

Koomen, Miriam; Backes-Gellner, Uschi

Working Paper

Occupational Tasks and Wage Inequality in Germany: A Decomposition Analysis

IZA Discussion Papers, No. 15702

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Koomen, Miriam; Backes-Gellner, Uschi (2022) : Occupational Tasks and Wage Inequality in Germany: A Decomposition Analysis, IZA Discussion Papers, No. 15702, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/267439>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 15702

**Occupational Tasks and Wage Inequality
in Germany: A Decomposition Analysis**

Miriam Koomen
Uschi Backes-Gellner

NOVEMBER 2022

DISCUSSION PAPER SERIES

IZA DP No. 15702

Occupational Tasks and Wage Inequality in Germany: A Decomposition Analysis

Miriam Koomen

Swiss National Bank

Uschi Backes-Gellner

University of Zurich and IZA

NOVEMBER 2022

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Occupational Tasks and Wage Inequality in Germany: A Decomposition Analysis*

We study the role of occupational tasks as drivers of West German wage inequality. We match administrative wage data with longitudinal task data, which allows us to account for within-occupation changes in task content over time. We run RIF regression-based decompositions to quantify the contribution of changes in the returns to tasks to overall changes in the wage distribution from 1978 to 2006. We find that changes in the returns to tasks explain up to half of the increase in wage inequality since the 1990s, both at the top and the bottom of the wage distribution. Specifically, abstract tasks drive the upper wage gap, while interactive and routine tasks drive the lower wage gap. Importantly, we find low-wage occupations to have the highest routine task intensity. The association between occupational tasks and West German wage inequality is thus both stronger and different than prior research has found.

JEL Classification: C55, D63, E24, J31

Keywords: wage inequality, skills, tasks, routine-biased technical change, decomposition analysis, RIF regression

Corresponding author:

Miriam Koomen
Swiss National Bank
P.O. Box
8022 Zurich
Switzerland
E-mail: miriam.koomen@snb.ch

* We thank two anonymous referees for helpful comments and suggestions. We also thank Eric Bettinger, Tor Eriksson, participants of the Colloquium on Personnel Economics, the VFS Annual Conference, as well as seminar participants at the University of California, Berkeley and the University of Zurich for valuable comments on this as well as earlier versions of the study. We would also like to acknowledge the inspiring and helpful discussions with the late Ed Lazear. Miriam Koomen is grateful to the Swiss National Science Foundation for its financial assistance through the Doc.Mobility scholarship (project number: P1ZHP1 155498). The views expressed in this study are the authors' and do not necessarily reflect those of the Swiss National Bank. All remaining errors are our own.

1 Introduction

High levels of wage inequality continue to persist and, in some countries, they have been exacerbated by the Covid-19 pandemic (e.g. Palomino et al., 2020; Bonacini et al., 2021). High levels of inequality are not only incompatible with widely held norms of social justice, they can also fuel social tensions and might pose a threat to economic growth. Not surprisingly, the drivers of wage inequality and possible remedies have become some of the most hotly debated issues by policymakers and researchers alike.

In this paper, we investigate the role of occupational tasks in driving male wage inequality in West Germany.¹ A distinctive feature of our paper is the use of administrative wage data matched with task data from employee surveys which capture changes in task content within occupations. Economists have long recognised the importance of human capital in explaining workers' wage differentials (Becker, 1962; Mincer, 1974). In the early 2000s, researchers proposed a new approach to measuring human capital using data on tasks performed on the job (Autor et al., 2003). This task-based approach defines occupations as bundles of tasks and links workers' wages to occupational tasks. The main hypothesis is that recent technological change is biased toward replacing labour in routine tasks, leading to a decrease in wages and employment shares in routine-task intensive occupations.

Several studies have drawn on this hypothesis to explain the polarisation of the U.S. labour market since the late 1990s (e.g. Autor et al., 2006, 2008;

¹ A separate study on female wage inequality is outside the scope of this paper. The importance of using gender-disaggregated data in this line of research—highlighted by contributions such as Black and Spitz-Oener (2010) and Cortes et al. (2018)—stems from two major aspects: First, due to a strong occupational gender segregation, differential selection into work and stark differences in fulltime work patterns across genders, wage determinants and patterns will starkly differ across genders. Second, activities performed on the job vary significantly across gender. In line with previous research, we therefore focus on males.

Autor and Dorn, 2013). Some researchers argue that the routinisation hypothesis also (partly) explains polarisation tendencies in Western European countries, most notably in Germany and in the UK (e.g. Goos and Manning, 2007; Dustmann et al., 2009; Goos et al., 2009, 2014). Contributions in this field focus on employment or wage polarisation, or both. The term employment polarisation describes the phenomenon that the employment shares of high- and low-skill occupations have been increasing at the expense of employment shares in middle-skill occupations. In the same vein, wage polarisation means that wages at the top and the bottom have been increasing faster than those in the middle of the distribution. In other words, the pay difference between high-wage and middle-wage jobs has increased, while the difference between middle- and lower-wages has been shrinking. In this paper, we focus on wage polarisation.

However, the majority of studies investigating the routinisation hypothesis suffer from several shortcomings. First, most are of a descriptive nature: They simply document changes in wage and employment patterns of occupations and juxtapose them to the routine-task intensity of those occupations. Second, most are silent on whether tasks *within* occupations have changed.² Third, most studies lack a discussion on how to measure routine tasks. However, such a discussion is nontrivial and wrong classifications may lead to wrong conclusions (Autor, 2013).

In this paper, we try to overcome these shortcomings. First, we apply a rigorous econometric framework and move beyond merely comparing patterns. Specifically, we use a two-stage decomposition procedure introduced by Firpo et al. (2009), which allows for (i) accounting of multiple explana-

² Notable exceptions are a few recent contributions by Atalay et al. (2018, 2019, 2020), Deming and Kahn (2018), Hershbein and Kahn (2018), and Michaels et al. (2019), who use (online) vacancy postings or contemporary and historical versions of the DOT to detect changes in the composition of tasks.

tory factors in driving wage inequality in a regression framework, and (ii) investigating changes along the entire wage distribution. The approach is similar to the Oaxaca-Blinder (OB) decomposition (Oaxaca, 1973; Blinder, 1973) in that it decomposes wage differentials in a counterfactual manner. However, it is more flexible, allowing for the decomposition of any distributional statistic. We may thus uncover important differences at different parts of the distribution, which the mean analysis of the OB decomposition would overlook. Second, we account explicitly for changes in task content within occupations. Third, we introduce a new classification of tasks, focusing the discussion on how to identify routine tasks.

Our task data are the BIBB/IAB and BIBB/BAuA Employment Surveys, representative cross-section surveys conducted roughly every seven years since 1979. For each survey wave, we can directly measure the task content of occupations; that is, survey participants indicate the activities they perform on the job. We characterise occupational tasks by four categories: abstract (such as researching), interactive (such as negotiating), manual (such as repairing) and routine (such as operating machinery). We aggregate this worker-level data into occupational cells and use the occupational code to match our task data to our wage data, which is the Sample of Integrated Labour Market Biographies (SIAB). The SIAB is a 2% random sample of all social security records in Germany, covering the employment histories of about 1.5 million individuals from 1975 through 2008.

We provide descriptive evidence that the task content within occupations varies over time. This evidence provides support for Autor's (2013) assessment that new technologies have not simply wiped out certain occupations, but fundamentally changed the tasks performed in these occupations. Generally, we find that the shares of routine and manual tasks decrease over

time, while the shares of abstract and interactive tasks increase. This trend holds across wage deciles, while its magnitude differs at different parts of the distribution. We thus speak to the recent literature investigating *within-occupation* task changes in the US labour market (Atalay et al., 2018, 2019, 2020; Deming and Kahn, 2018; Hershbein and Kahn, 2018; Michaels et al., 2019). In a departure from prior research, we find that the share of routine tasks is largest in occupations at the lowest tail of the wage distribution and that this share slowly declines with increasing wage. We can thus not confirm that occupations in the middle of the wage distribution have the largest routine-task intensity.

Related to this point, we show descriptively that changes in the West German wage structure are different to those observed in Anglo-Saxon countries. Instead of a U-shaped polarisation of wages since the early 1980s (Goos and Manning, 2007; Autor et al., 2008), we find a uniform increase in wage inequality with widening wage gaps at the top starting in the early 1980s, and at the bottom starting only in the 1990s. Our findings are in line with studies for West Germany that use register data (Fitzenberger, 1999; Kohn, 2006; Dustmann et al., 2009, 2014; Riphahn and Schnitzlein, 2016; Biewen et al., 2017; Antonczyk et al., 2018; Baumgarten et al., 2020). Changes in the wage structure are thus concentrated both at the top and the bottom, while the middle remains relatively stable.

Our decomposition analyses reveal several additional findings. First, we find composition effects, i.e. changes due to varying worker characteristics, not to matter much. This result somewhat contradicts previous studies (Antonczyk et al., 2009; Dustmann et al., 2009; Biewen et al., 2017), which point toward composition effects explaining some of the inequality increase. However, those studies use either different data or look at different time horizons

compared to us, which may explain the differing findings. Contrarily, using the Integrated Employment Biographies, which contains nearly all private sector employees in Germany, Card et al. (2013) also find that compositional effects do not matter greatly.

Second, we find wage structure effects, i.e. changes in the returns to those characteristics, to be important drivers of wage inequality. Specifically, from the late 1970s until the early 1990s, we identify changes in the returns to education and experience as major drivers of wage inequality. Occupational tasks start to matter in the 1990s only. For that period, we find changing returns to abstract tasks to drive upper-tail inequality. Technological change may have increased the demand for and thus the returns to abstract tasks. An example could be advancements in communication technology that benefit workers in occupations with large shares of abstract tasks. This assessment is in line Dustmann et al. (2009), who argue that technological change is an important driver for upper-tail wage inequality.

For lower-tail inequality, we find changing returns to routine tasks as a major driver. The intuition here is that workers at the bottom of the wage distribution have a comparative advantage in occupations that are intensive in routine tasks, while the demand for these tasks has declined. This result differs from earlier work focusing on labour markets where occupations in the middle of the distribution are characterised by high routine intensity. This difference may be the result of our more rigorous definition of routine tasks. For example, Katz (2014) highlights for the US that many crafts occupations that are commonly classified as routine-manual have been faring well in terms of labour market performance. This could point to an inappropriate classification of routine-manual tasks.

In the 2000s, we find abstract and routine tasks not to matter greatly

anymore, while changing returns to interactive tasks are now driving lower-tail inequality. This finding could be related to previous empirical findings that some tasks are not easily offshored and thus less affected by decreasing wages (Firpo et al., 2011; Autor and Dorn, 2013). Occupations with large shares of ‘non-offshorable’ tasks may thus have experienced an increasing demand and thus increasing returns during the 2000s. Our results also speak to the relatively few studies that focus on the rising complementarity between cognitive and social skills (Borghans et al., 2014; Weinberger, 2014; Deming, 2017).

Furthermore, in the 2000s, we find a role for firm characteristics to explain increases in lower-tail wage inequality. We thus confirm previous findings, which point to the importance of firm heterogeneity in explaining increases in inequality (Dustmann et al., 2009; Antonczyk et al., 2010; Card et al., 2013).

Our study is closest to Antonczyk et al. (2009), who also run a decomposition analysis for Germany. However, our study differs in several important aspects. First, they use the classification of tasks introduced by Spitz-Oener (2006), while we develop a new task categorisation. Second, they investigate the years 1999 to 2006, while our time horizon spans from 1978 to 2006. This longer horizon provides the opportunity for a detailed analysis of task changes within occupations. Third, they solely rely on survey data, while we use wage register data.

Our study adds several new insights to the discussion on the drivers of wage inequality. First, we show how the time-varying characteristics of occupations in Germany matter. Any discussion on how technological change affects the nature of work should thus explicitly account for changing task requirements within occupations. Second, we show that a simple compari-

son of wage and task patterns may lead to wrong conclusions. For example, in the early 2000s increases in upper-tail wage inequality coincide with increases in the share of abstract tasks. However, our RIF regression-based decompositions reveal that these patterns are largely unrelated. Only a regression framework allows controlling for confounding factors and uncovering the contribution of changes in occupational tasks to changes in the wage distribution.

Our study thus allows for a better understanding of why some occupations have incurred sharp wage losses, while others have experienced strong gains. Such an understanding is also important for the policy debate on whether workers today are well equipped to face the challenges arising from technological change and on whether education curricula may need to be revised.

2 Technological change and wage setting in occupations

Economists have long been concerned with the question of how technological change affects the wage distribution. In the early 1990s, a series of papers introduced the concept of a skill-biased technological change (SBTC), defined as a shift in the production technology that favours skilled over unskilled workers (Bound and Johnson, 1992; Katz and Murphy, 1992; Levy and Murnane, 1992; Juhn et al., 1993; Berman et al., 1994). The SBTC hypothesis posits that in the 1980s a burst of new technology shifted the productivity gap between high- and low-skilled workers, causing a pronounced rise in wage inequality. Most of this early literature follows a Mincerian wage-setting model whereby wages are determined only by observed and unobserved skill. The relative demand for skill increases over time because changes in technology are assumed to be skill-biased.

However, the SBTC hypothesis has difficulties explaining developments in the U.S. wage structure since the late 1990s. Besides the relative growth of wages at the upper deciles, wages at the lower deciles started increasing as well. The leading explanation for this wage polarisation is the routinisation hypothesis, or routine-biased technological change (RBTC), which relies on the assumption that technological change substitutes middle-wage occupations, which are routine-task intensive, and complements high-wage occupations (starting with Autor et al., 2003).

The Mincerian wage-setting model is ill-suited in this case, as it cannot explain differential wage changes at different parts of the distribution. To address this shortcoming, Acemoglu and Autor (2011) propose a Ricardian wage-setting model, whereby workers with different skill endowments perform different types of tasks.³

The crucial innovation in their model is that a worker of a given skill can perform a variety of tasks. They define tasks as units of work activity that produce output and skills as workers' endowments for performing various tasks. Workers apply their skills to carry out tasks in exchange for wages. This distinction is unnecessary if workers of a given skill always perform the same set of tasks. However, it becomes relevant when shifts in technology change the assignment of skills to tasks.

In their model, the equilibrium allocation of workers to tasks is determined by two thresholds: a lower one where low-skill workers perform all tasks below the threshold, and a higher one where high-skill workers perform all tasks above the threshold. Medium-skill workers perform all intermediate tasks. As technological change may modify the assignment of skills to

³ In applying this assignment model between jobs and skills to the study of technological change and wage inequality, Acemoglu and Autor (2011) largely build on Saint-Paul (2001, 2008) and Costinot and Vogel (2010).

tasks, the simple distinction between high, middle and low skills allows for non-monotone movements in the wage distribution.

The law of one price is assumed to hold within each skill group, i.e. wages are equalised across occupations conditional on skill. This equilibrium implies a simple structure of comparative advantage with an endogenous allocation of workers to tasks. High-skill workers are more productive in more complex tasks than medium-skill workers and medium-skill workers in turn are more productive than low-skill workers.

Acemoglu and Autor's (2011) task assignment model provides an explanation for both the increase in relative wages of high-skill workers and for wage polarisation. First, if technological change is skill-biased, it increases the relative productivity of high-skill tasks. Accordingly, the wage of lower-skill workers increases less in relative terms than that of high-skill workers, thereby increasing inequality. Second, technological change may reduce the wage of middle-skill workers through task reallocation. Specifically, some tasks performed by middle-skill workers may be shifted to high-skill workers, who thus expand their set of tasks. However, a corresponding shift of low-skill tasks to middle-skill workers does not take place. As a result, the relative wage paid to workers performing these formerly middle-skill tasks increases since they are now performed by the more productive high-skill workers. Conversely, the relative wage of medium-skill workers could fall as they may be reallocated to tasks for which they have lower comparative advantage.

Expanding on this model, Autor and Dorn (2013) clarify that the wage effects at the bottom are ambiguous, depending on whether low-skill occupations are complements to or substitutes of high-skill occupations. Technological change may thus also be consistent with rising wage inequality in

the middle and at the bottom.

Drawing on Acemoglu and Autor’s (2011) task assignment model, Firpo et al. (2011) propose a Roy-type wage-setting model and relax the assumption of the law of one price. They allow wages to vary across occupations conditional on the tasks workers perform. Firpo et al. (2011) argue that their model better fits the patterns observed in the data, namely that low-, middle- and high-skill workers overlap in the tasks they perform.

In the Firpo et al. (2011) model, technological change has a different impact on wages in different occupations, depending on the occupation’s relative task shares. For example, if automated machines decrease the marginal product of the task ‘calculating,’ then the level and the dispersion of wages for occupations with a high intensity of ‘calculating’ should decrease. However, if ‘calculating’ is not highly valued in some other occupations, then the change in the price for ‘calculating’ only marginally affects those occupations.

Formally, Firpo et al. (2011) assume that each worker i is characterised by a k -dimension set of skills $S_i = [S_{i1}, S_{i2}, \dots, S_{iK}]$. The occupation-specific output Y_{ijt} produced by worker i in occupation j at time t is assumed to depend linearly on skills. Each occupation requires different shares of different skills. Researchers, for example, require a large share of cognitive skills and a rather small share of manual skills. Machinists require the opposite.

Assuming that wages are set competitively, Firpo et al. (2011) derive the following wage equation:

$$w_{ijt} = \theta_{jt} + \sum_{k=1}^K r_{jkt} S_{ik} + u_{ijt} \quad (1)$$

where w_{ijt} is the wage of worker i in occupation j at time t , r_{jkt} are the

returns to skill component k specific to occupation j , S_{ik} are the skill components embodied in worker i , and θ_{jt} is a base payment that workers receive in occupation j regardless of their skills. Finally, u_{ijt} is the idiosyncratic error term.

We draw on Firpo et al.'s (2011) Roy-type model in our paper, but modify it in one important aspect. We include the assumption by Acemoglu and Autor (2011) that the assignment of skills to tasks changes over time. This assumption implies that the task content within occupations changes over time. Rather than the time-constant skill components S , we assume wages to linearly depend on the time-varying task component T . Wage equation (1) thus modifies to:

$$w_{ijt} = \theta_{jt} + \sum_{k=1}^K r_{jkt} T_{ikt} + u_{ijt} \quad (2)$$

where T_{ikt} are defined as the task components of worker i at time t . By adding this time dimension, we can explicitly account for changes in task composition.

Unlike Acemoglu and Autor (2011), Firpo et al. (2011) do not develop a full model of the labour market that would show the allocation of skills to tasks and the wage set in equilibrium. This distinction is also not necessary for the empirical analysis. However, the implicit assumption is that the assignment of skills to tasks is endogenous. One important caveat from this assumption is that a simple regression of wages on tasks may be difficult to interpret. Consequently, the task coefficients resulting from such a regression should not be interpreted causally. Furthermore, as the allocation of tasks depends on workers' human capital, which is assumed to be quasi-fix and determined prior to labour market entry, controlling for workers' human capital in such a regression framework is crucial.

3 Data

3.1 Task data

To measure occupational tasks, we use the BIBB/IAB and BIBB/BAuA Employment Surveys on Qualification and Working Conditions (hereafter: Employment Surveys), which are representative cross-section surveys that are conducted roughly every seven years. We use five survey waves (1979, 1985/86, 1991/92, 1998/99 and 2006), with each wave covering about 30,000 individuals.

These surveys have several advantages over the O*NET, the data most commonly used for studying occupational tasks. First, the occupational codes are constant across waves, which makes them particularly suitable to analyse changes in occupational tasks over time. Second, because workers themselves report on the activities performed on the job, the surveys should accurately report changes in task content. A frequently voiced concern with early versions of the O*NET, which rely on occupational experts, is that task requirements may be over- or underestimated. Third, contrary to most previous studies, our task data has a time dimension. For each occupation, we can compute five task portfolios at five different points in time.

A major disadvantage of the Employment Surveys is that the questions on tasks performed on the job change somewhat over time. These changes are small for consecutive waves, but they increase over time. For this reason, we refrain from running regressions across the whole observation period, i.e. from 1979 until 2006, and instead focus on year-pairs, comparing changes in wages and tasks between two consecutive survey waves.

Because our unit of analysis is the occupation, we aggregate the worker-level data into occupational cells and use group means for our decomposition

analyses. Following Spitz-Oener (2006) and Antonczyk et al. (2009), we first define tasks at the individual level and then aggregate them to the occupational level. For individual i in period p , we define task share t as:

$$t_{icp} = \frac{\text{number of tasks in category } c \text{ performed by } i \text{ in period } p}{\text{total number of tasks in category } c \text{ in period } p} \quad (3)$$

where $p = 1979, 1985/86, 1991/92, 1998/99$ and 2006 ; and task category $c = 1$ (abstract tasks), $c = 2$ (interactive tasks), $c = 3$ (manual tasks) and $c = 4$ (routine tasks). This definition measures the share of tasks a worker reports to perform among all tasks of type c . To generate a task category at the occupational level, we sum the individual task categories t_{icp} in each occupation and divide them by the number of observations in that occupation. We use the same 61 two-digit occupations across all survey waves, which are based on a classification system by the Federal Employment Office. The aggregation at the two-digit level decreases well-known measurement error issues of occupational classifications in survey data and allows us to match the task data to our wage data.

3.2 Wage data

Our wage data is the Sample of Integrated Labour Market Biographies (SIAB), a 2% random sample of administrative social security records in Germany from 1975 through 2008, covering the employment histories of more than 1.5 million individuals. The SIAB is representative of all individuals covered by social security, which is about 80% of the German workforce. The data includes gross daily wages, days worked at each job

in a year, as well as information on education, occupation, industry and employment status. The SIAB reports 2-digit occupational codes that are constant over time and similar to the classification used in the Employment Surveys. We can thus match the two data sets over the occupational codes, distinguishing 61 occupations.⁴

We analyse log wage differentials between two periods, relating changes in wages to changes in occupational tasks. To improve the precision of our estimates, we follow Lemieux (2006) and Firpo et al. (2011) and pool years of data at both the beginning and the end of the periods we analyse. We choose the years to match the waves of the Employment Surveys. Specifically, we pool the 1978 and 1979 data as the base, and 1985 and 1986 as the end years for the first period; 1985 and 1986 as the base, and 1991 and 1992 as the end for the second period; 1991 and 1992, and 1998 and 1999 for the third period; and, finally, 1998 and 1999, and 2005 and 2006 for the last period.

The main disadvantage of the SIAB is that, as in many administrative data sets, the data is top-coded at the highest level of earnings. Statutory pension insurance contributions are paid as a fixed earnings share only up to an earnings threshold, earnings exceeding this threshold only report the threshold. Each year between 5% and 15% of the wage distribution of our sample is censored. To solve the censoring problem, we follow a method proposed by Gartner (2005). Specifically, we use a series of Tobit models—fit by education, potential experience, industry, occupation and region—to stochastically impute the upper tail of the wage distribution. Using the estimated parameters from these models, we replace each censored wage value with a random draw from the upper tail of the appropriate conditional wage distribution.⁵

⁴ Figure 5 in Section 4 spells out these occupations.

⁵ The impact of this imputation procedure is illustrated in the Appendix. Figure A1 plots

Although we believe that this imputation procedure works reasonably well, a natural concern is that our analysis would yield different results if we used a different technique. Following Dustmann et al. (2009), we use the imputed wage for the descriptive analyses, but focus on the uncensored wage for our decomposition analyses. In the US and Europe, a large portion of the increase in wage inequality has been above the 85th percentile (Autor et al., 2008; Hoffmann et al., 2020). Thus, our results are likely to provide a lower-bound estimate of inequality.

3.3 Sample

Following previous studies (Spitz-Oener (2006) and Antonczyk et al. (2009) for the survey data; Dustmann et al. (2009), Card et al. (2013), and Riphahn and Schnitzlein (2016) for the register data), we apply several sampling restrictions. Importantly, to avoid introducing bias, we apply the same sampling restrictions to both data sources.

In particular, we focus on full-time employed males aged between 20 and 60, who are subject to social security contributions. We exclude females for several reasons. First, because a relatively high fraction work part-time. Second, as female labour force participation rose considerably during the 1980s and 1990s, it may have changed the selection of women into work, which may have had an independent impact on the female wage structure (Dustmann et al., 2009). Third, significant occupational gender segregation as well as gender trends in task performance exist (Black and Spitz-Oener, 2010; Cortes et al., 2018). This would warrant a separate analysis for men and women, which is beyond the scope of this paper.

In addition, we drop all observations for which daily wages are below 12 log wages using the imputed values for the censored observations versus the data without adjustment for censoring.

euros, because our wage data lacks information on hours worked.⁶ We also exclude the self-employed and the unemployed, and workers in agricultural occupations. Furthermore, we use West German data only, because both the level and the distribution of wages differ substantially between East and West Germany. Finally, because the sampling population is not uniform across waves of the Employment Surveys, we restrict our sample to German nationals.

3.4 Task classification

To develop task categories, we build on Spitz-Oener (2006), but make some important departures. Spitz-Oener (2006) first used the Employment Surveys to apply the task-based approach by Autor et al. (2003) to the German case. She was the first to document the changing task content of occupations, distinguishing routine manual, non-routine manual, routine cognitive, non-routine analytical, and non-routine interactive tasks. Her task classification has been highly influential and widely used in many subsequent studies.

However, this classification is not uncontroversial. Recent contributions caution that imposing the Autor et al. (2003) framework to data other than the O*NET may lead to misclassification, as most O*NET categories have no direct counterpart in other data (Green, 2012; Rohrbach-Schmidt and Tiemann, 2013). An example is ‘calculating,’ which is classified as routine cognitive in Spitz-Oener (2006), but could well be a non-routine cognitive task given that the GED Math measure in the Autor et al. (2003)

⁶ This restriction affects a very small number of observations in each year, with a maximum of 0.007% in 2005. Nonetheless, to investigate whether our results are sensitive to this cut-off, we rerun our decomposition analyses including observations with daily wages of less than 12 euros. We find somewhat stronger coefficients for lower-tail inequality, while our results for upper-tail wage inequality remain unaffected. By applying this restriction, we thus provide a lower bound estimate of wage inequality.

paper includes the item ‘adds and subtracts 2-digit numbers.’ Generally, while distinguishing manual from cognitive (or abstract) tasks is relatively straightforward, identifying routine tasks within these two task dimensions is particularly challenging (Rohrbach-Schmidt and Tiemann, 2013).

Autor (2013) points out a further difficulty: The common notion of a routine task may not necessarily mean that that task is codifiable or programmable. He provides an illustrative example whereby workers in service occupations report spending about as much time performing routine tasks as clerical workers. While the typical food service worker may find their primary job tasks repetitive, this does not necessarily mean that these tasks are programmable. Asking workers which of their tasks are routine may thus not prove helpful in identifying programmable tasks. Asking them which of their tasks are programmable may be equally unhelpful as it may be difficult to assess for someone who is not an expert in computer science.

One promising strategy is to use indirect questions that identify characteristics of codifiability while being easy to answer for the typical worker. To improve on existing measures of routine tasks, we take advantage of the richness of the Employment Surveys, in particular of the module focusing on working conditions. That module contains two questions—unaltered across survey waves—that may well capture the routine intensity of the tasks workers perform. The questions are:

Question 1: How often does it occur in your daily work that your tasks are prescribed precisely in every single step? (own translation)

Question 2: How often does it occur in your daily work that single tasks repeat themselves in every detail? (own translation)

These questions capture well Autor’s (2013) definition of what constitutes a routine task, i.e. ‘automating a task requires attaining a level of mastery beyond what is required for a worker to simply perform the task; it must be codified to the point where a relatively inflexible machine can perform the work semi-autonomously.’ (p. 187). These questions are also very similar in spirit to Deming (2017), namely (i) ‘how automated is the job?’ and (ii) ‘how important is repeating the same physical activities (e.g. key entry) or mental activities (e.g. checking entries in a ledger) over and over, without stopping, to performing this job?’ (p. 1614). Indeed, Rohrbach-Schmidt and Tiemann (2013), who provide an extensive analysis on the comparability of the Autor et al. (2003) task classification with the Employment Surveys, also use these two questions when evaluating the criterion validity of task items.

To identify routine tasks, we perform simple correlation analyses between these two questions and the task measures retrieved from the survey module on activities performed on the job. We find that the ‘routine manual’ tasks in Spitz-Oener (2006) correlate highly and statistically significantly with these two questions across all waves. The same holds for the task ‘measuring length/weight/temperature,’ which Spitz-Oener (2006) classifies as ‘routine cognitive.’ However, other ‘routine cognitive’ tasks such as ‘calculating’ and ‘writing texts’ do not correlate with these questions. Therefore, we classify those tasks as abstract. Furthermore, we find the task ‘cleaning and rubbish removal’ to significantly correlate with the survey questions on routineness. Therefore, we classify it as routine as well. Spitz-Oener (2006) ignores this task in her classification, while Dustmann et al. (2009) classify it as non-routine manual. Table 1 provides an overview of our task classification.

Table 1: Classification of tasks

Category	Tasks
Abstract tasks	calculating and correcting text/data; executing, interpreting, and advising on law/rules; planning, projecting, and designing; programming; researching, analysing, and evaluating
Interactive tasks	advertising, publishing, and public relations; coordinating and organising; negotiating and advising; teaching and training
Manual tasks	repairing, restoring, and renovating; securing; serving and accommodating guests
Routine tasks	cleaning and rubbish removal; equipping and operating machinery; manufacturing or producing; measuring length/weight/temperature

Notes: The classification comprises five survey waves of the Employment Surveys, from 1979 to 2006. Own translation.

4 Overview of trends in tasks and wages

4.1 Trends in tasks

This section provides some descriptive statistics on changes in task content from 1978 to 2006. Recall that we aggregate our individual-level task data to the occupation level to match them to our wage data. Figure 1 shows the task shares along the wage distribution at the beginning and end of our observation period. The figure shows a remarkably stable hierarchy

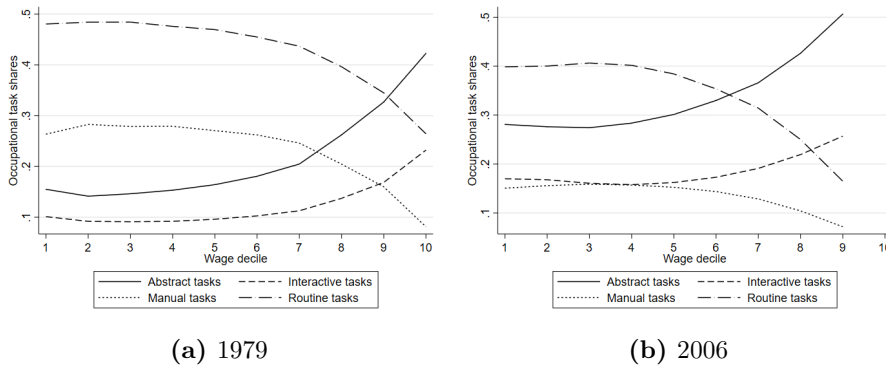
in the relative size of task shares over time. In the left panel (in 1979), workers perform around 45% routine and 30% manual tasks at the first decile. The shares of both routine and manual tasks are fairly stable along the lower parts of the distribution and only slowly decline until around the seventh wage decile. From there on, the decline is more pronounced and sharpest at the eighth decile. At the top of the wage distribution, workers perform around 25% routine and 10% manual tasks. The shares of abstract and interactive tasks increase with wage. At the bottom, workers perform around 15% abstract and 10% interactive tasks. At the top, workers perform around 40% abstract and over 20% interactive tasks.

In the right panel (in 2006), the overall shares of routine and manual tasks have decreased along the entire wage distribution compared to 1979, while the shares of abstract and interactive tasks have increased. This finding follows recent evidence for the US that all occupations have become less routine over time (Michaels et al., 2019; Atalay et al., 2020). The relative distribution of tasks across the wage distribution barely changes: Workers in low-wage occupations perform mostly routine and manual tasks, while workers in high-wage occupations perform mostly abstract and interactive tasks. Specifically, at the first decile, workers perform around 40% routine, almost 30% abstract and just below 20% manual and interactive tasks. In contrast, at the tenth decile, workers perform over 50% abstract, almost 30% interactive, less than 20% routine and less than 10% manual tasks in 2006.

These patterns are in line with Black and Spitz-Oener (2010), who—using the same data until 1999—show that the share of routine tasks has decreased relatively little for men. They also document an increase in analytical and interactive tasks across occupation groups. In contrast, using the

1991/92 wave only, Dustmann et al. (2009) find the share of routine tasks to be highest around the eighth wage decile.

Figure 1: Task shares across wage deciles in 1979 and 2006




Notes: This figure contrasts the relative shares of tasks performed along the wage distribution in 1979 with those in 2006.

Next, we take a closer look at changes in tasks over time at specific parts of the wage distribution. The top panel of Figure 2 shows the shares of tasks for workers at the first decile, the middle panel those for workers at the fifth decile, and the bottom panel the tenth decile. Generally, we confirm the trends from Figure 1: For low-wage workers, routine and manual tasks decrease from 45% to 40% and from 30% to below 20% respectively. Abstract and interactive tasks increase from around 15% and 10% to almost 30% and almost 20% respectively. Workers in the middle experience similar trends. Their shares of routine and manual tasks decrease by about five and 15 percentage points respectively, while their shares of abstract and interactive tasks increase by about 12 and eight percentage points.


Finally, at the top, abstract and interactive tasks increase from around 25% to above 40% and from almost 15% to above 20%. Conversely, routine and manual tasks decrease from almost 40% to under 30% and from above

20% to around 10%.

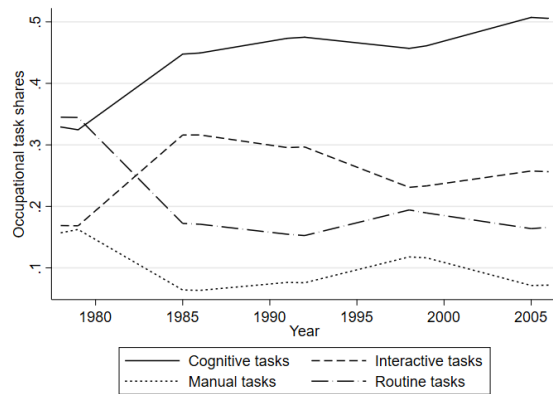
Figure 2: Task shares at different wage deciles over time



(a) first wage decile



(b) fifth wage decile



(c) ninth wage decile

Notes: This figure shows the shares of tasks performed over time at the first, fifth, and ninth wage decile.

4.2 Trends in wages

We now turn to the description of trends in wages. In a first step, we look at the overall wage distribution. Table 2 presents basic characteristics of our wage data across our four periods. The reported wages are log gross daily wages weighted by the number of days worked in a respective year. Wages are reported as 2006 wages in euros, adjusted for inflation using the Consumer Price Index.

The table shows that average real wages rose by about 6% between 1978 and 1986, by another 11% between 1986 and 1992, and then remained relatively stable over the next 20 years. The standard deviation of log wages rose by five log points between 1978 and 1986, then remained stable until the early 1990s, then surged over the next 15 years, rising by 11 log points.

Table 2: Summary statistics

Year	Obs.	Mean	St. Dev.
1978/79	349,806	4.462	0.293
1985/86	356,297	4.526	0.343
1991/92	384,994	4.636	0.349
1998/99	345,914	4.639	0.395
2005/06	307,821	4.652	0.460

Notes: The sample comprises full-time employed men, aged 20-60, who live in West Germany and are German nationals. The table shows the mean and standard deviation of log gross daily wages weighted by the number of days worked in a respective year. Wages are reported as 2006 wages in euros, adjusted for inflation using the Consumer Price Index.


To investigate distributional characteristics of changes in the wage distribution, Figure 3 plots the difference in log wages between two periods over the distribution. The vertical line shows the change in the median. These descriptive patterns give a first indication for why we need distributional analyses to understand wage dynamics. Focusing on the median or mean would overlook the large changes that occur at the top and bottom ends.

Specifically, the figure shows that from 1978 to 1986, wages were increasing overall, but more steeply so at the upper deciles. Conversely, between 1986 and 1992, wage increases were uniform over the wage distribution.⁷ The 1990s paint a different picture. While the median wage remained almost unchanged, wages below the fourth decile decreased, while wages above the sixth decile increased. From 1999 to 2006, the pattern of the 1990s continued: Wages continued decreasing below the median, while they continued increasing above the median. In sum, we observe a roughly symmetric widening of the upper and lower tails of the wage distribution for this time period.

These descriptive patterns indicate that, on average since the 1990s, low-wage workers suffered relatively higher wage losses than middle-wage workers. High-wage workers gained relatively more than those in the middle. These patterns are different from the ones observed in Anglo-Saxon countries, but are broadly in line with studies for Germany that use register data (Fitzenberger, 1999; Kohn, 2006; Gernandt and Pfeiffer, 2007; Dustmann et al., 2009, 2014; Riphahn and Schnitzlein, 2016; Biewen et al., 2017; Antonczyk et al., 2018; Baumgarten et al., 2020).

⁷ The 1986 to 1992 period stands out in this comparison, as it is the only one with uniform changes along the entire distribution. Section 6.3.1 discusses possible explanations for this development.

Figure 3: Changes in log wages over time



Notes: This figure plots the difference in log wages between two periods over the distribution. The vertical line shows the change in the median wage.

Figure 4 takes a closer look at the standard deviation of wages at the first, fifth, and ninth decile over time. While the standard deviation has increased across all deciles, the increase was more pronounced at the tails than in the middle. This descriptive pattern may indicate again that changes in wage inequality may be stronger at the top and the bottom of the distribution, while the middle remained more stable. Recent contributions by Eichhorst and Buhlmann (2015) and Möller (2015) reach a similar conclusion.

Figure 4: Standard deviation of wages at the first, fifth, and ninth decile



Notes: This figure shows the evolution over time of the standard deviation of wages at the first, fifth, and ninth wage decile.

We will now look more closely into the evolution of wages across occupations. Figure 5 shows the change in the real median log wage between 1978 and 1992, and 1992 and 2006, sorted by the highest paying job in 1978. A first visible trend is that while all occupations experienced real wage gains between 1978 and 1992, many incurred losses in real terms over the following period. Among those, unskilled workers experienced by far the largest losses with a magnitude of almost 50%. These developments in the second period are related to the unification shock and the 1992/93 recession, which will be discussed more in detail in Section 6.2. In particular, the exceptionally large losses experienced by unskilled workers may partly be driven by a supply shock as outlined in Dustmann et al. (2009). The authors posit that the breakdown of the communist regimes in Eastern Europe and the reunification of East and West Germany likely led to a relative increase in

the share of low-skilled workers. If we compare trends in employment shares for this occupational category in our wage data, we can confirm this supply shock explanation, as the employment share doubles from 0.94% to 1.9% during the second period.


What Figure 5 also shows is the large heterogeneity in terms of the evolution of wages across occupations: Many high-wage occupations experienced significant wage growth over both periods. Examples are chemists, physicians, and engineers. At the bottom, most occupations suffered wage losses during the second period, between 1992 and 2006. The most extreme case is that of hairdressers, who in 2006 earned less in real terms than in 1978. In the middle of the distribution, the picture is less clear-cut. While the majority of occupations still realised small wage gains in the second period, others suffered significant losses. These observations confirm our earlier assessment derived from the aggregate figures, i.e. that most changes in wage inequality are clustered at the top and bottom.

Next, we investigate the ranked order of wages, i.e. the question of whether wage changes within occupations would result in a change of the wage hierarchy of occupations. Table 3 shows the median log wage of the five highest, intermediate, and lowest-paying jobs at five different points in time. The tails of the distribution are remarkably stable in their ranked order, whereas the middle is somewhat more volatile. However, most middle-wage occupations appear at least twice in the table. The wage hierarchy of occupations thus seems to have changed little based on this descriptive evidence.

A further interesting feature of Table 3 is that it shows inequality in wage levels, not only in relative changes. In 1978, the lowest paying job in our data is that of a hairdresser, who receives a median daily wage of about

50 euros. The highest paying job is that of a chemist with a median daily wage of 122 euros. In 2006, the hairdresser's median wage is about 48 euros, whereas the highest-paying job, a physician, earns 172 euros.

Figure 5: Change in median real log daily wage by occupation



Notes: This figure shows the change in median real log daily wages for single occupations. The light grey bar shows the change between 1978/79 and 1991/92, the dark grey bar between 1991/92 and 2005/06.

Table 3: Wage hierarchy of occupations over time (log wages)

Lowest-paying occupations		Middle-paying occupations		Highest-paying occupations	
1978/79					
Hairdresser	3.90	Miner	4.40	Social scientist	4.79
Unskilled constr. worker	4.20	Metal producer	4.40	Engineer	4.85
Innkeeper	4.20	Truck driver, conductor	4.40	Entrepreneur	4.87
Tailor	4.26	Nurse	4.40	Physician	4.93
Weaver	4.28	Electrician	4.42	Chemist	4.94
1985/86					
Hairdresser	3.83	Metal producer	4.43	Social scientist	4.93
Innkeeper	4.11	Truck driver, conductor	4.43	Engineer	4.97
Baker	4.23	Nurse	4.43	Entrepreneur	4.99
Cook	4.23	Blacksmith	4.44	Physician	5.04
Unskilled constr. worker	4.25	Plasterer	4.44	Chemist	5.06
1991/92					
Hairdresser	3.91	Truck driver, conductor	4.53	Pub. admin., politician	5.01
Innkeeper	4.21	Blacksmith	4.53	Engineer	5.05
Baker	4.32	Bricklayer	4.53	Entrepreneur	5.05
Cook	4.35	Janitor	4.53	Physician	5.10
Tailor	4.38	Miner	4.55	Chemist	5.15
1998/99					
Hairdresser	3.90	Blacksmith	4.51	Pub. admin., politician	5.05
Innkeeper	4.17	Plasterer	4.51	Engineer	5.05
Baker	4.30	Miner	4.52	Entrepreneur	5.10
Cook	4.30	Metal presser & molder	4.52	Chemist	5.15
Tailor	4.33	Metal producer	4.54	Physician	5.17
2005/06					
Hairdresser	3.87	Paper & pulp processing	4.51	Pub. admin., politician	5.11
Unskilled worker	4.04	Mechanic	4.51	Engineer	5.14
Innkeeper	4.13	Miner	4.52	Entrepreneur	5.17
Baker	4.21	Product tester	4.54	Chemist	5.21
Cook	4.25	Assembler	4.54	Physician	5.22

Notes: This table shows the five highest-, middle-, and lowest-paying occupations in our sample over time with their respective real log daily wages.

5 Decomposing wage distributions

5.1 Econometric model

To quantify the contribution of occupational tasks to changes in the wage structure, we apply RIF-regression based decompositions, a method introduced by Firpo et al. (2009) (hereafter: FFL) and discussed in Fortin et al. (2011) and Firpo et al. (2018). Decomposition methods seek to explain the wage gap between two groups by decomposing it into a component attributable to differences in the observed characteristics of the groups (the composition effect) and a component attributable to differences in the returns to these characteristics (the wage structure effect). In our application, these two groups represent two periods.

Let $\nu(F_Y)$ denote a distributional statistic for the cumulative wage distribution F_Y . Let $F_{Y_{0|T=0}}$ denote the cumulative wage distribution observed in period 0 and $F_{Y_{1|T=1}}$ the distribution observed in period 1. In contrast, let $F_{Y_{0|T=1}}$ denote the counterfactual distribution that would have prevailed if workers in period 1 had been paid under the wage structure of period 0.

We can decompose the overall change in $\nu(F_Y)$ between two periods as:

$$\begin{aligned}\Delta_0^\nu &= \nu(F_{y_1|t_1}) - \nu(F_{y_0|t_0}) \\ &= [\nu(F_{y_1|t_1}) - \nu(F_{y_0|t_1})] + [\nu(F_{y_0|t_1}) - \nu(F_{y_0|t_0})] \\ &= \Delta_W^\nu + \Delta_C^\nu\end{aligned}\tag{4}$$

The first difference in this equation, Δ_W^ν , is the wage structure effect and the second difference, Δ_C^ν , the composition effect.

The challenge in retrieving the wage structure and composition effects

lies in estimating the counterfactual wage distribution. FFL apply the reweighting procedure by DiNardo et al. (1996), which replaces the marginal distribution of covariates X for workers in period 0 with the marginal distribution of X for workers in period 1 using a reweighting factor, $\Psi(X)$. The reweighting factor is defined as follows:

$$\Psi(X) = \frac{Pr(X | D_1 = 1)}{Pr(X | D_1 = 0)} = \frac{Pr(D_1 = 1 | X)/Pr(D_1 = 1)}{Pr(D_1 = 0 | X)/Pr(D_1 = 0)} \quad (5)$$

It is computed by estimating a probability model for $Pr(D_1 = 1 | X)$ and using the predicted probabilities to compute a value $\hat{\Psi}(X)$ for each observation. Distributional statistics can then be computed with $\hat{\Psi}(X)$ as weight.

The main advantage of this reweighting procedure is its simplicity. However, for distributional statistics besides the mean, it cannot be extended from aggregate to detailed decompositions. This means that it can only quantify the wage structure and the composition effects, but not further decompose the contribution of single covariates.

To go beyond the mean, FFL suggest using the recentered influence function (RIF) regression. An influence function quantifies how a target statistic changes in response to small changes in the data. For each value y , the influence function $IF(y; \nu; F_Y)$ provides an approximation of how the functional $\nu(F_Y)$ changes if a small probability mass is added at point y .

Influence functions are centred around zero. To centre an influence function around the statistic of interest, we can simply add the statistic to the influence function. The *recentered* influence function is then defined as:

$$RIF(y; \nu, F_Y) = \nu(F_Y) + IF(y; \nu, F_Y) \quad (6)$$

We can model the conditional expectation of $RIF(y; \nu, F_Y)$ using the following linear regression model:

$$E(RIF(y; \nu, F_Y | X)) = X\gamma \quad (7)$$

The coefficient γ thus provides an approximation of how $\nu(F_Y)$ reacts to changes in X .

Using the coefficients γ from the RIF regression in two groups, FFL show that one can run Oaxaca-Blinder (OB) decompositions on the reweighted data to retrieve the detailed decompositions. The equation then modifies to:

$$\hat{\Delta}^\nu = \hat{\Delta}_W^\nu + \hat{\Delta}_C^\nu = (\bar{X}_0 - \bar{X}_1)\hat{\gamma}_0 + \bar{X}_1(\hat{\gamma}_0 - \hat{\gamma}_1) \quad (8)$$

The FFL method can thus be considered an extension of the OB decomposition, as it permits decompositions of differences in functionals of distributions. The composition effect can be written as the sum of the true composition effect and a specification error component:

$$\Delta_{C,rew}^\nu = (\bar{X}_{01} - \bar{X}_0)\hat{\gamma}_1^\nu + \bar{X}_{01}(\hat{\gamma}_{01}^\nu - \hat{\gamma}_0^\nu) = \Delta_{C,t}^\nu + \Delta_{C,e}^\nu \quad (9)$$

The specification error occurs because the unweighted decomposition only provides an approximation of the composition effect. If the approximation is accurate, the specification error should be small.

Similarly, the wage structure effect can be written as:

$$\Delta_{W,rew}^\nu = \bar{X}_1(\hat{\gamma}_1^\nu - \hat{\gamma}_{01}^\nu) + (\bar{X}_1 - \bar{X}_{01})\hat{\gamma}_{01}^\nu = \Delta_{W,t}^\nu + \Delta_{W,e}^\nu \quad (10)$$

The reweighting error occurs because the reweighted mean \bar{X}_{01} is not ex-

actly equal to \bar{X}_1 . In large samples, i.e. $\bar{X}_{01} \rightarrow \bar{X}_1$, the reweighting error approaches zero.

5.2 Identifying assumptions

As is the case for many other methods, RIF-regression based decompositions rely on some assumptions to identify the composition and wage structure effects. The key assumptions are ignorability and overlapping support. First, we discuss the ignorability assumption.

Ignorability assumption: Let (T, X, ε) have a joint distribution. For all x in X : ε is independent of T given $X = x$.

We thus assume that the contribution of unobservables in the wage determination is the same across groups 1 and 0, which allows us to condition on a vector of observed components. Put differently, the error terms in the wage equation are ignorable: Conditional on X the distributions of the error terms are the same across groups. In our application, it implies that the contribution of unobservables remains stable across two periods.

For example, the residual in our wage equation contains the contribution to wages of workers' sorting on unobservable characteristics. If, between period 0 and 1, sorting dynamics change and workers with high unobserved ability increasingly sort into occupations at the upper tail, then the wage structure effects will underestimate the contribution of tasks to changes in wage inequality. The larger the time horizons under investigation, the larger will be the potential bias. Because we run decompositions in eight-year intervals, we are not particularly worried about this potential bias. Moreover, the bias would result in an underestimation of the effects.

Second, we discuss the overlapping support assumption.

Overlapping support assumption: For all x in X , $p(x) = Pr[T = 1|X = x] < 1$. Furthermore, $Pr[T = 1] > 0$.

This implies that there is no value of x in X such that it is only observed among individuals in Group 1. No one value of a characteristic can perfectly predict belonging to one group. This assumption may be quite restrictive in gender wage gap decompositions if some occupations are only held by men or by women. However, it is less restrictive when looking at changes in the wage distribution over time. Again, it mostly depends on the length of the time horizon considered. If group 1 consists of workers in 1970 and group 0 of workers in 2010, the difference in wages should consider that many occupations in the year 2010 did not exist in 1970. However, in our application this does not appear to be an issue. If these two assumptions are satisfied, then we are able to identify the parameters of the counterfactual distribution.

5.3 Estimation

To decompose changes in the distribution of wages into the contribution of tasks and other factors, we use the RIF regression-based decomposition. In the first step, we estimate the reweighting function by DiNardo et al. (1996). Using the pooled data for the periods 0 and 1, we run a probit regression to estimate the predicted probability of belonging to each group conditional on covariates X for each observation in the sample. Then, we use the predicted probabilities of belonging to group 0 and 1 to calculate the reweighting factor, $\Psi(X)$. Our probit specification considers the covariates

of the decomposition analysis as well as additional interaction terms. As individual characteristics, we include our four occupational task measures (abstract, interactive, manual and routine task shares), education, and years of potential experience. We also include occupational, region, and sector dummies.

In the second step, we calculate the RIF by wage quantile. The RIF of quantile Q_p is simply:

$$RIF(y; \nu, F_Y) = \nu(F_Y) + IF(y; \nu, F_Y) = \nu(F_Y) + \frac{p - I(y \leq Q_p)}{f_Y(Q_p)} \quad (11)$$

In practice, we first compute the sample quantile \hat{Q}_p and then use the kernel density estimation with a Epanechnikov kernel and a bandwidth of 0.65 to retrieve $\hat{f}(\hat{Q}_p)$, the density of Y at point \hat{Q}_p . The $RIF(Y_i; Q_p, F_Y)$ is computed for each observation by plugging these estimates in the above formula. Finally, we regress $RIF(Y_i, Q_p, F_Y)$ on X to get an estimate of γ .

Using the coefficients from the RIF regression in two groups, we perform two OB decompositions at each quantile by replacing the outcome variable (log daily wage) with the RIF. Performing OB decompositions on the RIFs allows further breaking down the wage structure and decomposition effects into the contribution of each explanatory variable of the model.

First, we perform an OB decomposition using the $T = 0$ sample and the counterfactual sample ($T = 0$ sample reweighted to be as in $T = 1$) to obtain the composition effect. The total unexplained effect in this decomposition corresponds to the specification error and allows assessing the importance of departures from the linearity assumption.

Second, we perform an OB decomposition using the $T = 1$ sample and

the counterfactual sample, using the counterfactual wage structure as reference, and obtain the wage structure effect. The total explained effect in this decomposition corresponds to the reweighting error, which should go to zero in large samples. It thus provides a straightforward way of assessing the quality of the reweighting.

Importantly, the standard regression assumption that the conditional expectation of the error is zero for all covariates applies to the OB decomposition as well. When this conditional independence assumption does not hold, the conventional solution is to use instrumental variable methods. For example, if we suspect occupational tasks to be correlated with the error term in the wage equation, we need a valid instrument for these tasks in order to estimate the model consistently.

In cases where the zero conditional mean assumption fails, the aggregate decomposition may remain valid provided that the ignorability assumption holds. This would be the case for example when unobserved ability is correlated with occupational tasks, but the correlation is the same in groups 0 and 1. While this will not allow us to identify the contribution of ability vs. tasks, we know that there are no systematic ability differences between 0 and 1, once we control for tasks.

Following the task assignment model, the larger potential worry is the endogenous relationship between tasks and skills. This is why we add years of education as a control variable in our regressions. Task assignments that are driven by years of education should thus be accounted for by our control variables.

Finding a valid instrument is difficult and lies outside the scope of this paper. We do not claim causality of our estimates in the returns to tasks, but are able to uncover statistical associations. While future research may need

to focus on establishing causality, our contribution is a great advancement compared to a simple visual inspection of patterns.

Finally, our decomposition analysis, similar to alternative approaches used in the literature (Fairlie, 2005; Machado and Mata, 2005), is not able to account for general equilibrium effects. This is because RIF regressions assume the invariance of the conditional distribution to retrieve a valid counterfactual wage distribution.

This point is related to the occupational sorting discussed previously. For example, if technological change affects the demand for tasks, worker sorting in response to such a shock will tend to limit changes in wages, but such sorting is ignored in our decomposition analysis.

Models that explicitly account for worker sorting (Atalay et al., 2019) may be superior to the decomposition procedure in this regard. Such models will, however, have to impose strong parametric assumptions on workers' unobserved ability to work in each potential occupation, i.e. sources of sorting beyond what is measurable with task content. As such, these approaches should be viewed as complementary rather than one being superior to the other.

6 Results

6.1 Aggregate decomposition results

Table 4 reports the aggregate decomposition results of the RIF regression-based decompositions, i.e. the composition and wage structure effects. Overall, the results are consistent with the descriptive analyses in Section 4, which indicates that our model provides a reasonable fit. To simplify the exposition, we focus on standard measures for wage inequality in the discussion of

results. We report changes in the 90-10 log wage differential as a measure for overall inequality, and changes in the 90-50 and the 50-10 log wage differentials as measures of upper-tail and lower-tail wage inequality respectively. A widening of these wage gaps implies an increase in wage inequality.

Table 4 reveals several important results. First, overall wage inequality has increased over time, with a particularly pronounced increase between the late 1990s and the mid-2000s.⁸ Second, in the beginning of our observation period, wage inequality was limited to the upper part of the distribution. Lower-tail inequality started increasing slowly in the mid-1980s and then surged in the 1990s. These results are consistent with previous literature that uses register data (Fitzenberger, 1999; Kohn, 2006; Dustmann et al., 2009, 2014; Riphahn and Schnitzlein, 2016; Biewen et al., 2017; Antonczyk et al., 2018; Baumgarten et al., 2020). However, we find smaller wage gaps than those reported in previous contributions. While this may be puzzling at first sight, note that most of those contributions look at larger time horizons than us.

Third, over the entire observation period, wage structure effects are mostly driving changes in wage inequality, whereas composition effects play a minor role. This finding suggests that changes in the distribution of the underlying characteristics of the population were hardly affecting inequality between any two periods, whereas changes in the returns to these underlying characteristics did. In contrast, Lemieux (2006) shows that a large part of the rise in the US residual wage inequality between 1973 and 2003 can be attributed to changes in the workforce composition. Furthermore, some contributions argue that composition effects play a role for German wage inequality as well (Dustmann et al., 2009; Antonczyk et al., 2010; Biewen

⁸ Note that the results reported in the decomposition analyses only include uncensored wages, which likely underreports upper-tail wage inequality.

et al., 2017). A reason why our findings differ may be that these studies use either other data or look at other (larger) time horizons. Notably, Card et al. (2013) also find that compositional effects do not matter greatly for Germany.

Fourth, the specification errors, the difference between the total composition effect obtained by reweighting and the RIF-regression, show that the RIF-regressions capture the overall trend in composition effects accurately. The size of the errors are very small. Exceptions are the lower tails in the first and the last period. However, even these larger errors are still similar in magnitude to the ones reported in Firpo et al. (2011). Overall, the reweighting thus appears to produce reliable results.

A clear limitation of the aggregate decomposition is that it provides neither information about the contribution of single covariates nor about the underlying mechanism. It only provides a broad picture of trends in wage inequality and points the researcher toward the economically significant effects. As suggested by FFL, we perform the second step of the RIF regression-based decomposition to quantify the contribution of single covariates.

Table 4: Aggregate decomposition results

Inequality measure	90-10	90-50	50-10
A. 1978/79 - 1985/86			
Total change	0.0477*** (0.0003)	0.0474*** (0.0003)	0.0003 (0.0003)
Composition effect	-0.0009*** (0.0000)	0.0008*** (0.0000)	-0.0017*** (0.0000)
Wage structure effect	0.0116*** (0.0005)	0.0403*** (0.0004)	-0.0287*** (0.0003)
Specification error	0.0344*** (0.0004)	0.0084*** (0.0003)	0.0260*** (0.0003)
B. 1985/86 - 1991/92			
Total change	-0.0000 (0.0003)	-0.0172*** (0.0002)	0.0172*** (0.0002)
Composition effect	-0.0021*** (0.0000)	0.0001*** (0.0000)	-0.0022*** (0.0000)
Wage structure effect	-0.0077*** (0.0002)	-0.0205*** (0.0002)	0.0128*** (0.0002)
Specification error	-0.0022*** (0.0000)	-0.0003*** (0.0000)	-0.0019*** (0.0000)
C. 1991/92 - 1998/99			
Total change	0.1032*** (0.0001)	0.0609*** (0.0001)	0.0423*** (0.0001)
Composition effect	0.0062*** (0.0001)	-0.0024*** (0.0000)	0.0086*** (0.0000)
Wage structure effect	0.1101*** (0.0007)	0.0926*** (0.0006)	0.0175*** (0.0007)
Specification error	-0.0057*** (0.0007)	-0.0055*** (0.0005)	-0.0001 (0.0007)
D. 1998/99 - 2005/06			
Total change	0.1334*** (0.0003)	0.0524*** (0.0001)	0.0811*** (0.0002)
Composition effect	0.0102*** (0.0001)	-0.0032*** (0.0000)	0.0134*** (0.0001)
Wage structure effect	0.0845*** (0.0003)	0.0468*** (0.0002)	0.0377*** (0.0003)
Specification error	0.0303*** (0.0002)	0.0040*** (0.0001)	0.0262*** (0.0002)

Notes: Bootstrapped standard errors in parentheses. 100 replications of the full procedure.
*** p<0.01, ** p<0.05, * p<0.1

6.2 Detailed decomposition results

Tables 5 to 8 present the detailed decomposition results, with Panel A showing the composition effects and Panel B the wage structure effects. To guide the discussion, we will first provide an overview of the findings, then focus on some general features of our RIF regressions, and lastly look at the specific results. When discussing detailed results, we will bundle the years from 1978 to 1992 in one subsection and those from 1992 to 2006 in another subsection due to similar trends.

Our main results can be summarised as follows: First, from the late 1970s until the early 1990s, increases in wage inequality were mostly limited to the upper part of the distribution. Hereby, changing returns to education and experience contribute to the rising inequality.

Second, starting in the 1990s, we find an important role for occupational tasks. Specifically, abstract tasks drive upper-tail wage inequality, while routine tasks are associated with the increasing lower-tail wage inequality. In terms of magnitude, abstract tasks explain about half of the increase at the top and routine tasks half of the increase at the bottom. These increases in inequality result from wage gains at the top and wage losses at the bottom, while wages in the middle remain relatively stable. These results can be viewed as partly confirming the routinisation hypothesis by Autor et al. (2003) in that new technology may be complementing abstract tasks and substituting routine tasks. It also follows recent evidence by Atalay et al. (2019).

Third, in the 2000s, we find that interactive tasks explain some of the increase in lower-tail inequality, accounting for about one fifth of the widen-

ing of the gap. This finding speaks to the relatively recent literature on the importance of social skills (Borghans et al., 2014; Weinberger, 2014; Deming, 2017).⁹

Next, before diving into the discussion of results, we review some important features of the estimated RIF-coefficients. A typical feature is that the covariates have non-monotonic effects. They may be increasing inequality in some parts of the distribution, while they decrease it in other parts. For example, Table 7 shows that firm characteristics are inequality-increasing at the upper part of the wage distribution, while they are inequality-decreasing at the lower part. One explanation may be that high-wage firms are early technology adopters. This feature underlines once again why an analysis of the median or mean will miss crucial information.

Furthermore, as in standard OB decompositions, the effect size of single coefficients can be larger than the overall composition or wage structure effects, as the effects of single covariates may offset each other. For example, Table 7 shows that the routine tasks effect in the 50-10 gap is more than the full size of the wage structure effect. Again, this underlines the importance of the disaggregation of effects.

Additionally, as in standard OB decompositions, the change in intercepts captures the part of the wage structure effect that cannot be accounted for by covariates. Finally, because of our large sample size, almost all results are highly statistically significant, although not all of them are economically

⁹ As sensitivity checks not reported in the paper, we have rerun our decomposition analyses (i) with time-invariant occupational tasks and (ii) without including any task measures. Regarding (i), we use the content of the tasks from the first period for all analyses. We find largely negligible task effects. This result reveals the importance of accounting for *within*-occupation changes in the task composition over time, while it also confirms earlier studies for Germany that did not find a relationship between tasks and changes in wage inequality, when keeping tasks constant (Antonczyk et al., 2009). Furthermore, both (i) and (ii) result in larger effects of education and firm-specific controls, which include occupation dummies.

relevant. We refer to the aggregate decomposition results as guidance for the economic significance.

6.2.1 Results for the period from 1978 to 1992

Table 5 presents the decomposition results for the period 1978/79 to 1985/86. The increase in wage inequality is limited to the upper part of the distribution during this time. The largest contributors are our education variables, which account for about half of the wage structure effect at the 90-50 wage gap. Furthermore, they also strongly contribute to a widening of the 50-10 wage gap. However, the overall wage structure effect at the 50-10 gap is actually negative, i.e. inequality-decreasing. This is because other factors, caught in the intercept and not accounted for in our model, are offsetting the inequality-increasing education effect. Additionally, the experience coefficients are large and inequality-increasing at the lower part of the distribution as well. In comparison to the magnitude of these control variables, our task measures do not play a role. In sum, changes in the wage structure in the 1978-1986 period can be explained mainly by changes in the returns to education at the 90-50 gap and factors outside of our model at the 50-10 gap.

Next, Table 6 presents the decomposition results for the period 1985/86 to 1991/92. Overall, wage inequality remains largely stable during this period. Indeed, the overall measure of wage inequality, the total change of the 90-10 wage gap, is close to zero and not significant (view Table 4). This result is mostly due to the 90-50 wage gap closing by a similar magnitude as the 50-10 wage gap is widening. The contribution of most covariates is close to zero. In particular, again, occupational tasks hardly contribute to changes in the wage structure.

One exception are our education variables, which are strongly inequality-increasing at the upper part. This is a continuation of the trend already observed in the previous period. These results at the upper tail for the full 1978-1992 period thus speak to the SBTC literature, favouring skilled workers at the upper parts of the wage distribution (Bound and Johnson, 1992; Katz and Murphy, 1992; Levy and Murnane, 1992; Juhn et al., 1993; Berman et al., 1994). This conclusion has also been underlined by Fitzenberger (1999) and Antonczyk et al. (2009). However, for the late 1980s, the inequality-increasing effect of education is offset by factors outside of our model, accounted for by the large inequality-reducing effects of the intercept. Therefore, overall, the 90-50 gap actually decreases in the late 1980s.

At the lower part of the distribution, the 50-10 wage gap is increasing in the late 1980s. With the upper wage gap closing and the lower wage gap widening, the 1985-1992 period stands out in our analysis. However, the variables included in our model can neither explain the closing of the 90-50 gap nor the widening of the 50-10 gap.

This period coincides with the German reunification in 1990, which has fundamentally affected the West German labour market. Dustmann et al. (2009) argue that a supply shock, specifically a massive migration of ethnic Germans from the former Soviet Union triggered by the German reunification, may explain the increase in lower-tail inequality. That upper-tail inequality did not increase further during that time may be related to the policy decision of imposing West German wage scales on the less productive East, which inevitably should have led to repercussions on West German wages as well (Burda, 2000).

Table 5: Detailed decomposition results, 1978/79 - 1985/86

Inequality measure	90-10	90-50	50-10
A. Detailed composition effects			
Abstract tasks	0.0000*** (0.0000)	0.0000*** (0.0000)	-0.0000*** (0.0000)
Interactive tasks	0.0004*** (0.0000)	0.0002*** (0.0000)	0.0002*** (0.0000)
Manual tasks	-0.0003*** (0.0000)	-0.0003*** (0.0000)	-0.0000*** (0.0000)
Routine tasks	-0.0000*** (0.0000)	0.0000*** (0.0000)	-0.0000*** (0.0000)
Education	-0.0017*** (0.0000)	-0.0012*** (0.0000)	-0.0005*** (0.0000)
Experience	0.0027*** (0.0000)	0.0032*** (0.0000)	-0.0005*** (0.0000)
Firm controls	-0.0020*** (0.0000)	-0.0011*** (0.0000)	-0.0009*** (0.0000)
Composition effect	-0.0009*** (0.0000)	0.0008*** (0.0000)	-0.0017*** (0.0000)
B. Detailed wage structure effects			
Abstract tasks	0.0097*** (0.0002)	0.0065*** (0.0001)	0.0032*** (0.0001)
Interactive tasks	-0.0093*** (0.0002)	-0.0021*** (0.0001)	-0.0072*** (0.0001)
Manual tasks	-0.0025*** (0.0001)	-0.0042*** (0.0000)	0.0018*** (0.0001)
Routine tasks	-0.0006*** (0.0002)	0.0054*** (0.0002)	-0.0060*** (0.0002)
Education	0.0560*** (0.0013)	0.0211*** (0.0009)	0.0349*** (0.0011)
Experience	0.0264*** (0.0003)	0.0064*** (0.0003)	0.0200*** (0.0003)
Firm controls	-0.0075*** (0.0002)	0.0002 (0.0002)	-0.0078*** (0.0002)
Constant	-0.0607*** (0.0013)	0.0070*** (0.0010)	-0.0676*** (0.0011)
Wage structure effect	0.0116*** (0.0005)	0.0403*** (0.0004)	-0.0287*** (0.0003)

Table 6: Detailed decomposition results, 1985/86 - 1991/92

Inequality measure	90-10	90-50	50-10
A. Detailed composition effects			
Abstract tasks	0.0008*** (0.0000)	0.0005*** (0.0000)	0.0003*** (0.0000)
Interactive tasks	0.0010*** (0.0000)	0.0006*** (0.0000)	0.0005*** (0.0000)
Manual tasks	0.0000 (0.0000)	-0.0000 (0.0000)	0.0000 (0.0000)
Routine tasks	-0.0001*** (0.0000)	-0.0001*** (0.0000)	0.0000*** (0.0000)
Education	-0.0011*** (0.0000)	-0.0008*** (0.0000)	-0.0003*** (0.0000)
Experience	-0.0051*** (0.0000)	-0.0009*** (0.0000)	-0.0043*** (0.0000)
Firm controls	0.0025*** (0.0000)	0.0009*** (0.0000)	0.0015*** (0.0000)
Composition effect	-0.0021*** (0.0000)	0.0001*** (0.0000)	-0.0022*** (0.0000)
B. Detailed wage structure effects			
Abstract tasks	-0.0013*** (0.0000)	-0.0007*** (0.0000)	-0.0006*** (0.0000)
Interactive tasks	-0.0002*** (0.0000)	0.0000 (0.0000)	-0.0002*** (0.0000)
Manual tasks	0.0001*** (0.0000)	-0.0007*** (0.0000)	0.0008*** (0.0000)
Routine tasks	-0.0003*** (0.0000)	0.0001*** (0.0000)	-0.0004*** (0.0000)
Education	0.0236*** (0.0012)	0.0225*** (0.0010)	0.0010 (0.0009)
Experience	0.0052*** (0.0003)	0.0018*** (0.0002)	0.0034*** (0.0003)
Firm controls	-0.0086*** (0.0003)	0.0051*** (0.0002)	-0.0137*** (0.0002)
Constant	-0.0262*** (0.0013)	-0.0486*** (0.0010)	0.0224*** (0.0009)
Wage structure effect	-0.0077*** (0.0002)	-0.0205*** (0.0002)	0.0128*** (0.0002)

6.2.2 Results for the period from 1992 to 2006

Table 7 presents the decomposition results for the period 1991/92 to 1998/99, when both upper-tail and lower-tail inequality increased significantly. Panel A shows the detailed composition effects. While the composition effects increase strongly compared to earlier periods, they are still much smaller than the wage structure effects and thus generally negligible. Panel B presents the detailed wage structure effects. Here, we find a pronounced role for abstract and routine tasks, while the coefficients for interactive and manual tasks remain small. Importantly, as in earlier periods, we continue to find a substantial role for changes in the returns to education and experience. This finding indicates that our task categories do not merely represent different degrees of educational attainment.

In particular, we find changes in the returns to abstract tasks to drive upper-tail wage inequality. They explain about one third of the wage structure effect at the 90-50 gap. In terms of the total effect ('total change'), they explain about half of the inequality increase. In our classification, abstract tasks comprise i.a. calculating, writing and planning and designing. These activities may be complemented with new technologies, which would raise productivity in these tasks. Therefore, it is plausible that workers in abstract-task intensive occupations could realise wage gains compared to workers in occupations with lower abstract task shares. This finding is (partly) in line with previous work by Spitz-Oener (2006), Dustmann et al. (2009), and Böhm et al. (2016), who argue that routine-biased technological change (RBTC) is an important driver for West German wage inequality, particularly at the upper part of the distribution.

At the lower part of the distribution, we find sizable effects for changes in the returns to routine tasks. In terms of magnitude, they explain just about the full wage structure effect at the 50-10 wage gap, and almost half of the total effect. Routine tasks comprise e.g. equipping and operating machinery and producing goods. It is plausible that these activities can—at least partly—be automatised and thus substituted with new technologies. This result is thus also consistent with Autor et al.’s (2003) hypothesis that computer technology decreases the demand for jobs requiring routine tasks.

However, the crucial difference is that we find routine-task intensity to be highest at the lowest percentiles and not in the middle of the distribution. Our results are thus in line with the assumptions of a RBTC, but do not explain a polarisation of wages but rather a widening gap both at the top and the bottom of the wage distribution. They also speak to a recent study by Atalay et al. (2019), who show that within-occupation changes in task composition significantly contribute to a rise in income inequality in the US.

Table 8 presents the decomposition results for our final period, 1998/99 to 2005/6. In contrast to the 1990s, we find abstract and routine tasks not to matter greatly anymore. Conversely, changes in the returns to interactive tasks explain about two fifths of the wage structure effect at the 50-10 gap. This finding implies that interactive tasks, such as advising and negotiating, became increasingly important during the 2000s. Because such tasks are likely to require face-to-face meetings and personal interactions, they may be less affected by technological change. This could explain why they have not played an important role in the 1990s.

That interactive tasks start to matter in the 2000s could be a result of their ‘non-offshorability.’ Occupations with a high interactive task-intensity should be more protected from offshoring, whereby firms carry out specific

subcomponents (or tasks) of their production processes abroad. Some studies suggest that the offshoring trends in the late 1990s and early 2000s may have increased the demand for interactive tasks, increasing wages for workers with large shares of interactive tasks (Firpo et al., 2011; Autor and Dorn, 2013; Goos et al., 2014). However, recent studies for Germany show that the effects of exporting on the wage structure are rather small (Baumgarten et al., 2013, 2020). Offshoring trends are thus likely to partly, but not fully explain the patterns we find in the decomposition.

Our finding also adds to the relatively few studies that focus on the rising importance of social skills and their complementarity with cognitive skills (Borghans et al., 2014; Weinberger, 2014; Deming, 2017). For example, Deming (2017) shows that the labour market return to social skills was much greater in the 2000s than in the mid-1980s and 1990s and that employment and wage growth were particularly strong for jobs requiring both social and cognitive skills. Similarly, Weinberger (2014) finds growing complementarity over time between cognitive skills and social skills, while Borghans et al. (2014) point to the increasing importance of ‘people skills’ in the workplace.

Among the included control variables, several points are noteworthy. First, for both observation periods, the contribution of education remains relatively large at both the upper and the lower tail. We thus confirm findings by Biewen et al. (2017), who find particularly strong effects of education for males. This is an important observation since it shows that our task measures are not simply reflecting workers’ human capital. Moreover, we find an important role for firm characteristics, which comprise occupation, sector and region dummies, to explain increases in upper and lower-tail wage inequality. We thus confirm previous findings by Antonczyk et al. (2010), Baumgarten et al. (2020) and, in particular, Card et al. (2013), who point

to the importance of firm heterogeneity in explaining increases in wage inequality.

Lastly, we should stress that a number of contributions suggest that labour market institutions such as unions and implicit minimum wages compressed the wage structure in Germany, particularly at the lower end of the wage distribution before the mid-1990s (Fitzenberger, 1999; Dustmann et al., 2009). In our decomposition results, this could be reflected by the relatively large intercept in the early periods, which accounts for residual wage inequality.

The evidence for more recent periods is more mixed. While Dustmann et al. (2009) and Biewen and Seckler (2017) find a strong role for deunionization in explaining wage inequality, Antonczyk et al. (2010) and Card et al. (2013) underline that the growing heterogeneity in wage setting at the firm level is a key driver. Finally, Dustmann et al. (2014) emphasize that after 1995, wage inequality has increased most among workers covered by collective bargaining. Overall, since the mid-1990s, it appears that firm heterogeneity and differing firm practices play a larger role than institutional factors as drivers in wage inequality. Our decomposition appears to support this conclusion, given that the intercepts in our later observation periods are much smaller than in the earlier two periods.

On a last note, we should emphasise that our analysis is restricted to the fulltime employed, while the last two decades have witnessed a growing prevalence of so-called atypical employment contracts such as temporary work, fixed-term contracts or ‘mini-jobs’ (Eichhorst and Buhlmann, 2015; Eichhorst and Marx, 2019). These more precarious forms of employment are undoubtedly contributing to larger levels of wage inequality than the ones we can observe in our study.

Table 7: Detailed decomposition results, 1991/92 - 1998/99

Inequality measure	90-10	90-50	50-10
A. Detailed composition effects			
Abstract tasks	-0.0005*** (0.0001)	-0.0003*** (0.0000)	-0.0002*** (0.0000)
Interactive tasks	0.0053*** (0.0000)	0.0022*** (0.0000)	0.0031*** (0.0000)
Manual tasks	0.0002*** (0.0000)	0.0001*** (0.0000)	0.0001*** (0.0000)
Routine tasks	-0.0006*** (0.0000)	-0.0000** (0.0000)	-0.0006*** (0.0000)
Education	-0.0022*** (0.0000)	-0.0020*** (0.0000)	-0.0002*** (0.0000)
Experience	-0.0064*** (0.0000)	-0.0048*** (0.0000)	-0.0016*** (0.0000)
Firm controls	0.0104*** (0.0000)	0.0024*** (0.0000)	0.0080*** (0.0000)
Composition effect	0.0062*** (0.0001)	-0.0024*** (0.0000)	0.0086*** (0.0000)
B. Detailed wage structure effects			
Abstract tasks	0.0292*** (0.0005)	0.0318*** (0.0004)	-0.0026*** (0.0005)
Interactive tasks	0.0004*** (0.0001)	-0.0012*** (0.0000)	0.0017*** (0.0000)
Manual tasks	-0.0061*** (0.0003)	-0.0027*** (0.0002)	-0.0034*** (0.0003)
Routine tasks	0.0168*** (0.0002)	-0.0012*** (0.0001)	0.0180*** (0.0001)
Education	0.0243*** (0.0008)	0.0060*** (0.0005)	0.0183*** (0.0007)
Experience	0.0145*** (0.0004)	0.0041*** (0.0003)	0.0103*** (0.0003)
Firm controls	-0.0061*** (0.0002)	0.0004** (0.0002)	-0.0065*** (0.0002)
Constant	0.0370*** (0.0011)	0.0553*** (0.0007)	-0.0184*** (0.0010)
Wage structure effect	0.1101*** (0.0007)	0.0926*** (0.0006)	0.0175*** (0.0007)

Notes: Bootstrapped standard errors in parentheses. 100 replications of the full procedure.
*** p<0.01, ** p<0.05, * p<0.1

Table 8: Detailed decomposition results, 1998/99 - 2005/06

Inequality measure	90-10	90-50	50-10
A. Detailed composition effects			
Abstract tasks	0.0011*** (0.0000)	0.0003*** (0.0000)	0.0008*** (0.0000)
Interactive tasks	0.0024*** (0.0000)	0.0004*** (0.0000)	0.0020*** (0.0000)
Manual tasks	0.0001*** (0.0000)	-0.0000*** (0.0000)	0.0001*** (0.0000)
Routine tasks	-0.0029*** (0.0000)	0.0000*** (0.0000)	-0.0029*** (0.0000)
Education	-0.0017*** (0.0000)	-0.0005*** (0.0000)	-0.0012*** (0.0000)
Experience	-0.0004*** (0.0000)	-0.0023*** (0.0000)	0.0019*** (0.0000)
Firm controls	0.0116*** (0.0001)	-0.0012*** (0.0000)	0.0128*** (0.0000)
Composition effect	0.0102*** (0.0001)	-0.0032*** (0.0000)	0.0134*** (0.0001)
B. Detailed wage structure effects			
Abstract tasks	-0.0006*** (0.0000)	-0.0023*** (0.0000)	0.0016*** (0.0000)
Interactive tasks	0.0230*** (0.0002)	0.0080*** (0.0001)	0.0150*** (0.0002)
Manual tasks	-0.0007*** (0.0002)	0.0002* (0.0001)	-0.0009*** (0.0002)
Routine tasks	0.0015*** (0.0000)	0.0005*** (0.0000)	0.0010*** (0.0000)
Education	0.0304*** (0.0009)	0.0026*** (0.0006)	0.0278*** (0.0007)
Experience	0.0058*** (0.0005)	-0.0066*** (0.0003)	0.0124*** (0.0004)
Firm controls	0.0132*** (0.0003)	-0.0007*** (0.0002)	0.0139*** (0.0003)
Constant	0.0120*** (0.0011)	0.0451*** (0.0007)	-0.0331*** (0.0009)
Wage structure effect	0.0845*** (0.0003)	0.0468*** (0.0002)	0.0377*** (0.0003)

Notes: Bootstrapped standard errors in parentheses. 100 replications of the full procedure.
*** p<0.01, ** p<0.05, * p<0.1

7 Conclusion

West Germany has experienced substantial increases in wage inequality over the past decades. During the same time, the task content of occupations has changed as well. In this paper, we relate changes in tasks to changes in the wage distribution. Our task data are the Employment Surveys, representative cross-section surveys conducted roughly every seven years since 1979. We describe occupations as task bundles, distinguishing between abstract, interactive, manual and routine tasks. We match this data to the Sample of Integrated Labour Market Biographies, a 2% random sample of all social security records in Germany, covering the employment histories of about 1.5 million individuals from 1975 through 2008.

We use a two-stage decomposition procedure introduced by Firpo et al. (2009), which allows to account for multiple explanatory factors in a regression framework and to investigate changes along the entire wage distribution. Just like in a standard OB decomposition, we can first break down observed wage changes into composition effects, i.e. changes due to varying worker characteristics, and wage structure effects, i.e. changes in the returns to those characteristics, and then look at the contribution of single covariates. Overall, we find wage structure effects to strongly dominate composition effects.

Our main findings can be summarised as follows. From the late 1970s until the early 1990s, education and experience are important drivers for both upper and lower-tail wage inequality. Since the early 1990s, occupational tasks start to matter as well. Specifically, through the 1990s, changes in the returns to abstract tasks explain about half of the increase in upper-

tail wage inequality, while changes in the returns to routine tasks explain almost half of the increase in lower-tail inequality. These results thus follow earlier evidence presented by Autor et al. (2003) that computer technology is complementing abstract tasks, leading to productivity and wage gains, and substituting routine tasks, leading to productivity and wage losses. In an important departure from the routine-biased technological change (RBTC) literature, we find routine-task intensity to be highest at the lowest wage percentiles and not in the middle of the distribution.

In the 2000s, we find abstract and routine tasks not to matter greatly anymore, while changes in the returns to interactive tasks can explain almost a fifth of the increase in lower-tail inequality. This result speaks to the relatively recent studies that focus on the rising complementarity between cognitive and social skills (Borghans et al., 2014; Weinberger, 2014; Deming, 2017). Interactive tasks comprise teaching and coordinating, tasks where face-to-face meetings and interactions are important. These tasks are less likely to be affected by technological change, and may thus not have played an important role in the 1990s.

Our study adds several new insights to the discussion on recent changes in wage inequality. First, we show descriptively that changes in the West German wage structure differ from those observed for Anglo-Saxon countries. The middle in West Germany remained largely stable, while workers at the lower part of the distribution suffered wage losses and workers at the upper part enjoyed wage gains. Second, we show that the task content within occupations varies substantially over time. This indicates that new technologies have not simply wiped out certain occupations, but fundamentally changed the tasks required in these occupations (Autor, 2013). This finding could imply that the institutionalised occupational structure preva-

lent in West Germany, where occupational task requirements are regularly revised and adapted, could play a crucial role for employment dynamics and the degree of wage polarisation.

Third, our decomposition results show the importance of changing task content in explaining changes in the wage distribution. That these effects are non-monotonic emphasises the necessity for a detailed decomposition at different parts of the distribution. Our results thus also speak to recent studies which show that within-occupation changes in task composition significantly contribute to a rise in income inequality in the US (Atalay et al., 2019; Michaels et al., 2019). We are the first to investigate this relationship for Germany. By quantifying the contribution of changes in tasks to changes in the wage structure, we allow for a better understanding of why some occupations have incurred sharp wage losses, while others have experienced strong gains.

References

- Acemoglu, D., and Autor, D. H. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, 4, 1043–1171.
- Antonczyk, D., DeLeire, T., and Fitzenberger, B. (2018). Polarization and rising wage inequality: Comparing the US and Germany. *Econometrics*, 6(2), 20.
- Antonczyk, D., Fitzenberger, B., and Leuschner, U. (2009). Can a task-based approach explain the recent changes in the German wage structure? *Jahrbücher für Nationalökonomie und Statistik*, 229(2-3), 214–238.
- Antonczyk, D., Fitzenberger, B., and Sommerfeld, K. (2010). Rising wage inequality, the decline of collective bargaining, and the gender wage gap. *Labour Economics*, 17(5), 835–847.
- Atalay, E., Phongthientham, P., Sotelo, S., and Tannenbaum, D. (2018). New technologies and the labor market. *Journal of Monetary Economics*, 97, 48–67.
- Atalay, E., Phongthientham, P., Sotelo, S., and Tannenbaum, D. (2019). The inequality implications of occupational evolution. *Unpublished Working Paper (University of Michigan)*.
- Atalay, E., Phongthientham, P., Sotelo, S., and Tannenbaum, D. (2020). The evolution of work in the United States. *American Economic Journal: Applied Economics*, 12(2), 1–34.
- Autor, D. H. (2013). The “task approach” to labor markets: An overview. *Journal for Labour Market Research*, 46(3), 185–199.
- Autor, D. H., and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the U.S. labor market. *The American Economic Review*, 103(5), 1553–1597.
- Autor, D. H., Katz, L. F., and Kearney, M. S. (2006). Measuring and interpreting trends in economic inequality. *AEA Papers and Proceedings*, 96(2), 189–194.
- Autor, D. H., Katz, L. F., and Kearney, M. S. (2008). Trends in U.S. wage inequality: Revising the revisionists. *The Review of Economics and Statistics*, 90(2), 300–323.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333.

- Baumgarten, D., Felbermayr, G., and Lehwald, S. (2020). Dissecting between-plant and within-plant wage dispersion: Evidence from Germany. *Industrial Relations: A Journal of Economy and Society*, 59(1), 85–122.
- Baumgarten, D., Geishecker, I., and Görg, H. (2013). Offshoring, tasks, and the skill-wage pattern. *European Economic Review*, 61, 132–152.
- Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy*, 70(5, Part 2), 9–49.
- Berman, E., Bound, J., and Griliches, Z. (1994). Changes in the demand for skilled labor within U.S. manufacturing: Evidence from the annual survey of manufactures. *The Quarterly Journal of Economics*, 109(2), 367–397.
- Biewen, M., Fitzenberger, B., and De Lazzer, J. (2017). Rising wage inequality in Germany: Increasing heterogeneity and changing selection into full-time work. *ZEW-Centre for European Economic Research Discussion Paper No. 17-048*.
- Biewen, M., and Seckler, M. (2017). Changes in the German wage structure: Unions, internationalization, tasks, firms, and worker characteristics. *IZA Discussion Paper No. 10763*.
- Black, S. E., and Spitz-Oener, A. (2010). Explaining women’s success: Technological change and the skill content of women’s work. *The Review of Economics and Statistics*, 92(1), 187–194.
- Blinder, A. S. (1973). Wage discrimination: Reduced form and structural estimates. *Journal of Human Resources*, 436–455.
- Böhm, M., von Gaudecker, H.-M., and Schran, F. (2016). The polarization of skill prices in Germany, 1985-2010. *Unpublished Working Paper (University of Bonn)*.
- Bonacini, L., Gallo, G., and Scicchitano, S. (2021). Working from home and income inequality: Risks of a ‘new normal’ with COVID-19. *Journal of Population Economics*, 34(1), 303–360.
- Borghans, L., Ter Weel, B., and Weinberg, B. A. (2014). People skills and the labor-market outcomes of underrepresented groups. *ILR Review*, 67(2), 287–334.
- Bound, J., and Johnson, G. E. (1992). Changes in the structure of wages in the 1980’s: An evaluation of alternative explanations. *The American Economic Review*, 82(3), 371–392.
- Burda, M. (2000). East-West German wage convergence after reunification: Migration or institutions? *Unpublished Working Paper (Humboldt University Berlin)*.

- Card, D. E., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of West German wage inequality. *The Quarterly Journal of Economics*, 128(3), 967–1015.
- Cortes, G. M., Jaimovich, N., and Siu, H. E. (2018). The “end of men” and rise of women in the high-skilled labor market. *NBER Working Paper No. 24274*.
- Costinot, A., and Vogel, J. (2010). Matching and inequality in the world economy. *Journal of Political Economy*, 118(4), 747–786.
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4), 1593–1640.
- Deming, D. J., and Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(1), 337–369.
- DiNardo, J. E., Fortin, N. M., and Lemieux, T. (1996). Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 64(5), 1001–1044.
- Dustmann, C., Fitzenberger, B., Schönberg, U., and Spitz-Oener, A. (2014). From sick man of Europe to economic superstar: Germany’s resurgent economy. *Journal of Economic Perspectives*, 28(1), 167–188.
- Dustmann, C., Ludsteck, J., and Schönberg, U. (2009). Revisiting the German wage structure. *The Quarterly Journal of Economics*, 124(2), 843–881.
- Eichhorst, W., and Buhlmann, F. (2015). Die Zukunft der Arbeit und der Wandel der Arbeitswelt. *IZA Standpunkte*, 77.
- Eichhorst, W., and Marx, P. (2019). Der Wandel der Arbeitswelt als Herausforderung für die Sozialpolitik. In H. Obinger and M. G. Schmidt (Eds.), *Handbuch Sozialpolitik* (pp. 409–430). Wiesbaden: Springer.
- Fairlie, R. W. (2005). An extension of the Blinder-Oaxaca decomposition technique to logit and probit models. *Journal of Economic & Social Measurement*, 30(4), 305–316.
- Firpo, S. P., Fortin, N. M., and Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, 77(3), 953–973.
- Firpo, S. P., Fortin, N. M., and Lemieux, T. (2011). Occupational tasks and changes in the wage structure. *IZA Discussion Papers No. 5542*.

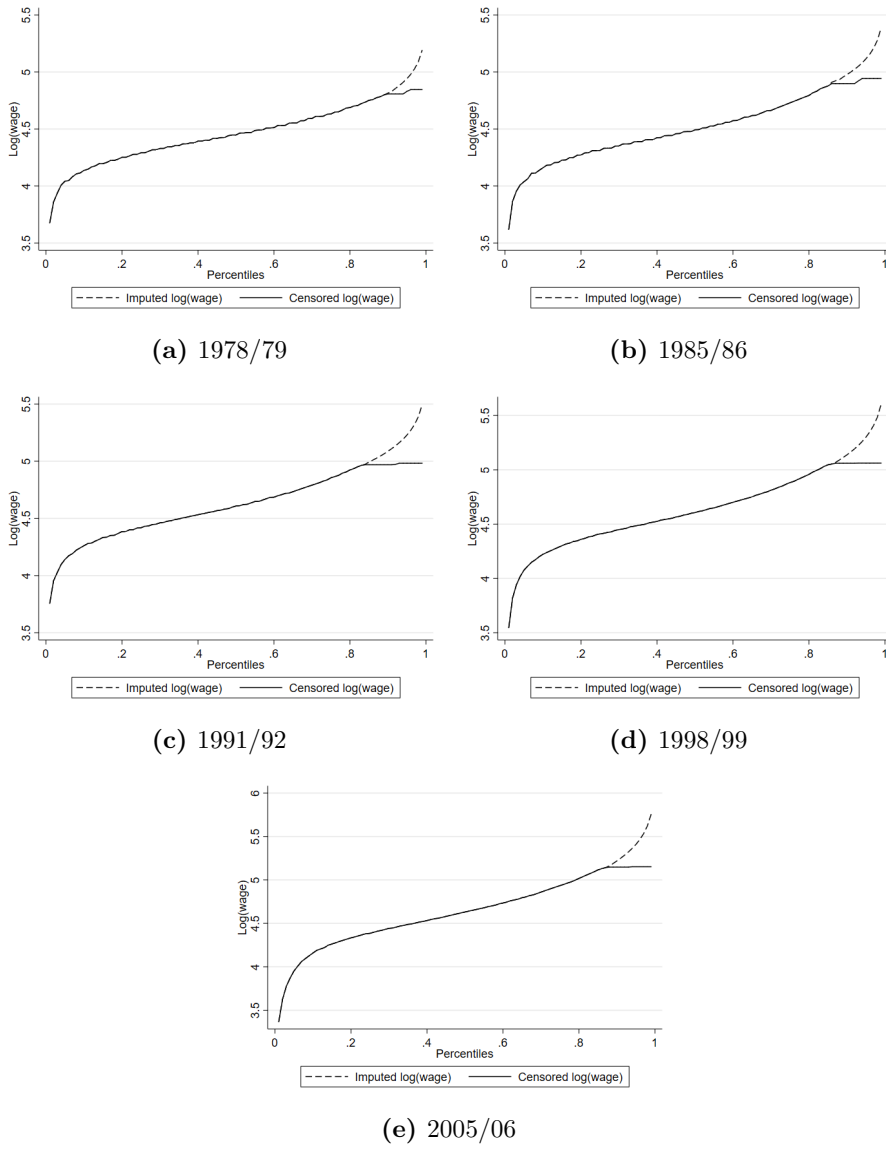
- Firpo, S. P., Fortin, N. M., and Lemieux, T. (2018). Decomposing wage distributions using recentered influence function regressions. *Econometrics*, 6(2), 28.
- Fitzenberger, B. (1999). *Wages and employment across skill groups: An analysis for West Germany*. Heidelberg: Physica-Verlag.
- Fortin, N. M., Lemieux, T., and Firpo, S. P. (2011). Decomposition methods in economics. *Handbook of Labor Economics*, 4, 1–102.
- Gartner, H. (2005). The imputation of wages above the contribution limit with the German IAB employment sample. *FDZ Methodenreport 02*.
- Gernandt, J., and Pfeiffer, F. (2007). Rising wage inequality in Germany. *Jahrbücher für Nationalökonomie und Statistik*, 227(4), 358–380.
- Goos, M., and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *The Review of Economics and Statistics*, 89(1), 118–133.
- Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in Europe. *The American Economic Review*, 99(2), 58–63.
- Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *The American Economic Review*, 104(8), 2509–2526.
- Green, F. (2012). Employee involvement, technology and evolution in job skills: A task-based analysis. *ILR Review*, 65(1), 36–67.
- Hershbein, B., and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *The American Economic Review*, 108(7), 1737–1772.
- Hoffmann, F., Lee, D. S., and Lemieux, T. (2020). Growing income inequality in the United States and other advanced economies. *Journal of Economic Perspectives*, 34(4), 52–78.
- Juhn, C., Murphy, K. M., and Pierce, B. (1993). Wage inequality and the rise in returns to skill. *Journal of Political Economy*, 101(3), 410–442.
- Katz, L. F. (2014). America’s jobs challenges and the continuing role of the US Department of Labor. *ILR Review*, 67(3-suppl), 578–583.
- Katz, L. F., and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1), 35–78.

- Kohn, K. (2006). Rising wage dispersion, after all! The German wage structure at the turn of the century. *ZEW-Centre for European Economic Research Discussion Paper No. 06-031*.
- Lemieux, T. (2006). Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? *American Economic Review*, *96*(3), 461–498.
- Levy, F., and Murnane, R. J. (1992). US earnings levels and earnings inequality: A review of recent trends and proposed explanations. *Journal of Economic Literature*, *30*(3), 1333–1381.
- Machado, J. A. F., and Mata, J. (2005). Counterfactual decomposition of changes in wage distributions using quantile regression. *Journal of Applied Econometrics*, *20*(4), 445–465.
- Michaels, G., Rauch, F., and Redding, S. J. (2019). Task specialization in U.S. cities from 1880 to 2000. *Journal of the European Economic Association*, *17*(3), 754–798.
- Mincer, J. A. (1974). *Schooling, experience, and earnings*. New York: Columbia University Press.
- Möller, J. (2015). Did the German model survive the labor market reforms? Hat das Modell Deutschland die Arbeitsmarktreforment überlebt? *Journal for Labour Market Research*, *48*(2), 151–168.
- Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review*, 693–709.
- Palomino, J. C., Rodríguez, J. G., and Sebastian, R. (2020). Wage inequality and poverty effects of lockdown and social distancing in Europe. *European Economic Review*, *129*, 103564.
- Riphahn, R. T., and Schnitzlein, D. D. (2016). Wage mobility in East and West Germany. *Labour Economics*, *39*, 11–34.
- Rohrbach-Schmidt, D., and Tiemann, M. (2013). Changes in workplace tasks in Germany – Evaluating skill and task measures. *Journal for Labour Market Research*, *46*(3), 215–237.
- Saint-Paul, G. (2001). On the distribution of income and worker assignment under intrafirm spillovers, with an application to ideas and networks. *Journal of Political Economy*, *109*(1), 1–37.
- Saint-Paul, G. (2008). *Innovation and inequality: How does technical progress affect workers?* Princeton, N.J. and Oxford: Princeton University Press.

- Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics*, 24(2), 235–270.
- Weinberger, C. J. (2014). The increasing complementarity between cognitive and social skills. *The Review of Economics and Statistics*, 96(4), 849–861.

Appendix

Figure A1: Imputed vs Censored log(wage)



Notes: This figure plots the values of the imputed and censored log wages.