.

As interest grew for the issue of income and earnings inequality, the controversy of the mid- to late-1980s was rapidly resolved as new data clearly showed a sharp and continuing growth in inequality. By the early 1990s, a set of highly influential papers by Bound and Johnson (1992), Katz and Murphy (1992), Levy and Murnane (1992), and Juhn, Murphy and Pierce (1993) established a wide consensus that *i*) inequality had been growing sharply in the 1980s, and *ii*) the primary factor behind the growth was the increase in the relative demand for skill. The consensus was also that the relative demand for skill had been growing in the 1970s, but had been outstripped by the steep growth in educational attainment --the relative supply for skills—linked to the entry of the highly-educated baby boom generation in the labor market (Freeman, 1976). For instance, Juhn, Murphy and Pierce (1993) argue that the within-group dimension of inequality had been growing throughout the 1970s and 1980s, thereby providing evidence that the relative demand for skills was already increasing in the 1970s.

While this first set of papers did not take a strong position on the source of the increase in the relative demand for skills, it became widely accepted during the early 1990s that skill-biased technical change (SBTC), driven by the computer revolution, was the primary source of growth in the relative demand for skill (Krueger, 1993, Bound, Berman, and Griliches, 1994). The leading alternative explanation, international trade or globalization, was generally rejected as the main source of the increase in the relative demand for skill.¹

¹ Berman, Bound, and Griliches (1994) dismiss the trade explanation on the basis that most of the increase in the use of skilled labor happens within industries, while a traditional Heckscher-Ohlin would predict that changes should happen between industries. More recently, however, authors such as Feenstra and Hanson (2003) have pointed out that trade in intermediate inputs, which can lead to within-industry changes, was a leading alternative explanation to SBTC for the increase in the relative demand for skill. Their focus on intermediate, as opposed to final goods and services, is consistent with popular perceptions that off-shoring is now the key way in which globalization affects domestic employment and inequality.

One last dimension of the 1990s consensus was that the increase in the relative demand for skill was pervasive, or ubiquitous, in the sense that all dimensions of inequality were growing. For example, in a simple human capital model, wage inequality can increase because returns to education and experience increase, or because residual or within-group inequality increases. Juhn, Murphy and Pierce (1993) show that all of these dimensions of inequality had been growing in the 1980s. They attributed this to a pervasive increase in the relative demand for all dimensions of skill (education, experience, unobserved ability, etc.), that was later interpreted as being driven by SBTC.

The consensus reached in the early 1990s about the importance of technological change in the growth in inequality has remained central in most of the subsequent research on inequality. For instance, ten years after Levy and Murnane (1992) was published in the *Journal of Economic Literature (JEL)*, the *JEL* published another survey (Acemoglu, 2002) on inequality. Acemoglu (2002) again concludes that technological change, or at least a more sophisticated form of endogenous technological change, was the leading explanation for inequality growth in the United States throughout the 1970s, 1980s and 1990s.

Even back in the 1990s, however, it was already clear that a number of facts were difficult to reconcile with the prevailing view that SBTC was driving the growth in inequality in the United States. A first major challenge to the 1990s consensus was that other large advanced economies like France, Japan, or Germany, had failed to experience any significant growth in inequality during the 1980s. If technological change is the explanation for growing inequality, how can it be that other advanced economies subject to the same technological change do not experience an increase in inequality? The response to this question suggested by Freeman and Katz (1995) was that supply and demand (SBTC) were only one part of the explanation for changes in inequality. The other part of the explanation was wage-setting institutions. Freeman and Katz (1995) suggest a richer SDI (supply, demand, institutions) explanation for diverging patterns of inequality across countries where common demand shock across countries are mitigated by supply and institutional factors.

For example, in “Anglo-Saxon” countries like the United States, Britain, or Canada where the wage-setting is decentralized and unions are weak, a negative demand

change depresses the wages of less-skilled workers. By contrast, in France or Germany where the wage-setting is more centralized and unions are strong, wages of less-skilled workers remain fixed despite the same negative demand change, and so does wage inequality.

Another set of papers shows that, even in the United States, wage-setting institutions appear to play a substantial role in the growth in inequality. For instance, both Card (1992) and Freeman (1993) conclude that the decline in unionization accounts for about 20 percent of the increase in the variance of log wages for U.S. males in the 1980s. Another important wage-setting institution is the minimum wage, which fell sharply (in real terms) during the 1980s. DiNardo, Fortin, and Lemieux (1996) show that the decline in the real value of the minimum wage had a very large impact on wage inequality among women, and a smaller but still substantial impact among men. This finding was later confirmed by Lee (1999) who concludes that all of the increase in inequality in the lower end of the wage distribution during the 1980s was due to the decline in the real value of the minimum wage. Taken together, these results are a serious challenge to the view that, even in the United States, inequality growth in the 1980s was due to a pervasive increase in the relative demand for skill, or SBTC. Perhaps what happened in the low- and high-end of the wage distribution in the 1980s was instead a coincidence. This alternative view is that different factors affecting inequality at different points of the distribution all happened to move in the same direction during the 1980s, giving the misleading impression that a single factor like SBTC was the sole driving force behind the increase in wage inequality.

A final, but more technical, point is that, even in the mid-1990s, one main stylized fact about inequality growth was known not to be robust to the choice of data. As mentioned above, Juhn, Murphy, and Pierce (1993) argue that within-group inequality had been growing in the 1970s, thereby providing evidence that the relative demand for skills was already increasing in the 1970s, and that the growth in inequality in the 1980s was not just an “episodic” phenomenon. This finding was based on data on weekly earnings of full-time male workers from the March supplement of the Current Population Survey (CPS). Using data on hourly wages from the May and Outgoing Rotation Group (ORG) Supplements of the CPS instead, both Card and Lemieux (1996) and DiNardo,

Fortin, and Lemieux (1996) find that within-group wage inequality did not increase for men (and declined for women) during the 1970s. Though this finding was the subject of controversy for some time, it is now widely accepted that the growth in within-group inequality in the 1970s is not robust to the choice of data.² By contrast, the growth in within-group inequality in the 1980s and the evolution of the college-high school wage gap are very similar in the March and May-ORG CPS.

In summary, even at the time when SBTC emerged as the leading explanation for the growth in inequality, the explanation was already known to have several important shortcomings. First, it could not account for the diverging patterns of inequality growth across advanced countries. Second, it failed to account for the fact that institutional factors were also an important part of the story, especially in the lower end of the wage distribution. Third, it relied in part on the unrobust fact that within-group inequality was already increasing in the 1970s.

3. Recent evidence and further challenges to the 1990s consensus.

a. Recent Studies

More recently, a number of further challenges to this “1990s consensus” have emerged. A first challenge has to do with the timing of the growth in wage inequality. In particular, Card and DiNardo (2002) and Beaudry and Green (2005) argue that much of the increase in the return to education was concentrated in the 1980s. In a similar vein, Lemieux (2006b) shows that the growth in residual wage inequality also appears to be concentrated in the 1980s once composition effects are controlled for. This evidence is difficult to reconcile with a simple SBTC story that would likely predict a continuing and steep growth in wage inequality throughout the 1990s. For instance, Card and DiNardo

² Katz and Autor (1999) re-examined the May-ORG data and concluded that within-group inequality had increased in the 1970s. Based on this evidence and the earlier findings from the March CPS data by Juhn, Murphy, and Pierce (1993), Acemoglu (2002) concluded that “...there is considerable evidence that residual and overall inequality started to increase during the 1970s”, while correctly noting that “there is less uniformity among data sources regarding the behavior of residual inequality than the returns to schooling.” (Acemoglu, 2002, page 65). This finding was based, however, on a comparison between May 1973 data where observations with missing wages were dropped from the sample, and May 1979 data where missing wages were replaced by their allocated values. Both Lemieux (2006b) and Autor, Katz and Kearney (2005) have now confirmed that within-group inequality does not increase in the 1970s when allocated wages are dropped in 1979 too (as in Card and Lemieux, 1996, or DiNardo, Fortin, and Lemieux, 1996).

(2002) argue that “IT-related technological change has been going on since at least the 1970s and has continued throughout the 1980s and 1990s. Moreover, there is some evidence ... that the rate of technological change accelerated in the 1990s, relative to the 1980s.” (Card and DiNardo, 2002, page 740). Indeed, casual evidence suggests that the internet revolution of the 1990s was at least as important, from the point of view of its impact on the workplace, as the introduction of personal computers in the 1980s.

Mounting evidence also suggests that, far from being “ubiquitous”, the growth in wage inequality is increasingly concentrated in the top end of the wage distribution. For example, Mincer (1997) and Deschênes (2002) show that (log) wages are an increasingly convex function of years of education. In other words, the wage gap between college post-graduates and college graduates has increased more than the wage gap between college graduates and high school graduates, which, in turn, increased more than the wage gap between high school graduates and high school dropouts. Looking more broadly at the distribution of taxable earnings, Piketty and Saez (2003) also find that relative wage gains are disproportionately concentrated in the very top of the earnings distribution. In more recent work, Piketty and Saez (2006) also show that the 1980s puzzle of why some advanced countries experienced a sharp increase in inequality while others did not has persisted over time. Once again, English-speaking countries (United States, Canada and Britain) experienced a steep growth in inequality, as measured by the share of income or earnings going to the top 10 or 1 percent of the distribution. By contrast, these shares have remained remarkably stable in France and Japan over the post-war period.

Changes in residual inequality also appear to be concentrated at the top end. For instance, Lemieux (2006b) shows that within-group inequality grew substantially among college-educated workers but changed little for most other groups. A related finding by Autor, Katz, and Kearney (2005) is that “top end” residual inequality (e.g. the difference between the 90th and 50th percentile of the distribution of residuals, or the “90-50” gap) increased substantially while residual inequality at the low end (the 50-10 gap) actually decreased.

b. Direct evidence from the Current Population Survey.

I now present some results from the CPS to better illustrate how the main inequality trends have changed over time. Figures 1 to 4 and Table 1 summarize the recent developments in the wage distribution and contrast them to earlier changes from the 1970s and 1980s. The table and figures are based on data on hourly wages from the May 1973-78 and ORG 1979-2005 supplements of the CPS. The data files are processed as in Lemieux (2006b) who provides detailed information on the relevant data issues.

Table 1 shows the evolution of broad measures of wage inequality, like the standard deviation and the 90-10 gap in log wages, as well as wage differentials across education groups. These education wage differentials are computed for each single year of potential labor market experience, and then averaged out over all experience groups using the average shares of experience over the whole 1973-2005 period.³ The measures of inequality are computed separately for men and women at five different points in time: 1973-75, 1978-80, 1988-90, 1998-2000, and 2003-2005. The last two columns show the changes in the different measures of wage dispersion over the first (1973-75 to 1988-90) and second half (1988-90 to 2003-05) of the sample period.

A strong pattern of results emerge from these simple comparisons. First, the change in broad measures of inequality (standard deviations and 90-10 gap) is three times larger in the 1970s and 1980s than in the 1990s and 2000s. Furthermore, looking more closely at the table indicates that all of the growth in inequality prior to 1988-90 is actually concentrated in the 1980s. The pattern holds for both men and women.

A second interesting pattern is that while both “low-end” (50-10 gap) and “top-end” (90-50 gap) inequality were growing prior to 1988-90, the situation changes radically in the 1990s and 2000s. The 50-10 gap stagnates for women and actually declines for men. By contrast, the 90-50 gap keeps growing as fast in the 1990s and 2000s as in the earlier period.

Education wage differentials (all measured relative to high school graduates) exhibit essentially the same pattern as overall measures of wage inequality. Generally speaking, there is a steep slowdown in the growth in these differentials after 1988-90, as

³ This is a more general way of adjusting for experience than in a standard Mincer-type regression where a low level polynomial in experience is included in the regression. The approach used here is equivalent to running separate regressions for each year with a full set of experience and education (five groups) dummies, and then averaging out (using a fixed weighted average across experience groups) the returns to education within each experience cell.

in the case of the standard deviation or the 90-10 gap. This general pattern hides important differences, however, across education groups. At the low end of the education distribution, the gap between high school dropouts and high school graduates expands by 0.045 and 0.085 in the 1970s and 1980s for men and women, respectively. As in the case of the 50-10 gap, however, the growth completely stops, and even reverts itself in the case of men after 1988-90. At the other end of the education spectrum, the gap between college post-graduates and high school graduates keeps growing as fast after as before 1988-90. The gap between college graduates (with exactly 16 years of education) and high school graduates falls somewhat in the middle, as the growth in the gap slows down by about a half after 1988-90.

The precise timing of these changes is illustrated in more detail in Figures 1 to 3. Figure 1 shows that, for men, both the 90-50 and the 50-10 gap were more or less stable in the 1970s, grew steeply in the 1980s, and diverged sharply after the 1980s. While the 50-10 gap returned to its level of the early 1970s by the mid-1990s (and stayed more or less constant since then), the 90-50 gap kept growing steadily and is now 40 percent larger than in 1973.

Figure 2 shows that, as in the case of men, the 90-50 gap for women grew steadily over most of the sample period, and is now 40 percent larger than back in 1973. The story for the 50-10 gap is different, however. The 50-10 gap declined sharply in the 1970s before increasing even more dramatically between 1979 and 1986. Since then, the 50-10 gap has remained more or less stable around 20 to 30 percent above its 1973 level.

Figure 3 shows the evolution of the various education wage differentials, for men, relative to their 1973 level. As in the case of the broader measures of inequality, all education wage gaps are stable in the 1970s, increase in the early 1980s, and diverge substantially after the mid-1980s. As a result, by 2005 the gap between high school dropouts and high school graduates is only 20 percent above its 1973 level. By contrast, the gap between college post-graduates and high school graduates in 2005 is more than twice as large as in 1973. Under the assumption that a post-graduate degree requires, on average, 18 years of education, the 0.36 wage gap between the two groups (Table 1) in 1973-75 corresponds to a yearly return to education of 6 percent, compared to 13 percent (0.77 wage gap) in 2003-2005. Lemieux (2006a) shows that this enormous increase in

the return to “post-secondary” education accounts for a large share of the overall increase in inequality.

In summary, the evidence from the CPS data strongly supports the findings of the aforementioned studies that the nature of the growth in wage inequality has changed very substantially over the last 15 years. All the recent growth in inequality is concentrated at the top end of the distribution, in sharp contrast with the situation that prevailed prior to 1990. Figure 4 illustrates this dramatic change by showing the growth in real male wages at each percentile of the wage distribution over the two sub-periods. As in Juhn, Murphy and Pierce (1993), the increase in inequality prior to 1989 is pervasive, or ubiquitous, in the sense that wage differentials increase at essentially all points of the distribution. Interestingly, however, the resulting curve obtained using the May-ORG data is not quite as linear as what Juhn, Murphy and Pierce (1993) found using the March CPS. In fact, the curve for 1974 to 1989 is already convex in the sense that wage differentials in the low end are not expanding quite as fast as in the top end.

The situation changes radically after 1989. First, real wages increase at each and every points of the distribution, unlike prior to 1989 when real wages decreased except for the top 20 percent of the distribution. More importantly, the curve is U-shaped as workers both in the lower and upper tails of the distribution experienced larger gains than workers in the middle of the distribution.

4. Wage-setting institutions and the growth in top-end wage inequality

Explaining the continuing growth in top-end inequality in the United States over the last 15 years, in a context where inequality elsewhere in the distribution is not increasing, is now attracting a lot of attention. In this section and the next, I discuss the main emerging explanations that have been suggested for these changes. I stress, however, that these explanations have to be evaluated in a broader comparative context where two important factors have to be kept in mind. First, the nature of the changes in inequality prior to 1990 was very different from the more recent period. A “good” explanation the post-1990 situation should also be able to account for the very different situation that prevailed prior to 1990. Second, a “good” explanation should also have something to say, if

possible, about why the evolution of wage inequality has been so different across different advanced countries.

In this section, I explore a number of possible explanations for the growth in top-end inequality that revolve around the theme of wage setting institutions. I then explore more standard demand and supply explanations in Section 5.

a. Wage setting institutions and social norms: top executives

The work by Thomas Piketty, Emmanuel Saez and their co-authors has been very influential in attracting attention to the issue of the growing inequality at the top end of the distribution.⁴ The most important contribution of this branch of the literature is to provide good quality and comparable tax data for several countries over a long period of time. As mentioned above, the main finding is that top-end inequality started increasing steadily in the 1980s in English-speaking countries, but remained relatively unchanged in other advanced countries like Japan and France.

Piketty and Saez (2006) discuss several possible explanations for the growth in top end inequality, focusing on the question of why the incomes of top end executives increased so much over time. One first explanation that they end up dismissing is that technological change has made managerial skills more general and less firm-specific, thus increasing the global competition for the best executives. The problem with this explanation is that it, presumably, represents a world wide phenomenon that fails to explain why top-end earnings increased much more in some countries than in others.

The second, and more promising, explanation explored by Piketty and Saez is that changes in pay-setting institutions and social norms have removed some implicit barriers to higher wages in the United States, but not in other countries like France where pay-setting institutions such as union power remained constant. The third, and related, explanation is that executives got better at setting their own pay or extracting rents at the expense of shareholders.

While all these explanations are plausible, they are difficult to test for a number of reasons. First, it is not clear how to measure the social norms and general pay-setting

⁴ See, for example, Piketty (2003) for France, Piketty and Saez (2003) for the United States, Saez and Veall (2005) for Canada, as well as the summary piece by Piketty and Saez (2006).

institutions that affect the pay of high executives, or to measure the capacity of executives to set their pay or extract rents. Second, the major disadvantage of tax data is that they are not based on detailed micro-data that can then be used to test for interesting hypotheses. Other data sources have to be used to make progress on these issues.

For example, Bebchuk and Grinstein (2005) use detailed data from the ExecuComp database for 1993 to 2003 to look explicitly at explanations for the large increase in executive pay over this period. Using standard regression analysis, they find that changes in firm size, performance and industry classification fail to explain most of the growth in executive pay over this period. They conclude that mean compensation in 2003 would have only been half of its actual size had the relationship between compensation and firm size, performance, and industry classification remained the same as in 1993. The problem is that though large (or more performing) firms pay their executives more than smaller firms, the elasticity is not nearly large enough for the growth in firm size to account for the growth in executive pay.

Gabaix and Landier (2008) reach the very different conclusion that “The six-fold increase of CEO pay between 1980 and 2003 can be fully attributed to the six-fold increase in market capitalization of large U.S. companies during that period.” In other words, they find that the elasticity of pay with respect to firm size (market capitalization in their paper) is roughly equal to one, which is much larger than the elasticity estimated by Bebchuk and Grinstein (2005). Their estimates are mostly based, however, on the time-series relationship between executive pay and market capitalization, a questionable identification strategy.⁵

b. Wage setting institutions for the broader workforce: unions and minimum wages

Although the focus on top executives can help explain some of the growth in earnings at the very top of the distribution, executives are by no means the only workers driving earnings growth at the top end. For instance, it was recently reported that Goldman and

⁵ Gabaix and Lanthier (2008) report a more standard elasticity of compensation with respect to market capitalization of .3 in a standard panel data regression, which is more or less comparable to Bebchuk and Grinstein (2005). They include an aggregate measure of market capitalization instead of standard year dummies, however, and add the effect of that measure (.7) to get a “full” elasticity of about 1. The problem, however, is that this purely time-series variable could also capture other proposed explanations (social norms, etc.) for the change in executive pay. Without this questionable approach for estimating the elasticity, their results would be very much in line with those of Bebchuk and Grinstein (2005).

Sachs average compensation for 2006 reached over \$600,000.⁶ Literally thousands of traders, investment bankers, and managers other than the top executives of this single firm (those included in executive data bases) likely earned enough to put them in the top 0.1 percent of the earnings distribution.⁷ One must thus clearly broaden the analysis beyond top executives to understand the sources of growing inequality at the very top end of the distribution.

One major limitation of most micro data sets for understanding top end inequality is top coding. For example, in the March CPS, reported wages and salaries were until recently top-coded at \$150,000 a year, which is barely above the 95th percentile of the distribution of earnings in the tax data of Piketty and Saez (\$125,471 in 2004).⁸ One well known data set for which top-coding is not an issue is the Panel Study of Income Dynamics (PSID), which is unfortunately not ideal either for studying top-end inequality because of smaller sample sizes.⁹ But though these data sets are not very useful for studying changes in inequality at the very top end, they remain extremely useful for understanding the sources of wage movements up to, at least, the 95th percentile of the wage distribution. In the remainder of the paper I will thus focus on analyses of inequality based on standard data sets such as the CPS.

As mentioned above, Piketty and Saez (2003, 2006) argue that social norms and institutional factors are a promising explanation for why top end inequality increased much more in the United States than in many other countries. Since social norms and other broad institutional factors presumably affect the whole labor market in a country, it is not clear how one can successfully use micro data to test for the effect of these factors on the wage distribution. By contrast, the more standard labor economics literature

⁶ International Herald Tribune, Dec. 13, 2006.

⁷ The updated numbers of the Piketty and Saez (2003) indicate that the 99.9th centile of the salary distribution was \$837,892 in 2004 (data accessed on Emmanuel Saez's website at: <http://elsa.berkeley.edu/~saez/TabFig2004.xls>). See also Kaplan and Rauh (2007) for some estimates of the contribution of high earners other than CEOs to the top end of the distribution.

⁸ This is the top-code on annual earnings on the main job in the March CPS that prevailed until 2002. The top-code has been revised up to \$200,000 for 2002 earnings which is still, however, below the 99th percentile of the distribution of earnings in the tax data of Piketty and Saez (\$251,000 in 2004). The top-code on earnings in the ORG supplement of the CPS is also \$150,000 a year for individuals who have worked for all 52 weeks of the year (top-code of \$2884.61 a week), while the top-code in the 2000 census is only slightly larger (\$175,000 a year) than in the pre-2002 March CPS.

⁹ The top-code on annual earnings in the PSID was \$99,999 until 1982, which was well within the top one percent of the earnings distribution back then. It was subsequently increased to very large values that only bind for a handful of individuals.

focuses on more narrowly defined institutions such as unions, minimum wages, and deregulation that affect some workers but not others. As a result, the importance of these factors in the recent evolution of wage inequality can be readily tested for using micro data.

One institutional factor that can safely be ruled out as an explanation for the continuing growth in top end inequality is the minimum wage. As mentioned in Section 2, there is strong evidence that the decline in the real value of the minimum wage played a major role in the increase in low-end inequality in the 1980s. Figure 5 also shows the remarkable coincidence between the evolution of the minimum wage and the 50-10 gap for women from 1973 to 2005. DiNardo, Fortin, and Lemieux (1996) and Lee (1999) only look at data for up to about 1990, and conclude that most of the growth in the 50-10 gap shown in the figure was due to the decline in the minimum wage. Figure 5 shows that out-of-sample predictions from these studies are highly consistent with the post-1990 data, since both the minimum wage and the 50-10 gap remained more or less constant over that period. Interestingly, both series still move together around this relatively flat trend for 1990 to 2005. But since the minimum wage did not change much after 1990 and unlikely affects top-end inequality anyway, it simply cannot account for the growth in top-end inequality over the last 15 years.

A second institutional factor that unlikely accounts for recent changes in top end inequality after 1990 is deregulation. Fortin and Lemieux (1997) show that industries that were deregulated in the late 1970s (airlines, trucking, etc.) did not experience more growth in inequality than industries that were unaffected by deregulation. Given the lack of effects of deregulation on inequality, and the fact that deregulation mostly happened 25 to 30 years ago, it is not a plausible explanation for recent changes in top-end inequality.

A more promising explanation, at least in the case of men, is de-unionization. As mentioned in Section 2, both Card (1992) and Freeman (1993) show that de-unionization accounts for a sizable share (up to 20 percent) of the increase in the variance of wages in the 1980s. DiNardo, Fortin, and Lemieux (1996) reach a similar conclusion but also explore in more detail the distributional impacts of unions on wages. Interestingly, they find that the 10.8 percentage point decline in male unionization accounts for 0.040, or a

third, of the 0.119 increase in the 90-50 gap between 1979 and 1988. By contrast, they find that de-unionization has a negative effect on the 50-10 gap, reducing it by 0.019 during a period where the gap expanded by 0.076. In other words, the de-unionization story was going the wrong way at the low end during the 1980s, though institutions as a whole accounted for some of the increase in low-end inequality because of the contribution of the declining minimum wage that accounted for 0.050 of the 0.076 increase in the 50-10 gap. The results from DiNardo, Fortin and Lemieux (1996) thus imply that the negative effect of de-unionization on low-end inequality was more than offset by the large effect of the declining minimum wage over this period.

Since 1988, however, the rate of male unionization kept decreasing (by 8.6 percentage points between 1988 and 2005 in the ORG CPS) while the minimum wage remained more or less constant (Figure 5). Extrapolating the DiNardo, Fortin, and Lemieux (1996) results to the 1988-2005 period suggests that de-unionization can both explain some of the increase in top-end inequality over this period, as well as some of the decline in low-end inequality. Qualitatively speaking, de-unionization works very well as an explanation for the changing nature of inequality since 1990 as in can both account for increasing inequality at the top end, and decreasing inequality at the low end. From a quantitative point of view, however, the question is whether the effect of unionization is large enough to account for the magnitude of the recent changes in inequality?

To answer this type of question, one could simply redo a DiNardo, Fortin, and Lemieux type analysis for the 1988-2005 period. A simpler “back-of-the-envelope” calculation could also be performed if one knew what is the effect of unions on wages at different quantiles of the wage distribution. As it turns out, Firpo, Fortin and Lemieux (2007) recently proposed a new unconditional quantile regression method that precisely yields these types of estimates. They also apply their method to the case of unions for men in 1983-85, controlling for standard factors such as experience and education. Their estimated effects are reproduced in Figure 6. Since it well known since Freeman (1980) that unions tend to reduce wage dispersion among male workers, it is not surprising that the union effects in Figure 6 are larger and the low end (say 10th percentile) than at the high end (say 90th percentile) of the wage distribution. The standard “within” effect of

unions can account for this pattern of results.¹⁰ More interestingly, the effect is non-monotonic. It first increases to reach the 0.30 to 0.40 range between the 20th and the 50th percentiles, declines steadily after the 50th percentile, and eventually reaches negative values above the 80th percentile. The effect is very large and negative (up to -0.30) by the time we reach the highest percentile shown in the figure (98th percentile).

Visually speaking, the effect of unions plotted in Figure 6 looks like a mirror image of the observed 1989-2004 change in wages (by percentiles) shown in Figure 4. The back-of-the-envelope calculation presented in Table 2 indeed indicates that de-unionization can account for about of third of both the increase in top-end inequality, and of the decline in the low-end inequality. Remarkably, this is very similar to DiNardo, Fortin, and Lemieux (1996) who also found that de-unionization accounts for a third of the increase in the 90-50 gap between 1979 and 1988.

The other appealing feature of this “union” story is that the English-speaking countries (U.S., U.K. and Canada) where top-end inequality increased the most dramatically (Piketty and Saez, 2006), are also the ones where the role of collective bargaining in wage setting decreased the most sharply (Card, Lemieux, and Riddell, 2003). Unlike alternative explanations like SBTC that should, presumably, affect all advanced countries equally, de-unionization provides of powerful way of accounting for some of the changing nature of inequality within the United States, and the differential experience of different countries in this regard.

c. Another wage-setting institution: Performance Pay

Traditional labor market institutions like the minimum wage and unions tend to affect workers predominantly at the bottom or middle of the wage distribution. By contrast, pay-setting mechanisms of CEOs discussed by Piketty and Saez (2003, 2006) concern workers at the very top of the distribution. This raises the question of whether some pay-setting institutions may also be affecting workers in between these two polar cases.

¹⁰ The overall effect of unions on wage inequality is usually understood as the interplay between the effect of the inequality enhancing “between” effect (union raise average wages of similar workers) and the inequality reducing “within” effect (union reduce wage dispersion among similar workers). At the bottom of the distribution, both the within and between effects go in the same direction (raise wages of low wage workers), which explains why the overall effect is large of positive. At high enough percentiles, however, (negative) within effects can eventually dominate (positive) between effects, explaining why the overall effect eventually turns negative in Figure 6.

When looking at the way compensation of CEOs is set, it is clear that straight salaries represent only a fraction of total compensation. Bonuses, stock options and other forms of compensation linked (in principle) to CEO performance now account for most of the compensation of top CEOs (Piketty and Saez, 2003). Interestingly, this movement towards performance-pay schemes extends well beyond the small number of top CEOs. Using data from the PSID, Lemieux, MacLeod, and Parent (2007) show that the fraction of U.S. male workers on performance-pay jobs increased from about 30 percent in the late 1970s to over 40 percent in the late 1990s.¹¹ They also show that wages are less equally distributed on performance-pay than non performance-pay jobs, in particular because returns to education are higher in performance-pay jobs. They conclude that the growth of performance-pay has contributed to about 25 percent of the increase in the variance of log wages between the late 1970s and the early 1990s.

Interestingly, Lemieux, MacLeod, and Parent show that most of the contribution of performance-pay to the growth in wage inequality is concentrated at the top end of the distribution. This is consistent with their observation that workers on performance-pay jobs tend to be more educated and to work in higher paid occupations (professionals, managers, and sales workers) than workers not paid for performance. Note also that since there is essentially no top-coding in the PSID (see above), that data set is well suited for looking at what is happening even in the very top percentiles of the wage distribution. The findings of Lemieux, MacLeod, and Parent are reproduced in Figure 7, which shows that performance-pay accounts for very little of the changes in wage dispersion up to the 80th percentile. By contrast, performance-pay accounts for remarkably large share of the change in inequality above the 80th percentile. Simple calculations based on the results reported in the Figure 7 indicate that performance-pay accounts for only 3 percent of the growth in the 80-10 wage gap, but for 71 percent of the growth in the 99-80 wage gap.

d. Institutions: An Overall Assessment

¹¹ Performance pay is measured by looking at whether a worker received piece rates, commissions, or bonuses on his current job (looking at all observations of a given worker-employer match). Bonuses are by far the most important component (in terms of workers affected) of this performance-pay measure.

Wage setting institutions are a fairly successful explanation for recent changes in inequality in the United States. De-unionization implies increasing inequality at the top end but decreasing inequality at the low end, which is consistent with changes in the wage distribution observed over the last 15 year. Adding another institutional factor, the minimum wage, can also account for the fact that inequality also expanded in the low end of the distribution in the 1980s, when the real value of the minimum fell sharply. Finally, changes in the way wages are set due to the growth in performance-pay jobs can also help explain a large share of the growth in inequality above the 80th percentile of the wage distribution.

Two other advantages of the institutional explanation are that it is also consistent with long-run trends in inequality in the United States, and with cross-country differences in inequality changes over countries. On the historical side, Levy and Temin (2007) provide both historical and quantitative evidence that institutional changes played both a major role in the decline in inequality around and after World War II (the “Detroit Treaty” era) and in the more recent growth in inequality since 1980 (the “Washington Consensus” era).

In terms of international comparisons, Freeman and Katz (1995)'s view that institutional factors are an important factor for explaining cross-country differences in inequality changes has been confirmed in a number of recent studies. As mentioned above, countries (U.S., U.K. and Canada) where top-end inequality increased the most also experienced a decline in the importance of unions in wage setting (Card, Lemieux, and Riddell, 2003). This being said, it would be incorrect to simply divide large advanced countries in a group of English-speaking countries (U.S., U.K. and Canada) that experienced institutional changes and growing wage inequality, and a group of other countries where both of these factors remained stable.

For instance, Manacorda (2004) shows that inequality first declined in Italy in the 1970s, but subsequently increased starting in the mid-1980s. Manacorda shows compelling evidence that these changes are linked to the “rise and fall” of the Scala Mobile, a country-wide indexation system that compressed wages by awarding larger cost-of-living adjustments to low wage than high wage workers. A more controversial case is the one of Germany where it is generally believed that unions remained strong and

inequality remained stable over the last few decades. Dustmann, Ludsteck and Schönberg (2007) challenge both of these views using new and better data indicating that inequality in fact increased while the unionization rate declined substantially over the last twenty years. They also show that these two factors are clearly linked and that de-unionization accounts for a substantial part of the growth in inequality in Germany over this period. So though the pattern of change in inequality and institutions across countries is more complex than was previously thought, the evidence remains that institutions are an important determinant of wage and income inequality.

It should be stressed, however, that even though wage setting institutions account well for the timing and the nature of changes in wage inequality, they typically do not quantitatively account for most of the observed changes in wage dispersion. For example, the effect of de-unionization in the United States discussed above accounts for no more than a third of the change in (low-end or top-end) wage inequality. Similarly, the decline in the minimum wage in the 1980s only affected the low end of the wage distribution.

The evidence reported here is, thus, not inconsistent with Autor, Katz and Kearney (2007)'s assessment that other factors beyond institutions have to account for the bulk of changes in inequality. Their main argument is that neither changes in the minimum wage nor composition effects can account for the continuing growth in top-end wage inequality. One source of difference between this assessment and the evidence presented here is that, like Lemieux (2006b), Autor, Katz and Kearney (2007) only focus on composition effects due to standard human capital variables that indeed play mostly a role in the low end of the distribution. By contrast, de-unionization and the growth in performance-pay jobs can be viewed as a broader set of composition effects that also effect top-end inequality.

A related point is that when Lemieux, MacLeod, and Parent (2007) separate the pure composition effect (higher incidence of performance pay, holding the wage structure constant) from the wage structure effect (faster growth in the return to skills in the performance-pay sector) of performance pay, they conclude that the latter effect account for most of the contribution of performance pay to the growth in top-end inequality. This means that another explanation such as SBTC needs to be introduced to explain why

inequality grew faster in performance-pay than in non performance-pay jobs. For all these reasons, I will explore in the next section other possible explanations for the changing nature of wage inequality that are required to complement the institutional explanation.

A potentially more serious criticism of the role of institutional factors is that they may themselves be an endogenous response to more fundamental changes, such as SBTC. In particular, Acemoglu, Aghion and Violante (2001) show that if unions maintain strong wage compression policies in the face of declining demand for low skill workers due to SBTC, this will result in de-unionization since the outside option of skilled workers will grow, thus undermining the coalition between skilled and unskilled workers in support of unions. A similar argument could be made that the growth in performance pay is also an endogenous response to underlying technological change.

The main problem with this explanation is that it fails to account for long-run trends and international differences in institutional change. For instance, it is difficult to believe that the growth in the U.S. unionization rate in the 1930s and 1940s was due to technological change, or that unions remained strong in countries like France over the last 20 years because there was no SBTC there. Furthermore, one could make the reverse argument that institutions are the underlying exogenous factor, and the technological change responds to constraints imposed by these institutions. For instance, firms where less-skilled workers are overpriced because of unions have an incentive to invest in new technologies that can effectively replace these less-skilled workers.

So while it would be important and useful to better understand why institutions or other determinants of wage inequality differ over time and countries, these questions are unfortunately beyond the scope of this paper. Following the bulk of the literature, I simply assume that supply, demand, or institutional changes are exogenous.

5. Demand Side Explanations

As mentioned earlier, a simple and parsimonious explanation for the pervasive change in wage inequality of the 1980s is that workers vary according to their skill level, and that the price of skill increased in response to a growing relative demand for skill due, for example, to SBTC. This simple model cannot account, however, for the more complex

change in inequality observed over the last 15 years. I now explore extensions to the simple model that have been suggested to explain the changing nature of wage inequality.

a. Refining technological change: routine vs. skilled tasks.

Autor, Katz and Kearney (2006) propose a richer model for explaining why workers at the top and bottom of the distribution did well since 1990, while workers in the “middle” did not do so well. Their explanation builds on Autor, Levy and Murnane (2003) who convincingly argue that a simple distinction between skilled and unskilled worker is not rich enough to capture the effect of technological change, and in particular the computer revolution, on the labor market. Accordingly, they introduce an important distinction between routine tasks that can now be executed by computers and non-routine tasks that still require human labor.

To illustrate the workings of the model, consider three types of jobs. Type 1 jobs are unskilled non-routine jobs such as truck drivers or nannies that do not require a high level of schooling (skills) but cannot be replaced computer operated machineries. Type 2 jobs are more skilled but routine jobs such as traditional blue-collar jobs in the manufacturing industry that can, to a large extent, be replaced by computer-operated machinery. Type 3 jobs are non-routine high skill jobs such as doctors or executives that are highly skilled and cannot be performed by computers. Since computers can be used as substitutes for type 2 jobs but not for the two other job types, the introduction of computers should depress the wages for these kinds of jobs relative to type 1 and type 3 jobs.

To see how the model can account for the changes observed over that last 15 years, assume that wages for type 2 jobs initially lie in-between those on type 1 (low wage) and type 3 (high wage) jobs. The introduction of computers increases top-end wage inequality by expanding the wage gap between middle-wage jobs (type 2) and high-wage jobs (type 3). By contrast, low-end inequality decreases as wages in middle-wage jobs (type 2) decline relative to those at the bottom (type 1), since computers cannot be used as substitutes for type 1 jobs.

While this story based on the distinction between skilled and routine tasks is appealing, there is not much direct empirical evidence on the precise contribution of this

explanation to the recent changes in wage inequality. Following Autor, Levy and Murnane (2003), a natural way of approaching this issue empirically is to look at detailed occupations. Using data from the Dictionary of Occupational Titles (DOT), they show that the share of the workforce in routine occupations has declined over time, confirming the view that technological change is biased against these types of jobs (as opposed to unskilled jobs).

Autor, Katz and Kearney (2006) also look at detailed occupations and show that, since 1990, changes in occupational shares are U-shaped in the sense that it is occupations in the middle of the skill (or wage) distribution that declined relative to occupations both at the bottom and top end. Goos and Manning (2007) also observe the same changes in the distribution of occupational employment in the U.K., and refer to these changes as a polarization of the labor market into low- and high-paying occupations.

Autor, Katz and Kearney then show that changes in occupational shares were positively sloped (and not U-shaped) in the 1980s, suggesting that technological changes was of a more traditional skill-biased nature during that earlier period. As the shapes of the occupational change curves in the two sub-periods closely match the shapes of the wage change curves (Figure 4), it is tempting to conclude that the technological change imbedded in the occupational change curves can account for the changing nature of wage inequality. This does not say, however, what is the quantitative contribution of this proposed explanation to changes in wage inequality. As pointed out in Section 4, the effect of de-unionization is also U-shaped. When the contribution of unionization is directly estimated, however, it is found to account for about a third of the changes in inequality. So although it is desirable for any explanation to account for the qualitative features of inequality changes, whether it explains a small or a large fraction of the observed changes in inequality can only be assessed using more direct empirical tests.¹²

¹² One major criticism of the earlier SBTC literature is that it often tended to simply attribute changes unexplained by other factors to technological change, instead of coming up with precise measures of technological change that can then be tested on equal footing against other proposed explanations such as de-unionization (see Card and DiNardo, 2002, for more discussion of this). So even though the more sophisticated explanation of the effect of technological change based on the distinction between skilled and routine tasks is intuitively appealing, there is not enough direct evidence yet to conclude that it is the main source of change in wage inequality over the last 15 years.

While it is beyond the scope of this paper to estimate the contribution of technological change to changes in inequality over the last fifteen year, Figures 8 and 9 present some simple evidence from the ORG CPS data on the contribution of different occupations. This provides a rough way of seeing to what extent the pattern of wage changes across occupations is consistent with the skill/routine technological change story. Figure 8 shows relative wage changes by two-digit occupation as a function of average education in the occupation in the base period, as well as a fitted regression of wage changes on a quartic in education. Note that the base period is 1983-85 while the end period is 2000-2002. These two periods are chosen because major changes in occupational classification before 1983 and after 2002 make it difficult to compare occupations over a longer time period.

Interestingly, the fitted regression function is U-shaped, as one would expect on the basis of Autor, Katz, and Kearney (2006) who show that changes in relative demand is U-shaped. Occupations that do the worst are those in the lower middle range of the skill distribution, while occupations in the lower- and, especially, upper-end of the skill distribution do much better.

Occupations at the higher end of the skill distribution (doctors, lawyers, professors, scientists, etc.) almost all do better than average. It is also highly plausible to argue that workers in these occupations mostly perform non-routine work. A more disturbing pattern is that among high-skill occupations, the only ones that experience negative relative wage changes are precisely the ones most closely linked to the computer revolution. In particular computer programmers and engineers both experience negative relative wage changes, despite the fact –indicated by the regression line—that occupations with similar skills generally experience a large positive relative wage change.¹³ The case of computer programmers is particularly striking since this is the two-digit occupation that experienced by far the largest increase in “relative demand” as measured by its employment share that more than tripled (from 0.7 to 2.2 percent) over this period. So while Figure 8 shows that high skill occupations where workers perform non-routine work did very well in terms of relative wages, it is difficult to understand

¹³ This is consistent with Card and DiNardo (2002) who also argue that the relative poor wage performance of technology-related occupations is hard to reconcile with the SBTC story.

why occupations at the core of the computer/technological revolution did relatively badly compared to occupations like lawyers or even professors that are relatively more peripheral to the computer revolution. One possible explanation for this phenomenon discussed below in Section 5c is offshoring.

Figure 9 complements Figure 8 by looking more directly at which occupations contribute the most to the growth in inequality at the top end of the distribution. This is done by first computing the fraction of workers in each occupation who are in the top five percent of the overall (all occupations) wage distribution, and then graphing the change in this proportion over time.¹⁴ Note that this curve cannot be U-shaped since only a very small fraction of workers in middle- and low-wage occupations earn more than the 95th percentile. The change in the fraction of workers in the top five percent for these occupations is thus very close, when not exactly equal, to zero.

As in the case of mean wages in Figure 8, science and technology occupations fare relatively badly according to this alternative measure. Computer scientists, engineers and natural scientists all experience a substantial decline in the fraction of their workers in the top 5 percent of the distribution. By contrast, none of the four occupations that experience the largest growth in the fraction of their workers in the top five percent (doctors, lawyers, health treatment occupations and sales workers in finance, insurance, and real estate (FIRE)) appears to be at the core of the computer/information technology revolution.

In summary, although the more sophisticated form of technological change linked to the distinction between skilled and routine tasks is an intuitively appealing explanation for the polarization of the labor market, it is not yet clear how much it explains of recent changes in wage inequality. Two important challenges are to come up with more direct empirical tests of the effect of this type of technological change on wages, and to account for the puzzling behavior of the wages of computer scientists and other technology workers. A related point that needs to be explored in more detail has to do with the timing of technological changes. In particular, it is not clear why it took until the 1990s for technological change to be biased against routine (as opposed to skilled) tasks and

¹⁴ Redoing the graph for other wage thresholds such as the 90th percentile yield very similar results.

result in some polarization of the labor market, since the computer revolution is typically believed to have played a major role in the 1980s as well.¹⁵

b. Heterogenous returns to human capital

An alternative demand-based explanation suggested by Lemieux (2006a) relies on the idea that increases in the demand for skill can result in more inequality growth in the top end than in the low end of the distribution when returns to skills are heterogenous across workers. Another implication of this model is that the within-group variance increases more for skilled than unskilled workers, an empirical regularity noted in Section 3 (see Lemieux, 2006b).

To fix ideas, consider the standard human capital pricing equation:

$$(1) \quad w_{it} = \alpha_t a_i + \beta_t S_i + \gamma_t X_i + e_{it},$$

where a_i is unobserved ability, S_i is education, X_i is experience, and e_{it} is a measurement error. S_i and X_i are entered linearly to simplify the exposition but could be replaced by more flexible specifications.

This simple model is not consistent with the fact that residual wage dispersion increases at higher values of education (suggesting that α_t is increasing), but does not increase at lower values of education (suggesting that α_t is not increasing). A richer model is clearly required to capture the changes in the wage structure observed over the last 15 years.

As it turns out, equation (1) is a popular but highly restrictive version of a human capital model as it imposes that the return to education (β_t) and experience (γ_t) is the same for all workers.¹⁶ Incorporating heterogeneity in the return to experience and education to equation (1) yields the random coefficient model:

$$(2) \quad w_{it} = \alpha_t a_i + (\beta_t b_i) S_i + (\gamma_t c_i) X_i + e_{it},$$

where b_i and c_i are the person-specific return to education and experience, respectively.

The key implication of this model is that, consistent with empirical evidence, an increase

¹⁵ By contrast, Goos and Manning (2007) find that the polarization of the labor market grew smoothly over both the 1980s and the 1990s in Britain.

¹⁶ Mincer (1974) argues that returns to potential experience are higher for individuals who invest more in on-the-job training than for workers who invest less. Similarly, Becker (1967) develops a human capital investment model where workers have heterogenous returns to education and discount rates. The fact that different people face different returns to education is also front and central in the literature on the estimation of the causal effect of education on earnings (see, for example, Card, 2001).

in the price of education, β_t , increases both the level and the wage dispersion of highly- relative to less-educated workers. To see this, consider the conditional mean and variance of wages under the assumption that the random effects a_i , b_i , and c_i are uncorrelated and that they have a mean of one (normalization):

$$(3) \quad E(w_{it} | S_i, X_i) = \alpha_t + \beta_t S_i + \gamma_t X_i,$$

$$(4) \quad \text{Var}(w_{it} | S_i, X_i) = \alpha_t^2 \sigma_a^2 + (\beta_t^2 \sigma_b^2) S_i^2 + (\gamma_t^2 \sigma_c^2) X_i^2 + \sigma_t^2.$$

Consider the variance term linked to education, $(\beta_t^2 \sigma_b^2) S_i^2$. Since this term depends on education squared, it follows that *i*) the variance should be larger for more-educated workers, and *ii*) there should be a larger increase in the variance of more-educated workers when the price of education, β_t , increases. In this model, an increase in the price of education can, therefore, potentially explain the pattern of change in both between- and within-group differentials without resorting to an increase in the price of unobserved ability, α_t .

Since the link between (log) wages and schooling has become increasingly convex over the 1980s and 1990s (Mincer, 1997, and Deschênes, 2002), Lemieux (2006a) implements empirically the model of equation (3) and (4) using either a quadratic specification or a linear spline in education. His findings indicate that the increase in the return to postsecondary education (more than high school) accounts for a substantial part of the increase in both the between- and within-group components of wage inequality. The main contribution to the within-group component is to explain why this component increases much more among more- than less- educated workers.

While this approach is another appealing explanation for the fact that inequality increased more at the top end than at the low end of the distribution over the last 15 years, it is not clear why the return to post-secondary education has increased so much while returns to other dimensions of skill did not. One potential explanation is again the model of technological change of Autor, Levy, and Murnane (2003). If jobs that require post-secondary education are typically “non-routine” jobs, while jobs that only require high school education tend to be more of the “routine” type, then computerization can result in an increase in the demand for post-secondary (non-routine) skills.

Note that even if this more sophisticated source of technological change was indeed the driving force behind recent changes in inequality, it would still be useful to

combine this explanation with the heterogenous returns story to account for changes in the within-group variance across skill groups. To see this, consider Figure 10 that shows the change in the within-group variance of wages for the same two-digit occupations used in Figures 8 and 9. The figure shows a very strong link between the growth in the within-occupation variance and the average level of education in the occupation. In fact, the R-square of the fitted regression shown in Figure 10 is larger than in corresponding regression (shown in Figure 8) for the change in occupation mean wages. Understanding why the pattern of change in the variance is so closely linked to education must be an important element of any credible explanation for the observed changes in wage inequality.

c. Other demand explanations

Some studies have suggested that Rosen (1981)'s famous theory of "Superstars" could explain why some highly talented individuals at the very top end have done so well over the last few decades.¹⁷ This is obviously a natural explanation for classic cases of superstar markets like professional athletes, movies actors, or singers, though the most important technological revolutions in these sectors (e.g. television) probably happened quite some time ago. The one sector where more recent technological change may have resulted in superstar effect is software development, for example the case of computer game developers. As seen in Figures 8 and 9, however, computer programmers don't appear to have made much of a contribution to the growth in inequality at the top end. By contrast, doctors did much better despite the fact that Rosen (1981) mentions that profession as a leading example where superstar effects are unlikely to happen. The paper by Gabaix and Landier mentioned earlier discusses another possible source of superstar effects, this time in the global market for executives. One possible argument is that globalization leads to much larger firms where it is more critical than before to have a "star" CEOs. In addition to the empirical issues raised earlier, however, it is not clear why this explanation is particularly well suited for the last 15 years since large multinational firms have been around for a long time.

¹⁷ Dew-Becker and Gordon (2005) mention the superstar phenomena as a possible source of growth in top-end inequality.

A perhaps more promising explanation linked to globalization is the more recent phenomena of offshoring. In fact, Levy and Murnane (2006) argue that the offshoring explanation to growing inequality is closely linked to the technological change explanation based on the model of Autor, Levy, and Murnane (2003). The idea is that routine jobs that can be “coded up” and replaced by computers are also typically the easiest ones to offshore. Offshoring could also potentially explain some of the anomalies documented in Figures 8 and 9. For instance, even if computer programmers and scientists do skilled non-routine work, it may be relatively easy to offshore this R&D work to India or Eastern Europe. By contrast, one cannot offshore the services performed by doctors, lawyers, and professors nearly so easily.

Feenstra and Hanson (2003) show that offshoring and, more generally, trade in intermediate goods provides a more compelling empirical role for the effect of trade on wage inequality than more traditional explanations based on the trade in final goods. It would be interesting to see whether this “new” trade channel can also help account for the fact that the more recent growth in inequality has been concentrated at the top end of the distribution.

One final demand side explanation suggested by Manning (2004) is that the growth in earnings at the top end of the distribution due, for instance, to SBTC, has resulted in an increased demand for services consumed by high-income people. To the extent that these services are provided by low-income workers (gardeners, nannies, etc.), the growing demand for these services may have increased wages at the bottom end, but not in the middle end, of the distribution. The implied wage changes under this scenario thus closely parallel those predicted by Autor, Levy, and Murnane (2003)’s model of technological change.

5. Conclusion

During the 1980s, there was a pervasive growth in wage inequality throughout the whole wage distribution. By contrast, inequality growth since 1990 has been concentrated in the top end of the distribution, while inequality in the low end of the distribution declined, at least for men. These recent developments are not consistent with standard models of SBTC that were suggested as the leading explanation for the growth in

inequality in the 1980s. Another limitation of the SBTC explanation is that it cannot explain very well why, unlike the United States and the United Kingdom, some advanced countries experience little change in wage inequality during this period.

An alternative explanation that was suggested to explain inequality growth in the 1980s is based on changes in labor market institutions. I argue in this paper that, unlike SBTC, the institutional change explanation can help explain why inequality changes became concentrated in the top end after 1990, and why inequality grew more in the United States and the United Kingdom than in other advanced countries. This being said, just like in the 1980s, available estimates indicate that institutional change can only account for about a third of the observed recent changes in wage inequality. I also show, however, that broadening the traditional institutional explanation to include pay setting mechanisms such as performance pay can help explain more of the growth in inequality at the top end.

A number of other explanations have also been suggested for explaining the changing nature of wage inequality. One possible way of “rescuing” the SBTC explanation is to introduce a richer form of technological change based on the distinction between skilled and routine tasks. A closely related explanation is linked to offshoring. In both cases, however, we do not yet have detailed decompositions showing exactly how much of the recent changes in inequality can be attributed to these factors. It will be important to estimate the precise explanatory power of these explanations in future research. To be empirically successful, these explanations will also have to explain why the within-group variance increased so much more among highly-educated workers than among less-educated workers. A promising way of addressing this issue is to use a model where returns to skill are heterogeneous across workers.

For the time being, however, most of the growth in top-end inequality over the last 15 years remains unaccounted for. It may be due to more general institutions and social norms, technological change, offshoring, or something else. Short of coming up with more direct empirical tests, there is no reason to privilege one explanation over another a priori. I provide a number of possible research directions in the paper that may hopefully help direct future empirical work. One of the most fruitful research direction would be to study in more detail the recent experience of countries other than the United

States. International comparisons are a very powerful way of separating explanations that have similar implications in all countries, like technological change, from those that have more country-specific implications, like social norms and institutions.

References

Acemoglu D (2002) Technical Change, Inequality, and the Labor Market. *Journal of Economic Literature* 40(1): 7-72

Acemoglu D, Aghion P, and Violante G (2001) Technical Change, Deunionization, and Inequality. *Carnegie-Rochester Conference Series On Public Policy* 55(1): 229-264

Autor DH, Katz LF, Kearney MS (2005) Rising Wage Inequality: The Role of Composition and Prices. NBER Working Paper No. 11628

Autor DH, Katz LF, Kearney MS (2006) The Polarization of the U.S. Labor Market. *American Economic Review Papers and Proceedings* 96(2): 189-194

Autor DH, Katz LF, Kearney MS (2007) Trends in U.S. Wage Inequality: Revising the Revisionists. Harvard University mimeo.

Autor DH, Levy F, Murnane RJ (2003) The Skill Content of Recent Technological Change: An Empirical Investigation. *Quarterly Journal of Economics* 118(4): 1279-1333

Beaudry P, Green DA (2005) Changes in US Wages, 1976-2000: Ongoing Skill Bias or Major Technological Change? *Journal of Labor Economics* 23(3): 609-48

Bebchuk, L, Grinstein Y (2005) The Growth of Executive Pay. NBER Working Paper No. 11443

Becker GS (1967). Woytinski Lecture, University of Michigan. Reproduced in *Human Capital (2nd Edition)*. University of Chicago Press, Chicago

Berman E, Bound J, Griliches Z (1994) Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures. *Quarterly Journal of Economics* 109(2): 367-97

Blackburn ML, Bloom DE (1987) Earnings and Income Inequality in the United States. *Population and Development Review* 13(4): 575-609

Blinder, AS (1980) The Level and Distribution of Economic Well-Being. In: Feldstein M (ed) *The American Economy in Transition*. University of Chicago Press, Chicago, 415-499

Bluestone B, Harrison B (1988) The Growth of Low-Wage Employment: 1963-86. *American Economic Review, Papers and Proceedings* 78(2): 124-128

Bound J, Johnson G (1992) Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations. *American Economic Review* 82(3): 371-92

Card, D (1992) The Effects of Unions on the Distribution of Wages: Redistribution or Relabelling? NBER Working Paper No. 4195

Card D (2001) Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems. *Econometrica* 69(5): 1127-1160

Card D, DiNardo J (2002) Skill Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles. *Journal of Labor Economics* 20(4): 733-83

- Card D, Lemieux T (1996) Wage Dispersion, Returns to Skill, and Black-White Wage Differentials. *Journal of Econometrics* 74(2): 319-361
- Card D, Lemieux T, Riddell WC (2003) Unions and the Wage Structure. In: Addison JT, Schnabel C (eds) *The International Handbook of Trade Unions*. Edward Elgar, Cheltenham, 246-92
- Deschênes O (2002) Unobserved Ability, Comparative Advantage and the Rising Return to Education in the United States: A Cohort-Based Approach. UC Santa Barbara mimeo
- Dew-Becker I, Gordon RJ (2005) Where Did the Productivity Growth Go? Inflation Dynamics and the Distribution of Income. *Brookings Papers on Economic Activity* 2005(2): 67-127
- DiNardo J, Fortin NM, Lemieux T (1996) Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach. *Econometrica* 64(5): 1001-46
- Dustmann C, Ludsteck J, Schönberg U (2007) Revisiting the German Wage Structure. IZA Discussion Paper No. 2685
- Feenstra RC, Hanson GH (2003): Global Production Sharing and Inequality: A Survey of Trade and Wages. In: Choi EK, Harrigan J (eds) *Handbook of International Trade*, Blackwell, Malden Oxford Victoria, 146-185
- Firpo, S, Fortin NM, Lemieux T (2007) Unconditional Quantile Regressions. NBER Technical Working Paper No. 339
- Fortin NM, Lemieux T (1997) Institutional Changes and Wage Inequality: Is there a Linkage? *Journal of Economic Perspective* 11(2): 75-96
- Freeman, RB (1976) *The Overeducated American*. Academic Press, New York
- Freeman, RB (1980) Unionism and the Dispersion of Wages. *Industrial and Labor Relations Review* 34(1), 3-23
- Freeman, RB (1993) How Much has Deunionization Contributed to the Rise of Male Earnings Inequality? In: Danziger S, Gottschalk P (eds) *Uneven Tides: Rising Income Inequality in America*, Russell Sage Foundation, New York, 133-163
- Freeman, RB, Katz LF (1995) *Differences and Changes in Wage Structures*. University of Chicago Press, Chicago.
- Gabaix X, Landier A (2008) Why Has CEO Pay Increased So Much? *Quarterly Journal of Economics*, forthcoming.
- Goos M, Manning A (2007) Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *The Review of Economics and Statistics* 89(1): 118-133.
- Juhn C, Murphy KM, Pierce B. Wage Inequality and the Rise in Returns to Skill. *Journal of Political Economy* 101(3): 410-42
- Kaplan SN, Rauh J (2007) Wall Street and Main Street: What Contributes to the Rise in the Highest Incomes? NBER Working Paper No. 13270
- Katz, LF, Autor DH (1999) Changes in the Wage Structure and Earnings Inequality. In: Ashenfelter O, Card D (eds) *Handbook of Labor Economics, volume 3A*. Elsevier Science, Amsterdam, 1463-1555
- Katz LF, Murphy KM (1992) Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *Quarterly Journal of Economics* 107(1): 35-78
- Krueger AB (1993) How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989. *Quarterly Journal of Economics* 108(1): 33-60

- Lee DS (1999) Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage. *Quarterly Journal of Economics* 114(3): 977-1023
- Lemieux T (2006a) Post-secondary Education and Increasing Wage Inequality. *American Economic Review Papers and Proceedings* 96(2): 195-199
- Lemieux T (2006b) Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill? *American Economic Review* 96(3): 461-498
- Lemieux, T, MacLeod WB, Parent D (2007) Performance Pay and Wage Inequality. NBER Working Paper No. 13128
- Levy F, Murnane RJ (1992) U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations. *Journal of Economic Literature* 30(3): 1333-81
- Levy F, Murnane RJ (2006) How Computerized Work and Globalization Shape Human Skill Demands. MIT and Harvard mimeo
- Levy F, Temin P (2007) Inequality and Institutions in 20th Century America, NBER Working Paper No. 13106
- Manacorda M (2004) Can the Scala Mobile Explain the Fall and Rise of Earnings Inequality in Italy? A Semiparametric Analysis, 1977-1993. *Journal of Labor Economics* 22(3): 585-613
- Manning A (2004) We Can Work It Out: The Impact of Technological Change on the Demand for Low-Skill Workers. *Scottish Journal of Political Economy* 51(5); 581-608
- Mincer, J (1974) *Schooling, Experience, and Earnings*. Columbia University Press, New York
- Mincer J (1997) Changes in Wage Inequality, 1970-1990. In: Polachek S (ed) *Research in Labor Economics* 16. JAI Press, Greenwich 1-18
- Piketty T (2003) Income inequality in France, 1901-1998. *Journal of Political Economy* 111(5): 1004-1042
- Piketty T, Saez E (2003) Income Inequality in the United States, 1913-98. *Quarterly Journal of Economics* 118(1): 1-39
- Piketty T, Saez E (2006) The Evolution of Top Incomes: A Historical and International Perspective. *American Economic Review, Papers and Proceedings* 96(2): 200-205
- Rosen, S (1981) The Economics of Superstars. *American Economic Review* 71(5): 845-58
- Saez E, Veall M (2005) The Evolution of High Incomes in Northern America: Lessons from Canadian Evidence. *American Economic Review* 95(3): 831-849

Table 1: Summary Measures of Wage Inequality

	1974	1979	1989	1999	2004	Change:	
						1974 to 1989	1989 to 2004
A. MEN							
Standard deviation	0.501	0.499	0.571	0.580	0.598	0.072	0.027
90-10 gap	1.266	1.262	1.486	1.487	1.557	0.224	0.071
50-10 gap	0.681	0.693	0.787	0.721	0.729	0.094	-0.058
90-50 gap	0.585	0.569	0.699	0.767	0.828	0.130	0.129
Gap relative to HS graduates							
High school dropouts	-0.246	-0.253	-0.297	-0.295	-0.287	-0.045	0.010
Some college	0.116	0.108	0.165	0.168	0.183	0.058	0.018
College graduates	0.369	0.328	0.466	0.521	0.545	0.138	0.079
Post-graduates	0.365	0.361	0.568	0.715	0.766	0.207	0.198
B. WOMEN							
Standard deviation	0.435	0.414	0.512	0.531	0.549	0.098	0.037
90-10 gap	1.033	0.989	1.314	1.353	1.404	0.325	0.090
50-10 gap	0.470	0.409	0.633	0.605	0.633	0.223	0.001
90-50 gap	0.564	0.580	0.681	0.748	0.770	0.101	0.089
Gap relative to HS graduates							
High school dropouts	-0.221	-0.191	-0.275	-0.275	-0.281	-0.085	-0.005
Some college	0.164	0.141	0.212	0.193	0.204	0.070	-0.008
College graduates	0.376	0.311	0.459	0.534	0.547	0.149	0.088
Post-graduates	0.558	0.489	0.629	0.770	0.790	0.140	0.161

Notes: Measures computed by pooling groups of three year centered around the year listed in the table. For example, “1974” corresponds to year 1973 to 1975, etc.

Table 2: De-unionization and 1988-2005 Changes in Male Wage Inequality

	Wage gap:		
	50-10	90-50	90-10
Effect of Unions:	0.139	-0.469	-0.330
Change in wage gap	-0.037	0.119	0.082
Change in unionization rate	-0.086	-0.086	-0.086
Predicted effect of de-unionization	-0.012	0.040	0.028
(effect in %)	32.3	33.9	34.7

Notes: Changes in wage gap taken from Figure 4. Effect of unions computed using estimates provided in Firpo, Fortin, and Lemieux (2007). Change in unionization rate computed using 1988 and 2005 ORG supplements of the CPS. The predicted effect of de-unionization is the product of the union effect (first column) and of the change in the unionization rate (third column).

Figure 1: Low-end and Top-end Wage Inequality, Men

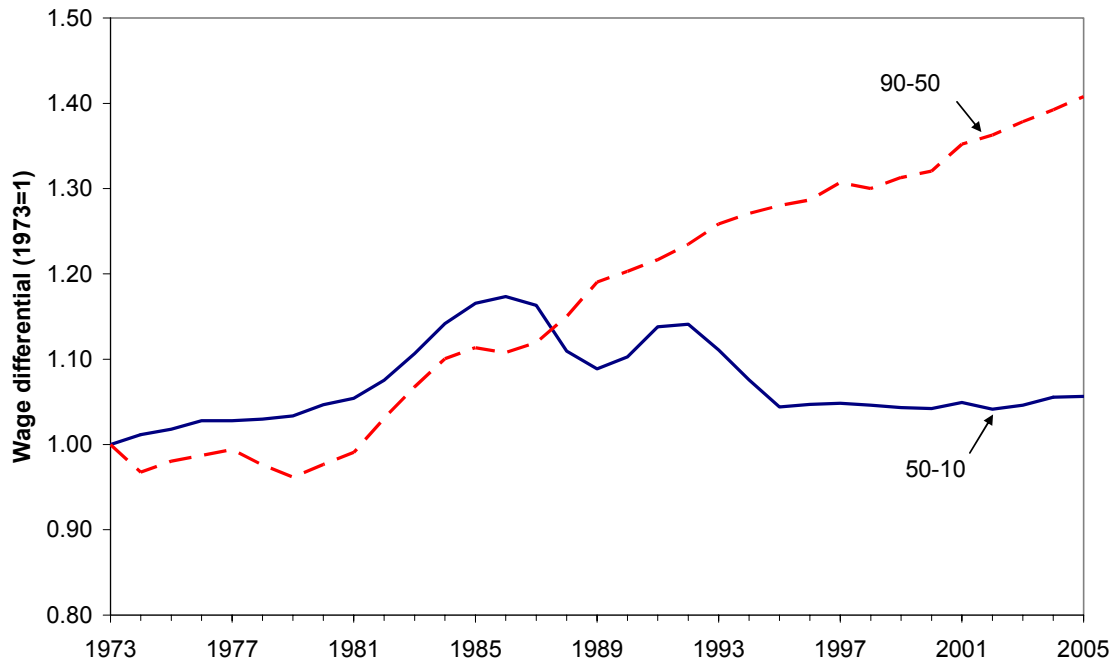


Figure 2: Low-end and Top-end Wage Inequality, Women

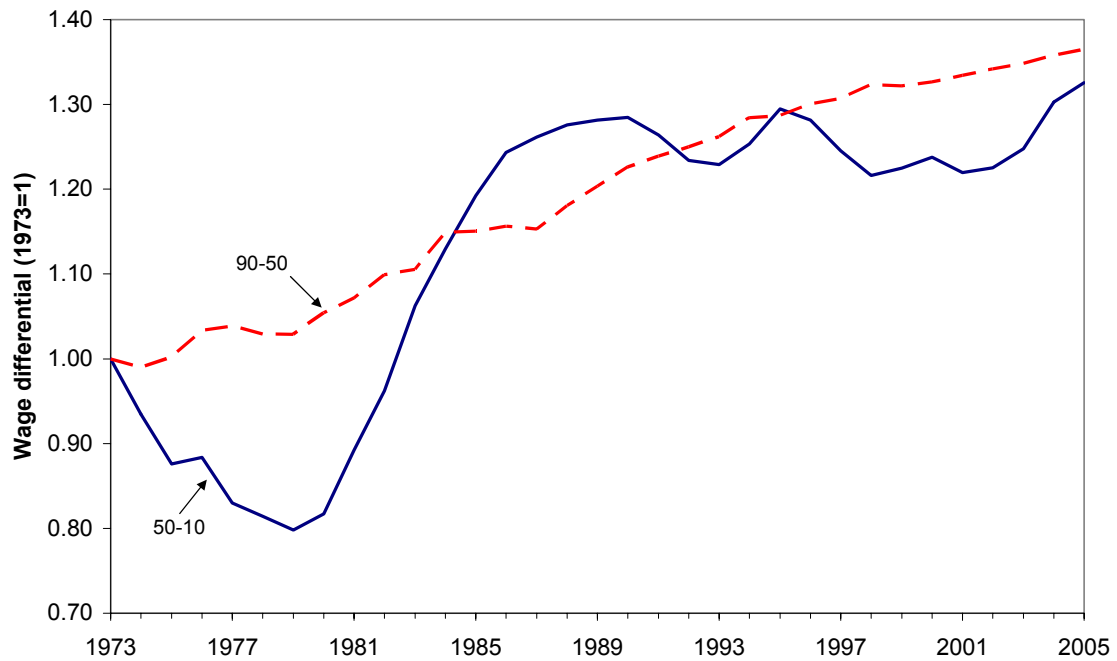


Figure 3: Education Wage Differentials Relative to High School Graduates, Men

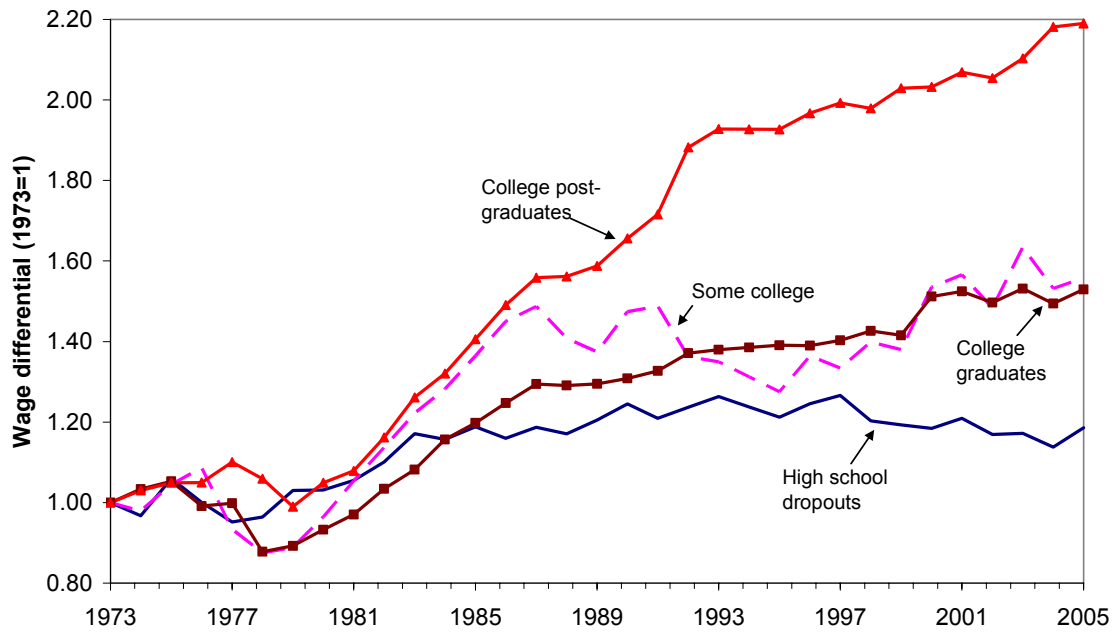


Figure 4: Change in Real Wages by Percentile, Men

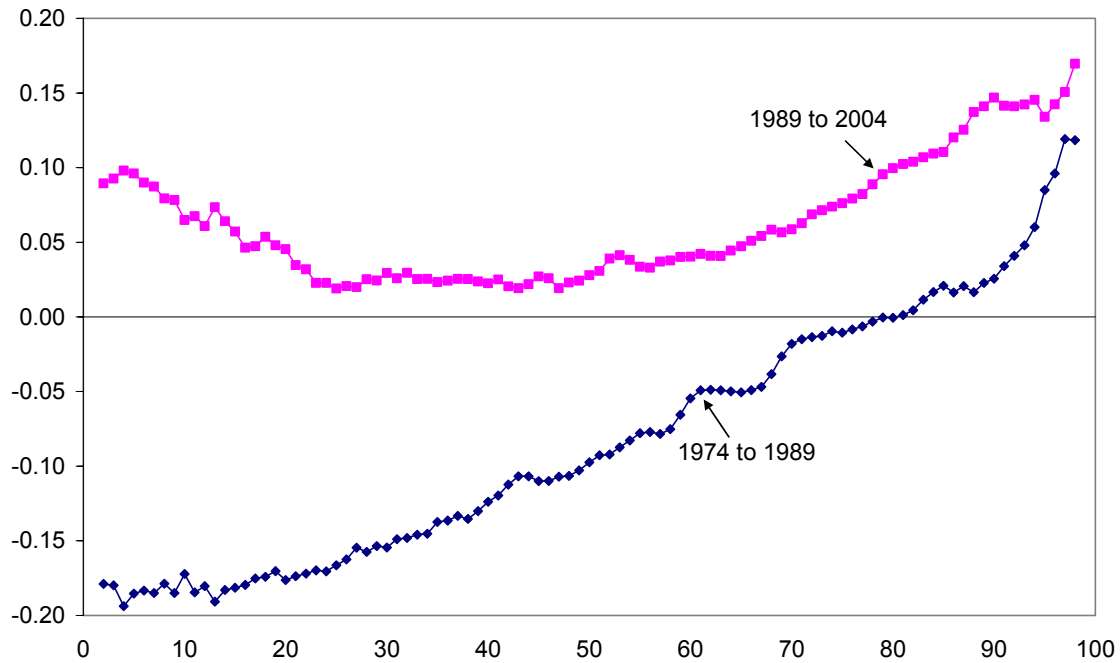


Figure 5: 50-10 Gap for Women vs. Minimum Wage

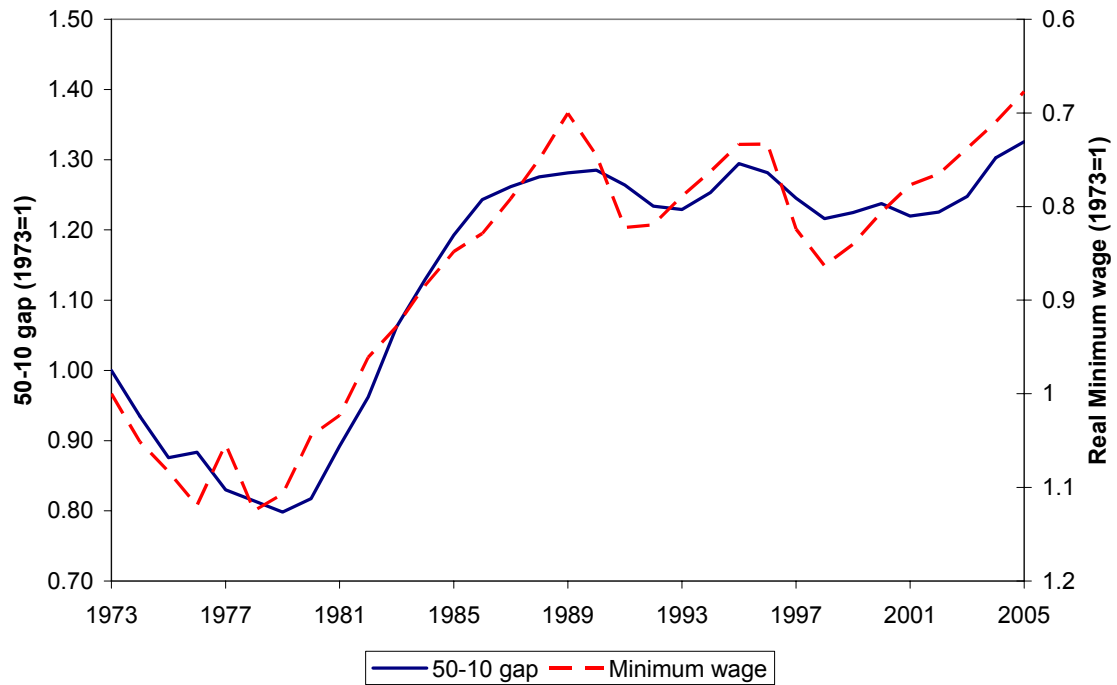


Figure 6: Union Effect by Wage Percentile, Men

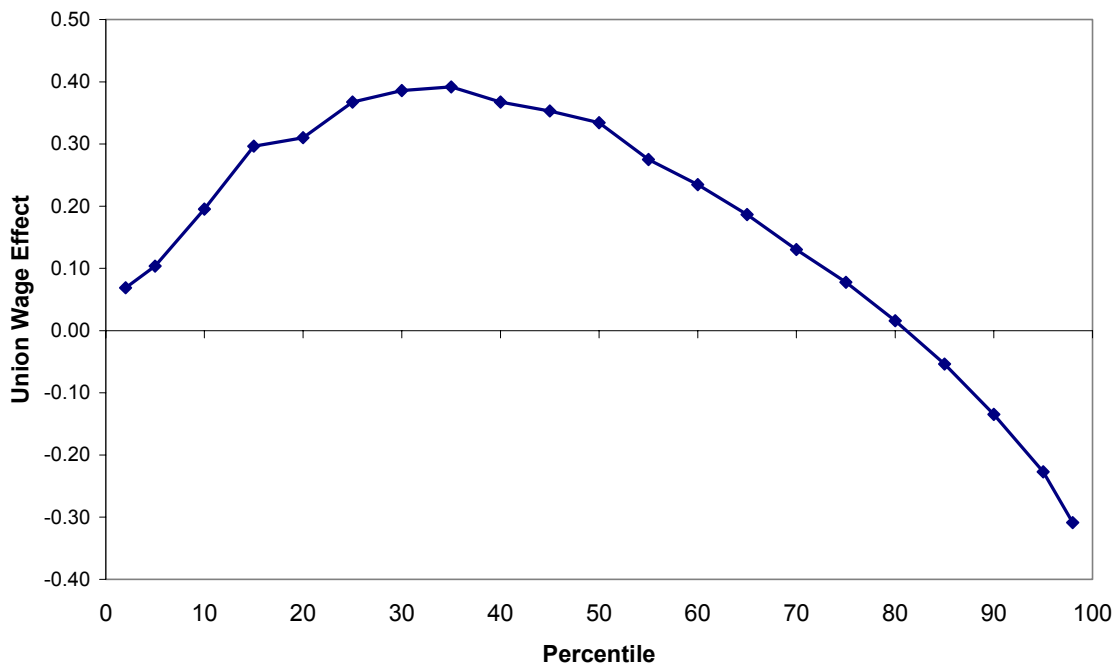


Figure 7: Pay-for-Performance and the 1976-79 to 1990-93 Change in Real Wages by Percentile, Male Heads in the PSID

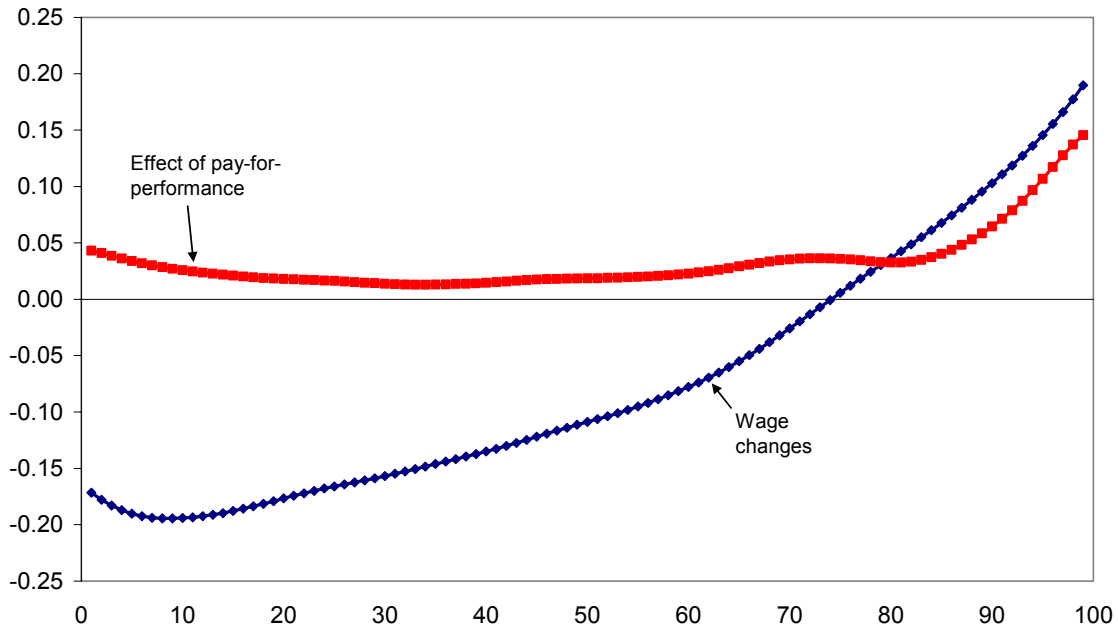


Figure 8: Relative Wage Change by Two-Digit occupation, 1983-85 to 2000-02

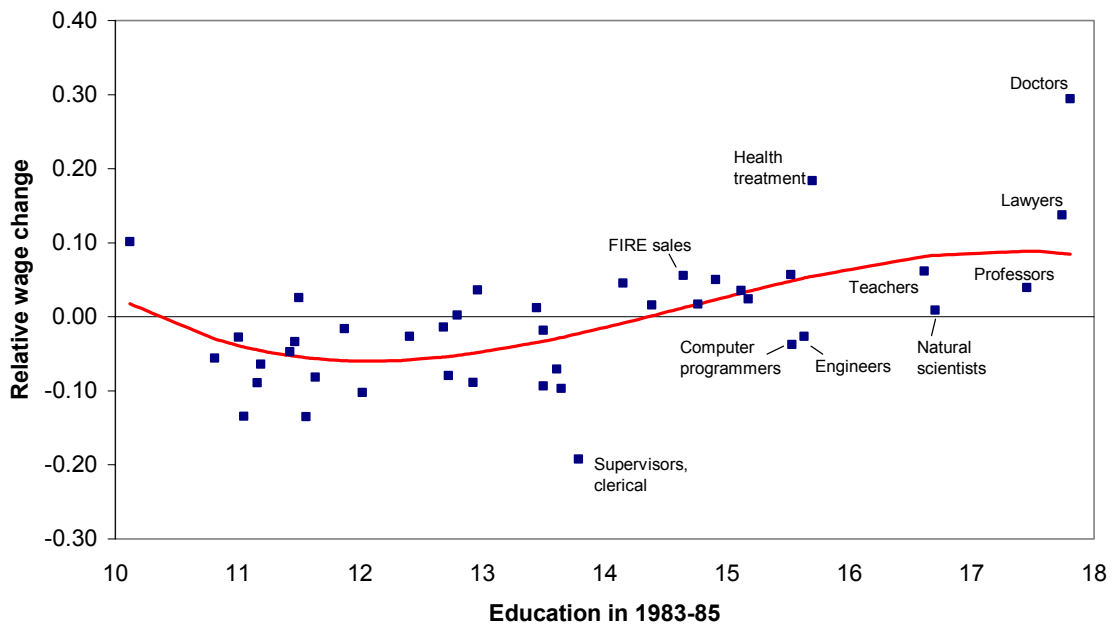


Figure 9: Change in the Fraction in the Top 5 percent by Two-Digit Occupation, 1983-85 to 2000-2002

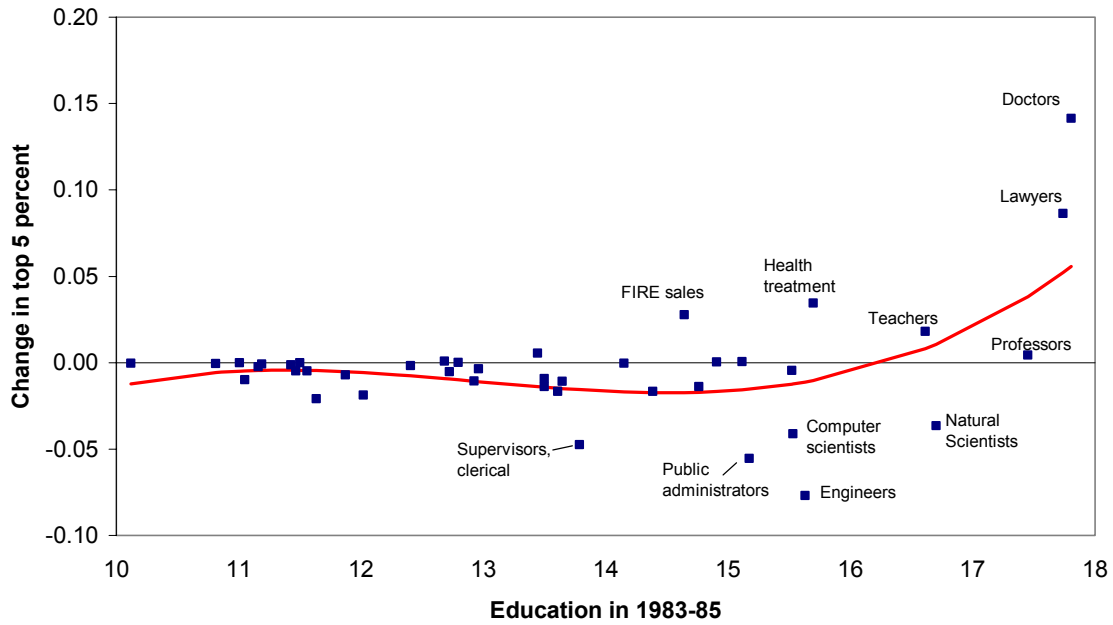
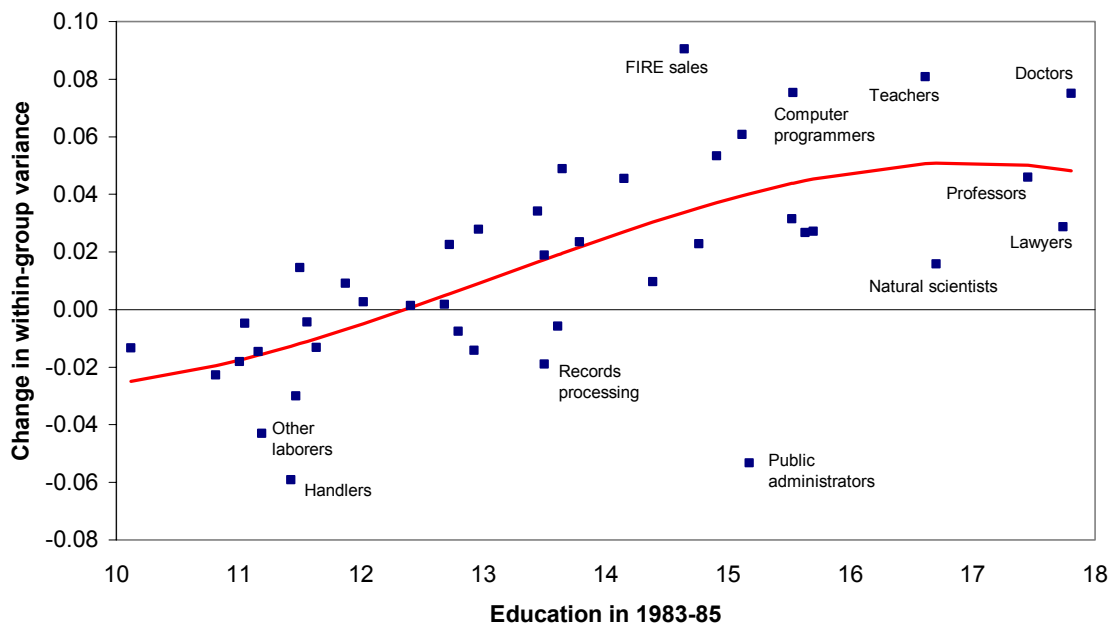


Figure 10: Change in Within-Group Variance by Two-Digit Occupation, 1983-85 to 2000-2002



ENDNOTES