Labor Supply and the Pension Contribution-Benefit Link*

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

June 26, 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 pension 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, Emma Scandolo, 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. In contrast, 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 substantially impacted 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 accrual. A key reason for this difference is that, in the DB system, pension benefits are calculated using an individual's "best" years of their 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 their life cycle earnings profile. For those with steeper earnings profiles, DB pension benefits 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 substantially reduces pension benefits, providing strong labor supply incentives at that age. Conversely, for individuals with flatter wage profiles, incentives under DB rules differ less across the lifecycle. We exploit differences in DB and NDC pension incentives to identify the effect of changing the link between contributions and future benefits on labor supply. In particular, we 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 reduced pension wealth. 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 differed substantially 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 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 a funded DC scheme, the rate of return reflects changes in economic prospects and growth.

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 $\frac{2.29}{4.46} = 0.51$ (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 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

[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.

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 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

[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 Klejdysz and Zawisza (2026). 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.

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 smaller 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 pension wealth (Ye, 2022) and the pension eligibility age (Jean-Olivier et al., 2010; Carta and De Philippis, 2019) affects labor supply prior to the eligibility age. 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 pensions 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 $k = \{DB, NDC\}$:

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

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 pension accrual, i.e. 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 the difference between $PV_{it}^{\text{Employed}_t,k}$ (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). In the next section, we discuss the two pension schemes in detail and illustrate how the reform changed pension accrual and thus 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:

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

where $P_t^{NDC} - P_t^{DB}$ represents the change in the employment rate and Δnrw_t represents the change in the net return to work at age t from switching from the DB to the NDC scheme.

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 the 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]

Importantly for our empirical strategy, the date-of-birth discontinuity is sharp only for

[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.

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 τ^{pi} is the personal income tax rate, $\Big(\bar{y}_{65}(1-\tau^{ss})\Big)$ 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 benefits, 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 $E[T|t = 65]$ is remaining life expectancy at age 65 and A^{NDC}_{i65} is the value of accumulated capital in the notional account at 65. This capital accrues according to the formula:

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

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^{NDC}_{it} 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]

Contribution rates. The total social security contribution rate to the pension system

[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:

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

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 $A^{NDC}_{i50} = b^{start}_{i50} \times E[T|t = 62]$, where $E[T|t = 62]$ is remaining life-span at 62.

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, 1948 lost the option of early retirement, even if they were unfit for work. However, they could still claim

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

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) (Polish Central Statistical Office, 2017) and the Polish Labor Force Survey (LFS) (Polish Central Statistical Office , 2012). 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 1948 and 1949 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

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.

already in that age range. These workers are unaffected by the reform and so our main empirical results are relevant for workers not employed in these special occupations.

- 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.
- 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:

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

$$= \left(\frac{1}{1+r}\right)^{65-t} \sum_{s=65}^{T_{death}} S_{s|t} \left(\frac{1+r^{index}}{1+r}\right)^{s-65} (b_{i65}^{\text{Employed}_t,k} - b_{i65}^{\text{Not employed}_t,k}),$$

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 Klejdysz and Zawisza (2026) 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**

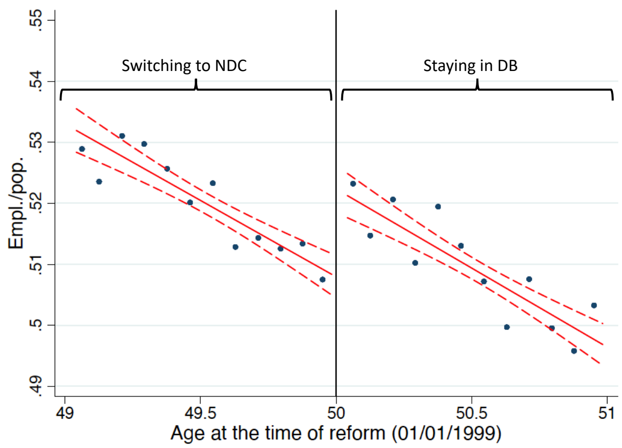

(b) Low-growth regions: **reform cohort**

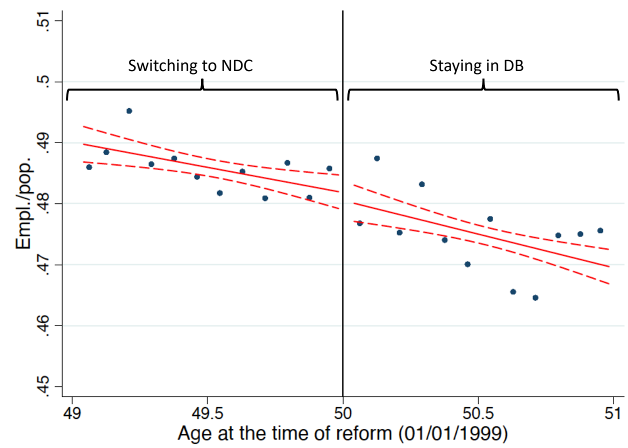

(c) High-growth regions: **placebo cohort**

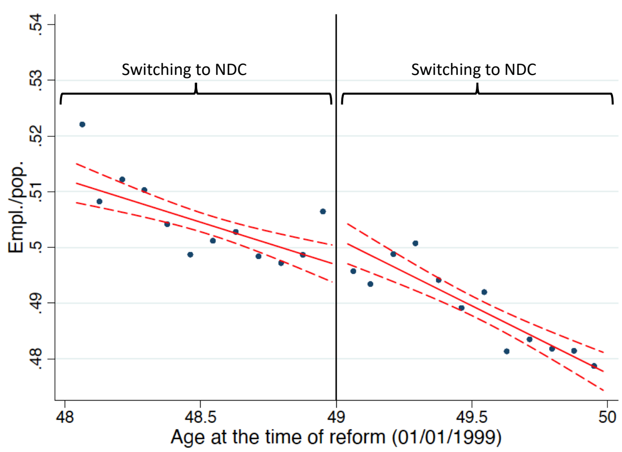

(d) Low-growth regions: **placebo cohort**

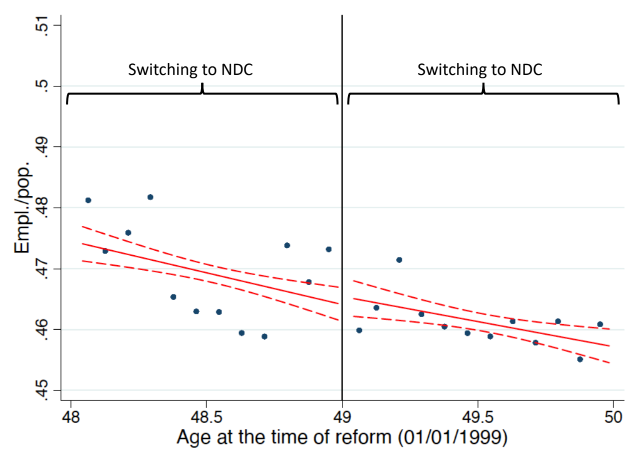

Notes: This figure plots the fraction of individuals who have positive earnings in a given year by month of birth (measured as the age on 01/01/1999). The top two panels, (a) and (b), show our treatment results. Individuals younger than age 50 on 01/01/1999 are in the new NDC scheme, while older individuals 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

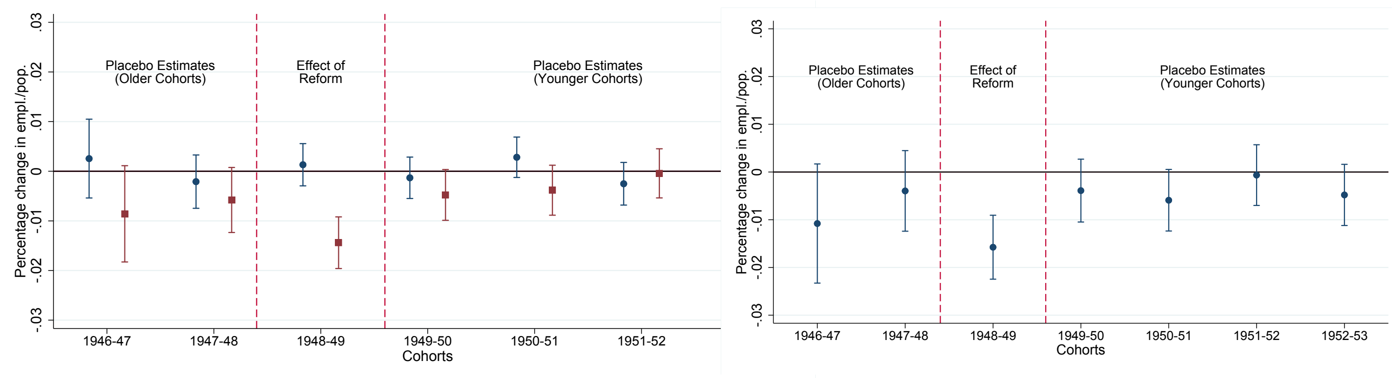

(a) Employment change for high and low growth regions (b) Employment change, difference (high-low)

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. For instance, we find a statistically insignificant 0.39 (s.e. 0.33) percentage point difference for the 1949-50 placebo cohort. This is evidence that

[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.

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. 0.30) 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 13.4% among high wage individuals and 14.2% among low wage individuals in high wage growth regions. Likewise, pension wealth falls by 11.9% among high wage individuals and 15.0% 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 0.28% increase in employment. We attribute the 2.50% (i.e., the difference between -2.22% and 0.28%) 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

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

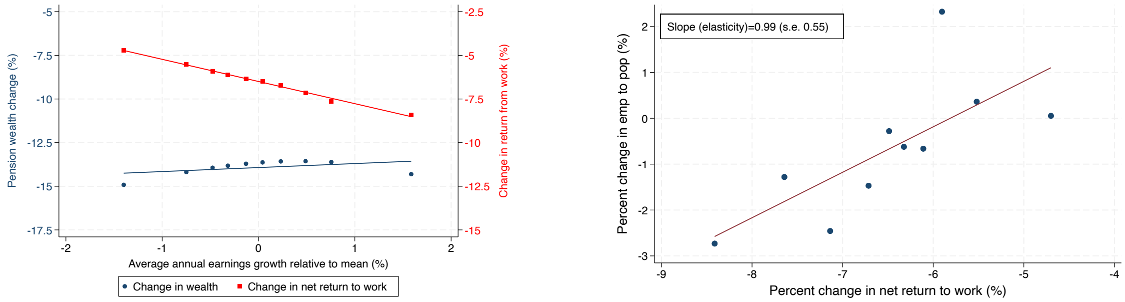

(a) Percent Change in Work Incentives and Pension Wealth Across Locations

(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

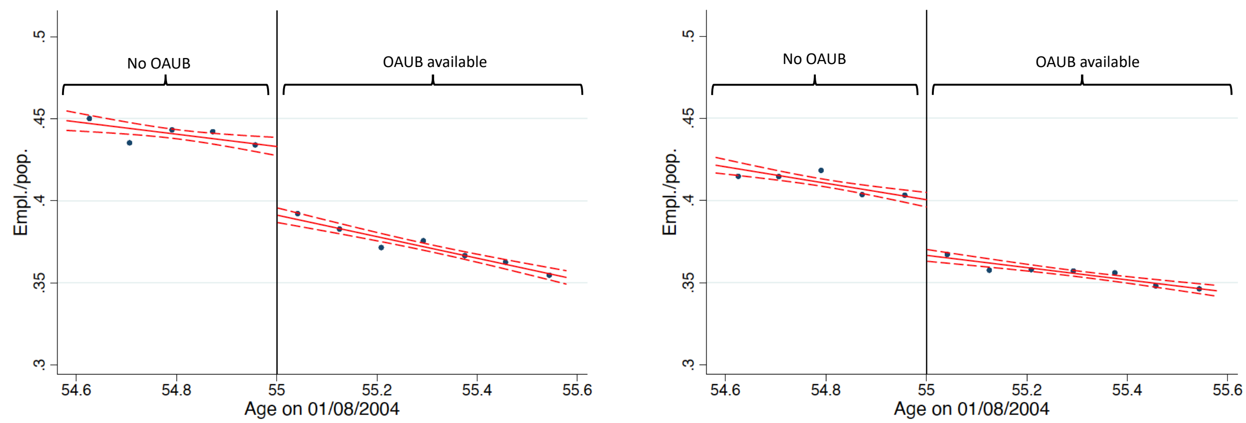

(a) High earnings growth regions

(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 0.99-1.12, 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 difference in the change in incentives (between the high- and low-growth regions) is 4.42%, while the difference in 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:

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

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

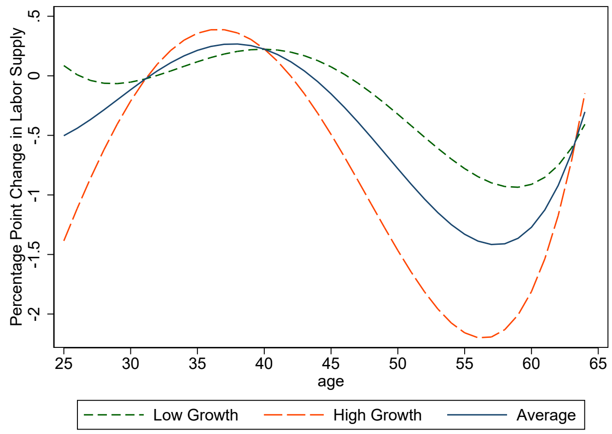

Notes: This figure plots the percentage point change in the employment-to-population ratio at each age coming from switching from the DB pension scheme to the NDC scheme. We predict behavior for a cohort that spent their whole working life under DB work incentives and compare it to a cohort who spent 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.

Replication Package The replication package for this research is available on Zenodo at https://doi.org/10.5281/zenodo.19415764.

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.

Klejdysz, Justyna and Tomasz Zawisza, "Taxation and Business Entry: Evidence from the Polish Self-Employment "Flat" Tax," *Journal of Public Economics, forthcoming*, 2026.

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.

Polish Central Statistical Office , "Labour Force Survey," 2012.

Polish Central Statistical Office, "Household Budget Survey," 2017.

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.

Ye, Han, "The effect of pension subsidies on the retirement timing of older women," *Journal of the European Economic Association*, 2022, *20* (3), 1048–1094.