

# Are senior workers overpriced?\*

Steinar Holden<sup>†</sup>

Simen Markussen<sup>‡</sup>

Knut Røed<sup>§</sup>

April 27, 2026

## Abstract

TBW

**Keywords:** Payroll tax, tax incidence, labor demand, labor supply, differences in differences

**JEL codes:** ? ?

---

\*This is a preliminary draft. Please do not cite without permission from the authors. The work has received funding from the Norwegian Ministry of Labor. Data is provided by Statistics Norway. Declarations of interest: none.

<sup>†</sup>Department of economics University of Oslo, Norway

<sup>‡</sup>Ragnar Frisch Centre for Economic Research, Oslo, Norway

<sup>§</sup>Ragnar Frisch Centre for Economic research, Oslo, Norway

# 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 considered 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 highly successful in raising the retirement age. 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). Labor supply among the elderly has later been boosted by extending employment protection legislation to cover workers up to age 72 and by ensuring full accumulation of additional pension entitlements up to age 75. As a result, employment and hours worked have increased considerably.

However, the employment of elderly workers may not only be limited from the supply side. Although there is no overwhelming empirical evidence that individual productivity systematically decreases with age (Hernæs et al., 2023), the *heterogeneity* in individual productivity is likely to increase at high age, 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 that offer steeper wage increases with tenure than what is rooted in productivity improvements (Lazear, 1979), giving employers an incentive to lay them off or (in the presence of employment protection barriers) encourage them to quit voluntarily.

One way of stimulating the demand for elderly workers is to reduce the payroll tax for employees exceeding a certain age threshold. This is exactly what was done in Norway in 2002 when the employer-paid 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. This was perhaps not the most obvious explanation, as the pension system of the early 2000s strongly subsidized early retirement and

thus offered workers weak incentives for continued participation in the labor market after the age of 62 (Hernæs et al., 2016). However, casual evidence at the time suggested that employers often took advantage of the generous early retirement conditions to ease downsizing processes by convincing senior workers to quit "voluntarily".

In the present paper, we evaluate the effects of the reform on employment (including both extensive and intensive margin) and hourly wages for the treated age group. The analysis is based on a flexible triple difference strategy in which we compare treated and not-yet-treated workers in a way that allows the inclusion of flexible age and calendar time controls. The strategy takes into account that although the reform was implemented at a specific point in time (July 1, 2002), the way it directly affected workers was staggered, both in terms of age and calendar year.

Our results suggest that the reform had the intended effect of raising employment of affected cohorts. Our most reliable effect estimate indicates an increase in the number of hours worked during age 62-64 around 2-3%. This effect is entirely accounted for by extensive margin responses (the number of days worked), which presumably arise through small postponements of (early) 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 considerable part of the tax cut (approximately 25%) was passed on to directly affected employees in the form of higher wage growth. Hence, the effect on employment may have been driven by a combination of labor supply and labor demand responses.

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 that only captured the first year of the reform and was also based on a methodology (event history analysis) different from the one used in the present paper. An evaluation of a similar reform implemented in Hungary in 2014 indicates more positive employment effects (Bíró et al., 2022). Here, a payroll tax cut that reduced wage costs for private-sector workers over 55 years of age by 5.3% is estimated to have 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. A plausible explanation is that the population we study consists of people with relatively safe jobs, for which we can assume that a small postpone of labor market exit is a real opportunity for almost all.

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 a Norwegian setting where wages are largely determined through collective bargaining with an explicit aim of wage compression (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 at the time varied from 14.1% of the total wage bill in the central parts of the country to zero in the most Northern part of the country (Nord-Troms and Finnmark). In the latter areas (which include only XX% of the labor force), the tax reduction had no effect, as it was not made negative. 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.

The reform of the payroll tax came approximately four years after the completion of an early retirement reform that for most workers had gradually reduced the age at which a full

pension could be claimed from age 67 in 1988 to age 62 in 1998.<sup>1</sup> For our analysis, this has the important implication that cohorts born before 1936 were subjected to a different pension system than the cohorts partly or fully affected by the payroll tax reduction. Hence, the menu of potential control cohorts, i.e., cohorts that are both unaffected and comparable, is strictly limited and depends on the age at which the effect is measured.

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 earnings-by-age path of a particular birth cohort, from the 1934-cohort (bottom line) to the 1949-cohort (top line). 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 they were exposed to a reduced payroll tax from age 62 through 66. Other cohorts were directly affected only in parts of the age 62-66 interval.

We focus on labor earnings rather than employment in Figure 1 for the reason that labor earnings are (third-party) reported without measurement error in Norwegian administrative registers, whereas employment status or hours worked must either be inferred from earnings data or be based on imperfect employer-employee registers. Identifying effects of the payroll tax reform on the earnings-by-age profiles may be considered a first-order approximation to its effects on labor demand, provided that not much of the tax cut is passed on to the workers in the form of higher wages (a premise we evaluate empirically later in this paper).

In Figure 1 we show average labor earnings for complete birth cohorts, measured in "base amounts" (BA), 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 report earnings measured in real terms, without conditioning on employment.

There are some important points to note from the earnings profiles shown in Figure 1. The first is that real labor earnings decline monotonously with age for all cohorts and that the steps of the decline become larger with age (with some prominent exceptions related to age-thresholds in the pension system). The second point to note is that age-specific earnings have increased over time, except at ages 63 and 64 during the period of early retirement expansion (1997-

---

<sup>1</sup>This reduction was also part of a tripartite tariff agreement between the associations of employers and employees and the state. The early retirement age was first reduced to 66 years in 1989, then to 65 in 1990, to 65 in 1993, to 63 in 1997, and then finally to 62 in 1998




Figure 1: **Labor earnings by cohort, age and year**

*Notes:* Labor earnings are taken from tax-records and deflated according to average wage growth in Norway using the "basic amount" of the social security system. Each line follows a specific birth cohort, from the 1934-cohort (bottom line) to the 1949-cohort (top line). The sample includes the full population of residents belonging to the respective cohorts without conditioning on employment

1998), reflecting a trend toward higher labor force participation among persons of mature age in Norway. Figure 1 can also be used to visually scan for reform effects. What we will then look for is a conspicuously large growth in earnings for the treated observations (the red dots in the figure), potentially followed by a conspicuously small growth (or decline) in the period after the second vertical line (when the tax cut was removed). However, in this case such an inspection does not give a clear answer. No doubt, the age-specific earnings levels are rising among the treated during the treatment period, but they are rising for the non-treated as well, and it is not obvious which are the most appropriate comparison groups. Hence, the reform effect "is not in the data" in that sense, and we need more refined tools to identify it. That is what we turn to in the next section.

It may be noted that all outcomes reported in this paper 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.<sup>2</sup>

<sup>2</sup>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.

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.

Our analysis is based on administrative register data that contain the entire population of residents in Norway, with annual accounts of all labor-related earnings. These data are then linked to employer-employee registers containing start and stopping dates for each employment spell and information on contracted work hours (measured in categories). We use these data to compute estimates of the total hours worked during each year for each employee.

Given that the population in Nord-Troms and Finnmark (with zero payroll tax and thus not exposed to the reform) is too small to be meaningfully used as a control population, we drop all observations from this part of the country.

### 3 Empirical strategy

There is almost no entry into the labor market at ages above 60 years.<sup>3</sup> Hence, an analysis of employment at this high age is in practice an analysis of labor market exit patterns. In our empirical analysis, we thus focus on persons that were employed at a slightly lower age (i.e., at age 58).<sup>4</sup> Our empirical strategy is based on a triple-difference event-study model where we compare the trajectories of treated cohorts, before and after treatment, to three control cohorts included to distinguish the treatment effects from the effects of aging, calendar time (e.g., business cycles) and cohort differences. We start by recognizing that we have a number of different combinations of treatment ages and treatment years, as indicated by the red dots in Figure 1. Each of these combinations may be associated with distinct causal effects, and possibly also distinct identification challenges. To evaluate the overall effects of the reform, e.g., as characterized by exposure time and/or age, we can thus start by estimating separate effects for each combination of age and calendar year, and then aggregate over those estimates to obtain interpretable and potentially policy relevant parameter estimates.

To estimate a treatment effect at age  $a$  in year  $t$  with both age- and time-fixed effects, we

---

<sup>3</sup>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%

<sup>4</sup>We present main results for the complete population in the online Appendix

need a same-age control group (which by construction cannot be observed at the same time) and same-time control groups (which by construction cannot be observed at the same age) for both the treated and the same-age controls. Figure 1 illustrates the rich set of potential treatment-control designs. This provides opportunities for examining model uncertainty and robustness, but it also highlights the need for some pre-defined principles for model selection in order to determine a preferred estimator.

We estimate the effects applying for a particular age-year combination  $r = a \times t$  with the following two-way fixed effects model:

$$y_{ij} = \alpha_i + \sum_{\substack{j=-3 \\ j \neq 0}}^J (\beta_{1j} + \beta_{2j} \textit{Age}_i + \beta_{3j} \textit{Year}_i + \beta_{4j} (\textit{Age}_i \times \textit{Year}_i)) + \varepsilon_{ij} \quad (1)$$

Here,  $\textit{Age}$  indicates observations at the age relevant for treatment, whereas  $\textit{Year}$  indicates observations in a calendar year relevant for treatment. We then have four groups, characterized by the following characteristics: i)  $\textit{Age} = 0, \textit{Year} = 0$ , ii)  $\textit{Age} = 1, \textit{Year} = 0$ , iii)  $\textit{Age} = 1, \textit{Year} = 0$ , and iv)  $\textit{Age} = 1, \textit{Year} = 1$ . The actually treated observations are those in the latter group. The sum operator works over relative time  $j = -3, \dots, J$ , with  $j = 0$  corresponding to the last year prior to treatment for the treatment group.<sup>5</sup>  $\beta_{1rj}$  is thus a common relative time effect.  $\beta_{2rj}$  is a relative time-effect for the two older cohorts vs. the two younger cohorts, while  $\beta_{3rj}$  is the relative time-effect for the two late cohorts vs. the two early (in calendar time) control cohorts.  $\beta_{4rj}$  captures the remaining difference between the treatment group and the control groups; i.e., the effects of interest. Finally, the model also contains an individual fixed effect ( $\alpha_{ri}$ ), allowing the four cohorts to be on different levels of the outcome. The individual fixed effect adds the last difference in the triple difference estimator.

A natural principle for the selection of control groups is to minimize age and time differences relative to the treatment group. To fix ideas, consider the case of estimating the effects of the reform on persons aged 62 years in 2002 (the first reform year). The closest-in-time not-yet-treated same-age control group in this case is persons aged 62 years in 2001 and the closest-in-age not-yet-treated same-year control group is persons aged 61 in 2002. To complete the difference-in-difference setup, we also need a same-year control group for the same-age controls, which will then be persons aged 61 in 2001. This implies that members of the treatment group in this particular case will serve as controls the year before they become treated, and thus help identify

---

<sup>5</sup>The choice of -3 as the starting point for individuals' outcome profiles ensures that we avoid the inclusion of profiles influenced by the reform of the early retirement system ending in 1998.

pre-reform time effects.

We estimate the model separately for each combination  $r$  of  $a = 62, \dots, 64$  and  $t = 2002, \dots, 2005$ , in total 12 combinations.<sup>6</sup> It can, however, also be joined into a common estimator involving all or a subset of  $r$  by stacking the data and estimate the model jointly using Ordinary Least Squares (OLS). Assuming that the sizes of the effects depend on exposure time only (and not on age or calendar year), the estimation equation becomes 2, where all coefficients are estimated separately by  $r$  except the coefficients capturing effects by exposure time  $\beta_{4,-3}, \dots, \beta_{4,t}$ . The person-fixed effects are estimated separately for each  $r$  by role (where "role" refers to the concepts of old, late, and treated), as each role exploits different segments of of a person's work-history.

$$\begin{aligned}
 y_{ijr} = & \alpha_{ir} + \sum_r \sum_{\substack{j=-3 \\ j \neq 0}}^J (\beta_{1rj} + \beta_{2rj} \textit{Age}_i + \beta_{3rj} \textit{Year}_i) \mathbb{1}[r] \\
 & + \sum_{\substack{j=-3 \\ j \neq 0}}^J \beta_{4j} (\textit{Age}_i \times \textit{Year}_i) + \varepsilon_{irj}
 \end{aligned} \tag{2}$$

Although our estimation strategy controls flexibly for both age, time, and cohort effects (the latter through individual fixed effects) in outcome levels, identification can still be threatened if the age *profiles* of the outcome changes over time; e.g., through a trend toward later retirement not captured by the individual-fixed effects. While it is difficult to rule out such underlying trends, the event-time setup makes it possible to assess their presence over a pre-treatment period. We can also to some extent evaluate their potential influence by choosing alternative control cohorts. A potential weakness with the estimation approach is that our choice of closest-in-age same-year controls implies that a few specific cohorts are re-used several times. In particular, the 1940-cohort will serve as the same-year control for the same-age control for all the estimates shown in Figure 1. This makes our estimate vulnerable to any circumstances that are specifically relevant for this cohort. Although we are not aware of any such circumstances, we report in the online appendix results built on an alternative strategy, where we chose different controls for each each age-time combination (with a three-year difference).

Given that the tax cut was abolished in January 2007, it is in principle also possible to use the never-treated instead of the not-yet-treated as controls (or interpret the abolishment

---

<sup>6</sup>The model can also be estimated for age 65, but we consider the data to be too thin for this exercise to be meaningful

as a new reform amounting to a tax *increase*). However, the first never-treated cohort consists of persons who reached the age of 62 in 2007, and since the earlier neighboring cohorts were treated to different degrees, it is difficult to find appropriate comparison cohorts. Attempts to use models similar to Equation 2 to estimate the effects of the abolishment of the reform have thus failed to exhibit satisfactory "pre-trend" properties. Moreover, the post-treatment period encompasses a rather unstable labor market environment resulting from the financial crisis. We have thus chosen to present results based only 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 shown in Figure 1) 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 per day worked (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) 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 standard errors and confidence intervals is based on a bootstrap procedure (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. While estimates reported to the left of the vertical lines (in black) refer to pre-treatment years (included to facilitate an assessment of

pre-trends), the estimates to the right (in red) refer to treatment effects by years of exposure. Hence, year 1 on the x-axis refers to the treatment effect in the first year of treatment, and year 2 to the treatment effect in the second year conditional on also being exposed in year 1.

The top panels of Figure 2 report estimates measured at age 62. Here, we see little sign of any causal effects, but instead some borderline-significant effects in the pre-period. Moving on to rows two and three (ages 63 and 64), there are some clear 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, and we cannot claim the complete absence of spurious effect estimates. This may call for alternative strategies for the selection of control groups, a point to which we return in later sections.

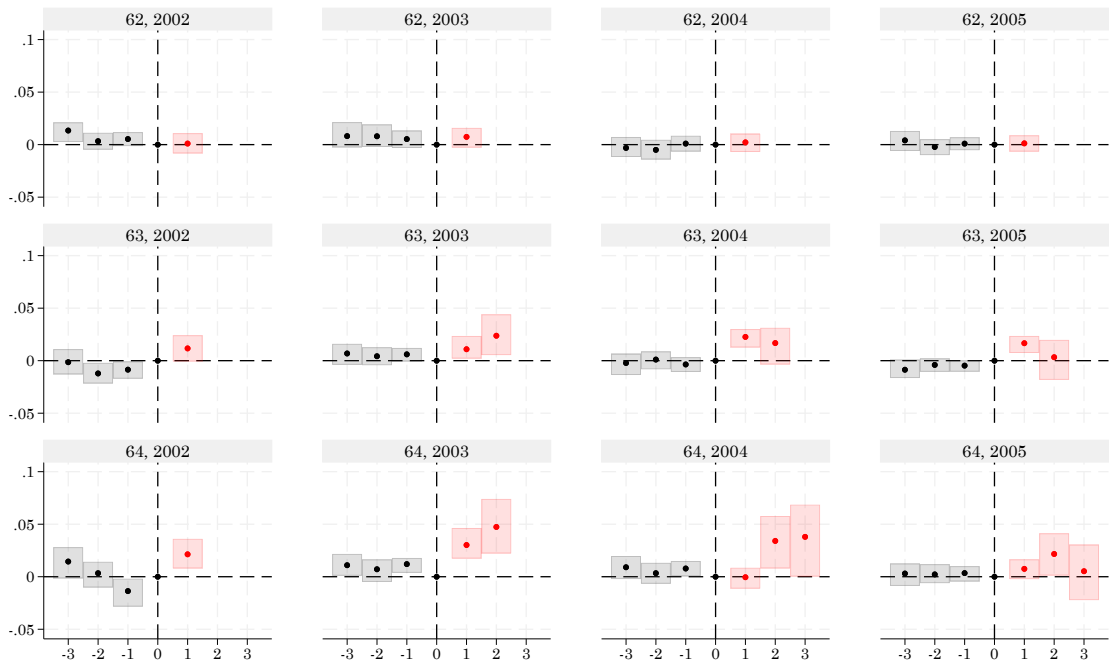


Figure 2: **Single estimators, labor income, relative responses**

*Notes:* The graph displays the estimated dynamic treatment effects with 95% confidence intervals. The effects are 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 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.

Our next step is to aggregate all the estimates in Figure 2 into a single set using Equation (2), assuming effect heterogeneity along the dimension of treatment exposure time only. The resultant estimates are shown in panel (a). They indicate that the 4 percentage point reduction in the payroll tax caused labor earnings to rise by approximately 1.5% in the first year of

exposure and by 3-4% in the two subsequent years. Whereas the first- and second-year effects are estimated with a reasonable degree of precision, the third-year effects are less informative, as they are estimated with considerable statistical uncertainty.



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

*Notes:* The graph displays the estimated dynamic treatment effects with 95% confidence intervals. The effects are 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, except in panel (b) where the estimation sample includes persons with full-time job in both the base-year and the outcome-year. 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.

## 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 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 by estimating the effect on total wage income for persons registered to have worked full-time all relevant years. For this strategy to identify the reform’s effect on average wages, we must assume that the scope for hours-variation for full-time employees is limited and that the reform’s effect on their wage growth does not deviate from the effect on other workers.

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 *and* having earnings exceeding 3 Basic amounts. In this case, we follow the standard practice of specifying the outcome as  $\log(\text{earnings})$  to obtain effects measured in relative terms.

The estimated wage effects are displayed in panel (b) of Figure 3. They 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 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.

### 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. Based on the (imperfect) data collected from the employer-employee register, we make an attempt to decompose the estimated hours effects into these two margins. Since hours worked per day is

a meaningful concept only for those who actually work, we perform this 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 comparing with the effects on total (fractional) employment. We thus use the following approximation:

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

The results are shown in Figure 3, panels (d) and (e). They suggest that the effect on employment is fully accounted for by the extensive margin, as defined by the number of days worked. There is no effect on hours worked per day. A natural interpretation of this result is that the effect on employment arises exclusively from a slightly 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 not formally a part of public expenses but instead paid for by a separate fund financed by employers through a fixed "tax" 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 the 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. However, for this estimate to represent the total fiscal consequences of the reform, we must rule out any spillover effects on other employees. If the extra hours worked by the treated workers substitute for hours worked by non-treated (younger) workers, the degree of self-financing will be smaller than indicated by our estimate.

#### 4.5 Robustness and heterogeneity

Table 1 reports the estimated treatment effects for a number of different groups, defined on the basis of employment status three 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), 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.

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 larger and the effects on employment were smaller for employees with high earnings/wages.
- The effects were roughly the same for full-time and part-time workers.
- In the private sector, the reform increased employment, but not the wage level.
- In the public sector, the reform increased both employment and the wage level.
- The reform increased the wage level for men and women equally much, but the effects on employment were much larger for men.

In the online appendix, we report results from two additional robustness exercises, one where we use the total population belonging to the relevant age groups as the foundation for analysis (i.e., without conditioning on initial employment) and one where we select same-year controls

Table 1: Estimated effects, by subgroup

	Labor earnings		Wage rate (FT)		Employment
	<i>N</i>	Estimate [95% CI]	<i>N</i>	Estimate [95% CI]	Estimate [95% CI]
All	148,506	0.035 [0.027, 0.043]	35,040	0.009 [ 0.007, 0.011]	0.026 [ 0.018, 0.034]
<i>Working time</i>					
Full-time	66,453	0.032 [0.021, 0.042]	20,958	0.008 [ 0.006, 0.011]	0.024 [ 0.013, 0.034]
Part-time	33,216	0.034 [0.016, 0.057]	2,418	0.008 [ 0.001, 0.017]	0.026 [ 0.008, 0.050]
<i>Earnings level</i>					
High earnings	72,954	0.038 [0.027, 0.047]	26,202	0.012 [ 0.010, 0.015]	0.026 [ 0.014, 0.035]
Low earnings	75,551	0.037 [0.027, 0.046]	8,838	0.001 [−0.003, 0.004]	0.036 [ 0.027, 0.047]
<i>Wage level</i>					
High wage	29,125	0.024 [0.008, 0.039]	11,933	0.011 [ 0.007, 0.014]	0.013 [−0.003, 0.028]
Low wage	30,005	0.047 [0.031, 0.061]	8,612	0.006 [ 0.004, 0.009]	0.041 [ 0.023, 0.055]
<i>Gender</i>					
Men	77,292	0.049 [0.037, 0.062]	22,873	0.010 [ 0.007, 0.012]	0.039 [ 0.027, 0.052]
Women	71,214	0.010 [0.001, 0.019]	12,167	0.009 [ 0.007, 0.012]	0.001 [−0.008, 0.008]
<i>Sector</i>					
Private	56,326	0.027 [0.014, 0.040]	14,919	−0.002 [−0.005, 0.002]	0.029 [ 0.015, 0.041]
Public	60,306	0.041 [0.032, 0.051]	14,472	0.022 [ 0.018, 0.024]	0.019 [ 0.012, 0.030]
<i>AFP eligibility (private sector)</i>					
AFP	36,084	0.021 [0.005, 0.042]	8,039	−0.003 [−0.007, 0.000]	0.024 [ 0.008, 0.045]
No AFP	30,207	0.034 [0.020, 0.056]	8,907	0.002 [−0.002, 0.007]	0.031 [ 0.018, 0.049]

*Note:* Point estimates with 95% confidence intervals in brackets. High and low earnings levels and high and low wage levels are defined as levels above or below the median observation. As employment effects are derived from the estimates on labor earnings and the wage rate, we do not report observation numbers for these estimates.

differently to avoid excess reliance on a single cohort (the 1940 birth cohort). The main 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 raising the demand for senior labor. The implicit presumption was that workers of high age in many cases had become too expensive for employers to justify their continued employment. In the present paper, we have provided evidence that the policy had the intended effect of raising employment among senior (above age 62) workers by 2-3%, and did so primarily by delaying exit from the labor market. Most of the effect came about through the realization of excess supply; ie., without raising the take-home wages for workers. We can thus conclude that there are some senior workers, particularly in the private sector, whose employment options are constrained from the demand side and that this constraint to some extent can be alleviated by reducing the payroll tax. But there

were apparently also some responses along the labor supply curve. Our estimates indicate that approximately 25% of the tax cut was passed on to workers in the form of higher wage growth.

Our findings suggest that differentiation of the employer-paid payroll tax may be a useful tool to offset imperfections in the wage formation if these imperfections imply that particular worker groups are entitled to higher wages than what can be supported by 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.
- 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.
- 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.
- Lazear, E. P. (1979). Why is there mandatory retirement? *Journal of political economy*, *87*(6), 1261–1284.
- 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.

## Appendix

ADD APPENDIX TABLE AND FIGURES HERE