

REVISITING THE GERMAN WAGE STRUCTURE¹

Christian Dustmann

Johannes Ludsteck

Uta Schönberg

Abstract

This paper shows that wage inequality in West Germany has increased over the past three decades, contrary to common perceptions based on less reliable data. During the 1980s, the increase was concentrated at the top of the distribution; in the 1990s, it occurred at the bottom end as well. Changes in the workforce composition only partially explain these increases. During the 1990s, Germany experienced a sharp decline in union coverage as well as a slowdown in the skill upgrading of medium- relative to the low-skilled workers. These events increased both lower-tail inequality and the medium-low wage differential over the same period. These findings are consistent with the view that technological change is responsible for the widening of the wage distribution at the top. The widening of the wage distribution at the bottom, however, may be better explained by episodic events, such as supply shocks and changes in labor market institutions.

Keywords: inequality, polarization, institutions

JEL: J3, D3, O3

¹Acknowledgments: For helpful comments, we would like to thank our editor, three referees, David Autor, David Card, Bernd Fitzenberger, Thomas Lemieux, Alexandra Spitz-Oener, and seminar participants at the Australian National University, ESPE, Frankfurt University, the Institute for Employment Research (IAB), Mannheim University, and the University of Melbourne for comments and suggestions. We gratefully acknowledge financial support from the German Research Foundation (DFG) and the Anglo-German Foundation (AGF). We thank Bernd Fitzenberger, Alexandra Spitz-Oener, and Joachim Wagner for sharing their programs and/or data with us.

1 Introduction

The United States witnessed a sharp increase in wage and earnings inequality during the 1980s (e.g. Bound and Johnson 1992, Katz and Murphy 1992, Levy and Murnane 1992, and Acemoglu 2002). Upper-tail inequality, measured as the 90-50 wage gap, continued to rise at a similar pace during the 1990s, whereas lower-tail inequality, measured as the 50/10 wage gap, has been falling or flat since the late 1980s (e.g. Autor, Katz, and Kearney 2008).² A similar increase in inequality in the 1980s has also been observed in other Anglo-Saxon countries, such as the United Kingdom (e.g. Gosling, Machin, and Meghir 2000) and Canada (e.g. Boudarbat et al. 2006).

In contrast, most countries in Continental Europe seem to have witnessed much smaller increases in inequality in the 1980s, or no increases at all (see, for example, Freeman and Katz 1995 and OECD 1996 for a summary of trends in inequality in European countries). In particular, West Germany, the third largest economy and the largest exporter in the world, has been singled out as a country characterized by a stable wage distribution during the 1980s (see, for example, Steiner and Wagner 1998; Prasad 2004). Numerous scholars cite this stability as evidence against the hypothesis that the growth of inequality observed in the United States and United Kingdom is primarily due to skill-biased technological change, as firms in Continental Europe had access to the same technologies as firms in the United States or United Kingdom (e.g. Card, Kramarz, and Lemieux 1999; Piketty and Saez 2003; and Saez and Veall 2005). Possible explanations for this puzzle include a larger expansion in the relative supply of the high-skilled in Germany (e.g. Acemoglu 2003; Abraham and Houseman 1995), unions and other labor market institutions (e.g. Krugman 1994; Abraham and Houseman 1995)³, and more recently social norms (e.g. Piketty and Saez 2003).

²Lemieux (2006a, 2007) also emphasizes that the increase in inequality in the United States is increasingly concentrated at the top of the wage distribution.

³Acemoglu (2002) emphasizes an interesting link between technological change and institutions. If unions compress wages, then firms have greater incentives to adopt labor-complementary technologies, which will reinforce wage compression.

This paper revisits the changes in the wage structure in (West) Germany over the past three decades, between 1975 and 2004. Most existing studies on the German wage structure, such as Steiner and Wagner (1998) and OECD (1996), are based on the German Socio-Economic Panel (GSOEP). We instead use a 2% random sample of social security records, the IABS. Our data have several advantages over the GSOEP, most importantly a much larger sample size. Using this data, we show that the common perception that Germany's wage structure has remained largely stable during the 1980s is inaccurate. We find that wage inequality has increased in the 1980s, but mostly at the top half of the distribution. In the early 1990s, wage inequality started to rise also at the bottom half of the distribution. This pattern holds for both men and women.

The first of these findings has also been documented by Fitzenberger (1999), using an earlier version of our data for the years 1975 to 1990. The second finding is in line with recent papers by Kohn (2006) and Gernandt and Pfeiffer (2006) who document a similar increase in inequality in lower-tail inequality in the IABS and GSOEP, respectively. However, we are not aware of any paper that jointly analyzes changes in inequality in both the 1980s and 1990s, and compares these trends with those in the United States. Our analysis highlights that, while the United States and Germany experienced similar changes at the top of the distribution during the 1980s and 1990s, the two countries markedly differ with respect to the lower end of the wage distribution: The rise in lower tail inequality happened in the 1980s in the United States, but in the 1990s in Germany.

We investigate several explanations for the changes in wage inequality in Germany. First, we use the kernel re-weighting procedure first proposed by DiNardo, Fortin, and Lemieux (1996) to analyze whether the changes in inequality are explained by mechanical changes in the workforce composition, or whether they reflect changes in skill prices. In line with Lemieux (2006a), we show that it is important to account for changes in workforce composition, in particular at the upper end of the wage distribution. However, these changes cannot fully account for the divergent path

of upper and lower tail inequality in the 1980s, or for the divergent path of lower tail inequality in the 1980s and 1990s.

Second, we document a sharp decline in unionization rates in the late 1990s: The share of workers covered by union agreements declined from 87.3% in 1995 to 72.8% in 2004. There is little evidence of a similar decline during the 1980s. Using the same decomposition method as above, we find that between 1995 and 2004 de-unionization can account for 28% of the rise in inequality at the lower end of the wage distribution, but only 11% at the upper end.

Third, we document a rise in the wage differential of medium-skilled workers (i.e. those with an apprenticeship degree) relative to the low-skilled (i.e. those with no post-secondary education) starting in the late 1980s, around the same time lower-tail inequality started to increase. There is, however, no clear trend in the wage differential of high-skilled workers (i.e. those with a college degree) relative to the medium-skilled. We also document that the decline in the share of the low-skilled started to slow down in the late 1980s, whereas the share of the high-skilled increased at a roughly linear rate from 4.7% in 1975 to 14.8% in 2004. Using a nested CES production framework based on that by Goldin and Katz (2007a, 2008), we show that fluctuations in relative supply explain the evolution of the wage differential between the low- and medium-skilled very well, but do a poor job in predicting the evolution of the wage differential between the medium- and high-skilled.

Fourth, building on the analysis by Spitz-Oener (2006), we provide evidence that is consistent with a polarization of work: during the 1980s and 1990s, occupations with high median wages in 1980 experienced the largest growth rate, while occupations in the middle of the 1980 wage distribution lost ground relative to occupations at the bottom. Moreover, occupations at the high end of the 1980 wage distribution predominantly use non-routine analytic and interactive skills, while routine task usage is highest in the upper middle of the wage distribution. This is consistent with Autor,

Levy, and Murnane's (2003) hypothesis that computer technology decreases the demand for jobs that require routine manual or clerical skills (and are found in the middle of the wage distribution), and increases the demand for jobs that require non-routine cognitive and interpersonal skills (and are found at the top of the wage distribution). This paper thus adds to the growing evidence that technology does not simply increase the demand for skilled labor relative to that of unskilled labor, but instead asymmetrically affects the bottom and the top of the wage distribution (see, for example, Autor, Katz, and Kearney 2006, 2008 for the United States and Goos and Manning 2007 for the United Kingdom). This may begin to supply the unifying international evidence on technological change that so far has been absent.

The evidence provided in this paper is consistent with the idea that technological change is an important driving force behind the widening of the wage distribution, particularly at the top. This conclusion is reinforced by our finding that for occupations above the median, employment and wage changes by wage percentile are positively correlated. In contrast, below-median employment and wage changes are negatively correlated. The rise in lower-tail inequality may therefore be better explained by episodic events, such as supply shocks and changes in labor market institutions. We argue that these shocks happened a decade later in Germany than in the United States.

The plan of this paper is as follows. Section 2 describes the data used for the analysis. Section 3 documents the major changes in the German wage structure between 1975 and 2004. We then analyze four possible explanations for the increase in inequality: changes in the workforce composition (Section 4.1), a potential decline in unionization (Section 4.2), supply shocks (Section 4.3), and polarization (Section 4.4). We conclude with a discussion of our findings in Section 5.

2 Data Description

Our empirical analysis is based on two data sets: the IABS, a 2% random sample of social security records, and the LIAB, a linked employer-employee data set. We describe each data set in turn.

2.1 IABS: 2% Random Sample of Social Security Records, 1975-2004

Our main data set is a 2% sample of administrative social security records in Germany for the years 1975 to 2004. The data is representative of all individuals covered by the social security system, roughly 80% of the German workforce. It excludes the self-employed, civil servants, individuals currently doing their (compulsory) military service, and individuals on so-called "marginal jobs" (i.e. jobs with at most 15 hours per week or temporary jobs that last no longer than 6 weeks). This data set (or earlier versions of it) has been used to study wage inequality by, among others, Steiner and Wagner (1998), Möller (2005), Fitzenberger (1999), Kohn (2006), and Fitzenberger and Kohn (2006).

The IABS has several advantages over the German Socio-Economic Panel, the data set most often used to analyze trends in inequality in Germany (e.g. Steiner and Wagner 1998; OECD 1996; Prasad 2004). First, the IABS is available from 1975 onward, as opposed to 1984 for the GSOEP. Second, the sample size is much larger (more than 200,000 observations per year, as opposed to around 2,000 in the GSOEP). Third, wages are likely to be measured much more precisely in the IABS than in the GSOEP, as misreporting by firms in the IABS is subject to severe penalties. Fourth, attrition rates in the GSOEP are large enough to worry that results are not representative for the population as a whole (see for example Spiess and Pannenberg 2003; and Haisken De-New and Frick 2005). In contrast, while workers can also be followed over time in the IABS, each year the original sample is supplemented by a random sample of new labor market entrants. This guarantees

that the IABS is representative of workers who pay social security contributions.

The main disadvantage of the IABS is that it is right-censored at the highest level of earnings that are subject to social security contributions. Overall, each year between 9.4% and 14.2% of the male wage distribution is censored. Because of censoring, this paper mostly focuses on the changes in the uncensored part of the wage distribution, up to the 85th percentile. In the United States, much of the action in rising wage inequality since the mid-1980s has been above the 85th percentile (e.g. Autor Katz and Kearney 2008; and Piketty and Saez 2003); consequently, top-coding in our data may lead us to substantially under-state inequality growth.

Another difficulty in our data is a structural break in the wage measure in 1984. From 1984 on, our measure includes bonus payments as well as other one-time payments (Steiner and Wagner 1998). We follow Fitzenberger (1999) and correct for the break (see Appendix A.1 for details). Further, our data set does not contain precise information on the number of hours worked; we only observe whether a worker is working full- or part-time (defined as working less than 30 hours per week). We therefore restrict the wage analysis to full-time workers and use the *daily* wage, averaged over the number of days the worker was working in the year, as our wage measure. Robustness checks against the GSOEP suggest that this does not affect our results.

From this data base, we select all men and women between 21 and 60 years of age. Since the level and structure of wages differ substantially between East and West Germany, we concentrate on West Germany (which we usually refer to simply as Germany). While we provide a descriptive overview of the evolution of inequality for both men and women, our main analysis focuses on men only. Further details on the sample selection and variable description can be found in Appendix A.2.

2.2 LIAB: Linked Employer-Employee Data, 1995-2004

The data set just described provides no information on union coverage, and can thus not be used to analyze the impact of de-unionization on the wage structure. Our analysis here is based on the LIAB, a linked employer-employee data set provided by the Institute for Employment Research (IAB). It combines information from the IAB Establishment Panel with information on all workers who were employed in one of these firms as of the 30th of June. The information on workers is drawn from the same social security records as our main data.

Although data is available from 1993 to 2004, we only use waves from 1995 onward for which information on union recognition is consistent. In Germany, a firm recognizes the union by either joining an employer federation (*Arbeitgeberverband*), or by engaging in bilateral negotiations with the union. In the first case, union wages are negotiated at a regional and industry level, typically on an annual basis. Our union variable distinguishes between firm- and industry-level agreements.

The IAB establishment panel over-samples large establishments. To make our results representative of the German economy as a whole, we weight our results using the cross-sectional weights provided by the LIAB. In Table B.1 in Appendix B, we compare median wages as well as interquartile differences for men in the LIAB and the IABS. Both data sources draw a very similar picture of the developments in the wage structure over this period.

3 Trends in Wage Inequality

In Section 3.1, we describe the major changes in wage inequality in Germany from 1975 to 2004. We compare our findings with those reported in other studies in Section 3.2. Because of wage censoring, we focus on the changes in the uncensored part of the wage distribution, and impose no distributional assumptions on the error term in the wage regression. However, some of our findings, such as

the evolution of the standard deviation of log-wages and log-wage residuals, require distributional assumptions. We assume that the error term is normally distributed, with different variances for each education group and each age group, and impute the censored part of the wage distribution under this assumption. We prefer to work with imputed wages rather than with censored wages because wage residuals can be computed in a straightforward manner. A comparison between OLS estimates based on imputed wages and tobit estimates based on censored wages shows that both the estimates and the standard errors are almost identical. More details on the imputation method can be found in Appendix A.3. We have conducted extensive robustness checks regarding alternative distributional assumptions, including an upper-tail pareto distribution. Our results are highly robust to alternative imputation methods. Findings for alternative imputation methods can be found in Section 1 in the on-line appendix.

3.1 Basic Facts

Standard Deviation of Log-Wages Figure I displays the evolution of the standard deviations of log-wages and log-wage residuals. Panel A refers to men, Panel B to women. The standard deviation is obtained from standard OLS regressions on imputed wages, estimated separately for each year. We control for three education categories, eight age categories, and all possible interactions between these two. For men, the figure shows a continuous rise in both overall and residual inequality during the 1980s, with an acceleration in the 1990s. A simple within-between decomposition indicates that the majority of the increase in inequality occurred within age and education groups (86% between 1975 and 1989, and 65% between 1990 and 2004).

For women, in contrast, the standard deviation of log-wages and log-wage residuals remained roughly constant during the 1980s, and started to increase only in the mid-1990s. A further difference between men and women is that age and education explain a smaller portion of the overall

variance of log-wages for women. As with men, most (82%) of the increase in overall inequality between 1990 and 2004 is due to a rise in within-group inequality.

The Top versus the Bottom Next, we separately analyze changes in inequality at the bottom and top of the wage distribution. Figure II displays the wage growth of the 15th, 50th, and 85th percentiles of the wage distribution. We distinguish between the pre- and post-unification period (1975 to 1989 and 1990 to 2004). For men, the 15th and 50th percentile evolved similarly between 1975 and 1989, and increased by about 16%. Over the same time period, the 85th percentile rose by 27.2% (Panel A). The picture looks very different during the 1990s (Panel C): between 1993 and 2004, the 15th percentile declined by almost 5%, while the 50th and 85th percentile increased by 4% and 13%, respectively.

The pattern for women is somewhat different: between 1975 and 1989, wage gains were highest for the 15th percentile (about 25%, compared to only 16% for men). Over the same time period, both the 50th and the 85th percentiles grew by about 22%, compared to 16% and 27% for men (Panel B). In the post-unification period, in contrast, wages at the 15th percentile stagnated, while the 85th percentile experienced the highest wage growth (17%, Panel D). Unlike to the 1980s, in the 1990s wages of women caught up to those of men throughout the entire wage distribution.

Figure III illustrates the divergent developments of the lower and upper ends of the wage distribution during the 1980s and 1990s in a different manner. It shows log real wage growth across the wage distribution, for the period between 1980 and 1990, as well as between 1990 and 2000. In the 1980s, male wages grew across the distribution, but substantially more so at the upper than at the lower tail. Wage growth accelerates beyond the 65th percentile. In contrast, between 1990 and 2000, wage growth has been negative below the 18th percentile, with wage losses at the 5th percentile of more than 10 log wage points. Starting from the 15th percentile, wage growth increases

roughly linearly along the wage distribution (Panel A).

For women (Panel B), the 1980s are characterized by wage compression at the lower tail of the wage distribution, whereas wage growth at the very top (i.e. 95th percentile) exceeds that at the median by about 6%.⁴ In the 1990s, in contrast, wage growth increases roughly linearly along the wage distribution.

How do these findings compare with developments in the United States? Both countries show an increase in inequality at the top of the wage distribution during the 1980s and 1990s, although in Germany the increase is more pronounced for men than for women. The two countries differ sharply with respect to the developments at the bottom of the wage distribution. In the United States, the 50/10 wage gap rose substantially in the 1980s, but ceased to increase in the 1990s. In Germany, the pattern is reversed. What about the magnitude of the changes? Since our wage measure is the full-time daily wage, our findings are probably most comparable to those based on the March CPS for weekly full-time earnings. Autor, Katz, and Kearney (2008) report that between 1975 and 2004, the difference between the 90th and 50th percentiles of the male earnings distribution increased by about one log point per year (their Figure 3). We find that over the same time period, the 85/50 wage gap in Germany rose by about 0.6 log points per year. However, it is important to bear in mind that in the United States, much of the action in rising wage inequality since the mid-1980s has been above the 85th percentile.⁵ Hence, top-coding of our data could lead us to substantially under-estimate the rise in upper-tail inequality during the 1980s and 1990s.

⁴Since for women less than 5% of wages are censored, we plot wage growth up to the 95th percentile.

⁵See for example Autor Katz and Kearney (2008), Dew-Becker and Gordon (2005), Gordon and Dew-Becker (2007), Goldin and Katz (2007b), and Piketty and Saez (2003).

3.2 Comparison with Existing Studies

These results seem to contradict the usual view that wage inequality in Germany has been largely stable over the past two decades, and in particular during the 1980s. What explains this discrepancy? The reason is that the majority of existing studies on inequality trends in Germany, such as Steiner and Wagner (1998), Prasad (2004) and OECD (1996), are based on a different data set, the German Socio-Economic Panel. Studies based on the IABS are generally consistent with our findings. In particular, Fitzenberger (1999) emphasizes that wage inequality rose during the 1980s, and that the increase was concentrated at the top of the distribution. His study uses data from 1975 to 1990 only, and was therefore not able to detect the large increase in lower-tail inequality in the 1990s.⁶ Existing studies based on the GSOEP and our study based on the IABS thus seem to draw a different picture of the trends in inequality in Germany.

We have investigated three possible explanations for the discrepancy between our findings and those based on the GSOEP. First, the GSOEP includes civil servants and the self-employed, but these workers are excluded in the IABS. Second, the wage measure in the IABS includes bonuses as well as other one-time annual payments. In contrast, studies based on the GSOEP typically do not include one-time payments although they are available. Third, and most importantly, most studies based on the GSOEP construct an hourly wage rate, whereas the wage measure in the IABS is a daily wage.

Here, we provide only a brief overview, focusing on men. A detailed comparison between the GSOEP and IABS can be found in Section 3 of the on-line appendix. Our findings indicate similar trends in inequality whether or not we include civil servants or the self-employed, or whether or not we include bonuses and other one-time payments in our wage measure. Importantly, inequality

⁶Other studies using the IABS focus on other aspects of the wage structure. For instance, Kohn (2006) concentrates on the recent developments in the 90s as well as differences between East and West Germany (see also Möller 2005), while Fitzenberger and Kohn (2006) analyze trends in the returns to education.

trends based on monthly wages are also similar to those based on hourly wages. Any differences between the GSOEP and IABS are therefore not adequately explained by differences in the sample used or by differences in the wage measure.

Our analysis further indicates that inequality rose during the 1990s in the GSOEP, in particular at the bottom of the wage distribution, which has also been stressed by Gernandt and Pfeiffer (2006). Our analysis also highlights that measures of inequality are very noisily estimated in the GSOEP. The change in the 50/15 and 85/50 wage gap as well as the change in the standard deviation of log-wages between two years observed in the IABS is almost always within the 95%-confidence interval of that observed in the GSOEP. For instance, using the specification that most closely resembles that in the IABS, the 95% confidence intervals for the changes in the 50/15 and 85/50 wage gaps between 1993 and 2002 are $[0.044, 0.154]$ and $[-0.039, 0.103]$, respectively. Over the same period, the 50/15 and 85/50 wage gaps rose by 0.059 and 0.058 in the IABS. Given the large standard errors in the GSOEP, it is not surprising that earlier studies, such as the 1996 OECD Employment Report, concluded that the German wage structure was largely stable between the mid-1980s to mid-1990s.

Next, we explore several explanations for the rising wage inequality in Germany. Here, we restrict the analysis to men, for two reasons. First, female labor force participation rates have risen considerably during the 1980s and 1990s; this is likely to have changed the selection of women into work, which may have had an independent impact on the female wage structure.⁷ Second, although the basic patterns in the wage structure (i.e., upper-tail inequality increased during the 1980s and 1990s, whereas lower-tail inequality mostly increased in the 1990s) are similar for both men and women, there are also important differences. For instance, wage gains are substantially larger for women than for men, especially in the 1990s. Moreover, the increase in upper-tail inequality is more pronounced for men than for women, especially in the 1980s. Explaining these differences between

⁷Mulligan and Rubinstein (2008) demonstrate that in the United States it is important to account for the changing selection of women into the workforce when computing male-female wage differentials.

men and women would be beyond the scope of this paper.

4 Why Did Wage Inequality Increase?

4.1 The Role of Composition and Prices

Is the increase in inequality described in the previous section explained by changes in the workforce composition, or do they reflect changes in skill prices? To see why it is important to account for compositional changes in the workforce, suppose that the variance of log-wages is increasing in education and age. If the employment share of educated and older workers increases over time, then this will lead to a *mechanical* rise in inequality, even if skill prices do not change. Lemieux (2006a) stresses that in the United States, a large fraction of the rise in residual wage inequality between 1973 and 2003—and all since 1988—can be attributed to such changes in the workforce composition. This section employs the kernel re-weighting approach developed by DiNardo, Fortin, and Lemieux (1996) to recover the counterfactual wage distribution that we would have observed if the workforce composition had remained unchanged. Like Autor, Katz, and Kearney (2008), we focus on the divergent path of upper- and lower-tail inequality in the 1980s and 1990s, rather than the variance of log-wage residuals, as Lemieux (2006).

The following expression decomposes the observed density of log-wages w in years t and t' into a "price" $g(\cdot)$ and a "composition" function $h(\cdot)$ (see also Autor, Katz, and Kearney 2008):

$$f(w|t) = \int g_t(w|x, T = t)h_t(x|T = t)dx \quad \text{and} \quad f(w|t') = \int g_{t'}(w|x, T = t')h_{t'}(x|T = t')dx.$$

Here, $g(w|x, T = t)$ is the density of log-wages in year t for observable characteristics x , and $h(x|T = t)$ is the density of characteristics x in year t . In order to compute the counterfactual

wage distribution in year t' that would have prevailed if the workforce composition were the same as in year t , we simply need to re-weight the price function $g_{t'}(\cdot)$ in year t' by the ratio $h_t(\cdot)/h_{t'}(\cdot)$ of the densities of characteristics x in years t and in year t' .⁸ In our application, all regressors (i.e. all possible interactions between three education and eight age groups) are categorical. The re-weighting function is therefore straightforward to compute, and we do not need to estimate logit models on pooled data.

This decomposition method applies to calculating counterfactuals for overall inequality. In order to recover counterfactuals for residual inequality, we replace the pricing function $g_t(w|x, T = t)$ with the residual pricing function $g_t(\epsilon|x, T = t)$. The residuals are obtained from OLS regressions on imputed wages that control for all possible interactions between three education and eight age groups. We would like to point out that we do not need to impose any distributional assumptions on the error term in order to obtain the uncensored part of the counterfactual distribution of *overall* inequality. However, distributional assumptions are required in order to compute the counterfactual distribution of *residual* inequality. Our results are robust to alternative imputation methods (see Section 1 of the on-line appendix for details). It is also important to stress that the decomposition ignores general equilibrium effects, as it is based on the assumption that changes in quantities do not affect changes in prices.

Table I provides a first overview about how wage dispersion, measured as the 50/15 and 85/50 wage gaps, and employment shares vary by age and education groups. We distinguish three education groups which we label low, medium and high. The low-skilled are workers who enter the labor market without post-secondary education. The medium-skilled are workers who completed an apprenticeship or a high school degree (*Abitur*). The high-skilled are workers who graduated from a university or college. Due to severe censoring for the high-skilled, we only report the 50/15

⁸This ratio can be calculated as $\frac{h(x|T=t)}{h(x|T=t')} = \frac{Pr(T=t|x)}{Pr(T=t'|x)} \cdot \frac{1-Pr(T=t)}{1-Pr(T=t')}$.

wage gap for this group. Note that this may lead us to understate the increase in within-group inequality, as in the United States much of the growth in inequality is found in the top half of the high skill group.

Results are based on imputed wages, and cells where the 85th or 50th percentile is censored are marked. Similar to the United States, wage dispersion is increasing in education and—with the exception of the low-skilled—in age. The share of the low-skilled decreased by 12 percentage points between 1976 and 1990, but only by 3.6 percentage points between 1990 and 2004. The share of the high-skilled rose monotonically from 4.7% in 1976 to 14.7% in 2004. The share of workers below the age of 36 rose from 38.9% in 1976 to 41.6% in 1990, and declined to 30.9% in 2004. Table I also highlights that wage dispersion rose *within* education and age groups, suggesting that mechanical changes in the workforce composition cannot fully account for the rise in inequality. Between 1976 and 1990, the medium-skilled above the age of 45 experienced the sharpest rise in inequality, while between 1990 and 2004, the rise in inequality is strongest for the young low-skilled. For this group, the increase in the 50/15 wage gap increases by more than 20 log points. Here, it is important to stress that our data include employees covered by the social security system only; if temporary and marginal employment were included in the data, the increase might be even larger.

Table II reports trends in observed and counterfactual overall and residual inequality. We distinguish three interquantile ranges: 85/15 (Panel A), 85/50 (Panel B), and 50/15 (Panel C). For each wage gap, the first row shows the observed change. The next rows show the counterfactual change that would have prevailed if the workforce composition were the same as in 1980, 1990, or 2000. The table shows that the overall 85/15 wage gap increased by about 8.2 log-points between 1980 and 1990, and by 10.8 log-points between 1990 and 2000. If the labor force composition had remained the same as in 1980, the 85/15 wage gap would have risen by 5.4 log-points between 1980 and 1990, and by 8.5 log points between 1990 and 2000. The results are similar when we use

the workforce composition in 1990 or 2000 to calculate the composition-constant increase in overall inequality.

Table II also illustrates that composition effects play a more important role for the upper tail than at the lower tail of the wage distribution. During both the 1980s and 1990s, changes in workforce composition can explain up to 50 percent of the increase in upper-tail overall inequality, but at most 15 percent of the increase in lower-tail overall inequality. This differs from findings for the United States where the impact of changes in workforce composition is concentrated at the lower end of the earnings distribution (Autor, Katz, and Kearney 2008). Turning to residual inequality, the qualitative patterns are very similar. However, composition effects account for a considerably smaller share of the rise in the residual 85/50 wage gap than in the overall 85/50 wage gap (e.g. 15% versus 37% for 1980 characteristics).

What are the principle factors that explain the role of composition in increasing upper-tail inequality, rising education or population aging? When we only account for changes in the education structure, but not in the age structure, the composition-adjusted increase in the 85/50 wage gap is similar to the one when we additionally account for changes in the age structure, during both the 1980s and 1990s. This suggests that rising education is the driving factor.

These results demonstrate that it is important to account for changes in the workforce composition, as emphasized by Lemieux (2006a). However, mechanical changes in the workforce composition do not fully explain the increase in upper-tail inequality in the 1980s, nor do they account for the divergent path of lower-tail inequality in the 1980s and 1990s.

4.2 Decline in Unionization

Several papers in the United States argue that part of the increase in inequality in the 1980s can be linked to a decline in the minimum wage and unionization (e.g. DiNardo, Fortin, and Lemieux

1996, Lee 1999, and Card and DiNardo 2002). We now explore this hypothesis for Germany using the LIAB data.

The German system of collective bargaining differs in several aspects from that in the United States. Most importantly, in Germany the recognition of trade unions for collective bargaining purposes is at the discretion of the employer. Once a firm has recognized a union, collective bargaining outcomes de facto apply to all workers in that firm, no matter whether they are union members or not. A firm recognizes union either by joining an employer federation (*Arbeitgeberverband*), or by engaging in bilateral negotiations with the union. In the first case, union wages are negotiated at a regional and industry level, typically on an annual basis. Another key difference from the United States is that there is no legal minimum wage in Germany. However, union contracts in Germany specify wage levels for specific groups in specific sectors, and can be considered an elaborate system of minimum wages.

Table III, based on the LIAB data set, shows a remarkable decline in union coverage during the mid 1990s and early 2000s: Between 1995 and 2004, the share of workers covered by an industry-level agreement declined by about 12 percentage points, and the share of workers covered by a firm-level agreement decreased by 3 percentage points. Unfortunately, comparable data on union coverage does not exist before 1995. For the 1980s, only data on union membership is available. Schnabel and Wagner (2006) report that throughout the 1980s, about 40% of men were union members.⁹ By 2000, however, union membership had dwindled to about 31%. This suggests that the decline in unionization in Germany is mostly a phenomenon of the 1990s.

There is strong evidence that unions compress the wage structure in Germany, and more so at the lower end of the wage distribution (see, for example, Gerlach and Stephan 2005, 2007; Fitzenberger and Kohn 2005; and Dustmann and Schönberg 2007). A natural question to ask

⁹Because in Germany collectively bargained agreements apply to all workers in a firm that recognizes the union, union membership is much smaller than union coverage.

is: did the de-unionization in the 1990s contribute to the rise in inequality over this period, in particular at the lower tail of the wage distribution? To test this hypothesis, we employ the same decomposition method as in Section 3.1, and include as regressors all possible interactions between the recognition of an industry- or firm-level agreement and the three education and eight age groups. It is again important to stress that the decomposition method ignores general equilibrium effects; in our application, this means that the union-non-union wage differential is assumed to be independent of union coverage. Moreover, the decomposition assumes that unionization is exogenous and not itself determined by the same factors that raise wage inequality. A further assumption behind the decomposition method is that there are no spillover effects from the unionized to the non-unionized sector.

Figure IV plots the observed wage changes between 1995 and 2004 as well the counterfactual wage changes that would have prevailed if unionization rates had remained at their 1995 level across the wage distribution. The figure illustrates that workers throughout the wage distribution would have experienced a higher wage growth if unionization rates had not declined. However, the impact of de-unionization is substantially stronger at the lower end of the wage distribution. For instance, wages in 2004 would have been 5.5% higher at the 5th percentile, but only 0.2% higher at the 85th percentile.

We provide more details in Table IV. The first set of columns refer to overall inequality, while the second set of columns refer to residual inequality. The residuals are obtained from OLS regressions on imputed wages. In each pair of columns, we first hold only unionization constant. We then additionally keep the age and education distribution constant. We again distinguish two interquantile differences: 85/50 and 50/15. We first report the observed change, and then the counterfactual change if the unionization, age, and education distribution were the same as in 1995 or 2004, respectively. Between 1995 and 2004, the overall 85/50 wage gap rose by 0.069 log points. If

unionization rates had remained at their 1995 level, the increase in upper-tail inequality would have been 0.061 log points—a reduction of 12%. Unionization plays a more important role at the lower end of the distribution: de-unionization can account for 28% of the increase in the overall 50/15 wage gap. The findings are similar for residual inequality. In line with the results in Table II, workforce characteristics also play an important role, particularly at the upper end of the distribution.

These results indicate that the decline in union recognition in the 1990s had a profound impact on the wage structure, especially at the lower end of the distribution. It is not surprising that de-unionization also affected the distribution above the median, since rather than a single minimum wage (as in the United States), union minimum wages are set for all levels of qualification.

4.3 The Role of Relative Skill Supplies

An important component of the rise in inequality in the United States is the remarkable increase in the return to education. We now provide evidence on the recent trends in the skill premium in Germany, and analyze the explanatory power of demand and supply factors. We focus on the wage differential between medium-skilled workers (i.e. those who completed an apprenticeship) and low-skilled workers (i.e. those who lack post-secondary education). For completeness, we also report results for the wage differential between high-skilled workers (i.e. those with a university degree) and the medium-skilled. However, due to the high incidence of censoring among the high-skilled, these results have to be viewed with considerable caution.

Panel A of Figure V plots the wage differential between the low- and medium-skilled (left y-axis) and the medium- and high-skilled (right y-axis). Our results are based on imputed wages. Our regressions control for three education and eight age groups as well as all possible interactions. The medium-low and the high-medium wage premiums are age-adjusted, and computed as a weighted average of the respective premium in each age group, where the weights are the employment-

weighted worker share in each age group, averaged over the entire sample period.

The medium-low wage differential declined slightly between 1975 and 1989, and then increased sharply by about 0.7 percentage points a year. This timing coincides with the sharp rise in lower-tail wage inequality. The medium-high wage differential declined between 1975 and 1980, remained roughly constant during the 1980s and mid-1990s, and started to increase in the late 1990s. In Panel B, we plot the (employment-weighted) share of low- and high-skilled workers over time; here, results refer to men and women together. During the late 1970s and 1980s, the share of low-skilled workers declined sharply from 26.2% in 1975 to 14.5% in 1989. This share then decreased by only 3.6 percentage points between 1990 and 2004, when the medium-low skill premium increased sharply. The share of university graduates rose at a roughly linear rate during the 1980s and 1990s, from 4.5% in 1975 to 14.7% in 2004, indicating that fluctuations in supply may have played an important role in the rise of medium-low skill premium.

In order to analyze the importance of fluctuations in labor supply more formally, we adopt a two-level CES production function framework (Goldin and Katz 2007a, 2008). Suppose that output Y only depends on quantities S and U of skilled and unskilled workers S , defined as workers with and without university degrees, respectively:

$$Y_t = A_t[\lambda_t S_t^\rho + (1 - \lambda_t)U_t^\rho]^{1/\rho}.$$

In this expression, A is total factor productivity and λ_t represents a shift in technology. The aggregate elasticity of substitution between skilled and unskilled workers is given by $\sigma_{SU} = \frac{1}{1-\rho}$.

Unskilled labor is itself a CES sub-aggregate that depends on the quantities L and M of low- and

medium-skilled workers:¹⁰

$$U_t = [\theta_t L_t^\eta + (1 - \theta_t) M_t^\eta]^{1/\eta}. \quad (1)$$

The elasticity of substitution between the medium- and the low-skilled is given by $\sigma_{ML} = \frac{1}{1-\eta}$. Under the assumption that labor is paid its marginal product, the medium-low and skilled-unskilled wage differentials satisfy (2) and (3):

$$\log\left(\frac{w_{S_t}}{w_{U_t}}\right) = \log\left(\frac{\lambda_t}{1 - \lambda_t}\right) - \frac{1}{\sigma_{SU}} \log\left(\frac{S_t}{U_t}\right), \text{ and} \quad (2)$$

$$\log\left(\frac{w_{M_t}}{w_{L_t}}\right) = \log\left(\frac{\theta_t}{1 - \theta_t}\right) - \frac{1}{\sigma_{ML}} \log\left(\frac{M_t}{L_t}\right). \quad (3)$$

Relative wages depend on demand shifters λ_t and θ_t , on relative supplies $\log(\frac{S_t}{U_t})$ and $\log(\frac{M_t}{L_t})$, and on the respective elasticities of substitution σ_{SU} and σ_{ML} . We estimate (2) and (3) in two steps. In the first step, we estimate (3) by OLS, and substitute for $\log(\frac{\theta_t}{1-\theta_t})$ with a linear time trend. We then use the estimate for σ_{ML} to compute the quantity U_t of the unskilled labor supplied in (1). In the second step, we estimate (2) by OLS, this time substituting $\log(\frac{\lambda_t}{1-\lambda_t})$ with a linear time trend. To account for generated regressor bias in the first step, we bootstrap standard errors in the second step. Although our wage differentials refer to men only, we include women in our supply measures.¹¹

Results are reported in Table V, Panel A. For the medium- versus low-skilled, we obtain an estimate for the elasticity of substitution of about $\sigma_{ML} = 5$ (1/0.206). This estimate is considerably larger than the estimate of around 1.4 typically found in the United States, but this is likely because the typical US estimate refers to the elasticity of substitution between low- and high-skilled, which

¹⁰This assumption implies that an increase in skilled labor relative to unskilled labor does not affect the wage premium of the medium-skilled relative to the low-skilled.

¹¹See Appendix A.1 for a detailed description how the wage premiums and the relative supply measures are computed.

are presumably less perfect substitutes than the low- and medium-skilled considered here.¹² This model can explain 94% of the time variation in the wage premium of the medium-skilled relative to the low-skilled. In contrast, for the skilled (i.e. university graduates) versus the unskilled (i.e. a mixture of the low- and medium-skilled), the model performs poorly: the relative supply coefficient estimate is positive, and the coefficient on the linear time trend is negative. The model can explain only 18% of the time variation in the relative wage premium between the high- and the combined medium/low-skilled. In Panel B, we jointly estimate equations (2) and (3), and restrict the elasticity of substitution to be the same for the low- and medium-skilled and the medium- and high-skilled and obtain an elasticity of substitution close to our previous estimate for the low- versus medium-skilled.

Figure V, Panels C to F, provide a visual illustration of the relationship between relative supplies and relative wages. Panels C and D plot the series of relative supplies and relative wages from 1975 to 2004, in deviations from a linear trend, for the medium- and low-skilled and for the skilled and the unskilled workers, respectively. For the medium- and low-skilled, the decline in the de-trended relative wage coincides with a rise in the de-trended relative supply, and vice versa. For skilled and unskilled, there is no clear pattern. Panels E and F in Figure V plot the observed relative wage gap as well as the gap predicted by the two-level CES production function (using the estimates in Table V, Panel A) against time. The figure confirms our previous conclusions: the model predicts trends in the wage differential between the medium- and low-skilled very well, but does a poor job in forecasting the evolution in the wage differential between the skilled and the unskilled.

These results suggest that the deceleration in the decline of low-skill employment shares in the 1990s had a profound impact on skill prices, and thus the wage structure—particularly at the lower tail of the wage distribution. What caused this deceleration? While more research is needed on this issue, the timing suggests that it is a consequence of the breakdown of the communist regimes

¹²A complementary explanation for the higher elasticity of substitution in Germany is that wages in Germany are less responsive to supply and demand shocks than in the United States, due to higher unionization rates.

in Eastern Europe, as well as the reunification of East and West Germany. These events lead to a large inflow of East Germans, Eastern Europeans, and ethnic Germans from Eastern Europe into the West German labor market. Many of these immigrants were low-skilled; see Glitz (2006) and Bauer et al. (2005) for more details.

It may seem puzzling that the wage gap between the high- and the medium-skilled does not increase except possibly in the late 1990s, while the overall and residual 85/50 wage gaps rise during both the 1980s and 1990s. We wish to stress that this seeming de-coupling should not be overemphasized, since the high incidence of wage censoring among the high-skilled casts some doubt as to whether the high/medium wage gap can be reliably estimated. Nevertheless, in an attempt to explain the divergent evolution of within- and between-group inequality, we next provide some evidence of a rising demand for the high-skilled, relative to the medium-skilled, by computing between-occupation demand shifts for each education group relative to a base year (see Katz and Murphy 1992):¹³

$$\Delta D_k = \sum_j \left(\frac{E_{jk}}{E_k} \right) \left(\frac{\Delta E_j}{E_j} \right).$$

Here k indexes skill groups and j indexes occupations, E_j is total labor input measured in efficiency units in occupation j , and $\frac{E_{jk}}{E_k}$ is group k 's employment share (in efficiency units) in occupation j in the base year. We prefer this measure of demand shifts over that implied by the CES production function framework, which performed very poorly for the high-skilled. This measure is also not based on relative wage differentials, which may be seriously compromised due to the high incidence of wage censoring among the high-skilled, and does require an estimate for the elasticity of substitution. It is important to stress, however, that this index does not account for the impact of price changes on observed employment shifts. Thus, if positive skill supply shocks cause expansions of high-skill occupations, this demand index will overstate the "demand" shock. Conversely, if skill premiums

¹³Our findings are similar when we compute between-industry demand shifts.

rise due to demand shifts, occupational shifts will be smaller than the price-constant counterfactual, thus leading the index to understate the "demand" shift. Figure VI plots the between-occupation demand shifts of the medium- versus low-skilled, and the high- versus the medium-skilled. The figure indicates a considerable demand shift favoring the high-skilled relative to the medium-skilled during the 1980s and 1990s. This demand shift was substantially larger than that favoring the medium-skilled relative to the low-skilled. To put these numbers into perspective, Katz and Murphy (1992) report a between-industry demand shift of college graduates relative to high school graduates of 0.067 between 1979 and 1987.¹⁴ Over the same period, we find a between-occupation demand shift of 0.157 and a between-industry demand shift of 0.084 for the high- relative to the medium-skilled. This suggests that the relative wage premium of the high-skilled did not increase much in Germany during the 1980s and 1990s because of off-setting shifts in the relative supply of the high-skilled.¹⁵

4.4 Polarization

Our findings highlight the importance of distinguishing between changes in lower- and upper-tail inequality. Moreover, Figure VI suggests that demand shifts for the high-skilled relative to the medium-skilled exceed those for the medium-skilled relative to the low-skilled. Autor, Katz, and Kearney (2005, 2006, 2008) provide a simple demand-based explanation for this pattern (see also Autor, Levy, and Murnane 2003). The idea is that technological change—in particular the implementation of computer technology—differently affects the bottom and top ends of the skill distribution. Suppose that computerization decreases the demand for jobs that require routine analytical or clerical skills, and increases the demand for abstract, non-routine cognitive and interpersonal skills.

¹⁴This number is based on the between industry demand shifts reported in Table 6, and computed as $\exp(0.029+0.036)-1$.

¹⁵This is in line with Acemoglu (2003), who first computes demand shifts implied by the CES production function for the United States. He then investigates whether the evolution of relative supply and skill premiums in Europe is consistent with US demand shifts, using the Luxemburg Income Study. This is the case for Germany. See also Abraham and Houseman (1995), who argue that, unlike the United States, there is no evidence for a slowdown in skill upgrading in Germany.

Computer technology neither strongly complements nor strongly substitutes manual skills. If routine analytical skills are predominantly used in the middle, and manual and interactive skills at the bottom and top of the wage distribution, then technological change may lead to ‘polarization’ (Goos and Manning 2007), and thus affect the wage distribution differently at the bottom and the top. For Germany, Spitz-Oener (2006) provides evidence that between 1979 and 1999, the demand for interactive and non-routine analytical skills increased, while the demand for routine-cognitive skills declined. Much of these changes can be linked to computerization. This section further investigates this hypothesis for Germany.

We first test a key assumption behind this approach: Non-routine cognitive tasks are predominantly used at the upper ends, while routine and manual tasks are more common at the middle and lower ends of the wage or skill distribution. To this end, we rank the 340 occupations in our data according to their median wages, and group them into 100 equally sized groups. Figure VII, Panel A, plots the smoothed share of men performing manual, abstract, and routine tasks in each occupation, using the 1991 wave of the German Qualification and Career Survey.¹⁶ The share of workers performing manual tasks declines monotonically with the occupational wage. Interestingly, abstract tasks are somewhat more common at the lower end of the occupational wage distribution than at the middle, and (as expected) most common at the high end. The relationship between routine task usage and occupational wages is likewise non-monotone, with the share of workers employing routine tasks being highest around the 80th wage percentile.

Following Goos and Manning (2007) and Autor, Katz, and Kearney (2008), we next test whether

¹⁶The survey is conducted jointly by the Federal Institute for Vocational Education and Training (BIBB) and the Institute for Employment (IAB), with the goal to track skill requirements of occupations. See Spitz-Oener (2006) for a detailed description of the data. Routine tasks include calculating and bookkeeping; correcting texts/data; equipping machines; shipping and transporting; and filing and archiving. Manual tasks include repairing or renovating houses/apartments/machines/vehicles; restoring of art/monuments; serving or accomodating; cleaning; and rubbish removal. Abstract tasks include research, evaluation, and planning; making plans, constructing, designing, and sketching; working out rules/prescriptions; using and interpreting rules; lobbying, coordinating, and organizing; teaching and training; selling, buying, and advertising; entertaining and presenting; employing and managing personnel.

occupations in the middle of the skill distribution experienced lower growth rates than occupations at the bottom and top of the skill distribution. We again rank occupations by their median wages, and group them into 100 equally sized groups. Figure VII, Panel B, plots the smoothed log changes in the employment shares by occupational skill for two periods: 1980-1990 (when wage inequality rose mostly at the top) and 1990-2000 (when wage inequality rose at both the top and the bottom). Both decades are characterized by polarization in employment growth: employment shares of occupations at the top of the wage distribution increased substantially in both periods. Employment shares of occupations in the middle of the wage distribution declined. Occupations at the low end of the wage distribution have experienced neither strong losses nor strong gains.

In order to more formally analyze the relationship between changes in occupational employment shares by skill percentile (measured as median wages) and wage changes by wage percentile in each decade, we estimate OLS regressions of the following form (see also Autor, Katz, and Kearney 2008):

$$\Delta E_{p\tau} = \alpha_\tau + \beta_\tau \Delta w_{p\tau} + \epsilon_{p\tau}.$$

In this expression, $\Delta E_{p\tau}$ denotes the change in log occupational employment at skill percentile p and decade τ , and $\Delta w_{p\tau}$ represents the change in the daily wage at wage percentile p in the same decade. We estimate $\beta_{80} = 0.43$ (t-value: 1.28) for the 1980s, and $\beta_{90} = 0.65$ (t-value: 1.56), suggesting that employment and wage changes are weakly positively correlated. However, in both the 1980s and the 1990s this masks important differences at the lower and upper ends of the wage distribution: below the median, the correlation between employment and wage changes is negative (we estimate $\beta_{80,p<50} = -3.05$ (3.09) for the 1980s, and $\beta_{90,p<50} = -0.97$ (1.81) for the 1990s); in contrast, above the median the correlation is positive (we estimate $\beta_{80,p>50} = 1.77$ (2.43) for the 1980s, and $\beta_{80,p>50} = 3.25$ (2.74) for the 1990s). These findings differ somewhat from those for the

United States, where changes in employment and changes in wages are strongly positively correlated in both decades throughout the entire distribution (Autor, Katz, and Kearney 2008).

These results are difficult to reconcile with a simple theory of skill-biased technological change, according to which technology symmetrically affects the bottom and the top of the wage distribution. They are consistent with a nuanced view of skill-biased technological change, according to which technology substitutes for routine tasks, but complements non-routine tasks, and thereby increases the demand for workers located mostly at the top of the wage distribution. Moreover, the negative correlation between employment and wage changes below the median speaks against a demand-based explanation for the rise in lower-tail inequality. Our previous findings are consistent with the view that the rise in lower-tail inequality may be better explained by episodic changes, most importantly the decline in unionization and the changes in the skill mix of the workforce in the 1990s.

5 Discussion and Conclusion

This paper challenges the common view that the rise in wage inequality is a phenomenon observed only in a handful of countries, such as the United States, United Kingdom, and Canada. In particular, we revisit trends in wage inequality in (West) Germany, a country that so far has been singled out as a country with a stable wage distribution. Based on a large administrative data set, we find that wage inequality in Germany has increased in the 1980s, but mostly at the top of the distribution. In the early 1990s, wage inequality started to rise at the bottom of the distribution as well. This holds for both men and women, although the rise in upper-tail inequality is somewhat more pronounced among men. Hence, while the United States and Germany experienced similar changes at the top of the distribution during the 1980s and 1990s, the two countries markedly

differ with respect to the lower end of the wage distribution: The rise in lower tail inequality that happened in the United States in the 1980s came a decade later in Germany.

We explore several explanations for the increase in inequality. Changes in workforce composition play an important role in explaining changes in the wage structure. However, they cannot fully account for the divergent path of upper and lower tail inequality in the 1980s, or for the divergent path of lower tail inequality in the 1980s and 1990s. Moreover, our results are consistent with a polarization of work: occupations that were at the top of the 1980 wage distribution experienced the largest growth rates, while occupations in the middle declined relative to occupations at the bottom. This speaks against a simple theory of skill-biased technological change, according to which technology increase the demand for skilled jobs relative to that of unskilled jobs. It is, however, consistent with a more nuanced view of technological change, according to which technology asymmetrically affect the bottom and the top of the wage distribution, by substituting for routine tasks and complementing non-routine tasks (e.g. Autor, Levy, and Murnane 2003). Since results consistent with a polarization of labor demand have now been found in three advanced countries¹⁷, this may begin to provide the unifying international evidence on technological change that so far has been absent—although more research for other advanced countries is needed to fully assess this hypothesis.

Can the polarization of work alone account for the finding that in the 1980s inequality mostly rose at the top, whereas in the 1990s inequality rose both at the top and the bottom? One piece of evidence against this hypothesis is that in Germany, employment and wage changes are negatively correlated for occupations below the median. The widening of the wage distribution at the bottom may therefore be better explained by episodic events, such as supply shocks and changes in labor market institutions. The hypothesis we put forward here is that these episodic events happened in

¹⁷See Autor, Katz, and Kearney (2006, 2008) for the United States, and Goos and Manning (2007) for the United Kingdom.

the 1980s in the United States, but in the 1990s in Germany.

First, the 1980s in the United States are characterized by an erosion of labor market institutions such as labor unions, as well as a declining minimum wage. In Germany, this process appears to have started only in the 1990s. Several papers show that these changes are important in explaining changes in inequality in the United States, in particular at the lower end of the wage distribution (e.g. DiNardo, Fortin, and Lemieux 1996; Lee 1999; and Card and DiNardo 2002). We find that between 1995 and 2004, de-unionization can explain 28% of the increase in lower-tail inequality.

Second, skill upgrading started to slow down in the United States in the early 1980s. In Germany, there is little evidence for a slowdown in skill upgrading during the 1980s. In the 1990s, however, the decline in the share of the low-skilled started to decelerate, whereas the share of the high-skilled kept increasing at a similar rate as during the 1980s. Several U.S. studies show that fluctuations in relative labor supply play an important role in explaining trends in the skill premium (e.g. Katz and Murphy 1992; Card and Lemieux 2001; Autor, Katz, and Kearney 2008). We find that fluctuations in relative supply go a long way in explaining trends in the wage differential between the medium- and low-skilled, but do not predict trends in the wage differential between the high- and the combined medium/low-skilled. Why did the slowdown in skill upgrading and the erosion in labor market institutions happen a decade earlier in the United States than in Germany? Although more research is needed on this topic, the relative increase in the share of the low-skilled that started in 1990 in Germany is likely to be a consequence of the breakdown of the communist regimes in Eastern Europe as well as the reunification of East and West Germany. These events lead to a large inflow of relatively low-skilled East Germans, Eastern Europeans, and ethnic Germans from Eastern Europe into the West-German labor market; many of these immigrants were low-skilled (see Glitz 2006 and Bauer et al. 2005 for more details).

Christian Dustmann, Univeristy College London.

Johannes Ludsteck, Institut für Arbeitsmarkt- und Berufsforschung, Nürnberg.

Uta Schönberg, University of Rochester.

A IABS, 1975-2004

A.1 The Structural Break in 1984

Starting in 1984, one-time payments, such as bonuses, are included in our wage measure (see Bender et al. (1996) for more details); as Steiner and Wagner (1998) point out, ignoring this structural break results in a spurious increase in inequality. We correct for this break in a similar way as Fitzenberger (1999). The correction is based on the idea that higher quantiles appear to be more affected by the structural break than lower quantiles, and thus have to be corrected upwards before 1984. To this end, we estimate locally weighted regressions of the wage ratio between 1982 and 1983 (i.e. before the break), and between 1983 and 1984 (i.e. after the break) on the wage percentiles in 1983 and 1984, respectively, using a bandwidth of 0.2. We then compute the correction factor as the difference between the smoothed values of the wage ratios in 1983 and 1984. In order to account for differential overall wage growth between the periods from 1982 to 1983 and from 1983 to 1984, we subtract from our correction factor the smoothed value of the wage ratio in 1983, averaged between the 2nd and 40th quantiles. In a final step, we correct wages prior to 1984 by multiplying them by 1 plus the correction factor.

A.2 Sample Selection and Variable Description

Sample Selection

In addition to the selection criteria described in Section 2, we drop wage spells of workers in apprenticeship training. We further impose the restriction that daily wages (in 1995 DM) have to be at least 20 DM. For the wage analysis, we use full-time spells only. Part-time spells are included in our relative supply measures, though with a lower weight (see below).

Wages

Our wage variable is the average daily wage. If a worker worked for more than one employer a year, we compute a weighted average, where the weights are the shares worked for each employer. Our results are employment duration-weighted: a worker who works 365 days a year gets a weight of 365, whereas a worker who works only 7 days a year gets a weight of 7. Wages are deflated by the Consumer Price Index, with 1995 as the base year. As of 1999, wages are measured in Euros; we use an exchange rate of 1 Euro = 1,95583 DM to convert Euros into Deutschmarks. Wages are censored at the highest level of earnings that are subject to social security contributions. More specifically, an individual may receive a wage increase from her current firm in the middle of the year that puts her wage above the censoring limit. In this case, the wage we observe is an employment-duration weighted average between the pre-censoring wage and the censoring limit.¹⁸ For this reason, we code wages as censored if they are three Deutschmarks below the censoring limit. Our results are very similar if we further reduce the censoring limit by a six Deutschmarks.

Education

Our education variable distinguishes three groups which we label low, medium and high. The low-skilled are workers who enter the labor market without post-secondary education. The medium-skilled are workers who completed an apprenticeship or a high school degree (*Abitur*). The high-skilled are workers who graduated from a university or college (*Universität* or *Fachhochschule*). In the raw data, the education variable is missing for 10.6% of observations. However, since our data is longitudinal, we can impute a value by looking at past and future values of the education variable. The analysis in this paper is based on the education variable IP1 provided by Fitzenberger et al. (2006). This variable is missing for only 1.3% of observations. We code missings as low-skilled. Our findings are similar if individuals with missing education are dropped.

¹⁸If the individual switches firms and earned a wage below the censoring limit at her old firm, but a wage above the censoring limit at her new firm, we observe the true wage at the old firm, and the censoring limit at the new firm.

Relative Supply Measures

Quantities supplied are measured in efficiency units. To compute the efficiency units, we calculate the mean real wage (based on imputed wages) by year, gender, education, and age. In each year, we normalize wages with the mean wage of medium-skilled men between 31 and 35. The efficiency unit for each group is computed as the arithmetic average over all years. The quantity supplied by each education group in a given year is the number of days worked in that year, multiplied by the respective efficiency unit, summed up over all workers in the education group. Part-time work is included, but weighted down by 0.67 ("long" part-time) or 0.5 ("short" part-time).

Education Wage Differentials

The medium-low and the high-medium wage differentials are based on imputed wages that assume that the error term is normally distributed with the same variance for each education and age group. This differs from our baseline imputation rule that allows for a different variance across education and age groups. We chose this restriction because the high-medium wage differential is now less sensitive to the chosen censoring limit.¹⁹ Our OLS regressions control for three education and eight age groups as well as all possible interactions; this specification therefore allows for a different wage premium by age group. The medium-low and high-medium wage premiums are computed as a weighted average of the respective premium in each age group, where the weights are the employment-weighted worker shares in each age group, averaged over the entire sample period. To compute the wage differential between the skilled (i.e. workers with a university degree) and unskilled workers (i.e. a mixture of the low- and medium-skilled), we predict age-adjusted wages for each education group in the same way. The wage of the unskilled is a weighted average of the wage of the low- and medium-skilled, where the weights are the employment-weighted worker

¹⁹For instance, in 1975 the high-low wage gap is 12 percentage points higher when the censoring limit is reduced by 6 Deutschmarks and the variance is allowed to vary by education and age. If the variance is restricted to be the same by education and age, lowering the censoring limit makes little difference.

shares of each education group, averaged over the entire sample period.

A.3 Imputation of Censored Wages

We impute censored wages in the IABS and LIAB under the assumption that the error term in the wage regression is normally distributed, with different variances for each education and each age group. To this end, we estimate censored regressions (allowing for a different variance for each education and age group) separately for each year (thereby allowing the variance in each group to vary across years). We control for all possible interactions between three education and eight age groups. For each year, we impute censored wages as the sum of the predicted wage and a random component, drawn from a normal distribution with mean zero and a separate variance for each education and age group, obtained from the standard error of the forecast. A comparison between OLS estimates based on imputed wages and tobit estimates based on censored wages shows that both the estimates and the standard errors are almost identical.

We have also imputed wages under four additional distributional assumptions. First, we continue to assume that the error term is normally distributed, but we restrict the variance to be the same across all education and age groups, or allow for a different variance for each education-age cell. Second, we assume that the upper tail of the unconditional wage distribution follows a Pareto distribution. Third, similar to United States studies based on the CPS, we replace censored observations by 1.5 times the censoring limit. Imputations based on the normal distribution, however, suggest that this imputation factor is too high, and is in fact closer to 1.2. As a fourth robustness check, we therefore replace censored wages by 1.2 times the censoring limit. Findings for these alternative imputation methods can be found in Section 1 in the on-line appendix.

We use a third data set, the GSES (German Structure of Earnings Survey), to evaluate which imputation method performs best at recovering the censored part of the wage distribution. The

GSES is a survey of about 27,000 establishments conducted by the German Federal Statistical Office. A scientific usefile is currently available for the year 2000. The main advantage of the GSES compared to the IABS is that wages are not censored. We find that the imputation method that assumes that the error term is normally distributed with a different variance by age and education works somewhat better than the other imputation methods. This method was therefore chosen for our baseline results.

B LIAB, 1995-2004

In addition to the selection criteria described in Section 2, we discard all firms for which the union variable is missing. The maximum number of establishments lost due to this restriction is 72 (around 0.8%) in 2001. We further restrict our sample to firms who employ at least one man working-full time between the ages of 21 and 60.

Most of the variables in the LIAB closely correspond to those in the IABS, but there are a few exceptions. First, the wage variable refers to the first of July in the LIAB as opposed to an annualized average, as in the IABS. Second, since the LIAB does not contain complete biographies of workers, it is impossible to impute missing observations in the education variable. We therefore code missing observations as an additional education category.

Table B.1 compares the median as well as the 85/50 and 50/15 wage gaps in the IABS and LIAB. Results refer to men only. The wage measure in the LIAB refers to the 1st of July. In order to facilitate the comparison between the two data sets, the wage measure in the IABS also refers to this date. Not surprisingly, the structure of wages is very similar in the two data sets. One reason for why the two data sets do not give identical results—although they are based on the same social security records—is that the weights in the LIAB were constructed to be representative for

the population as a whole, not for our estimation sample of prime-aged men.

References

- [1] Abraham, Katherine G., and Susan N. Houseman, "Earnings Inequality in Germany," in *Differences and Changes in Wage Structures*, Freeman, Richard B. and Lawrence F. Katz, eds. (The University of Chicago Press, Chicago and London, 1995).
- [2] Acemoglu, Daron, "Technical Change, Inequality and the Labor Market," *Journal of Economic Literature*, 40 (2002), 7-72.
- [3] Acemoglu, Daron, "Cross-Country Inequality Trends," *Economic Journal*, 113 (2003), F121-F149.
- [4] Autor, David H., Lawrence F. Katz, and Melissa S. Kearney, "Trends in U.S. Wage Inequality: Re-Assessing the Revisionists," *Review of Economics and Statistics*, forthcoming, 2008.
- [5] Autor, David H., Katz, Lawrence F. and Melissa S. Kearney, "The Polarization of the U.S. Labor Market," *American Economic Review Papers and Proceedings*, 96 (2006), 189-194.
- [6] Autor, David H., Levy, Frank, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Investigation," *Quarterly Journal of Economics*, 118 (2003), 1279-1333.
- [7] Bauer, Thomas, Dietz, Barbara, Zimmermann, Klaus F. and E. Zwintz, "German Migration: Development, Assimilation, and Labour Market Effects," in *European Migration. What Do We Know?*, Klaus F. Zimmermann, ed. (Oxford University Press, Oxford).
- [8] Bender, Stefan, Hilzendegen, Jürgen, Rohwer, Götz, and Helmut Rudolph, "Die Beschäftigtenstichprobe 1975-1990," *Beiträge zur Arbeitsmarkt- und Berufsforschung*, 197 (1996).
- [9] Boudarbat, Brahim, Lemieux, Thomas, and W. Craig Riddell, "Recent Trends in Wage Inequality and the Wage Structure in Canada," in *Dimensions of Inequality in Canada*, Green, David A. and Jonathan R. Kesselman, eds. (Vancouver: UBC Press, 2006).
- [10] Bound, John, and George Johnson, "Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations", *American Economic Review*, 82 (1992), 371-392.
- [11] Card, David, and John E. DiNardo, "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics*, 20 (2002), 733-783.
- [12] Card, David, Francis Kramarz, and Thomas Lemieux, "Changes in the Relative Structure of Wages and Employment: A Comparison of the United States, Canada, and France," *Canadian Journal of Economics*, 32 (1999), 843-877.
- [13] Card, David, and Thomas Lemieux, "Can Falling Supply Explain the Rising Return to College for Younger Men?" *Quarterly Journal of Economics*, 116 (2001), 705-746.
- [14] Card, David, Thomas Lemieux, and W. Craig Riddell, "Unions and Wage Inequality," *Journal of Labor Research* (2004), Symposium on "What Do Unions Do – A Retrospective After Two Decades".

- [15] DiNardo, John E., Nicole Fortin, and Thomas Lemieux, "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 64 (1996), 1001-1044.
- [16] DiNardo, John E., and David S. Lee, "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001", *Quarterly Journal of Economics*, 119 (2004), 1383-1441.
- [17] Dew-Becker, Ian, and Robert J. Gordon, "Where Did the Productivity Growth Go? Inflation Dynamics and the Distribution of Income", *Brookings Papers on Economic Activity*, 2 (2005), 67-127.
- [18] Dustmann, Christian, and Uta Schönberg, "Training and Union Wages", *Review of Economics and Statistics*, forthcoming, 2008.
- [19] Fitzenberger, Bernd, "Wages and Employment Across Skill Groups: An Analysis for West Germany," Heidelberg, Physica/Springer, 1999.
- [20] Fitzenberger, Bernd, and Karsten Kohn, "Gleicher Lohn für gleiche Arbeit? Zum Zusammenhang zwischen Gewerkschaftsmitgliedschaft und Lohnstruktur in Westdeutschland 1985-1997," *Journal for Labour Market Research*, 38 (2005), 125-146.
- [21] Fitzenberger, Bernd, and Karsten Kohn, "Skill Wage Premia, Employment, and Cohort Effects: Are Workers in Germany All of the Same Type?" IZA Discussion Paper No. 2185, 2006.
- [22] Fitzenberger, Bernd, Osikominu, Anderonke and Robert Völter, "Imputation Rules to Improve the Education Variable in the IAB Employment Subsample," *Schmollers Jahrbuch (Journal of the Applied Social Sciences)*, forthcoming, 2007.
- [23] Freeman, Richard B., and Lawrence F. Katz, "Introduction and Summary", in *Differences and Changes in Wage Structure*, Freeman, Richard B., and Lawrence F. Katz, eds. (Chicago: The University of Chicago Press, 1995).
- [24] Gerlach, Knut, and Gesine Stephan, "Bargaining Regimes and Wage Dispersion," *Jahrbücher fuer Nationalökonomie und Statistik*, forthcoming, 2007.
- [25] Gerlach, Knut, and Gesine Stephan, "Wage Settlements and Wage Setting – Results from a Multi-Level Model," *Applied Economics*, 37 (2005), 2297-2306.
- [26] Gernandt, Johannes, and Friedhelm Pfeiffer, "Rising Wage Inequality in Germany," ZEW Discussion Paper No. 06-019, 2006.
- [27] Glitz, Albrecht, "The Labour Market Impact of Immigration: Quasi-Experimental Evidence," University of Pompeu Fabra, mimeo, 2007.
- [28] Goldin, Claudia, and Lawrence F. Katz, "The Race Between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005", NBER Working Paper No. 12984, 2007a.
- [29] Goldin, Claudia, and Lawrence F. Katz, "Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing", *Brookings Papers on Economic Activity*, 2 (2007b), 135-165.

- [30] Goldin, Claudia, and Lawrence F. Katz, "The Race between Education and Technology", Harvard University Press, 2008.
- [31] Goos, Maarten, and Alan Manning, "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics*, 89 (2007), 118-133.
- [32] Robert J. Gordon, and Ian Dew-Becker, "Selected Issues in the Rise of Income Inequality", *Brookings Papers on Economic Activity*, 2 (2007), 169-192.
- [33] Gosling, Amanda, Machin, Stephen and Costas Meghir, "The Changing Distribution of Male Wages, 1966-1992," *Review of Economic Studies*, 67 (2000), 635-666.
- [34] Gottschalk, Peter, and T. Mary Joyce, "Cross-National Comparisons of Earnings and Income Inequality, Market and Institutional Factors", *Review of Economics and Statistics*, 80 (1998), 489-579.
- [35] Herrlinger, Dagmar, Müller, Dana, and Lutz Bellmann, "Codebuch zum IAB Betriebspanel. Version 1: Längsschnitt 1993-2001 (5. überarbeitete Auflage)," FDZ Datenreport Nr. 5 (2005), Nürnberg.
- [36] Juhn, Chinhui, Murphy, Kevin M., and Brooks Pierce, "Wage Inequality and the Rise in Returns to Skill", *Journal of Political Economy*, 101 (1993), 410-441.
- [37] Katz, Lawrence F., and David H. Autor, "Changes in the Wage Structure and Earnings Inequality," in *Handbook of Labor Economics*, Ashenfelter, Orley and David Card, eds. (Volume 3, North Holland, 1999).
- [38] Katz, Lawrence F. and Kevin M. Murphy, "Changes in Relative Wages, 1963-87: Supply and Demand Factors", *Quarterly Journal of Economics*, 107 (1993), 35-78.
- [39] Kohn, Karsten, "Rising Wage Dispersion, After All! The German Wage Structure at the Turn of the Century," IZA Discussion Paper No. 2098, 2006.
- [40] Krugman, Paul, "Past and Prospective Causes of Unemployment," *Economic Review*, Federal Reserve Bank of Kansas City (1994), 23-43.
- [41] Lee, David S., "Wage Inequality in the United States in the 1980s: Rising Dispersion or Falling Minimum Wage?" *Quarterly Journal of Economics*, 114 (1999), 977-1023.
- [42] Lemieux, Thomas, "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *American Economic Review*, 96 (2006a), 1-64.
- [43] Lemieux, Thomas, "Post-Secondary Education and Increasing Wage Inequality", *American Economic Review Papers and Proceedings*, 96 (2006b), 1-23.
- [44] Lemieux, Thomas, "The Changing Nature of Wage Inequality," *Journal of Population Economics*, forthcoming, 2008.
- [45] Levy, Frank and Richard J. Murnane, "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations," *Journal of Economic Literature*, 30 (1992), 1333-1381.

- [46] Möller, Joachim, "Die Entwicklung der Lohnspreizung in West- und Ostdeutschland," in *Institutionen, Löhne, und Beschäftigung*, Bellmann, Lutz, Olaf Hübler, Wolfgang Meyer, and Gesine Stephan, eds. (Beiträge zur Arbeitsmarkt- und Berufsforschung, Nürnberg, 294, 2005).
- [47] Mulligan, Casey B. and Yona Rubinstein, "Selection, Investment, and Women's Relative Wages over Time," *Quarterly Journal of Economics*, 123 (2008), forthcoming.
- [48] Murphy, Kevin M. and Finis Welch, "The Structure of Wages," *Quarterly Journal of Economics*, 107 (1992), 285-326.
- [49] OECD, "Earnings inequality, low-paid employment and earnings mobility," in *OECD Employment Outlook*, 1996, 59-107.
- [50] Piketty, Thomas and Emmanuel Saez, "Income Inequality in the United States, 1913-1998", *Quarterly Journal of Economics*, 118 (2003), 1-39.
- [51] Prasad, Eswar S., "The Unbearable Stability of the German Wage Structure: Evidence and Interpretation," *IMF Staff Papers*, 51 (2004), 354-385.
- [52] Riphahn, Regina T, "Bruttoeinkommensverteilung in Deutschland 1984-1999 und die wachsende Ungleichheit unter ausländischen Erwerbstätigen," in *Wechselwirkungen zwischen Arbeitsmarkt und Sozialer Sicherung II, Schriften des Vereins für Sozialpolitik*, Schmähl, Winfried, ed., 294 (2003), 135-174.
- [53] Saez, Emmanuel and Michael R. Veall, "The Evolution of High Incomes in Northern America: Lessons from the Canadian Experience," *American Economic Review*, 95 (2005), 831-849.
- [54] Schnabel, Claus and Joachim Wagner, "The Persistent Decline in Unionization in Western and Eastern Germany, 1980-2004: What Can We Learn From a Decomposition Analysis?" IZA Discussion Paper No. 2388, 2006.
- [55] Spitz-Oener, Alexandra, "Technical Change, Job Tasks and Rising Educational Demands: Looking Outside the Wage Structure," *Journal of Labor Economics*, 24 (2006), 235-270.
- [56] Steiner, Viktor and Kersten Wagner, "Has Earnings Inequality in Germany Changed in the 1980s?" *Zeitschrift für Wirtschafts- und Sozialwissenschaften*, 118 (1998), 29-59.