Labor Supply and the Pension Contribution-Benefit Link*

Eric French[†] Attila Lindner[‡] Cormac O'Dea[§] Tom Zawisza[¶]

January 6, 2026

Abstract

We estimate the impact of public pension incentives on labor supply far from the normal retirement age by exploiting Poland's switch from a Defined Benefit to a Notional Defined Contribution (NDC) scheme. This reform created a sharp cohort-based discontinuity in the link between current contributions and future benefits. Using this discontinuity and the universe of taxpayers, we estimate an employment elasticity with respect to the net return to work of 0.51 for men at ages 51-54. We estimate a lifecycle model to match these responses and discuss the broader implications of the reform. The shift to NDC reallocates work incentives over the lifecycle, strengthening incentives at younger ages, when labor supply is relatively inelastic, and weakening them at older ages, when labor supply is more elastic. This reallocation of work incentives tends to reduce aggregate lifecycle labor supply, which highlights the advantage of targeting pension incentives towards ages at which labor supply is most responsive.

*We thank Sreevidya Ayyar, Hanxiao Cui, Lucas Finamor, Tianyu Pang, Pawel Struski, and Yuan Wang for excellent research assistance. We thank our discussants Jeffrey Liebman at the NBER conference on "The Labor Market for Older Workers", Roozbeh Hosseini at the 12^{th} Annual Employment Conference, and Jaime Arellano-Bover, Alan Auerbach, Barbara Biasi, Agar Brugiavini, David Card, Agnieszka Chlon-Dominczak, Hedvig Horvath, Patrick Kline, Wojciech Kopczuk, Erzo Luttmer, Olivia Mitchell, Leszek Morawski, Emmanuel Saez, Andras Simonovits, Joanna Tyrowicz, Danny Yagan, Owen Zidar and seminar participants at UC Berkeley, Berlin, Birkbeck, Notre Dame, Delaware, Mannheim, Monash, Tilburg, Wharton, Ca' Foscari Venezia, GAINS Workshop, NBER Summer Institute, Barcelona Graduate School Summer Forum, FRB Atlanta Annual Employment Conference, Bocconi-CEPR Workshop on Social Security, "Conference on Social Insurance" at Goethe University Frankfurt, ASSA Meetings, "The Microeconomics of Aging" Conference at ETH-Zurich, MRRC Annual Conference, Pension Webinar, RFBerlin Workshop on the Economics of Aging, IFS/NORFACE conference on "Trends in Inequality: Sources and Policy", IFS/ESRC conference on "Inequality and the Insurance Value of Transfers Over the Life Cycle", and NBER conference on "The Labor Market for Older Workers" for their insightful comments. We also gratefully acknowledge funding from the Economic and Social Research Council (for the grants "Centre for Microeconomic Analysis of Public Policy" at the Institute for Fiscal Studies (RES-544-28-50001) and "Inequality and the Insurance Value of Transfers across the Life Cycle" (ES/P001831/1)) and also NORFACE for Grant 462-16-120: "Trends in Inequality: Sources and Policy". Lindner also acknowledges financial support from the Economic and Social Research Council (new investigator grant, ES/T008474/1.) and from the European Union's Horizon 2020 research and innovation programme (grant agreement number 949995). Finally, we thank the Polish Ministry of Finance for assistance with accessing administrative data. Any errors are our own.

[†]Cambridge and IFS

[‡]UCL and IFS

[§]Yale, NBER, IFS

[¶]OECD and IFS

JEL Codes: D15, H55, J22, J26

Keywords: labor supply, pensions, contribution-benefit link, defined benefit, defined contribution.

1 Introduction

In most OECD countries, social security contribution (SSC) revenue exceeds income tax revenue.[1] Since SSCs mostly finance public pensions for an aging population, they are projected to consume an increasing share of national income. Economists have long recognized that SSCs and other payroll taxes might be less distortionary than income taxes because there is often a link between SSCs paid in and pension and other benefits received (Burkhauser and Turner, 1985; Liebman et al., 2009). Furthermore, since the link between SSCs and future benefits is often tenuous, there is potential for substantial efficiency gains from tightening this link (Auerbach and Kotlikoff, 1985; Feldstein and Liebman, 2002; Lindbeck and Persson, 2003).

For this reason, many multinational organizations such as the World Bank (e.g. World Bank, 1994) and IMF have advocated tightening the link between current contributions and future benefits by switching from a Defined Benefit (DB) pension system to a Notional Defined Contribution (NDC) system. Many countries have followed these recommendations, including Italy in 1996, Hungary in 1998, and Poland and Sweden in 1999. However, so far no empirical evidence has established that changing the contribution-benefit link has an impact on labor supply many years before the retirement age. Although it is well established that the large implicit taxes from pension systems hinder older age employment (Gruber and Wise, 1999), there is little evidence on the impact of pensions on labor supply far from the retirement age. The responsiveness of labor supply at younger ages to pension incentives is an open question in part because people may not be fully informed of the details of pension systems (Mitchell, 1988; de Mesa et al., 2006). This lack of information potentially impacts the responses to pension incentives (Chan and Stevens, 2008; Bottazzi et al., 2006; Mastrobuoni, 2011; Liebman and Luttmer, 2015).

This paper provides what is, to the best of our knowledge, the first empirical assessment of how changing the link between SSCs and future benefits affects labor supply far from retirement age. We exploit a Polish pension reform in 1999 that introduced an NDC pension scheme. This new pension system retained the pay-as-you-go nature of the previous DB pension system and kept retirement ages constant but introduced many of the incentives associated with Defined Contribution (DC) systems.[2] As emphasized by the architects of

[1]In 19 out of the 37 OECD countries (including Poland), personal income taxes are less than one third of the total tax wedge (the sum of SSCs and personal income taxes).

[2]In particular, in the NDC system, working-age individuals contribute to the system and fund the benefits of current retirees, while the link between benefits and contributions is altered by introducing a proportionality between each individual's contributions and the benefits they receive. Furthermore, similarly

this policy change, one of the most important elements of the reform was to "introduce a strong link between contributions and benefits" (see page 59 in Chlon et al., 1999).

An advantage of the Polish reform we study is that Poland switched from a DB scheme to an NDC scheme, leaving other features of the pension scheme (such as the retirement age, private pensions and funding status) largely unaffected around the discontinuity, whereas reforms in other countries usually changed several pension features all at once. As a consequence, we can directly assess the incentive effects of future pension benefits on labor supply, holding other features of the pension system constant.

The reform had a considerable impact on work incentives. While work incentives generated by Poland's NDC system differ little throughout the life cycle, the DB system makes earnings at certain ages particularly valuable for pension wealth accrual. A key reason for this difference is that, in the DB system, pension benefits are calculated based on earnings over a selected subset of "best" years in an individual's earnings history. In the NDC system, benefits are roughly proportional to the (indexed) sum of earnings in all years; thus no prominence is given to earnings in any particular years. This feature of the DB system generates stronger work incentives than the NDC system at points in the lifecycle when wages are high, and weaker incentives when wages are low.

An individual's "best" years depend on that individual's life cycle earnings profile. For those with steeper earnings profiles, pension benefits in the DB system often depend heavily on earnings near age 50 since these are typically the highest earning years. For these individuals, reducing labor supply at age 50 has a large impact on pension benefits, providing strong labor incentives at that age. Conversely, for individuals with flatter wage profiles, incentives under DB rules differ less across the lifecycle. We exploit these differences in the change in incentives to identify the effect of changing the link between contributions and future benefits on labor supply. In particular, we separately study the effect of the reform in regions with high earnings growth, where individuals have steeper lifecycle earning profiles, and in low earnings growth regions, where individuals' earnings profiles are flatter.

Besides the change in incentives, the new NDC scheme led to a substantial reduction in pension wealth for individuals under the new rule. To isolate the effect of incentive changes from the effect of a reduction in pension wealth, we exploit the fact that the reform caused similar losses in pension wealth in high and low earnings growth regions, while the change in incentives was substantially different across regions. We calculate that the difference in the change in the net return to work between high and low growth regions induced by the

to a funded DC scheme, the rate of return reflects changes in economic prospects and growth.

policy change was 4.46 percent, while the difference in the induced change in pension wealth was only 0.64 percent. Therefore, by comparing responses to the policy in high wage-growth and low wage-growth regions we can capture the effect of changes in incentives net of the effect of changes in pension wealth.

We estimate the employment responses to the pension reform by exploiting the sharp cohort-based discontinuity created by the reform. For men, the reform applied only to those who were born after December 31st, 1948 and so were younger than 50 years old at the time of the policy's implementation.[3] This sharp cohort-based discontinuity implied that two individuals born just a few minutes apart faced radically different pension systems from age 50: the older one would still participate in the traditional DB system while his slightly younger counterpart was ushered into the new NDC system.

Using a regression discontinuity design (RDD) and the full population of tax returns linked to the Polish population registry, we estimate labor supply responses to the reform. Our empirical design identifies the effect of the policy change by comparing individuals who were born only a few days apart and face a similar labor market and economic environment but are assigned to different pension schemes.

We find that, as a result of the reform, the employment rate in the high-growth regions, which saw the largest decrease in work incentives, fell by over one percentage point (or 2.29 percent) more at ages 51-54 than that in the low-growth regions.[4] Importantly, given our interest in identifying effects at ages distant from typical retirement ages, these responses are observed between 15 and 11 years before these individuals reach the full retirement age.[5]

We use our estimates to assess the implied employment elasticity with respect to the net return to work. Since the difference in the change in the effective net return to work between high and low growth regions induced by the policy change is 4.46 percent, while the difference in the employment increase is 2.29 (s.e. 0.94) percent, the employment elasticity with respect to work incentives is $\frac{2.30}{4.46} = 0.51$ (s.e. 0.21). This elasticity is in the range of those typically estimated in the literature (see Chetty et al. 2013; Blundell et al. 2016 for reviews). The novelty of our paper lies in the fact that we estimate the labor supply response to benefits

[3]Since the introduction of the reform was more gradual for women, we focus throughout the paper on men. Nevertheless, in Table A.6 we report estimates for women. In line with the gradual introduction of the reform, we find more muted responses at each cohort discontinuity, but with a larger cumulative labor supply response.

[4]Our main analysis ends in 2002 due to an unanticipated and substantial change in the old age unemployment benefit program that complicates the labor supply estimates for the years in 2003 and 2004. We provide separate estimates for the years after 2004.

[5]The full retirement age is 65. This is close to the expected retirement age, which is 63 according to the 2005-2009 waves of SHARE data.

received many years in the future, whereas most of the literature estimates the labor supply response to the contemporaneous return to work. Our results provide constructive evidence that individuals' labor supply responds in a forward-looking way to incentives in the pension formula, suggesting that tightening the link between contributions and benefits has the potential to alleviate labor supply distortions caused by SSCs.

To disentangle the responsiveness of labor supply to incentives from people's valuation of future benefits, we estimate the labor supply response to an unanticipated radical reform in 2004 that impacted *contemporaneous* work incentives. The reform changed eligibility for a generous unemployment benefit available to individuals older than 55 years who were laid off from their jobs. This policy change affected the cohorts born in August 1949 and later but not the cohorts born before then. We estimate the labor supply response to the change in access to these old age unemployment benefits and find an employment elasticity of 0.68 (s.e. 0.04). This is only slightly larger than our elasticity estimated in response to the pension reform, suggesting that labor supply is only slightly less responsive to changes in discounted future pension benefits than to contemporaneous unemployment benefits. The old age unemployment benefit reform changed incentives similarly in high- and low-growth regions. We also find that estimated labor supply responses are similar across regions, suggesting that the populations in high- and low-growth regions respond similarly to contemporaneous incentives. This implies that the differential labor supply response to the pension reform in high- and low-growth regions likely reflects the differential change in incentives and not some fundamental difference across the two regions.[6]

We show that our results are robust to alternative ways of implementing the regression discontinuity design, alternative assumptions on calculating the change in incentives, and finer geographic disaggregation. We also study the differential incentive changes across the wage distribution. The observed changes in employment are larger for high wage workers, consistent with the larger change in work incentives they faced. Furthermore, the considerable differences in employment between the 1948 and 1949 cohorts are not found between "placebo" cohorts where there was no change in the policy (those born in 1946 versus 1947, 1947 versus 1948, and so on). In addition to that, we also provide estimates for older ages (56-59) for whom we find larger elasticities (0.83, s.e. 0.35).

In the final part of the paper, we estimate a lifecycle model to evaluate the impact

[6]Besides exploring the differential responses across regions to the old age unemployment benefit reform, we also studied responses to a large tax cut instituted in 2009 and studied in detail by Zawisza (2022). We find no statistically or economically significant differences in responses to the tax cut between the high- and low-growth regions. This again supports our key identifying assumption that labor supply elasticities are the same across two types of regions. Results are available on request.

of pension reforms on labor supply over the whole lifecycle. We estimate the structural parameters of this lifecycle model by matching the lifecycle employment profile in addition to regression discontinuity estimates from the reform. We use the estimated model to compare the effect of a move from a DB system to a NDC system on labor supply over the lifecycle. Adjusting the DB system so that government revenue under both systems is equivalent, we find that altering the age structure of incentives by switching to the NDC system causes overall labor supply across the lifecycle to fall by two months. This net fall is explained by the fact that the negative labor supply effects at age 50 are only partially offset by positive labor supply effects from earlier in the lifecycle. Contributing to this is the fact that labor supply is less responsive to incentives for those in their 30s than for those at older ages; thus the improved labor supply incentives at earlier ages yield less additional labor supply than that which is lost due to reduced labor supply incentives later in the lifecycle. This highlights that the labor supply effects of linking SSCs to future benefits will be largest at ages when labor supply is most responsive. Furthermore, these results also highlight that while policymakers often consider the direct (budgetary) impact of pension policies, there are significant indirect budgetary consequences that arise through changes in labor supply.

Our paper relates to several strands of the literature. A number of papers have examined the savings responses to changes in pension wealth, exploiting differences across cohorts, including Attanasio and Brugiavini (2003) for Italy, Attanasio and Rohwedder (2003) for the UK, and Lachowska and Myck (2018) for the 1999 Polish pension reform which we study. Instead, our paper focuses on labor supply and not savings decisions. Additionally, instead of relying on survey data sample sizes and/or a gradual implementation of a pension reform that compares the behavior of cohorts born in very different years, we combine population-level administrative data with a sharp discontinuity in changes in incentives.

A large literature examines labor supply responses to retirement incentives. That literature, however, focuses almost exclusively on labor supply responses close to the retirement age: see Feldstein and Liebman (2002); Coile (2015); Blundell et al. (2016) for reviews and Fetter and Lockwood (2018); Gelber et al. (2018); Manoli and Weber (2016) for examples. For instance, Liebman et al. (2009) studies the effect of SSC-pension benefit linkage on retirement decisions using the Health and Retirement Study, where the average respondent is almost 60 years old. In contrast, our paper studies the employment responses of individuals who are far from the retirement age, and so our results better reflect how incentives built into the pension system can distort labor supply throughout the lifecycle. Bovini (2019) finds evidence of forward-looking labor supply responses to a complex pension reform in

Italy that impacted accrual rates, pension wealth, and the early retirement age. Bovini (2019) finds a small labor supply response, but it is hard to compare those estimates to ours because the former relate to the full package of reforms, and do not separate the roles of pension wealth shocks and incentive changes. There is also some evidence that the pension eligibility age affects labor supply prior to the eligibility age (Jean-Olivier et al., 2010; Carta and De Philippis, 2019). Dean et al. (2020) find that self-employed workers increase their reported earnings in years significantly impact future pension benefits. Bozio et al. (2019) provide evidence suggesting that the incidence of SSCs on wages may differ from that of taxes, particularly when there is a strong connection between SSC contributions and future benefits.

This paper makes two key contributions to this literature. First, the paper estimates the impact of the pension formula on labor supply far from the retirement age by exploiting a switch from a DB to a NDC pension system, providing the first direct evidence on this matter. Second, this paper is the first that separates incentive effects from other features of the pension scheme (e.g. pension wealth changes) when estimating the labor supply responses to future pension benefits. As a result, our estimates can be used to study how pension reforms alter behavior throughout the life cycle.

Our paper is related to the literature studying the impact of taxes on labor supply. Most of the literature does not account for how social security contributions impact future benefits, and it thus treats social security contributions as just another tax creating the same type of tax wedge between market work and leisure as any other tax (see e.g. Blundell et al., 1998; Kleven, 2014; Ohanian et al., 2008). Another large labor supply literature goes to the opposite extreme and assumes that individuals fully internalize how their contributions impact future benefits. This includes studies using dynamic structural models of labor supply and retirement (see e.g. French, 2005; van der Klaauw and Wolpin, 2008; O'Dea, 2019; Borella et al., 2023) and an evolving literature on optimal tax policies in dynamic contexts (see e.g. Huggett and Parra, 2010; Golosov et al., 2016; Jones and Li, 2020).

The remainder of the paper is structured as follows. Section 2 presents a simple framework that can be used to measure changes in incentives. Section 3 presents the details of the 1999 reform and introduces the data we use. Section 4 assesses the changes in the contribution-benefit link which arose as a result of the reform. Section 5 presents the RDD empirical strategy. Section 6 presents our estimation results. Section 7 presents a dynamic model which rationalizes our RDD results in the context of a lifecycle model and simulates the effects of the Polish reform over the entire lifecycle. Finally, Section 8 concludes.

2 The Net Return to Work

Social Security Contributions (SSCs) and other payroll taxes differ from standard income taxes because SSC payments are often linked to future benefits (Burkhauser and Turner, 1985). This linkage, if recognized by the individual, could impact labor supply, mitigating the resultant distortions of payroll taxes.

This section describes a framework for evaluating the labor supply response to changes in the SSC benefit link caused by the NDC reform. To do this, we define the net return to work (relative to staying out of the labor force on out-of-work benefits) under pension scheme $k = \{DB, NDC\}$:

$$nrw_{it}^{k} = (1 - \tau(\tau^{pi}, \tau^{ss})) \cdot w_{it} - u_{it} + E_t(PV_{it}^{\text{Employed}_t,k} - PV_{it}^{\text{Not employed}_t,k}). \tag{1}$$

where w_{it} is individual i's wage at age t. The net return to work includes three components. The first component is the after-tax earnings $(1 - \tau(\tau^{pi}, \tau^{ss})) \cdot w_{it}$, which is a function of the personal income tax rate τ^{pi} and the Social Security tax rate τ^{ss}. The second component is u_{it}, which represents the welfare and unemployment benefits that are lost when the individual works (see Online Appendix C for further details on these). The third component is the increment to the present discounted value of expected pension benefits from work at age t for each pension scheme k. This last term, reflecting the contribution-pension benefit link, is calculated using $PV_{it}^{\text{Employed}_t,k}$, the present value of pension wealth if the individual works in period t given the current wage and entire earnings history, holding future labor supply constant, and $PV_{it}^{\text{Not employed}_t,k}$, the value if the individual does not work in period t, again holding future labor supply constant. The difference between the two is the increment to pension wealth that occurs as a result of working in period t. In the next section, we discuss the two pension schemes in detail and illustrate how the reform changed the pension contribution-benefit link. Throughout the text, nrw_{it}^{k} denotes the net return to work each individual faces, while nrw_{t}^{k} is the sample average net return to work at age t.

The reform we study switched the pension system from a Defined Benefit to a Defined Contribution scheme, so it changed the link between today's contributions and future benefits, and thus the net return to work. We can use this reform to calculate an employment elasticity with respect to the net return to work:

$$\eta = \frac{(P_t^{NDC} - P_t^{DB})/P_t^{DB}}{\Delta nrw_t / nrw_t^{DB}} \tag{2}$$

where $P_t^{NDC} - P_t^{DB}$ represents the change in the employment rate at age t which arises from changing the contribution-benefit link, and Δnrw_t represents the change in the net return to work from switching from DB to NDC.

Our definition of the employment elasticity is closely related to the standard formula. The main difference here is that the variation in the net return to work is coming from the change in the pension contribution-benefit link rather than from the change in the tax rate, which was held constant. Because there was no change in tax rates from the reform, the change in net return to work is: $\Delta nrw_t = (E_t[\Delta PV_{it}^{NDC}] - E_t[\Delta PV_{it}^{DB}])$, and so the percent change in the net return to work coming from the pension reform is:

$$\frac{\Delta nrw_t}{nrw_t^{DB}} = \frac{E_t[\Delta PV_{it}^{NDC}] - E_t[\Delta PV_{it}^{DB}]}{E_t[(1-\tau(\tau^{pi},\tau^{ss}))\cdot w_{it} - u_{it}] + E_t[\Delta PV_{it}^{DB}]} \tag{3}$$

where the expectations are taken over all individuals. We use equation (3) to assess the effect of the reform on incentives to work. Central to this formula is the change in (expected) present value coming from working at age t, $E_t(\Delta PV_{it}^k)$, in pension scheme k. The next section describes the new and old pension schemes and the calculation of the change in the present value of pension benefits.

3 Institutional setup and data

3.1 Institutional setup

The 1999 Polish pension reform. The 1999 pension reform in Poland introduced NDC pensions for those born after 31st December 1948. Those born in 1948 or earlier remained in the DB scheme. In the new system, a virtual account was opened for every individual and a record of all subsequent contributions to this account was kept by the Polish Social Security Administration, named ZUS.[7] These contributions predominantly fund current pension expenditures on a pay-as-you-go basis, as in the previous scheme. As a result, the new system can be described as a *notional* defined contribution system.[8]

[7] Polish name: Zaklad Ubezpieczen Spolecznych.

[8] The reform also gave the option to accumulate some of the contributions in capital funds managed by private pension funds. Those born between 1949 and 1969 could choose to either accumulate all of their contributions in the state-managed notional account or 38% in a private fund and 62% in the notional account. The default option was opting out from the private fund, and the government suggested that men (women) who were older than 45 (40) years at the time of the reform should not take the risk of opting in. As a result, 93% of the cohort born in 1948 chose to accumulate all their contributions in the state-managed notional account (Leifels et al. (2010)). In the paper, we assume that all workers are fully enrolled in the notional account.

Importantly for our empirical strategy, the date-of-birth discontinuity is sharp only for men. For women, the new system was introduced gradually. For instance, only 20% of the pension for women born in 1949 would come from the notional DC account, with gradually increasing amounts for each subsequent cohort. Only cohorts of women born in 1954 or after get the entirety of their benefit under the new rules. Due to this gradual introduction of the NDC system for women, we focus on men, for whom 100% of the pension for those born in 1949 would come from the NDC account.

DB Benefit formulae. The way in which past contributions translate into current pension benefits differs substantially between the DB and NDC systems. In the DB system, pensions are a function of two key variables: (1) the number of years in which an individual made contributions into the retirement system and (2) the individual's earnings relative to the economy-wide average in their best earnings years. At age 65, the monthly after-tax benefit for individual i is calculated according to the formula

$$b_{i65} = (1-\tau^{pi})\Big(\bar{y}_{65}(1-\tau^{ss})\Big)\Big(0.24 + \ 0.013 \cdot c_i \cdot aime_i + \ 0.007 \cdot n_i \cdot aime_i\Big), \qquad (4)$$

where $\Big(\bar{y}_{65}(1-\tau^{ss})\Big)$ is the average monthly salary for everyone in the economy in the year when the beneficiary turns 65 (net of the Social Security tax rate τ^{ss}), c_i is the number of contributory years on retirement, and n_i is the number of "non-contributory years". Non-contributory years are those in which the individual was not contributing for reasons such as being on disability benefit, in higher education, on maternity leave or on sickness leave. Contributory years are those in which the individual was working or receiving unemployment insurance benefits. The variable $aime_i = \frac{1}{\#\text{best}_i}\sum_{j\in\text{best}_i}\frac{y_{ij}}{\bar{y}_j}$ is "Average Indexed Monthly Earnings". To calculate this, we first take the ratio of individual i's annual earnings y_{ij} relative to the economy's average annual earnings $\bar{y}_j$ of the employed, for each year j. We then average this ratio over individual i's best years, best_i. The best years period is chosen by the individual as one of two periods, the best 10 consecutive years out of the last 20 prior to the official retirement age, or the 20 best earnings years over their working lives. Because individual earnings are divided by average economy-wide earnings when constructing $aime_i$, the DB formula contributions in the "best" years are implicitly indexed by average earnings.

NDC Benefit formulae. Under the new NDC system, the formula for pension benefits creates a much more direct link between past contributions and the monthly pension amount b_{65} at the retirement age of 65:

$$b_{i65} = (1-\tau^{pi})A^{NDC}_{i65}/(E[T|t=65]) \qquad (5)$$

where A_{i65}^{NDC} is the value accumulated in the notional account at 65, and $E[T|t = 65]$ is remaining life expectancy at the retirement age. In the NDC system, capital in the notional account is accrued according to the formula:

$$A_{it+1}^{NDC} = A_{it}^{NDC} \cdot (1 + r^{NDC}) + \tau^{tpcr} \cdot y_{it+1} \tag{6}$$

where $(1 + r^{NDC})$ is the real uprating factor on accumulated capital and $\tau^{tpcr} \cdot y_{it+1}$ is the contribution to the notional account.[9] The nominal uprating factor at the time of the reform was CPI inflation plus 0.75 times the growth in real aggregate earnings in the economy.[10]

Under the old DB system, the impact of contributions on pension benefits depends critically on whether an individual is in their best earnings years relative to others in the economy before retirement. In the NDC system, on the other hand, contributions from any year feed directly into the accumulated amount A_{it}^{NDC} in a given period.

Starting capital in the NDC system. Since the reform took effect on 1st January 1999 and affected individuals born in 1949 onward, many of those affected had made significant contributions under the old system. As compensation for these contributions, such individuals were given "starting capital" in their notional accounts, calculated based on their contributions history.[11]

[9]Here the amount τ^{tpcr} is sum of the "worker" and "employer" social security contributions, which is 0.1952. This is slightly different from the Social Security tax rate τ^{ss} of 0.1871 described in equation (1), which includes additional taxes to pay for disability and sickness benefits but does not include employer contributions.

[10]Specifically, in nominal terms:

$$1 + r^{NDC,nom} = \pi_{t-1} + 0.75 \cdot \left(\frac{WageBill_{t-1}}{WageBill_{t-2}} - \pi_{t-1} \right) \tag{7}$$

where π_{t-1} is one plus the rate of increase of the CPI in the year preceding indexation and $WageBill_{t-1} = \bar{y}_{t-1} \cdot e_{t-1} \cdot Pop_{t-1}$ is the total revenue collected by the social security administration in the year preceding uprating. We convert to real terms and take sample means. Unlike the DB formula, therefore, a fall in the total level of contributions coming from a fall in the number of workers in the economy would result in lower indexation of past contributions, even if average earnings in the economy remained constant. In the Appendix, we document that for the years 2000-2013, which are the focus of this study, the uprating factors were similar in both systems.

[11]The formula used was very similar to the DB pension formula:

$$b_{i50}^{start} = 0.24 \cdot \bar{y}_{50} \cdot p_{i50} + \ 0.013 \cdot c_{i50} \cdot aime_{i50} \cdot \bar{y}_{50} + \ 0.007 \cdot n_{i50} \cdot aime_{i50} \cdot \bar{y}_{50} \tag{8}$$

where $c_{i50}, n_{i50}, aime_{i50}$, are respectively the number of contributory years, non-contributory years, and Average Indexed Monthly Earnings at the time of the reform (which was age 50 for the cohort we study), and p_{i50} had the role of increasing starting capital with a weighted average of age and the total number of contributory and non-contributory years at the time of the reform: $p_{i50} = \sqrt{\frac{50-18}{65-18} \cdot \frac{n_{i50}+c_{i50}}{25}}$. Starting capital was then calculated as $A_{i50}^{NDC} = b_{i50}^{start} \times E[T|t = 62]$, where $E[T|t = 62]$ is remaining life-span at 62.

Contribution rates. The total social security contribution rate to the pension system τ^{tpcr} remained the same between the DB and the NDC system, at 0.1952 of the earnings bill, up to a cap of 2.5 times the average earnings in the economy. For those on employment contracts, half of these contributions were paid by the employer, and half were paid by the employee. The self-employed paid a lump sum of contributions equivalent to those paid by an employee earning approximately the minimum wage.

Information. The reform was widely discussed and highly publicized at the time in Poland. Furthermore, each participant in the new NDC system received an annual statement that included information on their capital account balance and an estimate of the monthly pension benefit under different assumptions about the retirement age (Chlon et al., 1999). Appendix Figure A.1 shows an example of this annual statement.

Exceptions. While most men born on or after 1st January 1949 faced the new NDC system, there were some important exceptions who remained in the DB system. For instance, those who worked in occupations outside of the main state social security system, such as farmers, members of the military, police, judges, teachers, and railway workers, were excluded. Also excluded were those in "special occupations", which included physically demanding jobs in sectors such as mining, energy, metallurgy, construction, logging, transport, the health sector, glass production, artists, and journalists. We estimate that 12% of the population was employed in agriculture and another 5% was employed in the other excluded or special occupations.[12] Although these exemptions could bias our estimated labor supply responses towards zero, we show below that accounting for this has only a small effect.

Minimum pension. For the cohorts we study here, all men were eligible for the minimum pension if they made contributions for at least 25 years and their lifetime earnings were very low. The level of this pension, which is the same for those in both the old and the new system, is set by statute every year and increases by at least the CPI inflation rate. The minimum pension, however, is only binding for a few: fewer than 3% of male pensioners received it in 2019.[13] Thus, the realized pension benefit would be the greater of the minimum and the benefits described in equations (4) and (5) for the DB and NDC schemes, respectively.

Other relevant institutional features. The normal retirement age for men was 65.[14] As

[12] Unfortunately, we are unable to observe whether someone belonged to an excluded sector or special occupation in our administrative data. We calculate the share of the labor force in special occupations from the Household Budget Survey, and the share of farmers from the Labor Force Survey. Our administrative data, described below, excludes those in agriculture. Thus, our estimated labor supply responses are only for the non-agricultural sector.

[13] In line with that, we find in our simulations that the minimum pension applies to only a very small fraction of men.

[14] The 1999 reform also affected some rules related to early retirement. Men born after December 31,

a consequence of a Constitutional Court decision in 2008, those born in December 1948 and earlier were eligible for early retirement at age 60, but those born in January 1949 or later were not. To insure that our estimated employment responses at the 1948-1949 cohort discontinuity are not capturing this change in the retirement rules, we drop years after 2007. Individuals were also eligible for Old Age Unemployment Benefit (OUAB) from the age of 60 if the termination of employment was caused by the employer.[15] In 2002, the age limit was lowered to 55, and in August 2004 it was raised again to 60. To address concerns that our reform discontinuity coincides with the attainment of age 55 for the 1948 cohort but not for the 1949 cohort (which is a concern for 2003 and 2004), we drop 2003 and 2004 from the analysis. Nevertheless, we examine the impact of the 1999 pension reform on labor supply after 2004, when both cohorts had already reached age 55. In a separate analysis, we also use the abolition of the (OUAB) in August 2004 to study employment responses to a contemporaneous change in incentives.

Regionally targeted programs. The only regionally targeted program we are aware of is the Extended Unemployment Benefit program. This program increased the potential duration of unemployment benefits from half a year to one year when the local unemployment rate was above 150% of the national average. This happened infrequently over the period of our study: only 12% of all localities were affected.[16] The quantitative impact of these programs is limited. The alternative to receiving these unemployment benefits is receiving welfare benefits. Because the level of welfare benefits is close to the level of unemployment benefits, longer UI durations will have only limited impacts on net incentives to work.

3.2 Data

Our data consists of the entire population of anonymized income tax records filed in Poland. All non-agricultural workers are required to file taxes if their annual income (including pension benefits) is above a certain threshold (2,296zł in 2000, which is equivalent to US $547). These reported earnings are used to calculate SSCs and pension benefits. Agricultural income is not included in the tax data. However, workers in agriculture belong to a separate pension fund and are unaffected by the pension reform. Our employment responses are

1948 lost the option of early retirement, even if they were unfit for work. However, they could still claim a disability benefit equal to the early retirement benefit. Therefore, this change was a relabelling of early retirees as disabled rather than an actual change in incentives.

[15] The Old Age Unemployment Benefit was called "swiadczenie przedemerytalne" in Polish. Those who received it were entitled to it until they reached the normal retirement age.

[16] This program was slightly more likely to be in place in high-growth regions than in low-growth regions (60% of all Extended UI programs were in high-growth regions). However, given the small share of regions with an Extended program, such differences will have little impact on our main results.

therefore estimated for the non-agricultural workers impacted by the reform.

We also have access to the population register in Poland, which we can merge into the administrative tax data. This allows us to identify, for each member of the population, whether he/she filed a tax return. Our measure of employment is an indicator for whether the individual had employment or self-employment income exceeding the earning threshold required to file a tax return. Since self-employed individuals might simply respond to the policy change by changing their reporting behavior, we also report estimates separately for the employed and self-employed.

In the main analysis, we use data for the years 2000-2002 for estimating the employment response to the switch from a DB to an NDC scheme. We end the analysis in 2002 to make sure that we do not pick up the effect of changes in eligibility for old age unemployment benefits. We nevertheless, provide additional labor supply estimates using the years after 2004. When we directly study the impact of the old age unemployment benefit, we use the data from 2005-2007. Finally, we exploit the full data range 2000-2013 when we estimate the earnings process, which we use for measuring incentives generated by the different pension systems across the lifecycle (we describe this procedure in the next section). Our administrative data covers information on date of birth, gender, marital status, residence,[17] as well as reported income from employment and self-employment. For the baseline regression discontinuity analysis, we have 1,363,922 individual-year observations between 2000 and 2002.

In Appendix Table B.1 we show that the employment rate calculated in our administrative data lines up well with the employment rate calculated using two representative household surveys: the Polish Household Budget Survey (HBS) and the Polish Labor Force Survey (LFS). The fraction of individuals in non-agricultural employment for the 1948 and 1949 cohorts is 48% in the LFS (based on 9,485 observations), which is very similar to our estimate in the administrative data (49% based on 1,669,539 observations). The estimated total employment rate (including agriculture workers) for the 1949 and 1950 cohorts is 60% in the LFS and 61% in the HBS.[18] Problems of underreporting do not appear to be of serious

[17] If an individual did not file taxes in a given year, we have access to the region they were most recently formally registered in, as well as the previous region in which they filed taxes.

[18] The employment rate between age 50-54 in Poland is around 65%, which is lower than the OECD average at 84% (source: OECD Dataset: LFS - Sex and Age composition). The lower employment rate is partially explained by the fact that the old-age unemployment to population ratio is a bit higher in Poland (12%) than in other OECD countries (4%) in this period. Furthermore, the participation rate is around 71%, which is lower than the OECD average (87%). The lower participation rate is due to the fact that some workers with a long working history in special occupations (such as metal workers or teachers) can retire already in that age range. These workers are unaffected by the reform and so our main empirical results are

concern in our administrative data.

4 The Effect of the Reform on the Net Return to Work

In this section, we describe how we calculate the net return to work in the DB and NDC pension systems. In the DB system, work incentives depend heavily on whether an individual was experiencing one of their best earnings years in the period preceding retirement. Conversely, in the NDC system, best years do not play such a prominent role.

To illustrate this, consider the change in the expected replacement rate for an individual in their early 50s who is deciding whether or not to work. Because the DB pension benefit formula uses average indexed earnings in the better of either the best 10 consecutive years or best 20 years, the increase in the replacement rate from working is small in all but the best years. However, in those best years, the increase in the replacement rate is potentially very large if wages in the best years are much higher than at other ages. On the other hand, an individual in the NDC system will experience a similar change in the replacement rate, irrespective of age. This highlights a key difference between DB and NDC schemes: both schemes provide work incentives, but at different ages. The DB scheme provides strong incentives to work in a narrow set of ages whereas the NDC scheme provides weaker incentives but at all ages.

To calculate the reform's impact on the net return to work for the full population, we calculate the increase in the present discounted value of pension benefits from working at each age following the existing literature (e.g., Attanasio and Rohwedder, 2003):

- We calculate retirement benefits according to the legislation in effect in the year of observation. We take into account any reforms and future uprating rules that have been legislated up to the time of observation. We assume that people expect the current legislation to persist.
- We assume that, when forming their expectations, people take their current residence as given and fixed.
- We account for longevity uncertainty using year, age, and gender specific survival probabilities for the cohort aged 50 in 1999. We assume age and gender specific mortality do not change after 2016. The maximum attainable age is fixed at $T_{death} = 100$.
- We assume that individuals expect to retire at the male normal retirement age of 65.
- We assume that aggregate wage growth, interest rates, and benefit uprating factors are constant over time. We estimate these by taking averages over the post-reform period.

relevant for workers not employed in these special occupations.

- We extend the framework in Attanasio and Rohwedder (2003) by also allowing for wage and unemployment risk. We estimate both from the data.

Making these assumptions, the change in present discounted value of benefits under each system k is:

$$
\begin{aligned}
E_t(\Delta PV_{it}^k) &= \left(\frac{1}{1+r}\right)^{65-t} \sum_{s=65}^{T_{death}} S_{s|t} \left(\frac{1}{1+r}\right)^{s-65} (b_{is}^{\text{Employed}_t,k} - b_{is}^{\text{Not employed}_t,k}) \qquad (9)\\
&= \left(\frac{1}{1+r}\right)^{65-t} \sum_{s=65}^{T_{death}} S_{s|t} \left(\frac{1+r^{index}}{1+r}\right)^{s-65} (b_{i65}^{\text{Employed}_t,k} - b_{i65}^{\text{Not employed}_t,k}),
\end{aligned}
$$

where $S_{s|t}$ is the probability of being alive at age s conditional on being alive at age t, $1+r$ is the risk free interest rate (and therefore $\left(\frac{1}{1+r}\right)^{s-t}$ discounts benefits earned at time s to time t), $1+r^{index}$ is the yearly indexation of pension benefits after age 65, and $(b_{i65}^{\text{Employed}_t,k} - b_{i65}^{\text{Not employed}_t,k})$ is the difference in age 65 pension benefits between working and not working at age t under the pension scheme k.

The change in (expected) present value of pension benefits has the following components. First, working at age t will increase age 65 pension benefits by $b_{i65}^{\text{Employed}_t,k} - b_{i65}^{\text{Not employed}_t,k}$, which depends on the whole lifecycle path of earnings and the pension formula. Second, once calculated, pensions are indexed by $1+r^{index}$ each year after age 65. Third, pension benefits are only received if still alive. As a result, the present discounted value depends on the probability of being alive at age s conditional on being alive at age t. Finally, all these future payouts are discounted to the present using the risk free interest rate $1+r$.

In the net present formula above, we observe the indexation factor ($r^{index} = 0.0116$), interest rate ($r = 0.0288$), the NDC uprating factor ($r^{NDC} = 0.0381$) and survival probabilities ($S_{s|t}$) in the data; see Online Appendix C for details. Since the pension benefit at age 65, $b_{i65}^{\text{Employed}_t,k}$, depends on earnings throughout the lifecycle, we simulate earnings profiles for individuals around the discontinuity (aged 49-50 on 1st January 1999). In the simulations, we deviate from the existing literature that assumes deterministic earnings profiles (see e.g. Attanasio and Brugiavini (2003), Attanasio and Rohwedder (2003), and Lachowska and Myck (2018)). Instead, we take into account wage and unemployment risks. These risks are important for the DB system as they affect which best years enter the benefit formula.

We estimate the earnings process in the following way. In the first step, we estimate the process for annual wages, w_{it}. In the second step, we estimate unemployment risk. Earnings, y_{it}, are equal to the offered wage w_{it} if working (P_{it}=1) and 0 if not working ($P_{it} = 0$). We assume wages are the sum of a deterministic and a stochastic component:

$$\log w_{it} = \mathbf{x}_{it}^T \kappa + \eta_{it} + \omega_{it} \tag{10}$$

where $\mathbf{x}_{it}$ consists of a fourth order polynomial in age, a linear time trend, an indicator for high growth region, and high growth region interacted with a fourth order polynomial in age and a time trend, and $\eta_{it} + \omega_{it}$ is the stochastic component that we describe below.[19]

Since pension benefits and the change in incentives to work depends on the shape of individuals' lifecycle profiles, we also exploit that individuals' lifecycle profiles vary across locations. In our benchmark specification, we divide the data into regions with below- and above-median wage growth in the years 2000-2013.[20] The time trend and age polynomial interacted with region in the wage equation (10) capture geographic variation in wage growth over time. This creates variation in the timing of individuals' best earnings years which is important in the DB formula but not the NDC formula. Below, we show our results are robust to finer levels of regional disaggregation. We also confirm that the time trend and age polynomial parameters are robust to alternative methods of estimating the wage equation, including controlling for individual person-effects.

Appendix Table C.4 reports estimates of wage parameters in equation (10). For the cohorts we study, these estimates imply annual real wages grow 0.75% faster in high-growth than low-growth regions. The implied difference in growth rates is robust to alternative specifications such as controlling for a full set of time dummies or individual person-effects (the latter is shown in Columns (3) and (4) of Table C.4).[21]

The stochastic components of wages is an AR(1) process η_{it} with an MA(1) innovation:

$$\eta_{it} = \rho\eta_{i,t-1} + \varepsilon_{it}, \quad \varepsilon_{it} \sim N(0, \sigma_\varepsilon^2) \qquad\qquad \omega_{it} = \xi_{it} + \theta\xi_{i,t-1}, \quad \xi_{it} \sim N(0, \sigma_\xi^2) \tag{11}$$

[19]We control for time but not cohort effects in the regression above. As we noted previously, pension benefits under the DB rules are calculated using individual earnings relative to other members of the economy at a point in time. By including time and age effects in our specification, we measure wages of an individual at a point in time relative to other members of the economy. If we were controlling for cohort but not time effects, we would compare wages at different points in an individual's life.

[20]Figure A.2 shows the distribution of low and high growth regions on a map of Poland. Table A.1 shows summary statistics for both regions. High growth regions tend to be more rural and have lower incomes in 2000. At the same time, there are only modest differences in the average age of the population and in the proportion of the population that is female or self-employed.

[21]The differential earnings growth between high- and low-growth regions reflects convergence in earnings across regions during the sample period. In 2000, average earnings in high-growth regions were 14.2 percent lower than in low-growth earnings, but by 2013, earnings in high-growth regions were only 3.6 percent lower than in low-growth regions.

The parameters of the age polynomial and time trend are estimated from the administrative data for the years 2000-2013 for men between ages 21-64. We estimate $\rho, \theta, \sigma_\varepsilon^2, \sigma_\xi^2$ using a minimum distance estimator, matching the variance-covariance matrix of wages. We estimate $\rho = .949$ (.001), $\theta = -.235$ (.013), $\sigma_\varepsilon^2 = .059$ (.001), $\sigma_\xi^2 = .027$ (.001). Although we are unaware of any estimates of the dynamic process for wages in Poland, the estimates are similar to those in the US (French, 2005) and many other countries (see the range of estimates cited in Krueger et al. (2010)). To account for unemployment risk, we estimate a first-order Markov process of unemployment spells.[22]

We use the estimated parameters to simulate wage, unemployment, and earnings (the product of wages and employment) histories, and thus benefits. If someone is employed at period t, we calculate $b_{i65}^{\text{Employed}_t,k}$ given the earnings throughout the lifecycle. We also calculate $b_{i65}^{\text{Not employed}_t,k}$ by assuming that the individual faces the same earnings and unemployment history as before but is not working in period t. In our calculations, we take into account all the details discussed in the institutional section, including the starting capital and the minimum pension.

Our simulations suggest that individuals in high earnings-growth regions were more likely to experience one of their best earnings years when aged 51-54, whereas individuals in low earnings-growth regions were more likely to experience their best earnings years at younger ages.[23] Thus, incentives to work at ages 51-54 under the DB system were greater in high earnings-growth regions than in low earnings-growth regions.

Table 2 presents the percent change in the net return to work caused by the NDC reform. Using the formula in equation (3), we calculate the average percent change in the net return to work at ages 51-54 for those in the 1949 cohort (who were impacted by the reform) relative to the 1948 cohort (who stayed in the DB system). We present the percent change in the net return to work in high earnings-growth and low earnings-growth regions. Since in high earnings-growth regions the best years were more likely to occur at ages 51-

[22] We estimate unemployment risk using the Polish Household Budget Survey that has detailed information on transitions from employment to unemployment and vice versa. An individual is considered to be in unemployment if he/she receives unemployment benefits. We estimate transition probabilities for individuals below age 50 and then we extrapolate those for all ages.

[23] The chance that ages 51-54 will be part of the best years calculations in the DB system is 49% in high growth regions and 45% in low growth regions. Moreover, working during the highest-growth years has a larger impact under the DB system than under the NDC system—particularly in high-growth regions. Specifically, the increase in future pension benefits from working in a high-growth region instead of a low-growth region is 42% higher under the DB system than under the NDC system. This is because the wages of those in their best years are much higher than other points in the life cycle, and this is especially true for those experiencing high wage growth. See Online Appendix C and Figure C.1 for more on age-specific differences in the return to work in the two systems.

54, the net return to work declined 8.82% in high-growth regions (vs. 4.35% in low-growth regions), a difference of 4.46%.

Besides calculating the change in net return to work, we also calculate the change in present value of pension wealth coming from switching from DB to NDC. For each individual, we take the simulated wage and unemployment shocks and calculate pension benefits (and, using equation (9), the present discounted value of those benefits) under the DB and NDC rules. On average, pension wealth dropped by about 13.5% in both the high- and low-growth regions. This pension wealth drop exceeds the one predicted by policy makers at the time of the reform, but is in line with simulations of Lachowska and Myck (2018), who studied the same reform. This discrepancy can be explained by the fact that projections at the time of the reform did not take into account the shape of the earnings profile over the lifecycle, which led to a lower than expected starting capital for many individuals.[24]

The reduction in pension wealth provides an incentive to work more, partly offsetting the reduced work incentive from the reform. Nevertheless, the size of the pension wealth drop was similar across locations (13.83% in high-growth vs. 13.19% in low-growth). Therefore, we isolate the effect of incentives from the change in pension wealth by focusing on the difference between high and low-growth regions.[25]

Incentives changed more in high-growth regions because individuals in those regions have lower wages and steeper wage growth. The higher wage growth in those regions means they are more likely to be in their best earnings years. The age polynomial part is very similar across locations and so it plays little role explaining the differential changes in incentives. In Appendix Table C.3, we provide further detail on what drives the differential incentives across regions. Notice that our regression discontinuity design filters out the effect of differential labor market trends on labor supply, as those should have a constant effect around the discontinuity.

To summarize, Table 2 shows that the difference in the change in the net return to work between high-growth and low-growth regions was 4.46%, while the difference in the change in pension wealth was only 0.64%. In the next section, we study the response of the labor

[24]The pension projections at the time applied a simple deterministic model that abstracts from the shape of the lifecycle earnings profile and wage and unemployment risk (see the assumptions they made on page 36 and 37 in Chlon et al. (1999)). Lachowska and Myck (2018) take into account the shape of the earnings profile, but abstract away from the wage and unemployment risks. We take into account both the shape of the earnings profile and the wage and unemployment risks.

[25]A key identifying assumption is that labor supply responses would be the same across the two types of region. In line with this assumption, we find that responses to the old age unemployment benefit are similar across locations. We also find that workers respond similarly to a large tax cut studied by Zawisza (2022) in the two types of region.

supply to these changes in incentives.

5 Empirical strategy

To identify the effect of the reform on labor supply, we exploit the sharp discontinuity created by the cohort-based nature of this reform. We apply a regression discontinuity design (RDD) where we compare individuals who were born a few weeks from each other but are covered by different pension schemes. More specifically, we estimate the following regression equation:

$$P_{it} = \alpha + \beta \mathbf{1}\{z_i < 50\} + f(z_i) + \varepsilon_{it}, \tag{12}$$

where P_{it} equals 1 if individual i is employed at time period t, and z_i is the age of the individual on 1st January 1999 (when the reform was introduced). Individuals younger than 50 years old at the time of the reform, $\mathbf{1}\{z_i < 50\}$, were ushered into the new NDC scheme, and so β assesses the impact of switching from the DB pension to the NDC scheme. We follow Hahn et al. (2001) and Lee and Lemieux (2010) and estimate two separate regressions of $f(z_i)$ on each side of the cutoff point. We report estimates with linear regressions and with kernel-weighted local-linear regressions using a triangular kernel. For the local-linear regression, we set the bandwidth at 150 days on either side of the discontinuity. In Online Appendix A, we show that our results are not sensitive to the chosen bandwidth values.

Since our simulations in Section 4 suggest that incentives changed differently for individuals in high earnings-growth and low earnings-growth regions, we also estimate the RDD regression specification separately for these two regions. The standard error on the differential response between high- and low-growth regions is obtained using the delta method, although we obtain the same standard errors if we estimate the differential response between high- and low-growth regions in one regression specification.

In our RDD, the running variable is birth date, which was determined many years before the policy change. Therefore, manipulation in the forcing variable is not possible. Nevertheless, there is a spike in reported births which occurs on the 1st of January of every cohort in our sample. This spike in reported birth is also observed in registry data in 1998, before the reform implementation. Thus, the spike is not driven by some policy-induced manipulation. Instead, the spike on January 1st likely reflects the fact that many in these cohorts were born at home (and not at hospital) and the dates of birth for these individuals were self-reported. While this reporting behaviour took place 50 years before the pension reform was announced, the characteristics of these switchers may be correlated with the labor-market outcomes we care about.

To deal with this issue, we exclude individuals born between December 16th and January 5th. We pick these thresholds because we see no evidence of under- or over-reporting of births outside of this narrow range. This is sometimes known as a "donut hole" regression-discontinuity design and has been used in other instances of systematic bunching around the cutoff (see e.g. Almond and Doyle, 2011; Barreca et al., 2011). For robustness, we also alternately perform our analysis using no donut hole at all, and using a broader donut hole where we drop all individuals who were born in January or December. Our results are not sensitive to various definitions of donut holes.

We also report estimates relative to the observed discontinuity in the "placebo" sample, born exactly one year later than our main estimation sample. In the placebo sample, we see a similar spike in births on January 1st. We estimate the regression discontinuity net of placebo in the following way. First, we create a stacked data set by appending the main sample containing the 1948 and 1949 cohorts observed in 2000-2002 (our main sample) with a dataset containing the 1949 and 1950 cohorts observed in 2001-2003 (our placebo sample).[26] We denote individuals belonging to the main sample in this stacked data with $M_i = 1$ and individuals belonging to the placebo sample with $M_i = 0$. For the placebo sample we assume that there is a discontinuity between those born on December 31st 1949 and those born on January 1st 1950. Therefore, the discontinuity threshold is age 50 at the time of the reform for individual i in the main sample, formally $k_{M(i)} = 50$, and it is age 49 for individual i in the placebo sample, $k_{M(i)} = 49$. The specification for the net-of-placebo RDD is then:

$$P_{it} = \alpha^P + \beta^P \mathbf{1}\{z_i < k_{M(i)}\} + f^P(z_i) + \left(\alpha^M + \beta^M \mathbf{1}\{z_i < k_{M(i)}\} + f^M(z_i)\right) \cdot M_i + \varepsilon_{it}, \quad (13)$$

where $f^P(z_i)$ and $f^M(z_i)$ are sample-specific controls for the forcing variable (age at the time of the reform) estimated separately on each side of the cut-off. In this regression, β^P estimates the change in employment at the placebo cut-off, while β^M shows the estimated employment change in the main sample relative to the placebo sample and is thus the parameter of interest.

Finally, we also check whether there is a noticeable discontinuity in individuals' observable characteristics at the cut-off. The results of this exercise are presented in Table D.1 in Online Appendix D for all individuals. While there is some evidence of a lower female share,

[26]For the placebo sample, we use years between 2001-2003 to make sure that we have the same age bands in the main and in the placebo samples.

higher rural share and higher local area employment rate among those born after January 1 in our main sample covering 1948-1949 cohorts, that discrepancy is very similar in the placebo sample covering 1949-1950 cohorts. As a result, in our net-of-placebo estimates, we find no indication of any unusual change in the covariates around the January 1st discontinuity. We also show the results for individual-level covariates in the low- and high-growth regions separately, and likewise find no change around the January 1st discontinuity for either.

6 Results

Employment responses. We start our analysis by evaluating the effect of the reform on the employment rate. Panels (a) and (b) of Figure 1 show the average employment rates over the years 2000-2002 by month of birth around the reform discontinuity. The x-axis shows the age of the individual on January 1st, 1999, the date the pension reform was introduced. Therefore, as we move along the x-axis, we show the employment-to-population ratio for cohorts that are successively older. The red vertical line shows the threshold of age 50 on January 1st, 1999, for which the new rules applied. Cohorts younger than the threshold (left of the red vertical line) were ushered into the new NDC scheme, while older cohorts (right of the red vertical line) stayed in the old DB system.

Figure 1 also plots the lines of best fit for individuals both below and above the discontinuity, as well as the 95% confidence intervals. The downward slope of these lines reflects declining employment rates at older ages. As we described above, we report non-agricultural employment rates. Agricultural workers were unaffected by the policy change and are excluded from the employment data (see Section 3 "Exceptions").

Since the change in incentives was different in high- and low-growth regions, we report estimates separately for the two regions. Panel (a) shows that in high growth regions, a 1.48 percentage point decline in the employment rate as a result of switching to the NDC scheme (left of the vertical red line) when using a "donut hole" RDD that excludes individuals born right around the discontinuity. Using the 52% baseline employment rate in the high-growth regions, this translates to a 2.8 percent drop. This fall reflects the decrease in these individuals' net return to work (shown in Table 2) as well as the effect of the reform on pension wealth. In Panel (b) of Figure 1, we also show the RDD result for the low-growth regions. In these regions, we do not find a significant difference between the DB and NDC cohorts. There is only a slight change in the employment rate, in line with the smaller decrease in the net return to work shown in our simulations.

Panel A of Table 1 presents the RDD estimates in tabular form. It presents the estimates

Figure 1: Effect of the Pension Reform on Employment: Treatment and Placebo Estimates.

(a) High-growth regions: **reform cohort**

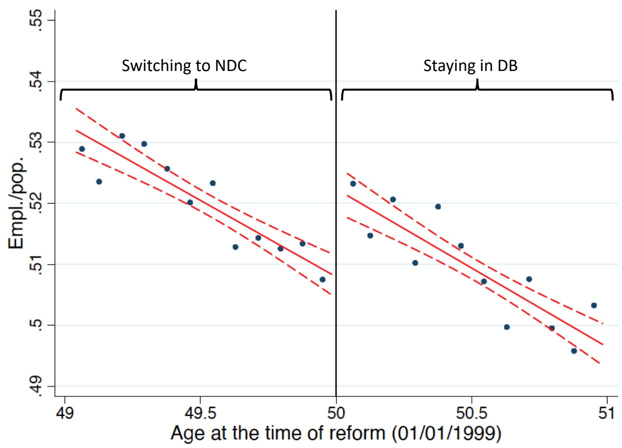

(b) Low-growth regions: **reform cohort**

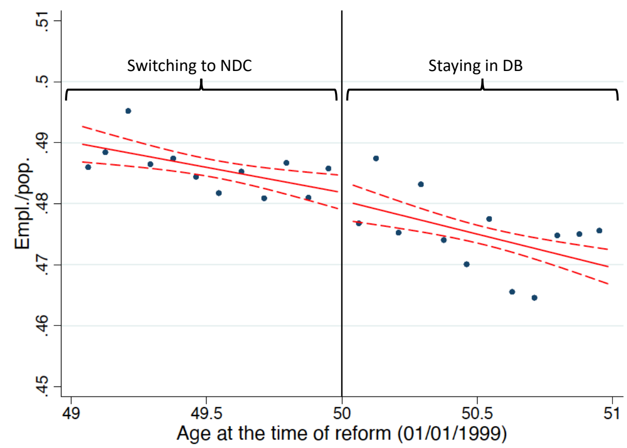

(c) High-growth regions: **placebo cohort**

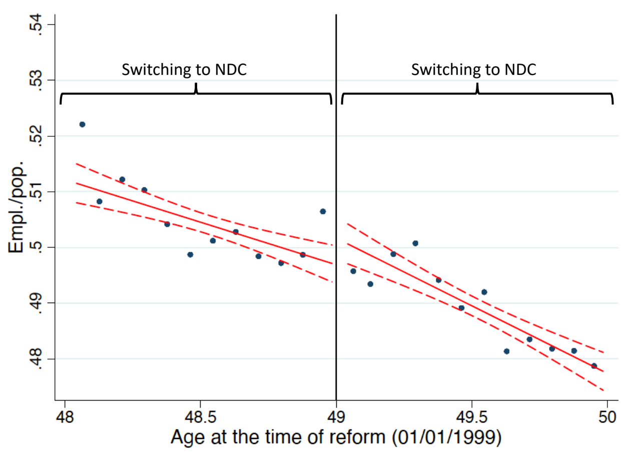

(d) Low-growth regions: **placebo cohort**

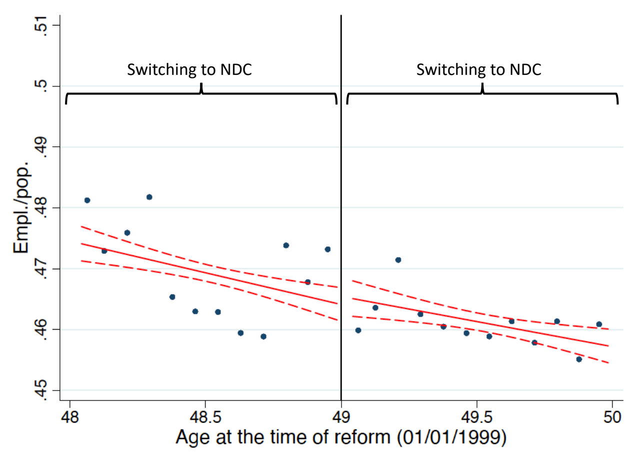

Notes: This figure plots the fraction of individuals who have positive earnings in a given year by month of birth (measured as the age on 01/01/1999). The top two panels, (a) and (b), show our treatment results. Individuals younger than age 50 on 01/01/1999 are in the new NDC scheme, while older individuals are in the DB scheme. We calculate the fraction having positive earnings for every year and then average them for the years 2000-2002. Panel (a) shows the fraction in high earnings growth regions, while panel (b) shows the fraction in low earnings growth regions. High earnings growth regions are regions with an above median earnings growth rate between 2000 and 2013, while low earning growth regions have below median earnings growth. To deal with the bunching in birth date at each year at January 1st, we apply a donut hole RD design and exclude those born between December 16th and January 5th. The solid lines are OLS lines of best-fit, allowing for different slopes and intercepts on both sides of the cutoff. The 95 percent confidence intervals are also shown. The bottom two panels, (c) and (d), show one of our placebo results. We plot the fraction having positive earnings for individuals born in 1949 and 1950 (age 48 and 49 at the time of the reform) where all individuals were ushered into the new NDC scheme and so there is no policy discontinuity. Panel (c) shows the fraction in high earnings growth regions, while panel (d) shows the fraction in low earnings growth regions. Otherwise, the plots are as in panels (a) and (b).

Table 1: The Effect of the Pension Reform on Employment And Wages

	(1)	(2)	(3)	(4)
Panel A: Change in employment probability				
High-growth	-0.0201***	-0.0148***	-0.0174***	-0.0105***
N = 545,435	(0.0024)	(0.0026)	(0.0050)	(0.0037)
Low-growth	-0.0022	0.0010	0.0027	0.0014
N = 818,487	(0.0020)	(0.0022)	(0.0041)	(0.0030)
Difference (High-Low)	-0.0179***	-0.0158***	-0.0201***	-0.0119***
	(0.0031)	(0.0034)	(0.0065)	(0.0048)
Panel B: Change in log wage				
High-growth	-0.001	-0.005	0.012	0.011
N = 313,720	(0.008)	(0.009)	(0.017)	(0.013)
Low-growth	-0.003	0.005	0.003	0.018
N = 439,545	(0.007)	(0.008)	(0.009)	(0.011)
Difference (High-Low)	0.002	-0.010	0.009	-0.007
	(0.011)	(0.012)	(0.023)	(0.021)
Sample	Full	Donut	Donut	Donut
$f(z_i)$	linear trend	linear trend	local linear	linear trend
net-of-placebo	no	no	no	yes

Notes: This table shows the estimated change in employment (panel A) and log wage (measured as earned income of workers) for those in work (panel B) at the reform discontinuity. Each cell in the table shows the β coefficients of the RDD specification shown in equation (12) (Columns (1)-(3)) or in equation (13) (Column (4)). The rows show the estimated employment and wage change for different regions. The first and second rows show the estimated effect in high and low-growth regions, respectively. High-growth regions are regions with above median earnings growth rate between 2000 and 2013, while low-growth regions have below median growth. The third row shows the difference between the high and low-growth regions. In Column (1) we use the full dataset. In Columns (2)-(4) we apply the donut hole RDD specification where we exclude those born between December 16th and January 5th. In Columns (1), (2) and (4) we estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. Column (3) estimates a kernel-weighted local linear regression, where we set the bandwidth at 150 days. Column (4) estimates the change in employment at the reform discontinuity relative to the change at the placebo discontinuity as in equation (13). The placebo discontinuity is estimated between the 1949 and 1950 cohorts, both of which switched to the NDC system. We report robust standard errors in parentheses. For the local-linear regression we calculate robust standard errors following Calonico et al. (2014). Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

of β from equation (12), which is the effect of being in the younger cohort at the discontinuity. The estimated effects are reported for both the high- and low-growth regions. We also calculate the difference between the two types of regions. The first column presents results from a specification using the main sample of all men born in 1948 or 1949 (who were age 49 or 50 at the time of the reform). The subsequent three columns show the "donut hole" RDD estimates where we exclude individuals born between December 16th and January 5th. The differences between Column 1 and the donut hole RDD estimates are small, suggesting that our results are robust to including individuals bunching right around December 31st.

Columns 2-3 explore alternative assumptions on the functional form of the running variable, $f(z)$, which is estimated separately on both sides of the discontinuity. In Column 2, we estimate a linear trend in birth date, while in Column 3 we estimate a local linear

Figure 2: Change in Employment for Various Cohorts

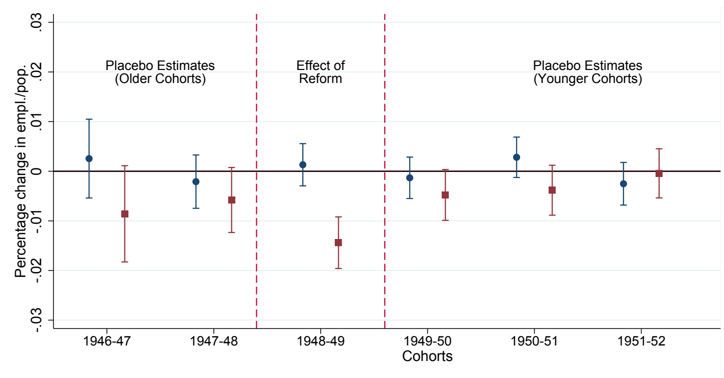

(a) Employment change for high and low growth regions

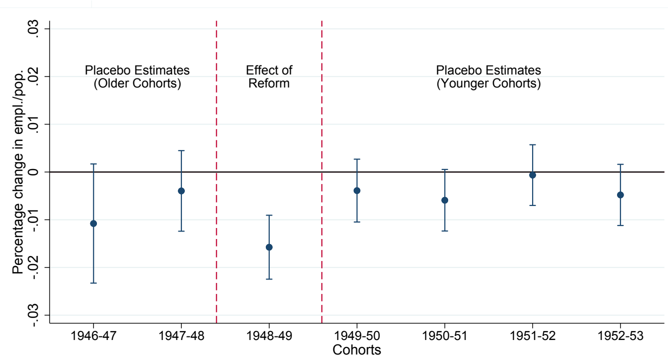

(b) Employment change, difference (high-low)

Notes: This figure plots the employment discontinuities estimated using the regression discontinuity design (see equation 12) for pairs of cohorts experiencing the policy discontinuity and for pairs of various "placebo" cohorts. The impact of the reform is estimated based on the 1948 and 1949 cohorts (individuals who were 49 and 50 years old on January 1st, 1999). Estimates for samples of older placebo cohorts are to the left of the 1948 and 1949 cohort estimates, while estimates for samples of younger placebo cohorts are to the right. We apply the same RDD specification as in Figure 1. Since we only have data from 2000 onwards, the estimates based on the 1946 and 1947 placebo cohorts are for individuals aged 53-54, while the estimates for the 1947 and 1948 placebo cohorts are for individuals who were 52-54 years old. The blue dots with the 95 percent confidence intervals on panel (a) show the estimated change in employment in low earnings growth regions, while the red squares show the estimates for high earnings growth regions. Panel (b) depicts the difference in employment change between high and low earnings growth regions.

polynomial, with a bandwidth of 150 days. In the latter specification, we apply Calonico et al. (2014)'s method to estimate bias-corrected robust confidence intervals. The estimates in the two specifications are very similar to each other. Moving from DB to NDC leads to a 1.5-1.7 (s.e. 0.3-0.5) percentage point decrease in employment in high-growth regions and a 0.1-0.3 (s.e. 0.2-0.4) percentage point increase in low-growth regions. In Table A.2 of the Appendix, we show that these results are robust to the chosen bandwidth.

As we discussed previously, the change in employment in high-growth and low-growth regions shows the combined effect of the pension wealth reduction and the change in net return (incentive effects). Nevertheless, since the reform affected pension wealth similarly in high- and low-growth regions, the difference in the employment change between high- and low-growth regions reflects the change in incentives net of any reform-induced changes in pension wealth. The difference between high- and low-growth in the last row of Panel A suggests that high-growth regions experienced a 1.6 (s.e. 0.3) percentage point drop in employment relative to low-growth regions in our specification with a linear trend and donut in Column (2).

We conduct a series of placebo analyses as a test of whether our estimates capture the effect of the reform and not of something else related to the timing of birth or age. Panels (c) and (d) of Figure 1 plots employment rates around a placebo discontinuity in high- and low-wage growth regions, respectively. Our placebo cohort consists of those aged 48 and 49 at the time of the reform (i.e., those born in the years 1949 and 1950). We analyze their labor supply in the years 2001-03 so that they are observed at the same age band as individuals in our main sample. Members of the placebo sample were all affected by the reform, regardless of whether they were born in late 1949 or early 1950. Similarly to the main estimates, we exclude individuals at the discontinuity, (i.e., these estimates incorporate the "donut hole"). Panels (c) and (d) of Figure 1 show that there is no significant difference in employment rates at the placebo discontinuity in either the high- or low-growth regions. The slightly higher employment rate for those born after January 1st in high-growth regions in the placebo sample is of an order of magnitude smaller than that for the cohorts in the main sample affected by the reform.

Figure 2 plots the estimated employment change using the 1948-49 cohort discontinuity, where individuals in the younger cohort were affected by the reform and individuals in the older cohort were unaffected. In addition, it plots a battery of placebo estimates using the cohorts 1946-47 and 1947-48 (none of whom were affected by the reform), and 1949-50, 1950-51, 1951-52 and 1952-53 (all of whom were affected by the reform). We also plot the 95% confidence interval for each estimate. We only use individuals aged 51-54 in all cohorts.

Panel (a) of Figure 2 shows estimates for both low- and high-growth regions. The estimate for the 1948-49 treatment cohort for the high-growth regions is statistically significant and large, while all of the placebo estimates for high-growth regions are smaller and statistically insignificant. For instance, our estimate for the 1948-49 treatment cohort is a 1.48 (s.e. 0.26) percentage point reduction in employment, whereas for the 1949-50 placebo cohort, we estimate a 0.43 (s.e. 0.26) percentage point reduction.[27] In the low-growth regions, the estimate for the 1948-49 treatment cohort is a 0.10 (s.e. 0.22) percentage point increase in employment, while the estimate for the 1949-50 placebo cohort is a 0.04 (s.e. 0.21) percentage point fall in employment. The placebo estimate for the 1946-47 cohort is larger than the other placebo estimates but is imprecisely estimated because it only uses data from 2000. By 2001, the cohort members reach age 55 and thus we do not use them.

Panel (b) of Figure 2 shows the difference in the estimate between the high- and low-growth regions for each cohort. For the treatment cohort, we estimate a statistically significant 1.58 (s.e. 0.34) percentage point difference, while all of the placebo estimates are smaller and statistically insignificant at the conventional levels. For instance, we find a statistically insignificant 0.39 (s.e. 0.33) percentage point difference for the 1949-50 placebo cohort. This

[27]These estimates are also reported in column (2) of Table 1 and column (2) of Table A.4.

Table 2: Employment Elasticity

	(1) High-growth	(2) Low-growth	(3) Difference (High-Low)
1. Change in net return to work (%)	-8.82	-4.35	-4.46
2. Change in pension wealth (%)	-13.83	-13.19	-0.64
3. Change in employment (%)	-2.01	0.28	-2.29
	(0.71)	(0.62)	(0.94)
4. Employment elasticity (Row 3) / (Row 1)	–	–	0.51
			(0.21)

Notes: This table shows the effect of the pension reform on the net return to work (row 1), on the pension wealth (row 2), on the change in employment (row 3) and on the resulting employment elasticity (row 4). The percent change in the net return to work is calculated using the formula in equation (3) and further details are in Section 2 and in Online Appendix C. To calculate the percent change in employment, we divide the net-of-placebo estimates of the change in employment from Panel A, Column (4) in Table 1, by the employment rate of the cohorts which were age 50 at the time of the reform and so stayed in the DB system. Columns (1) and (2) show the effects for high and low growth regions, respectively. High-growth regions are areas with above median earnings growth rate between 2000 and 2013, while low growth regions have below median earnings growth. The third column shows the difference between the high (Column (1)) and low-growth (Column (2)) regions. Row (4) shows the employment elasticity, which we calculate by dividing the percent change in employment (row 3) by the percent change in the net return to work (row 1). Robust standard errors are in parentheses.

is evidence that the main effects are only found where the policy discontinuity is present, and we find no indication of a differential effect in other cohorts. See Table A.4 in Online Appendix A for further details on the 1947-1948 and 1949-1950 placebo cohorts.

In Column 4 of Table 1, we present our main results relative to the placebo estimates. We report the β^M from RDD regression equation (13). The estimated impact of switching to NDC on employment is -1.05 (s.e. 0.37) percentage points in the high-growth regions and 0.14 (s.e. 0.32) percentage points in low-growth regions. The difference between the high- and low-growth regions is -1.19 (s.e. 0.48) percentage points. This is similar to the simple RDD estimates in Column 2 (-1.58 percentage points). We use these more conservative net-of-placebo estimates in the benchmark analysis when calculating elasticities.

Implied elasticity. Table 2 presents our estimated participation elasticity. Using the approach described in Section 4, row (1) shows the effect of the reform on net return to work in high- and low-growth regions, while row (2) shows the impact of the reform on pension wealth. Row (3) reports the percent change in employment as a result of the reform. We use our net-of-placebo estimates of the percentage point change (reported in Column 4 of Table 1) and divide it by baseline employment rates at the discontinuity for the cohort staying in the DB system.

To isolate the effect of the reform on incentives from the effect on pension wealth, we focus on differences between high- and low-growth regions (Column 3 of Table 2). Row (4) reports our estimated employment elasticity. We divide the percent change in employment (row 3) by the percent change in the net return to work (row 1). The estimated elasticity is 0.51 (s.e. 0.21), which is statistically significant. Since the regional differences in pension wealth changes are negligible (0.64%), this elasticity only captures the change in net return to work.

The change in net return to work induced by the reform has permanent and transitory components, and so our estimated reduced form elasticity is between the Marshallian (fully permanent change) and the Frisch (fully transitory change) employment elasticity. On the one hand, the change in the net return to work depends on whether earnings at certain ages belong to the "best years", which introduces some transitory component in the change in net return to work. On the other hand, being in the "best years" in a given year means that the next year will likely also be among the "best years". Therefore, there is some permanence in the change in incentives.

As a result, we can compare our estimates to the reduced form Marshallian and Frisch elasticities often reported in the literature. The Chetty et al., 2013 meta-analysis of these reduced-form micro studies suggest that the Frisch elasticity is around 0.32 while the Marshallian is around 0.25.[28] There is some evidence that older individuals tend to be more responsive to changes in incentives: Blundell et al., 2016 cite several studies where the reported elasticity is greater than 0.5.

Nevertheless, these estimated reduced form elasticities do not immediately translate to interpretable structural elasticities given that participation decisions are not simply governed by one single structural parameter or elasticity and may vary with other characteristics such as age. In Section 7, we use a lifecycle model, estimated to match the employment response to policy reform, in order to obtain a model-based estimate of the Frisch elasticity.

Wealth Effects. Our simulations show that the NDC reform delivered non-trivial declines in pension wealth in each region. Wealth declines are similar for individuals in high and low wage growth regions. Furthermore, wealth declines are similar for individuals who are high or low wage within a region. Thus the new NDC system was not significantly more or less redistributive than the DB system that preceded it.[29]

While it is not the focus of our paper, we can assess labor supply responses to these

[28] Chetty et al., 2013 report the Hicksian elasticity but acknowledge that their Hicksian elasticity is a mix of Marshallian and Hicksian elasticities, as many of the papers in their meta-analysis do not fully account for income effects.

[29] For example, when we stratify by whether an individual's wage is more or less than two times the minimum wage (which is relatively close to the median wage in the economy), pension wealth falls by 13.44% among high wage individuals and 14.23% among low wage individuals in high wage growth regions. Likewise, pension wealth falls by 11.88% among high wage individuals and 14.99% among low wage individuals in low wage growth regions.

wealth changes. To do this, we first net out the impact of the change in net return to work on labor supply. Table 2 shows that our estimate of the employment elasticity is 0.51. In low-growth regions the change in the net return to work is 4.35%. Together, these imply that employment should have fallen by (4.35%)(0.51)=2.22% in the absence of a pension wealth change. Nevertheless, we estimate a 0.29% increase in employment. We attribute the 2.51% (i.e., the difference between -2.22% and the 0.29%) higher observed employment change to the 13.19% reduction in pension wealth. An analogous calculation for the high wage region leads to similar numbers (we attribute the (8.82%)(0.51)-2.01% =2.49% increase in employment to the 13.83% reduction in pension wealth). This 2.5% change in the implied employment rate (absent pension wealth change) translates to 1.25 *percentage point* change in labor supply, given that the employment to population rate was around 0.5.

We can compare this magnitude to some recent estimates of labor supply responses to wealth shocks. Golosov et al. (2022) document employment responses to wealth shocks coming from lottery winnings. Their Table 3.2 indicates that a wealth shock that is 2.9 times of the average income in the economy[30] causes employment to decline by 3.7 percentage points. The wealth shocks we study are smaller – the 14% decline in pension wealth that our simulations show is equivalent to a wealth shock 1.5 times the average income in the economy.[31] Scaling our estimated 1.25 percentage point effect to the size of their wealth shock, our estimates imply an employment change by 2.4 percentage points, which is not far from the 3.7 percentage point effect estimated in Golosov et al. (2022). Furthermore, Lindqvist et al. (2020) finds considerably lower wealth elasticities in the Swedish context, while Giupponi (2024) finds somewhat larger wealth effects. This suggests that our estimated wealth effects are within the range of existing estimates in the literature.

[30]The shocks are worth $100,000, while their Table 2.1 indicates that average earnings are $34,500.

[31]Our simulation indicate that the present discounted value of pension wealth declined by 34,649 zloty, while average pre-tax earnings are 22,118 zloty.

Table 3: Employment Elasticity, Robustness

	(1) Change in *nrw* (%)	(2) Change in pension wealth (%)	(3) Change in emp(%)	(4) Employment elasticity
Panel A: Baseline				
1. Linear trend RDD, net-of-placebo,	-4.46	-0.64	-2.29	0.51
donut sample			(0.94)	(0.21)
Panel B: Estimation methods				
2. Linear trend RDD, net-of-placebo,	-4.46	-0.64	-2.99	0.67
full sample			(0.88)	(0.20)
3. Linear trend RDD, net-of-placebo,	-4.46	-0.64	-1.85	0.41
Jan-Dec donut sample			(1.04)	(0.23)
4. Local-linear RDD, net-of-placebo,	-4.46	-0.64	-3.40	0.76
donut sample			(1.80)	(0.40)
Panel C: Alternative Wage Processes				
5. AR(1) wage process	-4.80	-0.57	-2.29	0.48
(parameters from French (2005))			(0.94)	(0.20)
6. AR(1) + White Noise wage process	-4.46	-0.64	-2.29	0.51
			(0.94)	(0.21)
7. Region-specific wage process	-3.76	-0.76	-2.29	0.61
parameters			(0.94)	(0.25)
8. Wage process with individual	-5.08	-0.49	-2.29	0.45
fixed-effects			(0.94)	(0.19)
Panel D: Interest rates				
9. $r = 0.04$	-3.75	-0.64	-2.29	0.61
			(0.94)	(0.25)
10. $r = 0.06$	-2.74	-0.64	-2.29	0.84
			(0.94)	(0.34)
11. Actuarially fair uprating	-4.75	-0.64	-2.29	0.48
of NDC contributions			(0.94)	(0.20)
Panel E: Definition of employment				
12. $2000 p.a. threshold	-4.46	-0.64	-2.02	0.45
			(0.95)	(0.21)
Panel F: Elasticity Formula				
13. Alternative adjustment for	-4.46	-0.64	-2.54	0.57
unaffected workers			(1.05)	(0.24)
Panel G: Employment change between 2005-2007				
14. Estimate at age 56-59	-4.42	-0.64	-3.68	0.83
			(1.55)	(0.35)
Panel H: Elasticity by Wage				
15. Below 2x minimum wage	-2.23	0.76	-1.13	0.50
			(1.74)	(0.78)
16. Above 2x minimum wage	-4.38	-1.55	-3.76	0.86
			(1.46)	(0.33)

Notes: Panel A reports the benchmark elasticity in Table 2. Panel B shows robustness to various differences in the estimation methods. Panel C explores alternative ways to calculate the wage process. Row 5 uses the AR(1) parameterization from French (2005). Row 6 uses an AR(1) process with White Noise. Row 8 applies the benchmark AR(1) + MA(1) process but allows region specific parameter values. Row 9 estimate the deterministic components of wages with individual fixed effects (see Appendix Table C.3). In Panel D, rows 10 and 11 explore higher values of the real interest rate (instead of the baseline of 2.88%). Row 7 assumes that pension indexation in the NDC system is actuarially fair – an additional 1 zloty contribution leads to 1 zloty higher present discounted value of pension benefits in equation (9) instead of applying the actual risk free interest rate, pension indexation, and survival rates. Row 12 in Panel E shows robustness to increasing the threshold of earnings above which individuals are deemed to be in employment ($2000 p.a. instead of the baseline of $547). Panel F explores an alternative adjustment for unaffected workers described in detail in C.2. Panel G explores the effect of pension system at later ages. Panel H calculates incentives and employment change for low and high wage earners separately. Robust standard errors are in parentheses.

Robustness. Table 3 evaluates the robustness of the estimated benchmark elasticity. Panel A reports the baseline estimate of the implied elasticity derived in Table 2. Panel B shows the implied elasticity under alternative specifications of the regression discontinuity design using: a linear trend in birth date without applying the donut hole restriction, a linear trend applying a larger donut hole restriction (namely excluding all individuals born in January and December), and a local-linear specification applying the baseline donut. For each specification, we estimate the change in employment which we use to calculate the implied employment elasticity. We provide the net-of-placebo estimates on employment in these specifications, where the placebo estimates come from the 1949-1950 cohorts in 2001-2003. The implied elasticity estimates in all cases are statistically significant. The point estimates vary between 0.76 (local linear, with donut) and 0.41 (with December-January donut) in the various specifications, both of which are close to our baseline estimate of 0.51.

In Panel C, we assess the robustness of the results to alternative assumptions made in the wage process used for calculating the net return to work. In particular, we explore how the implied elasticities change if we apply alternative wage processes in our simulations. In row 5, we use the estimated wage process from French (2005), which uses data from the U.S. Panel Study of Income Dynamics. Row 6 in Panel C presents simulations where the estimated stochastic component of wages is modeled as an AR(1) process with White Noise, rather than the AR(1)+MA(1) process used in our benchmark specifications. In Row 7, we return to the benchmark parametrization but introduce separate parameter values for low- and high-wage regions. Finally, Row 8 includes individual fixed effects in the estimation of the deterministic component of the wages. Across all cases, the simulated changes in incentives remain very similar to our benchmark results, with elasticities ranging from 0.45 to 0.61. This demonstrates that our main findings are robust to alternative assumptions in the wage process.

Panel D explores robustness to different assumptions on the interest rates. In our baseline specification, we assume that individuals equally value the after-tax wage and the increase in the expected present discounted value from pension benefits from working, using an interest rate of 2.88% which is the rate for government bonds over the period 2000-19. However, households might discount future benefits more heavily if they are borrowing constrained, face high interest rates, are myopic, or lack full information about their pension incentives. In rows 9 and 10 we consider higher interest rates when calculating the net return to work to investigate sensitivity to assumptions about discounting. With higher interest rates, we find that the implied elasticity is somewhat larger. For instance, if the interest rate is 4%, then the implied elasticity is 0.61. If the interest rate is 6%, the implied elasticity is 0.84.

In row 11, labeled "actuarially fair", we explore an alternative assumption on the uprating factor, r^{NDC} (see equation 6). In the benchmark specification, we calculate the present discounted value of pension benefits in equation (9) by using the actual uprating factor on accumulated capital, r^{NDC}, the indexation factor r^{index}, interest rate r, and survival rate probabilities ($S_{65|t}$). This discounting implies that a contribution of 1 Polish złoty increases the present value of the NDC account by approximately 0.7 Polish złoty. As a robustness check, we increase the uprating factor on accumulated capital, r^{NDC}, to ensure that 1 Polish złoty contributed to the NDC account increases the present value of benefits by 1 złoty, and so the pension scheme is actuarially fair. The implied elasticity in this case is modestly smaller (0.48 instead of 0.51 in the benchmark case). This highlights that our estimates are not very sensitive to the actuarial unfairness of the NDC system.

In our baseline specification, we consider someone to be employed if he/she filed a tax return and so his/her annual income is at least \$547. In Panel E, we show that we get very similar results if we use alternative definitions of employment, namely earning above \$2,000 annually, which is equivalent to the annual earnings of someone earning the minimum wage. As can be seen, our estimates is 0.45, which is very close to our baseline estimate.

Some individuals at the age discontinuity were in either an excluded sector or a special occupation and were thus unaffected by the reform. Failure to account for the fact that not all men were affected by the reform could bias our estimate of the responsiveness of labor supply towards zero. In Panel F, we adjust our elasticity estimates by applying an upper bound on the number of workers who are in excluded sectors following the procedure described in Appendix C.2. Because the vast majority of men were affected by the reform, accounting for those who were not affected raises the estimated elasticity only slightly, to 0.57.

Elasticity estimates by finer regions. Our main estimates so far compared the employment change, the change in incentives and the change in pension wealth between high and low earnings-growth regions. In Figure 3, we assess the changes at a finer regional level. We calculate pension incentives over 2000 small administrative local areas in Poland.[32]

Panel (a) of Figure 3 shows that the change in incentives is tightly linked to local area-level earnings growth, while the change in pension wealth varies very little across regions.

[32]We calculate the change in incentives at each local area separately. For each local area, we use the economy-wide lifecycle earnings profile but adjust the earnings growth rate to reflect the local-area level earnings growth between 2000 and 2013.

Figure 3: The Percent Change in Employment and in Work Incentives Across Locations

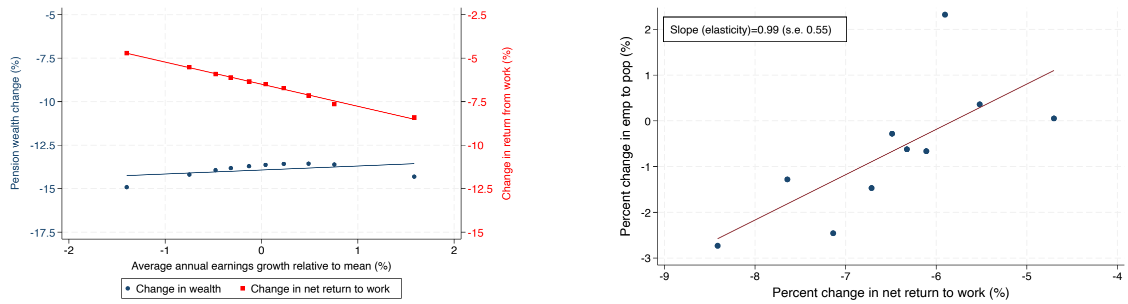

(a) Percent Change in Work Incentives and Pension Wealth Across Locations

(b) Percent Change in Employment and Work Incentives Across Locations

Notes: Panel (a) shows the non-parametric bin-scattered relationship between the percent change in net return to work and the average annual earnings growth (red squares) and between the percent change in pension wealth and the average annual earnings growth (in blue dots) across 2000 local areas. The annual earnings growth for each location is calculated between 2000-2013. We group the 2000 administrative local areas to 10 equally-sized bins based on their change in annual earnings growth (x-axis). The percent change in the net return to work is calculated according to equation (3). Panel (b) shows the non-parametric bin-scattered relationship between the estimated employment change at the reform discontinuity (net-of-placebo estimates based on equation (13)) and the percent change in the net return to work across the 2000 local areas. We estimate the employment changes for each local area separately by applying the local-linear RDD specification with the baseline donut (y-axis). We also plot a linear fit line using OLS. The slope of the linear fit in panel (b), reported in the top left of the panel, shows the relationship between the percent change in employment and the percent change in incentives across areas and so is an estimate for the employment elasticity.

Thus, as before, we use regional variation to identify the impact of changes in work incentives on employment separately from the effect of changes in pension wealth.[33]

In panel (b) of Figure 3 we plot the non-parametric bin-scattered relationship between the estimated RDD employment change at the reform discontinuity (net-of-placebo estimates based on equation (13)) and the percent change in net return to work. There is a clear positive relationship between the change in work incentives as a result of the pension reform and the estimated effect of the reform on employment outcomes. The figure also shows that the best linear fitting line is clearly upward sloping. The slope shows the relationship between the percent change in employment and the percent change in incentives and is therefore an estimate of the employment elasticity. We estimate a slope of 0.99 (s.e. 0.55), suggesting an even larger behavioral response to changes in incentives. Nevertheless, this design yields a noisier estimate, and the coefficient is not statistically different from our benchmark elasticity.

Overall, the finer regional-level analysis reinforces our benchmark results: incentives induced by the pension system matter for labour supply. This also highlights that the

[33]To be more precise, our calculation below assumes that there is no differential change in pension wealth between low and high average annual earnings growth regions. Accounting for the small differences in the magnitude of the decline in pension wealth between high and low average annual earnings growth regions will have a very small effect on the slope estimates in Figure 3(b), even if the wealth effects were large.

estimated difference between high and low earnings-growth regions are not sensitive to the specific cutoff used to define those regions.

Intensive margin responses. In Panel B of Table 1, we present the RDD results for log wage earnings among those reporting positive earnings. The estimates are small and insignificant, and the sign of the estimates is sensitive to the estimation method used. Furthermore, in Appendix Table A.5 we present results for the distribution of wage earnings using a Quantile Treatment Effects estimator. Consistent with the results in Table 1, at all quantiles the estimates are statistically insignificant when differencing between high and low growth regions. For this reason we focus on the extensive margin estimates, which are robust.

Women. As we explained in Section 3, women were more gradually moved from the DB into the NDC system. Women born prior to Dec. 31, 1948 were in a DB system. Those born Jan .1, 1949-Dec. 31, 1953 were in a hybrid system. For women born in 1949, 80% of the pension was calculated based on the old system, and 20% was based on the new system. For each subsequent birthyear cohort (beginning Jan. 1 of that year), the shares were 70/30, 55/45, 35/65, 20/80. In Table A.6 we report the estimated response to the reform for women. At each cohort discontinuity, we find a small employment response. Nevertheless, once we add up the estimated changes across the cohorts we find larger employment responses (albeit noisily estimated) than for men, consistent with female labor supply being more elastic than male labor supply.

Employment vs. self-employment. In Appendix Table A.7, we study separately the effect of the reform on self-employed workers and workers with employment contracts. A potential concern with our elasticity estimates is that they only pick up reporting responses that mainly affect self-employed workers (Kopczuk, 2012). Nevertheless, the results in Appendix Table A.7 highlight that the change in employment is mainly driven by changes in employment rates of people in paid employment, while the change in self-employment rates is limited. Employment responses for those with employment contracts feature third party reporting, and so tax evasion is less prevalent in those type of jobs (Kleven et al., 2011), suggesting we capture real responses. Moreover, the more limited response of the self-employed is consistent with the more muted changes in incentives.[34]

Role of firms. Because our administrative data do not include firm level information, we cannot study how the reform affected firm-level decisions. Nevertheless, our estimated

[34]Self-employed need to pay only a fixed amount of money for SSCs that is relatively low and thus the reform only induces small changes in work incentives for the self-employed, which is in line with the more muted response of this group.

employment responses should represent labor supply responses if firms cannot discriminate between workers who were born only a few days apart. Given the ample evidence that firms are reluctant to pay different wages to similar individuals even if they face different tax systems (see e.g. Saez et al. (2012); Saez et al. (2019)), we believe we are capturing labor supply responses.[35]

Future versus contemporaneous change in incentives. The labor supply responses shown above are to benefits received in the future. These responses depend on both the responsiveness of labor supply to incentives and the way that individuals value future benefits relative to current benefits. To disentangle the responsiveness of labor supply to incentives from people's valuation of future benefits, we estimate the labor supply response to a subsequent reform that impacted contemporaneous work incentives.

In particular, we exploit a radical change in eligibility for an old age unemployment benefit (OAUB) program which provided generous benefits to individuals whose employment was terminated by the employer. On 1st August 2004, a reform raised the eligibility age for this benefit from 55 to 60. Individuals could therefore take up the benefit if they demonstrated that their employment was terminated by the employer and they reached age 55 by 1st August 2004 and (and thus were born before 1st August 1949). This created a cohort-based discontinuity in access to the benefit: individuals born before 1st August 1949 were potentially eligible for the benefit, and individuals born after were not eligible. In Appendix Section E.1, we provide further details and analysis of the benefit program.

We exploit a RDD strategy to estimate the labor supply response to the reform. Figure 4 shows employment rates for men over the years 2005-2007 by age of the individual (in months) on 1st August 2004.[36] We compare individuals who were slightly younger than 55 on 1st August 2004 to individuals who were slightly older than 55. As we move along the x-axis, we show the employment-to-population ratio for increasingly older cohorts. The vertical line shows the eligibility threshold. Cohorts younger than the threshold (left of the vertical line) did not have access to the generous old age unemployment benefit program at age 55, while older cohorts (right of the vertical line) had access to the benefit. The figure shows a clear change in the employment rate around the discontinuity in both low-

[35] The decrease in employment could reflect involuntary separations. For example, in response to the reduced work incentives under the NDC rules, workers might lower their effort. If wages do not adjust accordingly (as shown in Table 1, Panel B) this could lead to an increase in the probability of layoffs. Unfortunately, we cannot empirically test the relevance of this productivity channel.

[36] In our data, we only observe yearly earnings. As a result, even if someone stops working in the middle of the year, we will see positive earnings for that individual in that year. That is why we focus in this analysis on the years between 2005 and 2007. Outcomes in 2004 are excluded because this would include information prior to the reform.

Figure 4: Effect of the Old Age Unemployment Benefit (OAUB) Program on Employment

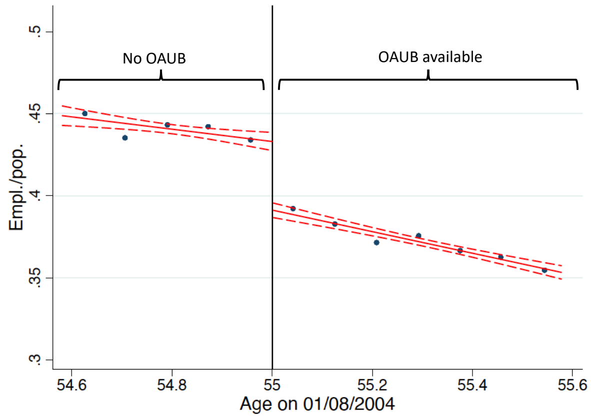

(a) High earnings growth regions

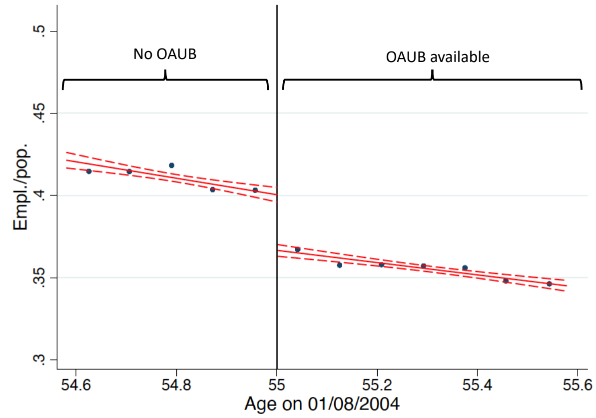

(b) Low earnings growth regions.

Notes: Fraction of individuals employed in a given year by month of birth (with age measured in months on the date of the OAUB reform, 01/08/2004). Individuals younger than age 55 on 01/08/2004 ceased to be eligible for the OAUB program, while older individuals could still claim the OAUB if they satisfied the eligibility criteria. We calculate the fraction employed in each year and then average this for the years 2005-2007. Panel (a) shows the fraction in high earnings growth regions, while panel (b) shows this for low earning growths regions. High growth regions are regions with an above median earnings growth rate between 2000 and 2013, while low growth regions have below median earnings growth. The solid lines are OLS lines of best-fit, allowing for different slopes and intercepts on both sides of the cutoff. The 95 percent confidence intervals are also shown.

and high-growth regions. We find that employment increases are similar across regions in response to the reform, with 3.3 and 4.1 percentage point changes in employment in low- and high-growth regions, respectively. When pooling all regions, the estimated drop is 3.8 percentage points. This large drop in employment is in sharp contrast to the years preceding the reform (2001-2003) when there is no discontinuous change in employment for those born around August 1st (see Table Appendix Table E.2).[37]

To calculate the change in incentives, row (1) of Table 4 shows the percent change in the net return to work as a result of the reform. The net return to work is:

$$nrw_{it}^{l} = (1 - \tau(\tau^{pi}, \tau^{ss})) \cdot w_{it} - u_{it}^{l} + E_t(PV_{it}^{\text{Employed}_t, NDC} - PV_{it}^{\text{Not employed}_t, NDC}),$$

where l reflects whether someone has access the OAUB program (l = OAUB) or not (l = NOAUB). The change in the net return to work comes from the change in the value of unemployment benefits when not working, $nrw_{it}^{\text{NOAUB}} - nrw_{it}^{\text{OAUB}} = -(u_{it}^{\text{NOAUB}} - u_{it}^{\text{OAUB}})$.

We calculate the implied elasticity as $\frac{(P_t^{\text{OAUB}} - P_t^{\text{NOAUB}})/P_t^{\text{OAUB}}}{-(u_t^{\text{OAUB}} - u_t^{\text{NOAUB}})/nrw_t^{\text{OAUB}}}$ where $P_t^{\text{OAUB}} - P_t^{\text{NOAUB}}$ is the change in employment among workers who were eligible for the OAUB program and

[37]This lack of an anticipation effect is consistent with the fact that the reform was passed on April 30, 2004 and was only announced a few months prior.

workers who were not eligible. This employment elasticity exploits the change in net return to work coming from the change in *contemporaneous* out-of-work benefits.[38]

Table 4: Elasticity Estimates using Contemporaneous Incentives

	(1)	(2)	(3)
Region	All regions	High-growth	Low-growth
1. Change in net return to work	-21.51	-21.49	-21.54
2. Change in net wealth	0.00	0.00	0.00
Fraction eligible: **40%**			
3. Change in employment (%)	-22.03	-24.03	-21.20
	(1.31)	(2.07)	(1.81)
4. Implied elasticity (Row 3) / (Row 1)	1.02	1.12	0.99
	(0.06)	(0.10)	(0.08)
Fraction eligible: **60%**			
5. Change in employment (%)	-14.68	-16.02	-14.13
	(0.87)	(1.38)	(1.21)
6. Implied elasticity (Row 5) / (Row 1)	0.68	0.74	0.66
	(0.04)	(0.06)	(0.06)

Notes: This table shows the effect of the old age unemployment benefit reform (OAUB) on the net return to work (row 1), on the net wealth (row 2), on the change in employment (row 3 and 5) and on the resulting employment elasticity (row 4 and 6). The change in the net return to work is a result of the change in out-of-work benefits at the policy discontinuity. The percent changes in employment when 40% and 60% were eligible are shown in row 3 and 5, respectively. The employment elasticity is shown in row 4 and 6, respectively.

There are multiple reasons why the labor supply responses to the NDC and OAUB reforms might be different. First, the NDC reform potentially impacts different individuals than the OAUB reform. Second, individuals may be unaware of the details of the NDC reform, whereas the incentives of the OAUB reform were more straightforward. Third, if individuals are impatient, myopic, or face liquidity constraints, then individuals may respond more to contemporaneous incentives.

With these caveats in mind, Table 4 presents estimates and the implied employment elasticity. OAUB eligibility caused the net return to work to fall 21.49% in high-growth and 21.54% in low-growth regions. The reform had no direct effect on individuals' wealth. To translate the estimated employment change around the discontinuity to an employment elasticity, we need to take into account the fact that, besides reaching age 55, there were other eligibility criteria for the OAUB program. Most notably, individuals needed to have a

[38]These individuals were age 49 at the time of the 1999 NDC pension reform, so they were covered by the NDC scheme.

sufficiently long employment history and the termination of the job must have been involuntary. While we do not directly observe the fraction of people who satisfy these criteria, we can use survey data to infer this information. We calculate that the eligible population in this age group in Poland would be between 40% and 60%, with 60% being our preferred rate. We provide details on how we arrive at these numbers in Online Appendix E.

When the eligible fraction of the population is 40%, then the implied estimated percent change in employment at the discontinuity among the eligible population is 22% when averaging over all regions, with a 21% change in low-growth regions and a 24% change in high-growth regions. The implied employment elasticity is between 1.12-0.99, depending on region. This employment elasticity is almost double the estimated elasticity coming from the pension reform. This would imply that even if workers are responsive to the incentives built into the pension formula, they are less responsive than to contemporaneous changes in incentives.

When the eligible fraction of the population is 60%, the implied employment elasticity is 0.66-0.74, depending on region. Our preferred estimate averages over the two regions, giving an estimate of 0.68. This is 1.33 times larger than our baseline estimate on the pension reform (though the estimates are not statistically different from each other). Overall, these findings suggest individuals are somewhat less responsive to changes in future pension benefits than to changes in the contemporaneous return to work, which is consistent with modest discounting of future benefits above and beyond standard discounting for time and mortality.

We calculate similar changes in incentives for high- and low-growth regions and also estimate similar changes in employment across the two regions. This suggests that labor supply reacts similarly to contemporaneous changes in incentives for the two types of regions. The fact that the estimated elasticity does not vary much by region supports our assumption that the differential employment responses to the pension reform for high- and low-growth regions documented before reflect the differential change in incentives, not a differential responsiveness of labor supply across regions.

Employment change between age 56-59. Panel G of Table 3 shows the estimated incentive change, employment change, and elasticity for those between age 56 and 59. To make sure that our estimates are not contaminated by the OUAB reform we focus here on the unaffected cohorts who were born before August 1st, 1949. We find that the change in incentives (between the high- and low-growth regions) is 4.42%, while the change in wealth is again very small. The estimated employment change is somewhat larger than the baseline estimates, leading to a 0.83 (s.e. 0.35) employment elasticity. As we show in section 7 below, the larger employment responses closer to the retirement age are consistent with the prediction of our structural model.

The employment responses after age 55 could in principle be influenced by the interaction between the OUAB benefit and the different pension systems. OUAB receipt affects pension accrual differently under the DB and the NDC rules. Nevertheless, these interactions are similar for both the high-growth and low-growth regions. Therefore, if we compare the incentive and employment changes between high and low-growth regions, we net out the interactions between the OAUB and the different pension systems, and isolate the effect of the differential incentives operating through the pension systems. Further discussion on this is provided in Online Appendix E.

Employment effects throughout the wage distribution. So far we have exploited variations in incentives for people living at high and low-wage growth locations. Here we exploit incentive differences across the wage distribution following the approach in Cengiz et al. (2019).

To do this we first calculate the incentive changes separately for individuals with high (low) realized earnings (in each region separately). To be more precise, we take every individual in our simulations whose earnings is two times above (below) the minimum wage, or would be above (below) if he were employed, and calculate the net return to work change for those individuals. We then estimate regression equation (13) where the left-hand side variable is an indicator whether the individual both works and earns above (below) twice the minimum wage. This estimate will yield the extensive employment responses of high (low) wage individuals if there are no intensive margin responses. Consistent with this assumption, we find no evidence of intensive margin responses, especially around the threshold, either if we look at averages wages (see Panel B in Table 1) or quantile regressions (see Table A.5). Therefore, our empirical estimates by high/low realized income likely to pick up the main extensive margin responses along the wage distribution.

The employment responses for below and above 2 times of the minimum wage are reported in Table A.8, while the implied elasticities are summarized in Table A.9 of the Appendix and Panel H of Table 3. The net return-to-work difference between high and low growth regions for low wage workers is 2.23%. This is in line with an employment response at the bottom of the wage distribution, which is also very small. At the same time, the difference in incentives between high and low growth regions is larger for high-wage individuals, where the net-return to work difference is 4.38%.[39] In line with the larger predicted change in

[39]Note that the calculated change in the net-return to work is larger for the full sample than in either the

incentives for high wage individuals, the employment response is larger at the higher part of the wage distribution. As a consequence, estimated employment elasticities for high- and low-wage workers are 0.50 and 0.86, highlighting that our estimated responses are driven by high wage individuals, for whom work incentives changed the most.

7 Effects over the lifecycle

Motivated by the sharp discontinuity in work incentives induced by the reform at age 50, our empirical work has focused on labor supply behavior near that age. However, pension reforms potentially impact labor supply at *all* ages. To place our results into a lifecycle context, in this section we develop a parsimonious lifecycle model, estimate its parameters using our quasi-experimental variation for identification, and use the estimated model to evaluate the effects of the reform across the whole lifecycle. In this section we outline the model, give our estimates, and show the implications of the reform for labor supply over the lifecycle. Further details are given in Online Appendix F.

7.1 Model

In the model, heterogeneous agents, who are subject to uncertainty over wages, unemployment, and survival, make consumption, saving and labor supply choices over their lifecycle.

Choices. Individuals make two choices each period – an extensive margin labor supply choice ($P_{it} = \{0, 1\}$), and a choice between consumption and saving.

Preferences. Individuals have preferences over consumption and leisure that can be represented by the following utility function:

$$U(c_{it}, l_{it}; \nu_i) = \frac{(c_{it}^{\nu_i} l_{it}^{1-\nu_i})^{1-\gamma}}{1-\gamma}, \tag{14}$$

where c and l are, respectively, consumption and leisure. The quantity of leisure consumed is given as $1 - (h(t) \cdot P)$, which is equal to an endowment of leisure (normalized to 1) less a share of that endowment foregone in periods when the agent works. Following French and Jones (2011), we allow that proportion to vary with age, reflecting the fact that work becomes costlier at older ages. The share of the leisure endowment lost when working varies with age according to: $h(t) = 0.3(\exp(t - 50))^{\zeta}$, and is normalized at age 50 to 0.3. The parameter ζ captures how rapidly the leisure cost of work grows with age. We assume that

high or low growth regions. This is due to how we calculate the net return-to-work. See the notes in Table A.9 for a discussion.

the weight each agent places on consumption in the utility function (ν) is constant across time but is heterogeneous in the population, and distributed as $\nu_i \sim N(\mu_\nu, \sigma_\nu^2)$.

Agents discount the future geometrically at rate β.

Demographics. Agents start working life at the age of 25 and face mortality risk. Conditional on surviving to t, the probability of surviving to period $t+1$ is s_{t+1}.

Wages, Income and Assets. In each period, each agent has a wage w_{it}. This has a deterministic component that evolves with age and a stochastic component that follows an AR(1) process. The deterministic component varies by region of residence, allowing us to study the implications of different pension schemes for different individuals with different wage growth over the lifecycle.

Agents get a job offer each period with a probability that follows a first-order Markov process. After observing the job offer and potential wage, they make a labor supply choice. If they work, their actual earnings, similarly to the treatment in Section 2, are $(1-\tau(\tau^{pi}, \tau^{ss}))w_{it}$: their potential wage net of taxation. If they do not work, their earnings are zero and they instead receive a welfare payment of u.

Pension Systems To be consistent with our previous simulations of the net return to work, we measure pension accrual as an increment to income in a method similar to that used in French and Jones (2011). In each period, agents accrue an increment to their income which is proportional to their wage and depends on each of a) which pension system is prevailing (DB or NDC), b) whether they work, c) their age and d) their region. This captures the key dimensions of variation relevant for pension accrual in our setting. Full details of how we model pension wealth accrual are given in Appendix F.1.

Agents save in a risk-free asset a_{it}, which earns a return of r. They cannot borrow. The budget constraint is:

$$a_{it+1} = (a_{it} + y_{it} - c_{it})(1+r), \quad a_{it+1} \geq 0 \tag{15}$$

where y_{it} is non-asset income. Prior to retirement (age 65), non-asset income is equal to after-tax earnings plus pension accrual if working and is equal to unemployment benefits plus (lower) pension accrual if not working. Thus, the difference between income from work and income from not work is the same concept used earlier in the paper, and outlined in Section 2. At age 65, the individual retires.

Model Solution and Summary. This model has six state variables ($\mathbf{X}_{it}$). Two of these – the agent's region of residence region$_i$ $\in$ {low growth, high growth} and their consumption weight (ν_i) – represent permanent heterogeneity, and four of which – age (t),

Table 5: Parameter Estimates, Model Fit, and Policy Evaluation

Panel A: Parameterization	Value	
Interest Rate (r)	0.0288	
Discount Factor (β)	0.972	
Risk aversion (γ)	4	
Panel B: Estimated parameters	Estimate	SE
Consumption Weight Mean (μ_ν)	0.511	(0.012)
Consumption Weight St. Dev (σ_ν)	0.077	(0.018)
Cost of work slope (ζ)	0.058	(0.002)
Panel C: Model fit		
Matched moments	Data	Model
Labor supply at age 30 (%)	84.1 (s.e. 1.3)	85.8
Labor supply at age 40 (%)	80.3 (s.e. 0.8)	82.1
Labor supply at age 50 (%)	72.3 (s.e. 0.7)	68.6
Labor supply at age 60 (%)	29.4 (s.e. 1.1)	34.4
Labor supply at age 64 (%)	16.2 (s.e. 1.2)	12.7
Reform labor supply effect, low-growth (%)	0.28 (s.e. 0.62)	0.13
Reform labor supply effect, high-growth (%)	-2.01 (s.e. 0.71)	-2.41
Implied difference in reform effect		
Reform labor supply effect, difference (high-low) (%)	-2.29 (s.e. 0.94)	-2.53
Panel D: Effect of switching from DB to NDC	Effect	
Net change in lifecycle labor supply, all	-2.2 months	
Net change in lifecycle labor supply, low-growth	-1.1 months	
Net change in lifecycle labor supply, high-growth	-3.8 months	
Panel E: Frisch Employment Elasticity	Elasticity	
Frisch Employment Elasticity at age 30	0.21	
Frisch Employment Elasticity at age 40	0.27	
Frisch Employment Elasticity at age 50	0.52	
Frisch Employment Elasticity at age 60	1.22	

Notes: Panel A reports parameters set prior to estimation. Panel B reports structural parameters estimated using Indirect Inference. Panel C shows the seven moments that the model targets. The final row of Panel C shows the difference in employment change between high and low growth regions. Panel D summarizes the overall net change in labor supply over the lifecycle as a result of switching from the DB system to the NDC system, using the estimated model to predict behavior for a cohort that spent their whole working life in the DB system and compare it to a cohort that spent their entire working life in the NDC system. To isolate the reform effect on changing the net return to work from the effect operating through a reduction in the overall generosity of the pension system, we scale down the accruals in the DB system (proportionally) to ensure that the two pension systems are revenue-equivalent. The average effects are weighted averages where we take the population share in the low-growth region at 60%. Panel E gives Frisch Employment Elasticities, calculated by perturbing individual wages at each age and calculating the percentage change in labor supply and dividing by the resulting percentage change in the net return to work.

assets (a_{it}), the presence (or otherwise) of an employment (offer_{it}) and wages (w_{it}) – vary across the lifecycle. Agents maximize:

$$V_t(\mathbf{X}_{it}) = \max_{\{c_{it},P_{it}\}} U(c_{it}, l_{it}; \nu_i) + \beta\Big(s_{t+1}\mathbb{E}_t V_{it+1}(\mathbf{X}_{it+1})\Big) \tag{16}$$

subject to the asset accumulation equation (15), leisure of $l_{it} = 1 - h(t) \cdot P_{it}$ and the determinants of income described in Appendix F.1. We solve the model using value function iteration. Appendixes F.2 and F.3 detail the decision problem and solution method respectively.

7.2 Parameterization and Estimation

Our approach to estimation of the model follows a two-step procedure. First, we set model parameters which can be identified external to the model or set with reference to the literature. These are given in Panel A of Table 5. We use the same interest rate r used previously in the paper, and set $\beta = \frac{1}{1+r}$. The coefficient of relative risk aversion on utility, which is set at 4, is a typical value (e.g., Conesa et al. (2009). The parameters of the deterministic and AR(1) components of the wage process are those in the simulation (see equation (10)). The Markov process governing unemployment is the same as in the simulations discussed in Section 4 and with parameter values given in Table C.2. In a second step, we estimate the three parameters of the model which are most directly linked to the labor supply decisions of our population. These are the mean and variance of the distribution of the consumption weight in the utility function, and the age trend in the leisure cost of work. We collect the parameters in the vector χ (where $\chi = (\mu_\nu, \sigma^2_\nu, \zeta)$). We estimate these using Indirect Inference, matching labor supply at ages across the lifecycle and our baseline estimated employment response to the reform.

To estimate these parameters, we first solve the model and simulate behavior for a cohort who stayed in the DB system. We then solve and simulate for a cohort who, like those born just after the year-of-birth discontinuity, were moved from the DB system to the new NDC system at age 50. The reform reduces the net return to work and wealth for our modeled agents differentially by region in a manner that mimics the falls described in Section 4. The solution to our model allows us to predict, at any candidate vector of parameters, what the effects of these changes will be on labor supply at each age for each individual. We choose parameters to best match the model-implied employment response to the reform in each region, in addition to matching the employment rate at ages 30, 40, 50, 60, and 64. We estimate these using data from the Labor Force Survey (that records all work, including agricultural work that is not observed in the administrative data) and methods that control for cohort effects, allowing us to obtain an estimate for employment for our cohort of interest throughout their lifecycle.[40] This gives us seven moment conditions and three parameters to estimate.

[40] We discuss this in Appendix F.4.2.

Panel B of Table 5 gives our structural estimates and their associated standard errors. Our identification of the parameters of the structural model leverages both the lifecycle profile of employment and our causal estimates of the policy reform. The mean consumption weight (at age 50) μ_ν is identified by the average level of employment over the life cycle. A higher consumption weight implies higher employment in order to fund that consumption. Our estimate of 0.51 sits at the mid-point of a range of papers that use utility functions similar to ours applied to data from other countries.[41] The variance of the leisure weight σ_ν^2 is identified by the responsiveness of labor supply to the reform. When the variance of the leisure weight is greater, there is more dispersion in reservation wages. More dispersion in reservation wages implies that there are fewer people near the employment margin, and thus labor supply is less responsive to the reform. One standard deviation in the consumption weight of 0.077 indicates a modest degree of heterogeneity in the population – implying that 95% of the population has a consumption weight of between 0.36 and 0.66. The parameter ζ captures the percent growth in the leisure cost of work with respect to age. This parameter is identified by the slope of the age profile in employent; if the leisure cost of work rises more rapidly with age, employment falls faster with age. We estimate that grows 5.8% per year.[42]

Panel C of Table 5 shows the moments that the model targets (with Appendix Figure F.2 showing the fit in labor supply across all ages). The model closely matches the difference in the employment effect between high- and low-growth regions, which we use in Section 6 to calculate our headline elasticity. The estimated difference is -2.29 versus a model implied value of -2.53. Recall that it is these cross-region differences that capture the intertemporal substitution effect, which is the focus of this study.

7.3 Model Validation

As two validation exercises, we implement the OAUB reform, and we evaluate labor supply responses at later ages, in both cases comparing the model-predicted outcome to the empirical results.

Modeled impact of the OAUB Reform We develop a version of the model that contains the OAUB reform. With this in hand, we assess whether the model is capable of replicating the labor supply response to the OAUB reform estimated previously.

To do this, we proceed as follows. First, we run the simulated individuals through the

[41] Conesa et al. (2009) estimate a value 0.377. Nishiyama and Smetters (2007) have between 0.45 and 0.50, O'Dea (2019) estimates a range from 0.42 to 0.52, French and Jones (2011) estimate heterogeneous groups with an average in the population of 0.62.

[42] French and Jones (2011) estimate that this leisure cost of work grows 6.6% per year at age 60.

benefit calculator (described in Section 4) to evaluate the effect of the OAUB benefit on the income that individuals receive if they choose to work and the income they receive if they choose not to work. We use this to calculate average replacement rates out of potential wages. Second, we solve the model assuming individuals cannot receive the OAUB benefit, allowing us to mimic the incentives of the cohort born after August 1, 1949. Third, we solve the model, giving modeled individuals the option of the OAUB replacement rate–calculated in the first step–at age 55 if they do not work, mimicking the incentives of those born before that date. In particular, we introduce the benefit as an unexpected shock and then re-solve the model. Fourth, we simulate optimal decisions both for those who received the option of the benefit at 55 and those who did not, calculating the effect of the OAUB reform.

We find that the model predicts a 13.9% (15.4%) decline in employment at age 55 in high (low) growth regions in response to receipt of the benefit. Table 4 shows (assuming 60% of all individuals are eligible for the OAUB benefit) that being eligible for the benefit reduces employment by 16.0% (14.1%) in high (low) growth regions among those eligible for the benefit. Thus, the structural model predicts broadly similar employment responses to the OAUB reform as those we estimated.

Employment effects at older ages Table 3, Panel G shows that employment response to the NDC reform is larger at ages 56-59 than just after the reform. We estimate a -3.69% difference between high and low growth regions for ages 56-59 compared to -2.29% for ages 51-54. Our model predicts a -3.1% difference for ages 56-59 and a -2.5% difference for ages 51-54. This indicates that the model not only closely matches the targeted differences between high- and low-growth regions at ages 51–54 but also accurately captures the untargeted responses at later ages.

7.4 Results

An advantage of complementing our regression discontinuity estimates with a lifecycle model is that we can evaluate the effects of the incentives from different pension systems on labor supply across the whole lifecycle. To do this, we first use our estimated model to predict behavior for a cohort that spent their whole life in a system with NDC work incentives. We then compare their employment to a cohort who spent their entire life with DB work incentives. To focus on the role played by changes in the net return to work, we ensure that the two pension systems are revenue-equivalent by finding a factor which scales down the value of pension accrual for those in work in the DB system proportionally at each age by a factor such that the present discounted value of net government spending is the same in

Figure 5: Effect of Switching to an NDC on Labor Supply Over the Lifecycle

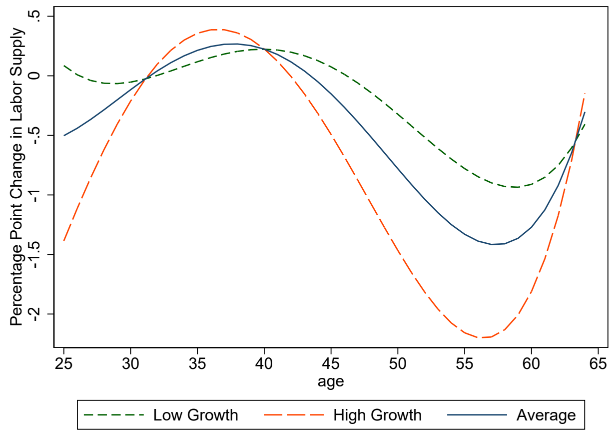

Notes: This figure plots the percentage point change in the employment-to-population ratio at each age coming from switching from the DB pension scheme to the NDC scheme. We predict behavior for a cohort that spent their whole working life under DB work incentives and compare it to a cohort who spent their entire working life under NDC work incentives. To isolate the reform effect on changing the net return to work from the effect operating through a reduction in the overall generosity of the pension system, we scale down the accruals in the DB system (proportionally) to ensure that the two pension systems are revenue-equivalent. The dashed green line shows the effect of the reform for individuals with low earnings growth, the red dashed line shows the effect for individuals with high earnings growth, and the blue line shows the average of the two. We smooth the responses using a 4^{th} order polynomial. The average effects are weighted averages with the weights representing the share of individuals in each region.

both pension systems (Online Appendix F has further details on the implementation of this counterfactual experiment). We thus compare behavior under two systems which are equally costly to implement but which provide different work incentives across the lifecycle.

Figure 5 shows the change in labor supply (in percentage points) across the lifecycle in each region caused by moving to the NDC system. Note that the model-predicted effect of the reform which was implemented as a surprise at age 50 when estimating the model will not be the same as the modeled difference between DB and NDC systems at that age in the experiment here. Here we study the impact of the implementation of an NDC system over the whole lifetime, and compare it to similarly-costly DB system. By design, this differs from the impact of the NDC reform on the transition cohorts. Hence, agents in the modeled 'steady-state' cohorts can adjust their behavior at all ages.

Figure 5 illustrates that switching to the NDC scheme increases labor supply at some ages and reduces labor supply at other ages. Panel D of Table 5 shows that, averaged over the lifecycle and over both low- and high-growth regions, employment would, on average, be reduced by around 2 months under the NDC scheme, compared to an equally-costly DB scheme. For those living in low-growth regions, there is a net fall in average labor supply of around 1 month. For individuals in high-growth regions, however, labor supply over the lifecycle falls nearly four months. This change for those in the high-growth regions is non-trivial. For instance, existing studies suggest that extending the early retirement age by one year extends work by roughly three months, depending on the country, institutional environment, gender, and data, with some studies suggesting more than three months (e.g., Lalive et al. 2023) and some studies suggesting less (e.g., Cribb et al. 2016).

A key reason why the negative effects dominate is that the estimated labor supply elasticities are greater at older ages (when the NDC reform reduced work incentives) than they are in their 30s (when the reform improved incentives). We calculate Frisch employment elasticities to be 0.21 at age 30, 0.27 at age 40, 0.52 at age 50, and 1.22 at age 60. This pattern of increasing labor supply elasticities across the lifecycle is a robust feature of structural models with an extensive margin choice.[43] The reason is that, at younger ages, male labor supply tends to be very high, and relatively few workers are close to the participation margin. Providing greater work incentives at those ages increases the net return to work for many workers who are inframarginal, and so has a modest effect, despite being costly. For those who are closer to retirement ages, there is a greater mass of agents close to the participation margin than for those their 30, and so providing work incentives will target those incentives towards workers whose participation behavior is more likely to be affected. Ages 30 is, therefore, an expensive time for the work incentives for men to be sharpened from a government revenue standpoint: labor supply is high, and so extra work incentives from higher pension accrual are expensive, and responsiveness to these incentives is lower, so revenue gains from increased labor supply are modest. This analysis highlights that targeting incentives at those ages where labor supply will be most responsive is a consideration that policy-makers need to consider carefully when designing policies that have implications for work incentives across the lifecycle.

8 Conclusion

This paper shows that individuals' labor supply is responsive to changes in the link between *current* social security contributions and *future* pension benefits, even 10-15 years before

[43]See Appendix F.6 for details on how we calculate these elasticities and Figure F.3(b) for a profile of elasticities across the lifecycle. These elasticities are calculated for an anticipated change in wages at those ages, implying opportunities for inter-temporal substitution of labor supply across the whole lifecycle. The elasticity we estimate using our quasi-experiment was for an unanticipated change in the net return to work, which limits the opportunities for inter-temporal substitution of labor supply and results in a smaller elasticity. This pattern of increasing elasticities across the lifecycle is also found by those studying extensive margin labor supply using structural models of labor supply in other settings. Figure F.3(b) compares our estimates of this elasticity to those found by French (2005), French and Jones (2011), Fan et al. (2024), Jones and Li (2023), Keane and Wasi (2016), Borella et al. (2023).

the expected retirement age. We demonstrate this by exploiting the 1999 Polish pension reform, which switched a DB system to an NDC system. Under the DB system, earnings in a small number of years – those in which earnings were at their peak – were particularly important in determining pension benefits. On the other hand, in the NDC system, all years are roughly equally important. In line with these changed work incentives, we find changes in labor supply that imply an employment elasticity with respect to the net return to work (which includes both the wage and the gain in expected pension benefits) of 0.51 (s.e. 0.21).

Our estimates, therefore, show that the incentives built into the pension calculation formula matter, and so the design of pension systems can have implications for labor supply throughout working life. Nevertheless, as demonstrated by our lifecycle model, how much these incentives matter varies over the lifecycle. Tightening the link between pension benefits and contributions might not have the desired impact if the changes in incentives do not target individuals for whom labor supply is most responsive.

It is worth emphasizing that pension design needs to consider more than labor supply. Tightening the link between current contributions and future benefits has implications for the distribution of living standards of retirees (Diamond and Gruber, 1999) although in the Polish context our simulations suggest that the NDC system was as redistributive as the previous DB system. Recent evidence highlights the efficiency-equity tradeoffs for pensions in other contexts (see e.g. Haller (2019); Kolsrud et al. (Forthcoming)). Therefore, even if our estimates can be used to quantify the potential efficiency gains from considering such reforms, the distributional aspects of such policies should be also taken into account. Balancing efficiency and distributional concerns should be a central focus for future research and policy discussion.

References

Almond, Douglas and Joseph J. Doyle, "After Midnight: A Regression Discontinuity Design in Length of Postpartum Hospital Stays," *American Economic Journal: Economic Policy*, August 2011, *3* (3), 1–34.

Attanasio, Orazio and Agar Brugiavini, "Social Security and Households' Saving," *The Quarterly Journal of Economics*, 2003, *118* (3), 1075–1119.

Attanasio, Orazio P. and Susann Rohwedder, "Pension Wealth and Household Saving: Evidence from Pension Reforms in the United Kingdom," *American Economic Review*, December 2003, *93* (5), 1499–1521.

Auerbach, Alan J and Laurence J Kotlikoff, "The efficiency gains from social security benefit-tax linkage," Technical Report, National Bureau of Economic Research 1985.

Barreca, Alan, Melanie Guldi, Jason Lindo, and Glen R. Waddell, "Saving Babies? Revisiting the effect of very low birth weight classification," *The Quarterly Journal of Economics*, 2011, *126* (4), 2117–2123.

Blundell, Richard, Alan Duncan, and Costas Meghir, "Estimating labor supply responses using tax reforms," *Econometrica*, 1998, pp. 827–861.

— , **Eric French, and Gemma Tetlow**, "Retirement incentives and labor supply," in "Handbook of the economics of population aging," Vol. 1, Elsevier, 2016, pp. 457–566.

Borella, Margherita, Mariacristina De Nardi, and Fang Yang, "Are marriage-related taxes and social security benefits holding back female labour supply?," *The Review of Economic Studies*, 2023, *90* (1), 102–131.

Bottazzi, Renata, Tullio Jappelli, and Mario Padula, "Retirement expectations, pension reforms, and their impact on private wealth accumulation," *Journal of Public Economics*, 2006, *90* (12), 2187–2212.

Bovini, Giulia, "Essays in applied economics." PhD dissertation, The London School of Economics and Political Science (LSE) 2019.

Bozio, Antoine, Thomas Breda, and Julien Grenet, "Does Tax-Benefit Linkage Matter for the Incidence of Social Security Contributions?," IZA Discussion Papers 12502, Institute of Labor Economics (IZA) July 2019.

Buckles, Kasey S and Daniel M Hungerman, "Season of birth and later outcomes: Old questions, new answers," *Review of Economics and Statistics*, 2013, *95* (3), 711–724.

Burkhauser, Richard V. and John A. Turner, "Is the Social Security Payroll Tax a Tax?," *Public Finance Quarterly*, 1985, *13(3)*, 253–267.

Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik, "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 2014, *82* (6), 2295–2326.

Carta, Francesca and Marta De Philippis, "Working horizon and labour supply: the effect of raising minimum retirement age on middle-aged individuals," *Banca d'Italia, Temi di discussione*, 2019, *26.*

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs*," *The Quarterly Journal of Economics*, 05 2019, *134* (3), 1405–1454.

Chan, Sewin and Ann Huff Stevens, "What you don't know can't help you: Pension knowledge and retirement decision-making," *The Review of Economics and Statistics*, 2008, *90* (2), 253–266.

Chetty, Raj, Adam Guren, Dayanand Manoli, and Andrea Weber, "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual*, 2013, *27* (1), 1 – 56.

Chlon, Agnieszka, Marek Góra, and Michal Rutkowski, "Shaping pension reform in Poland: security through diversity," Social Protection Discussion Papers and Notes 20852, The World Bank 1999.

Coile, Courtney C, "Economic Determinants of Workers' Retirement Decisions," *Journal of Economic Surveys*, 2015, *29* (4), 830–853.

Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger, "Taxing capital? Not a bad idea after all!," *American Economic Review*, 2009, *99* (1), 25–48.

Crawford, Rowena and Cormac O'Dea, "Household portfolios and financial preparedness for retirement," *Quantitative Economics*, 2020, *11* (2), 637–670.

Cribb, Jonathan, Carl Emmerson, and Gemma Tetlow, "Signals matter? Large retirement responses to limited financial incentives," *Labour Economics*, 2016, *42*, 203–212.

de Mesa, Alberto Arenas, David Bravo, Jere R Behrman, Olivia S Mitchell, and Petra E Todd, "The Chilean Pension Reform Turns 25: Lessons From the Social Protection Survey," Working Paper 12401, National Bureau of Economic Research August 2006.

Dean, Andrés, Sebastian Fleitas, and Mariana Zerpa, "Dynamic Incentives in Retirement Earnings-Replacement Benefits," 2020.

Diamond, Peter and Jonathan Gruber, "Social security and retirement in the United States," in "Social security and retirement around the world," University of Chicago Press, 1999, pp. 437–473.

Fan, Xiaodong, Ananth Seshadri, and Christopher Taber, "Estimation of a Life-Cycle Model with Human Capital, Labor Supply, and Retirement," *Journal of Political Economy*, 2024, *132* (1), 48–95.

Feldstein, Martin and Jeffrey B Liebman, "Social security," *Handbook of public economics*, 2002, *4*, 2245–2324.

Fetter, Daniel K. and Lee M. Lockwood, "Government Old-Age Support and Labor Supply: Evidence from the Old Age Assistance Program," *American Economic Review*, August 2018, *108* (8), 2174–2211.

French, Eric, "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior," *Review of Economic Studies*, 2005, *72(2)*, 395–427.

— **and John Bailey Jones**, "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica*, 2011, *79* (3), 693–732.

Gelber, Alexander, Adam Isen, and Jae Song, "The Effect of Pension Income on Elderly Earnings: Evidence from Social Security and Full Population Data," *Unpublished*, May 2018.

Giupponi, Giulia, "When Income Effects are Large: Labor Supply Responses and the Value of Welfare Transfers," 2024. Unpublished manuscript.

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski, "Redistribution and social insurance," *American Economic Review*, 2016, *106* (2), 359–86.

__ , Michael Graber, Magne Mogstad, and David Novgorodsky, "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income," *mimeo*, 2022.

Gruber, Jonathan and David A. Wise, *Social Security and Retirement around the World*, University of Chicago Press, 1999.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 2001, *69* (1), 201–209.

Haller, Andreas, "Welfare Effects of Pension Reforms," *Mimeo.*, 2019.

Huggett, Mark and Juan Carlos Parra, "How well does the US social insurance system provide social insurance?," *Journal of Political Economy*, 2010, *118* (1), 76–112.

Jean-Olivier, Hairault, François Langot, and Thepthida Sopraseuth, "Distance to Retirement and The Job Search of Older Workers: The Case For Delaying Retirement Age," *Journal of the European Economic Association*, 2010.

Jones, John Bailey and Yue Li, "Social Security Reform with Heterogeneous Mortality," Technical Report, Federal Reserve Bank of Richmond, 2020.

__ and __ , "Social security reform with heterogeneous mortality," *Review of Economic Dynamics*, 2023, *48*, 320–344.

Keane, Michael P and Nada Wasi, "Labour supply: the roles of human capital and the extensive margin," *The Economic Journal*, 2016, *126* (592), 578–617.

Kleven, Henrik Jacobsen, "How can Scandinavians tax so much?," *Journal of Economic Perspectives*, 2014, *28* (4), 77–98.

__ , Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez, "Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark," *Econometrica*, 2011, *79* (3), 651–692.

Kolsrud, Jonas, Camille Landais, Daniel Reck, and Johannes Spinnewijn, "Welfare Effects OF Pension Reforms," *American Economic Review*, Forthcoming.

Kopczuk, Wojciech, "The Polish Business flat tax and its effect on reported incomes: A Pareto improving tax reform?," 2012.

Krueger, Dirk, Fabrizio Perri, Luigi Pistaferri, and Giovanni L. Violante, "Cross-sectional facts for macroeconomists," *Review of Economic Dynamics*, 2010, *13* (1), 1–14. Special issue: Cross-Sectional Facts for Macroeconomists.

Lachowska, Marta and Michał Myck, "The Effect of Public Pension Wealth on Saving and Expenditure," *American Economic Journal: Economic Policy*, August 2018, *10* (3), 284–308.

Lalive, Rafael, Arvind Magesan, and Stefan Staubli, "How social security reform affects retirement and pension claiming," *American Economic Journal: Economic Policy*, 2023, *15* (3), 115–150.

Lee, David S and Thomas Lemieux, "Regression discontinuity designs in economics," *Journal of economic literature*, 2010, *48* (2), 281–355.

Leifels, Arne, Christoph Müller, and Bernd Raffelhüschen, "Mind the Pension Gap: On the Relationship Between Future Pensions and Pre-retirement Consumption in Poland," Technical Report, Research Center for Generational Contracts 2010.

Liebman, Jeffrey B and Erzo FP Luttmer, "Would people behave differently if they better understood social security? Evidence from a field experiment," *American Economic Journal: Economic Policy*, 2015, *7* (1), 275–99.

Liebman, Jeffrey B., Erzo F. P. Luttmer, and David G. Seif, "Labor Supply Responses to Marginal Social Security Benefits: Evidence from Discontinuities," *Journal of Public Economics*, December 2009, *93(11-12)*, 1208–1223.

Lindbeck, Assar and Mats Persson, "The Gains from Pension Reform," *Journal of Economic Literature*, March 2003, *52*, 33–42.

Lindqvist, Erik, Robert Ãstling, and David Cesarini, "Long-Run Effects of Lottery Wealth on Psychological Well-Being," *The Review of Economic Studies*, 2020.

Manoli, Day and Andrea Weber, "Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions," *American Economic Journal: Economic Policy*, November 2016, *8* (4), 160–182.

Mastrobuoni, Giovanni, "The role of information for retirement behavior: Evidence based on the stepwise introduction of the Social Security Statement," *Journal of Public Economics*, 2011, *95* (7-8), 913–925.

Mitchell, Olivia S., "Worker Knowledge of Pension Provisions," *Journal of Labor Economics*, 1988, *6* (1), 21–39.

Nishiyama, Shinichi and Kent Smetters, "Does social security privatization produce efficiency gains?," *The Quarterly Journal of Economics*, 2007, *122* (4), 1677–1719.

O'Dea, Cormac, "Inequality, Efficiency and the Design of Public Pensions," *Mimeo.*, 2019.

Ohanian, Lee, Andrea Raffo, and Richard Rogerson, "Long-term changes in labor supply and taxes: Evidence from OECD countries, 1956–2004," *Journal of Monetary Economics*, 2008, *55* (8), 1353–1362.

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*, May 2019, *109* (5), 1717–63.

__ , Manos Matsaganis, and Panos Tsakloglou, " Earnings Determination and Taxes: Evidence From a Cohort-Based Payroll Tax Reform in Greece," *The Quarterly Journal of Economics*, 01 2012, *127* (1), 493–533.

Tauchen, George, "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics letters*, 1986, *20* (2), 177–181.

van der Klaauw, Wilbert and Kenneth I. Wolpin, "Social security and the retirement and savings behavior of low-income households," *Journal of Econometrics*, 2008, *145* (1), 21–42. The use of econometrics in informing public policy makers.

World Bank, *Averting the Old Age Crisis*, Oxford University Press, 1994.

Zawisza, Tomasz, "The Taxation of Employment and Self-Employment: Evidence from Poland and Implications," *Mimeo.*, 2022.

Online Appendix A Additional Figures and Tables

Figure A.1: Polish Social Security Authority (ZUS) Annual Letter to those in NDC Scheme

ZUS

Informacja o stanie konta ubezpieczonego
w Zakładzie Ubezpieczeń Społecznych

III. Suma na Twoim koncie i subkoncie

1. Konto	294 601,34 zł
2. Subkonto	31 501,06 zł
Razem:	**326 102,40 zł**

IV. Składki należne i wpłacone (bez waloryzacji)

Subkonto:

1. Składki należne - czyli te, które wynikają z dokumentów przysłanych do ZUS	9 193,67 zł
2. Składki wpłacone na subkonto	9 193,67 zł

Otwarty fundusz emerytalny:

1. Składki należne - czyli te, które wynikają z dokumentów przysłanych do ZUS	18 400,19 zł
2. Składki odprowadzone do otwartego funduszu emerytalnego	18 400,19 zł

V. Hipotetyczna (prognozowana) wysokość Twojej emerytury

Twój wiek emerytalny to 65 lat

Emerytura obliczona na podstawie konta

Podajemy ją w dwóch wariantach:

Wiek, w którym przejdziesz na emeryturę	**Wariant I** Pracujesz do emerytury i składki wpływają średnio w tej samej wysokości co do tej pory	**Wariant II** W 2019 roku przestajesz pracować i do emerytury nie wpływają już nowe składki
65 lat	1 718,99 zł	1 348,91 zł

Ważne! Jeśli przekroczyłeś wiek emerytalny i nie masz jeszcze ustalonej emerytury, to jej prognozowaną wartość podajemy dla Twojego rzeczywistego wieku oraz dla kolejnych pięciu lat.

Emerytura obliczona na podstawie konta i subkonta

Podajemy ją w kolejnych wariantach:

Twój wiek emerytalny	**Wariant III** Pracujesz do emerytury i na konto i subkonto wpływają składki średnio w tej samej wysokości co do tej pory	**Wariant IV** W 2019 roku przestajesz pracować i do emerytury nie wpływają już nowe składki na konto i subkonto
65 lat	1 964,96 zł	1 493,14 zł

Ważne! Pamiętaj, że na 10 lat przed Twoim wiekiem emerytalnym środki z OFE będą co miesiąc wpłacane na Twoje subkonto w ZUS. Przy prognozie emerytury nie uwzględniliśmy tego, co masz jeszcze w OFE.

Co dają Ci ubezpieczenia społeczne już teraz

Ubezpieczenia społeczne poza emerytalnym obejmują także:

- ubezpieczenie chorobowe,
- ubezpieczenie wypadkowe,
- ubezpieczenia rentowe.

Ważne! Dzięki nim możesz uzyskać jednorazowe, okresowe lub stałe świadczenia w razie choroby, urodzenia dziecka i rodzicielstwa, wypadku przy pracy, niepełnosprawności, niezdolności do pracy.

Na Platformie Usług Elektronicznych na www.zus.pl możesz sprawdzić, jakim ubezpieczeniom podlegasz.

Świadczenia, które Ci przysługują, jeśli masz ubezpieczenie chorobowe

Zasiłek chorobowy w razie niezdolności do pracy z powodu choroby

2/3

Notes: An example of the annual letter sent to all individuals in the NDC system after the 1999 reform. It states the total accumulated funds in the notional account, the retirement age of the individual, and the anticipated pension at age 65 under two scenarios: 1) working until retirement age at current earnings level and 2) stopping work in the current year and not making further contributions to the NDC account.

Figure A.2: Map of Poland Showing Earnings Growth of Different Counties

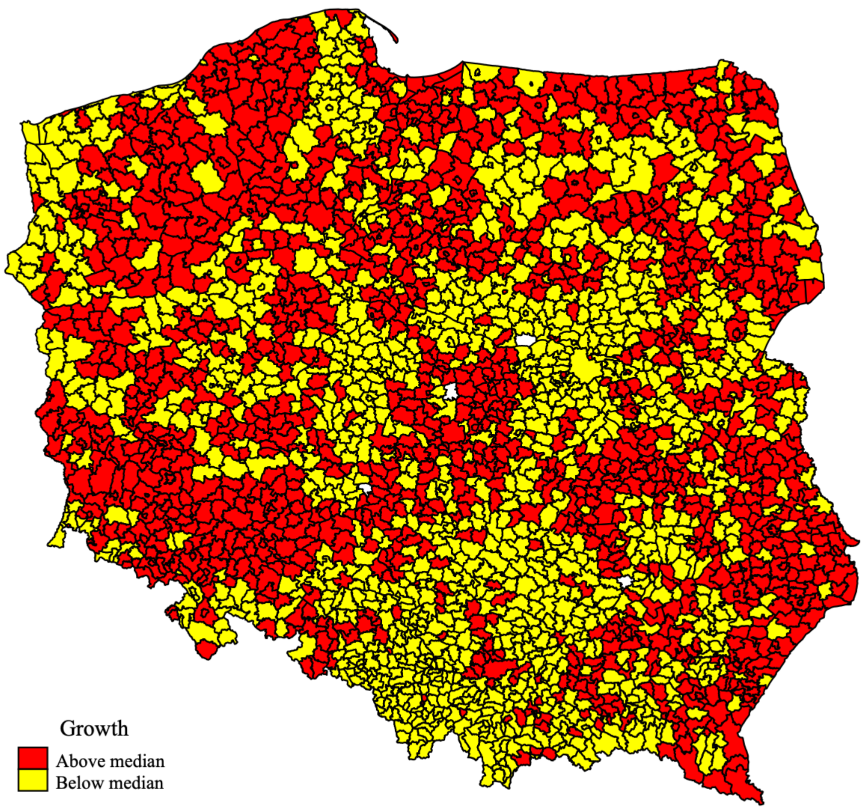

Notes: This figure shows on the map all the local districts in our analysis. For each local district we calculate average earnings growth between 2000 and 2013. Locations with below (above) median earnings growth shown in red (yellow). White regions are with missing data.

Table A.1: Summary statistics for regions

	High-growth		Low-growth	
	2000	2013	2000	2013
Lives in a rural area	29.9%	30.4%	22.2%	22.7%
Female	53.8%	54.4%	53.9%	54.4%
Employed	54.9%	58.4%	57.6%	59.1%
Sefl-employed	5.1%	4.8%	5.5%	5.0%
Avg. earnings	9,591	18,730	12,039	20,736
Avg. earnings conditional on work	17,487	32,054	20,904	35,083
Age	47.5	49.3	47.0	49.4

Notes: Earnings are in Polish zloty (in 2000 value).

Table A.2: The Effect of the Pension Reform on Employment And Wages, Robustness to Bandwidth Choice of the Local-Linear Regression

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Change in employment probability						
High-growth	-0.0174***	-0.0263**	-0.0208***	-0.0153***	-0.0155***	-0.0209***
N = 545,435	(0.0050)	(0.0130)	(0.0067)	(0.0041)	(0.0036)	(0.0058)
Low-growth	0.0027	0.0085	0.0032	0.0001	0.0004	0.0040
N = 818,487	(0.0041)	(0.0106)	(0.0055)	(0.0034)	(0.0029)	(0.0061)
Difference (High-Low)	-0.0201***	-0.0348***	-0.0240***	-0.0154***	-0.0159***	-0.0249***
	(0.0065)	(0.0168)	(0.0087)	(0.0053)	(0.0046)	(0.0084)
Panel B: Change in log wage						
High-growth	0.0119	0.0867*	0.0093	0.0073	-0.0012	0.0118
N = 545,435	(0.0174)	(0.0459)	(0.0233)	(0.0144)	(0.0125)	(0.0251)
Low-growth	0.0029	0.0117	0.0063	-0.0093	-0.0119	0.0135
N = 818,487	(0.0145)	(0.0378)	(0.0194)	(0.0120)	(0.0105)	(0.0260)
Difference (High-Low)	0.0090	0.0750	0.0030	0.0166	0.0107	-0.0017
	(0.0226)	(0.0595)	(0.0303)	(0.0187)	(0.0163)	(0.0361)
Sample	Donut	Donut	Donut	Donut	Donut	Donut
$f(z_i)$	local-linear	local-linear	local-linear	local-linear	local-linear	local-linear
net-of-placebo	no	no	no	no	no	no
Bandwidth	150	50	100	200	250	Calonico et al.

Notes: This table shows robustness of the local-linear regression estimate (shown in Column 3 in Table 1) for various choices of the bandwidth, measured in days, on each side of the discontinuity. We show the estimated change in employment (panel A) and log wage (measured as earned income of workers) for those in work (panel B) at the reform discontinuity. Each cell in the table shows the β coefficient of the RDD specification in equation (13). The rows show the estimated employment change for different regions. The first and second row show the estimated effect in high and low-growth regions, respectively. High-growth regions are regions with above median earnings growth rate between 2000 and 2013, while low-growth regions have below median growth. The third row shows the difference between the high and low-growth regions. In each column we report a kernel-weighted local linear regression, where we set the bandwidth at different levels. In Column (1) we set the bandwidth at 150 days as in Table 1. In Columns 2-5 we apply values of bandwidth between 50 and 250 days. Column (6) chooses bandwidth according to the methodology in Calonico et al. (2014), and it is 87 days for the low-growth and 120 days for the high-growth regions. In each column we apply the donut hole RDD specification where we exclude those born between December 16th and January 5th. We report robust standard errors in parentheses following the recommendation in Calonico et al. (2014). Significance levels are: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.3: The Effect of the Pension Reform on Employment And Wages, Robustness to Donut Size

	(1)	(2)	(3)	(4)
Panel A: Change in employment probability				
High-growth	-0.0148***	-0.0105***	-0.0145***	-0.0073
N = 545,435	(0.0026)	(0.0037)	(0.0031)	(0.0044)
Low-growth	0.0010	0.0014	-0.0027	0.0022
N = 818,487	(0.0022)	(0.0030)	(0.0025)	(0.0036)
Difference (High-Low)	-0.0158***	-0.0119***	-0.0119***	-0.095**
	(0.0034)	(0.0048)	(0.0040)	(0.0056)
Panel B: Change in log wage				
High-growth	-0.005	0.011	-0.020*	-0.001
N = 313,720	(0.009)	(0.0130)	(0.011)	(0.015)
Low-growth	0.005	0.018	-0.021**	0.022*
N = 439,545	(0.007)	(0.011)	(0.009)	(0.013)
Difference (High-Low)	-0.001	-0.007	0.001	-0.022
	(0.012)	(0.021)	(0.014)	(0.020)
Sample	Donut (base.)	Donut (base.)	Donut (Jan-Dec)	Donut (Jan-Dec)
$f(z_i)$	linear trend	linear trend	linear trend	linear trend
net-of-placebo	no	yes	no	yes

Notes: This table shows robustness to the change in the size of the donut in the benchmark regression specification (shown in Table 1). The table shows the estimated change in employment (panel A) and log wage (measured as earned income of workers) for those in work (panel B) at the reform discontinuity. Each cell in the table show the β coefficients of the RDD specification shown in equation (12) (Column 1 and 3) or in equation (13) (Column 2 and 4). The rows in Panel A show the estimated employment change for different regions. The first and second row show the estimated effect in high- and low-growth regions, respectively. High-growth regions are regions with above median earnings growth rate between 2000 and 2013, while low-growth regions have below median growth. The third row shows the difference between the high- and low-growth regions. In Column (1) and (2) we apply the benchmark donut hole RDD specification where we exclude those born between December 16th and January 5th. In Column (3) and (4) we apply a broader donut hole where we exclude everyone who was born in January or December. In all columns we estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. Column (2) and (4) estimates the change at the reform discontinuity relative to the change at the placebo discontinuity as in equation (13). The placebo discontinuity is estimated between the 1949 and 1950 cohorts, both of which switched to the NDC system. We report robust standard errors in parentheses. Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table A.4: Placebo Estimates on Employment and Wages for Men

	(1)	(2)
(3)		
	1947-1948	1949-1950
Panel A: Change in employment probability		
High-growth	-0.0058	-0.0043
	(0.0033)	(0.0026)
N	395,819	647,006
Low-growth	-0.0018	-0.0004
	(0.0027)	(0.0021)
N	590,418	977,555
Difference (High-Low)	-0.0040	-0.0039
	(0.0043)	(0.0033)
Panel B: Change in log wage		
High-growth	-0.013	-0.016
	(0.011)	(0.009)
N	290,184	314,948
Low-growth	-0.007	-0.022*
	(0.010)	(0.008)
N	281,750	445,816
Difference (High-Low)	-0.001	-0.008
	(0.015)	(0.012)
Sample	Donut	Donut
$f(z_i)$	linear trend	linear trend
net-of-placebo	no	no

Notes: This table shows the estimated change in employment (panel A) and log wage (measured as earned income of workers) for those in work (panel B) at two "placebo" discontinuities. Each cell in the table shows the β coefficients of the RDD specification shown in equation (12). Column (1) shows the employment and wage change between two cohorts (those born in 1947 and 1948), neither of whom were impacted by the reform, Column (2) shows between two cohorts (born in 1949 and in 1950) who were both impacted by the reform. We estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. We also apply the donut hole RDD specification where we exclude those born between December 16th and January 5th. We report robust standard errors in parentheses. Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table A.5: The Effect of the Pension Reform on Earnings Conditional on Employment

	(1) 10th percentile	(2) 25th percentile	(3) 50th percentile	(4) 75th percentile	(5) 90th percentile
High-growth	0.0870**	0.0131	0.0034	-0.0181*	0.0036
	(0.0402)	(0.0145)	(0.0112)	(0.0105)	(0.0154)
N					
Low-growth	0.0108	0.0085	0.0025	-0.0120	0.0060
	(0.0306)	(0.0162)	(0.0097)	(0.0087)	(0.0141)
N					
Difference (High-Low)	0.0761	0.0046	0.0009	-0.0061	-0.0024
	(0.0505)	(0.0218)	(0.0149)	(0.0136)	(0.0209)
Sample	Donut	Donut	Donut	Donut	Donut
$f(z_i)$	linear trend	linear trend	linear trend	linear trend	linear trend
net-of-placebo	yes	yes	yes	yes	yes

Notes: This table shows net-of-placebo estimates on the change in log-earnings distribution between the 1948 and 1949 cohorts (relative to the 1949 and 1950 cohorts). We estimate a quantile treatment effects where we condition on being employed. The columns show the estimated change in employment at different percentiles of the wage distribution. In the first and third row we report the estimated effect in high and low-growth regions. The fifth row shows the difference between the high and low-growth regions. We report robust standard errors in parentheses. Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table A.6: The Effect of the Pension Reform on Employment for Women

	(1) 1948-1949	(2) 1949-1950	(3) 1950-1951	(4) 1951-1952	(5) 1952-1953	(6) 1953-1954
Panel A: Cumulative estimate						
High-growth	0.0102***	0.0035	-0.0024	-0.0044*	-0.0120***	-0.0040
	(0.0025)	(0.0025)	(0.0025)	(0.0024)	(0.0024)	(0.0024)
Low-growth	0.0048**	0.0043**	0.0062***	0.0021	-0.0049**	0.0009
	(0.0021)	(0.0021)	(0.0020)	(0.0020)	(0.0020)	(0.0020)
Difference (High-Low)	0.0055*	-0.0008	-0.0086***	-0.0065**	-0.0071**	-0.0048
	(0.0033)	(0.0033)	(0.0032)	(0.0031)	(0.0032)	(0.0031)
Cumulative		0.0046	-0.0040	-0.0105	-0.0176**	-0.0225**
		(0.0054)	(0.0068)	(0.0078)	(0.0088)	(0.0099)
Panel B: Net-of-placebo cumulative estimate						
High-growth	-0.0020	-0.0087**	-0.0146***	-0.0166***	-0.0243***	-0.0162***
	(0.0036)	(0.0041)	(0.0038)	(0.0038)	(0.0041)	(0.0040)
Low-growth	-0.0049*	-0.0053*	-0.0035	-0.0075**	-0.0146***	-0.0088***
	(0.0029)	(0.0032)	(0.0031)	(0.0033)	(0.0031)	(0.0033)
Difference (High-Low)	0.0029	-0.0034	-0.0112**	-0.0091*	-0.0097*	-0.0074
	(0.0046)	(0.0051)	(0.0050)	(0.0051)	(0.0051)	(0.0052)
Cumulative		-0.0006	-0.0117	-0.0208	-0.0305	-0.0380
		(0.0087)	(0.0129)	(0.0169)	(0.0210)	(0.0251)
Sample	Donut	Donut	Donut	Donut	Donut	Donut
$f(z_i)$	linear trend	linear trend	linear trend	linear trend	linear trend	linear trend

Notes: This table shows the change in employment for each cohort discontinuity between 1948 and 1954 for women. Contrary to the men, the the Pension Reform for women was introduced gradually, and it was phased-in step-by-step between cohorts 1949 and 1954 by increasing the fraction of pension coming from NDC. Panel A shows the estimated change in employment at all the relevant cohorts where there is an increase in the NDC share (β coefficient of the RDD specification in equation (12)). Panel B show the employment change relative to the change at the placebo discontinuity (1948 and 1947 cohorts). We report the (β coefficient of the RDD specification in equation (13)). In all other respect we follow our benchmark RDD analysis (donut linear, column 4 in Table 1). The rows show the estimated employment change for different regions. The first and second row show the estimated effect in high and low-growth regions, respectively. High-growth regions are regions with above median earnings growth rate between 2000 and 2013, while low-growth regions have below median growth rate. The third row shows the difference between the high and low-growth regions. In each column we apply the donut hole RDD specification where we exclude those born between December 16th and January 5th. the cumulative effects add up the estimated employment response difference between high and low growth regions up to that cohorts. As a consequence, Column (6) shows to total effect of switching from DB to NDC. We report robust standard errors in parentheses. Significance levels are: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.7: The Effect of the Pension Reform on Employment And Wages, by Employment and Self-Employment

	(1)	(2)	(3)	(4)	(5)
Panel A: Change in employment probability					
High-growth	-0.0204***	-0.0156***	-0.0175***	-0.0098***	0.0005
N = 545,435	(0.0024)	(0.0026)	(0.0050)	(0.0037)	(0.0018)
Low-growth	0.0013	0.0048	0.0038	0.0056*	-0.0033
N = 818,487	(0.0020)	(0.0021)	(0.0040)	(0.0030)	(0.0015)
Difference (High-Low)	-0.0217***	-0.0204***	-0.0213***	-0.0154***	0.0038
	(0.0031)	(0.0033)	(0.0064)	(0.0048)	(0.0023)
Panel B: Change in log wage					
High-growth	-0.014*	-0.016*	-0.011	-0.001	0.061
N = 313,720	(0.008)	(0.009)	(0.016)	(0.012)	(0.042)
Low-growth	-0.014**	-0.017**	-0.005	0.009	0.027
N = 439,545	(0.007)	(0.007)	(0.014)	(0.010)	(0.035)
Difference (High-Low)	0.000	0.000	-0.006	0.009	0.035
	(0.010)	(0.011)	(0.021)	(0.016)	(0.054)
Sample	Full	Donut	Donut	Donut	Donut
$f(z_i)$	linear trend	linear trend	local linear	linear trend	linear trend
net-of-placebo	no	no	no	yes	yes
Type of employment	Empl.	Empl.	Empl.	Empl.	Self-empl.

Notes: This table shows the estimated change in employment (panel A) and log wage (measured as earned income of workers) for those in work (panel B) at the reform discontinuity for employed and self-employed separately. Columns 1-4 shows the estimated employment change when the outcome variable is 1 only if the worker is in employment (and not in self-employment). In column 5, the outcome variable is 1 if the worker is in self-employment. Each cell in the table shows the β coefficient of the RDD specification shown in equation (12) (Columns 1-3) or in equation (13) (Column 4-5). The rows show the estimated employment change for different regions. The first and second row show the estimated effect in high and low-growth regions, respectively. High-growth regions are regions with above median earnings growth rate between 2000 and 2013, while low-growth regions have below median growth. The third row shows the difference between the high and low-growth regions. In Column (1) we use the full dataset. In Columns (2)-(4) we apply the donut hole RDD specification where we exclude those born between December 16th and January 5th. In Columns (1), (2) and (4) we estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. Column (3) estimates a kernel-weighted local linear regression, where we set the bandwidth at 150 days. Column (4) and (5) estimate the change in employment at the reform discontinuity relative to the change at the placebo discontinuity as in equation (13). The placebo discontinuity is estimated between the 1949 and 1950 cohorts, both of which switched to the NDC system. We report robust standard errors in parentheses. For the local-linear regression we calculate robust standard errors following Calonico et al. (2014). Significance levels are: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.8: The Effect of the Pension Reform on Employment throughout the Wage Distribution

	(1) Employed All	(2) Employed Below 2x min wage	(3) Employed Above 2x min wage
High-growth	-0.0105***	-0.0048	-0.0065**
N = 497,474	(0.0037)	(0.0029)	(0.0031)
Low-growth	0.0014	-0.0019	-0.0039
N = 745,349	(0.0030)	(0.0022)	(0.0026)
Difference (High-Low)	-0.0119***	-0.0029	-0.0104***
	(0.0048)	(0.0037)	(0.0040)
Share of all employment, 2000-2002	100%	42.66%	57.34%
Sample	Donut	Donut	Donut
$f(z_i)$	linear trend	linear trend	linear trend
net-of-placebo	yes	yes	yes

Notes: This table shows the estimated change in employment separately for jobs below two times the minimum wage and above two times the minimum wage. Each cell in the table shows the β coefficients of the RDD specification shown in equation (12). Column 1 use a binary outcome variable of being employed just as in Column 4 in Table 1. Column 2 (3) use a binary outcome variable of being employed and earnings below (above) two times the earnings of a full-time, full-year minimum wage worker (around the median earnings in our sample). We estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. Note the parameters on the trend in birth date differ between between columns (1), (2), and (3) and so the sum of coefficients in column (2) and (3) does not neccessarily add up to column (1). We also apply the donut hole RDD specification where we exclude those born between December 16th and January 5th. We report robust standard errors in parentheses. Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table A.9: Employment Elasticity throughout the Wage Distribution

	(1)	(2)	(3)	(4)	(5)	(6)
	Below 2x min wage			Above 2x min wage		
	High-growth	Low-growth	Difference (H-L)	High-growth	Low-growth	Difference (H-L)
1. Change in net return to work (%)	-19.06	-16.83	-2.23	-7.54	-3.16	-4.38
2. Change in pension wealth (%)	-14.23	-14.99	0.76	-13.44	-11.88	-1.55
3. Change in employment (%)	-2.13	-1.01	-1.13	-2.39	1.37	-3.76
	(1.29)	(1.16)	(1.74)	(1.14)	(0.91)	(1.46)
4. Employment elasticity (Row 3)/(Row 1)	–	–	0.50	–	–	0.86
	–	–	(1.78)	–	–	(0.33)

Notes: This table shows the effect of the pension reform on the net return to work (row 1), on the pension wealth (row 2), on the change in employment (row 3) and on the resulting employment elasticity (row 4). The table follows a similar structure as Table 2, but we implement the analysis separately for individuals with earnings below two times the earnings of a full-time, full-year minimum wage worker in Columns 1-3 and above two times the earnings of a full-time, full-year minimum wage worker in Columns 4-6. The percent change in the net return to work is calculated separately for each group using the same procedure used in Table 2. To calculate the percent change in employment, we divide the net-of-placebo employment estimates from Table A.8, by the probability of being employed below/above two times of the minimum wage for the cohorts which were age 50 at the time of the reform and so stayed in the DB system. Robust standard errors are in parentheses. It is worth highlighting that these net-return to work changes are somewhat smaller than those estimated for the full sample in Table 2 (5.23% in Column 3). This comes from how we calculate the percent change in the net return to work. We calculate it by taking the actual zloty difference in the net return to work for each individual and then average this over all members of the group. Then we divide that by the average wage of that group to obtain the percent change in the net return to work. Thus we are calculating a ratio of averages (rather than an average of ratios) and thus there is no reason that the overall percent change in the net return to work is the average of the subgroup percent changes in the net return to work. For example, high wage individuals contribute much more to the overall level change in the net return to work, and thus it should be no surprise that the overall percent change in the net return to work is closer to the percent change in the net return to work for high wage individuals.

Online Appendix B Household Survey Data

Beyond our administrative data, we use two additional data sources: the Polish Household Budget Survey (HBS) and the Polish Labor Force Survey (LFS). Both of these data sources are large representative household surveys. However, they are many times smaller than our administrative data and have too few observations from each cohort for us to apply the RDD empirical strategy to study the labor supply responses to the NDC reform.

The HBS includes a rotating panel where a household is interviewed in two consecutive years. This allows us to observe individuals' transitions between unemployment and employment across years, where unemployment is measured as receipt of unemployment benefits. Moreover, the HBS has detailed information about the type of income individuals receive. We make use of this to estimate the fraction of individuals who receive retirement pensions before the full or early retirement ages, and who are therefore in excluded sectors and special occupations not affected by the pension reform.

The LFS is a large survey studying employment outcomes in the Polish population. Since the LFS captures all types of employment, including agricultural labor which does not appear in our administrative data, we are able to estimate employment rates with and without agricultural workers (see Appendix Table B.1). The data also contain some information on tenure, firm size and the type of employment contract for individuals in work. We exploit this information to estimate whether an individual was eligible for old age unemployment benefit(s) following the 2004 reform (see Online Appendix E).

Table B.1: Employment Rate for Men at Various Ages in the Administrative Data, Labor Force Survey and Household Budget Survey, 2000-2002.

	(1)	(2)	(3)	(4)	(5)
	Non-Agricultural Empl./Pop.		Total Empl./Pop.		Unempl./Pop.
	Admin. Data	LFS	LFS	HBS	LFS
1. Age 51 (cohorts: 1949-1951)	0.51	0.52	0.65	0.66	0.11
Number of Observations	890,620	5,141	5,141	2,299	5,054
2. Age 51-54 (cohorts: 1948-1949)	0.49	0.48	0.60	0.61	0.10
Number of Observations	1,669,539	9,485	9,485	4,254	9,410
3. Age 21-64 (cohorts: all)	0.56	0.56	0.68	0.68	0.12
Number of Observations	11,662,286	64,143	64,143	30,959	63,539
4. Age 50-54 (cohorts: all)	0.50	0.52	0.65	0.63	0.09
Number of Observations	3,907,140	7,771	7,771	10,811	7,678

Notes: This table compares the estimated employment rate for men in our administrative data with the employment rates in representative household surveys using the years between 2000 and 2002. Row 1 reports the employment to population rate for those who were 51 years old between 2000 and 2002, row 2 for those who were between 51 and 54 years old and born around the discontinuity (the 1948 and 1949 cohorts) and row 3 for those who were between age 21 and 64. We focus on non-agricultural employment as we do not observe agricultural employment in the administrative data. Note that agricultural workers belong to a separate pension system and are unaffected by the pension reform. Column (1) reports total non-agricultural employment divided by the total population for all males. Column (2) reports the same using the Labor Force Survey (LFS), the largest household study in Poland that provides the official measures of employment and unemployment. Columns (3) and (4) calculate total employment (including agricultural workers) divided by the total population in the LFS and Household Budget Survey (HBS) for the different age groups.

Online Appendix C Net Return to Work Calculation

C.1 Change in net return to work from DB to NDC switch

In this section we provide a detailed description of how we calculate the net return to work. The percent change in net return to work (which is also shown in equation (3)) is:

$$\frac{\Delta nrw_t}{nrw_t^{DB}} = \frac{E_t[\Delta PV_{it}^{NDC}] - E_t[\Delta PV_{it}^{DB}]}{E_t[(1-\tau(\tau^{pi},\tau^{ss}))\cdot w_{it} - u_{it}] + E_t[\Delta PV_{it}^{DB}]} \tag{C.1}$$

We calculate $E_t[(1-\tau(\tau^{pi},\tau^{ss}))\cdot w_{it} - u_{it}]$ as follows.

Poland has a flat income tax rate that has changed only slightly since before the reform we study, and so we set $\tau^{pi} = 0.19$. There is a higher marginal tax rate for those making approximately 2.5 $\times$ average earnings. However, at a similar earnings level the individual hits the cap on social security contributions, making the tax function effectively linear. Income taxes apply to income net of social security contributions (paid by the employees). As a result, the after-tax income if working is:

$$w_{it}(1-\tau(\tau^{pi},\tau^{ss})) = w_{it}(1-\tau^{pi}(1-\tau^{ss}) - \tau^{ss}). \tag{C.2}$$

We calculate the average after-tax wages using our administrative data. The average out of work benefit, $E_t[u_{it}]$, is calculated as follows. Those working last period are eligible for an unemployment benefit for the first six months of unemployment, which is a flat taxable benefit equaling approximately 40% of the average wage income of workers. After six months, the individual receives a flat untaxable welfare benefit equaling approximately 35% of the average wage income of workers. Given that our model is an annual one, we assume that if they worked last period, the individual receives unemployment for 50% of the year and welfare for 50% of the year, so $u_{it} = [(1-\tau^{pi})\cdot 0.5\cdot 0.4\cdot \overline{w}_t] + [0.5\cdot 0.35\cdot \overline{w}_t]$. Those not working last period receive a welfare benefit of $u_{it} = [0.35\cdot \overline{w}_t]$. The actual numbers of the tax formula that we use are presented in Table C.1.

We calculate the $E_t[\Delta PV_{it}^{NDC}]$ and $E_t[\Delta PV_{it}^{DB}]$ as follows. First, we simulate 4,000 life cycle earnings profiles following the procedure described in Section 4. In the first step, we simulate wages according to equation (10). This equation has a deterministic component capturing the age profile of wages and a stochastic component of AR(1) + MA(1) as specified in equation (11). In the second step, we simulate unemployment spells according to a Markov process. The parameter values of the benchmark simulation are shown in Table C.2. Table 3 shows that our estimated labor supply elasticity is robust to different parameter values for the wage process.

For each individual, we calculate the pension benefit at age 65 given their lifecycle earnings and the pension rules. To determine $b_{i65}^{\text{Employed}_t,NDC}$, we calculate the starting capital at age 50 and then apply the pension formulas described in equations (5) and (6). Panel C of Table C.1 shows the parameter values for the pension uprating rate r^{NDC} and life expectancy $E[T|t=65]$ that we used for the calculations. We calculate $b_{i65}^{\text{Not Employed}_t,NDC}$ using individual i's life cycle profile but assuming they were not working at age t.

To calculate $b_{i65}^{\text{Employed}_t,DB}$, we apply the pension formulas described in equation (4). As before, we calculate $b_{i65}^{\text{Not Employed}_t,DB}$ by taking the life cycle profile but assuming that the individual is not working at age t. For both pension systems, we apply the minimum pension if someone's pension level is below that threshold (Panel B of Table C.1 shows the value of the minimum pension).

Then we apply equation 9 to calculate the present value of $b_{i65}^{\text{Employed}_t,k} - b_{i65}^{\text{Not mployed}_t,k}$.[44] The real pension uprating rate r^{index} and the real risk free interest rate r are shown in Part D of Table C.1. The survival rates $S_{s|t}$ at age t are calculated from the survival probabilities in official life-tables. The figure shows at each age the probability of surviving one more year.

In Appendix Table C.3, we also provide further detail on what drives the differential incentives across regions. In Column (1) we simulate the change in the incentives if we keep the initial wage at age 30 and the trend (wage growth) the same across regions but allow the shape of age profiles to differ. Since the shape of age profiles is very similar across locations the differential change in incentives is close to zero. In Column 2, we also allow initial wages to differ, but assume the wage growth is the same across locations. We see that the initial differences in wages contribute to the incentives differences between locations. In Column (3) we report the benchmark estimates allowing differential initial wages, trend, and shape of age profile, where the parameters of the profiles are shown in Table C.2. The table shows that the difference in the change in net return to work is driven partly by different levels of wage between high- and low-growth regions, and partly by differences in the estimated time-trend parameters. Column (4) shows that the incentive change is similar if we estimate the wage process with individual fixed effects. Finally, column (5) includes regional differences in the stochastic component of wages where the estimated wage shock parameters are shown in Table C.2. It shows that adding region specific stochastic components modestly reduces differences in the net return to work between high and low growth regions.

In Figure C.1 we plot the pension accrual from working at each age between 22 and 64 for the steady state scenario where someone is in the DB or NDC system for their entire career. The figure shows that the DB system incentivizes employment in the later part of the life cycle, when earnings are higher and therefore they count more for the best years calculations. The differences between low and high-growth earnings profiles are also more pronounced for the DB system, especially in the later part of the life cycle.

For example, the chance that ages 51-54 will be part of the best years calculations in the DB system is 49% in high growth regions and 45% in low growth regions. In addition, the increase in future DB pension benefits as a result of working, conditional on being in one's best years, is much higher in high growth regions than in low growth regions. This is because the wages of those in their best years are much higher than other points in the life cycle, and this is especially true for those experiencing high wage growth. In particular, among those who are in their best years, the decision to work (instead of not work) increases future benefits 889 (784) Zl in the DB system in high (low) growth regions. In contrast, NDC system benefits are proportional to earnings. Because high growth regions have lower average earnings than low growth regions, the increase in future pension benefits as a result of working is smaller. In particular, among those who are in their best years in high (low) growth regions, the decision to work increases future benefits 468 (540) Zl in the NDC system. Therefore, the earnings gain conditional on that it is one of the best years, is 889-468= 421 Zl in the high-growth region, while it is 42% lower in low-growth regions (784-540=244 Zl). For those not in their best years, pension benefit accrual is similar across regions in both the DB and NDC systems. In the DB system, the increase in pension benefits from work are 211 (245) Zl, and in the NDC system, the increase in benefits is 172 (208) Zl for high (low) growth regions.

[44] If someone is unemployed we assume that $b_{i65}^{\text{Employed}_t,k} - b_{i65}^{\text{Not mployed}_t,k} = 0$.

Table C.1: Parameter Values Used for Calculating the Net Return to Work

Panel A: parameters of the tax function	**Value**	**Source**
τ^{pi} (income tax)	0.19	Official tax
τ^{ss} in eq. (4) (soc. sec. worker cont. rate)	0.1871	and social security rates
τ^{tpcr} in eq. (6) (total pension cont. rates)	0.1952	
u_{it} (unemployment benefit and welfare, working previous period)	$(1-\tau^{pi}) \cdot 0.5 \cdot 0.4 \cdot \bar{w}_t$ $+0.5 \cdot 0.35 \cdot \bar{w}_t$	Official rates
u_{it} (welfare, not working previous period)	$0.35 \cdot \bar{w}_t$	
Panel B: parameters of the pension system		
r^{index} (pension uprating rate)	0.0116	Off. rates, (2000-20)
Minimum pension	$0.2 \times \bar{w}_{65}$	Off. rates & administrative tax data
Panel C: NDC-specific parameters		
r^{NDC} (NDC contribution uprating rate)		Calculated from the aggregate earnings growth
$E[T \vert t = 65]$ (life expectancy at age 65)	212.4 months	Off. life-tables
Panel D: discounting parameters		
r (risk-free rate)	0.0288	10-year gov. bonds (2000-19)
$S_{s \vert t}$ (survival probability)		Off. life-tables

Notes: This table shows the parameter values used for calculating the net return to work defined in equation (1).

Table C.2: Parameter Values Used for Simulating the Earnings Process

Parameter	Value		Source
Panel A: region-specific deterministic wage component			
	Low	High	
constant	-2.4688	-3.2599	Estimation of (10)
age	1.0433	1.1344	(dataset: administrative tax data)
age^2	-0.0328	-0.0370	
age^3	4.58E-04	5.34E-04	
age^4	-2.41E-06	-2.89E-06	
t	0.04338	0.0517	
Panel B: stochastic wage component			
ρ	0.9496		GMM estimation of (11)
θ	-0.2353		(dataset: administrative tax data)
σ^2_ε	0.0591		
σ^2_ξ	0.0276		
Panel C: region-specific stochastic wage component			
	Low	High	
ρ	0.9610	0.9535	GMM estimation of (11)
θ	-0.2795	-0.2801	(dataset: administrative tax data)
σ^2_ε	0.0476	0.0463	
σ^2_ξ	0.0146	0.0115	
Panel D: unemployment process			
$Pr(UI_{it} = 1 \vert P_{it} = 1)$	0.0340		Estimation of Markov process
$Pr(P_{it} = 1 \vert UI_{it} = 1)$	0.4031		(dataset: HBS)

Notes: This table shows the parameter values used for simulating the lifecycle earnings profiles. Panel A shows the region-specific deterministic wage component of the wage process (see equation (10)). We estimate these values from the administrative tax data between 2000 and 2013. Panel B shows the parameter values of the stochastic components in the wage equation (see equation (11)). We estimate these values using a GMM estimator from the administrative tax data between 2000 and 2013. Panel C shows the parameter values of the unemployment process. These values are estimated from the Household Budget Survey (HBS), where we directly observe unemployment.

Table C.3: Change in net return to work (%) under different parameterizations of the wage process.

	(1)	(2)	(3)	(4)	(5)
1. High-growth	-7.30	-7.30	-8.82	-9.53	-8.31
2. Low-growth	-6.99	-4.35	-4.35	-4.46	-4.55
3. Difference (High-Low)	-0.31	-2.95	-4.46	-5.08	-3.76
Differential shape of age-profile	Yes	Yes	Yes	Yes	Yes
Differential level		Yes	Yes	Yes	Yes
Differential shape of time-trend			Yes	Yes	Yes
Individual fixed-effects				Yes	
Differential wage shock process					Yes

Notes: This table shows the effect of the pension reform on the net return to work for different hypothetical parameterizations of the deterministic component of the wage process. The purpose of columns (1)-(3) of this table is to illustrate what part of the wage process drives the difference in the change in the net return to work between the high and low-growth regions. Column (1) shows the net return to work in a scenario in which the high-growth and low-growth regions share the same trend rate of growth - assumed to be that estimated for the low-growth region in both cases - but have the different parameterizations of the age polynomial, which is set to their estimated values. The intercept of the low-growth wage process is adjusted such that individuals in the two regions have the same wages on average over the lifecycle. Column (2) shows the net return to work in which the high-growth and low-growth regions share the same trend rate of growth but have the estimated intercept and parameters of the age polynomial. Column (3) shows the net return to work in which the high-growth and low-growth regions differ in terms of time trend, the intercept, and parameters of the age polynomial, which are assigned their estimated values. This is also our baseline wage process. Column (4) shows the net return to work in which the high-growth and low-growth region parameters were estimated using an individual fixed-effects specification, where the level of fixed effects is chosen to match the average observed earnings of the 1948-1949 cohorts in 2000-2002. Finally, column (5) adds to the specification in column (3) region-specific stochastic component to the wages (with parameter estimates shown in the previous table).

Table C.4: Region-specific deterministic wage component parameters for different specifications of the wage regression.

Parameter	(1)	(2)	(3)	(4)
	Baseline		Individual FE	
constant	-2.4688	-3.2599	–	–
	(0.0626)	(0.0714)		
age	1.0433	1.1344	0.9362	0.9901
	(0.0066)	(0.0076)	(0.0080)	(0.0093)
age^2	-0.0328	-0.0370	-0.0286	-0.0309
	(0.0003)	(0.0003)	(0.0003)	(0.0004)
age^3	4.58E-04	5.34E-04	4.07E-04	4.51E-04
	(4.19e-06)	(4.80e-06)	(4.94e-06)	(5.70e-06)
age^4	-2.41E-06	-2.89E-06	-2.20e-06	-2.49e-06
	(2.51e-08)	(2.93e-08)	(2.87e-08)	(3.38e-08)
t	0.04338	0.0517	–	–
	(0.0001)	(0.0001)		
Region growth type	Low	High	Low	High
Implied annual wage growth	3.34	4.09	2.93	3.56
Difference (High-Low)	0.75		0.63	
Individual FE	No	No	Yes	Yes
N	6,692,578	4,887,578	6,689,495	4,885,520

Notes: This table shows estimates of the parameter values of the deterministic component of the wage process for different estimation methods. Columns (1) and (2) show the parameter estimates for the low and high-growth regions, respectively, under our baseline specification (see equation (10)), while columns (3) and (4) show the parameter estimates for the low and high-growth regions, respectively, from a fixed-effects regression. Here, the coefficient on the time trend is not identified separately from the coefficient on the first element of the age polynomial, and so only the latter is reported. The table also shows the real annual wage growth for the 1948-1949 cohorts between ages 21 and 64 implied by these parameter estimates. This shows that the implied annual wage growth is higher by a similar amount both under the levels and fixed effects specification, demonstrating that higher wage growth in high-growth regions is not driven by changes in the composition of the labor force with age. Standard errors are reported in parentheses.

Figure C.1: Change in the present value of pension benefits by age

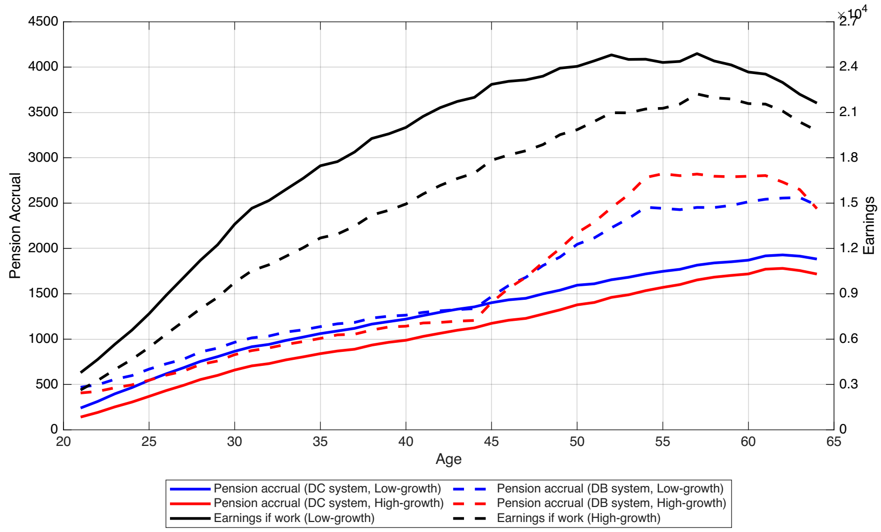

Notes: This figure shows the average pension accrual from working at every age between ages 22 to 64 for those in the DB and NDC system, by region type (left-hand axis). Pension accrual is defined as the increase in present value of pension benefits as a result of working in a given year. It is calculated using our baseline simulations for the steady-state scenario where individuals are in the DB or NDC systems for their entire working life. Additionally, the figure shows the average earnings if individuals work for the high and low-growth regions (right-hand axis). All values are expressed in Polish zloty (valued in 2000) and use 10,000 simulations.

C.2 Employment Elasticity Calculations: Accounting for those Unaffected by the Reform

We are interested in the employment response to the NDC reform for those affected by the reform. However, some individuals at the reform discontinuity were unaffected by the reform because they were either in an excluded sector or special occupation, and thus remained in the DB scheme. In our benchmark elasticity formula (presented in Equation (2)) we abstract away from this issue. Here we discuss how the presence of unaffected workers affects the calculation of the employment elasticity.

There are two groups of workers unaffected by the reform. First, agricultural workers are unaffected by the reform and their employment is not observed in our administrative data. The indicator variable $A_{it} = 1$ denotes that someone is working in the agriculture sector. Besides agricultural workers, we also have individuals working in excluded occupations. While these workers' employment is observed in our administrative data, we do not observe whether someone is in an exempted occupation or not. We denote (non-agricultural) workers in excluded occupations with $E_{it} = 1$. The pension reform affects workers in the non-agricultural sector, $A_{it} = 0$, and in non-exempted occupations, $E_{it} = 0$.

The employment elasticity for the affected groups can be written as:

$$\eta \quad = \quad \frac{\frac{P_t^{NDC} - P_t^{DB}}{P_t^{DB}}}{\Delta nrw_t / nrw_t^{DB}} \tag{C.3}$$

where now $P_t^{NDC} = Pr(P_{it} = 1 | A_{it} = 0, E_{it} = 0, z_i < 50)$ is the probability of work among those now facing the NDC scheme, $P_t^{DB} = Pr(P_{it} = 1 | A_{it} = 0, E_{it} = 0, z_i > 50)$ the probability of work facing the DB scheme (but also not in an excluded occupation), P_{it} is a binary variable (working or not working) and so we can express its (conditional) expected value as a conditional probability, $E[P_{it} | A_{it} = 0, E_{it} = 0, z_i] = Pr(P_{it} | A_{it} = 0, E_{it} = 0, z_i)$, for $z_i \in \{z_i < 50, z_i > 50\}$. The numerator of equation (C.3) is thus the percentage change in employment conditional on being affected by the reform.

The baseline employment probability under the DB system conditional on being in sectors or in occupations affected by the reform, $Pr(P_{it} | A_{it} = 0, E_{it} = 0, z_i > 50)$, is the following:

$$Pr(P_{it} = 1 | A_{it} = 0, E_{it} = 0, z_i > 50) = \frac{Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 0) | z_i > 50)}{Pr(A_{it} = 0 \cap E_{it} = 0 | z_i > 50)}, \tag{C.4}$$

where $Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 0) | z_i > 50)$ shows the probability that individual i is employed at time t ($P_{it} = 1$) and is in a non-agricultural and non-exempted occupation ($A_{it} = 0 \cap E_{it} = 0$) and older than age 50.

Our regression discontinuity estimates in equation (12) identify the change in non-agricultural employment at the discontinuity, formally:

$$\beta = \Pr(P_{it} = 1 \cap A_{it} = 0 | z_i < 50) - \Pr(P_{it} = 1 \cap A_{it} = 0 | z_i > 50), \tag{C.5}$$

where β is the estimated parameter from the regression equation (12), $\Pr(P_{it} = 1 \cap A_{it} = 0 | z_i < 50)$ is the probability of working in the non-agricultural sector for the younger cohorts (which contain both affected workers ushered into the new NDC pension system and exempted workers who stayed in the DB pension scheme) and $\Pr(P_{it} = 1 \cap A_{it} = 0 | z_i > 50)$ is the probability of working in the non-agricultural sector for the older cohorts, all of whom stayed in the DB pension scheme. The change in non-agricultural employment at the discontinuity can be rewritten as

$$
\begin{aligned}
&\Pr(P_{it} = 1 \cap A_{it} = 0 | z_i < 50) - \Pr(P_{it} = 1 \cap A_{it} = 0 | z_i > 50) = \\
&\quad = \quad \Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 0) | z_i < 50) + \Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 1) | z_i < 50) - \\
&\qquad\quad - [Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 0) | z_i > 50) + \Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 1) | z_i > 50)] \\
&\quad = \quad \Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 0) | z_i < 50) - \Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 0) | z_i > 50) \\
&\quad = \quad [Pr(P_{it} = 1 | A_{it} = 0, E_{it} = 0, z_i < 50) - Pr(P_{it} = 1 | A_{it} = 0, E_{it} = 0, z_i > 50)] \times \\
&\qquad\quad \times Pr(A_{it} = 0 \cap E_{it} = 0), \qquad \text{(C.6)}
\end{aligned}
$$

where the second equality uses the fact that young workers stayed in the DB system only if they worked in an exempted occupation and so for them $\Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 1) | z_i < 50) = \Pr(P_{it} = 1 \cap (A_{it} = 0 \cap E_{it} = 1) | z_i > 50)$. The third equality uses the definition of conditional probability (similarly to Equation (C.4)) and that the probability of being in an exempt occupation or of being an agricultural worker is unlikely to differ at the discontinuity and so $Pr(A_{it} = 0 \cap E_{it} = 0 | z_i < 50) = Pr(A_{it} = 0 \cap E_{it} = 0) Pr(A_{it} = 0 \cap E_{it} = 0 | z_i > 50)$.

Inserting equations (C.6) and (C.4) into the percentage change in employment conditional on being in sectors or in occupations affected by the reform (the numerator of equation (C.3)) yields:

$$
\begin{aligned}
&\frac{Pr(P_{it} = 1 | A_{it} = 0, E_{it} = 0, z_i < 50) - Pr(P_{it} = 1 | A_{it} = 0, E_{it} = 0, z_i > 50)}{Pr(P_{it} | A_{it} = 0, E_{it} = 0, z_i > 50)} = \\
&\qquad = \frac{Pr(P_{it} = 1 \cap A_{it} = 0 | z_i < 50) - Pr(P_{it} = 1 \cap A_{it} = 0) | z_i > 50)}{Pr(P_{it} | A_{it} = 0, E_{it} = 0, z_i > 50) \times Pr(A_{it} = 0 \cap E_{it} = 0)} \\
&\qquad = \frac{Pr(P_{it} = 1 \cap A_{it} = 0 | z_i < 50) - Pr(P_{it} = 1 \cap A_{it} = 0 | z_i > 50)}{Pr(P_{it}^{DB} = 1 \cap A_{it} = 0 \cap E_{it} = 0 | z_i > 50)} \\
&\qquad = \frac{\beta}{Pr(P_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 0 | z_i > 50)}. \qquad \text{(C.7)}
\end{aligned}
$$

Equation (C.7) shows that to obtain the employment elasticity for the affected population, we need to divide the estimated percentage point change in non-agricultural employment around the discontinuity, β (which is what is recovered by our RDD estimator) by the fraction of the population working in non-agricultural and non-exempted occupations, $Pr(P_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 0 | z_i > 50)$. This latter object is not directly observed in the administrative data, as we do not know who is exempt because they work in a special occupation. We calculate this object by first using the Law of Total Probability and rearranging:

$$
\begin{aligned}
&Pr(P_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 0 | z_i > 50) = \qquad \text{(C.8)} \\
&= Pr(P_{it} = 1 \cap A_{it} = 0 | z_i > 50) - Pr(P_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 1 | z_i > 50).
\end{aligned}
$$

The administrative data shows that the fraction of people working in the non-agricultural sector is $Pr(P_{it} = 1 \cap A_{it} = 0|z_i > 50) = 49\%$.

We infer the fraction of the population that is working in the non-agricultural sector in an excluded occupation (and thus exempt from the reform, regardless of age), $Pr(P_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 1|z_i > 50) = Pr(P_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 1|z_i < 50)$, in the following way. Workers born in 1949 were exempt from the switch to the NDC system if they worked long enough in a special occupation (e.g. metal workers, teachers etc.) and claimed retirement benefits before 2008. Using the Household Budget Survey, we estimate that 11% of the 1949 cohort (who were 49 years old at the time of the reform) claimed retirement benefits by 2008. These are the individuals who were exempt from the pension reform since they would have otherwise been unable to collect benefits. We then determine what fraction of these individuals were employed over the 2000-2002 time period – the period used for estimating employment response to the policy change. We estimate that 6% of the younger cohort were already drawing retirement benefits and were not employed in 2002. We do not know if the remaining 5% who began drawing retirement benefits between 2002 and 2008 were employed or not in 2002. In the benchmark case, we assume that *none* of these workers were employed, meaning that $Pr(P_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 1|z_i < 50) = 0$, and so we get $Pr(P^{DB}_{it} = 1 \cap A_{it} = 0\% \cap E_{it} = 0|z_i < 50) = 49\%$ according to equation (C.8). This leads to an estimated employment elasticity of 0.51 (see Table 2). As a robustness exercise, we assume that *all* individuals who began drawing retirement benefits between 2002 and 2008 were employed in 2002 and thus $Pr(P^{DB}_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 1|z_i < 50) = 5\%$. In that case, $Pr(P_{it} = 1 \cap A_{it} = 0 \cap E_{it} = 0|z_i < 50) = 49\% - 5\% = 44\%$ and so the employment elasticity is 0.57 (see Panel E of Table 3).

Online Appendix D Number of Observations by Birth Date and Covariance Balance

Our empirical strategy exploits date of birth: individuals born before January 1st, 1949 stayed in the DB system, while younger individuals switched to NDC. Since their birth dates were determined many years before the policy change, individuals close to the discontinuity could not have manipulated their eligibility in response to the policy.

In Panel (a) of Figure D.2 we plot the number of observations by birth month for the 1946-1953 cohorts in the years 2000-2002. Even though manipulation is not possible, there is a clear spike in reported births which occurs on the 1st of January of every year. Nevertheless, the spikes at January 1st were also present before the policy change, as demonstrated in Figure D.1, which shows the frequencies observed in the pre-reform year 1998. In fact, this spike at January 1st most likely reflects that many in these cohorts were born at home and self reported their date of birth, not that there was a higher coincidence of hospital births on the first day of the year. Since the cut-off enrollment at schools was 31st December/1st January, some parents strategically reported their children at the beginning of the calendar year so that their child would be among the oldest in the class when they started school. While this reporting behaviour took place 50 years before the pension reform was announced, the characteristics of these switchers may be correlated with the labor-market outcomes we care about.

To alleviate this concern about the bunching at January 1st, we apply various donut hole RDD estimates throughout the paper. In particular, our baseline specification drops individuals who were born between December 16th and January 5th. We picked these thresholds visually; there is an excess of births between December 16-31, and an absence of births between January 1-5, but very little evidence of excess births or an absence of births outside of this range. For robustness we also apply a broader donut hole where we drop all individuals who were born in January and December (see Table A.3). Panels (d) and (e) of Figure D.2 and Figure D.1 show the number of observations for the baseline and the broader donut hole, respectively. The unusual January and December effects disappear for both definitions of donut and we observe a seasonal pattern in birth rates observed in other countries (see e.g. Buckles and Hungerman (2013)). There is also a time trend which reflects a post-World War 2 baby boom in Poland.

In panels b, d, and f of Figure D.2 and Figure D.1, we investigate whether there is an abnormal mass of observations at our reform discontinuity above the predictable pattern of monthly births and cohort trend. We regress the number of observations at each birth month on the month-of-birth dummies and a linear trend for the cohorts born between 1944 and 1953. We plot the residuals under each donut assumption. The residuals vary across birth months, but there is no unusual jump or drop in the number of observations between the 1948 and 1949 cohorts.

We also study whether covariates substantially differ between cohorts. Table D.1 shows the regression discontinuity estimates based on equation (12), where the outcome variables are various covariates. In panel A, we assess whether there is a change in rural/urban share and share of women. These are the only covariates that we observe in our administrative data. In panel B, we show the change in the region characteristics where individuals live. We present two characteristics: regional level log household income and employment rate in 1998. We measure these outcomes in 1998, before the reform took place.

Column (1) shows the difference in covariates for the 1948 cohort and the first treated cohorts (1949). Panel A shows that the rural share in the 1948 December cohort is slightly larger than the rural share for the January 1949 cohort, and the female share is slightly lower. The 1948 December cohort has a lower employment rate in 1998 than the January 1st, 1949 cohort, although there is no differences in log household income. Nevertheless, these differences between the December and January cohorts are very similar to other cohorts, such as the 1949 and 1950 cohorts (see column 2). Reassuringly, when we estimate the discontinuity in covariates between the treated and untreated cohorts relative to two placebo cohorts (similarly to the estimates in the main text, see equation (13)), we find no significant differences in covariates (see column 3). This highlights that our net of placebo estimates pass the covariate balance tests for all outcomes in panels A and B. This can also be seen for individual-level covariates for both types of regions pooled together, and for low and high-growth regions separately.

Table D.1: Covariate Balance between Treatment and Control

	(1)	(2)	(3)
	Baseline RDD		Net-of-placebo RDD
Cohorts of interest	1948-1949	1949-1950	1948-1949
Panel A: individual-level covariates			
Rural, overall	0.0115***	0.0128***	-0.0013
	(0.0014)	(0.0014)	(0.0020)
Rural, low-growth regions	0.0102***	0.0111***	-0.0009
	(0.0017)	(0.0016)	(0.0023)
Rural, high-growth regions	0.0140***	0.0155***	-0.0015
	(0.0024)	(0.0024)	(0.0034)
Female, overall	-0.0078***	-0.0083***	0.0005
	(0.0012)	(0.0011)	(0.0016)
Female, low-growth regions	-0.0069***	-0.0102***	0.0033
	(0.0015)	(0.0015)	(0.0021)
Female, high-growth regions	-0.0091***	-0.0056*	-0.0035
	(0.0018)	(0.0018)	(0.0026)
Panel B: regional-level covariates			
Log household income in 1998	-0.0076	0.0078	-0.0045
($\times$100)	(0.0319)	(0.0384)	(0.0462)
Employment in 1998	0.0009***	0.0017***	-0.0005*
	(0.0002)	(0.0002)	(0.0003)
N	1,438,545	1,499,039	2,937,584

Notes: This table shows covariate balance between the treated and control cohorts. For each covariate we assess whether there is a discontinuity in its value between birth cohorts. Column (1) shows the estimated discontinuity in the covariate between the 1948 and 1949 cohorts, while column 2 shows it for the 1949 and 1950 cohorts. We estimate the (linear trend) RDD specification described in equation (12) with the reported covariate as an outcome variable. In column (3) we present net-of-placebo estimates where we compare the estimated discontinuity between the 1948 and 1949 cohorts to the estimated discontinuity between the 1949 and 1950 cohorts. There, we estimate the (linear) RDD specification described in equation (13) with the covariate as an outcome variable. In Panel A, the covariate is whether the individual lives in a rural or urban area. In panel B, we use the characteristics of the region where the individual lives. We estimate the regional-level log household income and the regional-level employment rate in 1998, before the policy change using the Household Budget Survey. In each specification, we apply a "donut hole" by dropping individuals born between December 16th and January 5th. Robust standard errors are reported in parentheses. Significance levels are: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Figure D.1: Frequencies by Month of Birth, Pre-reform Population Registry (1998)

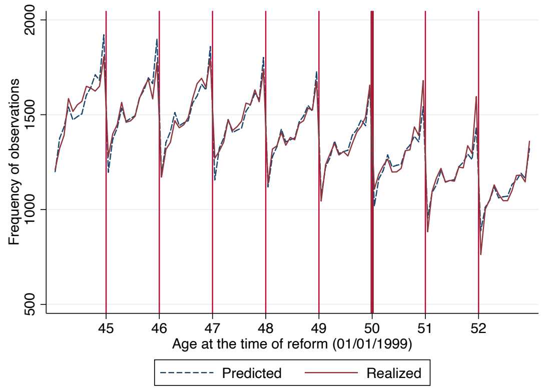

(a) No donut

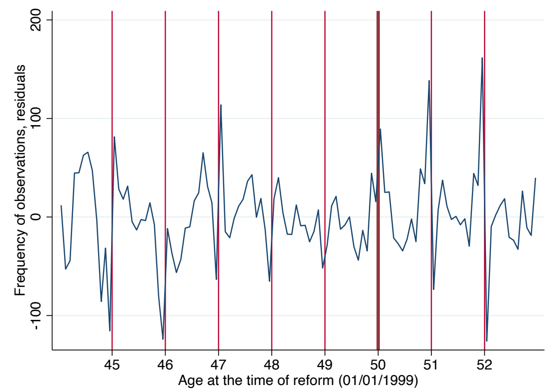

(b) No donut, residuals

(c) Baseline donut

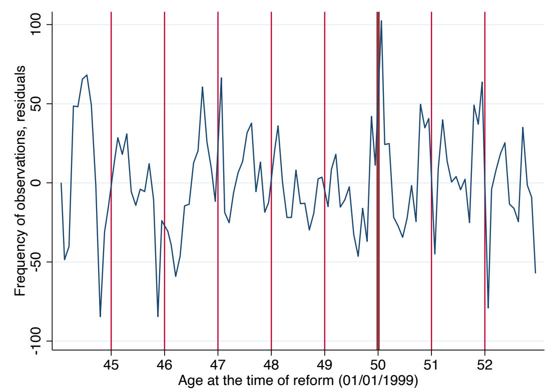

(d) Baseline donut, residuals

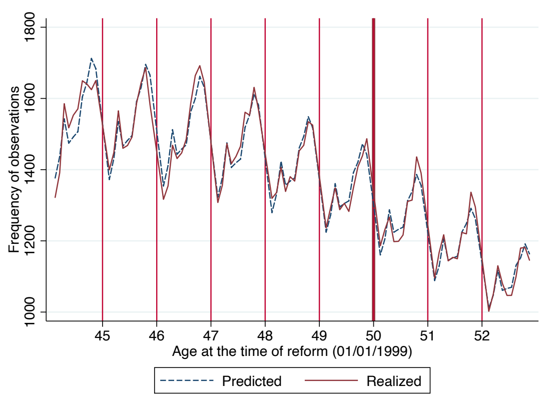

(e) Jan-Dec donut

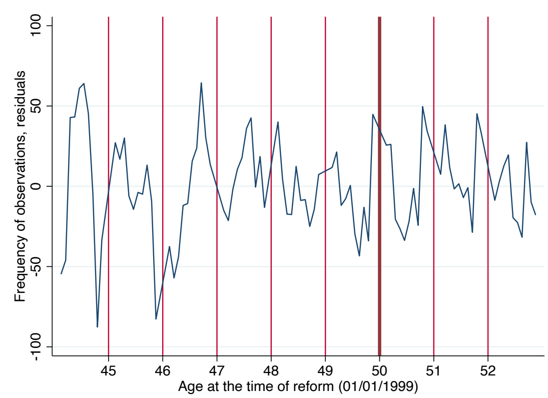

(f) Jan-Dec donut, residuals

Notes: This figure shows the frequencies by month of birth for the cohorts between 1944 and 1953. The frequencies are calculated using the population registry for 1998. We adjust the frequencies with the average number of days in each month. Panels a), c) and e) show the actual number of frequencies (red line) and the predicted one that is based on month of dummies and linear year trend (blue line). Panels (b), (d) and (f) show the difference between the predicted and the actual frequencies – the residuals from the prediction. Panels (a) and (b) show the actual data. Panels (c) and (d) show the frequencies when we drop individuals born between December 16th and January 5th (and adjust with the number of days remained in the month). Panels (e) and (f) show frequencies when we drop everybody born in January or December. Each red vertical line marks the December 31st/January 1st cohort threshold. The thick vertical red line shows the eligibility cut-off of the reform.

Figure D.2: Frequencies by Month of Birth, Post-reform Population Registry (2000-2002)

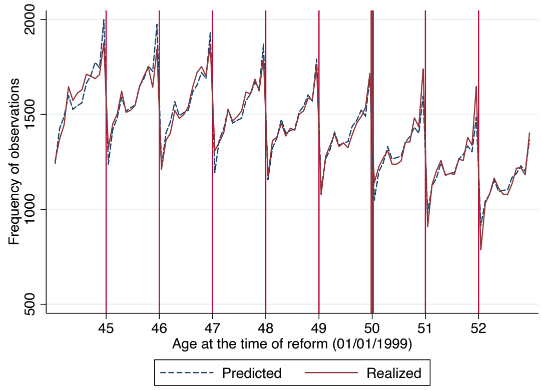

(a) No donut

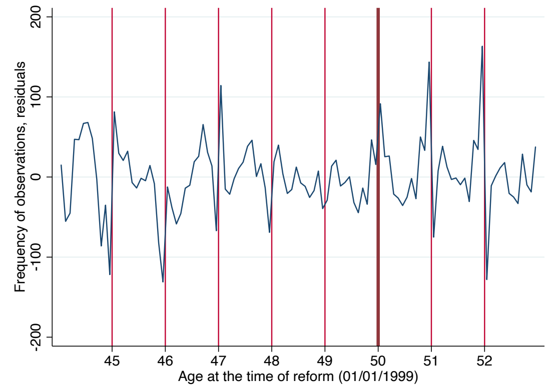

(b) No donut, residuals

(c) Baseline donut

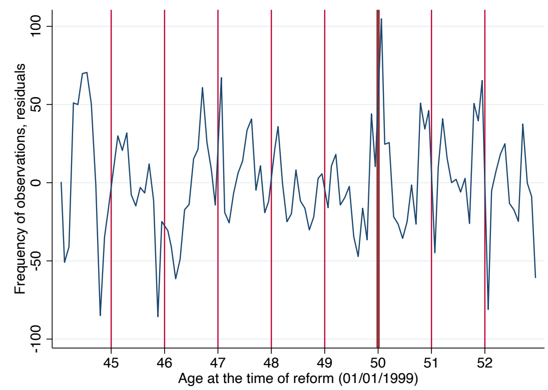

(d) Baseline donut, residuals

(e) Jan-Dec donut

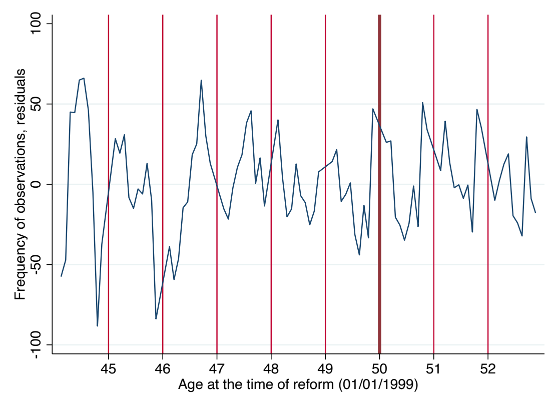

(f) Jan-Dec donut, residuals

Notes: This figure shows the number of observations by month of birth for the cohorts between 1944 and 1953. The frequencies are calculated using the population registry between 2000-2002. We adjust the frequencies with the average number of days in each month. Panels a), c) and e) show the actual number of frequencies (red line) and the predicted one that is based on month of dummies and linear year trend (blue line). Panels (b), (d) and (f) show the difference between the predicted and the actual frequencies – the residuals from the prediction. Panels (a) and (b) show the actual data. Panels (c) and (d) show the frequencies when we drop individuals born between December 16th and January 5th (and adjust with the number of days remaining in the month). Panels (e) and (f) show frequencies when we drop everybody born in January or December. Each red vertical line marks the December 31st/January 1st cohort threshold. The thick vertical red line shows the eligibility cut-off of the reform.

Online Appendix E The Effect of Old Age Unemployment Benefit

E.1 Institutional Details and the Change in Net Return to Work

To study the employment response to contemporaneous work incentives, we study the employment response to a radical change in eligibility for an old age unemployment benefit program (OAUB). This program provided generous benefits to older individuals who were laid off from their jobs. The benefit entitled individuals to 80 percent of a hypothetical pension, with a minimum benefit of 120 percent of the unemployment benefit and a cap at 200 percent of the unemployment benefit. The benefit was stopped if the combined earnings from employment and the allowance were more than 200 percent of the unemployment benefit.[45]

On 30th April 2004 a reform raised the eligibility age from 55 to 60, effective starting 1st August 2004. Individuals could therefore take up the benefit if they reached age 55 by the 1st August 2004, and could demonstrate that they were laid off. This created a cohort-based discontinuity in access to this benefit, with individuals born before 1st August 1949 being potentially eligible for the benefit, and individuals born after not eligible.

The net return to work for individuals who were slightly older than 55 on 1st August 2004 and were thus eligible for the OAUB program can be calculated as:

$$nrw_{it}^{\text{OAUB}} = (1-\tau(\tau^{pi},\tau^{ss}))\cdot w_{it} - u_{it}^{\text{OAUB}} + E_t(PV_{it}^{\text{Employed}_t,NDC} - PV_{it}^{\text{Not employed}_t,NDC}). \tag{E.1}$$

where $(1-\tau(\tau^{pi},\tau^{ss}))\cdot w_{it}$ is the after-tax wage, u_{it}^{OAUB} is the value of old-age unemployment benefit, and $E_t(PV_{it}^{\text{Employed}_t,NDC} - PV_{it}^{\text{Not employed}_t,NDC})$ is the increase in the present-value of pensions as a result of working under the NDC rules. We apply the NDC rules here as the individuals who were slightly older than 55 at 1st August 2004, but younger than 55 at 31st of December 2004, were 49 years old at the time of the 1999 pension reform and hence were ushered into the NDC system.

At the same time the net return to work for individuals who were younger than 55 at 1st of August and are not eligible for the OAUB is as follows :

$$nrw_{it}^{\text{NOAUB}} = (1-\tau(\tau^{pi},\tau^{ss}))\cdot w_{it} - u_{it}^{\text{NOAUB}} + E_t(PV_{it}^{\text{Employed}_t,NDC} - PV_{it}^{\text{Not employed}_t,NDC}). \tag{E.2}$$

where u_{it}^{NOAUB} is the combination of the standard welfare and unemployment programs that every individual in the society has access to. The change in return from work as a result of increasing the eligibility age of the OAUB program is therefore:

$$nrw_{it}^{\text{OAUB}} - nrw_{it}^{\text{NOAUB}} = -(u_{it}^{\text{OAUB}} - u_{it}^{\text{NOAUB}}). \tag{E.3}$$

Following our notation studying the impact of the pension reform (see Section 4), we denote nrw_t^{OAUB},

[45] The Polish name of the program is "zasiłek przedemerytalny" (or "pre-retirement allowance" in English). This name reflects the tight link of these benefits to the DB pension benefit. In that sense the program has some features of early-retirement program. Nevertheless, the program is more similar to a standard unemployment insurance as an individual is required to be laid-off to be eligible for the program.

u_t^{OAUB} and u_t^{NOAUB} as the sample averages of $nrw_i t^{\text{OAUB}}$, u_{it}^{OAUB} and u_{it}^{NOAUB} at age t, respectively. We then define the employment elasticity as follows:

$$\eta^{Contemp} = \frac{(P_t^{\text{OAUB}} - P_t^{\text{NOAUB}})/P_t^{\text{OAUB}}}{-(u_t^{\text{OAUB}} - u_t^{\text{NOAUB}})/nrw_t^{\text{OAUB}}}. \tag{E.4}$$

where the numerator, $(P_t^{\text{OAUB}} - P_t^{\text{NOAUB}})/P_t^{\text{OAUB}}$, shows the employment impact of the OAUB program among those who were eligible. We will describe in the detail how we calculate this next.

E.2 Empirical Strategy and Results

We apply a regression discontinuity design for estimating the effect of the OAUB program. We estimate the following regression equation:

$$P_{it} = \alpha + \beta^{OAUB}\mathbf{1}\{z_i \geq 55\} + f(z_i) + \varepsilon_{it} \tag{E.5}$$

where P_{it} equals to 1 if individual i is employed at time period t and z_i is the age of the individual on 1st August 2004. Those individuals who were younger than 55 years old at the time of the reform, $\mathbf{1}\{z_i < 55\}$, were not eligible for the OAUB program, while those individuals who were older than 55 were eligible for the OAUB if they met certain non-age eligibility requirements described in greater detail below. As such, the regression coefficient β^{OAUB} captures the effect of having access to the OAUB program. We follow Lee and Lemieux (2010) and estimate two separate regressions of $f(z_i)$ on each side of the cutoff point. We use a linear trend in the birth date.

A key assumption in such a regression discontinuity design is that there is no manipulation of the running variable z_i. Such a manipulation would induce bunching in the data. However, there is no such bunching observed in the frequency distribution around the discontinuity (figure available on request).

In Table E.1 we report the estimated β^{OAUB} from the regression equation (E.5). We report estimates using data from all regions (first row), and separately for high (second row) and low earnings growth regions (third row). Column (1) shows estimates with linear trend, while columns (2) - (6) shows the estimates with local-linear estimation by applying various bandwidth choices. The results are very similar across specifications and all specifications. There is a 3.1-4.2 percentage point drop in employment across the specifications and the estimated change in employment in high and low growth regions are similar to each other.

In Column (7) we also implement a net of placebo exercise following a similar strategy as in the main analysis on the pension reform. In particular, we compare the estimated post 2004 employment differences between individuals who were born just before and after 1st August 1949, to the pre 2004 employment differences. The results are virtually the same, which highlights that there is no change in employment preceding the reform. This is shown in Table E.2 where we directly show the employment differences around the discontinuity in the years preceding the reform.

E.3 Calculation of the Employment Elasticity

Our aim is to calculate the employment elasticity as defined in equation (E.4). Our empirical exercise recovers the percentage point change in employment at the policy discontinuity. To recover the percentage change in employment caused by losing access to OAUB program, $(P_t^{\text{OAUB}} - P_t^{\text{NOAUB}})/P_t^{\text{OAUB}}$, we need to take into account that not everyone who was older than age 55 was eligible to the program. Denote $Elig_i = 1$ if the individual satisfies the non-age related program requirements, such as having a long enough working history, and involuntary job loss and $Elig_i = 0$ if the individual does not satisfy the requirements.

The regression in equation (E.5) identifies the following difference in employment:[46]

$$\beta^{OAUB} = Pr(P_{it} = 1|z_i < 55) - Pr(P_{it} = 1|z_i > 55)$$

Whereas the above equation identifies the effect over both eligible and ineligible individuals, our elasticity formula shows the percentage change in employment for the eligible population, formally:

$$\frac{P_t^{\text{OAUB}} - P_t^{\text{NOAUB}}}{P_t^{\text{OAUB}}} = \frac{Pr(P_{it} = 1|Elig_i = 1, z_i < 55) - Pr(P_{it} = 1|Elig = 1, z_i > 55)}{Pr(P_{it} = 1|Elig_i = 1, z_i > 55)}. \tag{E.6}$$

We can apply similar steps as in C.2 to get the percentage change in employment for the eligible population. Non-age related eligibility (from layoff and tenure) is unlikely to vary much with age at the discontinuity so we assume $Pr(Elig_i|z_i < 55) = Pr(Elig_i|z_i > 55) = Pr(Elig_i)$ and so

$$\begin{aligned}
& Pr(P_{it} = 1|z_i < 55) - Pr(P_{it} = 1|z_i > 55) = \\
= \quad & \big(Pr(P_{it} = 1|Elig_i = 1, z_i < 55) - Pr(P_{it} = 1|Elig_i = 1, z_i > 55)\big) \times Pr(Elig_i = 1) - \\
& -\big(Pr(P_{it} = 1|Elig_i = 0, z_i < 55) - Pr(P_{it} = 1|Elig_i = 0, z_i > 55)\big) \times Pr(Elig_i = 0) \\
= \quad & (Pr(P_{it} = 1|Elig_i = 1, z_i < 55) - Pr(P_{it} = 1|Elig_i = 1, z_i > 55) \times Pr(Elig_i = 1)
\end{aligned} \tag{E.7}$$

where in the second equality we have used that there is no employment change around the discontinuity for non-eligible workers $Pr(P_{it} = 1|Elig_i = 0, z_i < 55) - Pr(P_{it}^{\text{NOAUB}} = 1|Elig_i = 0, z_i > 55) = 0$. We also have that

$$Pr(P_{it} = 1|Elig_i = 1, z_i > 55) = \frac{Pr(P_{it} = 1 \cap Elig_i = 1|z_i > 55)}{Pr(Elig_i = 1)} \tag{E.8}$$

and plugging the previous two equations into equation (E.6) implies that

$$\frac{P_t^{\text{OAUB}} - P_t^{\text{NOAUB}}}{P_t^{\text{OAUB}}} = \frac{Pr(P_{it} = 1|z_i < 55) - Pr(P_{it} = 1|z_i > 55)}{Pr(P_{it} = 1 \cap Elig = 1|z_i > 55)} = \frac{\beta^{\text{OAUB}}}{Pr(P_{it} = 1 \cap Elig = 1|z_i > 55)}$$

where the numerator is the percentage point change in employment at the discontinuity while the denominator is the fraction of the population that is working and eligible to the OAUB program. The denominator can be rewritten as

$$Pr(P_{it} = 1 \cap Elig_i = 1|z_i > 55) = Pr(P_{it} = 1|z_i > 55) \times Pr(Elig_i = 1|z_i > 55, P_{it} = 1) \tag{E.9}$$

[46] For simplicity, we abstract away from the issue that in our administrative data we only observe non-agricultural workers. Our empirical exercise recovers only the elasticity for non-agricultural workers, but this is exactly what we need for comparing the employment responses to the contemporaneous change in incentives to the changes coming through the pension reform.

where $Pr(P_{it} = 1|z_i > 55)$ is the employment rate to the right of the discontinuity, and $Pr(Elig_i = 1|z_i > 55, P_{it} = 1)$ is the fraction of individuals who are working and eligible to the pension reform.

We can directly observe $Pr(P_{it} = 1|z_i > 55)$ in our data, which is between 36-39%, depending on the region. Nevertheless, we do not directly observe eligibility in our data, and so we need to infer $Pr(Elig_i = 1|z_i > 55, P_{it} = 1)$ indirectly from the data. The key eligibility criteria were whether the worker has long enough previous employment history and the job separation was involuntary termination. In practice, involuntary job separations are not common in Poland. However, the key is whether someone can "engineer" a severance such that makes the worker eligible for old age unemployment benefit program.[47] This "engineering" is likely to be easier for workers at smaller firms with long tenure. We corroborate this conjecture by exploiting the Labor Force Survey (LFS) data where we observe job tenure and firm size. We estimate the employment response around the reform discontinuity by estimating separately the change in employment using equation (E.5) for four groups of workers: workers with short tenure (at or below 8 years) at large firms (more than 10 workers), workers with long tenure (longer than 8 years) at large firms, workers with short tenure at small firms (10 or less workers), and workers with long tenure at small firms.

We report the key estimates in column (2) in Table 4. In most of the groups there is only a limited change in employment, except for workers in small firms and with tenure of long duration. Since the fraction of individuals employed in small firms with long-tenure is around 41.6% (see column (1) in Table 4), we assume 40% is a lower bound for the fraction of workers who were eligible for the unemployment program. In that case applying our formula leads to an elasticity of 1.02 (or 1.12 in high earnings growth regions and 0.99 in low earnings growth regions).

Nevertheless, it is also possible that some of the workers in the other categories can also "engineer" an involuntary separation. We infer an upper bound of the fraction of eligible population in each subgroup as follows. We assume that all of the individuals in the most responsive category of workers (with long tenure and in small firms) are eligible for OAUB. From column (3) of Table E.3, it can be seen that this category of workers experiences a change in employment of 28.4% as a result of the reform. We then assume that the ratio of the estimated percent change for each of the other groups to the change in the most affected group corresponds to the fraction of workers eligible in each employment group. For instance, since the short-tenure and large-firm group experience a 20% change in employment as a result of the reform, we infer that 20.0/28.4 = 70.4% of individuals in this group are eligible for OAUB. And since workers in the short-tenure and large-firm group constitute 17.4% of the workforce, eligible workers in this group constitute 12.3% of the overall workforce. If we apply this exercise to each group and sum up to the total fraction of workers eligible, we arrive at an overall eligibility rate of 60.4%. This is our preferred eligibility rate, as it is likely that some fraction of all employment categories was eligible for the benefit. The corresponding elasticity for this eligibility rate is 0.68 (or 0.74 in high earnings growth regions and 0.66 in low-earnings growth regions).

[47] Employment laws in Poland make firings potentially costly for the employer, since they require severance payments and if the firing is not justified then they are legally responsible for compensation.

E.4 Exploring employment changes between age 56-59: the interaction between OAUB and the pension reform

As explained in Section E.1, the OAUB benefit could be claimed by those age 55+ who faced an involuntary job separation until August 1st, 2004, when the age threshold increased to age 60. Thus individuals born Aug. 1, 1949 and earlier could potentially receive the OAUB benefit at age 55 but those born afterwards could not. People born around our main pension reform discontinuity (born before or after January 1st, 1949) are unaffected by this policy change as people on both sides of the discontinuity could claim OAUB benefits from age 55 if they were eligible (i.e. they faced an involuntary separation). Taking up OAUB benefits nevertheless has different consequences for future pension benefits for individuals in the NDC and the DB system. Under the DB rules, OAUB benefit periods contributed to the "non-contributory years" component of the DB benefit formula, increasing pension benefits. However, under the NDC rules OAUB benefits did not affect future pension benefits. This interaction between the OAUB benefit and the pension system caused a change in incentives in 2004 (at age 55): since it is more beneficial to receive OAUB benefits for those whose pension benefits were based on the DB rules we expect higher OAUB benefit take-up and a larger reduction in employment at age 55 for them (relative to those whose pension benefits were based on the NDC rules). Nevertheless, the decline in work incentives should be the same in both high- and low-growth regions. Therefore, by comparing the differential employment changes around the discontinuity between high and low growth regions, we can net out the extra incentive changes coming from OAUB.[48]

Employment change between age 56-59. Taking into account these caveats we estimate the employment effects of the pension reform for ages 56-59 by applying a similar empirical strategy as for the estimates for ages 50-54. We apply the same RDD strategy described in Section 5 (see equation 12). However, to make sure that the 2004 changes in the OUAB do not contaminate our estimates, we focus on workers who were born before 1st August 1949, and so are unaffected by this policy. The main estimates for workers aged 56-59 are summarized in Table E.4, which has the structure as our benchmark estimates of Table 1 for workers aged 50-54.

Panel A shows that across various specifications we see negative employment effects of the NDC pension reform for high-wage growth regions and positive effects for low-wage growth regions. This is slightly different than what we saw in Table 1 for ages 50-54, where a larger negative employment effect was observed for high-wage growth regions. This difference is likely to reflect interaction between the pension system and the OAUB: employment in the DB cohorts lowered in general as there claiming OUAB is more advantageous. Nevertheless, once we focus on the differences between high- and low-growth regions, where the differential effect of OUAB should not play a role, we find that a similar pattern emerges as before. Employment for the NDC cohort is lower than for the DB cohort, which reflects the fact that the net return to work was 4.42% lower for the NDC cohort at these ages. We report the implied elasticity in Panel G of Table 3, which is calculated based on Column 4 of Table E.4 (the donut, linear trend, net of placebo specification). The implied employment elasticity is 0.83 (s.e. 0.35), which is slightly larger than our benchmark elasticity (0.51,

[48]The effect of OUAB on employment is the same in high- and low-growth regions. This is demonstrated, for example, in Figure 4, which shows the employment response to the OUAB for those in the NDC system. We also find no differential take-up rate in OUAB when we exploit the eligibility discontinuity at age 55 for those under the DB cohorts (e.g. born in 1948 and 1947).

s.e. 0.21).[49]

Panel B shows the estimates on the intensive margins. Once we focus on the high and low-growth difference we do not see a clear pattern emerging.

Table E.1: Effect of Old Age Unemployment Benefit (OAUB) Reform on Employment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1. All regions	0.0371***	0.0376***	0.0340***	0.0326***	0.0374***	0.0337***	0.0370***
N = 790,783	(0.0022)	(0.0026)	(0.0032)	(0.0033)	(0.0024)	(0.0053)	(0.0060)
2. High-growth	0.0418***	0.0421***	0.0325***	0.0337***	0.0414***	0.0333***	0.0400***
N = 313,470	(0.0036)	(0.0042)	(0.0051)	(0.0053)	(0.0039)	(0.0084)	(0.0096)
3. Low-growth	0.0339***	0.0347***	0.0351***	0.0316***	0.0346***	0.0353***	0.0349***
N = 474,131	(0.0029)	(0.0033)	(0.0060)	(0.0043)	(0.0031)	(0.0064)	(0.0041)
Sample	full	full	full	full	full	full	full
$f(z_i)$	Linear trend	Local-Linear	Local-Linear	Local-Linear	Local-Linear	Local-Linear	Linear trend
Bandwidth	–	150	50	100	200	Calonico et al.	–
net of placebo	no	no	no	no	no	no	yes

Notes: This table shows the estimated change in employment at the old age unemployment benefit program discontinuity. Each cell in the table shows the β^{OAUB} coefficients of the RDD specification shown in equation (E.5). The rows show the estimated employment change for different regions. The first row shows it for all regions, while the second and third rows show the estimated effect in high and low-growth regions, respectively. High-growth regions are regions with above median earnings growth rate between 2000 and 2013, while low-growth regions have below median growth. In all columns we use the full dataset (no donut hole applied). In columns (1), (7) we estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. Columns (2)-(6) apply a kernel-weighted local linear regressions with various bandwidth choices. Column (7) estimates the change in employment at the reform discontinuity relative to the change at the placebo discontinuity by applying an analogous empirical strategy to that described in equation (13). The placebo discontinuity is estimated for the 1949 cohort in 2001-2002, before the OAUB reform was announced. We report standard errors in parentheses. For the local-linear regression we calculate robust standard errors following Calonico et al. (2014). Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table E.2: Placebo Estimates on Employment Effect of the Old Age Unemployment Benefit (OAUB) Program

Wage growth region	Placebo (2001-2003)
All regions	0.0001
N = 833,934	(0.0023)
High-growth	0.0018
N = 333,173	(0.0036)
Low-growth	0.0010
N = 499,872	(0.0029)
Sample	full
$f(z_i)$	Linear trend

Notes: This table shows the estimated change in employment at the placebo discontinuity. We use the 1949 cohort in 2001-2003, before the OAUB reform was announced. We estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the placebo cutoff. We report robust standard errors in parentheses. Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

[49]Note that when we estimate the employment effect for the 56-59 age group, we only use individuals who were born before August 1st, 1949. If we apply the same restriction for the 50-54 age group, we get an employment elasticity of 0.46 (s.e. 0.25), which is very close to our benchmark estimate of 0.51 (s.e. 0.21) reported in Table 2.

Table E.3: Effects of Old Age Unemployment Benefit (OAUB) by Employment Type

	(1) Share of Workforce (%)	(2) RDD estimate	(3) Percent change (%)
Short tenure and large firm	17.42	0.0209	20.00
		(0.0154)	
Long tenure and large firm	32.58	0.0053	2.71
		(0.0238)	
Short tenure and small firm	8.39	0.0059	11.72
		(0.0125)	
Long tenure and small firm	41.61	0.0709***	28.40
		(0.0261)	
Number of observations		2,430	2,430

Notes: Column (1) shows the fraction of workforce in four employment categories: workers with short tenure (at or below 8 years) at large firms (more than 10 workers), workers with long tenure (longer than 8 years) at large firms, workers with short tenure at small firms (10 or less workers), and workers with long tenure at small firms. Column (2) reports the β^{OAUB} coefficients of the RDD specification shown in equation (E.5) estimating separately the employment discontinuity for each of the four categories. Each row corresponds to a separate regression with the outcome being employed in the particular employment category. For instance, in the 1st row, the outcome variable is equal to 1 if an individual is employed at a large firm with short tenure, and 0 otherwise. We use the Labor Force Survey for this analysis and the years for which we have month-of-birth information in the data, namely 2004 (post-1st August) and 2005. We estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. Column (3) reports the corresponding percent change in employment in each employment category (dividing column 2 by column 1 multiplied by 60%, the share assumed to be eligible for OAUB in our baseline). Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table E.4: The Effect of the Pension Reform on Employment And Wages at Ages 56-59

	(1)	(2)	(3)	(4)
Panel A: Change in employment probability				
High-growth	-0.0103***	-0.0039	-0.0057	-0.0037
N = 497,474	(0.0027)	(0.0030)	(0.0080)	0.0050)
Low-growth	0.0043	0.0095***	0.0110**	0.0111***
N = 745,349	(0.0022)	(0.0024)	(0.0065)	0.0040)
Difference (High-Low)	-0.0146***	-0.0134***	-0.0167***	-0.0148***
	(0.0035)	(0.0039)	(0.0103)	(0.0064)
Panel B: Change in log wage				
High-growth	-0.0401	-0.0241	0.0272	-0.0194
N = 174,001	(0.0146)	(0.0161)	(0.0390)	(0.0229)
Low-growth	-0.0049	-0.0098	-0.0442	-0.0046
N = 248,968	(0.0120)	(0.0132)	(0.0323)	(0.0190)
Difference (High-Low)	-0.0352	-0.0143	0.0714	-0.0149
	(0.0190)	(0.0208)	(0.0506)	(0.0297)
Sample	Full	Donut	Donut	Donut
$f(z_i)$	linear trend	linear trend	local linear	linear trend
net-of-placebo	no	no	no	yes

Notes: This table shows the estimated change in employment (panel A) and log wage (measured as earned income of workers) for those in work (panel B) at the reform discontinuity, for years 2005-2007 (ages 56-59). Each cell in the table shows the β coefficients of the RDD specification shown in equation (12) (Columns (1)-(3)) or in equation (13) (Column (4)). The rows show the estimated employment and wage change for different regions. The first and second rows show the estimated effect in high and low-growth regions, respectively. High-growth regions are regions with above median earnings growth rate between 2000 and 2013, while low-growth regions have below median growth. The third row shows the difference between the high and low-growth regions. In Column (1) we use the full dataset. In Columns (2)-(4) we apply the donut hole RDD specification where we exclude those born between December 16th and January 5th. In Columns (1), (2) and (4) we estimate a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. Column (3) estimates a kernel-weighted local linear regression, where we set the bandwidth at 150 days. Column (4) estimates the change in employment at the reform discontinuity relative to the change at the placebo discontinuity as in equation (13). The placebo discontinuity is estimated between the 1949 and 1950 cohorts, both of which switched to the NDC system. We report robust standard errors in parentheses. For the local-linear regression we calculate robust standard errors following Calonico et al. (2014). Significance levels are: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Online Appendix F Appendix to Lifecycle Model

F.1 Model Pension Systems

Our approach to modeling pension accrual closely follows the approach in the simulations in Section 4. To be consistent with our treatment of the net return to work there, we measure pension accrual as a wage increment, using an approach similar to that of French and Jones (2011). Our approach here differs in one important regard from the approach in the simulations. In our simulations we can allow pension wealth to depend on the entire history of wage and unemployment shocks, whereas in the dynamic model this would be computationally infeasible, and so we must approximate pension accrual given the model's state variables (in particular, age, wage and region).

We consider three scenarios when comparing DB and NDC systems. In the first, we consider DB and NDC systems which would apply to cohorts of individuals who work their entire careers under one of the two systems. We refer to these as 'steady-state' DB and NDC pension systems. In the second we compare a steady state DB system to a composite system where everyone begins in the DB system, but some people transition into the NDC system at age 50, in a way similar to the impacted cohorts in reality. Finally, we consider 'steady state' DB and NDC systems where we construct the two to be revenue-equivalent, allowing us to isolate the effect of work incentives. We now discuss each of these in turn.

F.1.1 Steady state pension systems

DB Pension Our approximated DB system has three components: (1) a universal payment to all individuals in retirement; (2) an increment to contemporaneous wages received in periods in which an individual works; and (3) an increment to contemporaneous income received whether or not an individual works. Taking each in turn, the DB system comprises:

1. A universal payment to all those over the pension age, α_0^{DB}, which represents the universal component of the DB pension (the first term in the last bracket in equation (4)).
2. An increment to contemporaneous earnings in period t received if the agent chooses to work, which is proportional to earnings by a factor ($\alpha_{1,t}^{DB}$) that varies by age and region (we suppress region subscripts here). This is estimated using our simulations in three steps:
   i) We first calculate the change in the present value of DB pensions as a result of choosing to work in a particular year for each simulated agent (in exactly the same manner as we do in Section 4).
   ii) We then calculate earnings in work at time t for each simulated agent.
   iii) Finally, we calculate the ratio of average accrual (from step (i) to average earnings (from step (ii)) at each age t, to yield our measure of the increment to DB wealth from work. We smooth this with a polynomial regression to obtain $\alpha_{1,t}^{DB}$.
3. Finally, under the DB system, unemployed individuals accumulated DB benefits by accumulating "contributory years" while on unemployment benefits – a period of 6 months in our simulations. To account for this, under the DB system, we assume individuals receive a contribution which is

proportional (by a factor $\alpha^{DB}_{2,t}$) to the earnings that they would have received if they had chosen to work.

We calculate this component as the change in the DB pension when being out of work and *not* receiving unemployment benefits, compared with being out of work and receiving unemployment benefits. This is estimated using our simulations as follows:

i) We calculate what the increment to DB wealth would be if the agent worked under the (modeled) scenario where there is 6 months of unemployment benefits received when not working. This is the average accrual factor $\alpha^{DB}_{1,t}$, the calculation of which is described above.

ii) We calculate what the increment to DB wealth would be if the agent worked under the (counterfactual) scenario where there would be no unemployment benefit receipt for each individual in our simulations. We average this and divide by average potential earnings at each age t. Potential earnings are the wage – i.e. earnings for workers, and the earnings that those who do not work would receive if they worked.

iii) We take the difference of the objects calculated in (ii) and (i). We take this to be the increment to DB wealth which results in not working but receiving six months of unemployment insurance. We smooth this with a polynomial regression and label it as $\alpha^{DB}_{2,t}$.

The first of these three components enters the budget sets directly when an individual is retired, as shown in equation (F.14) in the next subsection. The second and third of these components enter through income during working life, which can be characterized as follows:

$$\begin{aligned} y^{DB}_{it} &= (1-\tau(\tau^{pi},\tau^{ss})+\alpha^{DB}_{1,t}+\alpha^{DB}_{2,t})w_{it} \quad \text{if} \quad P_{it}=1 \\ y^{DB}_{it} &= u+\alpha^{DB}_{2,t}w_{it} \quad \text{if} \quad P_{it}=0 \end{aligned} \tag{F.10}$$

where $\tau(.)$ gives tax due, u is an out-of-work welfare payment, and $\alpha^{DB}_{1,t}$ and $\alpha^{DB}_{2,t}$ are pension wealth accrual factors which individuals receive, respectively, if they work and whether they work or not. w_{it} are potential earnings. The calculation of $\alpha^{DB}_{1,t}$ and $\alpha^{DB}_{2,t}$ accounts for the taxation of pensions.

NDC Pension We calculate NDC wealth accrual in a similar fashion. There is no universal component to the pension, so the analogy to the DB universal component is $\alpha^{NDC}_0=0$. $\alpha^{NDC}_{1,t}$ is a factor which, when multiplied by earnings, gives an increment to wages if the agent chooses to work. $\alpha^{NDC}_{2,t}$ is a factor which, when multiplied by the wage, gives the increment to income whether or not the agent chooses to work.

The form of income under the NDC system has the same form as that in equation F.10 but with $\alpha^{NDC}_{1,t}$ and $\alpha^{NDC}_{2,t}$ in the place of $\alpha^{DB}_{1,t}$ and $\alpha^{DB}_{2,t}$.

Summary The two steady-state pension systems can therefore be characterized by the following vectors of parameters:

$$\left(\alpha^k_0, \{\alpha^k_{1,t}\}^T_{t=1}, \{\alpha^k_{2,t}\}^T_{t=1}\right), \tag{F.11}$$

where $k\in\{DB,NDC\}$ indexes pension systems. Figure F.1 illustrates the accrual factors $\{\alpha^k_{1,t}\}^T_{t=1}, \{\alpha^k_{2,t}\}^T_{t=1}$ for each of the two pension systems. Note that the return to work arising from the DB system $\{\alpha^{DB}_{1,t}\}^T_{t=1}$ (in the left of the left panel of the figure) is higher at all ages for those in high-growth regions than for those in low-growth regions. Later in the lifecycle, this predominantly occurs due to more frequent (and larger) recalculations of AIME for those in high-growth regions. Early in the lifecycle, this predominantly occurs due to the fact that the ultimate pension will be calculated based on an AIME that is large relative to contemporaneous earnings (a fact which is true for both types of individuals, but which is quantitatively more important for those in high-growth regions).

Figure F.1: Pension Accrual Parameters

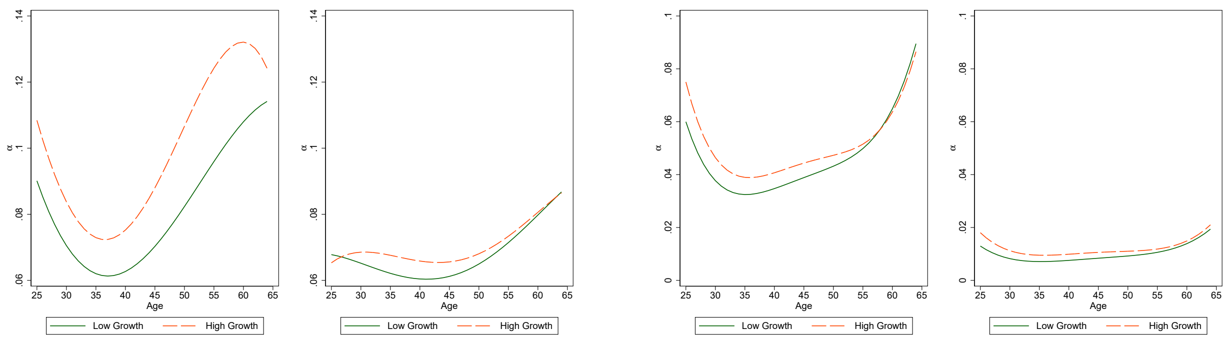

(a) $\{\alpha_{1,t}^{DB}\}_{t=1}^{T}$ and $\{\alpha_{1,t}^{NDC}\}_{t=1}^{T}$ (b) $\{\alpha_{2,t}^{DB}\}_{t=1}^{T}$ and $\{\alpha_{2,t}^{NDC}\}_{t=1}^{T}$

Notes: These graphs show the pension accrual parameters. The construction of these is described in Section F.1.1. The accrual parameters are smoothed by regressing on a quartic in age and taking the predicted values.

F.1.2 Pension system for transition cohort

The systems described above characterize the pensions relevant when considering the behaviours of cohorts who lived their entire working lives under the DB and NDC regimes respectively. Our approach to estimating the causal impact of the reform at the cohort discontinuity in Section 6, however, involves comparing the behavior of a cohort who worked their entire career under the DB regime with a transition cohort who worked until age 50 under the DB regime before being switched to the NDC regime. We need to model the behavior of the latter cohort to obtain model predictions of the change in labor supply, which are used in estimation. To do this we model the pension system for that cohort as:

$$\left(\alpha_0^{trans}, \{\alpha_{1,t}^{trans}\}_{t=1}^{T}, \{\alpha_{2,t}^{trans}\}_{t=1}^{T}\right) \tag{F.12}$$

$$
\begin{aligned}
\alpha_0^{trans} &= \varphi\, \alpha_0^{DB} && && \text{(F.13)} \\
\\
\alpha_{1,t}^{trans} &= \alpha_{1,t}^{DB} && t < 50 \\
\alpha_{2,t}^{trans} &= \alpha_{2,t}^{DB} && t < 50 \\
\\
\alpha_{1,t}^{trans} &= \alpha_{1,t}^{NDC} && t \geq 50 \\
\alpha_{2,t}^{trans} &= \alpha_{2,t}^{NDC} && t \geq 50
\end{aligned}
$$

where the transition cohort, unlike those who spend their whole careers under the NDC regime, *do* receive a universal payment ($\alpha_0^{trans} = \varphi\, \alpha_0^{DB}$), which is unrelated to their earnings history and is lower, by a factor φ, than that received by those who stayed in the DB system. This comes through starting capital, the formula for which is outlined in equation (8). Using that formula, we approximate $\varphi = 0.8995$.[50]

In summary, the parameters of the pension accrual increments that are proportional to earnings ($\alpha_{1,t}^{trans}, \alpha_{2,t}^{trans}$) are those of the DB scheme for up to age 50 (before the reform) and those of the NDC scheme from the age of 50 onwards.

F.1.3 Pension systems used in counterfactual exercise

Once the model is estimated, we use it to evaluate the likely effects on lifecycle labor supply of the changes in the net return to work implied by the reform. For this analysis (the results of which are illustrated in Figure 5), in order to isolate the changes in the net return to work from the changes induced by the fact that the payments became less generous overall, we construct two systems – one with NDC work incentives and one with DB work incentives – which are adjusted to yield the same government balance. We hold the two components which *do not* depend on work choices (α_0^k and $\{\alpha_{2,t}^k\}_{t=1}^T$) constant across systems, set at their levels in the DB system ($k = DB$). Pension accrual when working in the modelled NDC system is that of the new system ($\{\alpha_{1,t}^{NDC}\}_{t=1}^T$). Pension accrual when working in the modelled DB system is that prevailing in the old DB system scaled by a factor ϕ ($\{\phi\alpha_{1,t}^{DB}\}_{t=1}^T$). The value of ϕ is chosen such that the present discounted value of the government balance (the sum of taxes and social security contributions, less welfare and pension payments, adjusted for timing of payment using the interest rate) is the same in both cases.

Formally, the two pension systems we compare are a system with accrual in work following the NDC work incentives:

$$
\left(\alpha_0^{DB}, \{\alpha_{\mathbf{1,t}}^{\mathbf{NDC}}\}_{t=1}^T, \{\alpha_{2,t}^{DB}\}_{t=1}^T\right),
$$

and a modified DB system where the component that the agents get only if they work is scaled by a factor ϕ,

$$
\left(\alpha_0^{DB}, \{\phi\alpha_{\mathbf{1,t}}^{\mathbf{DB}}\}_{t=1}^T, \{\alpha_{2,t}^{DB}\}_{t=1}^T\right).
$$

We highlight in boldface the component that differs across these two systems. In our counterfactual analysis, ϕ is found to be 0.84 to obtain revenue neutrality.

[50] That is, the transition cohort has a universal component of the pension worth approximately 89.95% of that of the cohort who remained in the DB pension.

F.2 Recursive Specification of the model

Equation F.14 gives the recursive specification of the agent's problem. Our vector of state variables is $\mathbf{X}_{it} = \{\text{region}_i, \nu_i, t, a_{it}, \text{offer}_{it}, w_{it}\}$. Two of these variables represent permanent heterogeneity, namely, region of residence $\in$ {low growth, high growth}, and agent's consumption weight (ν_i – their 'type'). The four which vary across the lifecycle are age (t), assets (a_{it}), the presence (or otherwise) of an employment offer (offer_{it}) and wages (w_{it}).

$$V_t(\mathbf{X}_{it}) = \max_{\{c_{it}, P_{it}\}} U(c_{it}, l_{it}; \nu_i) + \beta\Big(s_{t+1}\mathbb{E}_t V_{it+1}(\mathbf{X}_{it+1})\Big) \tag{F.14}$$

$$s.t. \begin{cases} a_{it+1} = (a_{it} + y_{it}^k - c_{it})(1+r) & \text{if } t < R \\ a_{it+1} = (a_{it} + \alpha_0^k - c_{it})(1+r) & \text{if } t \geq R \\ l_{it} = 1 - hP_{it} \\ a_{it+1} \geq 0 \end{cases}$$

where c and P are the model's choice variables: consumption and participation, $U(.)$ is the utility function, defined over consumption and leisure, given in equation (14), β is the discount factor, s_{t+1} is survival probability to $t+1$, conditional on having survived to t, R is the age at which agents must stop work (65), a is the level of assets held, h is the share of the leisure endowment given up when an agent chooses to work, y^k is income, inclusive of pension wealth accrual under the prevailing pension system k (as defined in equation F.10 for the $k = DB$ case), α_0^k is the universal component of the prevailing pension system (k) and the expectation operator integrates over future period wage employment offer uncertainty.

F.3 Solution of the Model

The model outlined has no analytical solution and must be solved numerically. We do this using standard methods, which we summarize very briefly here. See the Appendix to Crawford and O'Dea (2020) for a more detailed description of the solution to a similar model.

We first select a discrete subset of the state space ($\mathbf{X}_t$). Of our six state variables, three are naturally discrete (region of residence, age, and employment offer). The other three are continuous variables (assets, wages and consumption weight) that need to be discretized. For assets, we form a grid of 50 points, spaced equally in log terms (so that more points are concentrated at the bottom of the grid where curvature will be greatest), except at the very top where we have 5 equally-spaced gridpoints. For wages, which are distributed log-normally, we create a grid by dividing the distribution in each period into 20 equi-probable regions, with the grid being formed of the expected value in each of those regions. For preference types, we divide the Normal distribution into 10 equi-probable regions, with the grid being formed of the expected value in each of those regions. These three discretized sets of our continuous state variables, along with our naturally discrete state variables, form our discretized state space.

Solving the model involves solving the objective function at each point in the discretized state space starting in the final period of life T. The solution in this period yields the value function at each point in the discretized state space: $V_T(\mathbf{X}_T)$. Using this, the optimization problem for V_{T-1} and the problems for all earlier periods can be carried out iteratively. Interpolation is used to approximate the value function outside of the discrete subset of the state points we evaluate, and integration over wage shocks is carried out using the procedure in Tauchen (1986).

Solution of the model yields decision-rules for each of consumption and labor supply. These give optimal behavior at each point in the state space. Using these decision rules, random draws for earnings which follow the model's earnings process, and an initializing of the asset distribution (all households start with no assets), we can simulate behavior and obtain a simulated data set. We simulate 8,000 individuals, and use the simulated data for estimation (see Section F.4.2 below). A simulated data set for a version solved under counterfactual policies can be used to assess what behavior would be under those policies (as in done in Figure 5).

F.4 Parameterization and Estimation

F.4.1 Parameterization

In our first step we estimate or calibrate model parameters which can be identified outside the model or calibrated to values commonly used in the literature. This includes the interest rate and survival probabilities. We use the same survival probabilities and the same interest rate used previously in the paper. We also set two preference parameters to values commonly used in the literature. One is the coefficient of relative risk aversion on utility, which is set at 4, a typical value for use in a non-separable utility function (e.g., Conesa et al. (2009)), and the other is the discount rate, which we set equal to the interest rate ($\beta = \frac{1}{1+r}$). We normalize the leisure endowment to one and assume that when agents work they forego a share of their leisure. We set this share, at age 50, to be 30% of their endowment (a standard value in the literature), and we estimate how the share changes with age (as discussed below).

The parameters of the deterministic and AR(1) components of the earnings process are those in the simulation (see equation (10)).[51] We initialize the earnings variance by assuming that shocks accumulate between ages 21 and 25, at which point the labor supply decision is modeled. The Markov process governing unemployment risk is also that in the simulation, discussed in Section 4 and with parameter values given in Table C.2.

F.4.2 Estimation

We estimate the parameters ($\chi = (\mu_\nu, \sigma_\nu^2, \zeta)$) to minimize the following criterion function using Indirect Inference:

$$\hat{\chi} = \arg\min_{\chi} \; (\mathbf{m}(\chi) - \hat{\mathbf{m}})' \mathbf{W} (\mathbf{m}(\chi) - \hat{\mathbf{m}}), \tag{F.15}$$

where $\hat{\mathbf{m}}$ is the vector of moments we match, $\mathbf{m}(\chi)$ is a vector of the values of those moments implied by the model at parameter vector χ, and $\mathbf{W}$ is a weighting function.

Moments We target seven moments: five labor supply moments (proportion of men in work at ages 30, 40, 50, 60, and 64), and the employment response to the reform in each region. We estimate these labor supply moments using the Labor Force Survey. Our aim is to obtain a lifecycle profile of labor supply

[51] The model does not include the MA(1) component process, which would involve the inclusion of another state variable. However, note from Panel C in Table 3, that our simulations are robust to changes in the earnings process and so this difference between the model and simulation is very unlikely to be consequential.

moments for a single cohort. Our data covers the period 1998 to 2010; thus we do not have a full lifecycle for any cohort. However, we are able to track our cohort of interest (born in the 1940s) from age 49 to 64 inclusive. Furthermore, for the earlier part of the lifecycle (up through age 50) we use synthetic cohort methods and use data on those born in the 1950s and 1960s. To estimate a lifecycle profile, we estimate the following regression:

$$l_{it} = \alpha + \beta^{age}\mathbf{age_i} + \beta^{cohort}\mathbf{cohort_i} + u_{it} \tag{F.16}$$

where $\mathbf{age_i}$ is a vector of age dummies and $\mathbf{cohort_i}$ is a vector of cohort dummies (one for each birthyear). This allows us to identify average labor supply for ages 29 to to 64. Our omitted cohort category is those born in 1950. The moments we target are coefficients on ages 30, 40, 50, 60, and 64. We use a selection of age moments rather than the full lifecycle, as, with only 2 treatment effect estimates to target we do not want the number of moments we use in estimation of the lifecycle to have overwhelming weight relative to the number of treatment effect parameters. We do, however, show the fit of labor supply across the whole lifecycle (in Figure F.2), with the targeted moments additionally shown in Table 5.

Weighting matrix For the weighting matrix in the Method of Simulated Moment criterion function, we use the diagonal of the variance covariance matrix of the moment conditions, with one adaptation. As there are 5 labor supply moments, and two moments conditions which target the reform treatment effects, we increase the weight on the latter by 2.5—the latter moments are no less important economically, and we do not wish the *number* of labor supply moments (necessary to target the lifecycle) to be playing a role in down-weighting the extent to which the reduced form estimates affect the estimation.[52] Standard errors are calculated using the formula:

$$(\hat{\chi} - \chi_0) \rightsquigarrow N(0, \mathbf{V}), \qquad \mathbf{V} = (1+\tau)(\mathbf{D'WD})^{-1}(\mathbf{D'W\Omega WD})(\mathbf{D'WD})^{-1}$$

where $\mathbf{D} = \frac{\partial \mathbf{m}(\chi)}{\partial \chi'}\Big|_{\chi=\chi_0}$ is the gradient matrix of the population moment vector, $\mathbf{\Omega}$ is the variance covariance matrix of the targeted moments, and τ is the number of observations per simulation. In practice $\mathbf{W}$, $\mathbf{D}$, and $\mathbf{\Omega}$ are replaced by estimated values: see French (2005) for details.

Model fit is given in Table 5. We show the lifecycle profile of labor supply in Figure F.2, which matches the data well, though the model understates slightly the fall in labor supply through the 50s.

Parameter estimates are also given in Table 5 (panel B). The estimate of mean consumption weight (at age 50) of approximately 0.51 sits at the mid-point of a range of papers that use utility functions similar to ours applied to data from other countries,[53] and the standard deviation in the consumption weight of 0.077 indicates a modest degree of heterogeneity in the population – implying that 95% of the population has a consumption weight of between 0.36 and 0.66.

[52] 2.5 is chosen as there are 2.5 times the labor supply moments as treatment effect moments. In practice this makes very little difference to our estimates.

[53] Conesa et al. (2009) estimate a value 0.377, Nishiyama and Smetters (2007) have between 0.45 and 0.50, O'Dea (2019) estimates a range from 0.42 to 0.52, French and Jones (2011) estimate heterogeneous groups with an average in the population of 0.62.

Figure F.2: Labor Supply fit Over the lifecycle

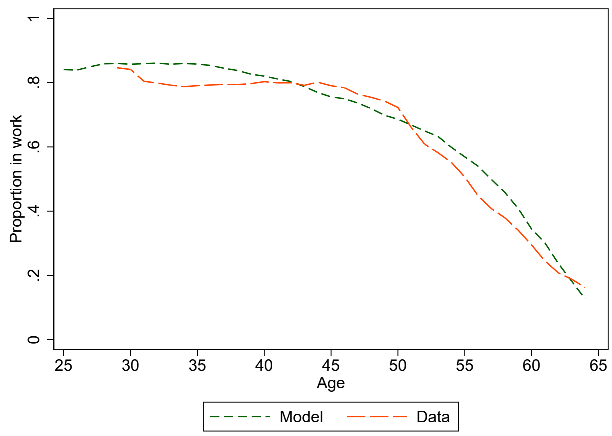

F.5 Counterfactual

The comparison of labor supply across the lifecycle summarized in Figure 5 compares a system with NDC work incentives with a system with DB work incentives designed to deliver an equivalent government balance to that prevailing under the NDC system. This is described in Section F.1.3.

The results in Figure 5 are smoothed using a fourth-order polynomial in age. To illustrate that there is very little difference between these smoothed profiles and the underlying unsmoothed profiles, Figure F.3(a) shows both together. The modest differences between the two figures are caused by approximation error. Even with fine grids and a large number of simulations, the object that we are analyzing – the difference between two labor supply averages – will be subject to modest approximation error.

F.6 Elasticities

Figure F.3(b) gives our employment elasticities across the lifecycle. At any given age, these are calculated by perturbing wages by 20% at (only) that age, calculating the percentage change in the proportion in work and dividing it by the size of the perturbation.[54] We show two versions of the elasticity, which differ in the denominator. One divides the percentage change in labor supply by 20% (the perturbation in gross wages). We label this as "Ours (Gross)". The second has the same numerator: the percentage change in labor supply. The denominator though is not 20% (the increase in gross wages), but the change in the net return to work induced by that change. We calculate this as the average change in the net return to work when wages change, divided by average net return to work without the perturbation. This more closely matches the concept of elasticity calculated in our analysis of the reform. We label this as "Ours (NRW)".

The rising labor supply elasticity with age is an implication of many structural models that contain an extensive margin choice (e.g. French and Jones (2011), Keane and Wasi (2016), Jones and Li (2023),

[54]In practice, we do this for both negative and positive perturbations. We increase and we decrease wages by 20% in turn, and average the absolute percentage point change in employment. We then divide this average by baseline employment to obtain the numerator for the elasticity calculation.

Figure F.3: Model Outputs

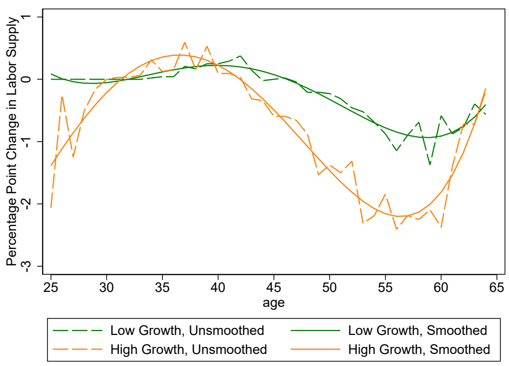

(a) Effect of Switching to an NDC on Labor Supply Over the Lifecycle

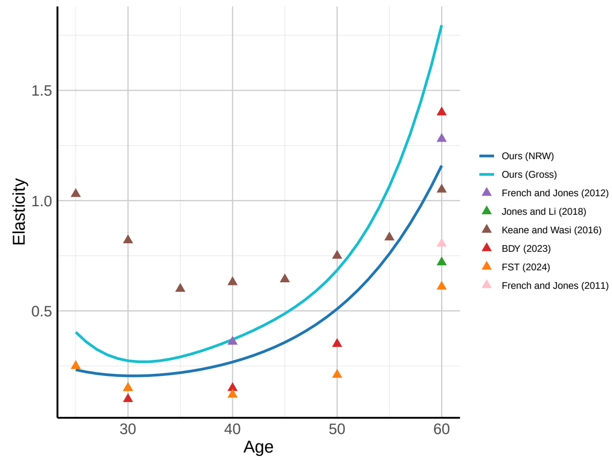

(b) Elasticities across the lifecyle

Notes: Panel (a) shows percentage point change in labor supply at each age between the DB and NDC pension schemes. Profiles are smoothed using a polynomial of order 4. To isolate the reform effect on changing the net return to work from the effect operating through a reduction in the overall generosity of the pension system, we scale down the accruals in the DB system (proportionally) to ensure that the two pension systems are revenue- equivalent. Panel (b) shows elasticities which are calculated by perturbing wages by 20% at each age, calculating the percentage change in the proportion in work and dividing it by labor supply at the unperturbed wage profile. We smooth the elasticities shown in this figure by regressing on a quartic in age. There are two versions of our elasticities. One divides the percentage change in labor supply by 20% (the perturbation in gross wages) - labeled as"Ours (Gross)". The second is the percentage change in labor supply when wages are changed by 20%, divided by the change in net return to work arising from such a change - labeled as"Ours (NRW)".

Borella et al. (2023), and Fan et al. (2024)). We show in Appendix Figure F.3 Panel (b) a comparison of our elasticities and those from the literature for ages at which there is an overlap.