Impact of Raising the Retirement Age on Firms

Stefan Hut*
Amazon

February 2023

Abstract

This paper studies the impact of raising the retirement age on firms. I rely on a policy change in the Netherlands that sharply increased the effective retirement age for workers born after 1950. For every additional older worker due to the reform, firms employ 0.6 fewer younger workers and reduce investments by 6,000 EUR annually. Adjustments are concentrated in low free cash flow firms, and only these firms experience declines in revenue and profitability. Lastly, I show that more time to anticipate the policy helps firms smooth adjustments and reduce negative profit effects.

*Senior Economist, Amazon. I am extremely grateful to my advisors Emily Oster, John Friedman, Jesse Shapiro and Neil Thakral for their valuable guidance and feedback. I thank Anna Aizer, Daniel Bjorkegren, Kenneth Chay, Andrew Foster, Simon Jaeger, Ryan Kessler, Brian Knight, Nicholas Reynolds, Bryce Steinberg, Matthew Turner, Juan Pablo Uribe and participants in the Brown Applied Microeconomics Seminar, the Applied Breakfast and the Health Breakfast for their comments and suggestions. I am grateful for generous financial support from the Alfred P. Sloan Foundation Pre-doctoral Fellowship on the Economics of an Aging Workforce, awarded through the NBER, the Inequality Project at the Brown Department of Economics and the Brown University Population Studies and Training Center. Results based on calculations by Stefan Hut, Brown University, using non-public microdata from Statistics Netherlands. Under certain conditions, these microdata are accessible for statistical and scientific research. For further information: microdata@cbs.nl.

1 Introduction

Many countries have increased their retirement age in response to demographic pressures, including increases in life expectancy, low birth rates, and rising dependency ratios.[1] While there is a large literature studying the effects of these reforms on individual retirement decisions (e.g. Staubli and Zweimuller, 2013; Vestad, 2013), there is limited evidence on how firms respond. In this paper, I examine the impact of a retirement policy on firm demand for labor and capital, revenue and profits. How do increases in employment among older workers due to a retirement policy affect firm demand for younger workers and capital? What kind of firms primarily adjust labor and capital, and how does this affect their bottom line?

To answer these questions, I use detailed administrative worker and firm data to analyze the effect of a policy change in the Netherlands that severely limited early retirement options for individuals born in 1950 or after but not those born before. Rather than starting retirement as early as 60, affected workers were more likely to work until the legal retirement age of 65, shifting the average retirement age by approximately two years. I rely on quasi-random variation in the birth year of workers at a firm level, comparing firms that had more treated workers, born in 1950 or 1951, to firms that had more untreated workers, born in 1948 or 1949. The analysis relies on a generalized difference-in-difference approach which assumes that close enough to the 1950 cohort boundary, the number of workers with a given birth year is as good as random. Using linked employment records and firm balance sheets for 2001-2018 made available through Statistics Netherlands, I examine firm worker composition and investment over time and assess how these adjustments translated into changes in firm revenue and profits.

First, the results show that the policy had a substantial effect on firm retention of older workers. On average, each additional treated worker caused firms to retain 0.25 additional older workers in the main years that the policy had an effect. This reflects the impact of the policy on the likelihood of early retirement – i.e., retiring before the legal retirement age of 65 – for individuals born in 1950 versus 1949. Absent other adjustments, the average firm's wage bill increased by approximately 24,000 EUR per retained older worker, or about 10 percent of the median firm-level total wage bill. Falsification checks show that the impact is unique to the 1949-1950 cohort cut-off, as opposed to being driven by other factors such as the Great Recession.

I further find that for each older worker retained, firms employed approximately 0.6 fewer younger workers. These adjustments are driven by reductions in hiring, rather than increases in firing or voluntary exits. I do not find adjustments in the number of middle-aged workers, nor intensive margin adjustments in hours worked or hourly wages among either young or middle-aged co-workers. Additionally, firms significantly reduced their investments in response to the policy. For each additional treated worker, firms on average reduced investments in equipment and machines by approximately 1,500 EUR annually, a reduction of 3 percent of total investment, or about 25% of the increase in the wage bill.

Second, I investigate mechanisms underlying the observed response. I find that reductions in younger worker employment and investments are larger for firms with low free cash flow prior to the policy. On

[1] The US raised the Social Security retirement age for people born in 1938 or later, by a few months for each birth year, until it reaches 67 for people born in 1960 or later (Social Security Administration, 2022). Most EU member states have similarly carried out substantial pension reforms over the last decades to maintain fiscal sustainability (Carone et al., 2016).

average, for every 1 EUR increase in the wage bill of older workers due to the policy, these firms reduce their younger workforce by about 0.7 EUR in payroll terms, and investments by about 0.2 EUR. Firms appear to absorb the remaining 0.1 EUR using external finance. The finding that low free cash flow firms exactly offset their increased older worker wage costs is consistent with cash constraints being a primary driver of the observed response. To investigate this further, I rely on proximity to the Great Recession by comparing adjustments among low free cash flow firms who relied primarily on bank vs. non-bank (government) funding prior to the policy. I find suggestive evidence that firms relying on bank funding – who were arguably more affected by the crisis – reduced their younger workforce more in response to the policy. In general, however, I cannot rule out that other investment or growth opportunity differences between high and low free cash flow firms may be driving some of the observed differences.

To understand how the observed adjustments impact a firm's bottom line, I next examine firm revenue and profitability. The median treated firm experienced a modest 2 percent increase in revenue and a 1 percentage point decline in the probability of reporting a positive profit (a 7 percent decline relative to baseline). Higher revenue is driven by an increase in the size of the firm's work force, and lower profitability is largely due to higher labor costs for retained older workers. I find that the negative profit impacts, too, are concentrated among low free cash flow firms, especially those that were shrinking prior to the policy. Shrinking firms can only make arguably costlier adjustments to labor through firing as opposed to reductions hiring, and such firms primary adjust investments rather than reducing their younger workforce. The median firm in this group experiences about a 3 percentage point increase in the probability of reporting a loss, which is a 20 percent increase relative to baseline.

Lastly, I show that more time to anticipate the retirement reform can help mitigate negative impacts of the policy on firm profitability. I rely on differences in the exact birth year of workers as a source of variation in the time to anticipate the policy. Intuitively, absent the reform, workers born in 1951 would have started early retirement one year later than workers born in 1950.[2] Each consecutive birth year therefore shifts the treatment window – the years in which workers would have retired early absent the policy but now continue to work – by one year. I use this variation to show that with three additional years to anticipate the policy, firm younger workforce adjustments are cut approximately in half. In addition, I find that each additional year of anticipation reduces the negative effect on the probability of reporting a positive profit by 1 percentage point, a reduction of about 20 percent. I find suggestive evidence in the data that firms with more time to anticipate may have built up cash reserves, given that they appear to have higher free free cash flow the years before the policy took effect.

This paper contributes to a growing literature studying the economic impact of retirement policies (e.g. Staubli and Zweimuller, 2013; Coile and Gruber, 2007; Gruber and Wise, 2000; Manoli and Weber, 2016). Most of these papers study the impact on individual retirement decisions, with the exception of a handful of papers that examine the impact of retirement policies on youth employment.[3] The paper's primary

[2]In either case workers start early retirement beginning at 60 years old, which occurs in 2010 for workers born in 1950 but in 2011 for workers born in 1951.

[3]Notable exceptions are recent papers by Boeri et al. (2016) and Bovini and Paradisi (2019) which both study firm adjustments in terms of younger workers employed and wages in the context of a public pension reform in Italy. In addition, Mohen (2019) studies the impact of delayed retirement on youth employment in the US, exploiting variation in the age distribution across US commuting zones.

contribution to this literature is to estimate the impact of these retirement policies on firm capital adjustments and firm performance. The results suggest that retention of older workers can negatively impact firm investment and profitability. In addition, I show that firm response and performance impacts is concentrated among low free cash flow firms. This is consistent with a growing body of literature that demonstrates that financial constraints affect labor decisions (see e.g., Schoefer (2015) and Giroud and Mueller (2016), Saez et al. (2017)).[4] I build on this literature by showing the financial impacts of constraints and demonstrating that additional time to anticipate a policy change can help firms smooth adjustments to their younger workforce and reduce impacts to firm profitability. In addition, the findings have implications for studies using observed firm response to shocks to estimate firm-level substitutability between workers of different cohorts (or other inputs).[5] The differential response between high- and low free cash flow firms found in this paper suggests that one cannot naively interpret changes in firm demand for labor or capital as direct evidence of input substitutability. In fact, assuming all of the difference in observed impact between low and high free cash flow firms were due to financial frictions, the implied elasticity of substitution between older and younger workers would be biased by up to a factor of seven when failing to take into account this mechanism.

The remainder of this paper is organized as follows. Section 2 describes the early retirement form and administrative data. Section 3 includes the empirical strategy and identification assumptions used in the analysis. Section 4 presents the overall effects of the policy. Section 5 studies heterogeneity in firm adjustments, between high and low free cash flow firms and depending on the time to anticipate the policy. Section 6 concludes.

2 Institutional Setting and Data

2.1 Policy Change: Early Retirement Reform

I study a 2005 policy change in the Netherlands that severely limited early retirement options for individuals born in 1950 or after, but not individuals born before. The early retirement scheme was first introduced in the 1980s, with the purpose of allowing older workers to retire at a younger age. The notion was that this would "free up" space in the labor market, allowing firms to hire more younger workers.[6] The schemes gave individuals the opportunity to retire in their early 60s, rather than at the legal retirement age of 65. The exact early retirement age varied by sector and industry (see Appendix B for more details).

On January 1, 2005, the Dutch government introduced a law that abolished the early retirement scheme for individuals born on January 1, 1950, or after. The main reason for this policy change is that from a public finance perspective the plan became unsustainable due to population aging.[7] As a result of the law, affected cohorts could no longer use the plan nor its fiscal advantages. In addition, the government imposed

[4] In addition, there is a large literature studying the relationship between investment and firm cash flow starting with Fazzari, Hubbard, and Petersen (1988, 2000), Kaplan and Zingales (1997, 2000) and between labor demand and firm cash flow in Benmelech et al. (2015).

[5] Including, for example, Jager (2016) who uses firm adjustment to labor in response to worker death as evidence of labor substitutability.

[6] This idea is not unique to the Netherlands: Gruber and Wise (2000) highlight that this is a common motivation for promoting early retirement among older cohorts of workers.

[7] A transcript of the law in Dutch can be found here: https://wetten.overheid.nl/BWBR0018053/2014-12-20.

a 52% tax penalty on any alternative early retirement arrangements offered by firms to workers. The main implication for workers born in 1950 or after is that the abolishment of the early retirement scheme made it much costlier to retire before the legal retirement age of 65.[8] As a result, the share of workers retiring before 65 dropped considerably.

Impact on individual retirement rates

Figure 1A shows that individuals born in 1950 are about 20 percentage points less likely to retire before the legal retirement age of 65 than individuals born in 1949. While the policy change did not fully eliminate early retirement among affected workers, it substantially reduced early retirement rates.

Figure 1B further breaks down retirement across different ages by cohort. It shows the cumulative retirement rates of workers born in 1948, 1949, 1950 and 1951. These figures further demonstrate that individuals born in 1950 or 1951 are substantially less likely to retire before the age of 65. The gap starts to open up at 60, and the largest cohort difference can be found around ages 62-64. Over 40 percent of individuals born in 1948 or 1949 will have retired by then, compared to only about 20 percent among 1950 or 1951 cohorts.

Representativeness for retirement reforms in other countries

This policy change is just one example of a host of reforms that have led to an increase in the retirement age across Europe and the US. Most EU Member States have carried out substantial pension reforms over the last decades in order to enhance fiscal sustainability, while maintaining adequate pension income (see Carone et al., 2016). For example, Italy implemented a similar reform called the "Fornero" pension reform in 2011. The reform raised the legal retirement age, reducing the number of new retirees and increasing the average age at retirement.[9] In the US, too, several Social Security reforms have led to an increase in the retirement age. A primary example is the Full Retirement Age (FRA) which has historically been at age 65, but is rising over time to age 67 in a series of steps, with the first increase for those born in 1938 (who turned 65 in 2003).[10]

How firms respond to these policies depends importantly on employment protection legislation in the country. While the Netherlands on average has stricter employment protection laws than the U.S., the OECD indicators for employment protection legislation indicate it scores similarly to other OECD countries in terms of procedural requirements, notice and severance pay.[11] In all OECD countries, notice period and severance pay are increasing with longer job tenure: on average, severance pay is seven times as high at 20 years of job tenure as at 9 months of job tenure in the OECD. As a result, dismissal of older workers tends

[8]Aside from increasing the effective retirement age, abolishing the early retirement scheme also had financial consequences for both affected workers and firms. These effects are small relative to the primary employment effect I exploit in this paper, however. I further discuss this in Appendix section B3.

[9]This reform has been studied by Bovini and Paradisi (2019) who find impacts on youth employment but do not study other firm outcomes.

[10]Mastrobuoni (2009) uses this discontinuity to show that these reforms had an impact on actual individual retirement ages: each two-month increase in the FRA is associated with a one-month increase in the average age of retirement.

[11]Including, for instance, Czech Republic, Israel, Portugal, Turkey, Belgium, Italy, Latvia, Greece and Luxembourg. See Chapter 3 of the 2020 Edition of the OECD Employment Outlook.

to be extremely costly not just in the context of the Dutch reform studied in this paper, but across most OECD countries.

2.2 Data

Main administrative data

To study the impact of the policy reform, I use administrative linked worker-firm data for the full population of the Netherlands. These data were made available through Statistics Netherlands (CBS). I combine employment and demographic data at the individual level with firm-level investment records, balance sheets and income statements. The data used in this paper span 2001 to 2018. The novel combination of worker and firm level data allows me to observe changes in firm worker composition, investment and capital over time and assess how these adjustments translate into changes in firm revenue and profits.

At the individual worker level, the data consist of the universe of employment contracts in the Netherlands, which are based on employment records from the tax authorities. These include the contract start- and end date, monthly earnings, hours worked, hourly wage, and benefits. Each employment contract has a person identifier and a firm identifier. I use these to link the files to demographic data, allowing me to observe an individual's birth year, birth month, and gender.

I aggregate these employment records to the firm-year level, to construct annual measures of the number of total- and full time workers, new hires, worker exits, mean hours worked, mean hourly wage and the firm's total wage bill. I construct both overall measures and measures by cohort, for each birth year separately and combined into three broader cohort groups: younger, middle-aged and older. Younger workers are defined as being born after 1965, middle-aged workers are defined as being born in 1955-1965, and older workers are defined as being born 1945-1955.[12] The goal of this sub-division is to examine how firm worker composition changes after the policy change and study adjustments made by firms to each cohort of workers separately.

I then apply several restrictions to the data, with the goal of ensuring that I capture actual firms and not self-employed individuals, temporary or seasonal businesses, etc. In addition, I only consider firms that were operational in 2005, the year that the policy change was announced. Hence, I restrict attention to firms that: (1) have least 5 full-time employees in 2005, (2) first appear in the data before 2005, (3) are in the data at least until 2005, (4) are observed in the data each year and (5) are in the private sector. There are a total of 99,240 such firms in the data.

Panel A of Table 1 shows descriptive statistics for these firms. A large share of the firms considered are relatively small. The median firm size is 18 employees, of which 10 are full time employees, and the median wage bill is 229K EUR. There is large variation across firms along each dimension, however, representing the broad set of firms considered in the analysis.

To study investment and firm production, I rely on firm-level investment, balance sheet and income statement data. These data cover about 80% of firms in the non-financial private sector of the Netherlands. The data are based on corporate tax records supplemented with data from surveys sent out by the tax authorities to all firms that are subject to corporate taxation. Coverage is almost universal among larger

[12]Equivalently, these workers were 40 or younger, 40-50, or 50-60 years old when the policy change was announced in 2005.

firms, but is worse when considering smaller firms. The data cover a wide range of firm-level outcomes. The investment data contain firm-level investments in material assets broken down by type. The balance sheet data contain assets, liabilities and equity. And the income statement consists of revenue, labor cost, other input cost and profits. In the analysis I primarily examine adjustments of investments in material assets, as well as effects on revenue and profits.

I link the firm-level data with the employment data using a national firm register. A challenge in linking these data is that the firm balance sheets and income statements are at the company level, whereas the employment and investment data are at the firm level. Companies are at a higher aggregation level, and can consist of multiple firms. Approximately 17 percent of establishments are part of a multi-establishment firm. For the firm production and profit analysis, since finances are shared across firms, I aggregate the corresponding employment data to the company level. About 8 percent of companies in the data have multiple balance sheet and income statement entries corresponding to separate legal entities. Summing such entries would not yield consolidated company-level data, because the balance sheets also contain within-company transfers. As a result I drop these 8 percent of companies from the analysis. In practice, this implies dropping mainly very large companies which have multiple legal entities. Altogether, there are a total of 63,469 companies in the production and profit data.

Panel B of Table 1 shows descriptive statistics for these companies. First, note that these companies tend to be slightly larger than the firms in panel A, due to the fact that coverage is incomplete for smaller firms and that some firms are aggregated into single companies. The median company employs 20 workers, of which 12 are full time employed, and faces a wage bill of 326K EUR. Revenues for the median company are 1.4 million EUR and profits about 77K EUR annually.

Measure of free cash flow

To capture the differential impact of the policy for low and high free cash flow firms, I measure firm-level free cash flow prior to the policy. I calculate free cash flow directly from the balance sheet and income statement for each firm. To compute free cash flow I use the following accounting definition:

$$\begin{aligned} \textit{Free Cash Flow} &= \text{Net Income} + \text{Depreciation} + \text{Change in Provisions} \\ &\quad \text{- Change in Working Capital - Net Investment} \end{aligned}$$

Table 1 demonstrates that the median firm has approximately 29,000 EUR in free cash flow in 2001-2004, the years before the policy. There is large variation across firms, however. About 20 percent of firms have negative free cash flow.

For the purpose of the analysis I scale free cash flow by assets in the previous year, following the literature on investment sensitivity to free cash flow (e.g. Fazzari et al., 1988). This is important because firm free cash flow in levels is correlated with firm size. Smaller firms mechanically will have free cash flows closer to zero. I then take the mean of this scaled measure of free cash flow by firm for 2002-2004 to construct firm-level free cash flow prior to the policy. The measure of free cash flow used in the analysis is a binary measure capturing whether firms have above or below median free cash flow at baseline.

Note that firm free cash flow is correlated with other firm characteristics. While this measure has been used widely in the literature studying sensitivity of investment to financing constraints (e.g. Fazzari et al., 1988; Alti, 2003; Richardson, 2006), I cannot rule out that other drivers – such as differential growth or investment opportunities – correlated with free cash flow drive the observed difference in firm response. In fact, other factors such as age and size have been used as proxies for financial constraints in the corporate finance (see Farre-Mensa and Ljunqvist, 2016) and economics literature (e.g. Saez et al., 2019). Appendix table A1 indeed shows that firms with higher free cash flow tend to be larger in size and older. As a second proxy for liquidity constraints, I compare firm response depending on exposure to the Great Recession, which partially overlaps with the analysis period. While I do not have access to firm-bank linked data, I can observe the share of external finance coming from bank vs. non-bank (government) financing. I use this to compare firm response between low free cash flow firms that had higher vs. lower dependencies on bank financing. This approach has its own set of limitations, however, and similar to the free cash flow measure, there may be other systematic differences between firms relying on bank vs. non-bank financing that may impact employment and investment opportunities other than through financial constraints. I further discuss this in Section 5.

Lastly, to examine the importance of asymmetric labor adjustment costs I establish whether a firm is growing or shrinking between 2001 and 2005. I consider changes in the number of full time employees between 2001 and 2005 to create a binary measure of whether the firm has experienced positive or negative growth in the size of its work force. By this measure approximately 54 percent of firms were growing in the baseline period.

3 Empirical Strategy

In the analysis, I rely on the 1950 cut-off to estimate the impact of the policy reform on firm outcomes, including labor demand, investment, revenue and profits. As shown in Section 2.1, the policy change generates a sharp cutoff at the 1950 cohort at the worker level. To understand the impact of the policy at the firm level, I rely on the same cutoff but I compare firm-level outcomes and adjustments depending on the firm-level share of older workers born either in 1950 or after, who are ineligible for early retirement, or before 1950, who are still eligible. The key assumption is that conditional on having the same share of workers in a narrow set of cohorts around the cohort boundary, the share of workers who are born after 1950 – and hence treated – is as good as random.

I define treatment at a firm level depending on the firm's workforce in 2005. This is the year in which the policy change was announced. I focus on the firm's work force in this year, rather than after, because firm composition after 2005 is endogenous to the policy change. Hence, I count the number of individuals working at a firm in 2005 of each birth cohort, and classify any worker born in 1950 or after as treated.

I first illustrate this idea using a firm-level Regression Discontinuity (RD) for firms with only one worker in a narrow set of cohorts around the cutoff. I then generalize the approach to firms with more than one such worker using a generalized Difference-in-Difference (DiD) approach. The generalized DiD relies on a similar logic to an RD, and assumes that holding constant the share of workers in a narrow bandwidth around the cutoff, firms with a higher share of treated vs. untreated workers had parallel trends prior to the policy.

Simple firm-level Regression Discontinuity (RD)

First, to demonstrate how I use the worker-level cutoff to estimate firm-level impacts, I restrict to the simple case where there is only one worker per firm who is born in a narrow set of cohorts around the cutoff. I define this narrow set of cohorts as workers born in 4 birth years around 1950, i.e. workers born 1948-1951. [13] I then use an RD design that leverages the 1950 cohort boundary at the worker level. Because I restrict to firms with only one worker born in 1948-1951, the RD design also works at the firm level.

In order to estimate firm-level effects using an RD, I aggregate outcome variables y_j to be the mean firm-level outcome between 2010 and 2015, the main years that the policy had an effect. I then estimate the following firm-level RD regression equation:

$$y_j = \alpha + \beta \mathbb{1}(C_j \geq 1950) + \gamma(C_j - 1950) + \epsilon_j \tag{1}$$

where j represents a firm. The coefficient of interest is β, which is the estimate of the treatment effect: the firm-level impact of having a worker born after 1950. I control for the threshold-deviated running variable C_j, which is the birth year or cohort of the worker at firm j.

Figure 2 shows the RD plot graphically. I examine the first-stage relationship, where y_j is the number of older workers at the firm (born 1945-1955). Generally, the number of older workers at the firm is increasing in the worker's birth year C_j. This makes sense intuitively: the earlier the birthyear in which a worker was born, the more likely it is – even absent the early retirement policy change – that they retired by 2010-2015 (simply because such individuals are older). However, we observe a discrete jump at the cohort boundary of 1950: firms with a worker born just to the right of the cutoff employ approximately 0.24 additional older workers during the 2010-2015 window than firms with a worker born just to the left of the cutoff. [14]

The key assumption underlying this analysis is that only the worker's eligibility for early retirement changes around the cutoff, and all other factors that determine y_j are continuous with respect to C_j. I show in Appendix Figure A2 that there is no discontinuous jump in two pre-policy outcomes, namely the number of older workers (born 1945-1955) and firm revenue prior to the policy.

Extending to all firms: Generalized Difference-in-Difference Approach

To extend the analysis to all firms, including those with multiple workers born around the 1950 cutoff, I leverage a Generalized Difference-in-Difference (DiD) design. This approach has a similar logic to the RD, in that it compares firms with a higher share of treated workers born to the right of the 1950 cutoff to firms with a higher share of control workers born to the left of the cutoff. A generalized DiD has two main advantages. First, it allows me to analyze the impact of the policy at a firm-level even for firms that have

[13] Note that the width of this narrow set of cohorts trades off bias and variance: a larger set of cohorts will mean workers are less comparable (leading to bias) but increases sample size (leading to lower variance). In the generalized DiD results, I show that the results are robust to using different window sizes, see Figure A19.

[14] I estimated the RD treatment effect using the estimation approach in Calconico, Cattaneo and Titiunik (2014), relying on a local polynomial sharp RD.

potentially a large number of workers born around the cutoff. Second, it allows me to examine dynamics over time. Understanding dynamics is helpful both to assess how long firms take to respond and identify potential anticipatory behavior, but also allows for natural falsification checks.

To study the impact of the policy on the firm's workforce using a generalized DiD, I run a regression of firm level outcomes on the number of treated workers, controlling for the number of old workers and the total number of workers at a firm:

$$y_{jt} = \alpha_j + \tau_t + \sum_{t=2001}^{2017} \beta_t \times \mathbb{1}(year = t) \times N_{j,\,treat} + \sum_{t=2001}^{2017} \gamma_t \times \mathbb{1}(year = t) \times N_{j,\,old} + \sum_{t=2001}^{2017} \delta_t \times \mathbb{1}(year = t) \times N_{j,\,all} + \epsilon_{jt} \quad (2)$$

where y_{jt} is a firm j's outcome in year t, α_j are firm fixed effects, τ_t are year fixed effects, $N_{j,\,treat}$ are the number of workers born in 1950 or 1951, $N_{j,\,old}$ are the number of workers born between 1948 and 1951, and $N_{j,\,all}$ are the total number of workers in 2005.

The coefficients of interest are β_t: the effect of having an additional treated worker in 2005 on the outcome in each year t, holding fixed the number of old and total workers. I cluster standard errors at the firm level to address potential serial correlation of outcomes across periods.

The key assumption underlying the generalized DiD is that holding constant the number of workers born in a narrow set of cohorts around 1950, firms with a number share of treated vs. untreated workers had parallel trends prior to the policy. This assumption is testable empirically: we can assess whether the treatment coefficients β_t are indistinguishable from zero for all years prior to the policy. In addition, in Section 4, I run a set of falsification checks where all cohorts are either treated (1950-1953) or untreated (1946-1949) and show that there is no impact in these cases. This further helps rule out concerns that the impacts are driven by differential (macro-economic) trends between firms with different treatment intensity, as opposed to the impact of the policy itself.

In addition, the parallel trends assumption may be violated if firms anticipated the policy change, and hired workers differentially around the cohort boundary prior to 2005. This would mean that treated and untreated firms were not similar prior to the policy. While there was some general expectation in the years leading up to 2005 that eligibility for early retirement would change, the speed with which this was implemented – the policy was announced January 1, 2005, and took effect immediately – and the sharp cohort boundary imposed were not anticipated. Furthermore, I show in Appendix Figure A3 that among firms with workers born in 1949 and 1950, the share born in 1950 is 50 percent – and that this is true for firms of different sizes. This further suggests that firms did not differentially employ 1950 versus 1949 workers before the policy was announced.

Estimating equation for firm investment, production and profits

For firm production outcomes, I modify the above estimating equation slightly by estimating the impact of the *share* of workers born in 1950-1951 vs. 1948-1949. The reason is that for such outcomes, effects are more likely to be homogenous in percentage terms. Specifically, an additional treated worker has a different effect for a small firm, where that worker represents a relatively large share of the work force, than for a large firm, where one additional worker is likely insignificant.[15] For these outcomes I examine the impact of an additional percent of the workforce being treated, holding constant the share of the workforce that is older.

To examine these effects the estimating equation is:

$$ln\,(y_{jt}) = \alpha_j + \tau_t + \sum_{t=2001}^{2017} \beta_t \times \mathbb{1}(year = t) \times Share_{j,\,treat} + \sum_{t=2001}^{2017} \gamma_t \times \mathbb{1}(year = t) \times Share_{j,\,old} + \epsilon_{jt} \quad (3)$$

where y_{jt} is a firm j's outcome in year t, α_j are firm fixed effects, τ_t are year fixed effects, $Share_{j,\,treat}$ is the share of workers born in 1950 or 1951, $Share_{j,\,old}$ is the share of workers born between 1948 and 1951.

The coefficient of interest is β_t, which represents the effect of an additional percent of the workforce being treated in 2005 on the outcome in each year t, holding fixed the share of old workers.

Summary statistics on firm-level treatment

Among treated firms, the share of treated workers is on average around 8% of the total workforce and 10% of the total firm wage bill. The wage bill share is higher on average because the treated workers are relatively older, which means their wages tend to be higher than the average worker. Appendix Figure A4 shows the firm-level distribution of treated workers as a share of total employment and the wage bill. There are some firms where treated workers represent a higher share, although firms in the analysis are restricted to have at least 5 employees (see Section 2.2), which means there few firms where treated workers represent more than 20% of the workforce or wage bill.

Treatment effect heterogeneity

Lastly, in order to examine heterogeneity in adjustments depending on how constrained firms are, I augment equations (2) and (3) with variables interacting the treatment variables with baseline measures of firm constraints. These are described in section 2.2. The interaction term gives the difference in adjustments made by firms depending on the extent to which the firm is either cash constrained or faces greater labor adjustment costs due to firing constraints.

4 Overall Effects of the Policy

In this section I first document the average effect of the policy on firm worker composition and investments. The policy increased the retention of older workers, and in response firms reduced the number of younger workers employed and investments in machines and equipment. Altogether, the policy led to small increases in treated firms' overall wage bill and revenue, but for a subset of firms decreased profitability.

[15] Another way of stating this is that, depending on the outcome, the units in which we expect the treatment effects to be homogenous are different. For labor adjustments, this is in levels, but for firm production outcomes, this in shares. Appendix C2 further justifies these approaches by showing that the corresponding effects are homogeneous across firms of different sizes.

First Stage: impact on retention of older workers

Figure 3A demonstrates the main effect of the policy on firm worker composition. The policy caused firms to retain more older workers over time. The plotted coefficients correspond to β_t from equation (2), which represent the effect of having an additional treated worker at the firm in 2005, controlling for the number of old workers and total firm size in 2005. Each of the lines represents a separate regression, examining how an additional treated worker affects the number of older workers, middle-aged workers and younger workers at the firm over time.

First, each additional treated worker employed at a firm in 2005 caused firms to retain 0.25 more older workers by 2014.[16] The pre-trend before 2005 is flat, suggesting that there were no pre-existing differences between firms that are more or less treated. The effect of having an additional treated worker leads to an increase in the number of older workers starting in 2009-2010. This is when untreated workers (born in 1948 or 1949) start to use the early retirement scheme, but treated workers (born in 1950 or 1951) are more likely to continue to work. By 2015, when treated workers reach the legal retirement age of 65, the number of older workers at the firm starts to fall. By 2017, the effect of the treatment is gone. By this time all treated workers have reached the statutory retirement age of 65 and retirement rates across the cohorts are the same again.

I then conduct several robustness checks to validate that the Generalized DiD reflects the impact of the retirement policy change, as opposed to other factors that differentially impacted treated and untreated firms. Notably, the impact of the policy partially overlapped with the Great Recession, which raises the concern that the estimated treatment effect reflects differential impacts from the recession on treated firms with more older workers. To address this I conduct two falsification checks. I estimate the same treatment effect coefficients but using cohorts that were (1) all untreated, born 1946-1949, with an imaginary threshold at 1948; and (2) all treated, born 1950-1953, with an imaginary threshold at 1952. Appendix Figure A5 shows that there is no impact of the treatment using these placebo cutoffs. This addresses concerns that treated and untreated firms were differentially impacted by other factors, such as the Great Recession, which would have also impacted the placebo groups. Instead, the observed impact is unique to the 1950 cutoff.

Impacts on younger worker employment and wage bill

In response to the policy, firms made substantial adjustments to the number of younger workers employed. For each additional treated worker employed in 2005, firms employ 0.15 fewer younger workers by 2014. This implies that for each retained older worker firms employ approximately 0.6 fewer younger workers, partially but not fully offsetting the additional older workers employed. Adjustments to the young work force start to occur in 2010, and over time treated firms gradually employ fewer younger workers. By 2016, when treated older workers have all reached 65, the firms increase their younger work force size again. By 2018 firms on average firms employ the same number of young workers regardless of treatment.

I do not find economically or statistically significant adjustments to the middle-aged workforce. For every treated worker, firms employ at most 0.05 fewer middle-aged in 2013. Generally, however, adjustments in

[16]Note that the measured impact on older worker retention from the Generalized DiD is very similar in magnitude to the simple RD in Section 1.

the number of middle-aged workers are not significant and close to zero.

All of these adjustments in the younger workforce are driven by changes in new hires, rather than an increase in firing. Appendix Figure A6 plots the β_t coefficients from equation (2) with the number of new hires at the firm in each year as the outcome variable. For each treated worker, firms hire 0.04 fewer workers annually in 2010-2014. Integrating over these effects across years, differences in hiring can explain virtually all of the reductions in employment of younger and middle-aged workers observed for firms with treated workers. This suggests that firing is not a common margin of adjustment.

Furthermore, I do not find considerable adjustments on the intensive margin, in terms of hours worked or hourly wages among either young or middle-aged co-workers. Appendix Figure A7 plots the β_t coefficients from equation (2) with mean weekly hours worked (panel A) and mean hourly wages (panel B) among young and middle-aged co-workers as outcome variables. Note that the series starts in 2006, because I only observe hours worked or hourly wages since that year. The treatment effect on both hourly wages and hours worked are a precisely estimated zero. I can reject changes in hourly wages exceeding 0.2 EUR/hour, which is small relative to the median hourly wage of 17 EUR/hour. In addition, I reject changes in hours worked exceeding 0.3 hours/week. Note that hours worked and hourly wages are often set in collective labor agreements that are set for a fixed period of time and require union bargaining to change. So it may not be surprising that it is hard for firms to make adjustments along these margins. This is in line with evidence in other papers suggesting wage rigidity may prevent adjustments along this margin (Bewley 2002; Kaur 2014; Cahuc, Carcillo, and Le Barbanchon 2018; Johnston, 2018).

Taken together, these changes in firm composition lead to a slight increase in the average wage bill among firms with treated workers. Each additional treated worker raises a firms' wage bill by 6,025 EUR annually, as shown in Figure 3. This is is equivalent to approximately 24,000 EUR per retained older worker, or about 10 percent of the median wage bill at these firms. The wage bill at treated firms is higher both because such firms employ slightly more workers in total, and because older workers tend to earn higher wages.

Investment Effects

A second margin of adjustment I consider is investment. On average, firms do not make significant adjustments to total investments in response to the policy. Appendix Figure A8 demonstrates that for each percent of the workforce that is treated at a firm, overall investments are lower by an insignificant 0.15 percent each year.

The lack of significant adjustments to overall investment masks changes to investment in more specific types of capital, however. Total material assets include: (1) owned real estate, (2) infrastructure and related equipment, (3) transportation equipment, and (4) machines and equipment. Some of these assets, such as real estate or infrastructure, may be difficult to adjust in the short-run. Existing literature has instead found – mainly in the context of business cycles – that investments in machines in particular are more volatile and subject to adjustments over time (see e.g. Cooper et al., 1999). Investments meant to replace existing machines may particularly be easier to adjust, because firms can decide upon the timing of those investments quite flexibly.[17]

[17] Note that in the data I am unable to distinguish between replacements of existing machines and installations, or investments

Figure 4 demonstrates that firms indeed make statistically significant adjustments to investments in machines and installations. For each percent of the workforce that is treated, firms reduce the amount of investments in machines and installations by 0.28 percent. At the median treated firm in the sample, treated individuals represent about 13 percent of the full time work force, which implies that investments in machines and installations decline by approximately 3 percent, or 1,477 EUR per year (Table 2).

Firm Production and Profits

Lastly, the policy change led to a small increase in firm revenue, but for a subset of firms led to a reduction in profits.

Figure 5A shows that the policy change led to a modest increase in firm revenue. The figure plots the β_t coefficients from equation (3), which represent the effect on log revenue of having an additional percent of the firm's work force treated in 2005, controlling for the share of old workers in 2005. The coefficients imply that for each percent of the work force that is treated, firm revenue increased by 0.11 percent in 2011-2014. At the median treated firm in the sample, this implies a modest median revenue increase of about 2 percentage points. These revenue effects are driven by an increase in the size of the firm's work force: Appendix Figure A9 shows that there is no statistically or economically significant change in revenue per worker.

Figure 5B reports the effect of the policy on the probability of reporting positive accounting profit in any given year.[18] The coefficients correspond to β_t from equation (3) which represent the effect on the probability of reporting a positive profit of having an additional percent of the firm's work force treated in 2005, controlling for the share of old workers in 2005. For each percent of the workforce that is treated, the probability of reporting a positive profit declines by 0.07 percentage points in 2011-2014. At baseline approximately 85 percent of firms in the subpopulation report a positive profit. For the median treated firm, the implied decline in the probability of making a profit is approximately 1 percentage point, or about a 7 percent decline relative to baseline.

5 Heterogeneity: mechanisms and role of anticipation

Section 4 shows that the retirement policy caused firms to employ fewer younger workers and reduce investments. This led to modest impacts on firm revenue and profitability. These findings raise two main questions. First, what caused firms to adjust their demand for younger workers and investments, and did certain firms make larger adjustments than others? Second, if we understand the mechanism driving the results, is there anything policy makers can do to reduce negative firm impacts from these kinds of retirement reforms?

5.1 Mechanism: input substitutability versus cash constraints

The observed firm response to the retirement reform may be driven by two main channels. The first is input substitutability or complementarity. Since the policy caused firms to retain more older workers, if

in new machines and installations.

[18] To deal with noise in the data I discretized this measure into whether firms make any profits at all. Figure A10 shows a similar, though noisier, patterns for a continuous measure of profits as a share of last year's assets.

older workers are substitutes with younger workers and capital, this would cause firms to decrease their labor demand and investment. Existing literature has similarly attributed firm adjustments to labor shocks as evidence of input substitutability. For instance Jager (2016) examines labor adjustments in response to worker deaths as evidence of worker substitutability, and Bovini and Paradisi (2019) look at demand for younger workers in the context of a similar pension reform in Italy to assess labor substitutability.

A second mechanism underlying the observed results, however, could be cash constraints. The retirement reform caused firms to retain costly older workers, which are hard to fire (see Section 2.1).[19] As a result, cash-constrained firms may have been forced to reduce demand for other workers or capital. The importance of cash constraints for labor decisions and investments has been demonstrated in a large body of literature (see e.g., Fazzari, Hubbard, and Petersen (1988, 2000), Kaplan and Zingales (1997, 2000) on investments and Schoefer (2015) and Giroud and Mueller (2016), Saez et al. (2019) on labor demand).

Given that firm-level exposure to the policy reform impacts both substitutability between firm inputs and their cash flows, the research design does not allow me to separately estimate these two effects. Instead, I investigate whether the observed response is stronger for firms that are more likely to be cash constrained, as proxied for by pre–policy free cash flow (free cash flow). Note that in general, I cannot rule out that other factors correlated with pre-policy free cash flow could be driving the observed results. For instance, I cannot definitively rule out that low free cash flow may have had different growth or investment opportunities during the analysis period, which could cause firms to respond differentially to the policy reform.

Differential impact by pre-policy free cash flow

First, I find that firms with high free cash flow made virtually no labor and investment adjustments in response to the policy. Figure 6A displays this result for changes to the younger workforce, and Figure 6B for investments in machines and installations. For every additional treated worker, high free cash flow firms reduced their younger workforce by 0.02 workers, which is statistically indistinguishable from zero. Similarly, for every additional percent of the workforce that is treated, high free cash flow firms reduced investments by an insignificant 0.07 percent. This is equivalent to a 1 percent reduction at the median treated firm. Table 2 summarizes the response in terms of cash flow effects, showing that in euro terms the adjustments made are negligible. Every additional treated worker led to a reduction in the younger worker wage bill of 407 EUR and investments of 322 EUR – despite adding 8280 EUR in older worker wage costs.

Low free cash flow firms, on the other hand, make large adjustments to labor and investments. Figure 6A demonstrates that for every additional treated worker employed in 2005, such firms employ 0.25 fewer younger workers by 2015, offsetting the additional retainment of older workers at approximately a one-to-one ratio. Figure 6B additionally shows that the observed reductions in investments in response to the policy are also concentrated in low free cash flow firms. For every additional percent of the workforce that is treated, such firms experience a reduction in investment of up to 0.5 percent annually, about a 6 percent reduction relative to baseline at the median treated firm. Table 2 shows the same results in terms of total EUR adjustments. Each additional treated worker causes low free cash flow firms to reduce their younger worker

[19] Note that, as I highlight in Section 2.1, this is true in the Netherlands but holds more generally across OECD countries.

wage bill by 6,100 EUR and their investments by 1,800 EUR. The difference in younger worker adjustments between high and low free cash flow firms is significant at the 5% level, and the difference in investments at the 10% level.

In fact, low free cash flow firms fully offset any additional wage costs associated with retaining older workers through adjustments to their younger workforce and investments. Taking the ratio of the adjustments relative to the increase in older worker wages, I find that for every 1 EUR increase in the older worker wage bill, low free cash flow firms reduce their younger wage bill by 0.7 EUR and investments by 0.2 EUR. Figure 7 indeed shows that consistently across years low free cash flow firms do not experience any increase in total costs in response to the policy, where total costs are defined as total labor costs plus investments. High free cash flow firms, on the other hand, experience a total cost increase of 8,300 EUR in the main years that the policy had an effect, driven by the increased wage cost for older workers. Note that, as shown in Appendix Figure A13, both high and low free cash flow firms experienced the same first stage impact on the retention of older workers.

Taken together, the results are consistent with cash constraints being a primary driver of firm response. In order to investigate this further, I leverage pre-policy reliance on bank financing and proximity of the analysis period to the Great Recession as a second proxy for cash constraints. I examine whether firms that relied more heavily on bank financing – and presumably were more impacted by the recession – responded differently to the reform than firms that relied more on non-bank, primarily government, financing. Figure A15 shows that firms that firms that relied primarily on bank loans prior to the policy reduced their younger workforce more than firms relying on other sources of finance. While there may be other differences between these two sets of firms – for instance, I find that firms relying on bank financing are larger and more mature – that could drive the differential impact, results are consistent with cash constraints playing a role.

Benchmarking implied substitutability and cash effects

We can use the observed differential impacts to estimate an implied elasticity of substitutability and cash flow sensitivity. This can provide a benchmark of (1) whether the observed effects are quantitatively plausible if they were entirely driven by cash effects, and (2) how different the implied elasticity of substitution would be if we failed to take into account the cash effect mechanism. The latter is particularly relevant to understand the potential degree of bias if we naively estimated input substitutability based on observed firm response to labor shocks.[20]

First, in the context of this policy I find that for every 1 EUR increase in older worker wages, low free cash flow firms reduce their younger worker wage bill by 0.66 EUR (95% CI: 0.34 - 0.98) and investments by 0.16 EUR (95% CI: -0.08 - 0.4). This estimated cash flow sensitivity for employment is slightly higher than others in the literature. For example, Schoefer (2015) estimates an employment-cash flow sensitivity between 0.2 and 0.6. The estimated investment sensitivity, on the other hand, is lower than typically found. For example, Rauh (2006) estimates that every dollar decrease in cash flow causes a decrease in capital

[20]Note that, as stated above, this is a common approach in related literature that relies on firm response to labor shocks to estimate elasticities of substitution.

expenditures of \$0.6-\$0.7.[21] Combined, however, the implied cash flow effects are quantitatively plausible relative to other estimates in the literature.

Second, I find that the implied elasticity of substitution would be biased by at least a factor of seven when failing to account for the potential cash constraint channel from retaining older workers. Overall, for every retained older worker, low free cash flow firm reduce their younger workforce by about 0.6 workers (Section 4) whereas high free cash flow firms only reduce their younger workforce by 0.08 workers. If the difference between these two is indeed fully driven by cash constraints, this would mean the estimate of the substitutability between older and younger workers is off by a factor of 7. While other factors may also drive the difference in response between high and low free cash flow, this at least provides an upper bound on the potential bias from failing to take into account cash constraint channels.

Margin of adjustment: younger workforce or investment

Having observed both adjustments to labor and capital in response to the policy, I next study what drives the margin of adjustment. Theoretically, one factor that might drive the choice of margin is the magnitude of labor adjustment costs faced by firms. To examine this empirically I exploit asymmetries in adjustment costs faced depending on whether a firm is growing or shrinking at baseline. Intuitively, if firing costs exceed hiring costs, then firms that can reduce hiring might be more likely to adjust labor than firms that would need to fire workers. Growing firms can adjust labor by reducing the rate of hiring, whereas shrinking firms can only further downward adjust by firing – which is much costlier. Such firms may opt to cut investments instead.

Figure 9a shows that only low free cash flow firms that were growing adjusted their younger workforce in response to the policy. The figure shows the mean adjustment in the young workforce in 2011-2014, when the policy had the largest effect. For each additional treated worker employed in 2005, growing low free cash flow firms reduce their younger workforce by 0.35 workers. On the other hand, the coefficient for young workforce adjustments is 0.008 for low free cash flow firms that were shrinking, suggesting that labor is not a feasible margin of adjustment when the only option is to fire workers. Table 2 shows that in cash flow terms, growing firms reduce their younger workforce by 11,000 EUR whereas shrinking firms make an adjustment of only 1500 EUR, which is statistically indistinguishable from zero. The difference in response between growing and shrinking firms is significant at the 1% level.

Low free cash flow firms that are shrinking, then, can only make adjustments along other margins, notably by adjusting investments. Figure 9b shows that, indeed, the investment response to the policy is concentrated in low cash flow shrinking firms. In this set of firms on average investments decline by 0.5 percent for each percent of the workforce that is treated. Table 2 shows that this is equivalent to a 4,500 EUR reduction in investment for each additional treated worker. Growing firms only reduce their investment by an insignificant 500 EUR. This difference in investment response between growing and shrinking firms is

[21] There may be several reasons for this. First, other papers have typically relied on firm financial data from sources such as Compustat which include much larger firms than the firms in my dataset. It is possible that such large firms are more likely to respond by adjusting investment. Another possibility is that the nature of the cash flow shock affects the channel through which firms respond. The fact that in the context of this paper the cash flow shock came through an increase in older worker wages might make firms more likely to respond by adjusting other types of labor, rather than capital. I am unable to test this empirically, however.

significant at the 5% level.

These findings highlight the role of labor adjustment costs in firm response to the policy.

Role of external finance

In addition to reducing their hiring of younger workers and investments, firms could also potentially take on more debt in order to finance the retained older workers for the duration of the policy change. I examine this directly by assessing changes in firm debt levels as a result of the policy.

Appendix Figure A12 shows impact of the policy on the amount of debt taken on by firms, depending on their level of free cash flow prior to the policy. While somewhat noisy, the results suggest that low free cash flow firms increase their debt levels by around 4,000 EUR during the years that the policy has an effect. High free cash flow firms do not appear to increase their level of debt. To examine the magnitude of the debt increase, consider that the policy caused firms to face an increase in their older worker wage bill of approximately 8,500 EUR a year between 2011 and 2014. Since the level of debt is a stock variable, this suggests that low free cash flow firms absorbed about 10 percent of the increased older worker wage costs through external finance.[22]

Taken together, the results imply that firms absorb roughly 70 percent of the increase in older worker wages through reductions to their younger worker wage bill, 20 percent through reductions in investments and 10 percent through external finance. The small role for external finance is consistent with other evidence demonstrating that firms rely primarily on internal funds in response to shocks. For example, Schoefer (2015) finds that in the context of the US free cash flow is the dominant source of finance, accounting for more than 95% of total finance within a quarter.

Effects on firm production and profitability

Lastly, I separately study the impact of the retirement reform on firm production and profits for the above groups. I intersect pre-policy free cash flow with whether the firm was growing or shrinking, to assess what kind of firms were most impacted by the policy change. Figure 8A shows that the positive revenue effects observed in the reduced form results of section 5.1 are concentrated in high free cash flow firms. For each percent of the workforce that is treated, such firms see a revenue increase of 0.15 percent, which corresponds to about a 3 percentage point or 24,000 EUR increase annually at the median firm. Firms with lower cash flow at baseline, on the other hand, do not observe revenue growth. This translates into differences in firm profitability, as shown in Figure 8B. Firms with low free cash flow experience a decline of 0.14 percent in the probability of reporting a positive profit for each percent of the work force that is treated. For the median treated firm this corresponds to a decline of nearly 2 percentage points, or about a 14 percent increase in the probability of reporting a loss relative to the average baseline rate in this subpopulation of firms. High free cash flow firms, on the other hand, maintain their profits.

Among low free cash flow firms, those firms that were shrinking and had only one margin to adjust, namely investment, are the most affected by the policy. Figure 9c suggests that this set of firms is the only

[22] Firms increased their debt levels by about 4,000 EUR, compared to an increase of about 8,500*4=34,000 EUR in older worker wages over the total period.

that does not experience an increase in revenue, and in fact may see a small revenue decline (though not significantly so). In addition Figure 9d demonstrates that these firms are hardest hit in terms of profit. For each percent of the work force that is treated, such firms see a decline of 0.37 percent in any given year. For the median treated firm, this corresponds to about a 3 percentage point decline, which is equivalent to a 20 percent higher probability of making a loss in the main years that the policy had an effect.

5.2 Role of anticipation

Having demonstrated the impact of the policy on firms, and particularly on low free cash flow firms, a natural follow up question is: what can policy makers do to minimize negative firm impacts? This section explores a possible policy response: giving firms more time to anticipate. I examine this empirically by relying on variation in the timing of the treatment resulting from differences in the treated worker's exact birth year.

First, note that the results presented above suggest that firms did not make labor or capital adjustments before the policy first took effect. This is perhaps surprising given that all firms had at least 5 years to anticipate the policy change: the policy was announced in 2005, but the first affected cohort born in 1950 would not have reached the early retirement age until around 2010. Figure 3A shows that firms only start to make adjustments to their younger workforce by 2010, once the policy starts to have a substantial effect on the retainment of older workers. Even low free cash flow firms do not appear to exhibit any anticipatory behavior prior to 2010, as shown in Figure 6A.

Two potential explanations may underlie this result. First, it may be that firms in general are not forward looking and are myopic in their response to retaining an older worker. Second, it may be that firms were unaware of the policy until it had its first effect in 2010, or alternatively did not expect that the policy would actually be implemented until the first cohort was affected. If the latter is the case, it is possible that additional time to anticipate the policy might have a different effect.

Anecdotally there is some evidence that firms did not change their behavior in anticipation because retirement policies change frequently and, hence, there was a non-negligible probability that the proposed measures would be reversed before they first took effect. In addition, it appears as though most of the public attention to the reform occurred when it first took effect, not when it was announced: there is a large Google trends spike in search for the policy in 2010, and only a small spike in 2005 when it was announced. This suggests that once it became clear that the policy would actually be implemented, workers and firms paid more attention to it. From that moment on additional time to prepare for the the policy may have a different effect than the initial 5 years since the policy was announced.

To study how more time to anticipate affects firm adjustments in the data, I examine firm response depending on the exact birth year of treated workers. Each consecutive birth year shifts the treatment window by one year. Intuitively, consider that early retirement typically occurred between the ages of 60 and 64. A worker born in 1951 turns 60 in 2011, whereas a worker born in 1950 turns 60 in 2010. Hence, workers born in 1951 would have started early retirement one year later than workers born in 1950. Each consecutive birth year shifts the treatment window, and therefore the firm's time to anticipate the policy change, by one year.

I examine firm adjustments to the workforce depending on the exact birth year of treated workers born between 1950 and 1953.[23] The key object of interest in each case is how firms adjust their workforce relative to what their workforce *would have been* absent the policy change. To capture this counterfactual I rely on firms with 1949 workers, which were not affected by the policy change. I restrict my sample to firms with exactly one worker born in 1949-1953. I then compare outcomes for the firms with a 1950-1953 worker to outcomes for firms with a 1949 worker. To establish the correct counterfactual I shift the 1949 control firms by 1, 2, 3, or 4 years to serve as controls for the firms with 1950, 1951, 1952, 1953 workers, respectively. In effect, this means I compare firm outcomes for the firms with a 1950-1953 worker to outcomes for firms with a 1949 worker, *when* those workers were of the same age.

The estimating equation is:

$$y_{ja} = \alpha_j + \tau_a + \sum_{a=55}^{68} \sum_{c=1950}^{1953} \beta_{c,\,a} \times \mathbb{1}(age = a) \times \mathbb{1}(cohort = c) + \epsilon_{ja} \tag{4}$$

where y_{ja} is firm j's outcome when the treated worker is of age a, α_j are firm fixed effects and τ_a are age fixed effects. The coefficient of interest is $\beta_{c,\,a}$ which represents the difference in the outcome of a firm employing a treated worker of cohort $c = \{1950, 1951, 1952, 1953\}$ relative to the outcome of a firm employing a control worker born in 1949 when both workers were of the same age a.

Appendix Figure A18 shows the first stage plot demonstrating that each consecutive birth year shifts the treatment window by one year. Panel A shows the effect of having one more worker of each treated cohort 1950-1953 on the retention rate of older workers, relative to a firm with a 1949 worker. The plot demonstrates that the onset of the treatment at the firm level occurs one year later for each one-year shift in the treated worker's birth year. Panel B shows the same results but with age of the treated worker on the x-axis. This figure makes it clearer that the size of the first stage is similar for workers of different cohorts, and that it is just the timing that is different.

To examine anticipatory behavior, then, I compare adjustments in the number of younger workers employed depending on the birth year of the treated worker at each firm, and hence the time to anticipate the policy. Figure 10 shows these effects by age of the treated worker. The results demonstrate that firms that had more time to anticipate appear to make smaller and somewhat smoother adjustments to their younger workforce. The point estimates suggest that firms with a 1950 worker reduce the number of younger workers by approximately 0.2, which is in line with the results found in section 4. Firms with a 1953 worker, on the other hand, reduce their younger workforce by about half of that. Note that while somewhat noisy, the observed magnitude of response appears to be monotonic in the amount of time firms had to anticipate: each consecutive birth year – and hence additional year to anticipate – reduces the observed adjustment to the younger workforce.

In addition I find that the differential response due to anticipation translates into differences in firm profitability. Figure 11 shows that, while noisier, it also appears that firms that have more time to anticipate the policy experience smaller declines in their probability of reporting a positive profit. Similar to the labor

[23] Note that my data ends in 2018, so I run into data censoring issues when considering more recent cohorts. For this reason I only consider cohorts born up to and including 1953, as these workers turn 65 in the last year of my data.

adjustments in Figure 11, firms with a 1953 worker experience profit declines that are about half as large as firms with a 1950 worker.

To capture the effects of anticipation with more precision, I then estimate a linear cohort trend relying on variation at the firm level in the mean birth year of treated workers. This allows me to extend the analysis to all firms in my sample with workers born in 1950-1953, rather than firms with exactly one such worker. I again compare the outcomes of a firm with mean birth cohort c to a shifted version of a firm with a 1949 worker.[24] In this part of the analysis I treat the cohort variable c as a continuous variable and estimate how the firm's outcome varies with each one year shift in c:

$$y_{ja} = \alpha_j + \tau_a + \sum_{a=55}^{68} \sum_{c=1950}^{1953} \beta_a \times \mathbb{1}(age = a) \times c + \epsilon_{ja} \tag{5}$$

Table 3 reports the results, collapsing the β_a to the mean between ages 60-64 – when the policy was most binding. Column (1) shows that for each additional year to anticipate the policy, the estimated adjustments in the younger workforce are reduced by 0.04 workers. Given that firms with a 1950 worker adjusts its young workforce by -0.22 workers, this means that three additional years to anticipate approximately cuts the adjustment made by firms in half. In addition, column (2) suggests that for each additional year to anticipate, the negative effect on the probability of reporting a positive profit is reduced by 1 percentage point, relative to a reduction of 5 percentage points for firms with a 1950 worker.

Lastly, a puzzling fact about the anticipation effects observed here is that firms do not begin to adjust their younger workforce prior to the year in which the treated worker turns 60. Rather than affecting the timing of the response, anticipation appears to only affect the magnitude of the response.

How do firms smooth instead? I find suggestive evidence in the data that firms with workers born in later cohorts build up more cash reserves in anticipation of the policy. Examining free cash flow for these firms in the years leading up to the treatment, firms with 1953 workers appear to have approximately 32K EUR (95% CI: 8K, 56K) higher free cash flow each year than firms with 1950 workers. Given the large confidence interval on this estimate, however, this result is at best suggestive.[25]

6 Conclusion

In this paper I study the impact of a change in the legal retirement age in the Netherlands on firms. The policy reform led to an increase in older workers at firms. Relying on detailed administrative data that combines employment contracts with firm balance sheets, I demonstrate that the policy caused firms to make substantial adjustments to their younger workforce and investments. Firm adjustments – as well as negative impacts on profitability – are concentrated in low free cash flow firms, which is consistent with cash

[24] Note that Appendix C3 validates that the mean is the correct transformation by showing that the plots by birth year obtained from equation (9) are similar to those one would get when rounding mean birth year at a firm to the nearest integer. Other transformations, notably taking the minimum birth year of treated workers (which might capture the notion that the first treated worker is what matters, and not the average timing of the policy) do not yield similar adjustment patterns. Figures available upon request.

[25] Note that this finding is consistent with a result in Midrigan and Xu (2014) that firms can mitigate financing constraints by accumulating savings during good economic times.

constraints being a driver of firm response. From a policy perspective, the findings highlight that some firms are more affected by these types of labor market policies than others, depending on their ability to adjust firm inputs. The results also suggest a potential avenue to alleviating negative firm impacts, by giving firms more time to anticipate policies. I find that more time to anticipate the retention of older workers helped firms smooth adjustments to their younger workforce and reduced negative profit effects.

References

Acemoglu, Daron and David Autor, "Skills, tasks and technologies: Implications for employment and earnings," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1043–1171.

___ **and Pascual Restrepo**, "Demographics and Robots: Theory and Evidence," *Work in progress*, 2017.

___ **and** ___ , "Demographics and automation," Technical Report, National Bureau of Economic Research 2018.

Administration, Social Security, "Retirement Age Calculator," *https://www.ssa.gov/benefits/retirement/planner/ageincrease.html*, 2022.

Alti, Aydoğan, "How sensitive is investment to cash flow when financing is frictionless?," *The journal of finance*, 2003, *58* (2), 707–722.

Autor, David H, Frank Levy, and Richard J Murnane, "The skill content of recent technological change: An empirical exploration," *The Quarterly journal of economics*, 2003, *118* (4), 1279–1333.

Avolio, Bruce J, David A Waldman, and Michael A McDaniel, "Age and work performance in nonmanagerial jobs: The effects of experience and occupational type," *Academy of Management Journal*, 1990, *33* (2), 407–422.

Baker, George, Michael Gibbs, and Bengt Holmstrom, "The internal economics of the firm: Evidence from personnel data," *The Quarterly Journal of Economics*, 1994, *109* (4), 881–919.

Behaghel, Luc and David M Blau, "Framing social security reform: Behavioral responses to changes in the full retirement age," *American Economic Journal: Economic Policy*, 2012, *4* (4), 41–67.

Benmelech, Efraim, Nittai K Bergman, and Amit Seru, "Financing labor," Technical Report, National Bureau of Economic Research 2011.

Bewley, Truman F and Truman F Bewley, *Why wages don't fall during a recession*, Harvard university press, 2009.

Bloom, Nick, Stephen Bond, and John Van Reenen, "Uncertainty and investment dynamics," *The review of economic studies*, 2007, *74* (2), 391–415.

Boeri, Tito, Pietro Garibaldi, and Espen R Moen, "A clash of generations? Increase in Retirement Age and Labor Demand for Youth," 2016.

Bond, Stephen and John Van Reenen, "Microeconometric models of investment and employment," *Handbook of econometrics*, 2007, *6*, 4417–4498.

Börsch-Supan, Axel and Matthias Weiss, "Productivity and age: Evidence from work teams at the assembly line," *The Journal of the Economics of Ageing*, 2016, *7*, 30–42.

Bovini, Giulia and Matteo Paradisi, "The Transitional Labor Market Consequences of a Pension Reform," *mimeo*, 2017.

Cahuc, Pierre, Stéphane Carcillo, and Thomas Le Barbanchon, "The effectiveness of hiring credits," *The Review of Economic Studies*, 2019, *86* (2), 593–626.

Carone, Giuseppe, Per Eckefeldt, Luigi Giamboni, Veli Laine, and Stephanie Pamies, "Pension reforms in the EU since the early 2000's: Achievements and challenges ahead," *European economy discussion paper*, 2016, (042).

Chaney, Thomas, David Sraer, and David Thesmar, "The collateral channel: How real estate shocks affect corporate investment," *American Economic Review*, 2012, *102* (6), 2381–2409.

Chodorow-Reich, Gabriel, "The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis," *The Quarterly Journal of Economics*, 2014, *129* (1), 1–59.

Coile, C, "Working longer in the US: Trends and explanations (No. w24576)," 2018.

David, H and David Dorn, "The growth of low-skill service jobs and the polarization of the US labor market," *American Economic Review*, 2013, *103* (5), 1553–97.

Dostie, Benoit, "Wages, productivity and aging," *De Economist*, 2011, *159* (2), 139–158.

Duchin, Ran, Oguzhan Ozbas, and Berk A Sensoy, "Costly external finance, corporate investment, and the subprime mortgage credit crisis," *Journal of financial economics*, 2010, *97* (3), 418–435.

Farre-Mensa, Joan and Alexander Ljungqvist, "Do measures of financial constraints measure financial constraints?," *The review of financial studies*, 2016, *29* (2), 271–308.

Fazzari, Steven, R Glenn Hubbard, and Bruce C Petersen, "Financing constraints and corporate investment," 1987.

Giroud, Xavier and Holger M Mueller, "Firm leverage, consumer demand, and employment losses during the Great Recession," *The Quarterly Journal of Economics*, 2017, *132* (1), 271–316.

Greenwald, Bruce C and Joseph E Stiglitz, "Financial market imperfections and business cycles," *The Quarterly Journal of Economics*, 1993, *108* (1), 77–114.

Gruber, Jonathan and David Wise, "Social security programs and retirement around the world," in "Research in Labor Economics," Emerald Group Publishing Limited, 2000, pp. 1–40.

___ **, Kevin Milligan, and David A Wise**, "Social security programs and retirement around the world: The relationship to youth employment, introduction and summary," Technical Report, National Bureau of Economic Research 2009.

Hamermesh, Daniel S and Gerard A Pfann, "Adjustment costs in factor demand," *Journal of Economic Literature*, 1996, *34* (3), 1264–1292.

Jäger, Simon, "How substitutable are workers? evidence from worker deaths," *Evidence from Worker Deaths (January 14, 2016)*, 2016.

Johnston, Andrew C, "Unemployment Insurance Taxes and Labor Demand: Quasi-Experimental Evidence from Administrative Data," *Available at SSRN 3062412*, 2018.

Kaplan, Steven N and Luigi Zingales, "Do investment-cash flow sensitivities provide useful measures of financing constraints?," *The quarterly journal of economics*, 1997, *112* (1), 169–215.

Kehoe, Patrick, Elena Pastorino, and Virgiliu Midrigan, "Debt constraints and employment," Technical Report, National Bureau of Economic Research 2016.

Manoli, Day and Andrea Weber, "Nonparametric evidence on the effects of financial incentives on retirement decisions," *American Economic Journal: Economic Policy*, 2016, *8* (4), 160–82.

Mastrobuoni, Giovanni, "Labor supply effects of the recent social security benefit cuts: Empirical estimates using cohort discontinuities," *Journal of public Economics*, 2009, *93* (11-12), 1224–1233.

Michaels, Ryan, T Beau Page, and Toni M Whited, "Labor and capital dynamics under financing frictions," *Review of Finance*, 2018, *23* (2), 279–323.

Midrigan, Virgiliu and Daniel Yi Xu, "Finance and misallocation: Evidence from plant-level data," *American economic review*, 2014, *104* (2), 422–58.

Rauh, Joshua D, "Investment and financing constraints: Evidence from the funding of corporate pension plans," *The Journal of Finance*, 2006, *61* (1), 33–71.

Richardson, Scott, "Over-investment of free cash flow," *Review of accounting studies*, 2006, *11* (2-3), 159–189.

Saez, Emmanuel, Benjamin Schoefer, and David Seim, "Payroll taxes, firm behavior, and rent sharing: Evidence from a young workers' tax cut in Sweden," *American Economic Review*, 2019, *109* (5), 1717–63.

Schoefer, Benjamin et al., "The financial channel of wage rigidity." PhD dissertation, Harvard University 2015.

Shimer, Robert, "The impact of young workers on the aggregate labor market," *The Quarterly Journal of Economics*, 2001, *116* (3), 969–1007.

Staubli, Stefan and Josef Zweimüller, "Does raising the early retirement age increase employment of older workers?," *Journal of public economics*, 2013, *108*, 17–32.

Vestad, Ola Lotherington, "Labour supply effects of early retirement provision," *Labour Economics*, 2013, *25*, 98–109.

Yashiv, Eran, "Hiring as investment behavior," *Review of Economic Dynamics*, 2000, *3* (3), 486–522.

——, "Frictions and the joint behavior of hiring and investment," 2012.

‘Tables and Figures

Table 1: **Summary Statistics: Firms**

	Median	*Standard Deviation*	*Sample Size*
Panel A: Labor and Investment			
Total Employees	18	55.72	99,240
Full Time Employees	10	17.89	99,240
Wage bill (1000 EUR)	229	534	99,240
Material Investments (1000 EUR)	42.16	93.89	85,732
Panel B: Production			
Total Employees	20	30.78	63,469
Full Time Employees	12	15.15	63,469
Wage bill (1000 EUR)	326	653	63,469
Revenue (1000 EUR)	1,414	4,107	63,469
Profits (1000 EUR)	77	306	63,469
Free Cash Flow (1000 EUR)	29	1,322	63,469

Note: This table shows summary statistics for 2005 for the firms included in the subpopulation studied in this paper. Panel A displayes the labor and investment data. Panel B displays the production data. The production data is at the company rather than the firm level (higher level of aggregation). In addition, both the investment and production data are more restrictive in terms of coverage because they do not contain the universe of smaller firms and exclude the financial sector. Sample sizes differ for these reasons. See Section 2.2 for more details.

Table 2: **Overview: Cash Flow Effects of Policy**

	(1)	(2)	(3)		(4)	(5)	
	Overall	High Cash Flow	Low Cash Flow	P-Value: (2)=(3)	Low Cash Flow Growing	Low Cash Flow Shrinking	P-Value: (4)=(5)
Older Worker Wage Bill	8497.74***	8280.94***	8614.36***	0.853	8743.57***	8256.51***	0.802
	(783.02)	(1512.18)	(915.43)		(1564.41)	(1153.51)	
Younger Worker Wage Bill	-3381.15***	-406.68	-6122.63***	0.021	-11536.46***	-1464.01	0.004
	(1342.04)	(2166.21)	(1652.37)		(3287.67)	(1850.16)	
Investments	-1477.17**	-321.78	-1847.32**	0.089	-478.56	-4586.65**	0.047
	(690.75)	(1273.27)	(830.78)		(2432.57)	(2000.87)	
N	144,851	73,566	71,285		38,493	32,792	

Note: This table summarizes the cash flow effect of the policy in terms of the older worker wage bill, younger worker wage bill and investments in machines and installations in euros. The reported effects are from separate regressions. The coefficients represent mean effects in 2011-2014 - when the policy was most binding - of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers at the firm in 2005. The outcomes are the older worker wage bill, the younger worker wage bill and investments in machines and equipment. The regressions are run for all firms (column 1), and run separately for above (column 2) and below (column 3) median free cash flow firms. In addition, among below median free cash flow firms those that were growing (column 4) - and hence could reduce labor by reducing hiring - versus shrinking (column 5) - and had to fire workers to adjust labor - prior to the policy. Standard errors are clustered at the firm level. P-values significance: * = 0.1, ** = 0.05, *** = 0.01.

Table 3: **Anticipation Effects**

	(1) Younger Workers	(2) Probability Reporting Profit
β: Effect of additional year to anticipate	0.0462**	0.0095*
	(0.0166)	(0.0053)
Mean adjustment/effect for 1950 worker	-0.2247***	-0.0524**
	(0.0591)	(0.0249)

Note: This table reports the estimated β coefficient from equation (5). The coefficient represents the difference in outcome for each one unit increase in mean birth cohort of treated workers c at a firm, when that mean birth cohort is of age 61-64. Conceptually, it captures the effect of having an additional year to anticipate the policy change, during the time frame when the policy has the largest impact. In addition, the table reports the mean difference in outcome between treated firms with a worker born in 1950 - firms with the least time to anticipate - relative to untreated firms with a 1949 worker. The outcome variable in column (1) is the number of younger workers employed at the firm when the treated worker was 61-64, and in column (2) the probablity of the firm reporting a profit when the treated worker was 61-64. Standard errors are in parentheses and are clustered at the firm level. P-values significance: * = 0.1, ** = 0.05, *** = 0.01.

Figure 1: **Individual Level First Stage: Retirement by Cohort**

Panel A: Share of Individuals Retiring Before 65

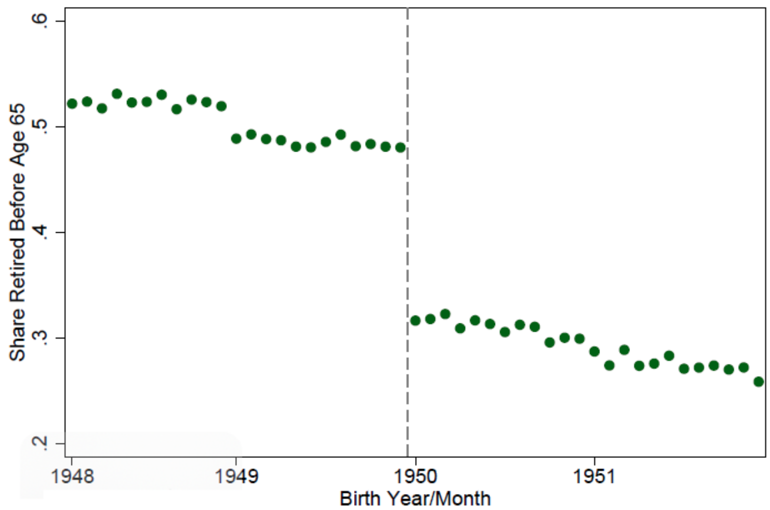

Panel B: Retirement by Age and Cohort

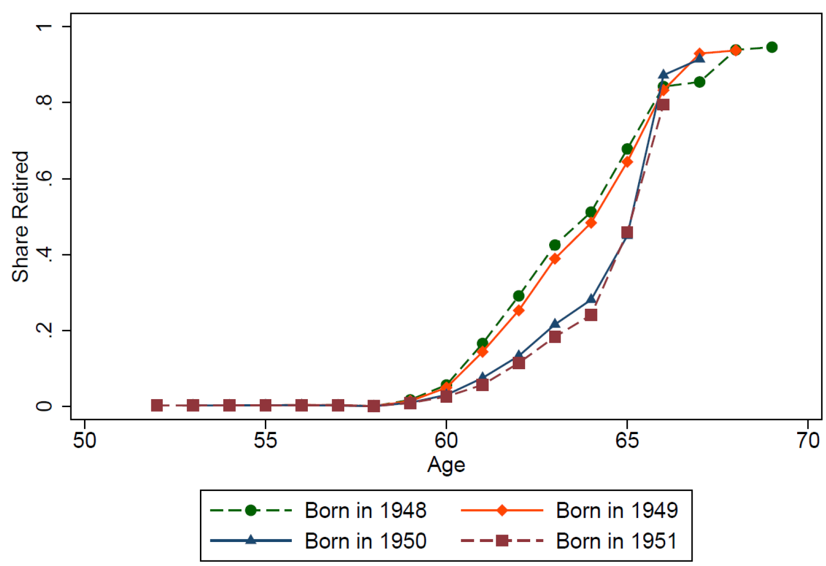

Notes: This figure shows individual-level retirement rates and highlights the sharp difference in retirement rates between workers born in 1949 and 1950. Panel A depicts the share of workers retiring before the age of 65, the legal retirement age in the Netherlands. Panel B depicts the cumulative retirement rate by age, for individuals born in 1948, 1949, 1950 and 1951. The population is restricted to all individuals observed working in the data at age 55 to capture the group affected by the policy: employed individuals.

Figure 2: **Firm-level RD: Simple Case with One Worker Born 1948-1951**

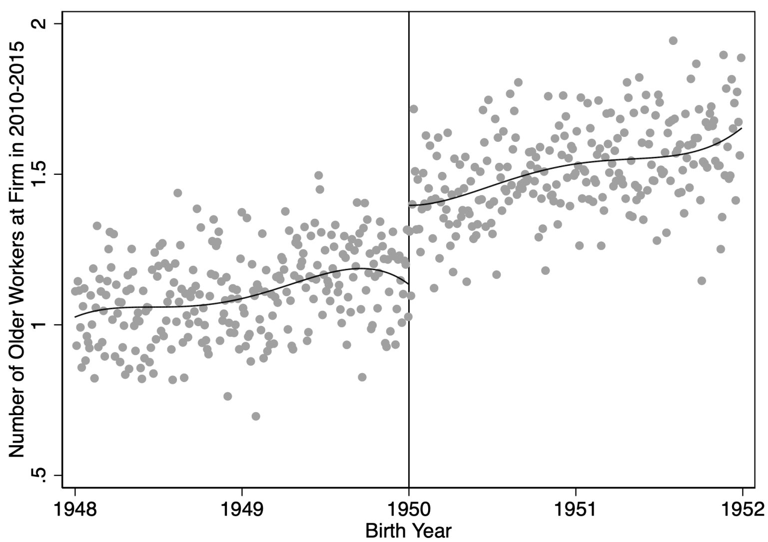

Notes: The unit of observation is the firm. I limit to firms with exactly one worker born in the 1948-1952 cohort window. The outcome variable is the mean number of older workers at the firm in 2010-2015, the main years that the policy change had an effect. The figure demonstrates that there is a substantial and significant impact of the 1950 threshold on the number of older workers employed. All workers considered are full time employed.

Figure 3: **Overall firm labor adjustments**

Panel A: Employment of workers of different cohorts

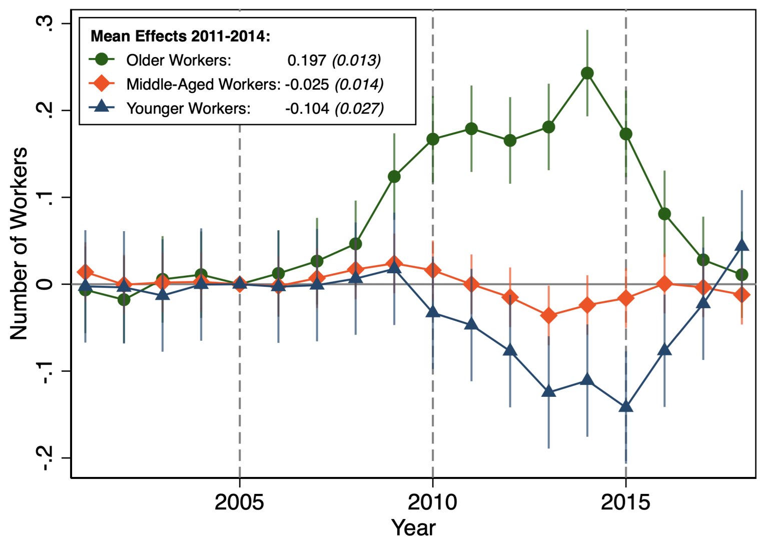

Panel B: Firm total wage bill

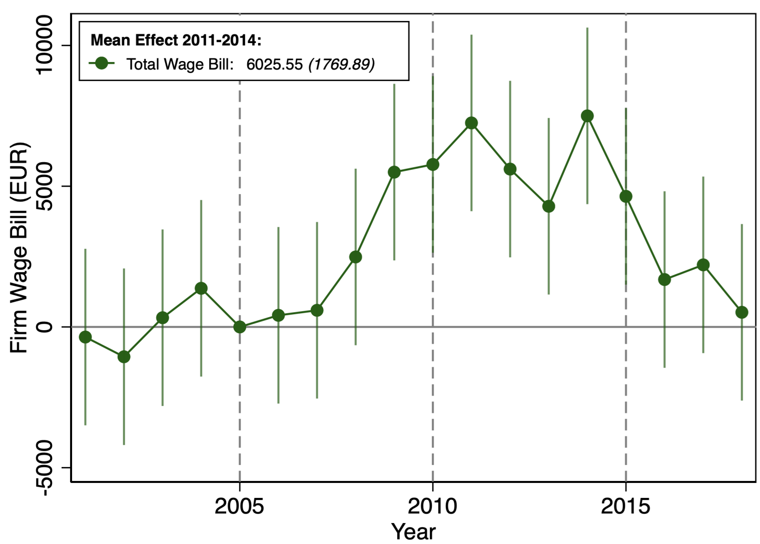

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on firm worker composition and the firm's wage bill. The coefficients plotted represent the effect in each year t of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. These correspond to β_t from equation (2). The coefficient for $t = 2005$ is normalized to 0. For panel A, the three lines represent three separate regressions. The outcome variables are the number of older workers (born 1945-1955), the number of middle-aged workers (born 1955-1965) and the number of younger workers (born 1965 and after) at the firm over time. All workers considered are full time employed. For panel B the outcome variable is the total wage bill for full time workers at the firm. The legend reports the mean of the coefficients between 2011-2014: $\overline{\beta}_{2011-2014}$ and the corresponding standard error in parentheses. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure 4: **Overall investment adjustments**

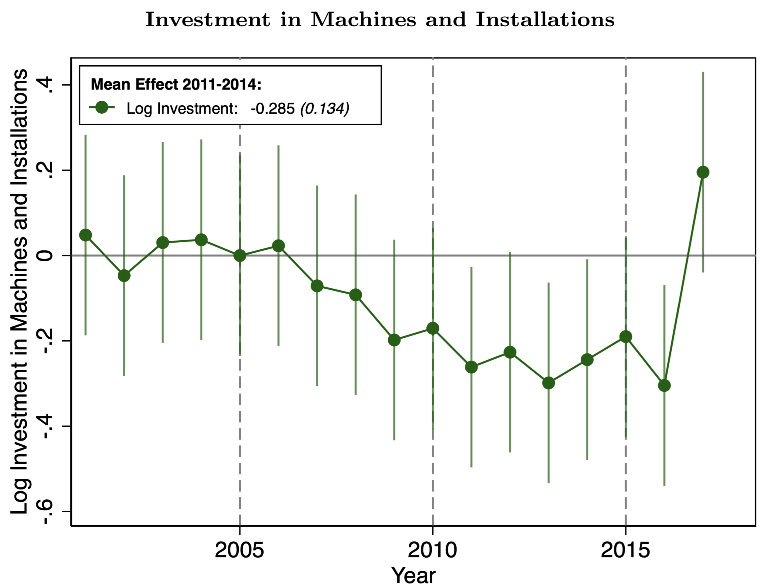

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on investments in machines and installations. The coefficients represent the effect in each year t of having an additional percent of the 2005 workforce treated (born 1950-1951), controlling for the share of old workers (born 1948-1951) at the firm. This corresponds to β_t from equation (3). The coefficient for $t = 2005$ is normalized to 0. The outcome variable is log investment in machines and installations. See section 2.2 for more details. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. The legend reports the mean of the coefficients between 2011-2014: $\overline{\beta}_{2011-2014}$ and the corresponding standard error in parentheses. Standard errors are clustered at the firm level.

Figure 5: **Overall firm production and profit effects**

Panel A: Revenue

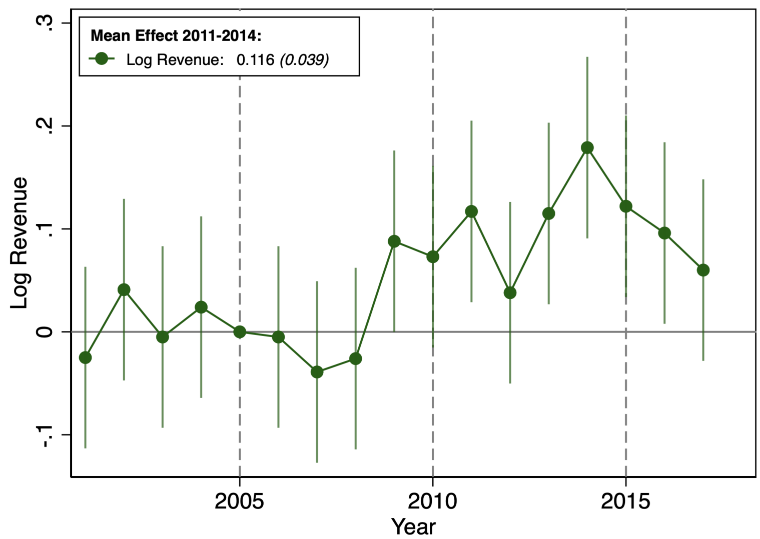

Panel B: Reporting Positive Accounting Profit

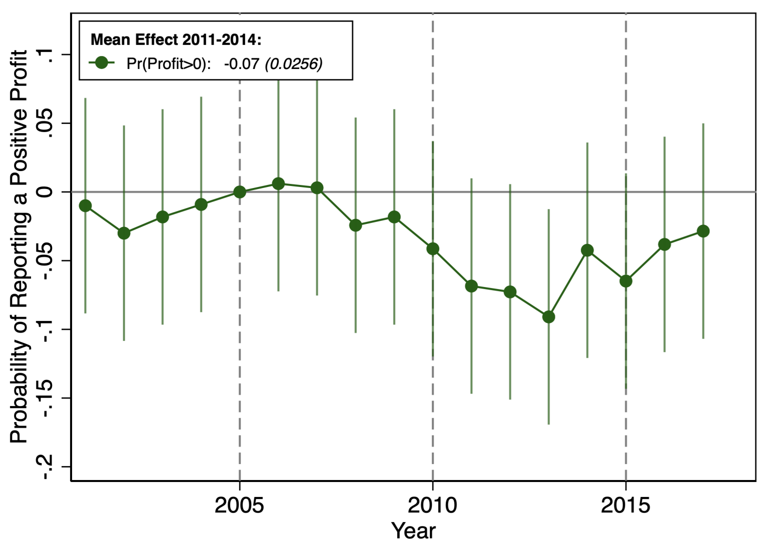

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on revenue and the probability of reporting a positive profit. The coefficients represent the effect in each year t of having an additional percent of the 2005 workforce treated (born 1950-1951), controlling for the share of old workers (born 1948-1951) at the firm. This corresponds to β_t from equation (3). The coefficient for $t = 2005$ is normalized to 0. In Panel A the outcome variable is log revenue. I windsorize revenue per worker at the 1st and 99th percentile to take out extreme outliers. In Panel B the outcome is the probability of reporting a positive profit in any given year. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. The legend reports the mean of the coefficients between 2011-2014: $\overline{\beta}_{2011-2014}$ and the corresponding standard error in parentheses. Standard errors are clustered at the firm level.

Figure 6: **Cash constraints: labor and investment adjustments**

Panel A: Younger Workers

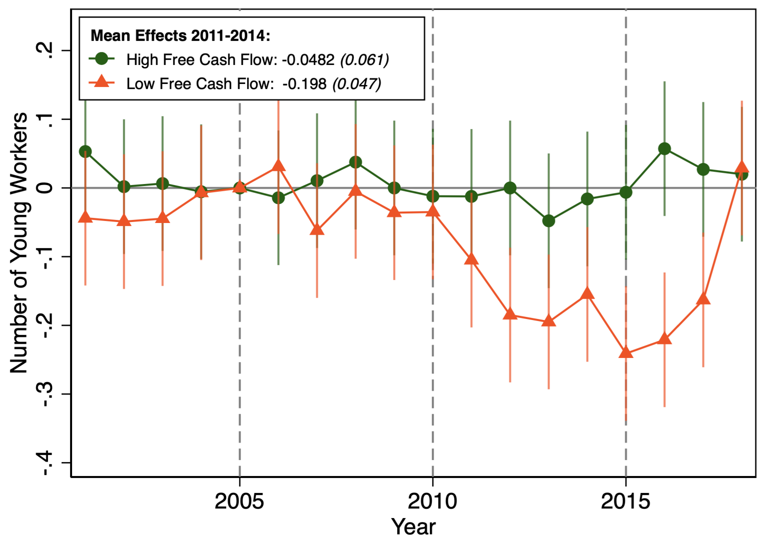

Panel B: Investment in Machines and Installations

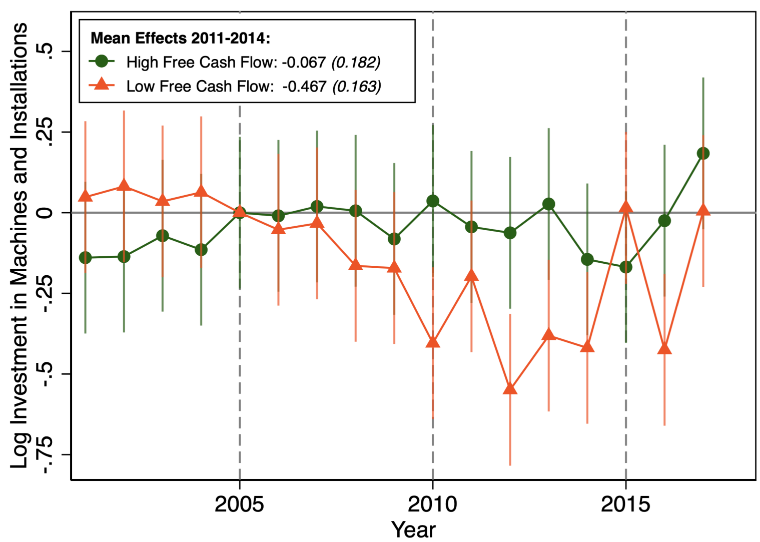

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on the number of younger workers and investment for high and low free cash flow firms, highlighting the role of cash constraints in firm adjustments. The coefficients in panel A represent the effect on the number of young workers in each year t of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. These correspond to β_t from equation (2). The coefficients in Panel B represent the effect on log investment in machines and installations in each year t of having an additional percent of the 2005 workforce treated (born 1950-1951), controlling for the share of old workers (born 1948-1951) at the firm. This corresponds to β_t from equation (3). The coefficient for $t = 2005$ is normalized to 0. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. The legend reports the mean of the coefficients between 2011-2014: $\overline{\beta}_{2011-2014}$ and the corresponding standard error in parentheses. Standard errors are clustered at the firm level.

Figure 7: **Cash constraints: total costs**

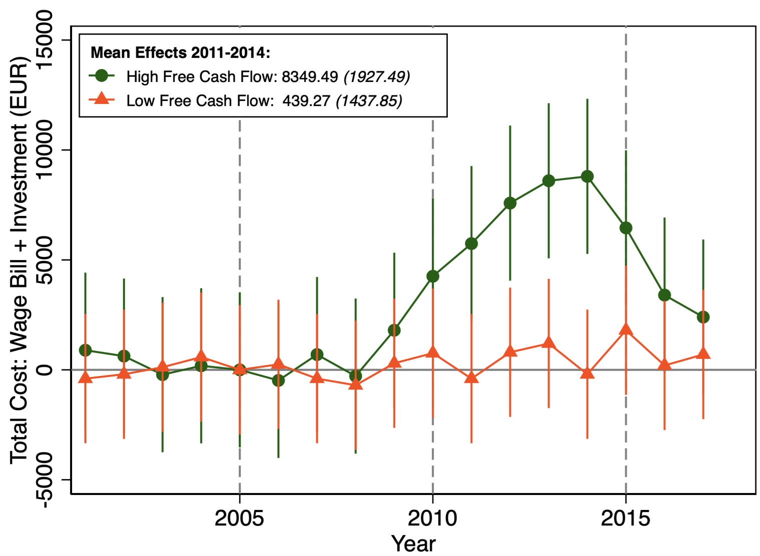

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on total firm cost (total wage bill + investments) for high and low free cash flow firms, highlighting that low free cash flow firms fully offset the additional older worker wage costs through reductions in the younger workforce and investments. The coefficients represent the effect on total firm labor and investment cost in each year t of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. These correspond to β_t from equation (2). The coefficient for $t = 2005$ is normalized to 0. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. The legend reports the mean of the coefficients between 2011-2014: $\overline{\beta}_{2011-2014}$ and the corresponding standard error in parentheses. Standard errors are clustered at the firm level.

Figure 8: **Cash constraints: firm production and profit**

Panel A: Revenue

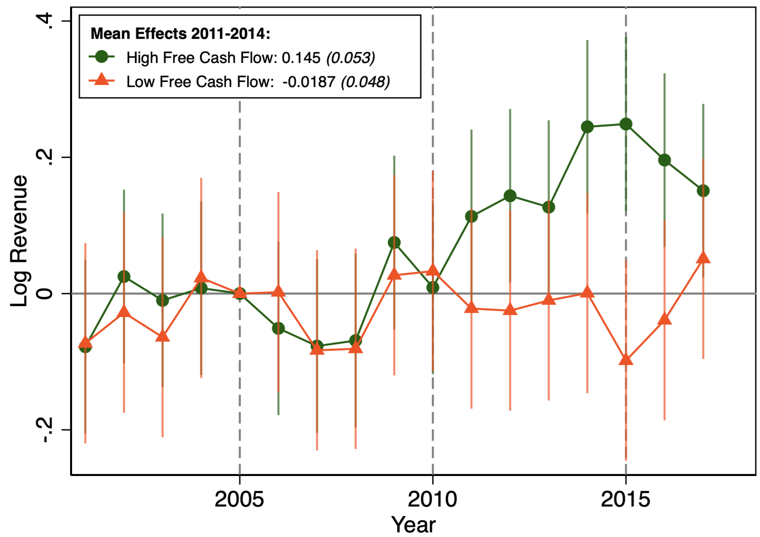

Panel B: Reporting Positive Accounting Profit

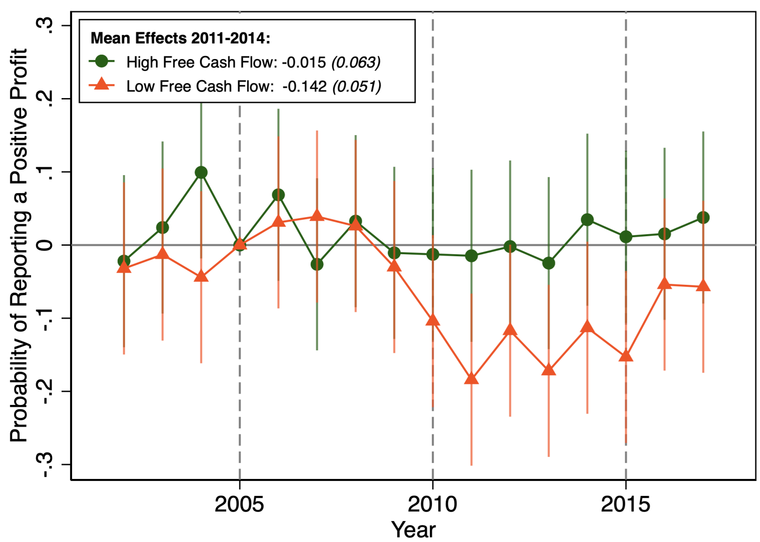

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on revenue and the probability of reporting a positive profit for high and low free cash flow firms, highlighting the role of cash constraints in firm adjustments. The coefficients represent the effect in each year t of having an additional percent of the 2005 workforce treated (born 1950-1951), controlling for the share of old workers (born 1948-1951) at the firm. This corresponds to β_t from equation (3). The coefficient for $t = 2005$ is normalized to 0. In Panel A the outcome variable is log revenue. I windsorize revenue per worker at the 1st and 99th percentile to take out extreme outliers. In Panel B the outcome is the probability of reporting a positive profit in any given year. In each case I run two separate regressions for high and below median free cash flow firms at baseline. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. The legend reports the mean of the coefficients between 2011-2014: $\overline{\beta}_{2011-2014}$ and the corresponding standard error in parentheses. Standard errors are clustered at the firm level.

Figure 9: **Interactions: Labor Adjustment Costs and Cash Constraints**

(a) **Young Worker Adjustments**

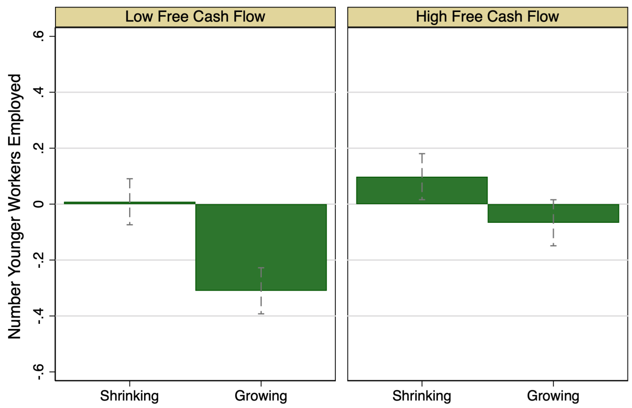

(b) **Investments in Machines and Installations**

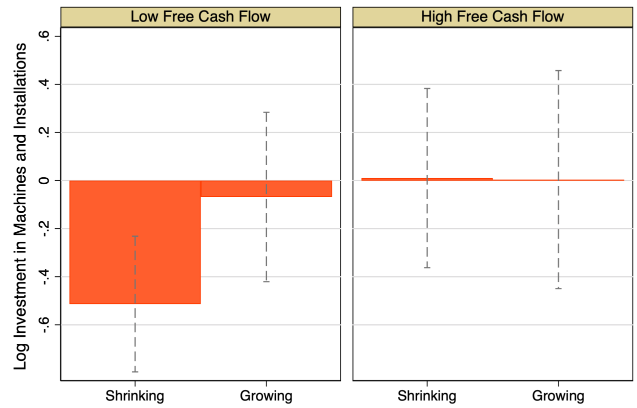

(c) **Revenue**

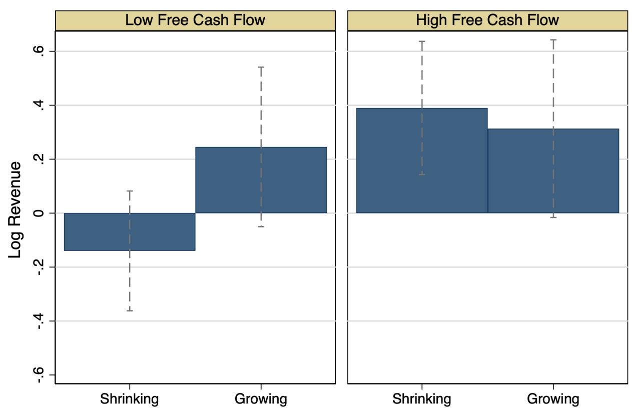

(d) **Reporting Positive Accounting Profit**

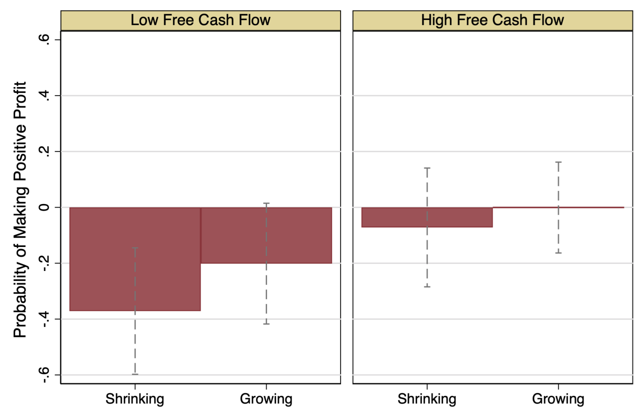

Notes: This figure demonstrates the effect of the policy on the number of younger workers, investments in machines and installations, revenue, and the probability of reporting a positive profit for the two way interaction between pre-policy free cash flow and employment growth. The reported effects represent the mean effect in 2011-2014, when the policy was most binding. The coefficients in panel A represent the average effect in 2010-2015 of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. The outcome is the number of younger workers employed. The coefficients in panel B-D represent the average effect in 2010-2015 of having an additional percent of the 2005 workforce treated (born 1950-1951), controlling for the share of old workers (born 1948-1951) at the firm. The outcomes are log investment in machines and installations (panel B), log revenue (panel C) and the probability of reporting a positive profit (panel D). Standard errors are clustered at the firm level.

Figure 10: **Anticipation: Young Worker Adjustments**

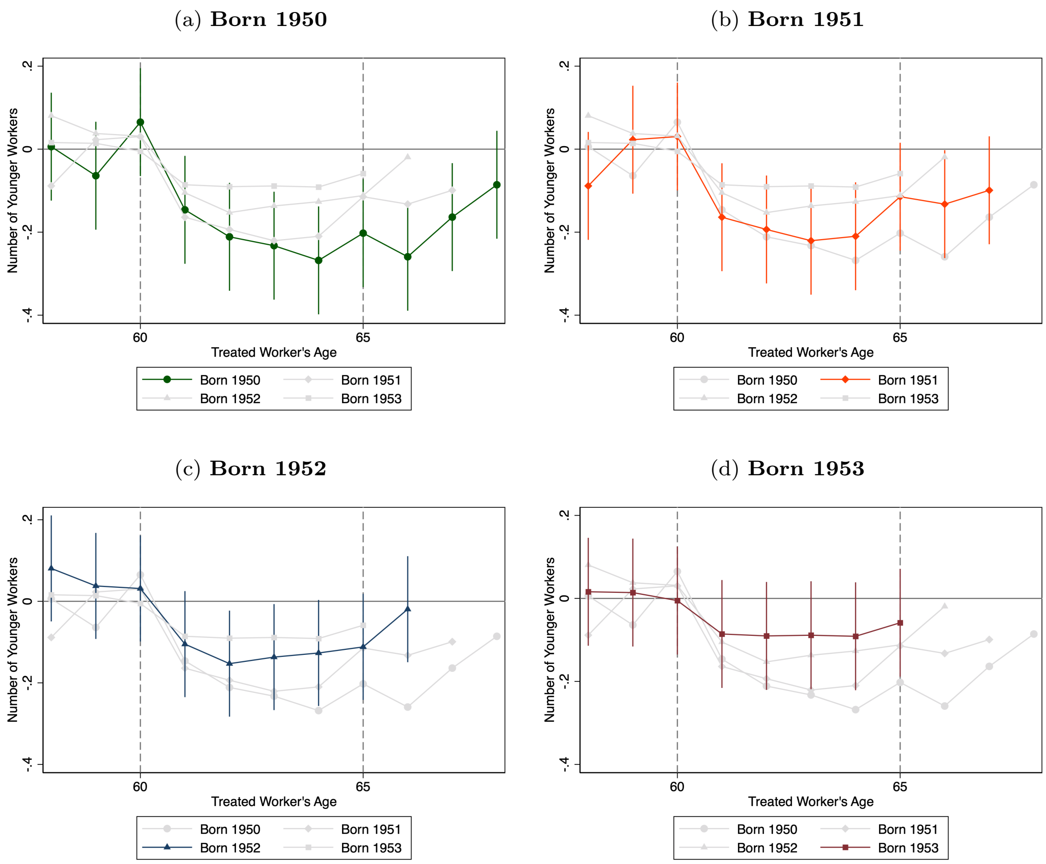

(a) **Born 1950** (b) **Born 1951** (c) **Born 1952** (d) **Born 1953**

Notes: The sample is restricted to firms with below median free cashflow at baseline that have exactly one worker born between 1949 and 1953. The figure demonstrates the effect of anticipation on the firm's labor adjustments. Each consecutive cohort c gives the firm an additional year to anticipate the policy change (see Appendix Figure A18 for the first stage plots by birth year). The coefficients represent the difference in the outcome of a firm employing a treated worker of cohort $c = \{1950, 1951, 1952, 1953\}$ relative to the outcome of a firm employing a control worker born in 1949, when both workers were of the same age a. These correspond to coefficients $\beta_{c,\,a}$ from equation (4), which represent firm adjustments in the number of younger workers depending on the treated worker's cohort c. The dotted lines represent the window within which a treated worker could have retired early absent the policy change, but is now more likely to work. Note that data are censored for younger cohorts. The last year in the data is 2018, for individuals born in 1953 are observed until they are 65 years old. Standard errors are clustered at the firm level.

Figure 11: **Anticipation: Profit Effects**

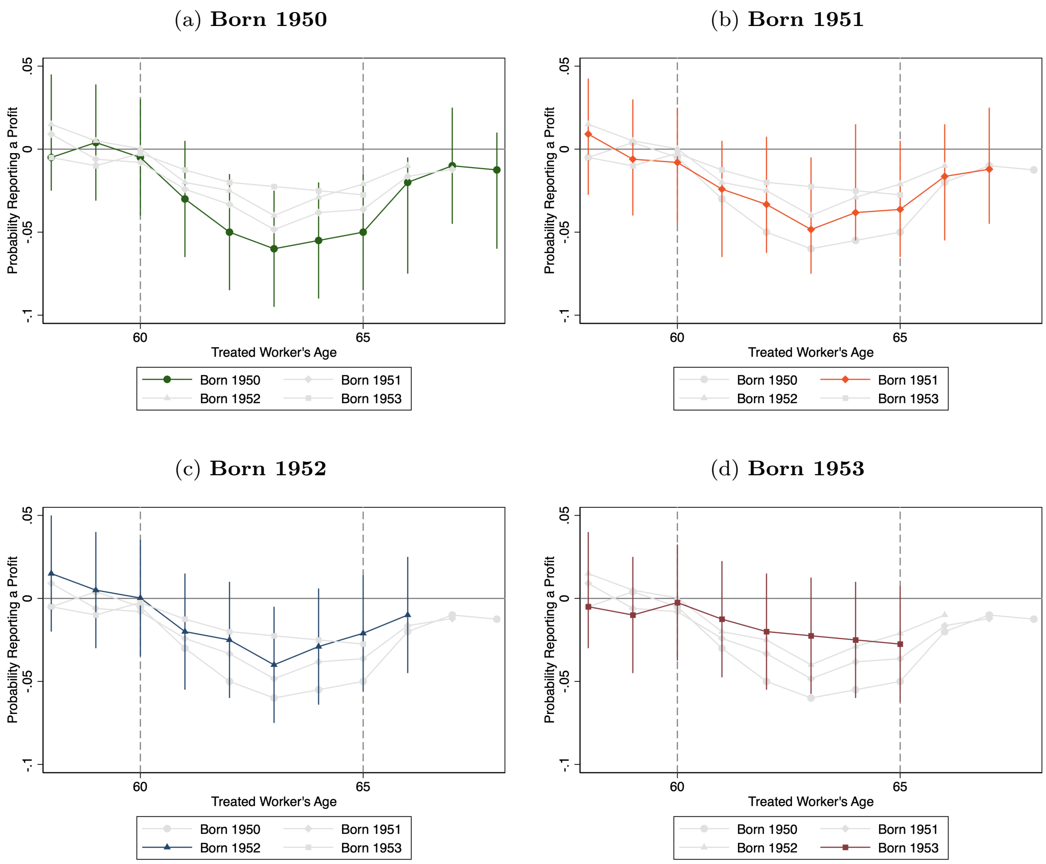

Notes: The sample is restricted to firms with below median free cashflow at baseline that have exactly one worker born between 1949 and 1953. The figure demonstrates the effect of anticipation on the firm's probability of reporting a positive profit. Each consecutive cohort c gives the firm an additional year to anticipate the policy change (see Appendix Figure A18 for the first stage plots by birth year). The coefficients represent the difference in the outcome of a firm employing a treated worker of cohort $c = \{1950, 1951, 1952, 1953\}$ relative to the outcome of a firm employing a control worker born in 1949, when both workers were of the same age a. These correspond to coefficients $\beta_{c,\,a}$ from equation (4), which represent the firm's probability of reporting a positive profit depending on the treated worker's cohort c. The dotted lines represent the window within which a treated worker could have retired early absent the policy change, but is now more likely to work. Note that data are censored for younger cohorts. The last year in the data is 2018, for individuals born in 1953 are observed until they are 65 years old. Standard errors are clustered at the firm level.

Appendix A: Additional Tables and Figures

Table A1: **Balance Table: High & Low Free Cash Flow Firms**

	(1)	(2)	
	High Cash Flow	Low Cash Flow	P-Value: (1)=(2)
Fulltime Workers	13.16	12.31	0.00
Fulltime Labor Costs (1000 EUR)	308.87	263.09	0.00
Other Input Costs (1000 EUR)	627.49	517.56	0.00
Material Assets	474.61	403.45	0.00
Share founded before 2001	0.21	0.18	0.00
Share of jobs in manufacturing	0.227	0.245	0.02
Share of jobs in services	0.297	0.298	0.29
N	73,566	71,285	

Note: This table shows summary statistics for firms that had high and low free cash flow prior to the policy change. High and low free cash flow firms are defined as having above and below median free cash flow in 2001-2004.

Figure A1: **Histogram of worker birth year around 1950 cut-off employed at firms**

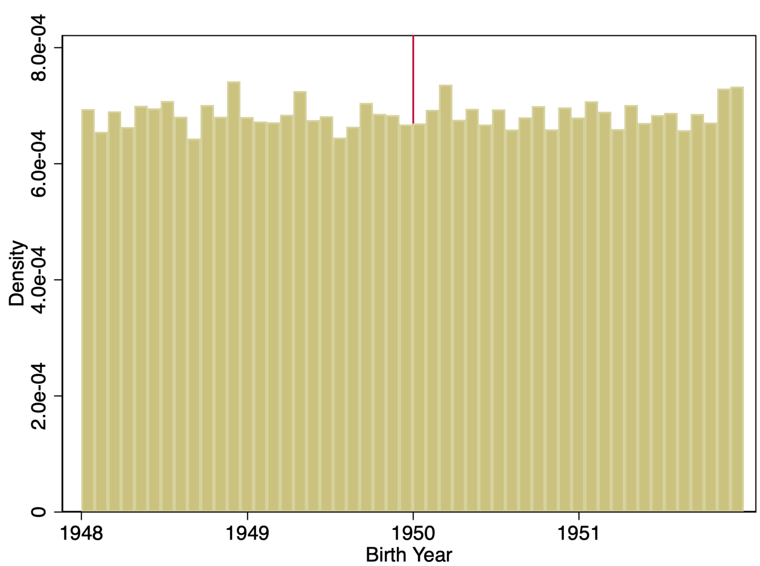

Notes: This figure shows there is no excess mass on either side of the 1950 cohort cutoff among workers employed at firms in 2005.

Figure A2: **Falsification Checks: Covariate Balance – Simple Firm RD**

Panel A: Employment of Older Workers Prior to Policy (2000-2005)

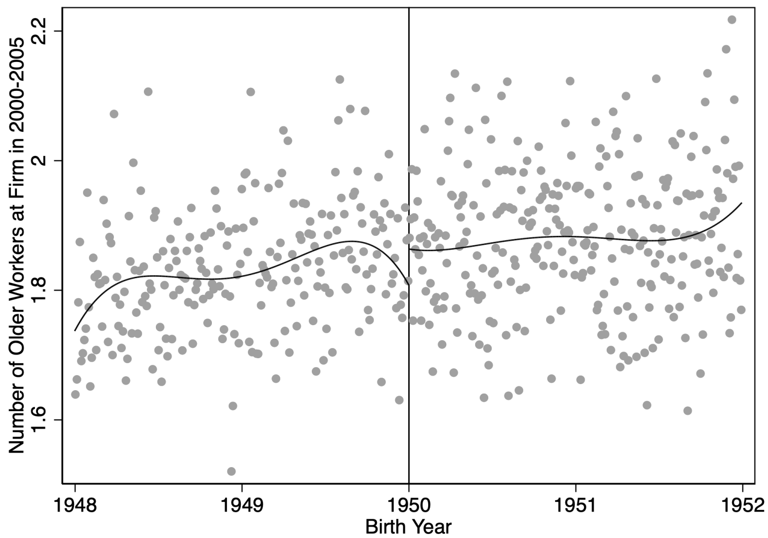

Panel B: Firm Revenue (in 1000s of EUR) Prior to Policy (2000-2005)

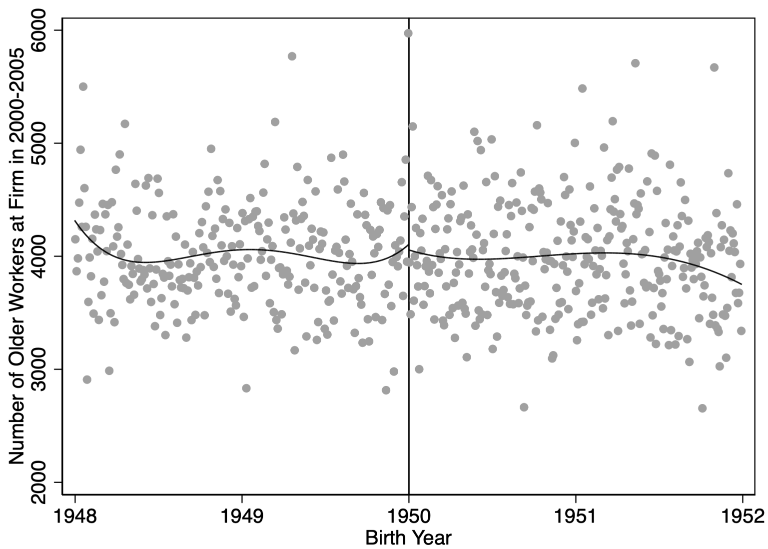

Notes: The unit of observation is the firm. I limit to firms with exactly one worker born in the 1948-1952 cohort window. The outcome variables are the mean number of older workers and mean revenue at the firm in 2000-2005, before the policy change had an effect. The figure demonstrates that there is no discontinuity in these pre-policy covariates at the threshold. All workers considered are full time employed.

Figure A3: **Firm-level Treatment: Share Born After Cutoff**

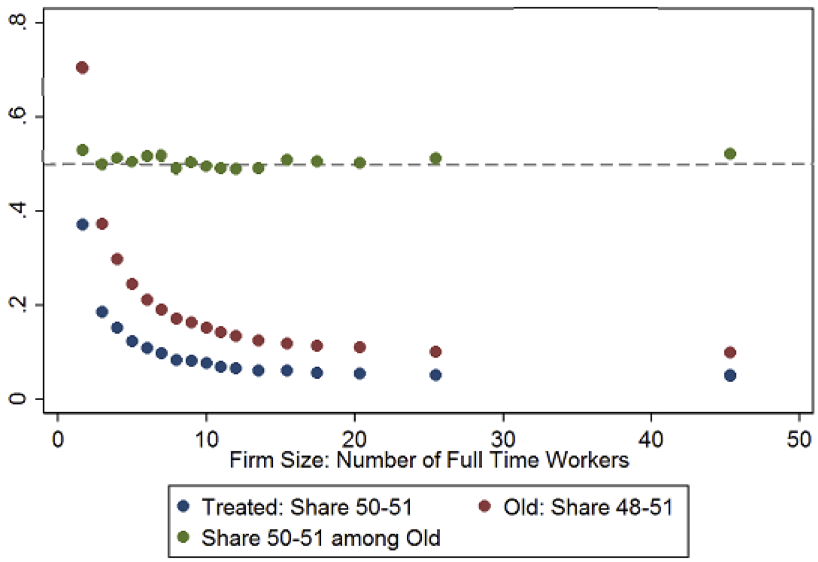

Notes: This figure shows the share of all 2005 full-time workers (1) born 1948-1951 (older workers), (2) born 1950 or 1951 (treated workers) by firm size. In addition, it shows the share of treated workers among old workers. The key observation from this figure is that the share of treated workers (born 1950-1951) among older workers (born 1948-1951) is consistently 0.5 for firms of different sizes. This suggests there is no excess mass of workers born either before or after the 1950 cohort boundary.

Figure A4: **Firm-level Distribution of % Treated Workers**

Panel A: Share of Treated Workers as % of Total Workforce

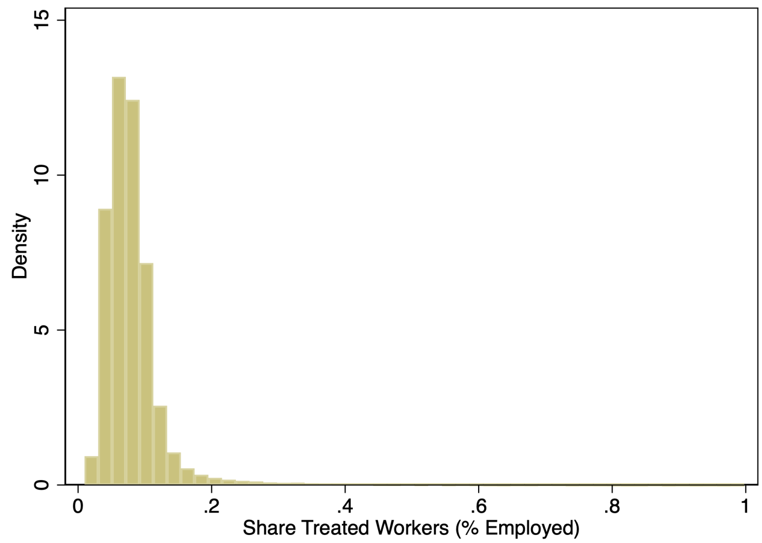

Panel B: Share Wages for Treated Workers as % of Total Wage Bill

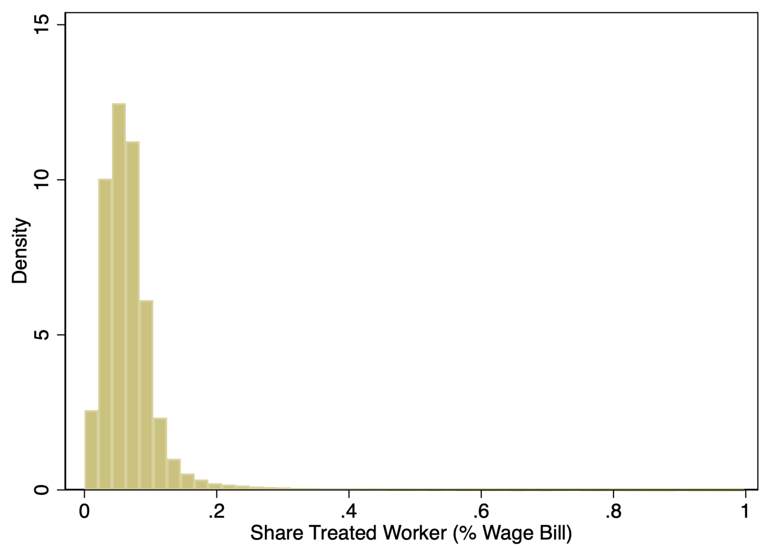

Notes: Limited to firms with at least one treated worker.

Figure A5: **Falsification: All Untreated or Treated**

Panel A: All Untreated (1946-1949, imaginary cutoff at 1948)

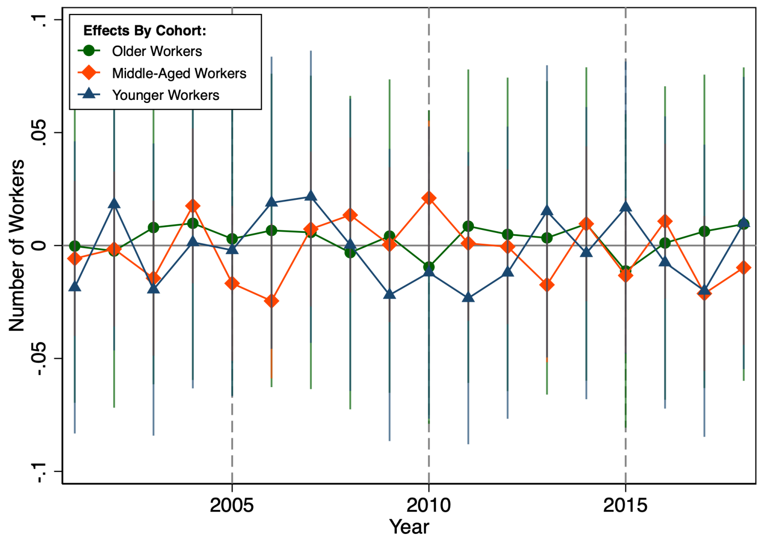

Panel A: All Treated (1950-1953, imaginary cutoff at 1952)

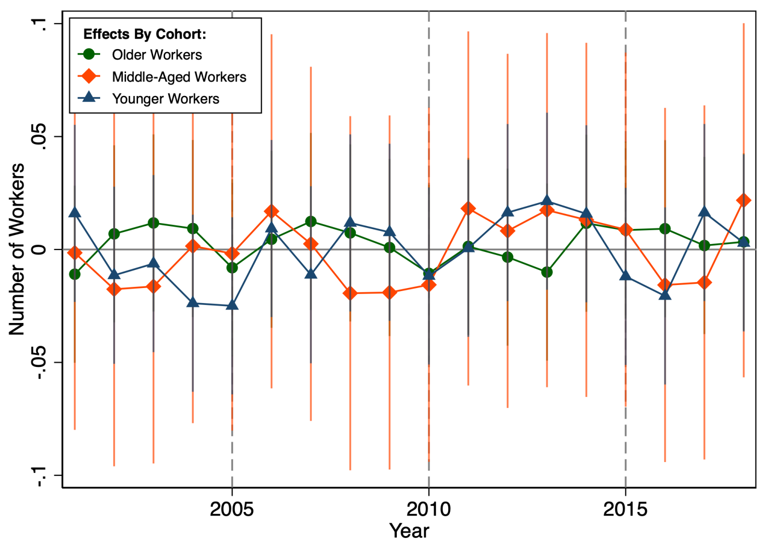

Notes: The unit of observation is the firm-year. Panel A is a falsification check of all untreated workers, that plots coefficients β_t from equation (2), which represents the effect in each year t of having an additional placebo treated worker (born 1948 or 1949) at the firm in 2005, controlling for the number of old workers (born 1946-1949) and the total number of workers. The outcome variable is the number of older workers. Panel B is a falsification check of all treated workers that plots coefficients β_t from equation (2), which represents the effect in each year t of having an additional placebo treated worker (born 1952 or 1953) at the firm in 2005, controlling for the number of old workers (born 1950-1953) and the total number of workers. The outcome variable is the number of older workers. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure A6: **Overall firm hiring effects**

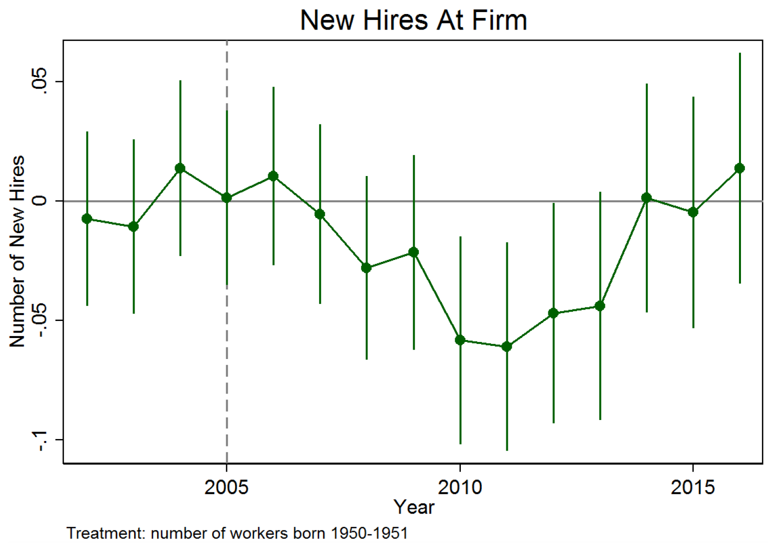

Notes: The unit of observation is the firm-year. The figure coefficients β_t from equation (2), which represents the effect in each year t of having an additional percent of the 2005 workforce treated, controlling for the share of old workers at the firm. The coefficient for $t = 2005$ is normalized to 0. The outcome variable is the number of new hires at the firm. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure A7: **Overall changes in hours worked and hourly wages - young and middle-aged**

Panel A: Hours Worked

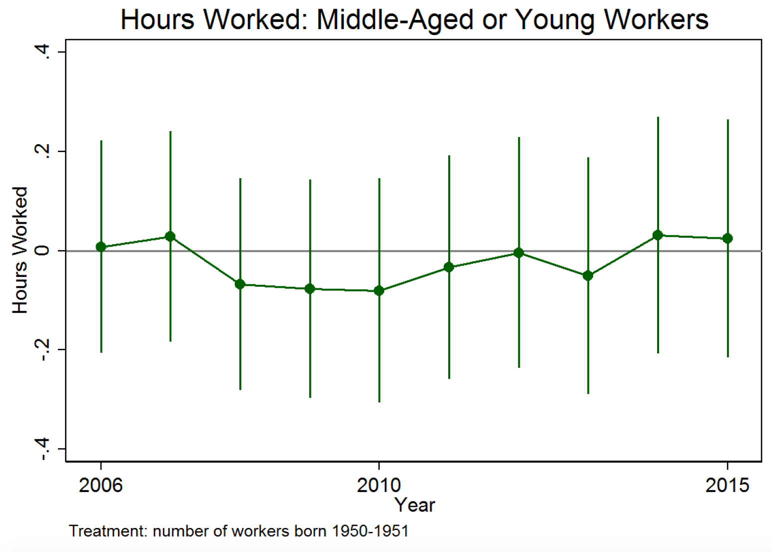

Panel B: Hourly Wages

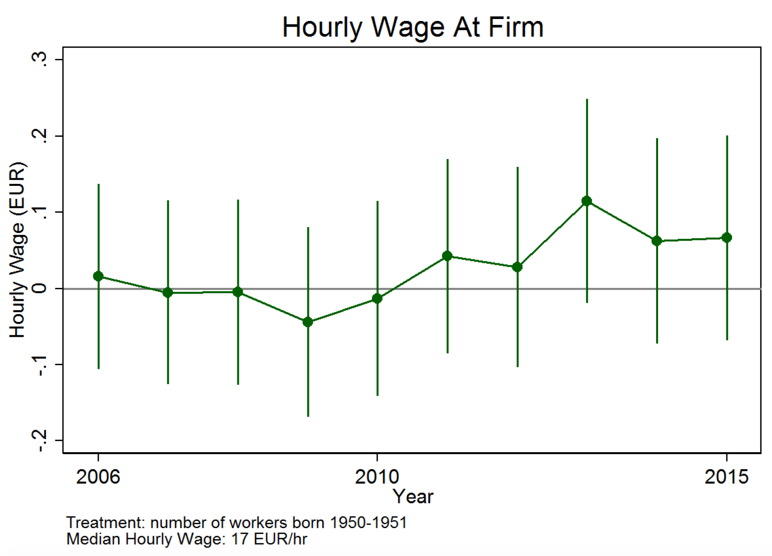

Notes: The unit of observation is the firm-year. The figure coefficients β_t from equation (2), which represents the effect in each year t of having an additional percent of the 2005 workforce treated, controlling for the share of old workers at the firm. The coefficient for $t = 2005$ is normalized to 0. The outcome variable is the hours worked and hourly wages at the firm. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure A8: **Overall firm material investment effects**

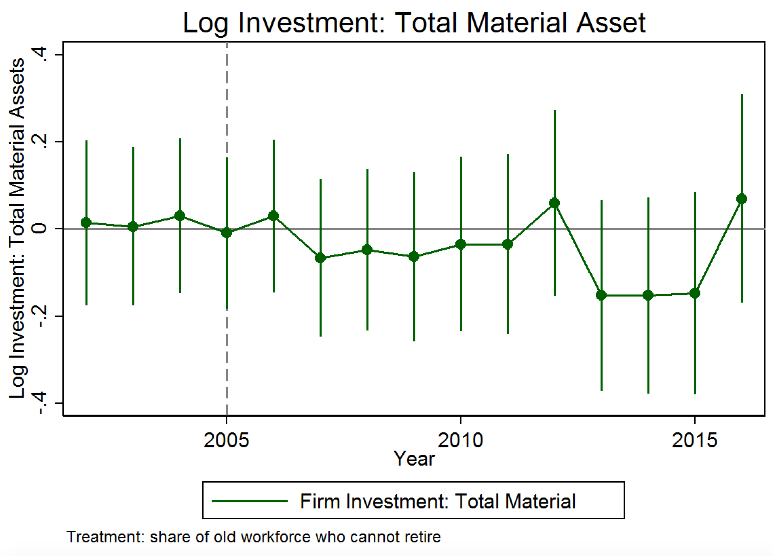

Notes: The unit of observation is the firm-year. The figure coefficients β_t from equation (3), which represents the effect in each year t of having an additional percent of the 2005 workforce treated, controlling for the share of old workers at the firm. The coefficient for $t = 2005$ is normalized to 0. The outcome variable is log investment in total material assets. Standard errors are clustered at the firm level.

Figure A9: **Overall firm revenue per worker effects**

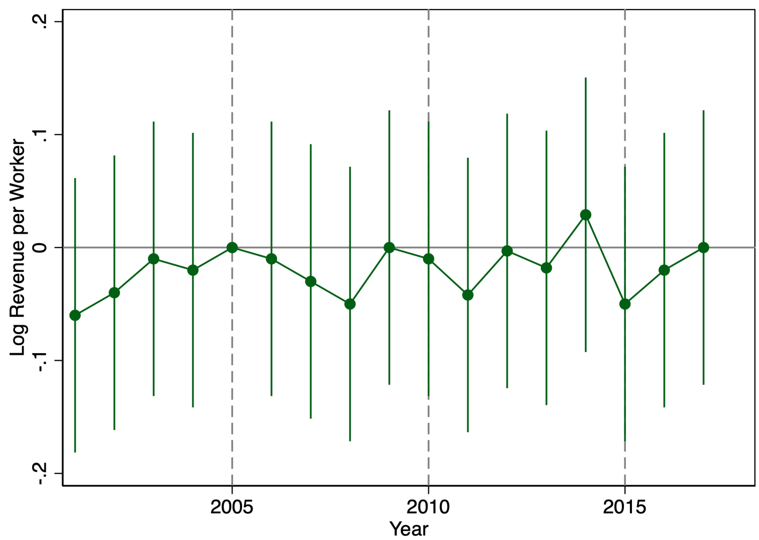

Notes: The unit of observation is the firm-year. The figure coefficients β_t from equation (3), which represents the effect in each year t of having an additional percent of the 2005 workforce treated, controlling for the share of old workers at the firm. The coefficient for $t = 2005$ is normalized to 0. The outcome variable is log revenue per worker. I windsorize revenue per worker at the 1st and 99th percentile to take out extreme outliers. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure A10: **Overall profit effects**

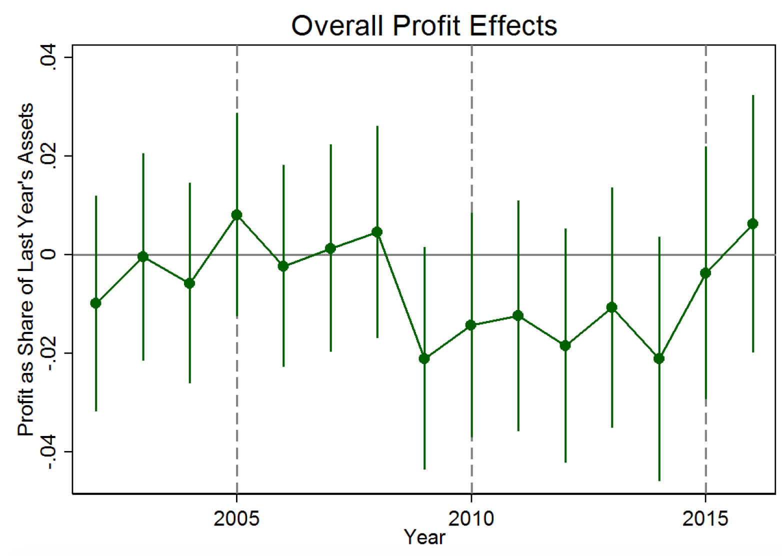

Notes: The unit of observation is the firm-year. The figure coefficients β_t from equation (3), which represents the effect in each year t of having an additional percent of the 2005 workforce treated, controlling for the share of old workers at the firm. The coefficient for $t = 2005$ is normalized to 0. The outcome variable is profits as a share of last year's assets. I windsorize this measure at the 1st and 99th percentile to take out extreme outliers. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure A11: **Overall effects on firm shutdown**

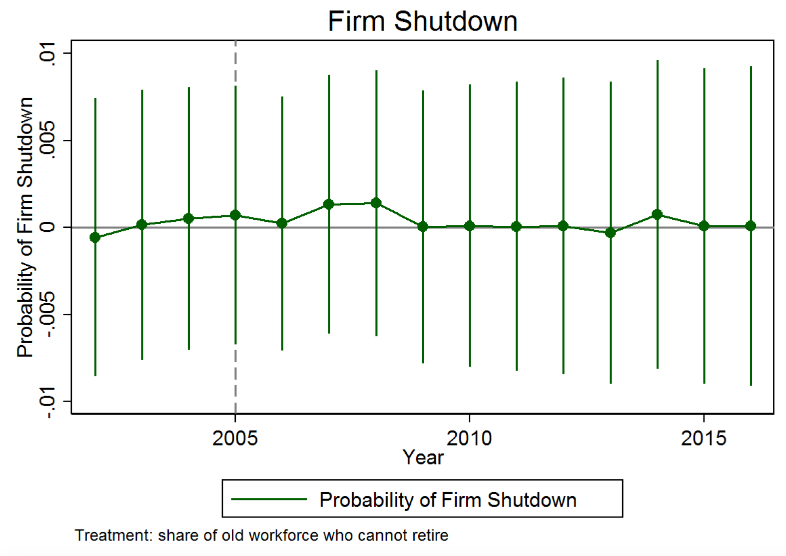

Notes: The unit of observation is the firm-year. The figure coefficients β_t from equation (3), which represents the effect in each year t of having an additional percent of the 2005 workforce treated, controlling for the share of old workers at the firm. The coefficient for $t = 2005$ is normalized to 0. Tthe outcome variable is firm shutdown. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure A12: **First Stage by FCF**

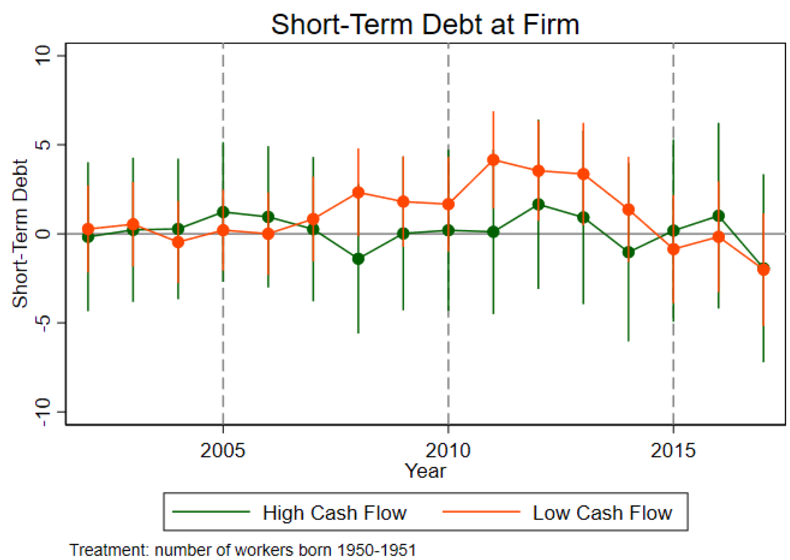

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on short-term debt levels at the firm, depending on whether firms had above or below median free cash flow prior to the policy. The coefficients plotted represent the effect in each year t of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. These correspond to β_t from equation (2). The coefficient for $t = 2005$ is normalized to 0. Standard errors are clustered at the firm level.

Figure A13: **First Stage by FCF**

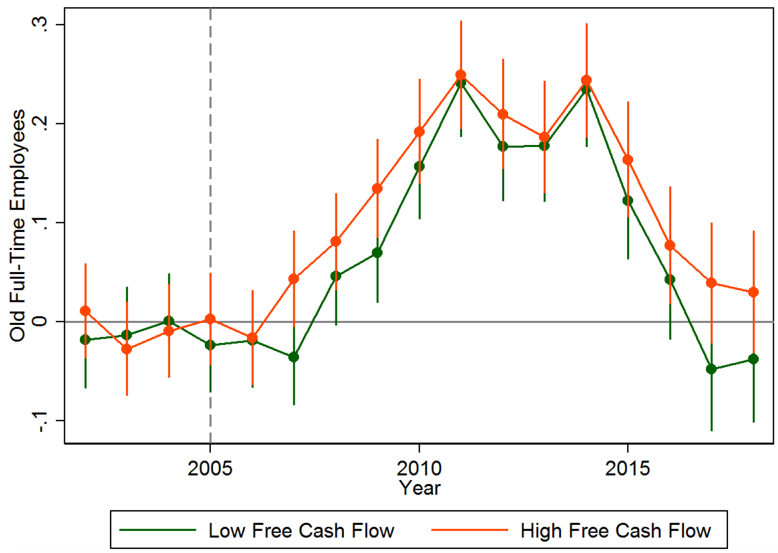

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on the number of older workers at the firm, depending on whether firms had above or below median free cash flow prior to the policy. The coefficients plotted represent the effect in each year t of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. These correspond to β_t from equation (2). The coefficient for $t = 2005$ is normalized to 0. Standard errors are clustered at the firm level.

Figure A14: **Young Worker Adjustments - By Baseline Productivity**

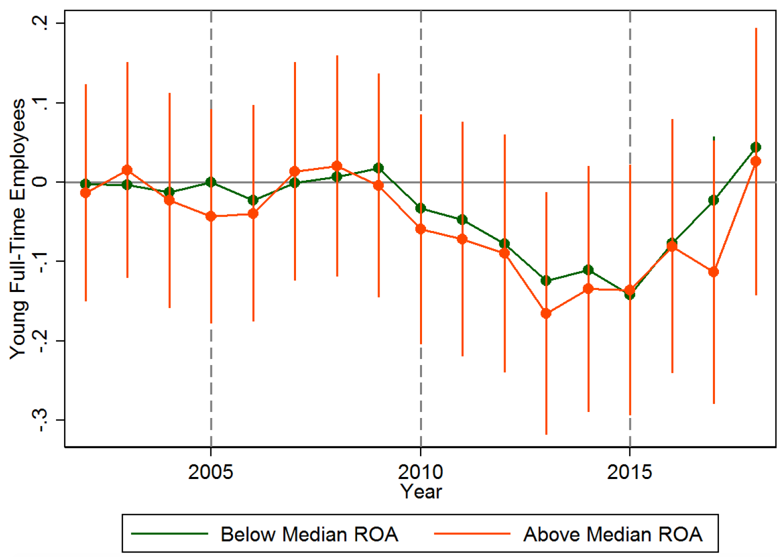

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on the number of younger workers at the firm, depending on whether firms had above or below median return on assets prior to the policy. The return on assets measure is residualized with respect to the firm's industry mean. The coefficients plotted represent the effect in each year t of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. These correspond to β_t from equation (2). The coefficient for $t = 2005$ is normalized to 0. Standard errors are clustered at the firm level.

Figure A15: **Young Worker Adjustments - By Pre-Policy Bank Funding Share**

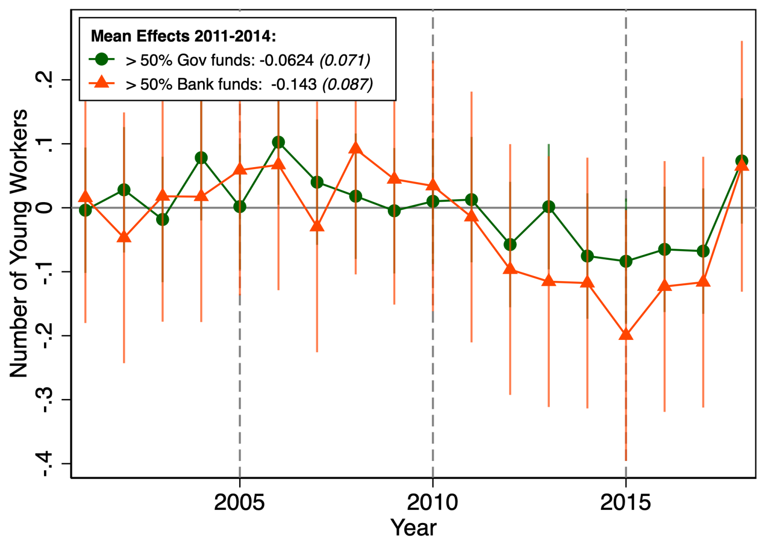

Notes: The unit of observation is the firm-year. The figure demonstrates the effect of the policy on the number of younger workers at the firm, depending on whether firms had above or below median return on assets prior to the policy. The analysis is run separately for firms that relied primarily on government funds vs. bank funds prior to the policy. The coefficients plotted represent the effect in each year t of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. These correspond to β_t from equation (2). The coefficient for $t = 2005$ is normalized to 0. Standard errors are clustered at the firm level.

Figure A16: **Interaction: Labor Adjustment Costs and FCF**

Panel A: Employment of Younger Workers

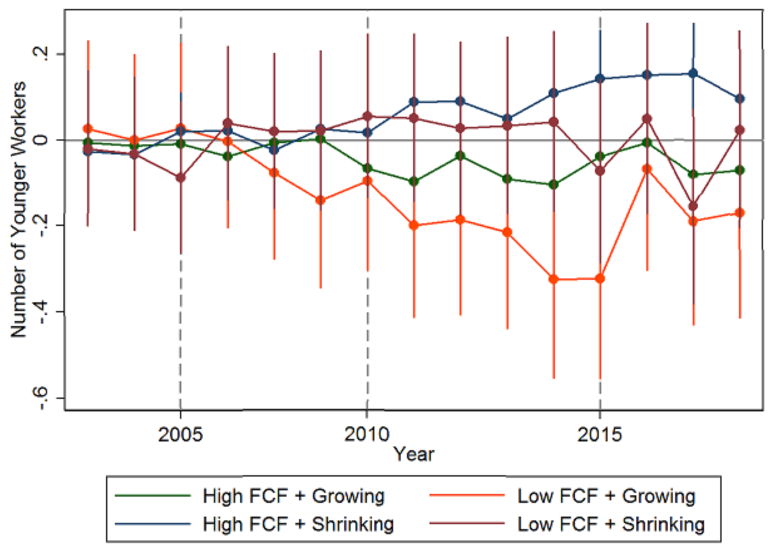

Panel B: Log Investment in Machines and Installations

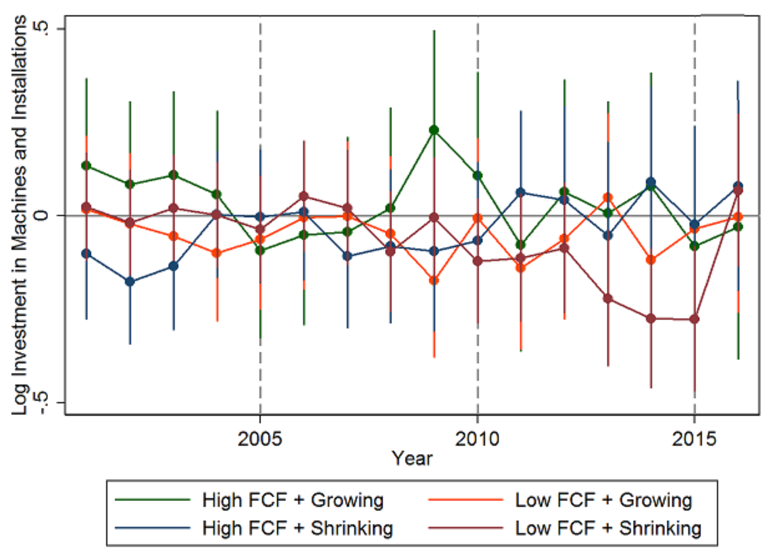

Notes: The unit of observation is the firm-year. Panel A plots coefficients β_t from equation (2), which represents the effect in each year t of having an additional treated worker at the firm in 2005, controlling for the number of old workers and the total number of workers. The outcome variable is the number of young workers. Panel B plots coefficients β_t from equation (3), which represents the effect in each year t of having an additional percent of the 2005 workforce treated, controlling for the share of old workers at the firm. The outcome variable is log investment in machines and equipment. The coefficient for $t = 2005$ is normalized to 0. In each case I run four separate regressions for the two-way interaction between high and below median free cash flow firms at baseline and growing/shrinking firms. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure A17: **Interaction: Labor Adjustment Costs and FCF**

Panel A: Log Revenue

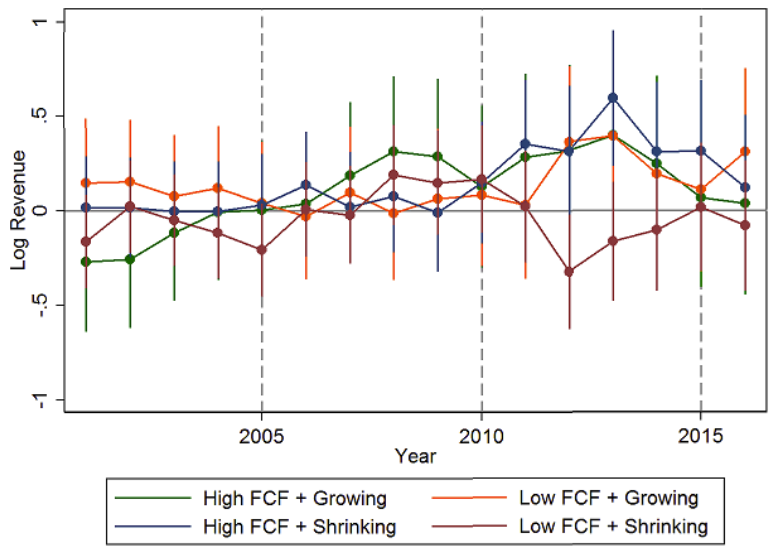

Panel B: Reporting Positive Accounting Profit

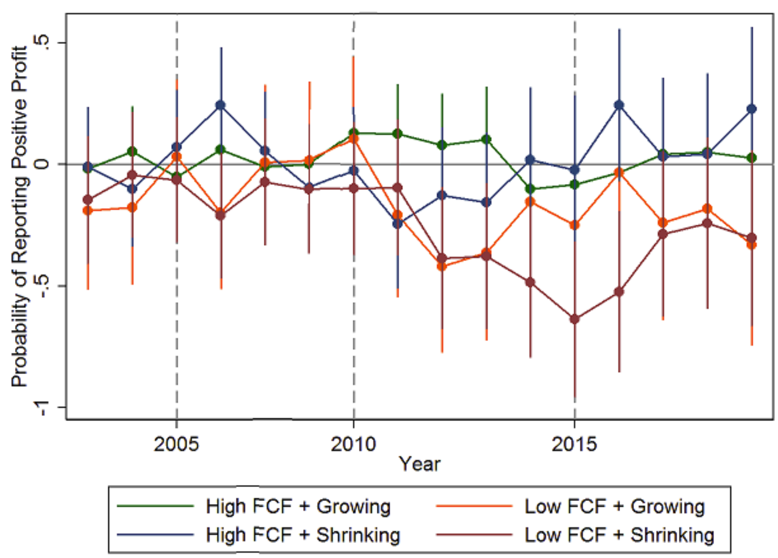

Notes: The unit of observation is the firm-year. The figure plots coefficients β_t from equation (3), which represents the effect in each year t of having an additional percent of the 2005 workforce treated, controlling for the share of old workers at the firm. The coefficient for t = 2005 is normalized to 0. In Panel A the outcome variable is log revenue. I windsorize revenue per worker at the 1st and 99th percentile to take out extreme outliers. In Panel B the outcome is the probability of reporting a positive profit in any given year. In each case I run four separate regressions for the two-way interaction between high and below median free cash flow firms at baseline and growing/shrinking firms. The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

Figure A18: **Single Worker Level: Anticipation First Stage**

Panel A: By Calendar Year

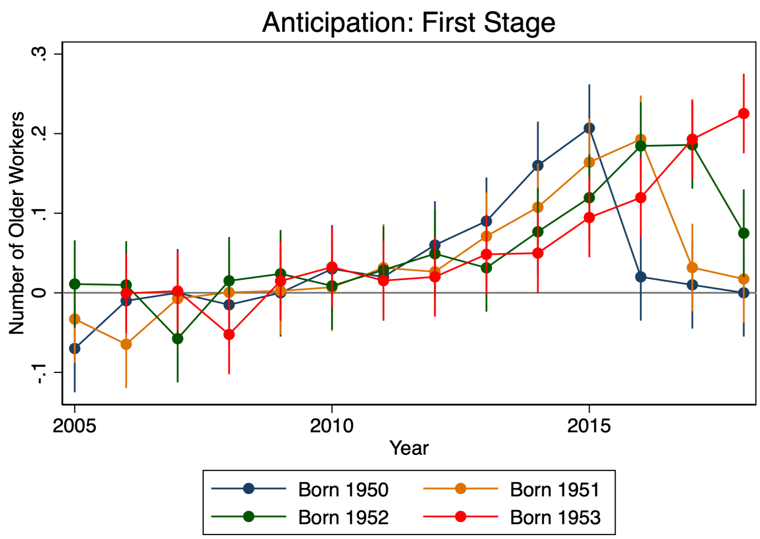

Panel B: By Age of the Treated Worker

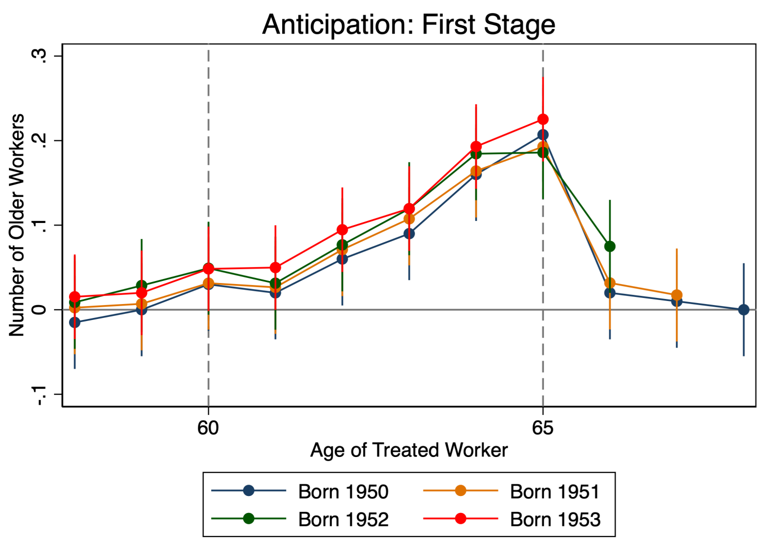

The sample is restricted to firms with below median free cashflow at baseline that have exactly one worker born between 1949 and 1953. The figure shows the first stage effect of anticipation on the firm's retention of older workers over time. Each consecutive cohort c gives the firm an additional year to anticipate the policy change. The coefficients represent the difference in the number of older workers of a firm employing a treated worker of cohort $c = \{1950, 1951, 1952, 1953\}$ relative to the outcome of a firm employing a control worker born in 1949. Panel A shows the first stage by calendar year, demonstrating that the onset of the treatment is shifted by a year for each consecutive cohort c. Panel B shows the first stage by age of the treated worker., demonstrating that the magnitude of the first stage is similar. Note that data are censored for younger cohorts. The last year in the data is 2018, for individuals born in 1953 are observed until they are 65 years old. Standard errors are clustered at the firm level.

Appendix B: Background on Dutch Pension System

B1. Dutch Pension System

The Dutch pension system consists of three pillars.

The first pillar is a public old age pension provided by the government to all residents of the Netherlands when they reach the legal retirement age of 65. This is a pay-as-you-go system financed through income taxation which provides pension benefits that are tagged to the minimum wage.

The second pillar of the pension system consists of employer-employee pensions. Pensions agreements are negotiated between unions and employers at the sector or firm level and are set in collective agreements. While membership of a pension fund is not mandatory by law, 90% of Dutch employees belong to a pension fund. Prior to 2005, workers were able to retire early (before the legal retirement age of 65) through the second pillar of the pension system. These early retirement schemes were tax deductible, and the tax advantage was approximately 25% of early retirement benefits. Typically, before 2005 workers could retire in their early 60s and get a large share of their final earned wage as pension benefits. For example, in the public sector a worker who had served for 40 years in the public sector could retire at the age of 62 and three months at a replacement rate of 70%. As a result, early retirement, was the social norm in the Netherlands.

The third pillar consists of individual pensions, used mainly by self-employed individuals or as supplemental pension. This pillar is relatively insignificant in the Netherlands. By comparison, the retirement benefits belonging to the second pillar are about twenty times as large as retirement benefits belonging to the third pillar (Bovenberg and Gradus, 2015).

B2. 2005 Early Retirement Reform

The early retirement scheme was first introduced in the 1980s. The scheme consisted of two main parts. The first was a pay as you go scheme ("VUT") in which employees pay for current early retirees. The second was an employer-employee savings plan ("pre-pension") which enabled workers to save for early retirement during their employment. Together, the schemes gave individuals the opportunity to retire in their early 60s, rather than at the legal retirement age of 65. The exact early retirement age varied by sector and industry.

The scheme was jointly funded by employers and employees. Contributions were a share of wages earned by employees, typically 8%. Employees and employers contributed 3.5% of wages to fund the pay as you go scheme and 4.5% of wages to fund the employer-employee savings plan. Approximately 60% of these contributions came from employers, and 40% from employees. There was a fiscal advantage to these schemes as well: contributions to the early retirement scheme were tax deductible, meaning that the government was an important funder of early retirement.

Aside from increasing the effective retirement age, abolishing the early retirement scheme also had financial consequences for both affected workers and firms. These effects are small relative to the primary employment effect I exploit in this paper, however.

B3. Other financial consequences of the policy for workers and firms

Aside from increasing the effective retirement age, abolishing the early retirement scheme also had financial consequences for both affected workers and firms. These effects are small relative to the primary employment effect, however.

First, affected individuals no longer had to make contributions to the employer-employee savings scheme, which accounted for approximately 1.8% of their wage. Until 2014 these cohorts still made contributions to the pay-as-you go component to fund current early retirees.[26] Second, the affected workers got to use any benefits they accumulated up to 2005 in the employer-employee savings plan for early retirement or to increase the benefits in their old age pension. Such contributions were no longer tax deductible, however. In practice this likely contributed to the fact that some affected workers still retired before the legal retirement age of 65, but not at as high a rate as they would have absent the policy change. The difference in post-retirement income is negligible, however. The tax data - described below - shows that the difference in mean income from pensions after retirement between individuals born in 1950 and 1949 is less than 200 euros (220 USD) annually (less than 1% of mean annual pension income).

Second, affected firms faced a small reduction in the early retirement benefits paid for by employers. Firms no longer had to contribute to the employer-employee savings scheme, which accounted for approximately 2.7 percent of each affected worker's wage. Similar to workers, firms continued to contribute to the pay as you go component until 2014. These reductions in contributions were small relative to the wages firms continued to pay when retaining an older worker.

Appendix C: Empirical Approach

C1. Robustness to window size

[26]Note that 2014 is the year in which the last eligible cohort, born in 1949, reached the statutory pension age of 65.

Figure A19: **Robustness: window size around 1950**

Panel A: 1949 vs. 1950

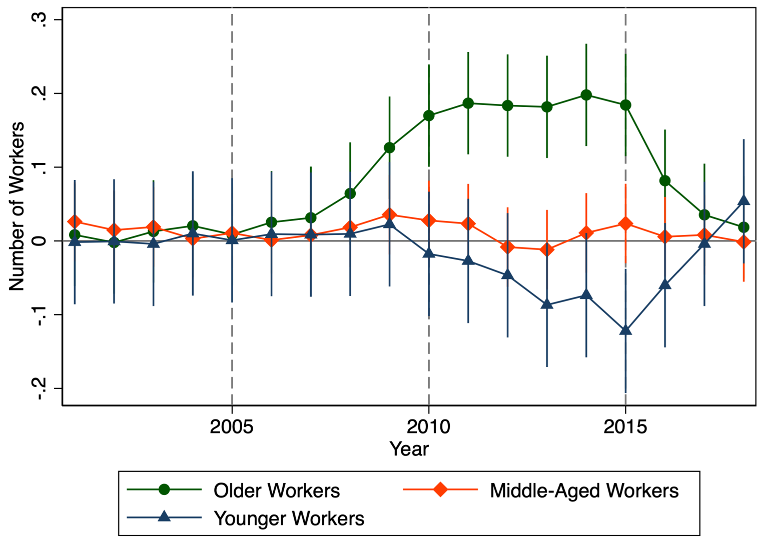

Panel B: 1947-1949 vs. 1950-1952

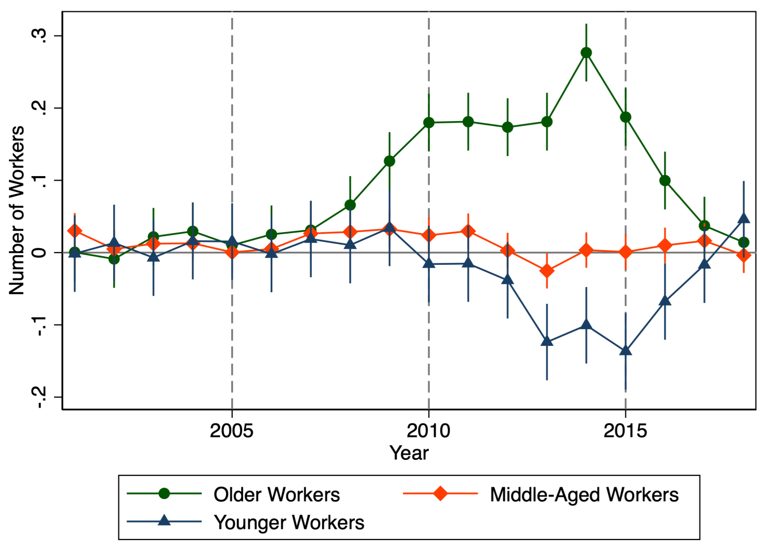

Notes: The unit of observation is the firm-year. The figures demonstrate robustness to adjusting the window size around 1950 which are used for identification. The main analysis takes a two year window on either side and compares the effect of having additional workers born in 1950/1951 (treated) to workers born in 1948/1949 (control). Panel A replicates this using only 1950 versus 1949 (one year window) and Panel B using 1950-1952 versus 1947-1949 (three year window). The coefficients plotted represent the effect in each year t of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. These correspond to β_t from equation (7). The coefficient for $t = 2005$ is normalized to 0. In each case the three lines represent three separate regressions. The outcome variables are the number of older workers (born 1945-1955), the number of middle-aged workers (born 1955-1965) and the number of younger workers (born 1965 and after) at the firm over time. All workers considered are full time employed.The dotted lines represent key timing of the policy: the policy was announced in 2005 and had its main impact in 2010-2015. Standard errors are clustered at the firm level.

C2. Labor and Production Specification: Levels and Shares

As specified in section **??**, the exact empirical specification differs slightly depending on the outcome studied. The primary reason for this is that I study firms of different sizes. We might expect that the units in which effects are homogenous across different firms might be different depending on the outcome.

For labor adjustments I study the effects in terms of the *number* of treated workers employed at a firm. Figure A20 shows that this approach yields roughly the same treatment effect across firms of different sizes.

Figure A20: **First Stage by Firm Size Decile**

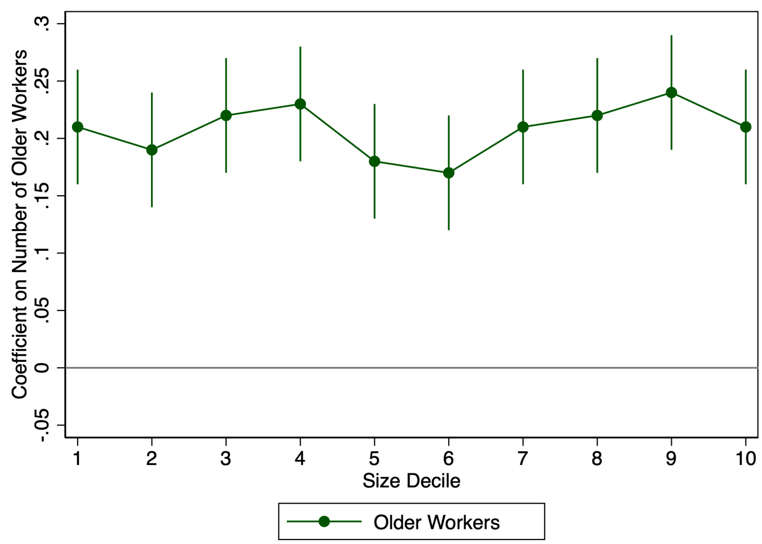

Notes: The figure demonstrates that the treatment effect of the policy in terms of the number of older workers retained is homogenous in the size of firms. The coefficients plotted represent the mean effect on the number of older workers in 2011-2014 of having an additional treated worker (born 1950 or 1951) at the firm in 2005, controlling for the number of old workers (born 1948-1951) and the total number of workers. Standard errors are clustered at the firm level.

Second, I examine the effect of having an additional *percent* of the workforce treated for firm investment, production and profits. Similarly, Figure A21 shows that this approach yields similar treatment effects across firms of different sizes.

Figure A21: **Revenue Effects by Firm Size Decile**

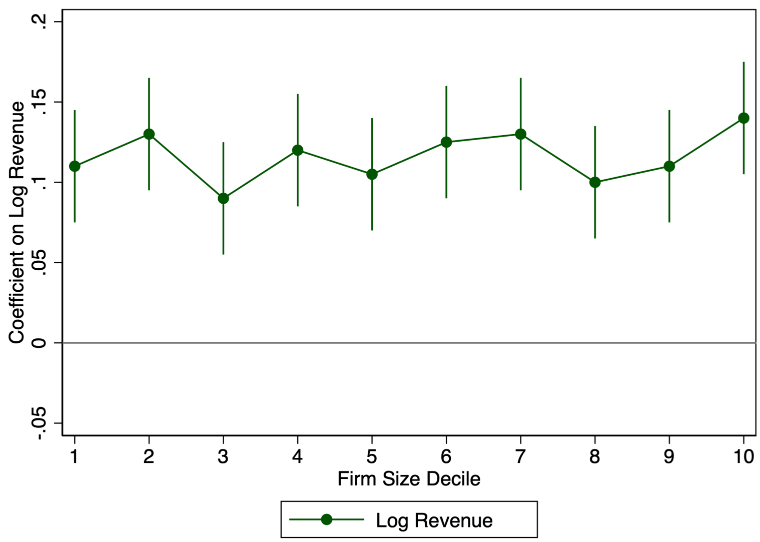

Notes: The figure demonstrates that the treatment effect of the policy in terms of log revenue is homogenous in the size of firms. The coefficients plotted represent the mean effect on log revenue in 2011-2014 of having an additional percent of the workforce treated (born 1950 or 1951) at the firm in 2005, controlling for the share of old workers (born 1948-1951) at the firm. Standard errors are clustered at the firm level.

Studying the effect of the *number* of treated workers for these outcomes would lead to mis-specification. This makes sense intuitively given that a levels approach would imply that adding a worker would have the same effect on production no matter the firm's size, which would only hold for the specific case in which firms faced constant returns to labor.

C3. Firm Level Anticipation Approach

In Section 5.2 I use the mean birth year among treated born 1950-1953 workers at firms. To validate that the mean is the appropriate transformation I compare rounded mean birth year plots to the plots in Figure 10. The latter plots show adjustments for firms with exactly one treated worker born 1950-1953. The Figure below shows that the rounded mean birth year approach yields adjustments very similar to those found in Figure 10. This suggests that the mean birth year of treated workers is the appropriate measure to use when examining the effect of anticipation at the firm level.

Figure A22: **Firm Level: Response by Rounded Mean Birth Year**

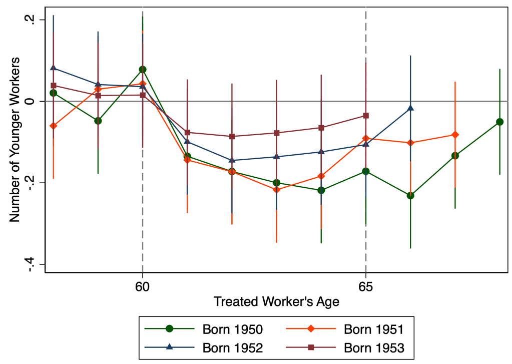

Notes: The sample is restricted to firms with below median free cashflow at baseline. The figure shows adjustments in the younger workforce by rounded mean birth year of treated workers at the firm. The plot demonstrates the effect of anticipation on the firm's labor adjustments. Each consecutive cohort c gives the firm an additional year to anticipate the policy change (see Appendix Figure A18 for the first stage plots by birth year). The coefficients represent the difference in the outcome of a firm employing a treated workers with rounded mean cohort $c = \{1950, 1951, 1952, 1953\}$ relative to the outcome of a firm employing a control worker born in 1949, when both workers were of the same age a. These correspond to coefficients $\beta_{c,\,a}$ from equation (9), which represent firm adjustments in the number of younger workers depending on the treated worker's cohort c. The dotted lines represent the window within which a treated worker could have retired early absent the policy change, but is now more likely to work. Note that data are censored for younger cohorts. The last year in the data is 2018, for individuals born in 1953 are observed until they are 65 years old. Standard errors are clustered at the firm level.