NBER WORKING PAPER SERIES

LABOR SUPPLY AND THE PENSION-CONTRIBUTION LINK

Eric French
Attila S. Lindner
Cormac O'Dea
Tom A. Zawisza

Working Paper 30184
http://www.nber.org/papers/w30184

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2022

We thank Sreevidya Ayyar, Hanxiao Cui, Lucas Finamor and Pawel Struski for excellent research assistance. We thank Jeffrey Liebman for discussing this paper 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 and Owen Zidar 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). Any errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Labor Supply and the Pension-Contribution Link
Eric French, Attila S. Lindner, Cormac O'Dea, and Tom A. Zawisza
NBER Working Paper No. 30184
June 2022
JEL No. D15,H55,J22,J26

ABSTRACT

We estimate the impact of public pension systems on labor supply far from the normal retirement age by exploiting Poland's switch from a Defined Benefit to a Notional Defined Contribution scheme for men born after 1948. Using the universe of taxpayers and this sharp cohort-based discontinuity in the link between current contributions and future benefits, we estimate an employment elasticity with respect to the return to work of 0.44 for ages 51-54. We estimate a lifecycle model that matches these results. The model implies that the change in the contribution-benefit link from the reform increases employment among those in their 30s but decreases it at older ages, reducing overall labor supply across the lifecycle by 2 months.

Eric French
Faculty of Economics, University of Cambridge
Austin Robinson Building
Sidgwick Avenue
Cambridge, CB3 9DD, UK
United Kingdom
Eric.French.econ@gmail.com

Attila S. Lindner
Department of Economics
University College London
30 Gordon Street
London
WC1H 0AX
United Kingdom
and CERS-HAS, IZA and IFS
a.lindner@ucl.ac.uk

Cormac O'Dea
Department of Economics
Yale University
37 Hillhouse Avenue
New Haven, CT 06511
and NBER
cormac.odea@yale.edu

Tom A. Zawisza
Department of Economics
University College London
30 Gordon Street
WC1H 0AX
United Kingdom
tazawisza@gmail.com

A data appendix is available at http://www.nber.org/data-appendix/w30184

1 Introduction

In most OECD countries, the share of labor income going to social security contributions (SSCs) exceeds the share going to income taxes.[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 (Browning, 1975; Liebman et al., 2009). A tightening of this link can further alleviate these distortions. Furthermore, since the link between SSCs and future benefits is often tenuous (see e.g. Blinder et al., 1980; Burkhauser and Turner, 1985; Prescott, 2004), it has been argued that there is potential for substantial efficiency gains from reforming pension systems (Auerbach and Kotlikoff, 1985; Feldstein and Liebman, 2002; Lindbeck and Persson, 2003).

For this reason, many governments and multinational organizations have considered proposals to tighten the link between current contributions and future benefits. For instance, the World Bank (e.g. World Bank, 1994) and IMF have advocated tightening this link by switching from a Defined Benefit (DB) pension system to a Notional Defined Contribution (NDC) system.[2] 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 – a key argument posited by studies that advocate tightening the link.[3] The responsiveness of labor supply at younger ages to pension incentives is an open question in part because people seem to not always 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).

---

[1] In 30 out of the 37 OECD countries, the share of personal income tax is less than half of the total tax wedge (the sum of social security contributions and personal income taxes). In 19 countries, including Poland, the share of personal income tax is less than one third of the total tax wedge (see Figure A.1).

[2] In particular, IMF research argued that DB systems create a "loose link between benefits and contributions and transform the contribution rate into tax, reducing employment" (see page 7 of de Castello Bronco, 1998).

[3] It is well established that retirement incentives have a large impact on retirement decisions, and it is often the case that large implicit taxes hinder older age employment (Gruber and Wise, 1999). Nevertheless, if incentives only matter close to retirement age, then tightening the link between benefits and contributions throughout may cause little efficiency gain.

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.[4] As emphasized by the architects of this policy change, one of the most important elements of the switch to the NDC scheme was to "introduce a strong link between contributions and benefits" (see page 59 in Chlon et al., 1999).

An advantage of the setting we study is that the Polish reform switched from a DB scheme to an NDC scheme, leaving other features of the pension scheme (such as the retirement age, private pensions and funding status) largely unaffected around the discontinuity, whereas reforms in other countries usually changed several pension features all at once. This allows us to isolate changes in incentives from other features of the pension system.

The reform had a considerable impact on incentives to work. While work incentives generated by the NDC system differ little throughout the life cycle, the old DB system in Poland made earnings at certain ages particularly valuable for pension wealth accrual. One of the key reasons for this difference is that, in the DB system, pensions are calculated based on average earnings over a selected subset of "best" years (usually 10) in an individual's earnings history, whereas the proportionality built into the NDC system means that no prominence is given to earnings in any particular years. This feature of the DB system generates stronger work incentives than the NDC system at points in the lifecycle when wages are high, and weaker incentives when wages are low.

An individual's "best" years depend on that individual's life cycle earnings profile. For those with steeper earnings profiles, pensions in the DB system often depend heavily on earnings near age 50 since these are typically the highest earning years in one's lifetime. For these individuals, reducing labor supply at age 50 has a large impact on the average earnings base that is used to calculate benefits, providing strong labor incentives at that

---

[4] 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 tightened by introducing a proportionality between each individual's contributions and the benefits they receive. Furthermore, similarly to the funded defined contribution scheme, the rate of return reflects changes in economic prospects and growth.

age. Conversely, for individuals with flatter wage profiles, incentives under DB rules diverge less across the lifecycle. We exploit these differences in the change in incentives to identify the effect of changing the link between contributions and future benefits. In particular, we will separately study the effect of the reform in regions with high earnings growth, where individuals have steeper lifecycle earning profiles, and in low earnings growth regions, where individuals' earnings profiles are flatter.

Besides the change in incentives, the new NDC scheme led to a substantial reduction in pension wealth for individuals under the new rule. To isolate the effect of incentive changes from the effect of a reduction in pension wealth, we exploit the fact that the reform caused similar losses in pension wealth in high and low earnings growth regions, while the change in incentives was substantially different across regions. We calculate that the difference in the change in the net return to work between high and low growth regions induced by the policy change was 5.2 percent, while the difference in the induced change in pension wealth was only 0.35 percent. Therefore, by comparing responses to the policy in high wage-growth and low wage-growth regions we can capture, the effect of changes in incentives net of the effect of changes in pension wealth.

We estimate the employment responses to the pension reform by exploiting the sharp cohort-based discontinuity created by the reform. 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.[5] 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 occurring between 2000 and 2002.[6] 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.

---

[5] Since the introduction of the reform was more gradual for women, we focus throughout the paper on men. Nevertheless, in Table A.4 we report estimates for women. In line with the gradual introduction of the reform, we find qualitatively similar results for women with more muted responses.

[6] The terminal year for our main analysis is 2002 due to an unanticipated and substantial change in the old age unemployment benefit program that differentially impacted labor supply incentives of the 1948 and 1949 cohorts.

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 around one percentage point (or two percent) more at ages 51-53 than that in the low-growth regions. 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.[7]

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 5.23 percent, while the difference in the employment increase is 2.29 (s.e. 0.95) percent, the employment elasticity with respect to work incentives is 0.44 (s.e. 0.18). 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 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 provide a benchmark against which to compare our forward-looking elasticity, we estimate the labor supply response to an unanticipated radical reform in 2004 that impacted *contemporaneous* work incentives for the same population of workers. 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 generous unemployment benefits and find an employment elasticity of 0.58 (s.e. 0.04). This is larger than our elasticity estimated in response to the pension reform, suggesting that people are somewhat less responsive to changes in future pension benefits than to the contemporaneous return to work, consistent with modest discounting of future benefits. The old age benefit reform changed incentives similarly in high- and low-growth regions. We also find that estimated labor supply

---

[7] On average, men ages 50-53 in Poland expect to start collecting their pension payment for the first time at age 63 according to the 2005-2009 waves of SHARE data.

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 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.

We present several robustness checks. We show that our results are robust to alternative ways of implementing the regression discontinuity design. We also show that 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, 1950 versus 1951 and so on). We additionally show that our results are robust both to alternative assumptions on calculating the change in incentives and pension wealth, and to applying finer geographic dis-aggregation instead of simply focusing on the difference between high and low earnings growth regions.

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 regression discontinuity estimates from the reform and use the estimated model to compare the effects on labor supply, at all ages, of a move from the DB system to the NDC system. When comparing labor supply under the NDC system to that under a DB system that is adjusted so that government revenue under both systems would be equivalent, we find that altering the age structure of incentives as a result of the switch to the NDC system caused a fall in overall labor supply across the lifecycle of 2 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 in the NDC system. Contributing to this is the fact that labor supply is less responsive for those in their 30s than for those at older ages, and so the improved labor supply incentives at earlier ages yield less additional labor supply than that which is lost due to reduced work incentives later in the lifecycle. This highlights that the link between SSC contributions and future benefits should be strongest at ages when labor supply is most responsive.

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 that we study in this paper. While we apply a similar methodology to calculate the impact of the pension reform on incentives and pension wealth, our paper focuses on labor supply and not savings decisions. Additionally, instead of relying on survey data sample sizes and/or a gradual implementation of a pension reform as in the previous literature, we can combine population-level administrative data with a sharp discontinuity in changes in incentives. This allows us to compare the behavior of very similar individuals under the DB and NDC rules and avoid identifying responses by comparing the behavior of cohorts which were distant from each other.

There is also an extensive literature examining labor supply responses to retirement incentives. That literature, however, focuses almost exclusively on labor supply responses close to the retirement age: see Lumsdaine and Mitchell (1999); Diamond and Gruber (1999); Feldstein and Liebman (2002); Krueger and Meyer (2002); Coile (2015); Blundell et al. (2016) for reviews and Fetter and Lockwood (2018); Gelber et al. (2018); Krueger and Pischke (1992); Manoli and Weber (2016) for notable 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. Our paper, instead, 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. A notable exception is Bovini (2019), who estimates significant labor supply responses of individuals aged 46-56 to Italian pension reforms. While the findings of that paper are consistent with ours, the labor supply responses reflect the combined effect of changes in the pension eligibility age, changes in pension generosity, and changes in contribution-benefit link, and the contribution of each of these factors cannot be evaluated.[8]

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

---

[8] Other papers have studied the effect of labor supply responses around age 50 to changes in the pension eligibility age (as opposed to the link between contributions and benefits) using survey data (see e.g. Jean-Olivier et al. (2010); Carta and De Philippis (2019)). Focussing instead on the reporting behavior of self-employed workers, Dean et al. (2020) show that self-employed workers increase their reported earnings in years that enter the base for pension calculations.

et al., 1998; Barro et al., 1986; Carey and Rabesona, 2003; Kleven, 2014; Mendoza et al., 1994; 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., 2019) and an evolving literature on optimal tax policies in dynamic contexts (see e.g. Huggett and Parra, 2010; Golosov et al., 2016; Kocherlakota, 2010; Jones and Li, 2020).

There is very little direct evidence on whether individuals internalize how their contributions impact future benefits. Disney (2004) provides cross-country evidence suggesting that pension incentives affect prime age labor supply. Gruber (1997) and Bozio et al. (2019) provide evidence that the incidence of SSCs differs from that of taxes whenever there is a tight link between SSC contributions and benefits. Here instead, we show the importance of this contribution-benefit link for labor supply. We do this using a large fundamental reform generating a sharp discontinuity. Furthermore, we use this new evidence and a lifecycle model to highlight how these incentives will affect labor supply across the whole of the lifecycle.

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 the results of our estimation exercise. 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

Taxes and social security contributions (SSCs) affect the net return to work and thereby labor supply. Nevertheless, SSCs and other payroll taxes differ from standard income taxes because SSC payments are often linked to future benefits (Browning, 1975). This linkage, if recognized by the individual, could mitigate the distortions resulting from SSCs.

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

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

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

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

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

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

Our definition of the employment elasticity is closely related to the standard formula.

The main difference here is that the variation in the net return to work is coming from the change in the pension contribution-benefit link rather than from the change in the tax rate, which was held constant. Because there was no change in tax rates from the reform, the change in net return to work can be written as follows:

$$\Delta nrw_t = (E_t[\Delta PV_{it}^{NDC}] - E_t[\Delta PV_{it}^{DB}]),$$

and so the percent change in the net return to work coming from the pension reform is given by the following formula:

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

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.[9] 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.[10]

---

[9] Polish name: Zaklad Ubezpieczen Spolecznych.

[10] 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

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

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

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

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

---

state-managed notional account (Leifels et al. (2010)). In the paper, we assume that all workers are fully enrolled in the notional account.

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:

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

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

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

where $(1 + r^{NDC})$ is the real uprating factor on accumulated capital and $\tau^{ss} \cdot y_{it+1}$ is the contribution to the notional account.[11] 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.[12]

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,

---

[11] Here the amount τ^{ss} is sum of the "worker" and "employer" social security contributions, which is 0.1952. This is slightly different from the Social Security tax rate of 0.1871 described in equation (1), which includes additional taxes to pay for disability and sickness benefits but does not include employer contributions. For simplicity, we avoid using separate notation here and in equation (1) for τ^{ss}, but when we calculate the net return to work we take into account that these two rates are not exactly the same.

[12] 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.

such individuals were given “starting capital” in their notional accounts, calculated based on their contributions history.[13]

Contribution rates. The social security contribution rate to the pension system τ^{ss} 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.2 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.[14] Although these exemptions could bias our estimated labor supply responses towards zero, we show below that accounting for this has

---

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

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

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

[14] 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.

only a small effect on our estimates.

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.[15] 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. Individuals were eligible for old age unemployment from age 60 if the termination of employment was caused by the employer.[16] In 2002, the age threshold was reduced to age 55, which was in turn repealed in August 2004, at which point it went up again to age 60. As a result, individuals who were born in 1948 and were 55 years old in 2004 were eligible for the old age unemployment benefit, but individuals who were born in 1949 and who only reached age 55 after August 2004 were not eligible anymore. This created a large discontinuity in eligibility for the old age unemployment benefit between the 1948 and 1949 cohorts from 2004 onward. Moreover, in 2003, the 1948 cohort was eligible for the old age unemployment benefit because of having reached age 55, while the 1949 cohort was not. To ensure that our estimates do not capture the differential effect of the old age unemployment benefit on the two cohorts, we focus on the years 2000-2002 in our main analysis.

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 (Zl 2,296 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 employ-

---

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

[16] The old age unemployment benefit was called “swiadczenie przedemerytalne” in Polish. Those who received were eligible to get the unemployment benefit until they reached the normal retirement age.

ment responses are 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.

We use data for the years 2000-2002 for estimating the change in the contribution-benefit link resulting from 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. 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 A.1 we show that the employment rate calculated in our administrative data lines up well with the employment rate calculated using two representative household surveys: the Polish Household Budget Survey (HBS) and the Polish Labor Force Survey (LFS). The fraction of individuals in non-agricultural employment for the 1948 and 1949 cohorts is 48% in the LFS (based on 9,485 observations), which is very similar to our estimate in the administrative data (49% based on 1,669,539 observations). The estimated total employment rate (including agriculture workers) for the 1949 and 1950 cohorts is 60% in the LFS and 61% in the HBS.[18] Problems of underreporting do

---

[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

not appear to be of serious 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.

Figure 1 highlights this intuition by showing the percentage point change in the replacement rate from working at a given age under the DB and NDC schemes. The top panel is for an individual with a hump-shaped wage profile that peaks in his early 50s, whereas the bottom panel is for someone with a relatively flat wage profile. Because pension benefits depend on an individual's earnings y_{it} relative to aggregate earnings $\bar{y}_t$ over multiple years, we express wages relative to the average earnings in the economy (right axis). To construct the change in the replacement rate from working, we first calculate the individual's pension benefit assuming that he is in the labor force and, if not unemployed, works in *all years* up to 65. Next, we calculate the pension benefit assuming that he is in the labor force and, if not unemployed, works in *all years but one* up to 65. We use the DB benefit formula (equation (4)) and the NDC formula (equation (5)) to calculate the change in the benefit level at age 65, $b_{i65}^{\text{Employed}_t,k} - b_{i65}^{\text{Not employed}_t,k}$, resulting from not working at a given age t. We divide these benefit changes by the age 64 earnings before retirement, which gives us the change in the replacement rate (left axis).

The figure shows that the percentage point increase in the replacement rate under the NDC pension parallels the wage profile regardless of the shape of the wage profile. The figure also shows that, because the DB pension benefit formula uses income in the best 10 consecutive years, the increase in the replacement rate from working is very small in all but the 10 highest wage years. However, in those best 10 years, the increase in the replacement rate is potentially very large if wages in the 10 best years are much higher than at other ages. 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

---

some workers with a long working history in special occupations (such as metal workers or teachers) can retire already in that age range. These workers are unaffected by the reform and so our main empirical results are relevant for workers not employed in these special occupations.

to work in a narrow set of ages whereas the NDC scheme provides weaker incentives but at all ages. These hypothetical examples highlight the importance of modeling the wage profile for measuring the impact of working on pension benefits.

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 using an approach that follows the existing literature (e.g., Attanasio and Rohwedder, 2003):

- We calculate retirement benefits according to the legislation in effect in the year of observation. We take into account any reforms and future uprating rules that have been legislated up to the time of observation. We assume that people expect the current legislation to persist.
- We assume that, when forming their expectations, people take their current residence as given and fixed.
- We account for longevity uncertainty using year, age, and gender specific survival probabilities for the cohort aged 50 in 1999. We assume age and gender specific mortality do not change after 2016. The maximum attainable age is fixed at $T_{death} = 100$.
- We assume that individuals expect to retire at the male normal retirement age of 65.
- We assume that aggregate wage growth, interest rates, and benefit uprating factors are constant over time. We estimate these by taking averages over the post-reform period.
- We extend the framework in Attanasio and Rohwedder (2003) by also allowing for wage and unemployment risk. We estimate both from the data.

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

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

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

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

In the net present formula above, we observe the indexation factor ($r^{index} = 0.0116$), interest rate ($r = 0.0288$), the NDC uprating factor ($r^{NDC} = 0.0381$) and survival probabilities ($S_{s,t}$) in the data; see Appendix Online Appendix C for details. Since the pension benefit at age 65, $b_{i65}^{\text{Employed}_t,k}$, depends on earnings throughout the whole 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 are equal to the offered wage w_{it} if working (P_{it}=1) and 0 if not working ($P_{it} = 0$):

$$y_{it} = \begin{cases} w_{it} & \text{if} \quad P_{it} = 1 \\ 0 & \text{if} \quad P_{it} = 0 \end{cases} . \tag{10}$$

We assume wages are the sum of a deterministic and a stochastic component:

$$\log w_{it} = \mathbf{x}_{it}^T \kappa + \eta_{it} + \omega_{it} \tag{11}$$

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 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. The time trend and age polynomial interacted with region (controlling for a time trend and age polynomial) in the wage equation (11) 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. As a robustness check, we exploit variation in the wage profile at the level of the local administrative areas, of which there are more than 2,000.

Appendix Table C.4 reports estimates of wage parameters in equation (11). 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). The robustness to controlling for person-effects demonstrates that our classification of regions is not driven by differential selection of individuals out of the labor force.[20]

---

[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] 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 stochastic components of wages includes an AR(1) process η_{it}:

$$\eta_{it} = \rho\eta_{i,t-1} + \varepsilon_{it}, \qquad \varepsilon_{it} \sim N(0, \sigma_{\varepsilon}^2), \tag{12}$$

and also a MA(1) process ω_{it}:

$$\omega_{it} = \xi_{it} + \theta\xi_{i,t-1}, \qquad \xi_{it} \sim N(0, \sigma_{\xi}^2). \tag{13}$$

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.[21]

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. If someone is unemployed at period t as a result of the simulated unemployment shocks, then we define the net return to work to be zero. 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.[22] 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.

---

[21] 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.

[22] The chance that working at age 51 will be part of the best years calculations in the DB system is 47% in high growth regions and 44% in low 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 51-54, the net return to work declined 11.17% in high-growth regions (vs. 5.94% in low-growth regions), a difference of 5.23%.

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 14% in both the high- and low-growth regions. This pension wealth drop is larger than the one predicted by policy makers at the time of the reform, but in line with simulations of Lachowska and Myck (2018), who studied the change in pension wealth for 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.[23]

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 (14.26% in high-growth vs. 13.91% 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. In the Appendix Table C.3, we provide further detail on what drives the differential incentives across regions. First, the age polynomial part is very similar across locations and so it plays little role explaining the differential changes in incentives. Second, individuals at high-growth regions have lower earnings and steeper wage growth, which together can explain why incentives have changed more at high-growth regions. Notice that our regression discontinuity design implemented

---

[23] 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.

separately for high- and low-growth regions filters out the effect of differential labor market trends on labor supply, as those should have constant effect around the discontinuity.[24]

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 5.23% change, while the difference in the change in pension wealth was only 0.35%. In the next section, we study the response of the labor 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:

$$P_{it} = \alpha + \beta \mathbf{1}\{z_i < 50\} + f(z_i) + \varepsilon_{it}, \tag{14}$$

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, while 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

---

[24] In Table C.3 we also show that that the simulated change in incentives is very similar if we use estimates parameters from a specification with individual fixed-effects.

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 17th 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).[25] 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

---

[25] 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.

the net-of-placebo RDD is then:

$$P_{it} = \alpha^P + \beta^P \mathbf{1}\{z_i < k_{M(i)}\} + f^P(z_i) + \left(\alpha^M + \beta^M \mathbf{1}\{z_i < k_{M(i)}\} + f^M(z_i)\right) \cdot M_i + \varepsilon_{it}, \quad (15)$$

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 E.1 in Online Appendix E for all individuals. While there is some evidence of a lower female share, 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. In Table E.2, we 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 region type.

6 Results

Employment responses. We start our analysis by evaluating the effect of the reform on the employment rate. Figure 2 shows 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 that were 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. As noted in the previous section, we implement a "donut hole" RDD, excluding individuals born right around the discontinuity.

Figure 2 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 the tendency of employment rates to fall with age at older ages. As we described above, non-agricultural employment rates for men aged 51-54 in Poland were 49% in the period under consideration. Agricultural workers were unaffected by the policy change and are excluded from the 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. Panel A shows a 1.5 percentage point decline in the employment rate as a result of switching to the NDC scheme (left to the vertical red line). 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), but it could also reflect the effect of the reform on pension wealth. In Panel B of Figure 2, 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 decrease 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 of β from equation (14), which show 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 17th 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 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 decrease in low-growth regions. In Table A.2 of the Appendix, we also explore the sensitivity of the results to the bandwidth choice and show that the results are robust to the applied bandwidth.

As we discussed previously, the change in employment in high-growth and low-growth regions shows the combined effect of the wealth change 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. Figure 3 plots employment rates around a placebo discontinuity in high- and low-wage growth regions, in Panel A and Panel B respectively. Our placebo sample 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 right at the discontinuity, (i.e., these estimates incorporate the "donut hole"). Figure 3 shows 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 4 plots our estimated treatment effect 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.

The top panel of Figure 4 shows estimates for both low- and high-growth regions. The estimate for the 1948-49 main sample 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 main sample is a 1.48 (s.e. 0.26) percentage point reduction in employment, whereas for the 1949-50 placebo sample, we estimate a 0.43 (s.e. 0.26) percentage point reduction.[26] 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 sample is a 0.04 (s.e. 0.21) percentage point fall in employment. The placebo estimate for the 1946-47 cohort is larger in magnitude 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.

The bottom panel of Figure 4 shows the difference in the estimate between the high- and low-growth regions for each cohort. For the main sample, we estimate a statistically significant 1.58 (s.e. 0.34) percentage point difference, while all of the placebo estimates are smaller and statistically insignificant at the conventional levels. For instance, we find a statistically insignificant 0.39 (s.e. 0.33) percentage point difference for the 1949-50 placebo sample. Moreover, for women, who experienced a much smaller alteration to their pension system as a result of the reform, we do find much smaller, statistically insignificant differences in the employment rate around the 1948-1949 cohort discontinuity (see column (3) of Table A.4). These pieces of evidence highlight that the main effects are only found where the policy discontinuity is present, and we find no indication of a differential effect in other cohorts. See Table A.4 in Online Appendix A for further details on the 1947-1948 and 1949-1950 placebo samples.

Nevertheless, in Column 4 of Table 1, we also present our main results relative to the placebo estimates. We report the β^M from RDD regression equation (15). The estimated impact of switching to NDC on employment is -1.1 (s.e. 0.4) percentage points in the high-growth regions and 0.1 (s.e. 0.3) percentage points in low-growth regions. The difference between the high- and low-growth regions is 1.2 (s.e. 0.5) percentage points. This is very

---

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

similar to the simple RDD estimates in Column 2 (-1.6 percentage points). We will use these more conservative estimates in the benchmark analysis when calculating employment elasticities.

Implied elasticity. Table 2 presents our estimated participation elasticity. As we explained 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 percentage change in employment as a result of the reform. We use our net-of-placebo estimates of the percentage point change in Column 4 of Table 1 and divide it by baseline employment rates found at the discontinuity for the older cohort still 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 employment change (row 3) by the percent change in the net return to work (row 1), formally,

$$\eta^{H-L} \quad = \quad \frac{\left(\frac{\Delta P_t^H}{P_t^{H,DB}}\right) - \left(\frac{\Delta P_t^L}{P_t^{L,DB}}\right)}{\left(\frac{\Delta nrw_t^H}{nrw_t^{H,DB}}\right) - \left(\frac{\Delta nrw_t^L)}{nrw_t^{L,DB}}\right)}, \tag{16}$$

where $P_t^{h,DB}$ and $nrw_t^{h,DB}$ denote baseline employment rates and return from work in region h. The estimated elasticity is 0.44 (s.e. 0.18), which is statistically significant at the conventional level. Since the differential change in pension wealth was negligible (0.35%), this elasticity only captures the change in net return to work.

The change in net return to work at a given age 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 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.[27] 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 illustrate the fact that the NDC reform delivered non-trivial declines in pension wealth in each region. While it is not the focus of our paper, we can assess labor supply responses to these wealth changes. To do this, we first need to net out the impact of the change net return to work on labor supply. Table 2 shows that our estimate of the employment elasticity is 0.44. In the low-growth regions the change in the net return to work is 5.94%. Together, these imply that employment would fall (5.94%)(0.44)=2.6% in the absence of wealth effects. Given the observed rise of 0.28% in employment, the implied change in employment due to the wealth effect is (5.94%)(0.44)+0.28%=2.9% (for high growth regions the analogous calculation is $(11.17\%)(0.44) - 2.01\% = 2.9\%$). This is a *percentage* change in labor supply; given an employment to population ratio of around 0.5, this corresponds to a 1.45 *percentage point* change in labor supply arising from the wealth effect.

We can compare this magnitude to some recent estimates of labor supply responses to wealth shocks. Graber 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[28], causes employment to decline by 3.7 percentage points. The wealth shocks we study are smaller – the 14% decline pension wealth that our simulations show is equivalent to a wealth shock 1.5 times the average

---

[27] 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.

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

income in the economy.[29] Scaling our estimated 1.45 percentage point effect to the size of their wealth shock, our estimates would imply an employment change by 2.8 percentage points, which is not far from the 3.7 percentage point estimated in Graber et al. (2022). Furthermore, Lindqvist et al. (2020) finds considerably lower wealth elasticities in the Swedish context, suggesting that the our estimated wealth effects are within the range of existing estimates in the literature.

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 implementations of the regression discontinuity design. We calculate the implied elasticity for alternative estimates on the change in employment: a specification with a linear trend in birth date without applying the donut hole restriction, a specification with 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. 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 from zero. The point estimates vary between 0.66 (local linear, with donut) and 0.35 (with December-January donut) in the various specifications, both of which are close to 0.44 – our baseline estimate.

In Panel C, we assess the robustness of the results to alternative assumptions made in 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 the first row, we use a wage process applying the parametrization and estimates in French (2005) using data from the U.S. Panel Study of Income Dynamics. The second row in Panel C uses simulations where the estimated stochastic component of the wage is an AR(1) process with White Noise instead of an AR(1)+MA(1) process as in our benchmark specifications. In both cases, the change in net return to work is almost identical to the change calculated in our benchmark specification and thus the implied employment elasticities (0.43 with the French (2005) parametrization and 0.45 with the AR(1) process with White Noise) are almost identical to the the benchmark estimate (0.44) also. These estimates highlight that our main results are robust to alternative assumptions made in our simulation of

---

[29] 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.

incentives.

In the final row of Panel C of Table 3, 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 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 (0.41 instead of 0.44 in the benchmark case). This highlights that our estimates are not very sensitive to alternative assumptions on the parameter values of the NDC system.

In our analysis, we assume that individuals equally value the after-tax wage and the increase in the expected present discounted value from pension benefits from working. Furthermore, we have assumed that individuals discount future benefits at 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 Panel D of Table 3, we consider alternative higher interest rates when calculating the net return to work (equation (3)) in order to explore the sensitivity of our results to assumptions about discounting. We explore how the implied elasticity varies when the interest rate is 4% and 6%. With higher interest rates, we find that the implied elasticity is somewhat larger. For instance, if 4%, then the implied elasticity is 0.51. If the interest rate is 6%, the implied elasticity is 0.69.

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 $\$1,000$ annually or $\$2,000$ annually, the latter of which is equivalent to the annual earnings of someone earning the minimum wage. As can be seen, our estimates of the elasticity are the same under both definitions at 0.39, slightly lower than 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 in excluded sector following the procedure described in Online Appendix D. 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 0.48.

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 5, we assess the changes at a finer regional level. We calculate pension incentives over 2000 small administrative local areas in Poland.[30] The change in incentives is tightly linked to local area-level earnings growth, while the change in pension wealth is unrelated to the regional growth rate (see Appendix Figure A.3). 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.

In Figure 5 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 (15)) 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 that the slope is 0.62 (s.e. 0.34), which is close to the benchmark 0.44 elasticity.

Overall, the finer regional-level analysis underscores our benchmark results. The employment changes are tightly linked to the percentage change in incentives. The elasticity obtained at the finer regional level is very close to the elasticities obtained from comparing the response in high-growth regions relative to low-growth ones, though the estimates are more imprecise. This highlights that the estimated difference between high and low

---

[30] 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.

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 observed log wages among those reporting positive earnings. The estimates are small and insignificant, and the sign of the estimates is sensitive to the estimation method used. For this reason we focus on the extensive margin estimates, which are robust.

Employment vs. self-employment. In Appendix Table A.5, 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.5 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 is limited. Jobs in paid employment feature third party reporting, and so tax evasion is less prevalent in those type of jobs (Kleven et al., 2011). Therefore, our estimates on employment change pick up real responses to the policy and not simply some change in reporting behavior.

Future versus contemporaneous change in incentives. The estimated 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 attempt to disentangle the responsiveness of labor supply from people's valuation of future benefits, we estimate the labor supply response to a subsequent reform that impacted contemporaneous work incentives for the same population of individuals.

We exploit a radical change in eligibility for an old age unemployment benefit 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 years. Individuals could therefore take up the benefit if they reached age 55 by 1st August 2004 and demonstrated that their employment was terminated by the employer. 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 F.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 6 shows employment rates for men over the years 2005-2007 by age of the individual (in months) on 1st August 2004.[31] 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 high-growth and low-growth regions. We find that employment changes are similar across regions, with 3.3 and 4.1 percentage point changes in employment in low- and high-growth regions, respectively. When pooling both low- and high-growth regions, the estimated drop is 3.8 percentage points.

To calculate the employment elasticity implied by these employment changes, we calculate the percent change in the net return to work as a result of the reform, which we show in row (2) of Table 4. The net return to work incentives is defined as

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

where l here reflects whether someone has access to old age UI (l = OAUB) or not (l = NOAUB). The change in net return to work is coming from the change in outside option, $nrw_{it}^{\text{NOAUB}} - nrw_{it}^{\text{OAUB}} = -(u_{it}^{\text{NOAUB}} - u_{it}^{\text{OAUB}})$.

We calculate the implied elasticity as follows:

$$\begin{aligned} \eta^{Contemp} &= \frac{(P_t^{\text{NOAUB}} - P_t^{\text{OAUB}})/P_t^{\text{OAUB}}}{(nrw_t^{\text{NOAUB}} - nrw_t^{\text{OAUB}})/nrw_t^{\text{OAUB}}} \\ &= \frac{(P_t^{\text{OAUB}} - P_t^{\text{NOAUB}})/P_t^{\text{OAUB}}}{-(u_t^{\text{NOAUB}} - u_t^{\text{OAUB}})/nrw_t^{\text{OAUB}}}, \end{aligned} \tag{17}$$

where $P_t^{\text{OAUB}} - P_t^{\text{NOAUB}}$ is the difference in employment among workers who were eligible for the OAUB program and workers who lost that eligibility as a result of the policy change. This employment elasticity exploits the change in net return to work coming

---

[31] 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.

from the change in *contemporaneous* out-of-work benefits.[32]

Table 4 summarizes the key estimates and calculates the implied employment elasticity defined in equation (17). We calculate that the percent change in incentives was around 24.97% in high-growth and 25.31% 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 sufficiently long employment work history and the termination of the job must have been involuntary. While we do not directly observe the fraction of people who satisfy these other eligibility 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 eligibility rate. We provide details on how we arrive at these numbers in Online Appendix F.

When the eligible fraction of the population is 40%, then the estimated change in employment at the discontinuity translates to a percent change in employment among the eligible population that 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 0.84-0.96, 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%, then the implied employment elasticity is 0.56-0.64, depending on region. Our preferred estimate averages over the two regions, giving an estimate of 0.58. This estimate is 1.32 times larger than our baseline estimate on the pension reform (though the estimates are not statistically different from each other). Overall, these findings suggest that 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.

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

---

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

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.

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 this broader 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 type of reform we study across the whole lifecycle. In the following three subsections, we outline the model, give our estimates, and show the implications of the reform for labor supply over all ages. Further details on implementation are given in Online Appendix G.

7.1 Model

The model is one in which heterogeneous agents, who are subject to uncertainty over earnings 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 of how to divide resources between consumption and saving.

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

$$U(c_{it}, l_{it}; \nu_i) = \frac{(c_{it}^{\nu_i} l_{it}^{1-\nu_i})^{1-\gamma}}{1-\gamma}, \tag{18}$$

where c and l are, respectively, consumption and leisure. The quantity of leisure consumed is given as $1-hP$, which is equal to an endowment of leisure (normalized to 1) less a share of that endowment (h, set to 0.3) foregone in periods when the agent works. We assume that 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:

$$\nu_i \sim N(\mu_\nu, \sigma_\nu^2). \tag{19}$$

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}.

Earnings, Income and Assets In each period, each agent has an earnings potential w_{it}. This has a deterministic component that evolves with age and a stochastic component that follows an AR(1) process. We allow the deterministic component to vary with region of residence. This allows the model to be used to study the implications of the differing change in incentives across 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 each period, agents accrue an increment to their income which is proportional to their wage and depends on each of a) which 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 wealth accrual in our setting. Full details of how we model pension wealth accrual are given in Appendix G.1.

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

$$a_{it+1} = (a_{it} + y_{it} - c_{it})(1 + r), \qquad a_{it+1} \geq 0 \tag{20}$$

where y_{it} is income. Prior to retirement (age 65), income is equal to after-tax wage income if working and is equal to unemployment benefits when not working. In addition, we model additional returns to work from pension accrual as an increment to their income in a method similar to that used in French and Jones (2011). Thus, the difference between income from work and income from not work is the same net return to work as in equation (1) and is the same concept used earlier in the paper. After age 65, the individual retires.

Model Solution and Summary This model contains six state variables which we collect in the vector: $\mathbf{X}_{it} = \{\text{region}_i, \nu_i, t, a_{it}, \text{offer}_{it}, w_{it}\}$. Two of these – the agent's region of residence $\text{region}_i \in \{\text{low growth, high growth}\}$ and their consumption weight (ν_i –their 'type') – represent permanent heterogeneity, and four of which – age (t), assets (a_{it}), the presence (or otherwise) of an employment (offer_{it}) and wages (w_{it}) – vary across the lifecycle. Agents maximize

$$V_t(\mathbf{X}_{it}) = \max_{\{c_{it}, P_{it}\}} U(c_{it}, l_{it}; \nu_i) + \beta \Big(s_{t+1} \mathbb{E}_t V_{it+1}(\mathbf{X}_{it+1}) \Big) \tag{21}$$

subject to the asset accumulation equation (20), leisure of $l_{it} = 1 - hP_{it}$ and the determinants of income described in Appendix G.1. We solve the model using value function iteration. Appendix G.2 gives more details on the decision problem and Appendix G.3 gives further details on our solution method.

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, and our choices are discussed in Section G.4. The parameters of the deterministic and AR(1) components of the earnings process are those in the simulation (see equation (11)). The Markov process governing unemployment is also that in the simulation, discussed in Section 4 and with parameter values given in Table C.1. In a second step, we estimate the two 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, which we denote by parameter vector χ (where $\chi = (\mu_\nu, \sigma^2_\nu)$). We estimate these using Indirect Inference, matching labor supply at the age of 50 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 discontinunity, were moved to the new 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. 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 age 51 (64.7%))[33]. This gives us three moment conditions and two parameters to estimate.

Our identification of the parameters of the structural model therefore leverages our causal estimates of the policy reform. The variance of the leisure weight 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. When reservation wages are more dispersed, there are fewer people near the employment margin, and thus labor supply is less responsive to the reform.

Panel B of Table 5 gives our structural estimates and their associated standard errors. Panel C shows the moments that the model targets, with all modeled moments lying within 95% confidence intervals of their empirical analogues. The model-implied wealth effects on labor supply are larger than the empirical effects, and so the model overpredicts the employment response to the loss of wealth due to the reform. However, the model very 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 -2.17. Recall that it is these cross-region differences that capture the intertemporal substitution effect, which is the focus of this

---

[33] See the discussion at the end of Section 3.2 and Table A.1.

study.

7.3 Results

To evaluate the effects of the incentives from different pension systems on labor supply across the whole of the lifecycle, we first use our estimated model to predict behavior for a cohort that spent their whole working life in a system with NDC work incentives. We then compare their employment to a cohort who spent their entire working life in a system with DB work incentives. To focus on the role played by changes in the net return to work we ensure that the two systems we are comparing are revenue-equivalent by scaling down the value of pension accrual for those in work under the DB system proportionally at each age by a factor that ensures revenue equivalence (Online Appendix G has further details on the implementation of this counterfactual experiment). We are comparing behavior under two systems which are equally costly to implement but which spread work incentives across the lifecycle in different manners.

Figure 7 shows the reform-induced change in labor supply (in percentage points) across the lifecycle in each region. The model-predicted effect of the reform at age 50 is similar to our estimated effect. This is partly by design (the model was estimated to target those), although the experiment here – the implementation of an NDC system over the whole lifetime here 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. The model also predicts negative effects on labor supply at the very start of working life. In these periods, when earnings are at their lowest, the DB system provides strong work incentives through the fact that a period of work adds to the number of 'contributory' years which multiplies the base for pension calculation (see the pension benefit formula in equation (4)). That base will be substantially higher than current earnings, which means that these years are a 'cheap' time to accrue DB benefits. In contrast, pension wealth accrual under the NDC system in these years is proportional to (low) earnings. When agents are in their 30s, on the other hand, the NDC system provides better incentives since NDC accrual is proportional to (now higher) earnings.

Figure 7 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 almost two months under the NDC scheme, compared to an equally-costly DB scheme. For those living in low-growth regions, increases in labor supply among those in their 30s and the falls at other ages almost exactly offset (with a net fall in average labor supply of less than half a month). For individuals in high-growth regions, however, labor supply over the lifecycle falls by over three months, with the large falls at the start and end of the career only partially offset by modest increases when those agents are in their 30s. 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 and Staubli 2015) and some studies suggesting less (e.g., Cribb et al. 2016).

An important 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 0.52 at age 30, 0.57 at age 40 and 0.68 at age 50.[34] For those in their 50s, there is a greater mass of agents close to the participation margin than for those their 30s, when (male) labor supply is very high.[35] Ages around 30 are, therefore, an expensive time for the work incentives to be sharpened from a government revenue standpoint: labor supply is high, and so extra work incentives from higher pension accrual have substantial revenue implications, and responsiveness is lower, so revenue gains from increased labor supply are modest. 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.

---

[34] See Appendix G.7 for details on how we calculate these elasticities and Figure G.4 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 applying extensive margin labor supply models to US data. For example, French (2005) finds employment elasticities of between 0.19 and 0.37 at age 40 and between 1.04 and 1.33 at age 60. Fan et al. (2019) (Figure 5) estimate elasticities that are lower than 0.2 for those at younger ages and that rise to 1 around retirement ages.

[35] At the very start of the lifecycle, when earnings are lowest and the therefore the value of the model's outside option (welfare) is highest, there is also a large mass of individuals close to the participation margin.

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 the expected retirement age. We demonstrate this by exploiting the 1999 Polish pension reform, which switched a Defined Benefit system to a Notional Defined Contribution 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 0.44 (s.e. 0.18).

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 also 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 and could increase inequality among pensioners (Orszag and Stiglitz, 1999; Diamond and Gruber, 1999). 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 gains with distributional concerns should be a central focus for both future research and for policy discussion.

References

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

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

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

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

_ and _ , "Evaluating fiscal policy with a dynamic simulation model," *The American Economic Review*, 1987, *77* (2), 49–55.

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.

Barro, Robert J, Chaipat Sahasakul et al., "Average Marginal Tax Rates from Social Security and the Individual Income Tax," *The Journal of Business*, 1986, *59* (4), 555–566.

Blinder, Alan S., Roger H. Gordon, and Donald E. Wise, "Reconsidering the Work Disincentive Effects of Social Security," *National Tax Journal*, 1980, *33(4)*, 431–442.

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 labor supply?," Technical Report, National Bureau of Economic Research, 2019.

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.

Browning, Edgar K., "Labor Supply Distortions of Social Security," *Southern Economic Journal*, October 1975, *42*, 243–252.

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.

Carey, David and Josette Rabesona, "Tax Ratios on Labour and Capital Income and on Consumption," 2003.

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.*

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 Castello Bronco, Marta, “Pension Reform in the Baltics, Russia, and Other Countries of the Former Soviet Union (Bro),” *IMF Working Paper*, 1998, *98* (11), 1–40.

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.

Disney, Richard, “Are Contributions to Public Pension Programmes a Tax on Employment?,” *Economic Policy*, July 2004, pp. 267–311.

Fan, Xiaodong, Ananth Seshadri, and Christopher Taber, “Estimation of a Life-Cycle Model with Human Capital, Labor Supply and Retirement,” 2019.

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.

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

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

Gruber, Jonathan, "The Incidence of Payroll Taxation: Evidence from Chile," *Journal of Labor Economics*, 1997, *15*, S72–S101.

_ 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.

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.

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.

Kocherlakota, Narayana R, *The new dynamic public finance*, Princeton University Press, 2010.

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

Krueger, Alan B and Bruce D Meyer, "Labor supply effects of social insurance," *Handbook of public economics*, 2002, *4*, 2327–2392.

_ and Jorn-Steffen Pischke, "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation," *Journal of Labor Economics*, October 1992, *10* (4), 412–437.

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 and Stefan Staubli, "How does raising women's full retirement age affect labor supply, income, and mortality," *NBER working paper*, 2015, *18660.*

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.

Lumsdaine, Robin L and Olivia S Mitchell, "New developments in the economic analysis of retirement," *Handbook of labor economics*, 1999, *3*, 3261–3307.

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.

Mendoza, Enrique G, Assaf Razin, and Linda L Tesar, "Effective tax rates in macroeconomics: Cross-country estimates of tax rates on factor incomes and consumption," *Journal of Monetary Economics*, 1994, *34* (3), 297–323.

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.

Orszag, Peter R. and Joseph E. Stiglitz, "Rethinking Pension Reform: Ten Myths About Social Security Systems," in Robert Holzmann, Estelle James Joseph E. Stiglitz, with Louise Fox, and Peter R. Orszag, eds., *New ideas about old age security: Toward sustainable pension systems in the 21st century*, Washington: World Bank, 1999, pp. 17–62.

Prescott, Edward C, "Why do Americans work so much more than Europeans?," Technical Report, National Bureau of Economic Research 2004.

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.

Figures

Figure 1: Incentives under the DB and the NDC Scheme for Different Wage Profiles

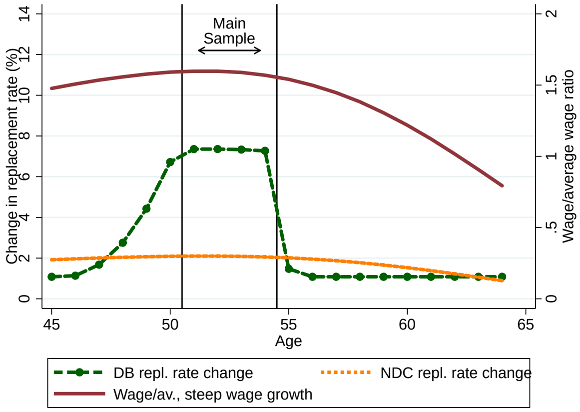

(a) Hump Shaped Wage Profile

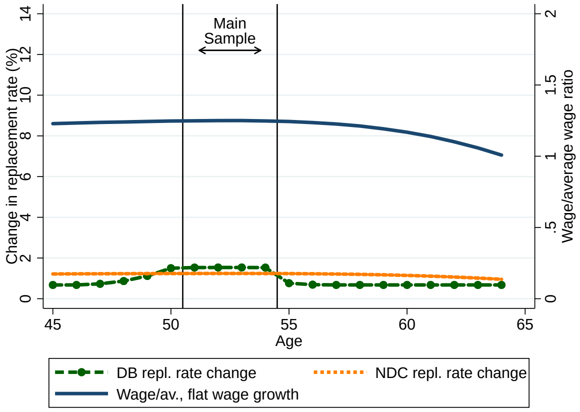

(b) Flat Wage Profile

Notes: This figure shows how incentives to work change over the life cycle for those facing the DB scheme and for those facing the NDC scheme. We highlight the change in incentives for two hypothetical individuals: one with a hump shaped wage profile (panel A) and one with a flat wage profile (panel B). The solid lines (red on panel A and blue on panel B) show a hypothetical individual's annual earnings, y_{it}, relative to the economy's average annual earnings, $\bar{y}_t$ (values shown in the secondary axis). We express individual earnings in terms of the economy's average annual earnings since this is what matters for $aime_i$ in the DB formula (see the text for details). Both panels show the percentage change in the replacement rate of pension benefits at age 65 coming from working at a given age, $(b_{65}^{\text{Employed}_t,k} - b_{65}^{\text{Not employed}_t,k})/y_{i,65}$, under the NDC and DB rules. In both panels we assume that individuals work until retirement age 65 except the age at which the change in employment status occurs.

Figure 2: Effect of the Pension Reform on Employment

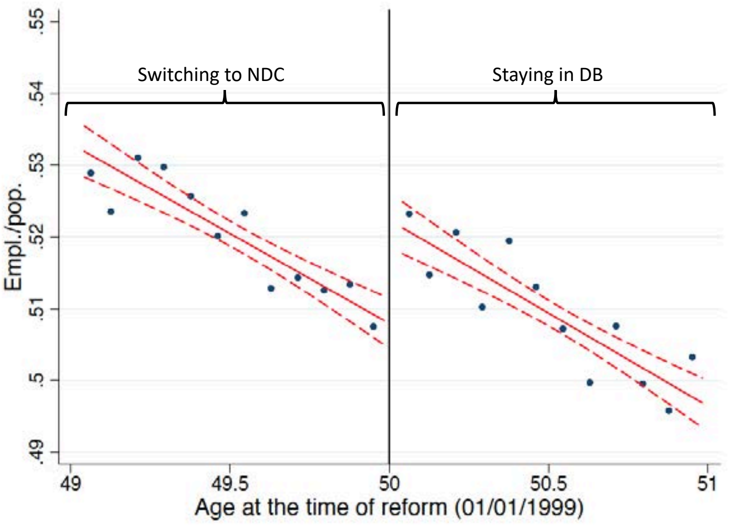

(a) High earnings growth regions

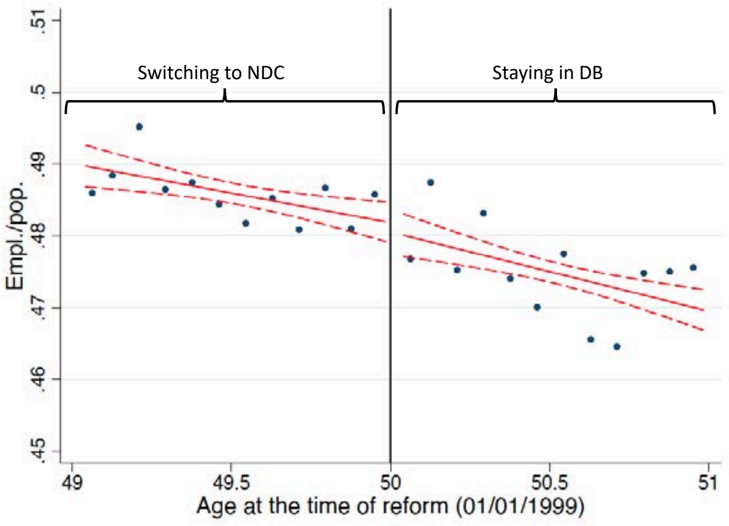

(b) Low earnings growth regions

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). 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.

Figure 3: Placebo Estimates of the Employment Effects

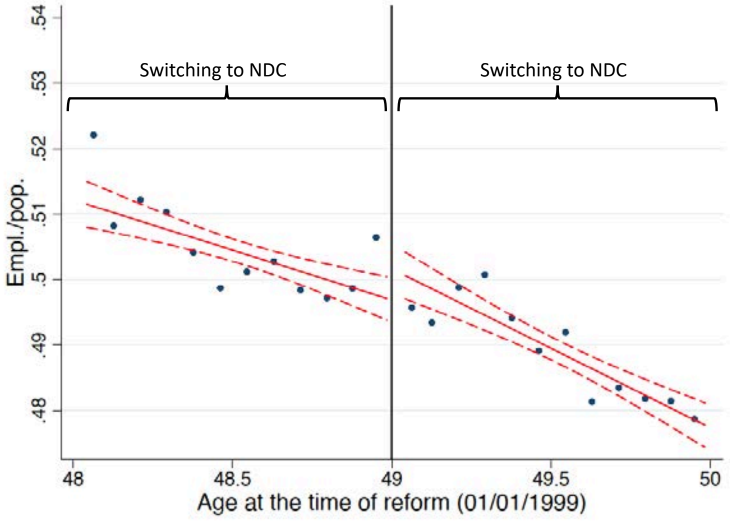

(a) High earnings growth regions

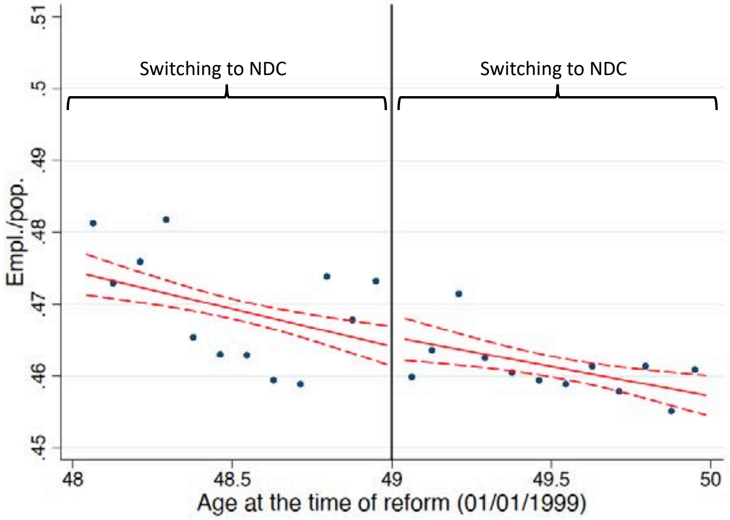

(b) Low earnings growth regions

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). 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. Otherwise, the plots are as in Figure 2. We calculate the fraction having positive earnings for every year and then average them for the years 2001-2003. 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 on January 1st, we apply a donut hole RD design and exclude those born between December 17th 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.

.

Figure 4: Change in Employment for Various Cohorts

(a) Employment change for high and low earnings growth regions

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

Notes: This figure plots the employment discontinuities estimated using the regression discontinuity design (see equation 14) 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 sample of 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. In each case, we estimate the donut RDD where we exclude those born between December 17th and January 5th. We apply a linear trend in birth date allowing for different slopes and intercepts at either side of the cutoff. The estimates based on the 1948-1949 and younger cohorts are for individuals when they were 51-54 years old. 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.

Figure 5: The Percent Change in Employment and in Work Incentives Across Locations

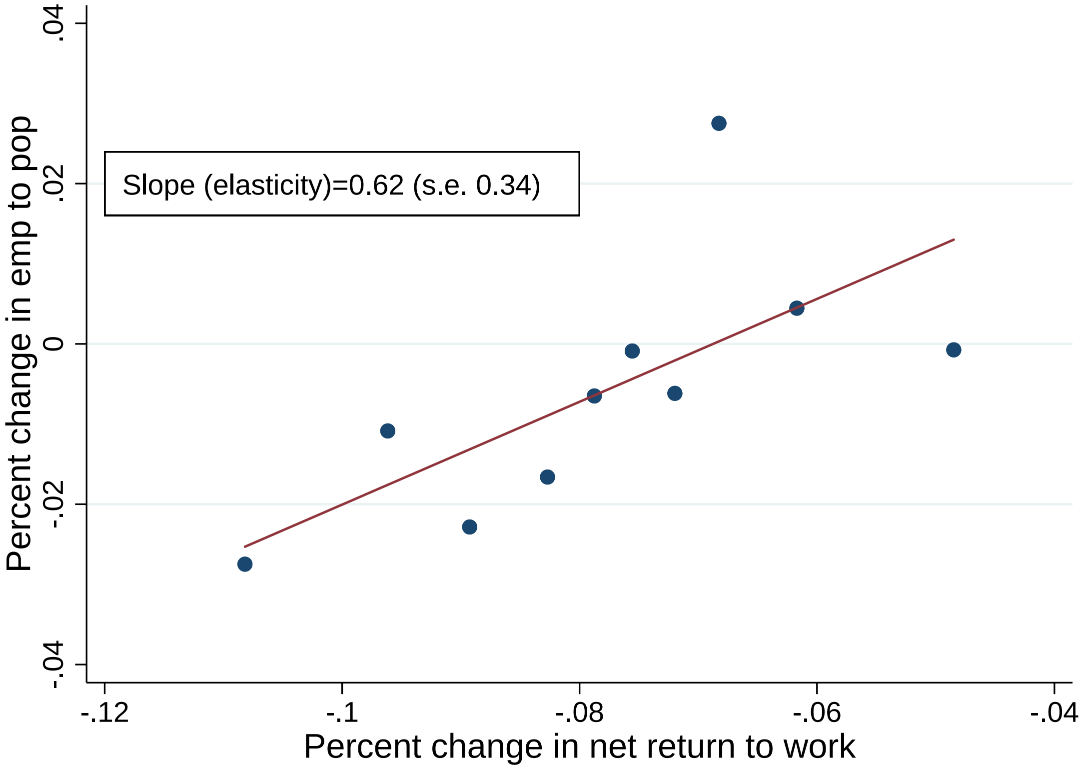

Notes: This figure shows the non-parametric bin-scattered relationship between the estimated employment change at the reform discontinuity (net-of-placebo estimates based on equation (15)) and the percent change in the net return to work across 2000 local areas. We estimate the employment changes for each local area separately by applying the local-linear RDD specification with the baseline donut. We group the 2000 administrative local areas into 10 equally-sized bins based on their change in net return to work and then compute the average change in incentives in the given bin (x-axis) and average estimated percent change in employment (y-axis). The percent change in the net return to work is calculated according to equation (3). We also plot a linear fit line using OLS, which represents the best linear approximation to the conditional expectation function. The slope of the linear fit, reported in the top left panel, shows the relationship between the percent change in employment and the percent change in incentives across areas and so it is an estimate of the employment elasticity.

Figure 6: Effect of the Old Age Unemployment Benefit (OAUB) Program on Employment

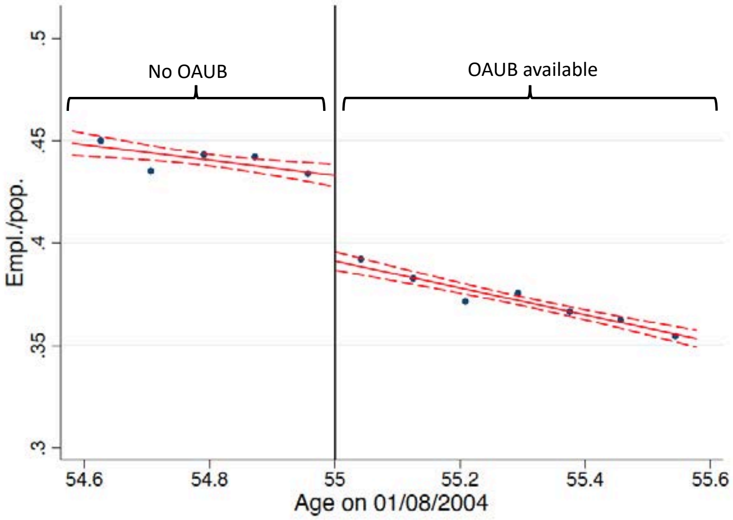

(a) High earnings growth regions

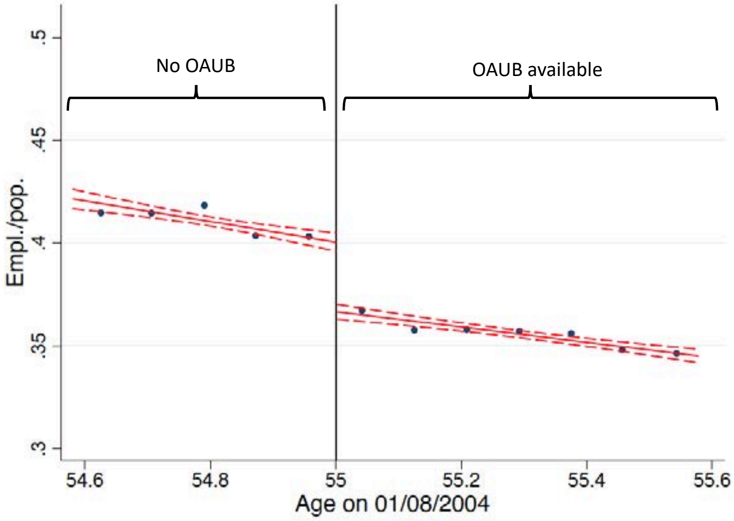

(b) Low earnings growth regions.

Notes: This figure plots the fraction of individuals who were 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 earnings growth regions are regions with an above median earnings growth rate between 2000 and 2013, while low earnings 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.

Figure 7: Effect of Switching to an NDC on Labor Supply Over the Lifecycle

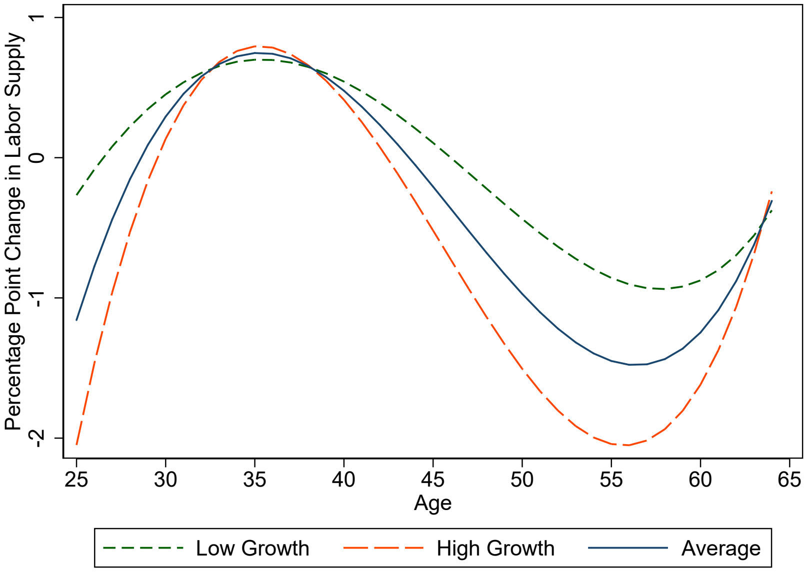

Notes: The figure shows the effects of the incentives from different pension systems on labor supply across the whole lifecycle. We use the model presented in Section 7 to predict behavior for a cohort that spent their whole working life under DB work incentives and compare it to a cohort who spent their 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 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. The dashed green line shows the effect of the reform for individuals for whom the earnings profiles are estimated based on the low earnings growth regions, the red dashed line shows the effect for individuals whose earnings profiles are estimated based on the high earnings growth regions, and the blue line shows the average of the two. The percentage point change in employment throughout the lifecycle is smoothed using a polynomial of order 4. Unsmoothed profiles are shown in Figure G.3 in Appendix G.6.

Tables

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)	(0.0032)
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 (14) (Columns (1)-(3)) or in equation (15) (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 17th 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 bandwith 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 (15). The placebo discontinuity is estimated between the 1949 and 1950 cohorts, both of which switched to the NDC system. We report robust standard errors in parentheses. For the local-linear regression we calculate robust standard errors following Calonico et al. (2014). Significance levels are: ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table 2: Employment Elasticity

	(1) High-growth	(2) Low-growth	(3) Difference (High-Low)
1. Change in net return to work (%)	-11.17	-5.94	-5.23
2. Change in pension wealth (%)	-14.26	-13.91	0.35
3. Change in employment (%)	-2.01	0.28	-2.29
	(0.71)	(0.63)	(0.95)
4. Employment elasticity (Row 3) / (Row 1)	–	–	0.44
			(0.18)

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). The details of calculating the change in the net return to work and pension wealth are described 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 based on equation (16): we divide the percent change in employment (row 3) by the percent change in the net return to work (row 1). To isolate the effect of the change in incentives from the effect of the reform on pension wealth, we calculate the employment elasticity only after we take the difference between the high and the low-growth regions (Column (3)). Standard errors are in parentheses.

Table 3: Employment Elasticity, Robustness

	(1) Change in net return to work (%)	(2) Change in pension wealth (%)	(3) Change in employment (%)	(4) Employment elasticity
Panel A: Baseline				
1. Linear trend RDD, net-of-placebo, donut sample	-5.23	0.06	-2.29	0.44
			(0.95)	(0.18)
Panel B: Estimation methods				
2. Linear trend RDD, net-of-placebo, full sample	-5.23	0.06	-2.99	0.57
			(0.88)	(0.17)
3. Linear trend RDD, net-of-placebo, Jan-Dec donut sample	-5.23	0.06	-1.85	0.35
			(1.04)	(0.21)
4. Local-linear RDD, net-of-placebo, donut sample	-5.23	0.06	-3.40	0.66
			(1.80)	(0.34)
Panel C: Calculation of net return to work				
5. AR(1) wage process (parameters from French (2005))	-5.31	0.40	-2.29	0.43
			(0.95)	(0.18)
6. AR(1) + White Noise wage process	-5.12	0.23	-2.29	0.45
			(0.95)	(0.19)
7. Actuarially fair uprating of NDC contributions	-5.57	0.06	-2.29	0.41
			(0.95)	(0.17)
Panel D: Interest rate				
8. $r = 0.04$	-4.46	0.06	-2.29	0.51
			(0.95)	(0.21)
9. $r = 0.06$	-3.30	0.06	-2.29	0.69
			(0.95)	(0.29)
Panel E: Definition of employment				
10. $1000 p.a. threshold	-5.23	0.06	-2.02	0.39
			(0.95)	(0.18)
11. $2000 p.a. threshold	-5.23	0.06	-2.02	0.39
			(0.95)	(0.18)
Panel F: Elasticity Formula				
12. Alternative adjustment for unaffacted workers	-5.23	0.06	-2.54	0.48
			(1.05)	(0.20)

Notes: This table shows robustness analysis for the estimated employment elasticity in Table 2. Panel A reports again the benchmark elasticity calculated based on the net-of-placebo estimates in Column (4) of Table 1. Panel B shows robustness of the elasticity to the estimation method of the change in employment. Row 2 reports estimates when we estimate the employment change at the discontinuity using a linear trend and not applying the donut hole restriction. Row 3 reports estimates when we apply a broader definition of the donut hole (omitting everyone born in January and December). Row 4 presents estimates using a local-linear RDD on the baseline donut sample. Panel C explores robustness to the calculation of the net return to work and pension wealth. Rows 5 and 6 show robustness to alternative ways of estimating the stochastic component in wages (see Equation (11)). Row 5 applies an AR(1) process where the parameterization comes from French (2005) instead of the AR(1) + MA(1) process used in our benchmark calculation. Row 6 explores instead an AR(1) process with White Noise, where we estimate the parameters by a GMM procedure described in Section 2. Row 7 explores an alternative assumption on calculating the net present value of pension benefits in equation (9). Instead of applying the actual risk free interest rate, pension indexation, and survival rates, we assume 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 Panel D, rows 8 and 9 explore the effect of changing the real interest rate from our baseline of 2.88% to 4% and 6%. Panel E explores changing the definition of employment we use as our dependent variable. Rows 10 and 11 show robustness to increasing the threshold of earnings above which individuals are deemed to be in employment to increase to $1000 and $2000 p.a., from our baseline of $547. Finally, Panel F explores an alternative adjustment for unaffected workers described in detail in Online Appendix D. All estimates in this table take the differences in the estimated employment change between high-growth and low-growth regions to isolate the effect of changes in net return to work. Robust standard errors are reported in parentheses.

Table 4: Elasticity Estimates using Contemporaneous Incentives

Region	(1) All regions	(2) High-growth	(3) Low-growth
1. Change in net return to work	25.13	24.97	25.31
2. Change in net wealth	0.0	0.0	0.0
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)	0.88	0.96	0.84
	(0.052)	(0.083)	(0.072)
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)	0.58	0.64	0.56
	(0.035)	(0.055)	(0.048)

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 (see Section F.1). The percent change in employment depends on the fraction of workers who were eligible for the OAUB program. In Online Appendix F we show that this eligibility is likely to fall between 40% and 60%. The percent changes in employment when 40% and 60% were eligible are shown in row 3 and 5, respectively. The implied employment elasticity calculated as in equation (F.4) is shown in row 4 and 6, respectively.

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)	0.485	(0.037)
Consumption Weight St. Dev (σ_ν)	0.145	(0.068)

Panel C: Model fit

Matched moments	Data	Model
Labor supply at age 50 (%)	64.7 (s.e. 0.01)	64.7
Reform labor supply effect, low-growth (%)	0.28 (s.e. 0.63)	1.37
Reform labor supply effect, high-growth (%)	-2.01 (s.e. 0.71)	-0.80
Implied difference in reform effect		
Reform labor supply effect, difference (high-low) (%)	-2.29 (s.e. 0.95)	-2.17

Panel D: Effect of switching from DB to NDC

	Effect
Net change in lifecycle labor supply, all	-1.8 months
Net change in lifecycle labor supply, low-growth	-0.4 months
Net change in lifecycle labor supply, high-growth	-3.3 months

Panel E: Frisch Employment Elasticity

	Elasticity
Frisch Employment Elasticity at age 30	0.52
Frisch Employment Elasticity at age 40	0.57
Frisch Employment Elasticity at age 50	0.68
Frisch Employment Elasticity at age 60	0.90

Notes: Panel A reports the key parameters which are set prior to estimation of the model presented in Section 7. We provide justification for the parameter values in the text and in Online Appendix G. Panel B provides the estimated structural parameters and their associated standard errors. We estimate these parameters using Indirect Inference (see the details in Online Appendix G). Panel C shows the three moments that the model targets, their standard errors and the model analogue predicted by the best fitting parameters. The final row of Panel C shows the model fit for the key object from which we estimate our headline result: 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. We use the parameter estimates and 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. Panel E gives Frisch Employment Elasticities at selected ages. These are estimated by perturbing wages (individually) at each age by 20% and calculating the percentage change in labor supply and dividing by 20%.