Discussion Paper Series

IZA DP No. 18727

June 2026

Are Senior Workers Overpriced? Evidence from an Age-Differentiated Payroll Tax in Norway

Steinar Holden
University of Oslo

Simen Markussen
Ragnar Frisch Centre for Economic Research
and IZA@LISER

Knut Røed
Ragnar Frisch Centre for Economic Research
and IZA@LISER

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18727 June 2026

Are Senior Workers Overpriced? Evidence from an Age-Differentiated Payroll Tax in Norway*

Abstract

Differentiation of the employer-borne payroll tax may be a tool to raise employment for groups whose wage entitlements are set above market clearing levels – e.g., through collective bargaining, minimum wage legislation or implicit contracts. We provide an empirical evaluation of a reform in Norway in 2002 whereby the payroll tax for mature (62+) workers was reduced by 4 percentage points. Our findings indicate that the reform led to a 2-3% increase in total hours worked by persons aged 62–64. Approximately 25% of the tax cut was passed on to the workers in the form of higher hourly wages.

JEL classification

H22, E24, J23, J26

Keywords

payroll tax, tax incidence, labor demand, labor supply, difference in differences

Corresponding author

Steinar Holden

steinar.holden@econ.uio.no

* The work has received funding from the Norwegian Ministry of Labor. The data that we use are provided by Statistics Norway. They are not publicly available but may be accessed through Statistics Norway with appropriate approvals. The authors declare that there are no conflicts of interest in relation to this paper.

1 Introduction

Making senior workers stay in employment up to a higher age is considered a major priority in virtually all developed countries. Longer working-lives are seen by policy-makers as a key to ensuring fiscal sustainability of pension systems and welfare state institutions. The main tool for making elderly workers postpone retirement has been to implement pension reform, either raising the statutory retirement ages directly and/or improving work incentives by eliminating income testing of pensions and raising the age at which additional pension points can be earned. The focus has thus been on the supply-side of the labor market.

Pension reforms have been successful in raising the retirement age; see, e.g., Baker and Benjamin (1999), Bönke et al. (2018), and Morgavi (2025). In Norway, a radical reform of the private sector early retirement system in 2011, which eliminated all pension income tests, has been shown to increase employment and earnings of affected workers aged 62-66 by approximately 30-40%; see Hernæs et al. (2016) and Andersen et al. (2021).

However, the employment of senior workers may not only be constrained from the supply side. Although there is no overwhelming empirical evidence that individual productivity systematically decreases with age in the relevant age intervals (Hernæs et al., 2023; Mahlberg et al., 2013), the *heterogeneity* in individual productivity is likely to increase at high age (Sharpe, 2011; Skirbekk, 2008), and the combination of employment protection, sticky wages, and high pension costs may discourage employers from hiring elderly workers. It is also possible that many older workers are paid more than their marginal product as a result of implicit contracts (Lazear, 1979).

One way of stimulating the demand for elderly workers is to reduce the payroll tax for employees exceeding a certain age threshold. This was done in Norway in 2002 when the employer-born part of the payroll tax was cut by 4 percentage points for all employees aged 62 or older, from a standard level of 14.1%. The tax cut was implemented on the employer-side because policy-makers considered insufficient demand for elderly workers to be part of the reason why many workers left employment prematurely. The pension system of the early 2000s strongly subsidized early retirement and anecdotal evidence at the time suggested that employers often took advantage of generous early retirement conditions to justify layoff of eligible senior workers during downsizing processes.

In the present paper, we evaluate the effects of the reform on employment and hourly wages.

We apply a flexible triple-difference event-study design, where we examine earnings trajectories in terms of differences relative to the year just before first treatment exposure, using not-yet-treated cohorts as controls. The analysis is built up from separate estimates by age and year of treatment, which are then aggregated into effects that are allowed to vary with the duration of treatment exposure.

Our results suggest that the reform had the intended effect of raising the employment of mature workers. The most reliable effect estimate indicates an increase in the number of hours worked at ages 62-64 about 2-3%. This effect is entirely accounted for by extensive margin responses (the number of days worked), which presumably arise through postponements of retirement. Interpreting the reform as a pure 3.5% reduction in firms' wage costs, the implied labor demand elasticity is between -0.6 and -0.8. However, our results also indicate that a part of the tax cut (approximately 25%) was passed on to directly affected employees in the form of higher wage growth. Exploring heterogeneous effects allows us to go one step further in distinguishing between demand and supply effects. For public sector employees, where strong employment protection makes retirement voluntary, we find that half of the tax reduction was passed through to the worker's wages, indicating that higher employment was realized along an upward-sloping labor supply curve. In contrast, for private sector employees with weaker employment protection, we found no effect on wages, indicating a clean demand response.

The reform has previously been evaluated by Ellingsen and Røed (2006), who found no convincing evidence of any effect of the reform on employment. However, this evaluation was based on data only capturing the first year of the reform and was also based on a different methodology (event history analysis). An evaluation of a similar reform implemented in Hungary in 2014 (Bíró et al., 2022) indicates that a payroll tax cut reducing wage costs for private-sector workers over 55 years of age by 5.3% increased the employment rate by 1.6%, implying an employment elasticity of around -0.3.

An age-differentiated payroll tax has also been used to boost the demand for young workers. In Sweden, a reform in 2007 introduced a 16 percentage point cut in the employer-born payroll tax for workers aged 19-25 (extended to age 26 in 2009). Saez et al. (2019) evaluate this reform and show that the tax cut was fully translated into lower labor costs and generated a presumably demand-driven increase in the employment-population ratio for the affected age group of approximately 1.4%. The implied elasticity of employment with respect to the wage cost was estimated to -0.23, i.e.; similar to the elasticity for older people reported by Bíró et al.

(2022). Our results indicate somewhat larger employment effects.

Our analysis speaks to the more general policy-issue of whether to use a differentiated payroll tax as a tool to offset presumed imperfections in wage setting behavior. In labor markets where wages are largely determined through collective bargaining with an explicit aim of wage compression, as in Norway (Bhuller et al., 2022), such imperfections may be particularly relevant for young/inexperienced workers as well as for low-productivity workers more generally, as the bargained wage for them may exceed expected productivity. The same may be the case for many mature workers, but for reasons more related to implicit contracts and heterogeneous productivity-by-age profiles. Our findings suggest that a differentiated payroll tax indeed has the potential to partially offset such mechanisms.

2 Institutional setting and data

From July 1, 2002, the payroll tax in Norway was cut by 4 percentage points for all employees aged 62 years or more. The standard rate was 14.1% of the total wage bill in the central parts of the country, with some discount in the most rural areas. The payroll tax is levied on employers and is not directly tied to benefits received by their employees. The reform was born as a result of a tripartite "inclusive worklife" cooperation between the state and the associations of employers and employees. The formal proposal was made by the Government on October 12, 2001 and it was approved by the parliament on November 27 the same year. The motivation for the reform was to “stimulate employers to keep and recruit elderly employees” (St.prp. nr. 1, Tillegg nr. 1 (2001-2002)), and if deemed successful, the government signaled willingness to propose further reductions. However, the policy was not deemed successful (Ellingsen & Røed, 2006), and the age-specific payroll tax was abolished from January 1, 2007.

Based on complete administrative register data, Figure 1 illustrates the extent to which different birth cohorts were affected by the reform and offers a first glimpse into its potential employment effects. Each line shows the observed labor-related earnings-by-age path of a particular birth cohort, unconditional on employment, from the 1934-cohort (bottom line) to the 1949-cohort (top line). Earnings are measured in "basic amounts" (G), a unit used to deflate benefits in the Norwegian social security system and intended to capture average wage growth. The profiles shown in the figure thus reflect changes in the degree of employment. The two dashed vertical lines mark the start and end of the payroll tax reduction period, and the red dots indicate which cohorts were affected at which age between 62 and 66. All cohorts born between 1936 and 1943 were affected by the tax cut to some extent, but only the 1940-cohort was fully affected in the sense that it was exposed to a reduced payroll tax from age 62 through 66.

We focus on labor earnings rather than employment in Figure 1 because labor earnings are (third-party) reported without measurement error in Norwegian administrative registers, while employment status or hours worked must either be inferred from earnings data or be based on imperfect employer-employee registers. In the empirical analysis, we use the most reliable elements of the employer-employee registers to disentangle the estimated earnings effect into its sources of the number of days worked, hours worked per day and the hourly wage.

In the most northern part of the country (Nord-Troms and Finnmark), the payroll tax was set to zero already before the reform, and since the tax was not made negative, there was no tax reduction in these regions. Since their population is too small to be meaningfully used as a control population, we drop these observations (2.1% of the observations in our data).

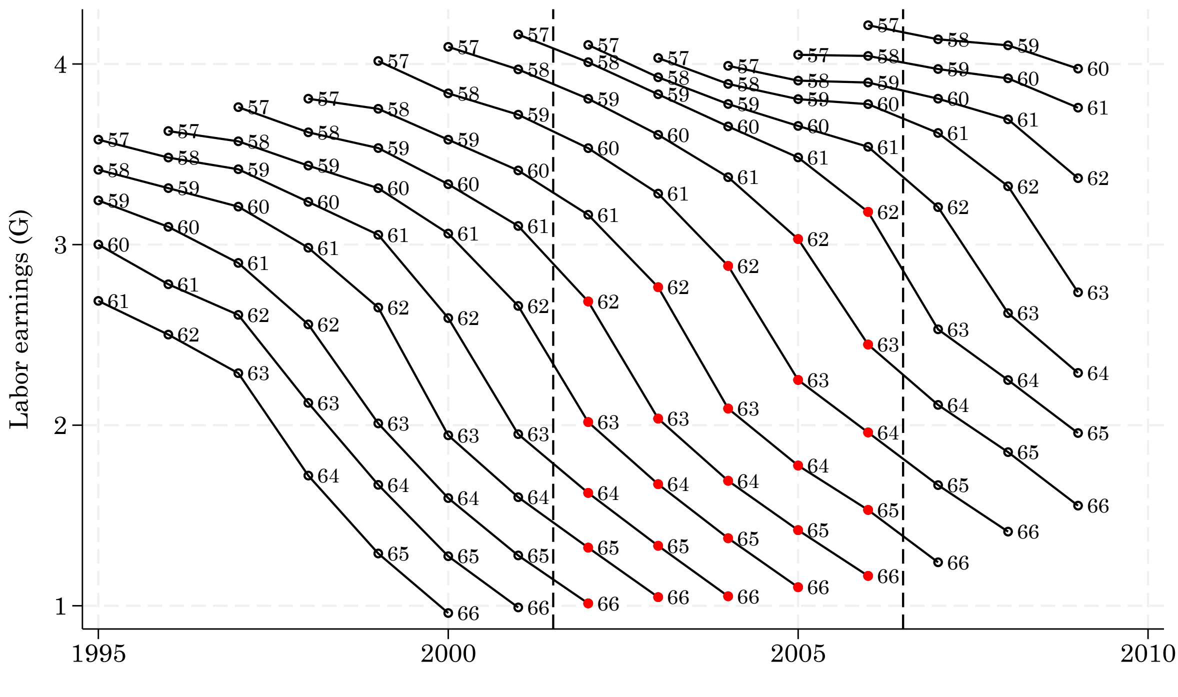

Figure 1: **Average labor earnings. By cohort, age and year**

Note: Labor earnings are taken from tax-records and measured in terms of the wage-growth-adjusted "basic amount" (G), which is used to annually adjust all social security transfers in Norway ($1G \approx$ 17% of average full-time-full-year earnings in Norway). Each line follows a specific birth cohort, from the 1934-cohort (bottom/leftmost line) to the 1949-cohort (top/rightmost line). The sample includes the full population of residents belonging to the respective cohorts without conditioning on employment

There are some important points to note from the earnings profiles shown in Figure 1.

First, real labor earnings decline monotonously with age for all cohorts, and more so as people move into their sixties. Second, age-specific earnings have generally increased over time since 2000, reflecting a trend toward higher labor force participation among persons of mature age. Although earnings of treated cohorts/years increased during the low-tax period, they did so for non-treated as well; hence, there is no noticeable treatment effect to be seen directly in Figure 1.

When considering possible reform effects, it seems reasonable to assume that when the tax cut was introduced in 2002, both employers and employees interpreted it as permanent in the sense that it would likely remain in place until retirement, at least for the first affected cohorts. However, in June 2006, the parties of the "Inclusive workplace" agreement decided to propose a reduction in the tax cut from 4% to 3%. Then, in December that year, the parliament instead decided to completely abolish it already from January 1, 2007.

3 Empirical strategy

There is almost no entry into the labor market at ages above 60 years.[1] An analysis of employment at this high age is therefore in practice an analysis of labor market exit patterns. We thus focus our empirical analysis on effects for persons who were employed four years before first exposure to the treatment and follow their earnings profiles from that point onward.[2]

Our empirical strategy combines the stacked event-study design (Cengiz et al., 2019) with the triple-difference logic formalized by Olden and Møen (2022). We examine earnings paths in terms of differences relative to the year just before each individual's first treatment exposure. It is clear from Figure 1 that we have a number of different combinations of treatment ages a and treatment years t, and we start by estimating separate treatment effects for each age/year combination r. We then aggregate across r to achieve average reform effects by year of exposure. To ensure sufficient power, the age-year-specific effects are estimated for $a \in \{62, 63, 64\}$ and $t \in \{2002, \ldots, 2005\}$, giving 12 r-specific estimates.

To net out the effects of age and calendar time trends/fluctuations, we compare each treated trajectory to three control trajectories, a same-age control (which by construction cannot be observed at the same time), a same-year control (which by construction cannot be observed

[1] For the potential treatment group in our data, the fraction of those who were not employed at age 58 who are still not employed at age 62 is 95%.

[2] We then avoid profiles influenced by the early-retirement reform that ended in 1998. We present main results without conditioning on employment in the Appendix.

at the same age) and a same-time control for the same-age control (control-for-controls). We select these control groups such that we minimize the age and time differences relative to the treatment group.

As noted by Olden and Møen (2022), a triple difference estimator can be computed as the difference between two difference-in-differences estimators, facilitating a more transparent exposition of its underlying identifying assumptions. To illustrate, let $\Delta y_a^b(c)$ indicate the change in average earnings from age a to age b for people who belong to birth cohort c, and consider the effect of treatment at age 62 in 2002 (the first treatment year). The DDD estimator can be written as follows:

$$DDD(62, 2002) = \Delta y_{61}^{62}(1940) - \Delta y_{61}^{62}(1939) - [\Delta y_{60}^{61}(1941) - \Delta y_{60}^{61}(1940)]$$

The first terms ($\Delta y_{61}^{62}(1940) - \Delta y_{61}^{62}(1939)$) give the change in average annual earnings from age 61 to age 62 for the treated 1940-cohort minus the corresponding change for the not-yet-treated 1939-cohort. This could be a valid difference-in-differences estimator, provided that there is no trend in the earnings change from age 61 to 62. But if there is a trend, the estimator is invalid. The second term of the DDD-estimator seeks to correct for that by subtracting the corresponding differences for earnings changes from age 60 to age 61 observed in the same time period ($\Delta y_{60}^{61}(1941) - \Delta y_{60}^{61}(1940)$). The DDD-estimator for the treatment effect at age 62 in 2002 is thus valid if the underlying 2002-trend in earnings *changes* from age 61 to age 62 is equal to the trend in earnings *changes* from age 60 to age 61. In this example, it is notable that the 1940-cohort serves a dual role, both as treated and as control-for-control.

When identifying effects in the first reform year (2002), the estimator involves differences between first-differences. However, when we move longer into the treatment period, comparisons must be made with larger time/age/cohort lags. For example, the effect for persons aged 64 in 2005 can be computed as:

$$DDD(64, 2005) = \Delta y_{61}^{64}(1941) - \Delta y_{61}^{64}(1937) - [\Delta y_{58}^{61}(1944) - \Delta y_{58}^{61}(1940)]$$

To embed the triple difference identification strategy within an event study design, we provide effect estimates for the entire earnings trajectory leading up to each treatment age/year in question, starting four years before the first treatment. Relative time runs from $j = -3$ (where the employment condition is imposed) to $j = J$ (the final year of the trajectory in question), with $j = 0$ constituting the reference year – the last year before treatment. Let $\mathbb{1}[\text{Tage}_i]$ and $\mathbb{1}[\text{Tyear}_i]$ denote indicators that a trajectory i extends to a treatment age or a treatment year, respectively. For each combination (r) and for each relative time $j \in \{-3, \ldots, J\}$, we estimate the following equation with OLS:

$$y_{ij} = \alpha_i + \beta_{1j} + \beta_{2j}\,\mathbb{1}[\text{Tage}_i] + \beta_{3j}\,\mathbb{1}[\text{Tyear}_i] + \beta_{4j}\,\mathbb{1}[\text{Tage}_i] \times \mathbb{1}[\text{Tyear}_i] + \varepsilon_{ij}, \tag{1}$$

where y_{ij} is the earnings outcome for trajectory i in relative time j.

The triple-difference logic is built into Equation 1. Trajectory-fixed effects remove level differences across trajectories and the controls for $Tage$ and $Tyear$ account for the fact that trajectories observed at the same relative time j differ in terms of age and calendar year. What remains is the treatment effects as captured by β_{4j}, which for $j = J$ is equal to the DDD estimator described above for the treatment age/year combination (r) in question.

To pool across these combinations we stack the r-specific trajectories and impose the assumption that the treatment effects depend on relative time only, i.e., the first, second, or third year with lower payroll tax for each worker, and not on the specific age or year of treatment:

$$y_{ijr} = \alpha_{ir} + \sum_r \mathbb{1}[r]\left[\beta_{1jr} + \beta_{2jr}\,\mathbb{1}[\text{Tage}_i] + \beta_{3jr}\,\mathbb{1}[\text{Tyear}_i]\right] + \beta_{4j}\,\mathbb{1}[\text{Tage}_i] \times \mathbb{1}[\text{Tyear}_i] + \varepsilon_{ijr}, \tag{2}$$

with $j \in \{-3, \ldots, J\}$. All control-group parameters β_{1jr}, β_{2jr}, β_{3jr} remain stack-specific; only the treatment-effect coefficients β_{4j} are restricted to be common across r. The trajectory-fixed effects α_{ir} are estimated separately for each stack, since the same trajectory may appear in different r-stacks.

Our choice of treatment ages/years and the selection of the closest-in-age/year control cohorts implies that our analysis sample comprises people born from 1937 through 1944. Appendix Table A1 provides an overview of the data structure and some descriptive statistics.

Although it is difficult to rule out distortive trends in the data, the event-time setup makes it possible to assess their presence over a pre-treatment period, and we can to some extent evaluate their potential influence by choosing alternative control cohorts. A potential weakness with our approach is that the choice of closest-in-age same-year controls implies that a few specific cohorts are re-used several times. In particular, the 1940 cohort serves as the same-year control for the same-age control for all estimates. This makes our estimates vulnerable to any circumstances that are specifically relevant for this cohort. Although we are not aware of any such circumstances, we report in the appendix results built on an alternative strategy, where we choose different controls for each age-time combination (with a three-year difference).

In principle, it is also possible to explore the effects of the abolishment of the reform in January 2007, treating it as a new reform with a tax increase. However, in this case it would be difficult to find appropriate comparison cohorts, since the earlier cohorts were treated to different degrees. Attempts to estimate the effects of the abolishment in a similar way have thus failed to exhibit satisfactory pre-trend properties. Moreover, the post-treatment period encompasses a rather unstable labor market environment during the financial crisis. Thus, we only present results based on the introduction of the reform.

4 Results

We start this section by presenting the estimation results obtained when we use annual labor earnings as the outcome. We move on to a decomposition exercise, where we first decompose the effect on earnings into hours worked and hourly wages, and then decompose the effect on hours worked into the number of days worked (referred to as the extensive margin) and the number of hours worked per day (intensive margin). We then evaluate the fiscal consequences of the tax cut and estimate the degree of "self-financing" caused by behavioral responses. Finally, we provide some robustness checks and explore the presence of heterogeneous responses along the dimensions of earnings/wage level, sector (private or public), gender, and initial work hours (part-time or full-time).

4.1 Effects on labor earnings

Effects on earnings are typically estimated using log-earnings as the outcome variable. In our case, we cannot do that, as a large fraction of the earnings-observations are equal to zero. However, it is clear from Figure 1 that average earnings levels decline rapidly with age, implying that a given level-effect will constitute a larger relative response the higher the age of measurement. To resolve this dilemma, we have chosen to estimate the individual effects in levels and then scale them to the age-specific observed average earnings level net of the estimated treatment effect for the treatment group in question. The calculation of confidence intervals is based on a bootstrap procedure clustered at individuals (120 trials).

To ensure transparency, Figure 2 first shows the estimated effects separately for each combination of age and calendar year, based on Equation 1. Estimates reported to the left of the vertical lines (in black) refer to pre-treatment years (included to facilitate an assessment of pre-trends), and estimates to the right (in red) refer to treatment effects by years of exposure. Hence, year 1 on the horizontal axis refers to the treatment effect in the first year of treatment, and year 2 to the effect in the second year conditional on being exposed in year 1. Note that the confidence intervals become much wider as we move to the right in each panel. One reason is that the effects are reported relative to age-specific averages that decline rapidly with age. In appendix Figure A1, we show results for the directly estimated level effects. The confidence intervals are then more stable, although they increase somewhat with distance from the reference year ($j = 0$) in both directions.

The top panels of Figure 2 report estimates measured at age 62, where we see little sign of any causal effects. Note, however, that earnings are measured at the calendar-year level, whereas the entitlement to reduced payroll tax is linked to the exact timing of reaching the age 62 threshold. Hence, in the year of the 62nd birthday, the average person would only have been 50% treated.[3] When we estimate separately for persons born in the first and last quarter, we find borderline-significant positive treatment effects for the former and zero effects for the latter (not shown). Moving on to rows two and three (ages 63 and 64), there are some indications of positive causal effects, particularly after 2-3 years of exposure. The estimated pre-treatment effects are in most cases close to zero, though some of them are (borderline) statistically significant. While we cannot claim the complete absence of spurious effect estimates, we see no signs of positive pre-trends that could plausibly account for the positive treatment effect estimates. We return to some robustness analyzes in later sections.

Note that in some cases, the estimation procedure yields more than one estimate for the same effect, based on different control groups. For example, we have two different estimates for the effect at age 63 in 2003, both the one to the right in the panel *"63,2003"* and the effect reported for $j = 2$ in the panel *"64,2004"*. As the former is based on a same-age control group that is one year closer in time (2001) than the latter (2000), it is natural to think of the former as the preferred estimator in this case.

Our next step is to aggregate all the estimates in Figure 2 into a single set using Equation (2), assuming homogeneous effects across cohorts and years, while allowing the effects to vary with exposure time; i.e., year 1, 2 and 3 with reduced payroll tax. The resulting estimates are

[3]Those reaching 62 in 2002 would on average only be 37.5% treated, while older workers would be 50% treated since the reform was implemented from July 1 that year.

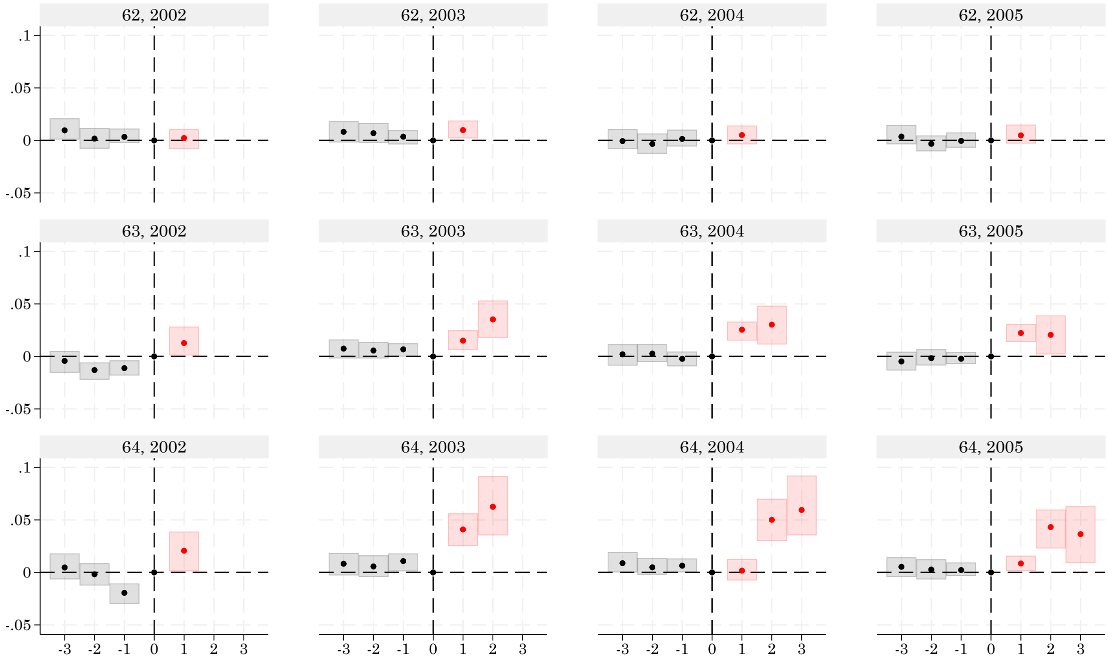

Figure 2: **Single estimators, labor income, relative responses**

Note: The graph displays the estimated dynamic treatment effects with 95% bootstrap-based confidence intervals (120 trials). The age and year entered in the panel labels refer to the rightmost estimate in the respective panels. The effects are first estimated in levels and the estimates are thereafter divided by observed earnings for the relevant treated age group net of the estimated effect. The time axis is centered on the last year prior to treatment (year 0) which is also the year excluded from the regressions. The estimation sample includes all treatment groups aged 62-64 in years 2002-2005. The same-year controls are aged 61 at the time of treatment. The same-age controls are the closest not-yet treated cohort of the same age.

shown in Figure 3, panel (a). They indicate that the 4 percentage point reduction in the payroll tax caused labor earnings to increase by approximately 1.5% in the first year of exposure and by 3-4% in the two subsequent years.

4.2 Effects on hours worked and on hourly wages

The effect of the reform on annual earnings may consist of two parts: an effect on hours worked and an effect on payment per hour. As explained in Section 2, while annual earnings are accurately measured in our data, hours worked and hourly wages are not. However, if we can find a way to estimate the effect on one of these margins, we can identify both by exploiting the following relationship:

$$\frac{\Delta Earnings}{Earnings} \approx \frac{\Delta Wage}{Wage} + \frac{\Delta Hours}{Hours}.$$

Since we already have estimates for the reform's influence on relative earnings (the left-hand side of the equation), it is sufficient to estimate its effects on either wages or on hours worked

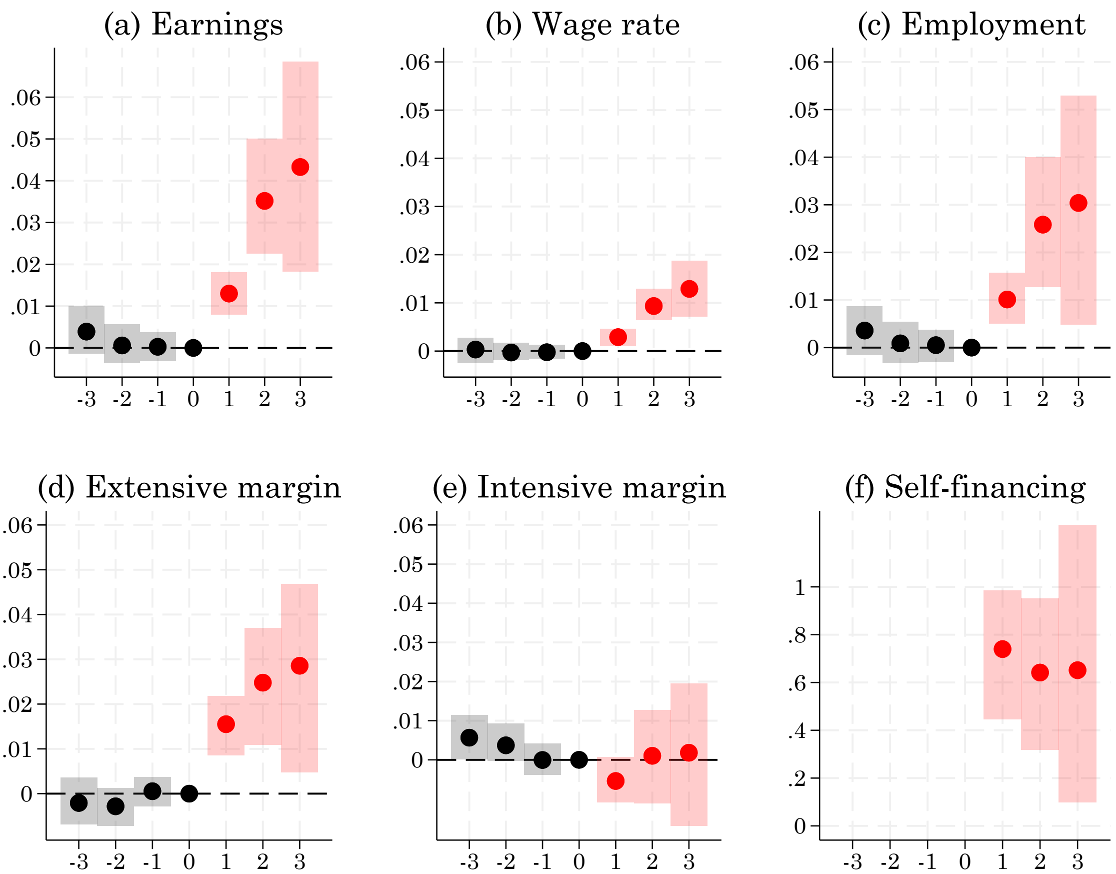

Figure 3: **Joint estimators, relative responses**

Note: The graph displays the estimated dynamic treatment effects with 95% bootstrap-based confidence intervals (120 trials). The effects are first estimated in levels and thereafter divided by observed earnings for the relevant treated age group net of the estimated effect. The time axis is centered on the last year prior to treatment (year 0) which is also the year excluded from the regressions. The coefficients reported in panels (a), (b), (d), and (f) are estimated directly, whereas the coefficients in panels (c) and (e) are inferred from other estimates. The estimation sample includes all treatment groups aged 62-64 in years 2002-2005 and their respective controls (305 734 unique individuals and 1 300 536 earnings trajectories), except in panel (b) where the estimation sample only includes persons with full-time job in both the base-year and the outcome-year (196 197 unique individuals and 807 397 earnings trajectories). The same-year controls are aged 61 at the time of treatment. The same-age controls are the closest not-yet treated cohort of the same age.

to ensure a full decomposition. Given that we do not have reliable data on hours worked except for full-time workers, we have chosen to follow the first route, but present results based on the alternative strategy in the appendix. We thus create an auxiliary subpopulation consisting of persons in our data who are full-time employed in all relevant years, defined as having an employment spell directly coded in the register as full-time.[4] In this case, we follow the standard practice of specifying the outcome as log(earnings) to obtain effects measured in relative terms. Provided that the scope for hours-variation is limited for full-time employees and that the effect

[4]For consistency, we also condition on earnings exceeding 3 Basic amounts (corresponding roughly to 50% of the average full-time-full-year earnings level in Norway).

on their wages is representative for all workers, this analysis will identify the reform's effect on hourly wages.

The estimated wage effects are displayed in panel (b) of Figure 3. The results indicate that the reform caused the wage rate of treated workers to increase by less than half a percent in the first treatment year, by approximately 1% in the second year, and then by a bit more than 1% in the third year. Hence, it appears that around 25% of the payroll tax cut was passed on to the workers. Panel (c) then shows the implication of these estimates for the total employment effect. Exploiting that the relative change in hours worked (including both extensive and intensive margins) approximately equals the difference between the relative change in earnings and the relative change in the wage rate, we estimate that the reduced payroll tax increased employment in the treated group by 1% the first year and by 2-3% in the two subsequent years. Since the reform amounted to a 3.5% reduction in wage costs, the implied elasticity of a 2-3% increase in employment is between -0.6 and -0.8.

4.3 Extensive and intensive margins

The effect on hours worked may arise from extensive or intensive margin responses. In the present case, it is natural to think of the extensive margin response as a postponement of labor market exit, whereas the intensive margin captures the number of hours worked per day. To decompose the estimated hours effects into these two margins, we exploit data from the Norwegian employer-employee register with relatively precise records on the number of days worked, but hours per day only coded in broad categories. We thus perform the decomposition exercise by using the number of days in employment as the outcome of interest and infer the effects on hours worked per day by exploiting the following approximation:

$$\frac{\Delta Hours}{Hours} \approx \frac{\Delta \text{Days in employment}}{\text{Days in employment}} + \frac{\Delta \text{Hours per day worked}}{\text{Hours per day worked}}$$

The results are shown in Figure 3, panels (d) and (e). There is no effect on hours worked per day, suggesting that the effect on employment is fully accounted for by the extensive margin (the number of days). A natural interpretation of this result is that the effect on employment arises exclusively from postponed retirement. It is worth highlighting that although the extensive margin responses reported in panel (d) and the total employment responses reported in panel (c) are very similar, they are based on completely different data sources. Whereas the overall employment effects (panel c) are based on annual third-party-reported earnings data, combined with information about full-time employees, the extensive margin effects (panel d) are based on the start and stopping dates for employment spells reported by employers in the administrative employer-employee register.

4.4 Fiscal consequences

Given that the payroll tax cut for senior workers had positive effects on both employment and wage levels, the reform contained some self-financing elements. To assess how much of the direct costs that are recouped through behavioral responses, we estimate Equation 2 with the following individual (annual) outcome: total payroll tax paid + total individual tax paid − social insurance transfers and early retirement benefits received. The latter is included despite the fact that it is only partly financed by the state (approximately 30% of the costs), with the rest coming from a fund financed by employers through a fixed "AFP-tax" levied on earnings. We compare the reform's estimated effect on this outcome with the effect it would have had without a behavioral response (the "mechanical" effect), which consists of the reduced payroll tax only. We finally calculate the estimated degree of self-financing (SF) as

$$SF = 1 - \frac{\text{Estimated effect}}{\text{Mechanical effect}}.$$

The results are shown in panel (f) of Figure 3. Our estimates indicate that approximately 60% of the direct fiscal loss is recouped through behavioral responses and their impacts on taxes and transfers. The fiscal consequences become even more favorable if one takes into account that the payroll tax for public sector employees is a transfer within the public budget. Moreover, previous evidence from Norway suggests considerable positive spillover effects within couples due to joint retirement (Kruse, 2021). On the other hand, there may also be negative spillover effects at the workplaces, such that the extra hours worked by the treated workers substitute for hours worked by non-treated (younger) workers.

4.5 Heterogeneity and robustness

Table 1 reports the estimated treatment effects for a number of different groups, defined on the basis of employment status four years before treatment exposure. This exercise serves the purpose of examining heterogeneity as well as robustness. We construct groups along the dimensions of initial work-hours (full-time/part-time), annual earnings level, hourly wage, sector of employment (private/public), access to subsidized early retirement, and gender.

We focus here on the estimated treatment effects in the second year after treatment, as this is the first year that likely captures the full effect of the treatment and, in contrast to later years, is also estimated with a degree of precision that facilitates analysis of smaller groups. We also concentrate on the three main outcomes of total earnings, hours worked, and hourly wage.

Table 1: Estimated effects in second treatment year ($j = 2$). By subgroup

	Labor earnings			Wage rate (FT)			Employment	
	N	Estimate	[95% CI]	N	Estimate	[95% CI]	Estimate	[95% CI]
All	148,506	0.035	[0.028, 0.042]	35,040	0.009	[0.007, 0.011]	0.026	[0.018, 0.033]
Working time								
Full-time	66,453	0.032	[0.022, 0.041]	20,958	0.008	[0.006, 0.011]	0.024	[0.014, 0.033]
Part-time	33,216	0.034	[0.013, 0.054]	2,418	0.008	[0.002, 0.016]	0.026	[0.003, 0.045]
Earnings level								
High earnings	72,954	0.038	[0.026, 0.048]	26,202	0.012	[0.010, 0.015]	0.026	[0.015, 0.036]
Low earnings	75,551	0.037	[0.025, 0.046]	8,838	0.001	[−0.003, 0.003]	0.036	[0.026, 0.046]
Wage level								
High wage	29,125	0.024	[0.012, 0.036]	11,933	0.011	[0.007, 0.014]	0.013	[0.002, 0.026]
Low wage	30,005	0.047	[0.032, 0.063]	8,612	0.006	[0.003, 0.009]	0.041	[0.026, 0.059]
Gender								
Men	77,292	0.049	[0.037, 0.060]	22,873	0.010	[0.006, 0.013]	0.039	[0.029, 0.050]
Women	71,214	0.010	[0.002, 0.020]	12,167	0.009	[0.007, 0.012]	0.001	[−0.007, 0.010]
Sector and gender								
Private	56,326	0.027	[0.010, 0.037]	14,919	−0.002	[−0.005, 0.002]	0.029	[0.014, 0.039]
Men	38,456	0.042	[0.022, 0.054]	11,834	−0.001	[−0.005, 0.003]	0.042	[0.021, 0.054]
Women	17,870	−0.011	[−0.031, 0.010]	3,085	−0.006	[−0.015, 0.000]	−0.006	[−0.025, 0.018]
Public	60,306	0.041	[0.034, 0.051]	14,472	0.022	[0.019, 0.024]	0.019	[0.011, 0.028]
Men	22,088	0.060	[0.049, 0.076]	7,424	0.026	[0.022, 0.030]	0.035	[0.024, 0.050]
Women	38,218	0.019	[0.009, 0.028]	7,048	0.018	[0.014, 0.021]	0.001	[−0.009, 0.010]
AFP eligibility (private sector)								
AFP	36,084	0.021	[0.002, 0.038]	8,039	−0.003	[−0.008, 0.001]	0.024	[0.006, 0.040]
No AFP	20,242	0.013	[−0.007, 0.032]	6,880	0.001	[−0.007, 0.006]	0.012	[−0.007, 0.031]

Note: Point estimates with 95% bootstrap-based confidence intervals in brackets (120 trials). High and low earnings levels and high and low wage levels are defined as levels above or below the median observation. AFP eligibility refers to eligibility to an early retirement option from the age of 62, with the annual pension calculated *as if* employment had continued until the statutory retirement age at 67. Observation numbers refer to the number of treated in $j = 2$. The numbers reported for subgroups do not always add to the total, due to missing information. For observations before 1996, we do not have data facilitating full-time/part-time distinction and information on the private/public distinction and the AFP/No-AFP distinction is missing for some observations. As employment effects are derived from the estimates on labor earnings and the wage rate using different samples, we do not report separate observation numbers for these estimates.

Overall, the group-specific estimates are similar, but there is a tendency for estimated employment effects to be larger the smaller are the estimated wage effects (and vice versa). The main takeaways from Table 1 can be summarized as follows:

- The effects on wages were smaller, and the effects on employment were larger for employees with low earnings/wages, indicating more initial excess supply ("overpricing") among the low-skilled.
- The effects were roughly the same for full-time and part-time workers, consistent with an extensive margin channel for the employment effects.
- The reform increased the wage level for men and women similarly, but the effects on employment were present only for men.
- In the private sector, the reform had no impact on wages for either men or women. It increased employment, but only for men. The employment effect was larger when an early retirement option (AFP) was available. This is consistent with the notion that access to AFP facilitated layoffs of senior workers if other conditions for layoff were met, giving more scope for a positive demand effect from a cut in payroll taxes.
- In the public sector, the reform raised wages for both men and women, indicating that the strong employment protection in this sector meant that higher employment could only be realized along an upward-sloping labor supply curve. Only men responded by increasing their labor supply.

In the appendix, we report results from two additional robustness exercises related to the sampling of data and the composition of control groups, and to the strategy for decomposition of earnings effects into employment and hourly wages. Including complete birth cohorts in our data (dropping the condition of being employed at $j = -3$) leads to lower, but still positive and statistically significant estimates of effects on earnings and employment. Selecting same-year controls three years younger (instead of the closest in age) to avoid repeated use of the same cohort changes almost nothing.

In a final exercise, we reverse the way in which we decompose the earnings effect into employment and wage responses; that is, instead of estimating wage-effects and then inferring employment responses, we estimate employment-effects directly (admittedly based on noisy data from the employer-employee register) and then infer the wage responses. The results remain similar to those already reported.

5 Conclusion

In 2002, the Norwegian government implemented a payroll tax cut of 4 percentage points for workers aged 62 or more, with the explicit aim of increasing the demand for senior labor. The implicit presumption was that senior workers in many cases were overpriced relative to expected productivity. In the present paper, we have provided evidence that the policy had the intended effect of raising employment among affected workers by 2-3%, and did so primarily by delaying exit from the labor market. Most of the effect occurred through the realization of excess supply; that is, without a corresponding increase in the take-home wages for workers. We can thus infer that there were some senior workers, particularly in the private sector, whose employment options were constrained from the demand side and that this constraint was alleviated by the reduced payroll tax. There were apparently also some responses along the labor supply curve, but these occurred exclusively in the public sector, where employment protection was particularly robust. Our estimates indicate that approximately 50% of the tax cut was passed on to public sector workers in the form of higher wage growth.

Although the results presented in this paper stem from a Norwegian reform implemented between 2002 and 2007, they are arguably relevant for policymaking in many countries today. More recent pension reforms have improved work incentives and potentially driven the labor supply for senior workers closer to its institutional maximum. The underlying mechanisms that give rise to demand-side constraints — collectively bargained wages, seniority pay, and the Lazear-type wedge between wage and marginal productivity — have probably changed less. In many countries, a differentiated payroll tax might therefore be a more potent policy instrument today than it was in Norway during the period we evaluate.

More broadly, our findings suggest that differentiation of the employer-born payroll tax can serve as an effective instrument to mitigate the consequences of imperfections in wage setting, where certain groups of workers receive wages that exceed their expected productivity.

References

Andersen, A. G., Markussen, S., & Røed, K. (2021). Pension reform and the efficiency-equity trade-off: Impacts of removing an early retirement subsidy. *Labour Economics*, *72*, 102050.

Baker, M., & Benjamin, D. (1999). How do retirement tests affect the labour supply of older men? *Journal of Public Economics*, *71*(1), 27–51.

Bhuller, M., Moene, K. O., Mogstad, M., & Vestad, O. L. (2022). Facts and fantasies about wage setting and collective bargaining. *Journal of Economic Perspectives*, *36*(4), 29–52.

Bíró, A., Lindner, A., Prinz, D., Branyiczki, R., & Márk, L. (2022). *Firm heterogeneity and the impact of payroll taxes* (tech. rep.). KRTK-KTI Working Papers.

Bönke, T., Kemptner, D., & Lüthen, H. (2018). Effectiveness of early retirement disincentives: Individual welfare, distributional and fiscal implications. *Labour economics*, *51*, 25–37.

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, *134*(3), 1405–1454.

Ellingsen, G., & Røed, K. (2006). Analyse av aldersdifferensiert arbeidsgiveravgift. *Frisch Centre Rapport*, *5*, 2006.

Hernæs, E., Kornstad, T., Markussen, S., & Røed, K. (2023). Ageing and labor productivity. *Labour Economics*, *82*, 102347.

Hernæs, E., Markussen, S., Piggott, J., & Røed, K. (2016). Pension reform and labor supply. *Journal of Public Economics*, *142*, 39–55.

Kruse, H. (2021). Joint retirement in couples: Evidence of complementarity in leisure. *The Scandinavian journal of economics*, *123*(3), 995–1024.

Lazear, E. P. (1979). Why is there mandatory retirement? *Journal of political economy*, *87*(6), 1261–1284.

Mahlberg, B., Freund, I., Cuaresma, J. C., & Prskawetz, A. (2013). Ageing, productivity and wages in Austria. *Labour economics*, *22*, 5–15.

Morgavi, H. (2025). Is it worth raising the normal retirement age? A new model to estimate the employment effects. *Public Sector Economics*, *49*(3), 339–367.

Olden, A., & Møen, J. (2022). The triple difference estimator. *The Econometrics Journal*, *25*(3), 531–553.

Saez, E., Schoefer, B., & Seim, D. (2019). Payroll taxes, firm behavior, and rent sharing: Evidence from a young workers' tax cut in Sweden. *American Economic Review*, *109*(5), 1717–1763.

Sharpe, A. (2011). Is ageing a drag on productivity growth? A review article on ageing, health and productivity: The economics of increased life expectancy. *International Productivity Monitor*, (21), 82.

Skirbekk, V. (2008). Age and productivity capacity: Descriptions, causes and policy options. *Ageing horizons*, *8*, 4–12.

Appendix

Table A1: Descriptive statistics by birth cohort

	Birth cohort							
	1937	1938	1939	1940	1941	1942	1943	1944
Panel A: By birth cohort								
Cohort population	35,670	37,191	38,539	39,567	37,965	43,189	46,905	52,747
Share female	0.510	0.509	0.507	0.501	0.504	0.500	0.498	0.493
Employment rate, age 57	0.635	0.649	0.657	0.668	0.676	0.689	0.701	0.703
Employment rate, age 61	0.560	0.571	0.574	0.577	0.587	0.597	0.606	0.610
Employment rate, age 63	0.403	0.394	0.400	0.404	0.412	0.428	0.449	0.459
Employed at 63, given employed at 61	0.691	0.663	0.667	0.675	0.679	0.692	0.711	0.725
Panel B: Analysis sample, measured at $j = 0$								
Persons	21,842	23,269	24,107	27,071	25,203	29,926	33,595	38,039
Share female	0.472	0.479	0.479	0.483	0.483	0.482	0.482	0.477
Employed at $j = 0$	0.608	0.618	0.653	0.695	0.694	0.699	0.706	0.730
Share private sector	0.282	0.290	0.308	0.320	0.335	0.334	0.344	0.361
Share public sector	0.322	0.323	0.344	0.349	0.340	0.344	0.337	0.340
Share full-time	0.456	0.468	0.499	0.556	0.551	0.556	0.566	0.587
Mean labor earnings (G)	4.03	4.10	4.35	4.65	4.66	4.77	4.84	5.03
Log hourly wage	1.867	1.868	1.869	1.855	1.859	1.869	1.875	1.872
Panel C: Stacked observations								
Observations	520,405	623,050	726,224	2,176,992	806,679	797,233	727,332	656,017
Observations per person	23.8	26.8	30.1	80.4	32.0	26.6	21.7	17.2
Share treated	–	0.171	0.349	0.204	0.540	0.383	0.209	–
Share same-age control	1.000	0.829	0.651	–	–	–	–	–
Share same-year control	–	–	–	–	0.460	0.617	0.791	1.000
Share control-for-controls	–	–	–	0.796	–	–	–	–

Note: Sample is the Norwegian resident population born 1937–1944, excluding the zero-payroll-tax zone (Nord-Troms and Finnmark). These are the birth cohorts used in the main part of our empirical analysis. *Panel A* reports population-level statistics for each birth cohort; employment rate is the share with positive labor earnings. *Panel B* reports statistics for the analysis sample (persons employed at relative time $j = -3$), with all characteristics measured at $j = 0$. Mean labor earnings are measured in basic amounts (G). *Panel C* reports the stacked regression dataset, in which each person enters under several (age, year, control-role) combinations. Each observation is one of four roles defined by Equation (1): treated ($T_{age} = T_{year} = 1$), same-age control ($T_{age} = 1$, $T_{year} = 0$), same-year control ($T_{age} = 0$, $T_{year} = 1$), or control-for-controls ($T_{age} = T_{year} = 0$); dashes indicate roles a cohort never serves in. The 1940 cohort serves as the common control-for-controls and therefore appears far more often than the others.

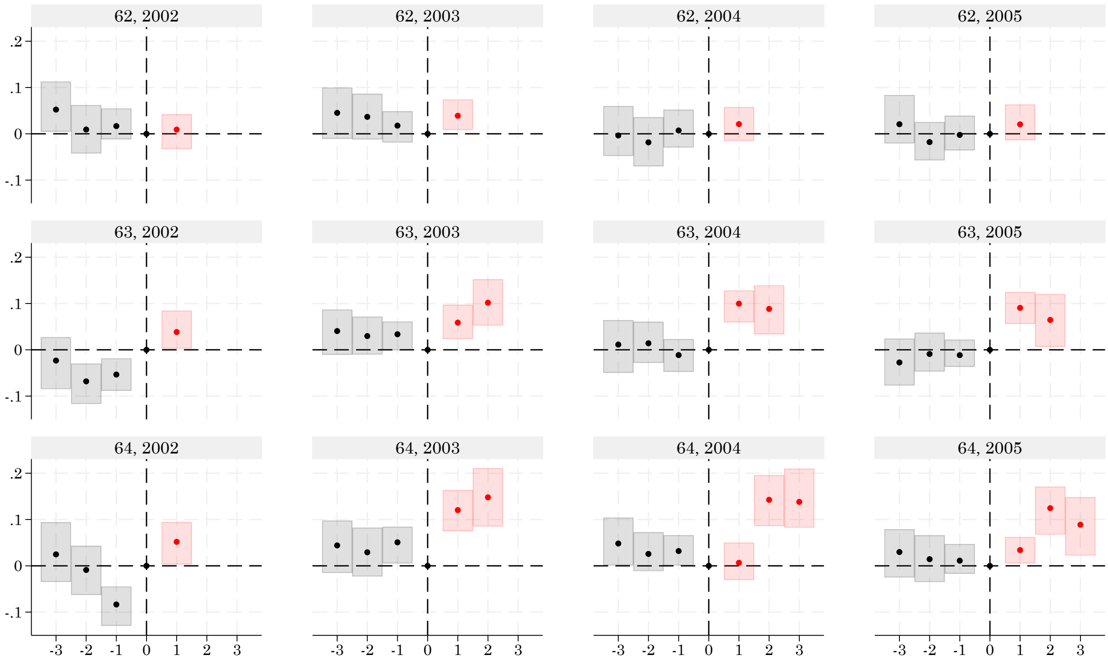

Figure A1: **Single estimators, labor earnings, in levels**

Note: The graph displays the estimated dynamic treatment effects with 95% bootstrap-based confidence intervals (120 trials). Labor earnings are measured in terms of the "basic amount" (G), which is used to annually adjust all social security transfers in Norway ($1G \approx 17\%$ of average full-time-full-year earnings in Norway).The age and year entered in the panel labels refer to the rightmost estimate in the respective panels. The time axis is centered on the last year prior to treatment (year 0) which is also the year excluded from the regressions. The estimation sample includes all treatment groups aged 62-64 in years 2002-2005. The same-year controls are aged 61 at the time of treatment. The same-age controls are the closest not-yet treated cohort of the same age.

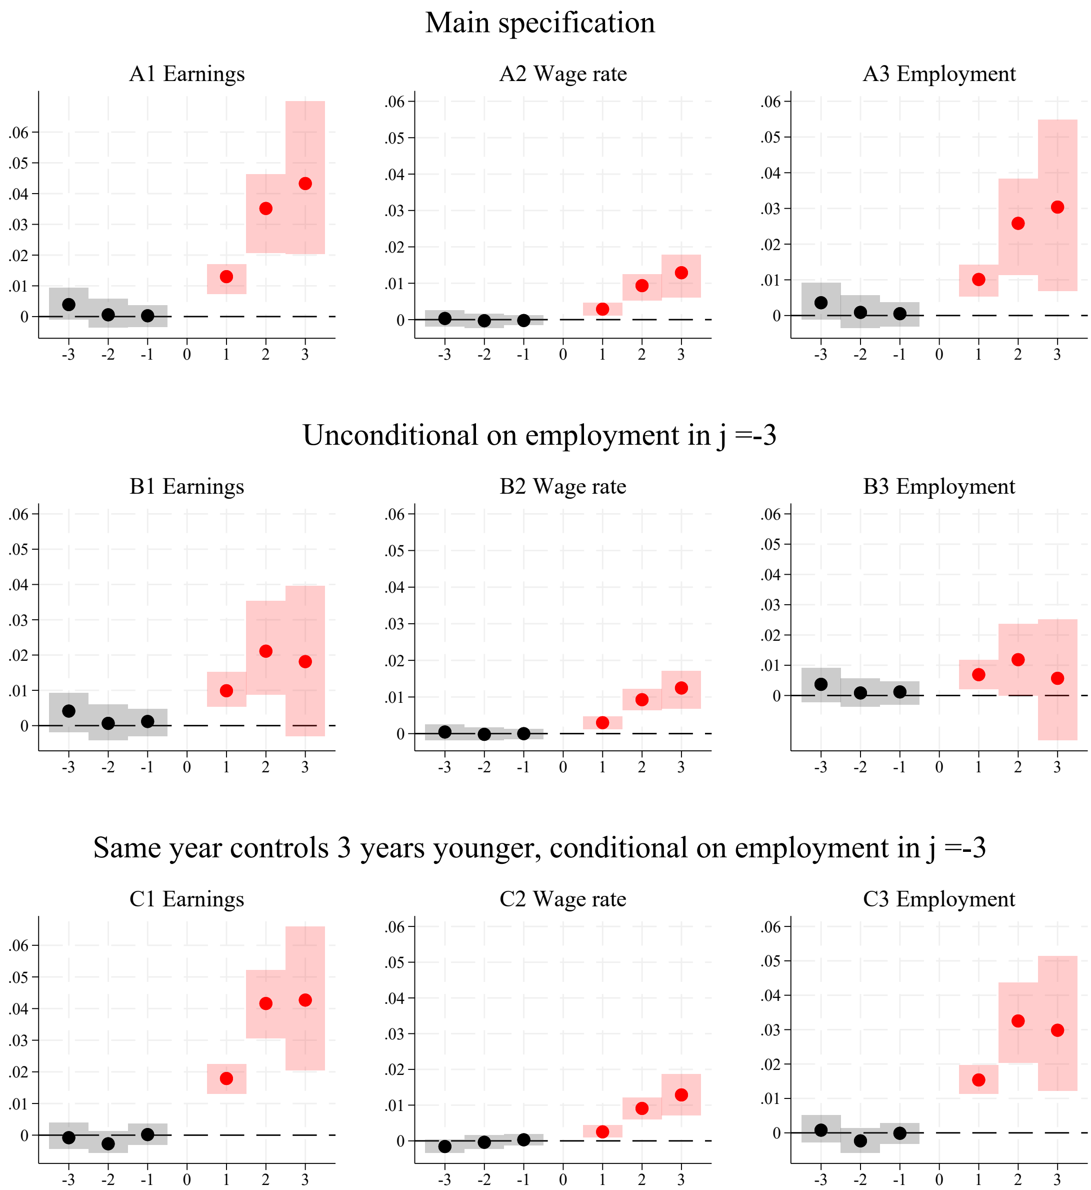

Figure A2: **Robustness: Sample and control group selection.**

Note: The graph displays the estimated dynamic treatment effects with 95% bootstrap-based confidence intervals (120 trials). The time axis is centered on the last year prior to treatment (year 0) which is also the year excluded from the regressions. The upper A-panels repeat the estimates from panels (a), (b) and (c) in Figure 3, the B-panels report estimates based on the complete population in the relevant birth cohorts, without conditioning on employment in $j = -3$, while the C-panels report estimates based on the main sample, but with the same-year control groups always being three years younger than the treated and the same-age controls.

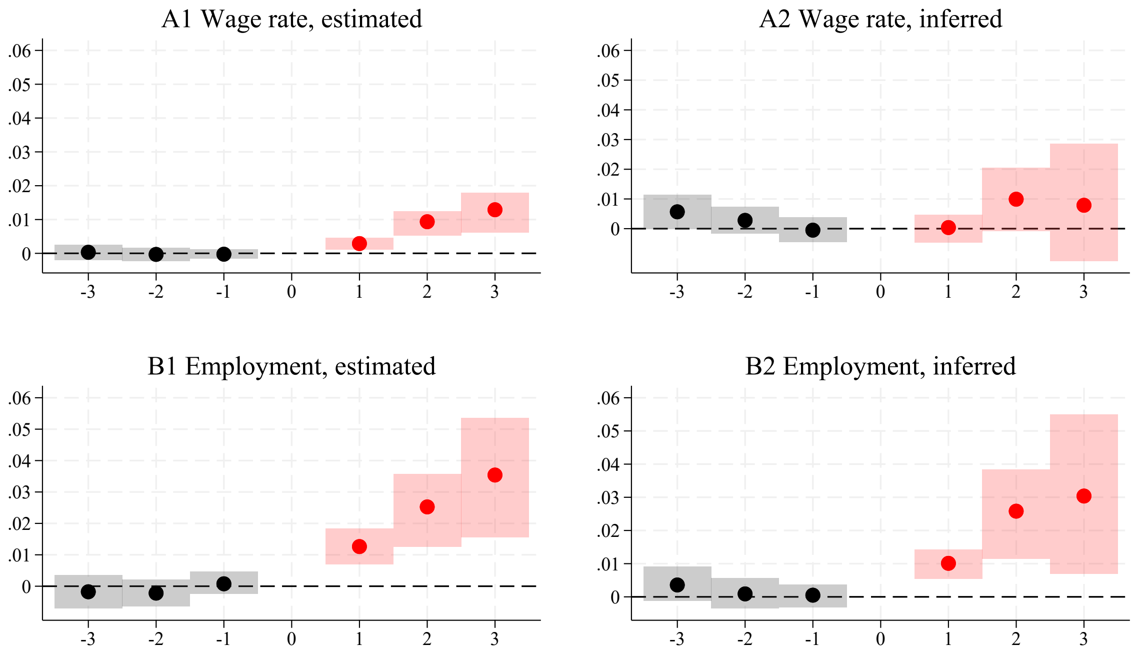

Figure A3: **Robustness: Earnings effect decomposition strategy**

Note: The graph displays the estimated dynamic treatment effects with 95% bootstrap-based confidence intervals (120 trials). Panels A1 and B1 repeat estimates from panels (b) and (c) in Figure 3, with wage effects estimated directly and employment effects inferred from earnings effects and wage effects. Panels A2 and B2 report estimates based on the reversal of the decomposition exercise, with employment effects estimated directly and wage effects inferred from earnings effects and employment effects.