Is it Bad to be Green in a Greying Firm? An Analysis of the Impact of Postponed Retirements on Younger Workers' Wage Growth☆

- Preliminary draft. Please don't cite without authors' permission -

Peter Berg[a], Marissa Eckrote[a], Mary Hamman[b], Daniela Hochfellner[c], Matthew M. Piszczek[d], Christopher Ruhm[e]

[a]*Michigan State University*
[b]*University of Wisconsin - La Crosse*
[c]*New York University*
[d]*Wayne State University*
[e]*University of Virginia*

Abstract

Most countries with pay-as-you-go public pension systems have raised the age of eligibility for full benefits and enacted other reforms to encourage longer working lives. A large body of literature shows that older individuals work longer when pension rules change in ways that make early retirements less attractive. However, we don't fully understand how this affects the careers of younger workers. This paper examines the effects of postponed retirements on the wage growth of younger coworkers employed in the same establishment following a policy reform that raised the pension eligibility-ages of older workers. Using administrative employment data, we show that the gradual increases in pensionable age introduced in a 1992 reform had highly variable impacts across establishments because of pre-policy differences in worker age distributions. This variation serves as the source of identification for examining the influence of delayed retirements on wage growth among younger colleagues. We find larger wage increases occur less frequently in establishments with higher shares of older workers. Specifically, when the share of workers over age 58 increases by 1 standard deviation (3.21 percentage points) the probability of receiving a wage increase of 10% or more falls by 3.9 percentage points among workers age 25 to 57, which is a 10% change relative to the average probability.

Keywords: *Pension Reform, Job trajectory, Administrative data, Shift-Share analysis*

☆Funding for this research was provided by the Alfred P. Sloan Foundation.
*Corresponding author
Email address: `eckrotem@msu.edu` (Marissa Eckrote)

1. Introduction

Workforces are aging across the globe. The number of people 65 or older is projected to nearly double in the United States from 2012 to 2050, rising to 83.7 million in the later year (Ortman et al., 2014) . Many European countries are already further along this trajectory and thus demonstrate the situation the U.S. will face in the near future. Aging threatens the solvency of public pension systems and, in response, a majority of OECD nations have increased the "pensionable age" for full old age benefits (OECD, 2011). For example, in the United States, the Social Security "full-retirement" age is gradually being increased from 65 to 67. A similar 1992 reform in Germany raised pensionable ages for most workers born after 1938 from 60 to 65. Prior research demonstrates individuals change retirement behavior in response to these increases in pensionable age (e.g. Stock and Wise 1988; Krueger and Pischke 1989; Börsch-Supan 2000; Duval 2003; Coile 2004; Barr and Diamond 2006).

The responses of firms to the resulting postponements of retirements is less well understood. For example after pension reforms, employers may limit promotion and compensation of younger workers. This potential linkage is the basis for the lump of labor "conjecture," which hypothesizes that earlier retirements among older workers create more job opportunities for younger individuals (Walker, 2007). Some economists are skeptical because this conjecture is predicated on the assumption of a fixed stock of jobs, with each retirement then opening up a new position (or a higher level position) that could be filled by a younger person. Economic theory predicts this is false for the economy as a whole and existing empirical studies support this prediction (Gruber and Wise, 2010; Brugiavini and Peracchi; Jousten et al., 2010; Munnell and Wu, 2012). Moreover, if there are complementarities between older and younger workers in the production process, delayed retirements could instead benefit those who are early in their careers. However, existing empirical evidence is based on analysis of aggregate labor market data or samples where it is not possible to identify workers with the same employer. So, underlying the apparent lack of association in the aggregate, there could be important effects that arise at the establishment level. This may be especially likely when the impacts of policies that encourage later retirements are unequally distributed across firms. One of the key strengths of analyses using matched employer-employee micro data has been to uncover economically important heterogeneity not fully revealed in prior analyses of aggregate and survey data (e.g., Davis, Faberman, & Haltiwanger 2006). Our work therefore seeks to revisit this question using matched employer-employee data.

The paper most similar to ours studies a similar reform in Italy and finds that when a policy that increases pensionable age is implemented, firms treat older workers and younger works as substitutes (Paradisi and Bovini, Working Paper). In addition, they also find that this change in pensionable age decreases labor earnings. One key difference between this study and ours is the context in which the reform took place. The Italian Fornero reform was implemented in 2011 as part of a much larger response to the financial crisis taking place at the time. This was not the case when Germany implemented their reform in 1992. Another key difference is that the decrease in wage earnings found from the Italian reform is including decreases in wages resulting from layoffs. Once these layoffs are accounted for the authors find much smaller effects specifically for young workers.

Using administrative employment data for German workers in more than 11,000 establishments, we show that the increases in pensionable age introduced by a 1992 reform had highly variable impacts across establishments because of pre-policy differences in worker age distributions. This variation serves as the source of identification for examining the influence of older workers' delayed retirements on wage growth among their colleagues. Germany is an advantageous laboratory for this research because the 1992 reform increased pensionable ages by 5 years across 7 birth cohorts. In addition, the vast majority of German retirement wealth comes from public pension benefits (Börsch-Supan and Wilke, 2004). This large change in the primary source of retirement wealth over a short period of time therefore had a substantial and plausibly exogenous impact on retirement timing. In this context it should be easier to separately identify any impact of working longer on younger workers wage growth than for the US or other countries where pension reforms were phased in more gradually or represent a smaller fraction of overall retirement savings.

Unike studies using aggregate data (Gruber and Wise, 2010; Brugiavini and Peracchi; Jousten et al., 2010; Munnell and Wu, 2012) and in keeping with (Paradisi and Bovini, Working Paper), we find some evidence of a negative impact of postponed retirements of older workers on their younger colleagues wage growth. Our estimates imply the likelihood of large wage increases (more than 10% relative to the prior year) among workers aged 25 to 57 decreases as the share of colleagues over age 58 increases. However, we find no change in the probability of smaller wage increases or the likelihood of wage reductions. The estimated negative impacts are approximately equal across the youngest (age 25 to 39) and oldest (age 40 to 57) workers. We conclude additional older workers may be creating a "blockage" in the flow of workers up job ladders, and this blockage affects the entire age distribution.

2. Institutional Background

During our study period, the German pension system was a pay-as-you-go scheme. Current tax contributions were used to fund payments to current retirees that provide a minimum standard of living in retirement for all private and public sector employees entitled to social security. Private savings were negligible for most workers retiring before 2010. Self-employed workers and civil servants are excluded from the pension system, which covers about 90 percent of the German workforce (Richter and Himmelreicher, 2008) and accounts for approximately 85 percent of retirement income (Börsch-Supan, 2000).

In 1992, Germany announced gradual increases in pensionable age. Before the reform, German women could retire with full pension benefits at age 60. Although the official pensionable age for men was 65, in practice many males qualified for full benefits at age 60 because they were classified as unemployed.

Workers could receive unemployment insurance payments starting at age 58 to bridge the gap until eligibility for full pension benefits at age 60 and this was a highly popular pathway into retiremnt. Prior to the reform, approximately 45 percent of 59-year old men self-identified as "retired" and only 20 percent of new pension claimants were age 65 (Börsch-Supan and Wilke, 2004). [1] The 1992 reform gradually increased ages of eligibility for full benefits to 65 for all except disabled workers beginning with the 1938 birth cohort. Because the 1938 cohort reached age 58 in 1996, these reforms first began to postpone claiming as early as 4 years after they were announced. All changes in eligibility associated with the 1992 reform were phased in between the 1938 and 1945 birth cohorts and thus were fully implemented by 2011 (Börsch-Supan and Wilke, 2004).

The reform created differences in pensionable age between adjacent birth cohorts of 6 months to 1.5 years, as shown in Figure 1. This heterogeneity across individuals created variation in the impact of the pension reform across employers and provides a source of identifying variation for estimating how delayed retirements among older workers influenced the wage growth of those earlier in their careers.

[1] This unemployment pathway to retirement remained open after the 1992 reform, but the duration of unemployment benefits was still only two years during our study period, so that the reform created an incentive to work beyond age 58 for affected cohorts.

Figure 1: Gap between the old and new effective pensionable ages by cohort and sex

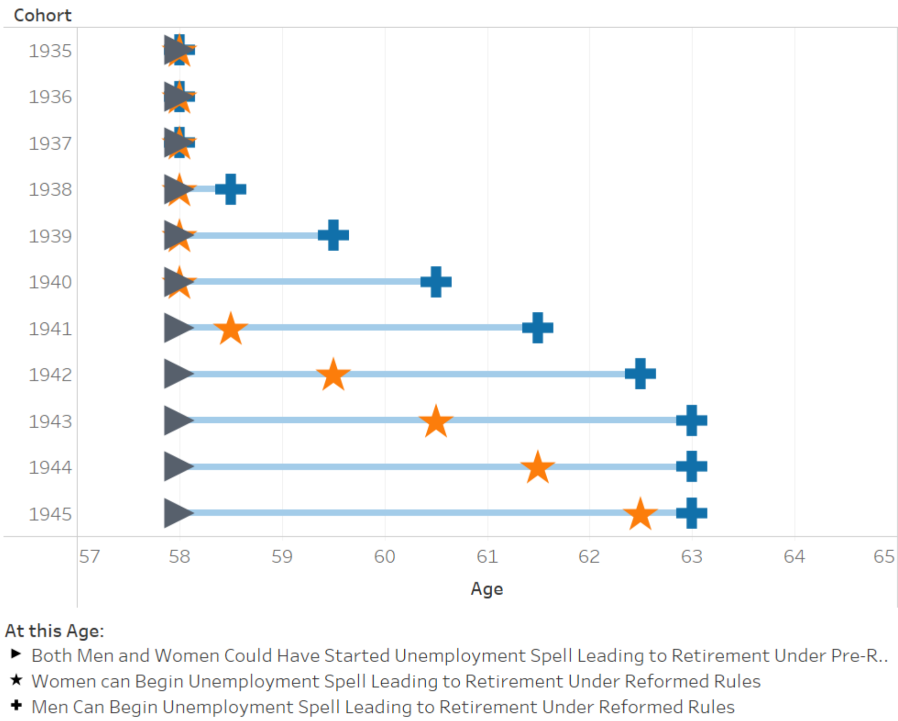

3. Data

Our data come from a rich source of matched employer-employee administrative data: the cross sectional Linked Employer Employee Data of the Institute for Employment Research (LIAB).[2] The LIAB offers the unique advantage of matching survey data from a national stratified random sample of German establishments to social security employment records for all establishment employees covered by the system. Because the LIAB data begins in 1993 and the pensionable age change was announced in 1992, we also use a custom made

[2]This study uses the Linked-Employer-Employee Data (LIAB) [cross-sectional model 2 1993-2010 (LIAB QM2 9310)] from the IAB. Data access was provided via on-site use at the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) and subsequently remote data access.

demographic file containing employer information in 1990.[3]

The LIAB matches administrative employment records to establishment survey information. The baseline for the sampling is the IAB Establishment Panel (IABBP), which collects data on about 15,500 establishments per year (Fischer et al., 2009). Our sample includes information for all individuals who work in these surveyed establishments on June 30th in each year. For each of these workers information is provided on employment status on June 30th, wage, and detailed occupation. Socio-demographic variables include sex, year of birth, and educational attainment (Klosterhuber et al., 2016). Because the data include all the workers in each establishment in a given year we can examine relationships between the share of older workers employed and wage growth among their younger colleagues within the same establishment.

The demographic file contains age distributions in 1990 by sex for each establishment that existed in 1990 and was later included in the LIAB in any year between 1993 and 2014. We use this information to construct instruments (see Section 4.2 for more details) that address the potential endogeneity of contemporaneous shares of older workers who are impacted by the 1992 pension reform. Doing so requires the exclusion of some firms where administrative data are lacking prior to the 1992 reform, including most East German firms.[4] Thus, our analytic sample is based on all West German establishments that existed in 1990 and in the LIAB. In total, we have data for 11,165 establishments.

Although the LIAB offers several unique advantages, it also has limitations. Most importantly, wages, are reported as the daily rate in Euros and may or may not contain bonuses or other incentive pay. This may lead to misclassification of temporary payments as earnings growth. Also, the data do not contain precise measures of work hours – we can observe only whether a work is employed full-

[3]The demographic file comes from the Employment History data (BeH) and was provided by the Research Data Center of the Institute for Employment Research (FDZ). We thank Andreas Ganzer for sampling the data for us and supporting us with de-identification of the data.

[4]Records are reliably complete for East German firms beginning in 1993.

time or part-time. This may lead us to misclassify changes in work hours as changes in wages. In addition, wages are censored at the contribution threshold for social security contributions. We discuss these issues in detail below and provide robustness checks to assess the implications for our results.

3.1. Sample Restrictions

Our main estimates impose several sample restrictions. First, we limit analysis to workers employed full-time in regular employment contracts (not temporary or trainee contracts). We do this to reduce misclassification of wage changes attributable to changes in hours of work. Second, to be included in our sample, a worker must have worked in the same establishment in the prior year. This restriction is needed to compute changes in wages. In addition, to limit the effects of potentially erroneous extreme values for year-over-year percentage changes in wages, we trim the wage change variable at the 1st and 99th percentiles of the distribution (-15% and 36%). We also drop people whose earnings are above the social security threshold (approximately 4% of the person-year observations) as we don't know exactly how much they earn. Finally, we restrict the sample to establishments with at least five workers, to avoid extreme variation in the share of workers over age 58 that occurs over time in very small firms where a single retirement or hire comprises a very large share of overall employment.

3.2. Measuring Wage Growth

We use within-person longitudinal variation in wages to measure wage growth. $wagegrowth_{ijt}$, is the percentage change in wages across two adjacent years of data for worker i in establishment j, computed with the first year's wages as the denominator:

$$wagegrowth_{ij,t} = 100 * \frac{wage_{ij,t} - wage_{ij,t-1}}{wage_{ij,t-1}}. \tag{1}$$

3.3. Measuring Working Longer

To measure the impact of older workers postponing retirement (working longer), we construct yearly shares of employees in each establishment age 58 and older, $share58pls_{jt}$. We use age 58 as the threshold because, as explained above, this was the effective pre-reform pensionable age.

4. Empirical Strategy

To examine the relationship between employment of older workers and wage growth of their younger colleagues, we estimate Equation [2]:

$$wagegrowth_{ij,t} = \beta_1 share58pls_{j,t} + \beta_2 X_{ij,t} + u_t + \epsilon_{i,jt}. \tag{2}$$

In [2], $share58pls_{jt}$ is the share of workers in establishment j who are age 58 and older in year t. X_{ijt} is a vector of controls including gender, education, experience, industry, inflows and outflows of workers, establishment size, occupation, year of hire fixed effects, and state.

4.1. Possible Selection and Endogeneity Bias

There are (at least) two obvious reasons to suspect the OLS estimate of $\hat{\beta}_1$ in Equation [2] will be biased. First, if there is a causal negative relationship between the share of older workers and wage growth, high ability younger workers may not apply to, or would leave, firms with a large share of older workers. This selection would exaggerate any negative impact of older workers on their younger colleagues' wage growth. Second, employers may attempt to mitigate any negative impact of workforce aging on opportunities for their younger employees through tactics that make earlier retirements more attractive to their older employees, like buyouts or partial retirement offers. Employers that can afford to pay these buyouts may be higher paying employers already, leading to negative bias. Conversely, employers offering these retirement incentives might be those that would struggle to pay competitive wages or retain their younger workers, leading to positive bias. Although some of these responses may be observable in the establishment survey data, informal offers and implicit pressure to retire likely exist and are not observable.

4.2. Identification Strategy

To address these potential biases, we use changes in pensionable age due to the 1992 reform as an instrument for $share58pls_{jt}$. A contemporaneous measure of the impact of the reform, like the number or share of workers eligible to retire in year t under the new rules, is itself a function of the endogenous number or share of older workers. Instead, we use pre-reform age distributions for each establishment to construct shares of workers who would have been in the gap between the pre- and post-reform pensionable ages in each subsequent year t if the establishment had the same age specific probabilities of entry and exit as others in their industry. This approach is based on the construction of a shift-share instrument, as detailed below.

An increase in the pensionable age creates a gap between the age at which retirement with full benefits was feasible under the old versus the new policy. Figure 1 depicts the gaps created by the 1992 reform. We use the 1990 demographic file to construct pre-reform counts of workers in each cohort shown in Figure 1 by sex for each of the establishments in our analytic sample. These counts comprise the "share" portion of the instrument.

The shifts are computed from the fitted values after estimating the following two regressions using 1993-2014 data separately for each of 11 industry sectors by sex (44 regressions in total).

$$begin_{ij,t} = \beta_0 + \beta_1 age_{ij,t} + \beta_2 year_t + \beta_3 age_{ij,t} * year_t + \epsilon_{ij,t}$$

$$end_{ij,t} = \beta_0 + \beta_1 age_{ij,t} + \beta_2 year_t + \beta_3 age_{ij,t} * year_t + \epsilon_{ij,t}$$

,

where $begin_{ij,t}$ is equal to 1 for employees in their first year of employment with establishment j in year t and equal to 0 for all subsequent years. Similarly, $end_{ij,t}$ is equal to 1 in the last year of employment with establishment j, which is indicated when the employer files an end of employment notification. $age_{ij,t}$ is a vector of single year of age dummy variables, and $year_t$ is a vector of dummy variables for years 1994 through 2014, with 1993 as the reference year.

We obtain fitted values $\widehat{begin_{ij,t}}$ and $\widehat{end_{ij,t}}$ at each age and year from the 44 sex- and industry-specific regressions. These fitted values are estimated probabilities of being in the first year of employment with establishment j in year t by age, sex, and industry, and of ending employment with establishment j in year t.

We use these estimated probabilities to "age" the 1990 workforce for each establishment as follows:

$$workers_{a,t} = workers_{a-1,t-1} * [1 - \widehat{end_{a-1,t-1}} + \widehat{begin_{a,t}}]$$

where the number of workers of age a in year t is equal to the number at age $a-1$ from the prior year adjusted by the probabilities of ending employment in year $t-1$ at age $a-1$ and of beginning employment at age a in year t.

$workers_{a,t}$ is computed for each age separately for men and women by industry sector. Using these projected $workers_{a,t}$ measures, we calculate the projected number of workers in the gap between the pre-reform effective pensionable age of 58 and the applicable post-reform effective pensionable age for their cohort for each establishment j in each year t. We then divide those counts by the size of the establishment workforce in 1990. The resulting $ingap_{j,t}$ is our instrument.

The $workers_{a,t}$ measure we create and build our instrument $ingap_{j,t}$ from is a shift-share instrument. Shift-share instruments, sometimes called "Bartik instruments" after Bartik (1991), have been widely used in the immigration and the regional growth literature but have many other applications (Goldsmith-Pinkham et al., 2018). Yet recent studies raise concerns about the validity of these instruments, that can be organized into four main points. First, to meet the exclusion restriction, the initial shares used to construct the instruments must be exogenous. Second, there must be sufficient variation in initial shares to ensure the instruments for units receiving the same shift will be different.

Third, there must be no spillover across establishments that would be analogous to spatial spillovers in studies of immigration. Fourth and finally, the shares observed must be steady state and not part of an ongoing dynamic adjustment process (Jaeger et al., 2018). In our setting, the first concern is addressed using data from 1990 to construct the initial shares. This was before the policy is announced and six years before the first cohorts affected reached age 58.The second concern, sufficient variation in the initial shares, is discussed below with empirical evidence of variation. The third concern, in our context, might be a problem if one establishment's change in the shares of older workers has a causal effect on wage growth in other establishments who did not experience the same change. This could occur if, for example, an establishment reduced pay for younger workers in response to larger shares of older workers and had sufficient market power or coordinated with other firms to bring down the market wage for younger workers. In the German context, this is a risk as many workers' wages are set through collective bargaining at the industry level. However, this should bias our estimates towards zero as it would lead differentially affected first to behave more similarly. To address the fourth concern, we focus on a reform that follows a period of general stability in retirement incentives dating back to 1972.

To demonstrate variation in the initial shares used to construct our instrument, we use the 1990 demographic file to summarize variation in the shares of workers within each firm born before 1937, between 1938 and 1945, and between 1946 and 1952. Our instrument exploits variation across single year cohorts and will exhibit even greater variation than we show across these broader bands. The bands correspond to the groups of cohorts in Figure 1 who were unaffected by the reform, affected but were part of the gradual increase in pensionable age, and those who experienced the full increase. We stop at the 1952 cohort as they are the last to reach age 58 during our sample period. We then compare these establishment level shares to the shares for their industrial sector by constructing the ratios of establishment shares to industry shares. A ratio of 1 would indicate the establishment employment share for that cohort group is exactly equal to the industry employment share. The means of these ratios are reported in Appendix Table A.2 .

To illustrate how the variation translates to differences across establishments, Figure 2 provides the establishment frequency distributions for the employment share ratios for the 1938 though 1945 cohort in the sector with the smallest standard deviation (Energy and Water Supply at 0.41) and the sector with the largest standard deviation (Transportation at 0.74). As expected, the distribution for Energy and Water Supply is more compact than for Transportation but there is still substantial variation. Approximately 50% of Energy and Water Supply establishments employ shares of workers in the 1938-45 cohort that are 50% to 99% of the industry employment share, but 8% have ratios below 0.5 and 5% have ratios above 1.5.

Figure 2: Distribution of Employment Share Ratios by Industry, 1938-45 Cohort

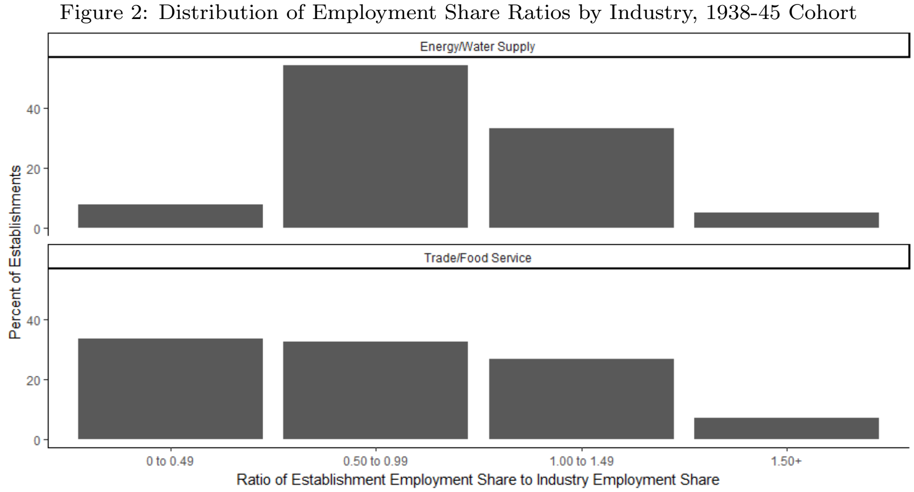

Finally, Figure 3 displays the variation in the resulting instrument overall (Panel a) and the distribution by year (Panel b). Both figures indicate there is substantial variation over time and across establishments. The between establishment standard deviation in the in-gap measure from 1996 forward is 2.79 percentage points and the within-establishment standard deviation is 4.26 percentage points.

Figure 3: Validity of Instruments

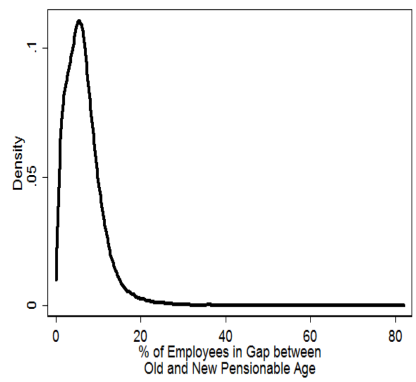

(a) Variation in Shares

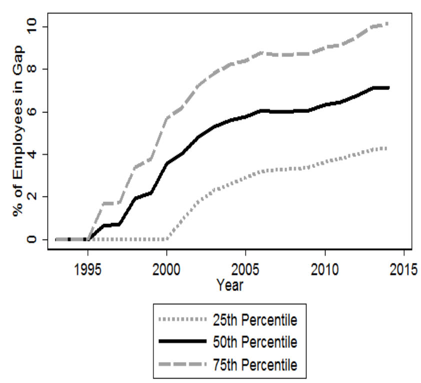

(b) Variation over Time

5. Results

Table 1 contains descriptive statistics for daily wages (in Euros) and establishment shares of workers over the age of 58 in our sample. On average, establishment shares of workers aged 58 and higher are around 5 percent. The average year over year change in wages is 12.4%.

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	N
$share58pls_{jt}$	4.41	3.21	15,027,568
Daily Wage	102.00	29.43	15,027,568
Wage Percent Change	12.4	15.5	15,027,568

Table 2 displays the results of the first stage for our IV estimates. The estimated coefficient for the instrument $ingap_{jt}$ has the expected positive sign and implies a one percentage point increase in the projected share of workers in the gap between the old and new pensionable ages is associated with a 0.27 percentage point increase in the share of the establishment's workforce aged 58 or older. This estimate is strongly statistically significant with a first stage F-Statistic of approximately 35. Coefficients for the industry dummy variables reflect expected heterogeneity in shares of older workers. Relative to the reference group (agriculture), the share of employees age 58 and above is 4.16 percentage points lower in mining and manufacturing establishments. This difference is expected because although both industries tend to require more manual labor and may feature earlier retirements on average, the inflows of younger workers to agricultural establishments may be falling faster than in manufacturing and mining. More generally, however, the overall volume of employee inflows and outflows is not associated with the share of workers over age 58. This indicates hires and departures are not happening in high enough volumes, relative to the overall size of the firm, to distort the shares from one year to the next. More educated workers are more likely to work in establishments with higher shares of workers

Table 2: First Stage Estimates

Covariate	$share58pls_{jt}$	Std. Error	p-Value
Instrument $ingap_{jt}$	0.270	0.046	0.000
Mining/Manufacturing	-4.158	0.555	0.000
Energy/Water Supply	-2.628	0.636	0.000
Construction	-2.523	0.801	0.002
Trade/Foodservice Industry	-3.191	0.573	0.000
Transportation	-3.976	0.650	0.000
Finance	-5.149	0.580	0.000
Real Estate	-3.421	0.636	0.000
Public Administration	-1.043	0.591	0.007
Administration	-0.892	0.645	0.167
STEM	-1.332	1.052	0.206
Inflows	-0.000	0.000	0.200
Outflows	-0.000	0.000	0.499
Establishment Total Employment	-0.000	0.000	0.001
Vocational Training	-0.109	0.051	0.034
University of Applied Sciences	0.155	0.086	0.071
University	0.180	0.101	0.077
Female	0.008	0.040	0.847
Years Labor Force Experience	0.001	0.006	0.808
Years Labor Force Experience^2	0.000	0.000	0.160
Intercept	9.480	0.618	0.000

The unit of observation is the person-year. Standard errors are clustered at the establishment level.

over age 58, as are female workers. Experience is not associated with higher or lower shares of workers over age 58.

In Table 3 we report the results of estimating Equation [4] both by OLS and IV, with wages in Euros ((1) and (2)), and the percentage changes in wages (this year relative to the prior year, (3) and (4)) as dependent variables. Based on the OLS estimate, a one percentage point increase in the share of workers over age 58 is associated with a €0.32 reduction in daily wages; conversely the

Table 3: Impact of Older Colleagues on Wages

	Wage*		Wage % Change*	
	(1)	(2)	(3)	(4)
Model	OLS	IV	OLS	IV
$share58pls_{jt}$	-0.324***	1.413***	0.000	-0.005***
	(0.077)	(0.374)	(0.000)	(0.002)
Mean	102.00		0.124	
SD	29.43		0.155	
N	15,027,568		15,027,568	

* Standard errors, clustered at the establishment level, are in parentheses. The unit of observation is person-year. Each regression includes a set of establishment characteristics (industry, inflows, outflow, firmsize, state), individual characteristics (education, sex, occupation, experience) and year dummies as controls. The instrumental variable regressions are estimated by two-stage least squares. One star, two stars, and three stars denote statistical significance at the 10-, 5-, and 1-percent confidence levels, respectively.

IV estimate indicates an increase of €1.41. Both are statistically significantly different from zero but, assuming approximately 250 working days in a year, the OLS estimate implies an annual decrease of €81 and the IV estimate implies an annual increase of €353.25. These are relatively small impacts. For percentage changes in wages, the OLS estimate is approximately zero and the IV estimate is negative, but both indicate that any impact is very small (0 and -0.5 percentage points, respectively). In total, these estimates indicate wage levels may be slightly higher for workers with more colleagues over age 58 but wage growth is either no different or slightly lower.

To compare our estimates to what could have been estimated without matched employer-employee data, Table 4 displays new estimates of Equation 4. In this specification the independent variable is the share of workers over age 58 in one's industrial sector.[5]

Although the estimated reduction in daily wages in Euros is larger in this

[5]This is the only sort of information that would be available from labor force survey data.

Table 4: Estimated Impact of Share of Older Workers in the Industry

	Wage*	Wage % Change*
	(1)	(2)
Model	OLS	OLS
$share58pls_{jt}$, Industry	-0.794***	-0.003
	(0.211)	(0.002)
Mean	102.00	0.124
SD	29.43	0.155
N	15,027,568	15,027,568

* Standard errors, clustered at the establishment level, are in parentheses. The unit of observation is person-year. Each regression includes a set of establishment characteristics (industry, inflows, outflow, firmsize, state), individual characteristics (education, sex, occupation, experience) and year dummies as controls. The instrumental variable regressions are estimated by two-stage least squares. One star, two stars, and three stars denote statistical significance at the 10-, 5-, and 1-percent confidence levels, respectively.

specification, it is still relatively small. Again assuming 250 working days in a year the annual wage decreases associated with a one percentage point increase in the share of workers over age 58 would be approximately €198.50. The impact on wage growth, although negative, is not statistically significantly different from zero. In total, the findings are not materially different using this approach from the results based on the share of older workers within the establishment.

6. Heterogeneity Analysis

Although there appears to be no evidence of any major negative associations between having more colleagues over age 58 and younger workers' wages in the foregoing estimates, we have defined "younger workers" to include all workers from age 25 through 57. It is possible, underlying the apparent small impacts, there is important heterogeneity by worker age. Also, these estimates assume a linear relationship between the share of workers over age 58 and wages or wage growth and potentially conceal substantial non-monotonic effects. For example, having more older workers might not lead to any measurable reduction in the size of the average wage increase but could substantially reduce the likelihood of large wage gains. To address these possibilities, we re-estimate Equation [4] separately for workers aged 25 to 39 and 40 to 57, and we specify sets of binary dependent variables equal to 1 for wage changes of specific directions and magnitudes. Figure 4 and Table A.1 contain these results. Each point in

Figure 4: Impact of Higher Shares of Older Colleagues on Wage Growth

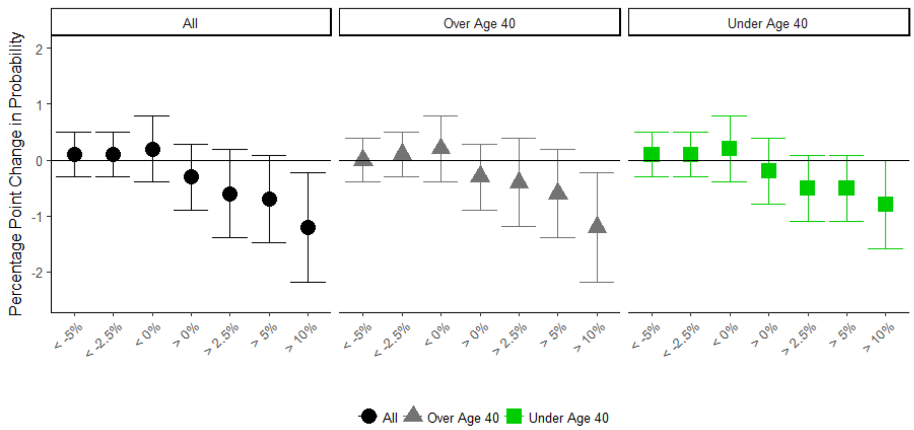

the figure represents a separate regression and plots the the IV estimate of the percentage point change in the probability of the wage increase (or decrease) described on the x-axis. Error bars denote 95% confidence interval around that estimate based on heteroskedasticity robust standard errors clustered at the establishment level. For example, the first point in each series provides the estimated percentage point change in the likelihood of having a wage decrease of 5% (or more) when the share of colleagues over age 58 rises by 1 percentage point. For workers in both age groups, this point estimate is very near zero, and not statistically significant.

The pooled estimates and estimates by age group all exhibit the same general pattern: there is no evidence of an effect on the likelihood of wage decreases but wage increases, and especially large wage increases, appear to be negatively associated with the share of workers over age 58. Although these estimates are somewhat imprecise, the estimates for large wage increases are statistically significant. Relative to the average probability of receiving these wage increases (reported in Appendix Table A.1) these estimates are relatively small. In the pooled sample, a 1.2 percentage point decline in the probability of a raise over 10% represents only a 3.3% change in probability relative to the average of 36.5%. However, this same estimate implies a 1 standard deviation change in the share of workers over age 58 (which would be a 3.21 percentage point change) would reduce the probability of a wage increase of 10% or more by 3.9%, which is a 10.6% change relative to the mean.

7. Robustness Checks

To check the robustness of our results to the model specification, sample restrictions, and dependent variable construction the following were performed. First, the dependent variable was changed to log wages. These results are consistent with those found for wages. Second, we estimate (2) by industry to investigate if the results found above are driven by specific industries. We find that there are similar patterns across most industries, but due to the highly reduced sample sizes, we lose precision in the estimates. The two industries that stand out as having a different pattern than what is observed in the estimates as a whole are administration and finance. In these cases it would make sense to see more complementarities among young and old workers due to the importance of institutional memory and experience. Figures plotting the coefficients of these regressions can be found in Appendix Tables A.1-A.10.

8. Conclusion

Using administrative micro data for workers in a large sample of German establishments, we examine how the earnings levels and growth of younger workers are affected by changes in the presence of older workers. Our key source of variation comes from a policy change that provided incentives for delayed retirement, and cross-firm heterogeneity in the share of older individuals affected by the policy change. We find some evidence of a negative impact of postponed retirements on wage growth of younger colleagues. When comparing the more "prime age workers" and the younger newer entrants into the labor force we find no difference in the level to which their wage growth is effected. Overall, our findings provide some evidence supporting the possibility that delayed retirements harm the wage growth of younger workers. The result raise the possibility of substitutablity in the production process such that establishments view both young and prime age workers as substitutes for those at the cusp of retirement.

It is interesting to compare our findings to those recently obtained by Jäger (2016) (Jäger, Working Paper) who uses similar German data but focusing on smaller firms (30 employees) to examine how the unexpected death of an employer affects the wage outcomes of other firm employees. He finds that deaths of coworkers lead to faster average wage growth but that the opposite is true when the death involves a manager or highly skilled worker in a different occupation dies. Unfortunately we are unable to differentiate between retirements by worker job level or skill using this identification strategy.

The relationships we identify in the German labor market seem likely to exist in other countries as well. However, the large differences in pensionable age between adjacent cohorts and scarcity of private pensions or retirement savings make the German labor market easier to study. As a result, the absence of an impact in the German context would presumably also apply to other countries, such as the U.S., where public pensions comprise a smaller share of retirement wealth.

Finally, we emphasize that our analysis was limited to full-time employees with regular employment contracts and that negative effects could potentially operate through other mechanisms. For instance, employers could reduce work hours, engage in more layoffs, or more extensively use temporary employment contracts. These are issues we hope to investigate in our continuing work.

References

Axel Börsch-Supan. Incentive effects of social security on labor force participation: Evidence in Germany and across Europe. *Journal of Public Economics*, 78(1–2):25–49, October 2000. ISSN 0047-2727. doi: 10.1016/S0047-2727(99)00110-3.

Axel Börsch-Supan and Christina B. Wilke. The German public pension system: How it was, how it will be. Technical report, National Bureau of Economic Research, 2004.

Agar Brugiavini and Franco Peracchi. Social Security Wealth and Retirement Decisions in Italy. *LABOUR*, 17(s1):79–114. ISSN 1467-9914. doi: 10.1111/1467-9914.17.specialissue.4.

Gabriele Fischer, Florian Janik, Dana Müller, and Alexandra Schmucker. The IAB Establishment Panel–things users should know. *Schmollers Jahrbuch*, 129(1):133–148, 2009.

Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik Instruments: What, When, Why, and How. Technical report, National Bureau of Economic Research, 2018.

Jonathan Gruber and David A. Wise. *Social Security Programs and Retirement around the World: The Relationship to Youth Employment.* National Bureau of Economic Research Conference Report. University of Chicago Press, Chicago, IL, 2010.

David A. Jaeger, Joakim Ruist, and Jan Stuhler. Shift-share instruments and the impact of immigration. Technical report, National Bureau of Economic Research, 2018.

Simon Jäger. How substitutable are workers? evidence from worker deaths, Working Paper. Job Market Paper.

Alain Jousten, Mathieu Lefebvre, Pierre Pestieau, and Sergio Perelman. The effects of early retirement on youth unemployment: The case of Belgium. 2010.

Wolfram Klosterhuber, Patrick Lehnert, and Stefan Seth. Linked Employer-Employee Data from the IAB: LIAB Cross-sectional Model 2. 2016.

Alicia H. Munnell and April Yanyuan Wu. Are Aging Baby Boomers Squeezing Young Workers out of Jobs? Technical Report 12-18, Center for Retirement Research, Boston College, October 2012.

OECD. Pensionable Age and Life Expectancy, 1950-2050. In *Pensions at a Glance 2011: Retirement-Income Systems in OECD and G20 Countries*. OECD Publishing, 2011.

Jennifer M. Ortman, Victoria A. Velkoff, and Howard Hogan. An Aging Nation: The Older Population in the United States. Technical Report P25-1140, U.S. Department of Commerce. Economics and Statistics Administration. U.S. Census Bureau, May 2014.

Matteo Paradisi and Giulia Bovini. Labor substitutability and the impact of raising the retirement age [job market paper]. Working Paper.

Maria Richter and Ralf K. Himmelreicher. Die Versicherungskontenstichprobe als Datengrundlage für Analysen von Versicherungsbiografien unterschiedlicher Altersjahrgänge. *DRV Schriften*, 79:34–61, 2008.

Tom Walker. Why economists dislike a lump of labor. *Review of Social Economy*, 65(3):279–291, 2007. ISSN 00346764, 14701162. URL `http://www.jstor.org/stable/29770416`.

Appendix

Figure A.1: Administration

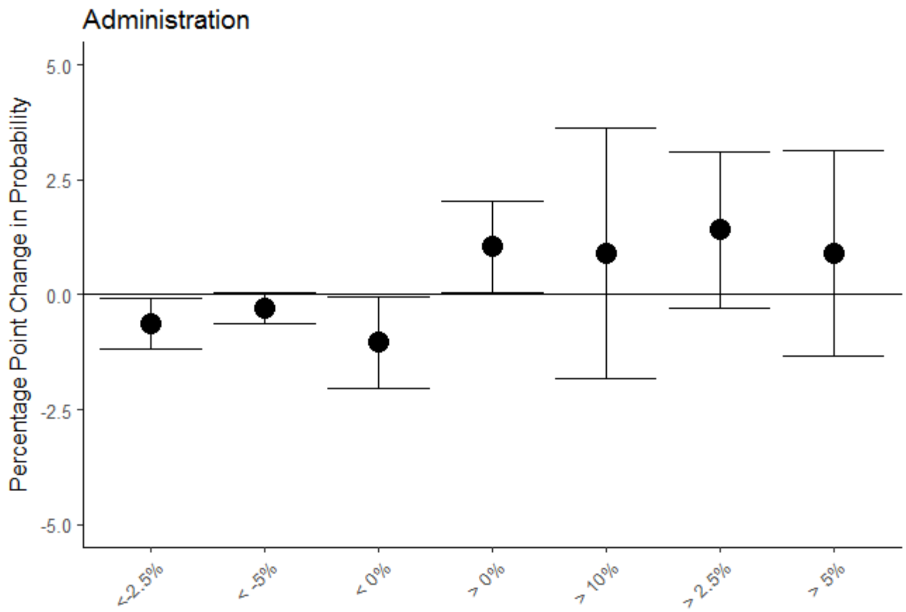

Figure A.2: Agriculture

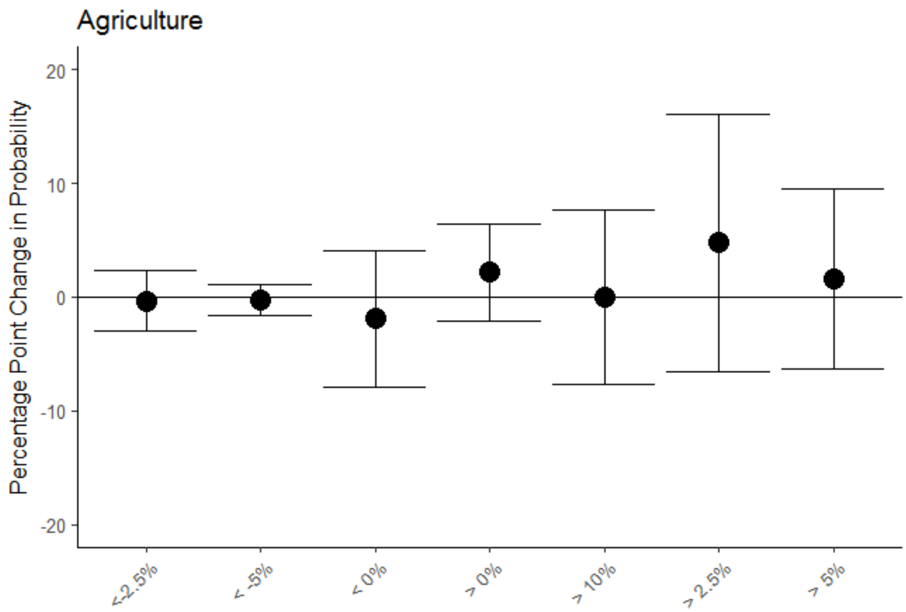

Figure A.3: Construction

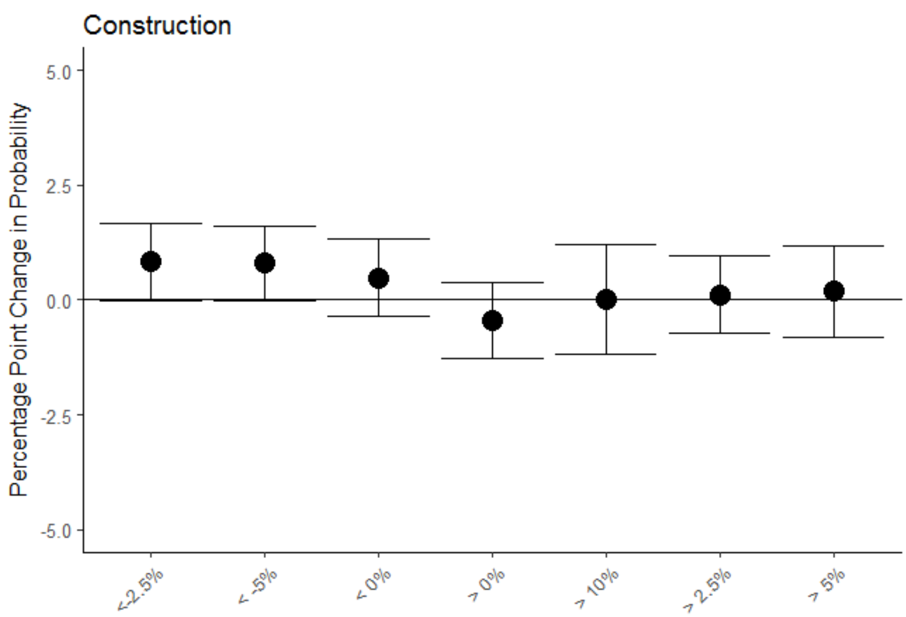

Figure A.4: Finance

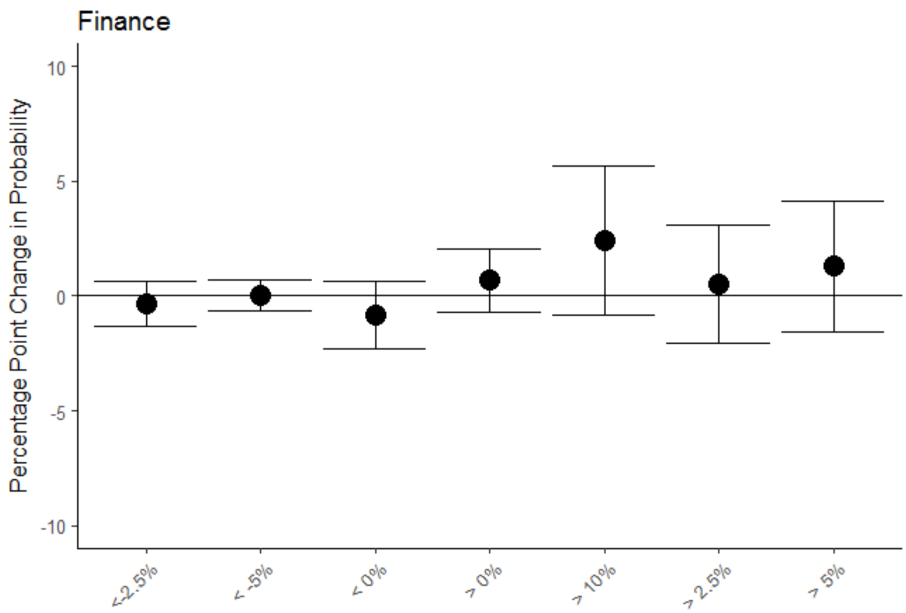

Figure A.5: Mining/Manufacturing

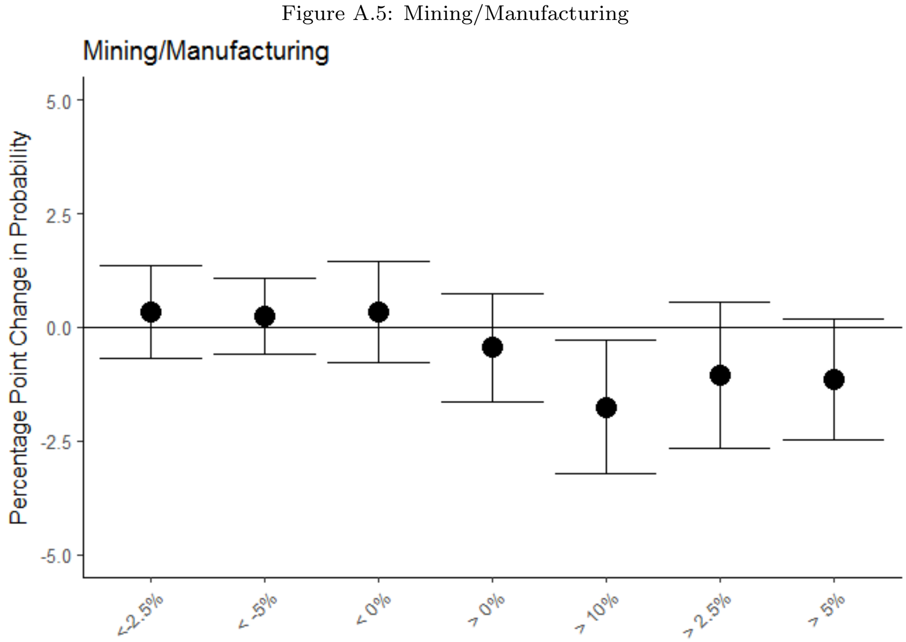

Figure A.6: Pubic Administration

Administration.png

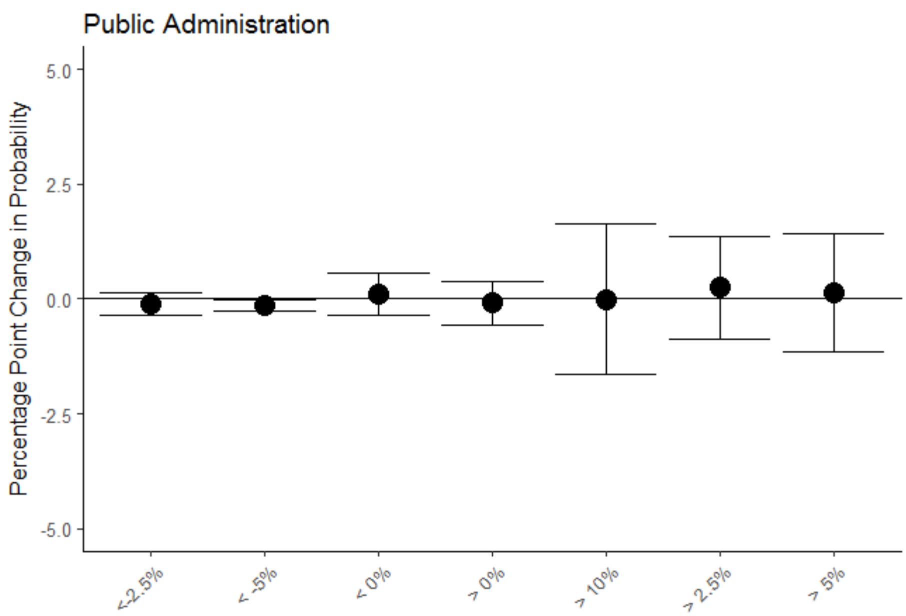

Figure A.7: Real Estate

Estate.png

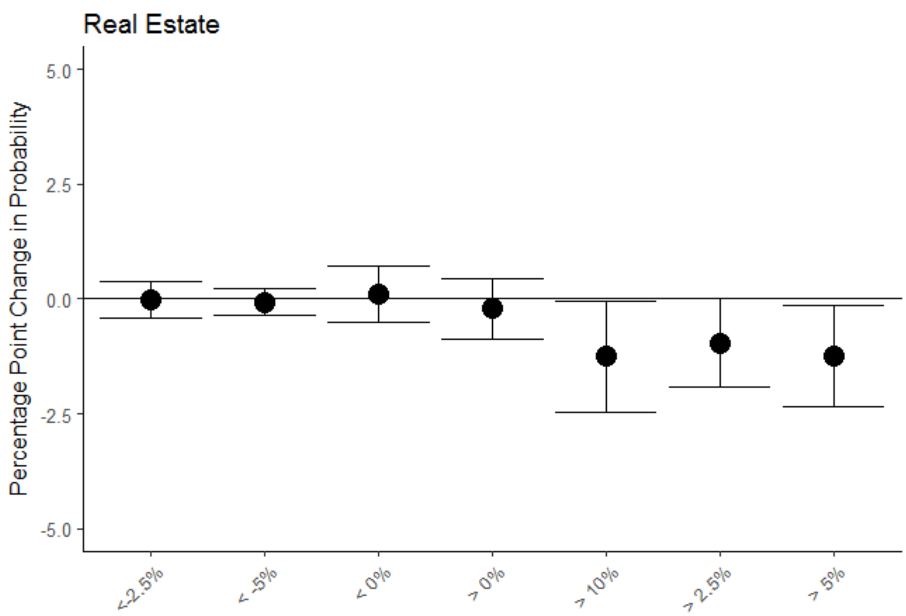

Figure A.8: STEM

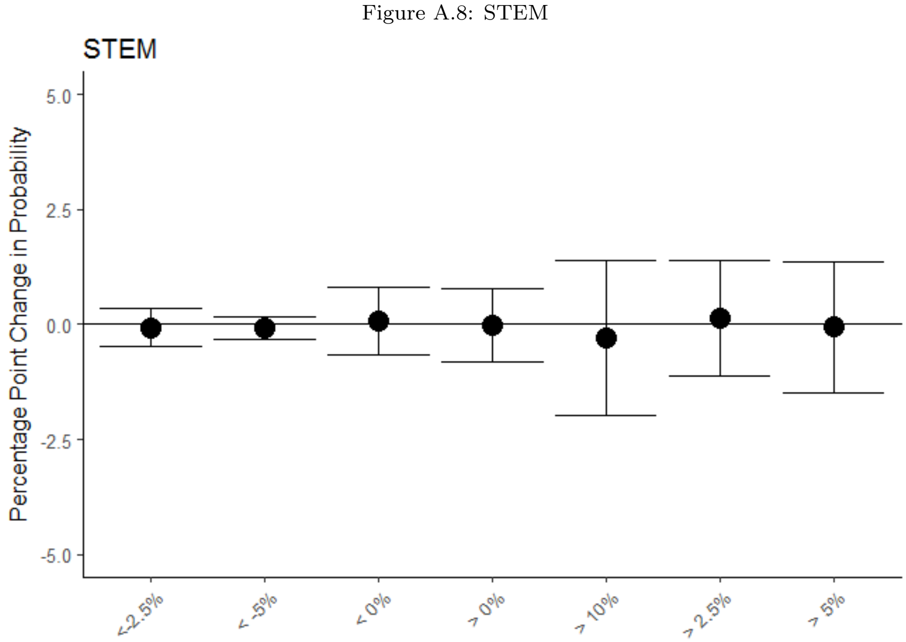

Figure A.9: Trade/Food Service

Service.png

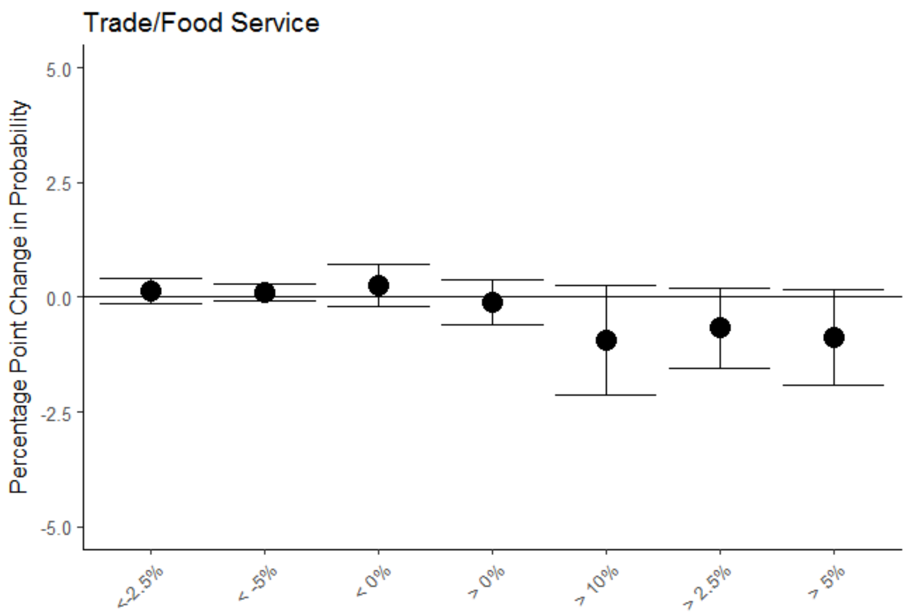

Figure A.10: Transportation

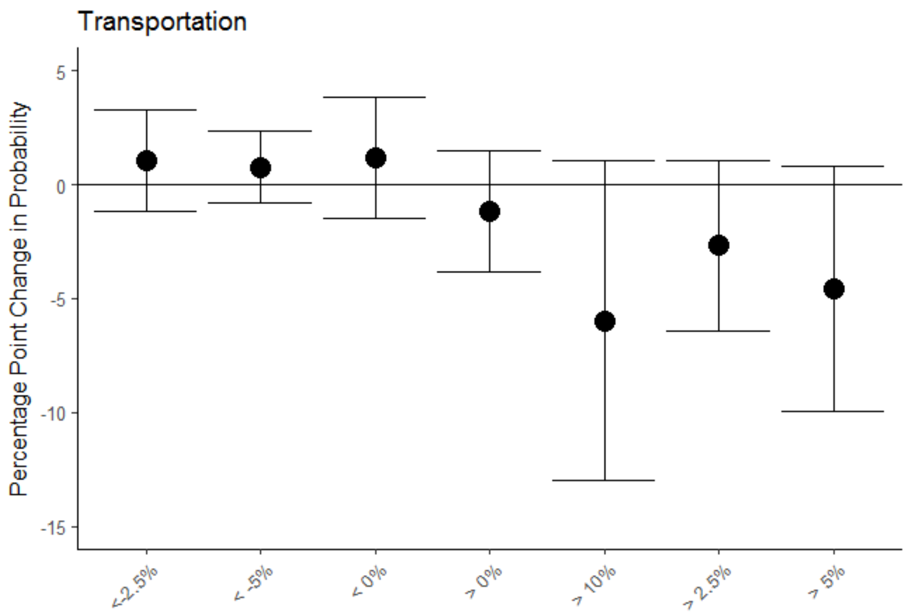

Table A.1: Impact of Older Colleagues on Wages

	Full Sample*						
	$< -5\%$	$< -2.5\%$	$< 0\%$	$> 0\%$	$> 2.5\%$	$> 5\%$	$> 10\%$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$share58pls_{jt}$	0.001	0.001	0.002	-0.003	-0.006	-0.007**	-0.012**
	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)
Mean	0.049	0.085	0.154	0.840	0.650	0.505	0.365
SD	0.216	0.279	0.361	0.366	0.477	0.450	0.481
N	15,027,568						

	Under Age 40*						
	$< -5\%$	$< -2.5\%$	$< 0\%$	$> 0\%$	$> 2.5\%$	$> 5\%$	$> 10\%$
	(8)	(9)	(10)	(11)	(12)	(13)	(14)
$share58pls_{jt}$	0.001	0.001	0.002	-0.002	-0.005	-0.005	-0.008**
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Mean	0.044	0.075	0.135	0.863	0.719	0.579	0.425
SD	0.205	0.264	0.341	0.344	0.449	0.494	0.494
N	6,951,928						

	Over Age 40*						
	$< -5\%$	$< -2.5\%$	$< 0\%$	$> 0\%$	$> 2.5\%$	$> 5\%$	$> 10\%$
	(15)	(16)	(17)	(18)	(19)	(20)	(21)
$share58pls_{jt}$	0.000	0.001	0.002	-0.003	-0.004	-0.006	-0.012**
	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)
Mean	0.054	0.093	0.171	0.827	0.590	0.440	0.313
SD	0.226	0.291	0.376	0.382	0.492	0.496	0.464
N	8,075,640						

* The unit of observation is the establishment-year. The dependent variables are binary. All estimates come from instrumental variable regressions by two-stage least squares. Standard errors, clustered at the establishment level, are in parentheses. Each regression includes a set of establishment characteristics (industry, inflows, outflow, firmsize, state), individual characteristics (education, sex, occupation, experience) and year dummies as controls. The instrumental variable regressions are estimated by two-stage least squares. One star, two stars, and three stars denote statistical significance at the 10-, 5-, and 1-percent confidence levels, respectively.

Table A.2: Mean Establishment to Industry Shares

Ratio	Establishments	Industry
0 to 0.49	0.079	Energy/Water Supply
0.50 to 0.99	0.540	Energy/Water Supply
1.00 to 1.49	0.331	Energy/Water Supply
1.50+	0.050	Energy/Water Supply
0 to 0.49	0.336	Trade/Food Service
0.50 to 0.99	0.325	Trade/Food Service
1.00 to 1.49	0.267	Trade/Food Service
1.50+	0.072	Trade/Food Service