WORKING PAPER NO. 585

The Life-Cycle Effects of Pension Reforms: A Structural Approach

Claudio Daminato and Mario Padula

October 2020

CSEF - Centre for Studies in Economics and Finance
DEPARTMENT OF ECONOMICS – UNIVERSITY OF NAPLES
80126 NAPLES - ITALY
Tel. and fax +39 081 675372 – e-mail: csef@unina.it
ISSN 2240-9696

WORKING PAPER NO. 585

The Life-Cycle Effects of Pension Reforms: A Structural Approach

Claudio Daminato[♦] **and Mario Padula**[*]

Abstract
To assess the life-cycle welfare effects of pension reforms, we provide a dynamic stochastic model of saving, portfolio choice and retirement with a pension system that operates according to the notional defined contribution principle. Relying on the exogenous variation from a sequence of Italian pension reforms, we identify and estimate the model, which is then used to draw implications of alternative pension policies. Our results also shed further light on the mechanisms behind the offset between social security and private wealth and show the importance of labor supply at retirement as an insurance mechanism against shocks to pension wealth.

JEL classification: D15, E21, H31, H55, J26.

Keywords: Pension reforms, Life-Cycle, Savings, Portfolio Choice, Retirement.

Acknowledgements: We thank Irina Gemmo, Fedor Ishakov and Cormac O'Dea for helpful comments and suggestions. All the errors are ours. The opinions expressed are those of the authors and do not reflect the views of COVIP.

[♦] Department of Management, Technology and Economics, ETH Zurich

[*] Commissione di Vigilanza sui Fondi Pensione (COVIP), Ca' Foscari University of Venice and CSEF.

Table of contents

1 Introduction

The population aging, and the related challenges to the pay-as-you go social security systems, caused profound changes to the pension legislation in several OECD economies. While there is a wide between-countries variation in the extent, the speed and the timing of these changes, a common trait are the heated policy debates over alternative interventions on social security systems. Most often, these debates lack the discipline of an operational economic model to understand the effects of pension legislation changes on households decisions. Understanding these effects is crucial to the design of pension reforms, most importantly before the reform takes place. This paper aims at providing a validated framework to study the possible consequences of alternative pension policies.

A body of economic research focuses on the evaluation of a social program before its actual introduction, as part of the problem of studying the effect of policy changes without the availability of ex-post information. The ex-ante evaluation of social programs sheds light into the understanding of which range of effects to expect from the introduction of alternative policy changes (Todd and Wolpin, 2006b, Heckman, 2010). It can then provide a number of useful prescriptions to the policy makers.

Todd and Wolpin (2006a) and Attanasio et al. (2012) follow this approach to develop and estimate two different dynamic models of education choices to study the impact of the PROGRESA program on children's schooling attendance; Blundell et al. (2016a) rely on tax and benefit reforms in the United Kingdom to estimate a dynamic model of employment, human capital accumulation and savings and to analyze the effects of welfare policies. In this paper, we also propose an ex-ante policy evaluation exercise and exploit the - arguably exogenous - variation induced by pension reforms. We focus on Italy, an interesting case because of the dramatic changes to the pension legislation occurred in the early 90's. For younger generations of workers, who entered the labor market after 1978, the contribution model replaced the earnings model for the computation of pension benefits, while the earnings model was kept for the older generation of workers. Furthermore, the reforms also affected the eligibility criteria, again drawing a line between younger and older workers. The extent of the changes and the policies discontinuity among workers made Italy an ideal "laboratory" to study the impact of pension policies on households behavior.

Miniaci and Weber (1999), Attanasio and Brugiavini (2003), Bottazzi et al. (2006) and Bottazzi et al. (2011) use the Italian laboratory to investigate the effect of these pension reforms on households decisions, looking at consumption, saving, wealth and portfolio choice. Our exercise shares with theirs the same quasi-experimental variation, but differs in exploiting such variation to estimate reduced form effects, which are then used to inform an economic model of households behavior about the most relevant channels of the pension reform.

The reduced form effects suggest substantial responses of households to the pension reforms in terms of discretionary wealth accumulation, participation in the financial markets as well as expectations about their future retirement age. These estimates represent the first stage in our estimation exercise. In the second stage, we develop and estimate a dynamic stochastic life-cycle model, in which households maximize expected lifetime utility choosing consumption, the allocation of wealth to risky assets and the age of retirement, while facing uncertainty with respect to income, returns from the risky assets and mortality. The model features a rich characterisation of the Italian pension system before and after the pension reforms, explicitly incorporating the transition from an earnings model to a notionally defined contribution scheme. To estimate the model, we target the first stage impacts of the reforms on discretionary wealth and participation in the financial markets. The structural approach carefully replicates the institutional setting, allowing for (ex-ante) heterogeneity with respect to the sector of employment, which in turn determines the treatment status under the pension reforms. Moreover, because the pension reform targeted only workers with less than 18 years of contribution in 1995, the year-of-birth cohorts span from 1940 to 1970. Years of contribution in 1995 depend on year-of-birth. Therefore, variation in cohort membership implies variation in the treatment status.

Since we match the model-driven impacts of major pension reforms to their data-driven counterpart, we provide an arguably credible tool to conduct ex-ante policy analysis. In particular, adopting a structural approach allows us to overcome two limitations inherently associated with the usage of a standard diff-in-diff strategy to study the effect of pension reforms: (i) the non-structural nature of the diff-in-diff parameter, which does not allow to draw conclusions on the long-run behavior of households and implications for the design of future pension policies; (ii) the concerns about the credibility of the identifying assumptions, upon which the diff-in-diff strategy relies (parallel trend and linearity of the functional form). By using an indirect inference approach to the structural estimation with a diff-in-diff regression as auxiliary model, we obtain unbiased estimates of the structural parameters independently of the unbiasedness of the diff-in-diff estimate of the causal effect of the reform.

Our main finding is that the estimated model, with reasonable values for the structural parameters, matches key pre-intervention statistics and the average effects of the pension reforms estimated from actual data exploiting the diff-in-diff identification strategy. By validating a life-cycle model with the reduced form effects of pension reforms, we contribute to the vast literature that studies intertemporal choices of consumption and savings. Further, we are the first to estimate a fully fledged life-cycle model of savings, portfolio allocation and retirement for the Italian economy. The structural approach provides a number of important novel insights about the consequences of the pension reforms. First, we shed further light on the displacement effect between private and pension wealth, in that contributing to the literature starting from Feldstein (1974).[1] In particular, we highlight the role of labor supply at retirement in shaping the interplay between private and pension wealth. Second, and related, the model predicts substantial wealth effects on retirement behavior, identifying labor supply at retirement as an important mechanism households use to insure against shocks to pension benefits generosity.[2] To the best of our knowledge, this is the first paper that provides a structural estimate for the wealth effect on labor supply at retirement exploiting variation in benefit generosity. Our findings complement the results in Manoli and Weber (2016), who provide nonparametric evidence of substantial retirement decisions response to financial incentives using data from Austria. Further, we add to the literature that studies retirement decisions in life-cycle models (see, e.g., Blau 2008, French and Jones 2011, Haan and Prowse 2014) by explicitly introducing the dynamic incentives individuals face to postpone retirement in a notional defined contribution pension system. Third, the model shows that older households in working age experience larger welfare losses for a given reduction in benefits' generosity, providing a quantification and a rationalization of the "life-cycle" welfare effects of pension reforms.

Finally, we use the estimated model to show how alternative pension policies can have different implications in terms of individuals' retirement and saving responses, as well as wealth inequality.

The rest of the work is organized as follows. Section 2 describes the main features of the institutional framework, presents some stylized facts from the data and the empirical strategy to estimate the reduced form effects of the pension reform. In Section 3 we present the dynamic life-cycle model used to capture the behavior of households before and after the introduction of the pension reforms. The estimation results of the model are presented in Section 4. In Section 5 we discuss the life-cycle consequences of pension reforms, the role of labor supply at retirement as an insurance mechanism against shocks to social security wealth, conduct welfare analysis and two policy experiments. Section 6 concludes.

[1] Private pension funds were introduced in Italy in 1993, but social security - the first pillar - with its 33% contribution rate still plays the lion's share in the Italian pension system. For most households, due to the still limited degree of coverage of private pensions - 1/3 of the working age population - , pension wealth coincides with social security wealth.

[2] French (2005) first introduced retirement decision in a life-cycle model of labor supply and saving behavior focusing on the role of health. In contrast to his work, where retirement is modeled as a labor supply decision and households can reenter the labor market, we explicitly introduce retirement as an absorbing state.

2 Institutional setting, data and reduced form evidence

2.1 Pension reform

Until the early nineties, pension spending was increasing in Italy on a steady basis to reach 16.2% as ratio to the GDP in 1992, at the time the highest value among the industrialized countries. The high pension spending was the consequence of high replacement rates, earnings based benefits, generous provision for early retirement, and a large number of social pensions. This trend fueled the growing alarm over the sustainability of the Italian pension system. As a result, the pension legislation was profoundly revised, with two major interventions in 1992 and 1995, which increased the minimum years of contributions for pension eligibility and progressively introduced a notionally defined contribution model for pension benefits.[3] The reforms targeted workers who had less than 18 years of contribution in 1995 (the so-called middle-aged workers), while keeping unchanged the provisions for workers who had at least 18 years of contribution in 1995 (older workers). Pension benefits are computed according to the *earnings model* for older workers and to the *pro-rata model* (earnings model form for working years before 1995, and *contributions model* afterwards) for middle-aged workers. In the earnings model pension benefits are obtained multiplying the average earnings in the last year (or in the last years, depending on the sector of employment) by the product of number of contribution years and the so-called accrual rate. The contributions model links pension benefits to the entire history of earnings, therefore providing incentive to postpone retirement. Pension contributions are proportional to earnings and capitalized on an annual basis using a five-years moving average of the GDP growth rate. Pension benefits are obtained multiplying the sum of capitalized contributions by an age-increasing transformation coefficient.

While introducing flexible retirement age, the pension reforms induced a dramatic decrease in pension wealth for middle-aged workers, for a given retirement age.[4] In a life-cycle setting, the decrease of pension wealth, an arguably exogenous shock to life-time resources, should make households consume less, save more and, accordingly, increase their non-pension wealth. Relying on the exogeneity of the shocks to pension benefits, previous studies have already analyzed these effects in a reduced-form framework: Miniaci and Weber (1999) focus on consumption, Attanasio and Brugiavini 2003 on saving, Bottazzi et al. (2006) on private wealth and Bottazzi et al. (2011) on portfolio choices. In this paper, we rely on the same shock but use the reduced form impacts to quantitatively inform a structural model of

[3]Attanasio and Brugiavini (2003) and Bottazzi et al. (2011) provide extensive details on how these interventions changed both the pension award formula and the eligibility rules.

[4]See Bottazzi et al. (2006) for a simulation of the replacement rate before and after the reform.

households' behavior about the most relevant mechanisms in the response of households to shocks to pension wealth.

2.2 The data

The data come from the Bank of Italy Survey on Household Income and Wealth (SHIW) for the years $1986 - 2008$, which provides a representative sample of the Italian population of households. The SHIW is not only a standard reference to analyze Italian households' balance sheets but also quite unique in recording the joint distribution of several demographics, labor market status, earnings, hours of work, years of contribution to the social security system, consumption and asset holdings variables.

Our definition of total assets includes real assets and financial assets, net of financial liabilities. We define households as participating in the financial markets if they have non zero investments in one of the following asset classes: mutual funds, equity, shares in private limited companies and partnerships.

In both the reduced form and the structural analysis, we consider a unitary model for households' behavior and then keep only observations referring to the head of household and household-level information data.[5] Notice that Italy shows a large gender gap in labor market participation, which is also reflected in the SHIW data, where the labor market participation rate among men is around 86 percent and among women less than 47 percent. We drop households whose heads were born before 1935 and after 1975. These households are not marginal to the reform. We also drop information on households whose heads are not married and employed in either the private or public sector. To identify the treatment from the control group, we follow Bottazzi et al. (2006) and use the information on whether the head of the household works in the private or in the public sector. In addition, based on the years of contributions of the households heads, we distinguish older (at least 18 years of contribution in 1995) from middle-aged workers (less than 18 years of contribution in 1995). Therefore, the treated households are those households whose heads are a middle-aged workers.

Life-cycle profiles We focus here on the life-cycle dynamics of consumption, assets, hours of work and participation in the financial markets for middle-aged and older workers. Figure 1 reports the age-profiles of mean consumption, assets, and hours of work and median financial

[5]The head of the households is defined as the person responsible for the economic management of the households, which often coincides with the main income earner. In couples, the main income earners tend to be the male member. Around 80 percent of the household heads identified by the survey correspond to the main income earner in the household. Among the latter, about 78 percent are men.

markets participation, separately for middle-aged and older workers. We cannot draw any conclusion about the effect of the reforms at this stage, but the comparison of the age profiles for the two groups of households provides some preliminary evidence.[6]

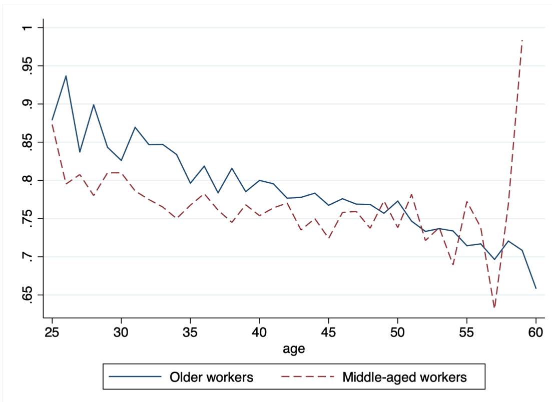

(a) Consumption to income ratio

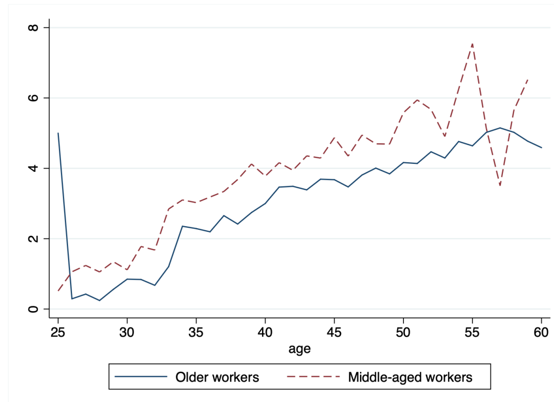

(b) wealth-to-income ratio

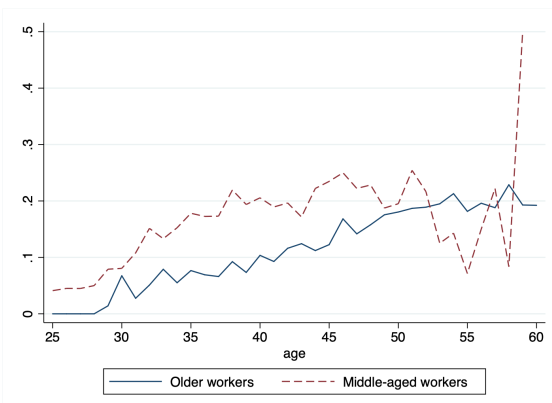

(c) Financial market participation

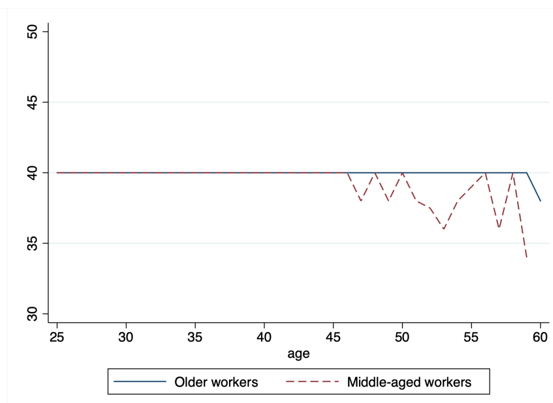

(d) Hours of work

Figure 1: Life cycle profiles for outcomes in the data by treatment status.

Notes: Panels (a), (b) and (d) plot the median consumption to income ratio, wealth-to-income ratio and hours of work, respectively, over the working life. Panel (c) plots the average financial market participation rates. Figures are computed from the SHIW data for the years 1986-2008 (values in 2010 euros), separately for older workers (at least 18 years of contribution in 1995) and middles-aged workers (less than 18 years of contribution in 1995).

The consumption to income ratio is lower for the treated than for the control households, below age 45, as shown in figure 1.a, implying higher saving rates for the group of households targeted by the reform. Similarly, figure 1.b shows a higher assets to income ratio for the middle-aged at all ages. Figure 1.c also documents the low propensity of Italian households

[6]Although it is not possible to disentangle year, age and cohort effects without additional information, the overlaps can still provide some insights about differences in life-cycle profiles of the two cohorts. However, notice that the data include few observations for the cohort of middle-aged workers above 50.

to hold risky assets, with the treated and control households showing substantial heterogeneity in their portfolio allocation choices. Indeed, the figure shows a remarkably higher participation rate before age 50 for middle-aged households: between ages 30 and 50, the average participation is as low as 10 percent on average for older workers while around 22 percent of middle-aged households have some positive share of their wealth invested in risky assets. In contrast, there seems to be no differences in the median working hours between the two groups, as shown in figure 1.d.

2.3 Reduced form evidence on the effects of the reform

Empirical strategy We inform the estimation of the structural model with the impacts of the pension reforms on households' private wealth, participation in the financial markets, hours of work and expected retirement age. These impacts are estimated here using a diff-in-diff (DID) identification strategy. Since the group of older private employees was untouched by the reform, while the groups of middle-aged private employees and public employees were targeted by the reform, older private employees are used as a the control group for the behavior of middle-aged private employees and public employees.

The DID strategy leads to the following empirical specification for the effect of the pension reform:

$$\begin{aligned} y_{it} = \delta_0 + \delta_1 POST_i + \delta_2 D_i + \delta_3 * PUB_i + \delta_3 POST_i * PUBi + \delta_4 D_i * PUB_i \\ + \delta_5 POST_i * D_i * PRIV_i + \delta_6 POST_i * D_i * PUB_i + \varepsilon_{it} \end{aligned} \tag{1}$$

where y is the relevant outcome variable (log of net wealth-to-income ratio, log hours of work, expected age of retirement and financial market participation), $POST$ is a time of intervention dummy taking value one in the period after the reform, PUB and $PRIV$ are sector of employment dummies taking value one if the household head is employed in the public or private sector, respectively, and D a treatment dummy, taking value one if the household head had more than 18 years of contributions in 1995.[7] The coefficients δ_5 and δ_6 associated to the interaction of time dummy $POST$, treatment dummy D and sector of activity dummy, $PRIV$ or PUB, represent the DID parameters of interest, capturing the variation in y induced by the reform to the group of middle-aged private and public employees under the DID assumptions. Notice that the departure from such assumptions (linearity and absence of pre-treatment trends) biases the estimation of the pension reform impacts, but is not an issue in our indirect inference approach.

[7] We define as pre-treatment period the years before 1992 ($POST = 0$), when the first intervention was made, and as post-treatment period the years after 1995, $POST = 1$.

We estimate equation (1) for the log of net wealth-to-income ratio, log hours of work, expected age of retirement using OLS, while we use a probit model for participation in the financial market.[8] We also include in the regressions, cohort dummies and various demographic variables to control for permanent differences in consumption and asset accumulation behavior induced by differences in earning profiles or preferences. Moreover, to capture macroeconomic shocks, we also allow for year dummies.

Reduced form results The diff-in-diff results are reported in Table 1 and seem to confirm previous findings in the literature on the effects of these pension reforms.

Table 1: Reduced form effects of the pension reforms, Diff-in-diff estimates

	(1) Log Net Wealth to income ratio	(2) Participation in financial market	(3) Log Hours of work	(4) Expected Age of Retirement
Private employees, middle-age,	0.188**	0.058**	0.006	0.597**
after the reform	(0.091)	(0.024)	(0.009)	(0.278)
Public employees, middle-age,	0.341***	0.059**	0.017	0.596*
after the reform	(0.091)	(0.028)	(0.014)	(0.352)
Controls	Yes	Yes	Yes	Yes
Cohort dummies	Yes	Yes	Yes	Yes
Time dummies	Yes	Yes	Yes	Yes
Observations	14,943	15,461	15,427	13,333
R-squared	0.105	0.111	0.112	0.105

Notes: Column 1 reports OLS estimates for log wealth-to-income ratio, column 2 marginal effects from a Probit model for financial market participation, column 3 OLS estimates for the log of total hours of work of both spouses, column 4 OLS for the expected age of retirement. The estimates also control for age, gender and education of the household head and household size. The data are obtained from the SHIW 1986-2008 waves. We drop the transitional years 1993-1995 as well as households whose heads are older than 60 years or out of the labor force. We keep married couples if the household head was born between 1930 and 1970. Standard errors for the estimated coefficients are reported in parenthesis, clustered at the household level. Three stars, two stars and one star indicate statistical significance at the 1 percent, 5 percent and at the 10 percent confidence level, respectively.

The results show that middle-aged private and public employees increase the private wealth-to-income ratio by around 21 percent and 30 percent, respectively. The larger response of public employees is consistent with the larger reduction of pension wealth for this group compared to private employees. Moreover, the reforms induced an increase in participation in the financial market for both middle-aged public and private employees, with

[8] We apply the baseline sample selection described in Section 2.2. Here, we also drop households whose heads are older than 60.

the latter increasing participation more. We find no effect on hours of work for both the middle-aged private and public employees, which suggests that Italian households do not vary hours of work to insure against shocks to pension wealth. However, consistently with previous evidence in Bottazzi et al. (2006), we show that these pension reforms induced instead a substantial revision in the age of expected retirement for the middle-aged workers. Therefore, the extensive margin of labor supply at retirement seems to represent an important mechanism used by Italian households to insure against shocks to pension wealth.

3 The life-cycle model

Structural models help to describe the mechanisms governing households decisions, which, in our context, is crucial to gain insights into alternative changes to the pension legislation. Even when the DID assumptions are satisfied, the reduced-form estimates are not enough to pin down such mechanisms, but can be used to inform a structural model. In what follows, we discuss how.

To match the reduced form effects, we model how households' decisions - before and after the reform - change according to their treatment status, which in turn depends on sector of employment and years of contribution when the reform has been introduced, as discussed in Section 2.2. Before and after the reform is relevant for how pension benefits are computed (earnings before, contributions after) and for whether the retirement age decision is available. To assign the treatment status, households are then grouped based on their year of birth (1940-1945; 1945-1950; 1950-1955; 1955-1960; 1960-1965; 1965-1970) and sector of employment (private and public sector). Households belonging to different cohorts have different years of contribution when the reform is introduced (determining the treatment status). Further, the year of birth matters for the life-cycle timing of the reform. When the new regime phases-in, older workers face a shorter time horizon, for a given retirement age, to adjust their asset accumulation and portfolio allocation, compared to younger households.

For each group of households defined by cohort-sector of employment, we solve the model in the pre- and post-reform regimes. We then simulate their counterfactual life-cycle behavior depending on their treatment status under the pension reform.[9] The model is then estimated through matching the simulated (structural-form) reform impacts to the reduced-form reform impacts.

[9]The simulated data are obtained to replicate the cohort composition of actual data. Appendix C provides details.

3.1 Model setup and assumptions

Our stochastic life-cycle model lends from Deaton (1991), Carroll (1997) Attanasio et al. (1999) and Gourinchas and Parker (2002). Households take decisions regarding consumption, the allocation of wealth between risky and riskless assets and age of retirement. As for labor supply, our modeling choice is driven by the data showing no effect of pension reforms on the intensive margin. Therefore, we focus on the extensive margin of labor supply at retirement. Because Italy's health-care system features universal coverage, we do not include uncertain medical expenditure, in contrast to the model of retirement behavior developed by French and Jones (2011).[10]

Households face uncertainty with respect to financial assets and human capital returns as well as length of life. Further modeling choices are driven by the institutional setting. Namely, the model is specified to replicate the salient features of the pension regimes faced by different cohorts before and after the pension legislation change, as we detail below. The time unit is a year.

Preferences The utility function is intertemporally separable. The period utility function is:

$$u(C_t, R; z_t) = q(z_t)\frac{\widetilde{C}_t^{1-\gamma}}{1-\gamma}e^{\phi_1(1-R)} - \phi_2(1-R)$$

where C_t is consumption, $q(z_t)$ is a function of demographic shifters to account for the evolution of households composition over the life-cycle, $\widetilde{C}_t = \frac{C_t}{q(z_t)}$, $R = \mathbb{1}\{t \geq N\}$ and N are the years of contribution at retirement. Following Attanasio et al. (2008), the parameters are constrained to deliver a not homothetic period utility function, with participation reducing utility directly and indirectly, through reducing the utility of consumption. Therefore, the coefficient of relative risk aversion γ is greater than one, ϕ_1 and ϕ_2 are greater than zero.

Bequests When households die at age t, the remaining wealth, A_t, is left to their heirs. As in De Nardi (2004), households value bequests according to the bequest function $b(A_t) = \theta\frac{(A_t+k)^{1-\gamma}}{1-\gamma}$, where θ is the intensity of the bequest motive and k the parameter controlling the curvature of the bequest function.

Length of life Households live at most until age T, but can die before. Therefore, length of life is uncertain. To model the uncertainty of the length of life, we denote as d_t, the probability that the household is alive in period $t+1$, conditional on being alive in period t.

[10] In this sense, we follow Haan and Prowse (2014) who also exclude medical expenses in a retirement model for Germany.

Financial assets returns Households allocate their wealth, A_t, between a riskless B_t and a risky S_t asset. The riskless and the risky assets earn returns equal to r_f and $r_f + \mu_S + \eta_{t+1}$, respectively, where the excess return from risky assets μ_S is greater than zero and η_{t+1} are independently and identically distributed according to $\mathcal{N}(0, \sigma_S^2)$.[11] As in Fagereng et al. (2017), we also allow for tail risk in the risky assets return distribution: the return in the tail event is r_{tail} and the probability is p_{tail}.

We assume costly collection and processing of the financial information needed to access the return from the risky asset. Since the access decision is made on a period basis, households pay a per-period fixed cost, ψ, to hold the risky assets.[12] Moreover, we assume borrowing (non-negative share of the riskless asset) and short-sale (non-negative share of risky asset) constraints: the share of risky assets, $\omega_t = S_t/A_t$, lies between zero and one. The return from a household's portfolio can then be written as:

$$r^p_{t+1} = r_f + \omega_t(\mu_S + \eta_{t+1}) \tag{2}$$

Earnings During the working life, households receive gross labor earnings Y_t:

$$Y_{t+1} = g_t Y_t v_{t+1} \tag{3}$$

where v_t are permanent i.i.d. shocks to earnings with constant variances, and g_t is the age-varying earnings growth factor.[13] We interpret v_t as productivity shocks, e.g., shocks to the value and price of a worker's skills. Both shocks variances and growth rates are then allowed to vary with the workers' sector of employment (private and public sector). In each period they work, individuals make social security contributions equal to $\frac{\tau}{3}Y_t$, where τ is the sum of employees' and employers' contribution rate.

Pension wealth and benefits In the pre-reform regime, pension benefits, PB, are computed according to an earnings model:

$$PB = \rho N H_N$$

[11] We impose zero correlation between shocks to risky returns and labor income. While complicating the analysis a non-zero correlation makes little difference in the portfolio rule and thus on the households' simulated portfolio allocations, see Gomes and Michaelides (2005).

[12] For a discussion on the various interpretations of the fixed cost of participation see Jappelli and Padula (2015).

[13] This income process corresponds to a standard permanent-transitory type earnings process setting the variance of the transitory error component to zero. We follow Scholz et al. (2006), among others, and choose to set the variance of the transitory shocks to zero as it mostly reflects measurement error in earnings.

where ρ is the accrual rate (which varies depending on the individual's sector of employment), N are years of contribution and H_N is a measure of average earnings at retirement. We approximate the evolution of average earnings using the following dynamic equation:

$$H_{t+1} = (1 - R)\left(h_1 H_t + h_2 Y_{t+1}\right) + R H_N \tag{4}$$

where h_1 and h_2 depend on the individual's sector of employment.[14]

In the post-reform regime, individuals retire under a pro-rata model, that is a combination of earnings and contributions models. Pension benefits are given by:

$$PB = \rho N_{1995} H_N + \Gamma_N$$

where N_{1995} is now the number of years of contribution in 1995 and Γ_N is the contributions model component of pension benefits, defined as:

$$\Gamma_N = \alpha_N \Xi_N$$

where Ξ_N is the amount of defined contribution wealth accumulated by the household at retirement and α_N is the so-called transformation coefficient, increasing with age of retirement. Each period they work, individuals contribute a non-contingent share, $\frac{\tau}{3}$, of their labor earnings to the defined contribution account, and receive an employer defined contribution equal to $\frac{2\tau}{3}$.[15] Defined contribution wealth earns each year a return factor, $\overline{G}_t$, that equals the five-years moving average of the GPD growth factor. During working life, defined contribution wealth thus evolves according to:

$$\Xi_{t+1} = (1 - R)\left(\overline{G}_t \Xi_t + \tau Y_{t+1}\right) + R \Xi_N \tag{5}$$

When the retirement decision becomes available for the working household (between ages 56 and 64), we can write the evolution of defined contribution benefits as:

$$\Gamma_{t+1} = \left(\frac{\overline{G}_t \Gamma_t}{\alpha_t} + \tau Y_{t+1}\right) \alpha_{t+1} \tag{6}$$

[14] In the pre-reform regime, H_N corresponds to the last year of earnings for individuals employed in the public sector ($h_1 = 0$ and $h_2 = 1$) and to the average of the last five years of earnings for individuals employed in the private sector ($h_1 = 0.8$ and $h_2 = 0.2$).

[15] These contribution rules reproduce those in the actual institutional setting. The Italian system identifies a list of contingencies, including unemployment, job mobility and disability, during which the State contributes on the individual defined contribution account in lieu of workers and employers. For this reason, we assume a working household to contribute to the defined contribution account in each period (up to age 56) in the post-reform regime.

3.2 The households' problem

Households choose consumption, the portfolio share of risky assets and retirement age to maximize:

$$E_0 \left\{ \sum_{t=0}^{T} \beta^t \left[d_t u(C_t, R; z_t) + (1-d_t) b(A_t) \right] \right\}$$

where $\beta < 1$ is the subjective discount factor. Before retirement, the dynamic budget constraint reads as:

$$A_{t+1} = (1 + r^p_{t+1}) \left[A_t + \left(1 - \frac{\tau}{3}\right) Y_{t+1} - C_t - \psi \times \mathbb{1}(\omega_t > 0) \right] \tag{7}$$

and after retirement as:

$$A_{t+1} = (1 + r^p_{t+1}) \left[A_t + PB - C_t - \psi \times \mathbb{1}(\omega_t > 0) \right] \tag{8}$$

In solving the model we assume that households retire at a fixed age in the pre-refom regime and decide the age of retirement in the $57 - 65$ window in the post-reform regime.[16] Moreover, in contrast to French (2005), where retirement is modeled as a participation choice and then households can reenter the labor market, we explicitly model the retirement choice as an absorbing state.[17]

State variables The dynamic optimization problem of the household is characterized by the state variables: age (t), assets (A), labour earnings (Y), average earnings (H), defined contribution wealth (Ξ) and pension benefits from defined contribution wealth (Γ).[18] Defined contribution wealth (Ξ) and benefits (Γ) are state variables only in the post-reform regime. We denote the set of state variables as $X_t = \{t, A, Y, H, \Xi, \Gamma\}$.

We provide below the recursive formulation of the households optimization problem. In each period the household consumes, chooses the portfolio composition and (between age 57 and 65) the extensive margin of labor supply, given the wealth available at the beginning of the period, the level of permanent income, average earnings, defined contribution wealth

[16] In the pre-refom regime, workers are indeed observed to retire as soon as they fulfill eligibility requirements, see Brugiavini (1999).

[17] Both the pre and post-reform pension legislation put limitations on the possibility to work after retirement.

[18] Similarly to O'Dea (2018), Ξ wealth is a state variable only up to age 56 (when the retirement decision becomes available), with defined contribution benefits Γ being a state variable only after age 56. The transition from defined contribution wealth to defined contribution benefits at age 56 is governed by the equation:

$$\Gamma_{57} = \left(\overline{G}_{56} \Xi_{56} + \tau Y_{57}\right) \alpha_{57}$$

and benefits. We present the household's problem in the periods after retirement and before retirement when the decision to retire is (between ages 57 and 65) or is not available.

The problem after retirement In both the pre and post-reform setting, the after retirement Bellman equation is:

$$V_t^R(X_t) = \max_{\{C_t,\omega_t\}} \left\{ u(C_t, 1; z_t) + \beta E_t \left[d_{t+1} V_{t+1}^R(X_{t+1}) + (1 - d_{t+1}) b(A_{t+1}) \right] \right\}$$

subject to (8).

The working household problem when the retirement decision is not available Before the reform (and after the reform for households younger than 57), the recursive formulation of the household problem is:

$$V_t(X_t) = \max_{\{C_t,\omega_t\}} \left\{ u(C_t, 0; z_t) + \beta E_t \left[d_{t+1} V_{t+1}(X_{t+1}) + (1 - d_{t+1}) b(A_{t+1}) \right] \right\}$$

subject to (3), (4), (5) and (7).

The working household problem when the retirement decision is available After the reform, households can decide when to retire between ages 57 and 65. In each $t \in [56, 64]$, the working household must then solve two problems, corresponding to the decision to retire ($R = 1$) or keep working ($R = 0$):

$$\begin{aligned} v_t(X_t, R = 1) &= \max_{\{C_t,\omega_t\}} \left\{ u(C_t, 1; z_t) + \beta E_t \left[d_{t+1} V_{t+1}^R(X_{t+1}) + (1 - d_{t+1}) b(A_{t+1}) \right] \right\} \\ v_t(X_t, R = 0) &= \max_{\{C_t,\omega_t\}} \left\{ u(C_t, 0; z_t) + \beta E_t \left[d_{t+1} V_{t+1}(X_{t+1}) + (1 - d_{t+1}) b(A_{t+1}) \right] \right\} \end{aligned}$$

subject to (3), (4), (6), (7) and (8), and where $v_t(X_t, R)$ is the *retirement choice-specific* value function.[19] The decision problem of the working household whether to retire at time t can then be expressed recursively as:

$$V_t(X_t) = \max \left\{ v_t(X_t, R = 0), v_t(X_t, R = 1) \right\}$$

The problem cannot be solved analytically and we derive the policy rules numerically by backward induction. The solution algorithm combines continuous and discrete choices based

[19] As noted in Iskhakov et al. (2017), because "retirement" is both a state and a choice in this setting, key is the distinction between state-specific (V_t and V_t^R) and choice-specific ($v_t(X_t, R = 0)$ and $v_t(X_t, R = 1)$) value functions (and policy rules).

on modifications of the algorithms in Iskhakov et al. (2017) and Druedahl and Jørgensen (2017). For each level of the discretized state space for permanent income, average earnings, defined contribution wealth and benefits, we then employ the Endogenous Grid Method proposed by Carroll (2006) to derive the optimal consumption function on an *exogenous* grid for cash-on-hand, and then compare the corresponding value-of-choice to compute the discrete retirement and portfolio choices.[20] Details on the solution algorithms are reported in Appendix B.

4 Estimation and results

The estimation is in two-steps, as in Gourinchas and Parker (2002). The first step focuses on the parameters that are estimated outside the structural model presented in Section 3: the earnings process, the risky asset returns distribution, the survival probabilities, the taste shifters parameters, the curvature of the bequest function, and the parameters characterising the pension rules.

Preferences (β, γ, θ, ϕ_1, ϕ_2), fixed-costs (ψ) and the probability of disastrous event (p_{tail}) parameters are estimated in the second step resting on the Gourieroux et al. (1993) indirect inference approach and using equation (1) as an auxiliary model. In our context, the approach implies minimizing the distance between the DID estimates on actual and simulated data and some additional target moments included to capture the age-profile of wealth and participation in the financial market during the working life in the pre-treatment period. Therefore, identification exploits the exogenous variation induced by the pension reforms as well as pre-treatment information.

The indirect inference approach delivers consistent estimates of the structural model parameters, independently of the credibility of the DID estimates as causal effects of the pension reforms. Hence, the threats to the validity of the DID assumptions are not an issue for the estimation of the structural model parameters and the analysis of the life-cycle effects of pension reforms.

4.1 Parameters estimated outside the model

Earnings process The age-profile of the earnings is a key determinant of asset accumulation over the life cycle (and the replacement rates at retirement). In addition, the extent of asset accumulation depends on the degree of earnings uncertainty. Therefore, obtaining

[20] Our solution algorithm is especially close to the *nested endogenous grid method* in Druedahl (2019).

consistent estimates of the earnings process parameters is crucial to credibly attempt to replicate on simulated data the DID estimates obtained with actual data.

To estimate the earnings process parameters we use data from the Survey on Household Income and Wealth for the period 1986-2008. The age-profile of earnings is estimated separately for each sector of employment using a third-order polynomial. The estimation controls for household size, time and cohort fixed effects.

Table 7 in Appendix A, panel A, reports the results. Figure 9 plots the estimated age-profiles of earnings, which show that private employees face a steeper age-profile of earnings than public employees.

The estimation (and identification) of the variances of permanent shocks to earnings focuses on the second order moments of the first differenced data and to use a GMM estimator, as in Blundell et al. (2016c). We estimate the variances of permanent shocks separately for each sector of employment from the residuals obtained projecting (log)earnings onto a third order polynomial in age, dummies for household size, time and cohort dummies. The estimation of the shocks variances uses the restrictions imposed by (3).

Table 7, panel B, shows the results. The variability of permanent shocks is greater for private employees than for public employees (the variance of the permanent component is estimated equal to around 0.0152 and 0.0109 for private and public employees, respectively), perhaps reflecting the higher exposure of private employees' marketable skills to technological shocks. Overall, our estimates are in line with previous evidence in the literature. In particular, the estimates for the permanent component variance fall within the confidence intervals of the estimates of Jappelli and Pistaferri (2010), obtained using SHIW data for the period 1989-2006 and the same definition of earnings.

Pension rules Next, we turn to the parametrization of the pension rules. The accrual rate under the earnings model, ρ, is set to 0.02 and 0.023 for private and public employees, respectively, in the pre-reform regime, and equal to 0.02 for both groups in the post-reform regime. In the earnings model, the relevant years of earnings for the computation of pension benefits are the last 5 years and the last year before the reform for private and public employees, respectively, and the last 10 years after the reform for both private and public employees. For this reason, in the pre-reform regime we set h_1 and h_2 equal to 0.8 and 0.2, respectively, for private employees and equal to 0 and 1, respectively, for public employees. In the post-reform regime, h_1 and h_2 are set equal to 0.9 and 0.1, respectively, for both groups of workers. We set the number of years of contribution in year 1995 in the pro-rata model (N_{1995}) equal to 10, which is the average value for middle-aged workers computed from the SHIW data.

The contribution rate, τ, is set to 0.33, its value in the Italian pension system. Table 8 in Appendix A reports the normal retirement age for older workers before and after the reform (which also applies to middle-aged workers in the pre-reform regime).[21] The table also reports the retirement age-dependent transformation coefficients, α, applying to middle-aged workers in the post-reform regime. Furthermore, we set $\overline{G}_t$ to 1.015, the annual real GDP growth factor hypothesized by the legislator at the time of the reform.

Other parameters Here, we focus on the other exogenous parameters used in the model. These are collected in Table 8 in Appendix A.

Changes in household composition affect the utility from consumption through the equivalence scale $q(z_t)$. To compute $q(z_t)$ at each age, we use the SHIW data for the period 1986-2008. As in Scholz et al. (2006) we set:

$$q(z_t) = \left(\frac{z_t^j + 0.7 z_t^k}{2.7} \right)^{\frac{3}{4}}$$

where z_t^j and z_t^k the age t mean number of adults and kids, respectively.

The parameter that determines the curvature of the bequest function, k, is set at euros 300,000.[22]

Households face mortality risk and we take the conditional survival probabilities of being alive at age $t+1$, given being alive at age t, from the National Institute of Statistics (Istat). The conditional survival probability at age 91 is set to zero, which implies a terminal age T equal to 90.

The model starts at age 25 for both private and public employees. To capture potential heterogeneity in the age at which households enter the labor force (together with other possible selection mechanisms), we allow for heterogeneity in the initial distributions of wealth and earnings, as well as for different contribution histories at retirement.[23]

[21] Minimum retirement age was progressively increased after the reform at 65 years of age for older workers, so that they could actually still enjoy an earlier retirement.

[22] This value approximately corresponds to the point estimate obtained structurally in De Nardi et al. (2010) using the version of their model without medical expenses (equal to 273,000 in 1998 dollars, converted to 2010 prices and to euros using the average exchange rate in that year). This value is consistent with earlier evidence (Guiso and Jappelli, 1991) that intergenerational transfers in Italy occur only above relatively high levels of the wealth distribution.

[23] Allowing for different contribution histories at retirement also accounts for contributions made by older workers that entered the labor force before the age of 25. We set N equal to 37 for older workers employed in both public and private sector, that is the average number of contribution years made by this cohort of workers in the SHIW data. Further, notice that the youngest cohort of middle-aged workers facing the introduction of the reform that we consider (born between 1965-1970) faces its introduction at 28 years of age.

Regarding the financial market returns, we average the Federal Reserve Bank of St. Louis data on "3-Month Interbank Rates" and (log changes in the) "Total Share Prices for All Shares" for Italy (after adjusting for inflation), for the period 1979-2008, to set r_f and μ_S to 0.0302 and 0.0496, respectively.[24] Using the same data, we find a value for the standard deviation of risky returns σ_S equal to 0.262. As in Fagereng et al. (2017), we set the return in the tail event, r_{tail}, to resemble the Italian stock market index returns during large crashes. In the period 1975-2010, the Italian stock market index experienced five large crashes, with losses between 45% and 70%, which implies a return in the tail event equal to -0.50, a value close to the average loss experienced by households investing in the stock market index during these crashes: the Italian stock market index lost around 54% in 1982, 45% in 1988, 45% in 1992, 47% in 1991, and 70% between March 2008 and March 2009.[25]

4.2 Structural parameters estimation

Estimation of the structural parameters exploits the exogenous variation induced by the pension reforms. To do this, we construct moments conditional on the treatment status and allow for heterogeneous policy variation between households cohorts. We explicitly target the DID estimates from the actual data and estimate equation (1) on the simulated data as an auxiliary model. As discussed above, the approach delivers consistent estimates of the structural parameter and of the reform impacts, even when DID estimates on actual data fail to do so.

The effects of a pension reform in a life-cycle framework predicted by the economic model also depend on the level of wealth accumulated by the households (and their portfolio allocation) prior to the introduction of the policy. Therefore, to lend further credibility to our validation exercise, we also target moments describing the evolution of wealth and participation over the working life in the pre-treatment period.

We simulate the behavior of $10,000$ households. For each sector of employment, we simulate the behavior of six year-of-birth cohorts of households (1940-1945; 1945-1950; 1950-1955; 1955-1960; 1960-1965; 1965-1970). Each household is assigned an initial level of income and wealth-to-income ratio as drawn from the empirical distributions of income and wealth (conditional on sector of employment) in the SHIW data for individuals aged 24-28.[26]

We take random draws from the cohort-specific earnings process, the risky asset returns and mortality distributions. These are then used, together with the cohort-specific policy functions (before and after the pension reform) and pension award formula, to simulate the

[24] The FRED short term rates time series for Italy is available starting from 1979.

[25] Again following Fagereng et al. (2017), we estimate p_{tail} within the structural exercise.

[26] We follow here French (2005).

behavior over the life-cycle, accounting for the pension policy change in 1995.[27] Notice that the age of the policy change varies between cohorts. Then, we pool the simulated series for a number of households in each year of birth-education-sector of employment cohort corresponding to that observed in the SHIW 1986-2008 data. Finally, applying the additional sample selection described in Section 2.2, we obtain a simulated sample that mimics the composition of the actual sample used to estimate (1) in Section 2.2. We construct the remaining variables (treatment status, pre- or post-treatment period and their interactions with the sector of employment) relying on year of birth (the median year in the year-of-birth range) and sector of employment information.

We use two sets of moment conditions. The first set collects the DID estimates from the regressions of the wealth-to-income ratio and of the financial market participation (marginal effects) for private and public employees.[28] The second describes the pre-treatment behavior of households: we target median wealth-to-income ratios and average participation rates for households in the age groups 25-35, 35-45 and 45-55, separately for private and public employees, as well as the (unconditional) median wealth-to-income ratio between 65 and 75 years of age. Specifically, we run median regressions of the wealth-to-income ratio on a third order polynomial of the household head's age, dummies for household size, cohort and year dummies, separately for private and public employees, and then take the predicted conditional median wealth-to-income ratio by age group.[29]

The model is overidentified since we estimate 7 parameters (β, γ, θ, ϕ_1, ϕ_2, ψ and p_{tail}) to match 17 moments in the data.[30] We minimize the weighted distance between the target moments in the simulated and actual data using a simulated annealing algorithm (Kirpatrick et al., 1983). Following the suggestion in Pischke (1995), we use the inverse of the diagonal of the bootstrapped variance-covariance matrix of the moments as a weighting matrix.[31] Details on the estimation procedure are reported in Appendix C.

[27] Table 9 reports average simulated replacement rates for selected sub-groups of households.

[28] We then estimate equation (1) on simulated data. We estimate all models with OLS on both actual and simulated data.

[29] A similar approach is adopted to estimate the age profile of stock market participation, employing a linear probability model.

[30] Since the domain of the parameters ϕ_1 and θ corresponding to a given level of utility cost of work and marginal propensity to bequeath, respectively, depends on the values of β and γ, to increase the efficiency of the estimation algorithm we estimate the utility cost of work $\widetilde{\phi_1}$ and the marginal propensity to bequeath $\widetilde{\theta}$ instead. For a given set of parameter values, we back out ϕ_1 and θ using $\phi_1 = \log\left(\beta(1+r_f)(1-\widetilde{\phi_1})^{-\gamma}\right)$ and $\theta = \frac{1}{\beta(1+r_f)}\left((\frac{1}{\widetilde{\theta}}-1)\frac{1}{1+r_f}\right)^{-\gamma}$.

[31] The approach addresses the small sample bias issues from the adoption of the optimal weighting matrix shown by Altonji and Segal (1996).

Identification The parameters are jointly identified and several sources of data variability contribute to identification, with a changing degree between parameters. Identification of the discount factor β mostly relies on the variation of the amount of private wealth accumulated over the life cycle.

The DID estimates for participation in the financial market help identify, γ, the degree of relative risk aversion, as higher concavity in the value function when working induces larger optimal responses of participation, for a given disutility from work. The observed profile of financial market participation over the working life also contribute to identify γ: a higher γ induces a larger share of older households to optimally exit the market and is then consistent with a flatter profile of participation. Human capital risk differs between private and public employees, which, through the intensity of the precautionary motive for saving, provides a further source of variation to identify γ.

The participation rates among the young households identify the fixed cost of stock market participation ψ (for a given degree of relative risk aversion).

For a given degree of relative risk aversion and participation costs, the probability of an extreme event p_{tail} is pinned down by the overall level of participation in the financial market.

The identification of the intensity of the bequest motive θ exploits the trajectory of private wealth late in life: the smaller the decrease in private wealth after retirement and the steeper the wealth accumulation before retirement, the higher the value of leaving a bequest.

As for the utility cost of work parameters, ϕ_1 is identified by the age profile of assets during the working life, with flatter profiles indicating, *ceteris paribus*, that households are optimally decreasing more strongly consumption at the time of retirement. For a given degree of substitutability between consumption and leisure, ϕ_2 is identified by the DID estimates of the wealth-to-income ratio (and participation, for a given degree of relative risk aversion). Given that the reforms raise the incentive to postpone retirement, the larger increase of the wealth-to-income ratio, the larger the fixed cost of working.

Finally, a crucial identifying assumption is that the pension reforms did not impact the preference parameters and the cost of financial market participation.

4.3 Estimation results

Table 2 reports the estimation results. The estimate of the time discount factor, β, is close to 0.99, which amply falls in the ballpark of previous estimates for dynamic life cycle models (see French, 2005).[32] We estimate a coefficient of relative risk aversion equal to around 1.62, a value close to those estimated by Attanasio and Weber (1995) and by Gourinchas and Parker (2002).[33] Our estimate for the marginal propensity to bequeath (0.88) corresponds to a value for the intensity of the bequest motive, θ, of 26, and is broadly in the range of those estimated in the literature. Specifically, it is close to that estimated by De Nardi et al. (2010) (0.88) and smaller than that estimated by Lockwood (2018) (0.96).

Table 2: Estimated structural parameters

Parameter		Value	Std. error
Time discount factor	β	0.9915	(0.0002)
Coefficient of relative risk aversion	γ	1.6168	(0.0085)
Cost of financial market participation	ψ	805.24	(0.8996)
Tail event probability	p_{tail}	0.0201	(0.086e-3)
Utility cost of work	$\widetilde{\phi}_1$	0.0502	(0.0019)
	ϕ_2	0.0005	(0.074e-3)
Marginal propensity to bequeath	$\tilde{\theta}$	0.8810	(0.0021)

Notes: The estimates are obtained using an indirect inference approach. The simulated annealing algorithm is used to minimize the distance between moments of actual and simulated data. The cost of financial market participation is expressed in 2010 euros. Asymptotic standard errors are reported in parentheses.

The estimated per-period cost of stock market participation is around 800 in 2010 euro. This value falls in the middle of the range of estimates for the median per-period cost of participation (650-930) obtained by Vissing-Jorgensen (2004).[34] For the median income earners, this estimate implies a cost between around 4.8% (for younger households) to around 2.7% (for older households) of the annual net household income, thus in the lower part of the range of estimates (4-6%) obtained structurally by Khorunzhina (2013). Our estimate for the tail event probability (2%) falls within the range of values (0.6% to 3.2%) that Fagereng et al. (2017) argue to be consistent with the stock market crashes history between 1920 and 2010.

As in Attanasio et al. (2008), the ψ_2 parameter reflects the utility cost of deciding to not retire and work one year longer. Our estimate for this parameter is 0.0005. We find a significant degree of substitutability between consumption and leisure in utility. Specifically,

[32] The estimate does not include mortality risk. In fact, as in French (2005), and described in Section 3, we explicitly introduce mortality risk in the model.

[33] Higher values have been employed in the calibration of life-cycle models when trying to match financial market participation rates over the life-cycle, as in Fagereng et al. (2017).

[34] We obtain this range converting the median estimates of \$350 (in 1983 dollars) for 1994 and \$500 for 1989 in 2010 euro.

we estimate a utility cost of working equivalent to around 5 percent of consumption, which compares to the 7.3 percent equivalent cost of working calibrated by Attanasio et al. (2008).[35]

Sensitivity To support the arguments on the identification of the model, we analyse the sensitivity of the parameter estimates to changes in the target moments using the measure proposed by Andrews et al. (2017). The sensitivity measure, reported in Figure 12 in Appendix C, confirms the main intuition for identification outlined above, and shows in particular the importance of the DID estimates for the effects of the reform to pin down the parameter values of the structural model.

Goodness of fit Table 3 reports the value of the auxiliary moments/parameters in the simulated data, in the SHIW data and the 95% confidence interval for the difference between data and simulations.[36]

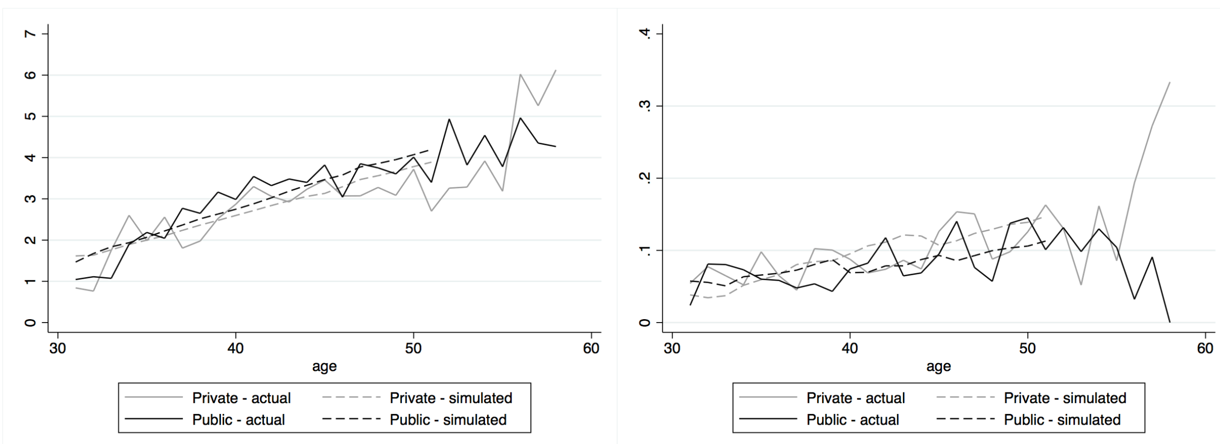

(a) wealth-to-income ratio

(b) Financial markets participation

Figure 2: Pre-treatment statistics: simulated and actual data

Notes: The straight lines plot the wealth-to-income ratio (panel a) and financial market participation rate (panel b) by age of the household head as computed from the SHIW data for the years 1986-1993 (values in 2010 euros), separately for private and public employees. The dashed lines plot the corresponding figures in the data simulated by the model (under the pre-reform pension regime).

The pre-reform age profile of both the wealth-to-income ratio and the financial market participation in the simulated data is close to that in the actual data. The model also mimics satisfactorily the heterogeneity between sectors of employment. Remarkably, for all age

[35]The estimate for the utility cost of working corresponds to a value of ϕ_1 of around 0.10.

[36]Details on the computation of the confidence intervals are reported in Appendix C. We follow Low and Pistaferri (2015) with this approach.

groups, the theoretical moments fall within the 95% confidence interval of the corresponding empirical moments. Figure 2 helps to visualize the comparison between simulated and actual data and confirms the ability of the model to replicate observed pre-treatment moments.

Table 3: Goodness of fit

Target moments						
Pre-treatment statistics	Sector	Age group	Model	Data	[95% CI	Diff.]
wealth-to-income ratio	Private	25-35	1.67*	2.08	-0.46	1.28
		35-45	2.81*	2.87	-0.14	0.26
		45-55	3.56*	3.45	-0.66	0.45
	Public	25-35	1.72*	1.68	-0.88	0.80
		35-45	2.99*	3.08	-0.11	0.30
		45-55	3.87*	4.19	-0.36	1.00
	All	65-75	6.70*	6.60	-0.34	0.13
Financial markets part.	Private	25-35	0.042*	0.067	-0.032	0.082
		35-45	0.100*	0.092	-0.024	0.009
		45-55	0.129*	0.120	-0.061	0.042
	Public	25-35	0.056*	0.056	-0.059	0.058
		35-45	0.082*	0.088	-0.010	0.022
		45-55	0.098*	0.102	-0.048	0.056
DID estimates	Sector					
(Log) wealth	Private		0.218*	0.199	-0.217	0.178
	Public		0.344*	0.352	-0.166	0.182
Participation	Private		0.047*	0.050	-0.036	0.044
(*Marginal effects*)	Public		0.044*	0.047	-0.040	0.045

Notes: The wealth-to-income ratio refers to the age-group median, the financial markets participation to the age-groups fraction of households holding risky assets in their portfolios. *indicates simulated moment falls within the 95% confidence interval of the empirical moment.

The model does also a good job in replicating the DID estimates obtained from estimating the equation (1) for the wealth-to-income ratio and the financial market participation. All the DID estimates (for private and public employees) obtained using the simulated data are close to the corresponding empirical estimates and fall within their 95% confidence interval.

5 Implications

5.1 Displacement effect

We can use simulated wealth and pension benefits generated by the model to investigate the offset between discretionary and social security wealth.[37] For each individual i in the simulated sample, we construct the effect of the reform on private wealth ($\Delta \log A_{i,t}$) and social security wealth ($\Delta \log PB_{i,t}$) at age t by taking the difference between the simulated *actual* behavior in the presence of the reform ($\log A_{i,t}^A$ and $\log PB_{i,t}^A$ for the log of private wealth and pension benefits, respectively) and the *counterfactual* behavior in the absence of the reform ($\log A_{i,t}^C$ and $\log PB_{i,t}^C$ for the log of private wealth and pension benefits, respectively).[38] To obtain an estimate of the model-implied substitutability between private and social security wealth, we estimate the following equation on simulated data:

$$\Delta \log A_{i,t} = \delta_0^A + \delta_1^A \Delta \log PB_{i,t} + \epsilon_{i,t} \tag{9}$$

where $\widehat{\delta_1^A}$ indicates the response of accumulated savings for retirement to a variation in expected pension benefits predicted by the structurally estimated economic model. To isolate the effect of changes in pension benefits on savings, we estimate eq. (9) on the simulated data for the cohorts of older workers, for whom the decision of when to retire is not available. Further, we focus on simulated discretionary wealth at the end of working life ($t = 60$).

The results, reported in Column (1) of Table 4, show that a 10 percent decrease in benefit generosity induces an average increase in discretionary wealth at retirement of around 6.5 percent. Clearly, this value should not be interpreted as an estimate for the offset between private savings and social security wealth because the policy change has been introduced during the working life of the cohorts of older workers (at ages 43, 48 and 53 for the cohort of workers born in the years 1950-1955, 1945-1950 and 1940-1945, respectively). As expected, results in Columns (2) to (4) of Table 4 show that the response of private wealth to changes in pension benefit generosity decreases with the age at which the cohort of workers faces the introduction of the policy change.

[37] Most of the previous studies (e.g., Attanasio and Brugiavini 2003, Bottazzi et al. 2006 and, more recently, Lachowska and Myck 2018) investigate the offset between discretionary and social security wealth using information on current private wealth and estimates of future pension wealth (together with the exogenous variation from pension reforms to instrument the latter). The advantage of using the data simulated by the economic model is that this allows to overcome the main threats to the validity of the empirical analysis from the presence of measurement error in wealth (and income) and unobserved heterogeneity affecting both discretionary and social security wealth.

[38] Notice that changes in pension benefits ($\Delta \log PB_{i,t}$) equal changes in social security wealth, provided that pension benefits are constant after retirement and assuming that the pension reform does not affect the individual survival probabilities.

Table 4: Response of discretionary wealth to changes in pension wealth

		Reform faced by workers at age			
	Older workers	53	48	43	25
	(1)	(2)	(3)	(4)	(5)
ΔPB	-0.653***	-0.527***	-0.676***	-0.738***	-0.840***
	(0.013)	(0.0123)	(0.019)	(0.020)	(0.032)
Constant	0.0298***	0.0173***	0.0310***	0.0409***	0.0532***
	(0.0012)	(0.0017)	(0.0026)	(0.0028)	(0.0036)
Observations	4929	1615	1590	1724	1273

Notes: The Table reports OLS estimates for eq.(9) on simulated data. Column 1 reports estimation results for the cohort of older workers. Columns 2,3 and 4 report estimation results for the cohorts of workers born in the years 1940-1945, 1945-1950 and 1950-1955, respectively. Column 5 reports the estimation results for households that face the introduction of the reform at the beginning of the working life. Three stars indicate statistical significance at the 1 percent, two stars at the 5 and one start at the 10 percent confidence level.

To rule out the bias coming from observing households with varying planning horizons (Gale, 1998) and obtain the offset between private wealth and social security wealth predicted by the model, we use the estimated model to simulate the introduction of the reform at the beginning of the working life (at age 25). We then run equation (9) using the simulated $A^A_{i,60}$. As shown in Column (5) of Table 4, the estimated model predicts an offset between discretionary wealth and social security wealth of around -0.84. This value falls in the ballpark of the estimates in the literature (see, e.g., Bottazzi et al. 2006 and Alessie et al. 2013), and suggests that private savings and pension wealth are not perfect substitutes even in a model with perfectly informed and financially sophisticated individuals.[39]

The results in Table 4 show the extent to which households find it optimal not to substitute, in the short-run, social security wealth with discretionary wealth, compared to what they would do were they facing the post-reform pension rules from the beginning of their working life. Figure 3 shows the percentage change in discretionary wealth before retirement (at age 60, "*long run*") and 5 years after the introduction of the reform ("*short run*"), for each decile of variation in expected pension benefits induced by the pension reform. While the short-run response obviously understates the total effect on the accumulation of savings for retirement, the figure shows the extent with which the bias increases with the (negative) variation in expected pension benefits.

[39] Several factors, including uncertainty and different implicit returns from social security wealth and financial wealth, may result in less than a full offset (Gale, 1998, Attanasio and Brugiavini, 2003).

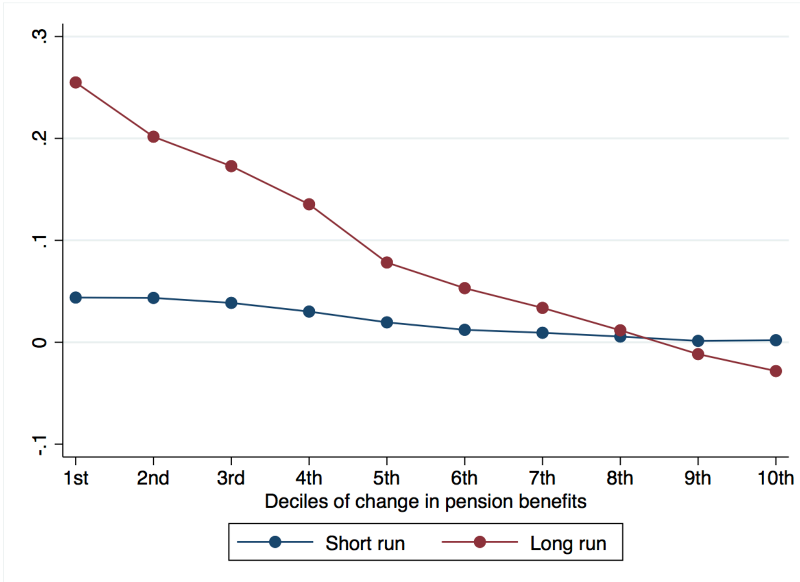

Figure 3: Short-run vs. long run response of private wealth to changes in pension wealth

Notes: Each point corresponds to the model-predicted response of private wealth to changes in pension wealth, in each decile of reform-induced variation in pension benefits. The response is expressed as percentage change in simulated discretionary wealth between post- and pre-reform regimes. The term "short-run" refers to the average response 5 years after the introduction of the reform, while "long-run" refers to the response at age 60.

5.2 The response of labor supply at retirement

In the post-reform regime, households can access higher pension benefits by retiring later. What is the role of labor supply at retirement in explaining the observed responses of savings and portfolio allocation? We start by considering the retirement behavior of the cohort of middle-aged workers predicted by the estimated model.[40]

The model suggests households revised expectations following the introduction of the pension reforms, with households retiring later under the contributive pension scheme than they did under the pre-reform regime. In particular, as shown in Table 5, the average retirement age implied by the estimated model is around 62.9 and 63.1 years for middle-aged private and public employees, respectively. Although we cannot observe the retirement choices of the cohort of middle-aged workers,[41] the model's prediction is consistent with the

[40]The simulated responses of discretionary wealth and financial markets participation to the introduction of the pension policy change depend on the expected age of retirement under the new contributive pension scheme with flexible retirement age. Because we match the effects of the reform, we can simulate the endogenous retirement behavior of these cohorts of workers predicted by the model that is consistent with the matched behavioral response in terms of asset accumulation and stock market participation.

[41]Middle-aged workers will become eligible to claim benefits for early retirement starting from around 2030.

increase in the *expected* age of retirement of the middle-aged workers in the SHIW data in the post treatment period. However, the simulated optimal age of retirement consistent with households responses of savings and portfolio allocation compares to an average expected retirement age of 62.5 years for middle-aged workers in the SHIW data, suggesting some expectation error in terms of future retirement age in the data.[42]

Table 5: Retirement in the model and expected retirement in the data

	Actual (*Expected*)	Simulated
Private	62.56	62.91
Public	62.43	63.07
All	62.51	62.97

Notes: Comparison between mean expected retirement age in the SHIW data and simulated by the economic model for middle-aged workers.

Hence the model shows that the response of households to the pension reforms in terms of savings for retirement and stock market participation is consistent with an increase in the labor supply at retirement. But what is the importance of labor supply at retirement against shocks to pension wealth? To answer this question, we exploit the simulated variation in pension benefits at retirement (for a given age of retirement) across middle-aged workers induced by the pension reform and the simulated effect on retirement age.[43] We estimate the following simple equation using simulated data:

$$\Delta \log RET_i = \delta_0^R + \delta_1^R \Delta \log PB_i + \epsilon_i$$

where $\Delta \log RET = \log RET^A - \log RET^C$ is the log change in the age of retirement (in years), at the individual level, in the presence and in the absence of the policy change. $\widehat{\delta_1^R}$ represents the simulated response of retirement to changes in the generosity of the social security system. We find a value of $\widehat{\delta_1^R} = -0.086$. The model predicts then households to

[42]Another interpretation of this result (under the alternative assumption of no retirement expectation error) is that part of the households in the data do not respond optimally to the pension reform, increasing their saving response less than it would have been optimal under their expectations of retirement age. Figure 10 in the Online Appendix reports a comparison of the distribution of expected retirement in the SHIW data and that predicted by the model.

[43]The simulated effect on retirement age is constructed as the difference between the simulated post-treatment retirement age and the retirement age in the pre-reform regime. Here, we consider the variation in pension benefits induced by the reform when individuals retire at 60 years of age (the retirement age in the pre-treatment period) as a measure of change in benefit generosity.

increase the retirement age by around 0.9% following a 10% decrease in the pension benefits they would receive for a given age of retirement. This is one of the first quantification of the importance of the extensive margin of labor supply at retirement as an insurance mechanism against (negative) shocks to benefit generosity.[44]

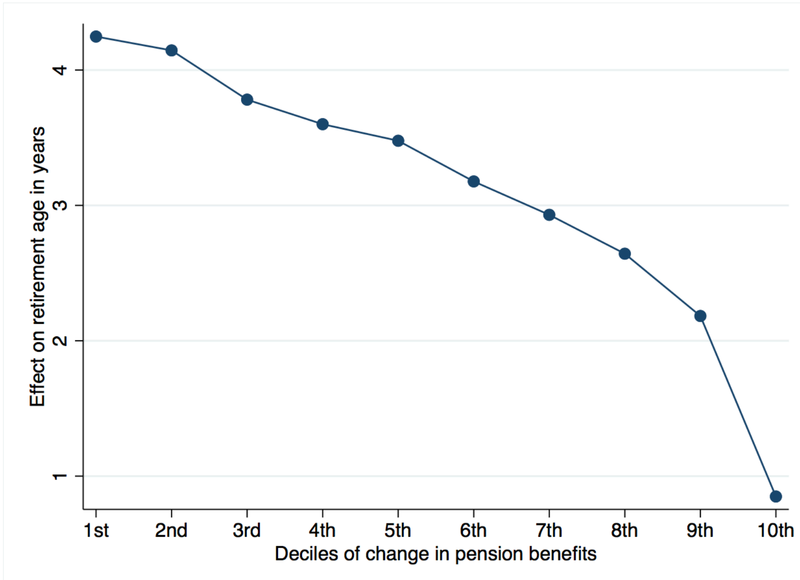

Figure 4: The extent of insurance through labor supply at retirement

Notes: Each point corresponds to the model-predicted response of retirement to changes in pension wealth, in each decile of reform-induced variation in pension benefits. The response is expressed as the difference between the simulated retirement age under the post-reform notional defined contribution scheme and that under the pre-reform defined benefit regime.

Figure 4 plots the increase in retirement age in years by deciles of variation in pension benefits induced by the pension reform, showing the extent of the labor supply responses to a reduction in expected pension benefits (for a given retirement age) predicted by the estimated model. The model predicts individuals to work up to three years longer to offset the reduction in benefit generosity.

[44]Few studies have documented the impact of benefit generosity on (the extensive margin of) labor supply. Krueger and Pischke (1992) and Snyder and Evans (2006) study the effect of benefit generosity on employment exploiting the lower benefits of the cohort of US individuals born in the period 1917-1922. While the first study finds no evidence of benefit generosity on employment, the latter finds considerable labor supply responses that might understate the true effects from increasing benefit generosity because of limited awareness about the policy change among the affected cohorts. More recently, Manoli and Weber (2016) provide nonparametric evidence of substantial retirement decisions response to financial incentives using data from Austria.

5.3 Life-cycle effects

We are particularly interested to study whether households facing the pension reform at different ages react differently to the same variation in expected pension wealth. The later in the working life households face a reduction in expected pension wealth the smaller their ability might be to close the gap between the optimal amount of discretionary wealth they accumulated in the pre-reform regime and that they would have accumulated had they faced the post-refom regime since the beginning of their working lives.

The reform introduced in 1995 hit the different cohorts of households in different moments of their life cycle. To study the potential distributional welfare effects across cohorts, we use a welfare metric similar to Low et al. (2010). We compute the counterfactual lifetime utility from t_{1995}, the (cohort-specific) age at which the pension legislation occurs, to T in two settings, had the change actually occurred ($m = A$) or not ($m = C$), using the corresponding simulated decision profiles:

$$E_{t_{1995}} U_i^m = \sum_{t=t_{1995}}^{T} \beta^{t-t_{1995}} u(C_{m;i,t}, R_{m;i,t}; z_{i,t})$$

Following Low et al. (2010), we define ζ_i as the fraction of consumption needed to make an individual indifferent between the $m = A$ and $m = C$ settings:

$$E_{t_{1995}} U_i^C = \sum_{t=t_{1995}}^{T} \beta^{t-t_{1995}} u\left((1-\zeta_i) C_{A;i,t}, R_{A;i,t}; z_{i,t}\right)$$

from which we can derive $\zeta_i = 1 - \left(\frac{E_{t_{1995}} U_i^C}{E_{t_{1995}} U_i^A}\right)^{\frac{1}{1-\gamma}}$. We interpret ζ_i as the consumption-equivalent welfare effect of the reform.

Table 6 reports the average value of ζ, for each cohort-sector of employment group. Because of the exogenous average reduction in expected lifetime resources, the introduction of the pension reform induced a welfare loss on average. Also, because of the larger reduction in expected pension benefits for a given retirement age faced by public employees (mainly due to the more generous pre-treatment pension provisions), workers employed in this sector experienced larger welfare losses.

Most importantly, the welfare analysis shows that households experienced a larger welfare loss the closer to the (pre-treatment) retirement age they are when the policy change is introduced, conditional on treatment status. Clearly, because different cohorts of households faced heterogeneous variations in expected pension wealth due to the reform, this heterogeneous welfare losses across cohorts may be just the consequence of the specific pension policy

Table 6: Welfare effects of the pension reforms by cohort

	Cohort	Private	Public	All
Older	1940-1945	-0.0069	-0.0297	-0.0173
	1945-1950	-0.0046	-0.0233	-0.0132
	1950-1955	-0.0036	-0.0195	-0.0111
Middle-aged	1955-1960	-0.0093	-0.0212	-0.0145
	1960-1965	0.0015	-0.0091	-0.0019
	1965-1970	0.0097	0.0014	0.0077
All		-0.0024	-0.0201	-0.0097

Notes: The Table reports the average simulated consumption-equivalent welfare effect of the reform ζ, by group.

design.[45] Figure 5 reports the welfare effects by deciles of variation in pension benefits due to the pension reform, separately for older (left panel) and middle-aged workers (right panel), for each cohort of households (facing the introduction of the reform at different ages).[46] The graphs provide a representation of the "life-cycle" welfare effects of the pension reform, with households that face the introduction of the reform at older ages experiencing larger welfare losses, conditional on the same (and for any level of) reduction in pension benefits generosity. These "life-cycle" effects are substantial. Among households experiencing a large reduction in benefit generosity (in the first 3 deciles of benefit change distribution), households belonging to the 1940-45 cohort (facing the reform at 53 years of age) would be willing to pay on average around 3.8% of annual consumption to face the reform at 33 years of age instead.

In the model, this effect quantitatively depends on the degree of decreasing marginal utility of consumption. When exposed to the same reduction in expected social security wealth, older households need indeed to increase the saving rate more than younger households to achieve the same level of accumulated private wealth at retirement.[47] However, in the presence of uncertainty and risk aversion, households find it costly to adjust their consumption response as much as it would be necessary to accumulate the amount of savings for retirement that they would have accumulated had they faced the same policy rules from the beginning of their working lives. The model quantifies the extent with which older households find it optimal to decrease current consumption more than younger households

[45]Under the pro-rata model, the later in the working life the middle-aged worker was in 1995 the larger the share of pension wealth computed with the earnings model.

[46]As above, to capture the change in benefit generosity and isolate it from the effect of endogenous labor supply responses, the variation in pension benefits is computed for a given age of retirement (60).

[47]Figure 11 in Appendix shows the extent of heterogeneity in the average consumption response to the pension reform (from the time of the introduction of policy change to retirement) between households belonging to the different cohorts.

(and then the greater welfare losses) following the decrease in expected pension benefits.

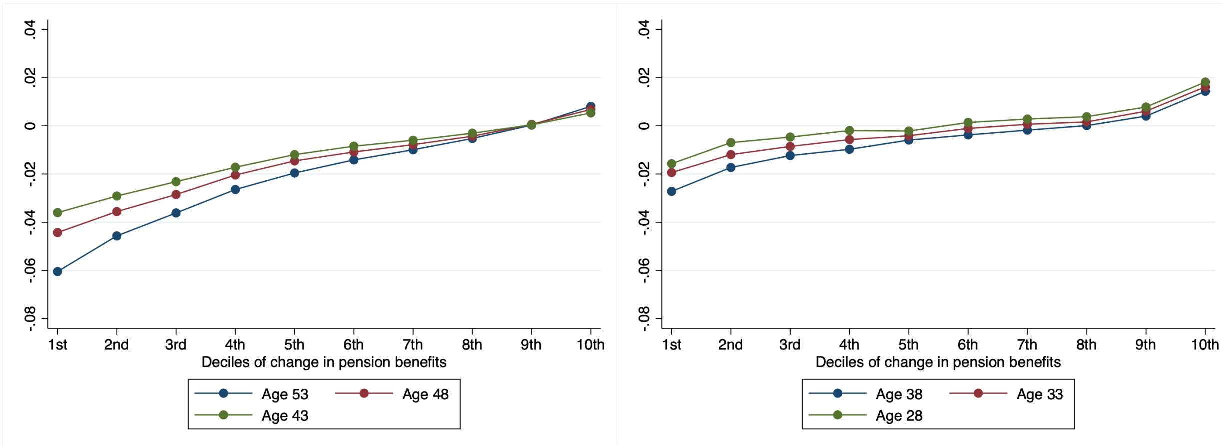

(a) Older workers

(b) Middle-aged workers

Figure 5: Welfare effects by changes in pension benefits and age at the time of reform

Notes: Each point corresponds to the model-predicted consumption-equivalent welfare effect of the reform ζ, in each decile of reform-induced variation in pension benefits. Panels (a) and (b) report the simulated welfare effects for older workers (facing a post-reform defined benefit pension scheme) and middle-aged workers (facing a post-reform notionally defined contribution regime). In each panel, the welfare effects are reported separately for different year-of-birth cohorts (facing the reform at different ages).

5.4 Early retirement age and benefit generosity

To shed further light onto the wealth effect on labor supply at retirement and the displacement effect between private and social security wealth, we use the estimated model to study the responses of households to two pension policies that a legislator may consider to increase the financial sustainability of the social security system during the demographic transition: (i) an increase in the minimum retirement age; (ii) a reduction in benefit generosity. We simulate the introduction of the policy reforms starting from the defined contribution regime in place in Italy in 2013, in which workers can choose to retire between 57 and 70 years of age with pension provisions increasing with age (for a given amount of contributions), separately for cohorts of workers aged 40 and 55 at the time of reform.[48] Specifically, the first policy change consists in an increase of the early retirement age to 62 while the second reform reduces the coefficients transforming total pension wealth contributions into annuities by 10%.[49]

[48]We conduct these policy experiments for private employees.

[49]Though transformation coefficients are typically linked to life expectancy, policy makers choose when (not) to update them following revisions in individuals life expectancy.

Panel (a) of Figure 6 plots the simulated labor supply effects of the increase in the early retirement age from 57 to 62, while panel (b) plots the predicted labor supply responses to the 10% percent decrease in benefit generosity. Similarly, Figure 7 compares the effects on the accumulation of discretionary wealth from the two reforms (panel (a) for the increase in early retirement age and panel (b) for the the decrease in benefit generosity).

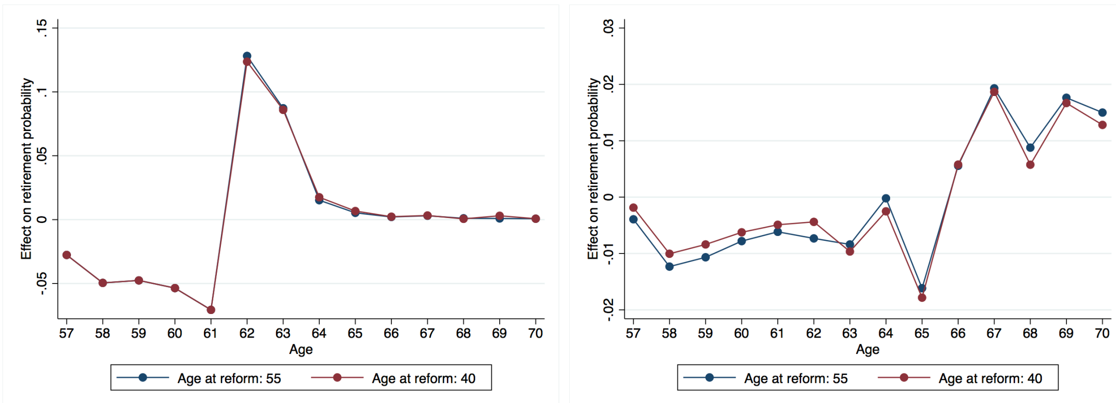

(a) Increase of minimum retirement age to 62 (b) 10 percent decrease in benefit generosity

Figure 6: Effects on the probability to retire at a given age of alternative reforms

Notes: The graphs plot the model-predicted retirement response to an increase in the minimum early retirement age (panel a) and a 10 percent decrease in pension benefit generosity (panel b). The response is expressed as the percentage point change in the probability to retire at a given age when switching from the baseline notionally defined contribution scheme to the post-intervention regime (panel a: the baseline regime when the minimum early retirement age is increased to 62; panel b: the baseline regime when benefit generosity is decreased by 10%).

The results show the substantially different implications of the two alternative policy changes. First, while increasing the early retirement age only induces a (mechanical) increase in the retirement age for those workers that would have retired before 62 in the pre-refom regime (without much differences between different cohorts of workers), decreasing the benefit generosity increases the labor supply at retirement, for any given level of pre-reform retirement age up to 70 years of age. Further, the model provides a quantitative assessment of the “life-cycle” wealth effects on labor supply, with younger cohorts of workers delaying retirement less than older workers to accumulate higher pension wealth. As shown in panel (b) of Figure 7, this is due to the fact that younger households optimally accumulate more wealth than older workers at retirement in response to the same decrease in benefit generosity.

Figure 7 also highlights the stark divergence in the response of households in terms of discretionary savings for retirement to the different pension policy changes. While decreasing benefit generosity induces households to increase savings for retirement, households’ average

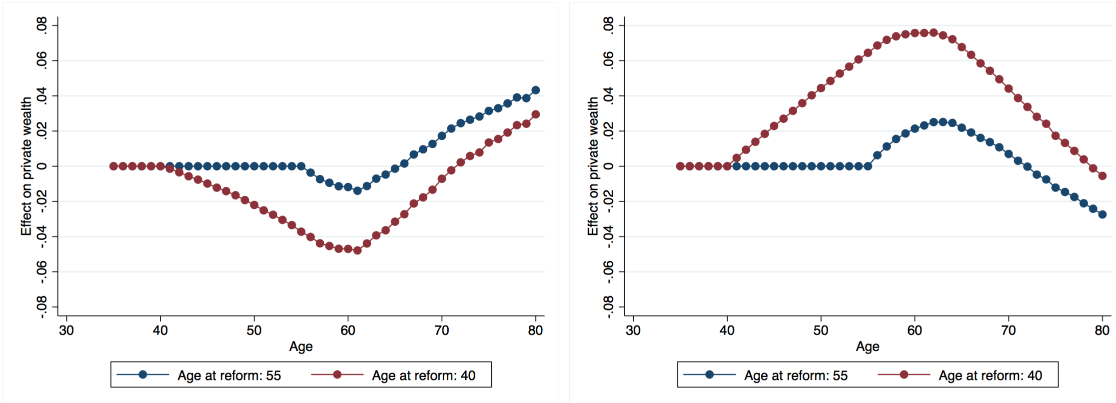

(a) Increase of minimum retirement age to 62 (b) 10 percent decrease in benefit generosity

Figure 7: Effects on the accumulation of private wealth of alternative reforms by age

Notes: The graphs plot the model-predicted effects on private wealth, by age of the household head, of an increase in the minimum early retirement age (panel a) and a 10 percent decrease in pension benefit generosity (panel b). The response is expressed as the percentage change in simulated private wealth between post-intervention (panel a: the baseline regime when the minimum early retirement age is increased to 62; panel b: the baseline regime when benefit generosity is decreased by 10%) and pre-intervention regimes (notionally defined contribution scheme, see details in Section 5.4).

response to the increase in the minimum age for early retirement depends on the age in which they face the reform. Because households that would have retired before 62 years of age are now constrained to retire later (and then access higher pension provisions once retired), those facing the reform at 40 years of age optimally choose to reduce the savings rate. The model shows then the importance of (negative) consumption wealth effects from increasing the minimum age for early retirement in the presence of benefit generosity increasing with retirement age. The model predicts the consumption wealth effect to be instead negligible for households whose head is close to retirement at the time of the policy change.

Heterogenous effects across the wealth distribution Panels (a) and (b) of Figure 8 show the effects of increasing the minimum age for early retirement and decreasing benefit generosity, respectively, on retirement, by wealth quintiles at the time of reform.

The model predicts substantial heterogeneity in the labor supply response at retirement to the increase in the early retirement age across the wealth distribution, with households in the top wealth quintile responding more (and with the heterogeneity in the response being larger when households face the policy change later in the working life). In contrast, decreasing the benefit generosity induces incentives to increase labor supply at retirement that are decreasing with the ranking of households in the wealth distribution.

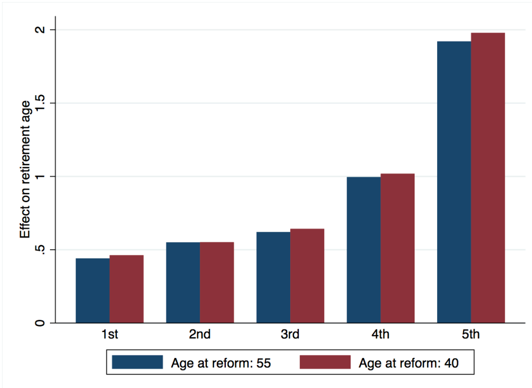

(a) Early retirement age - retirement

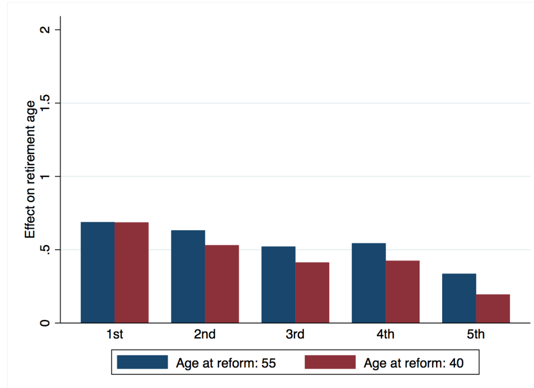

(b) Benefit generosity - retirement

(c) Early retirement age - consumption

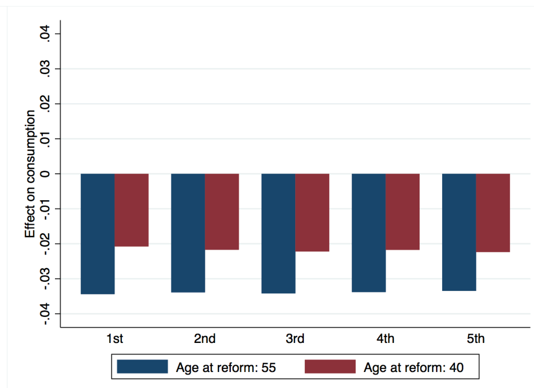

(d) Benefit generosity - consumption

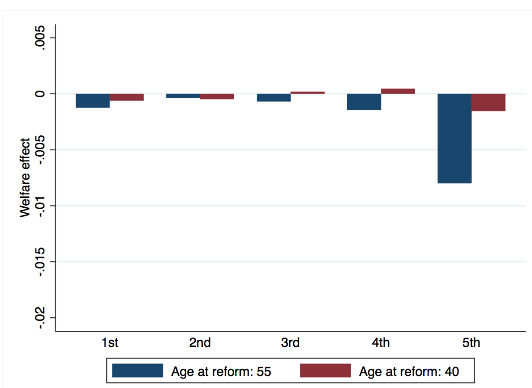

(e) Early retirement age - welfare

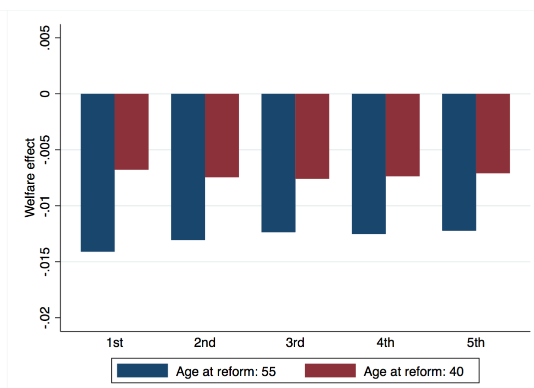

(f) Benefit generosity - welfare

Figure 8: Early retirement age and benefit generosity - effects across the wealth distribution

Notes: The graphs plot the model-predicted effects on retirement (top panels), consumption (middle panels) and welfare (bottom panels). The left (right) panels depict the effects of increasing the minimum early retirement age to 62 (decreasing pension benefit generosity by 10%). The simulated effects are reported separately for cohorts of workers facing the policy change at 40 and 55 years of age. Each bar corresponds to the simulated effect in each wealth quintile at the time of the reform.

In the model, this result reflects the fact that, prior to the policy change, households in the upper part of the wealth distribution were choosing to retire earlier due to the fixed costs of working, and the decreasing marginal gains from working with wealth.

The model also shows that the legislator's choice about the policy instrument to govern the demographic transition can have sizable consequences in terms of consumption and welfare inequality.[50] First, while increasing the early retirement age induces an increase in consumption in the upper part of the wealth distribution (panel c, Figure 8), the decrease in benefit generosity induces a rather homogeneous consumption response (panel d). Further, the model shows that the increase in consumption inequality from increasing the early retirement age is larger the later households face the introduction of the policy change.

Finally, the model predicts the two alternative pension policies to have different welfare implications. While a decrease in benefit generosity induces substantial welfare losses across the entire wealth distribution (slightly increasing the lower the initial level of wealth), as shown in panel (f) of Figure 8, increasing the minimum age for early retirement affects mainly the welfare of households at the top of the initial wealth distribution, as shown in panel (e) of Figure 8. Further, the estimated model shows that both options for pension reform entail substantial ”life-cycle” welfare effects.

6 Discussion and conclusions

This paper employs an evaluation approach to pension policies that combines ex-post and ex-ante evaluation methods. We estimate a dynamic life-cycle model of savings, portfolio choice and retirement using the reduced form effects of a wave of major pension reforms introduced in Italy in the nineties. The estimated model mimics pre-reform statistics and the effects on asset accumulation and participation rates to the financial markets estimated exploiting a difference in differences identification strategy, complementing previous evidence about the displacement effect between private and social security wealth. We also show that the estimated model predicts substantial retirement wealth effects, identifying labor supply at retirement as an important mechanism that households use to insure against shocks to social security wealth. Further, our framework suggests important life-cycle effects of the pension reforms, with older workers experiencing larger welfare losses, for any level of variation in benefit generosity. We use the estimated model to illustrate the substantially different consequences of alternative pension policies in terms of consumption and retirement wealth

[50] We take the median log variation in annual consumption (measured as the difference between consumption under the alternative option for pension reform and the counterfactual consumption in the baseline) in the years following the policy change. The welfare effects are compute as detailed in Section 5.3.

effects, as well as “life-cycle” welfare effects.

These results should be interpreted through the lenses of our assumptions on individual behavior. First, our model neglects the role of health. Clearly, there are several reasons why individuals' health may affect their retirement behavior. However, although a large literature (starting from French 2005) has studied the impact of health on retirement, the existing evidence shows that health explains only a small fraction of the variation in retirement (Blundell et al., 2016b). Importantly, because our focus is on the interplay between households decisions and the financial incentives brought about the social security system, neglecting the role of health is problematic for our findings to the extent that changes in the pension rules have an effect on health. Some recent research has shown that increasing the minimum retirement age has positive consequences on individuals' health (Bertoni et al., 2018). This finding suggests that our analysis about the consequences of raising the minimum age for early retirement may represent a lower bound for the effect on labor supply at retirement. While we believe that our findings represent first important evidence on the consequences of pension reforms for asset accumulation and retirement from a validated structural model, a promising avenue for future research is to extended this structural framework for pension policy analysis to incorporate the role of health. Second, our model assumes that all households are aware of the pension reforms and can optimally adjust their choices in response to the changing pension policy framework. As suggested by previous studies (e.g., Bottazzi et al. 2006), part of the population may have substantial expectation error in terms of future replacement rates. Moreover, there is large evidence that financial knowledge is generally low in the population (see, e.g., Lusardi and Mitchell 2014). Limited financial knowledge and awareness about the policy changes may lead to understate the true behavioral responses using reduced form approaches.[51] Hence, because we neglect the role of limited financial literacy, our structural estimation may be providing a lower bound for the cost of working with the consequence that the estimated model overstates (some) individuals' labor supply response at retirement to the pension reforms. While this argument may provide an explanation for the difference between the average simulated retirement behavior and the average *expected* retirement age in the data, financial literacy is more likely to have a role in the distributional effects of the pension reforms, with heterogeneous responses to the pension reforms across the financial literacy (and wealth) distribution. Future research should explore the implications of allowing for endogenous knowledge accumulation on the response of households to pension policy shocks.

Our findings inform models of asset accumulation and social security about the importance of labor supply at retirement in response to shocks to pension wealth, and are relevant

[51]On the role of frictions in the attenuation of estimated responses see, e.g., Chetty (2012).

to studies considering the role of financial incentives for retirement. Further, we highlight and quantify the different trade-offs that policy makers need to consider when designing future pension policies, with important implications in terms of their acceptability among the population.

References

Alessie, R., V. Angelini, and P. van Santen (2013): "Pension wealth and household savings in Europe: Evidence from SHARELIFE," *European Economic Review*, 63, 308–328.

Altonji, J. G. and L. M. Segal (1996): "Small-sample bias in GMM estimation of covariance structures," *Journal of Business & Economic Statistics*, 14, 353–366.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017): "Measuring the sensitivity of parameter estimates to estimation moments," *The Quarterly Journal of Economics*, 132, 1553–1592.

Attanasio, O. P., J. Banks, C. Meghir, and G. Weber (1999): "Humps and Bumps in Lifetime Consumption," *Journal of Business & Economic Statistics*, 17, 22–35.

Attanasio, O. P. and A. Brugiavini (2003): "Social Security And Households' Saving," *The Quarterly Journal of Economics*, 118, 1075–1119.

Attanasio, O. P., H. Low, and V. Sanchez-Marcos (2008): "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *American Economic Review*, 98, 1517–52.

Attanasio, O. P., C. Meghir, and A. Santiago (2012): "Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA," *Review of Economic Studies*, 79, 37–66.

Attanasio, O. P. and G. Weber (1995): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103, 1121–57.

Bertoni, M., G. Brunello, and G. Mazzarella (2018): "Does postponing minimum retirement age improve healthy behaviors before retirement? Evidence from middle-aged Italian workers," *Journal of health economics*, 58, 215–227.

Blau, D. M. (2008): "Retirement and consumption in a life cycle model," *Journal of Labor Economics*, 26, 35–71.

Blundell, R., M. Costa Dias, C. Meghir, and J. Shaw (2016a): "Female labor supply, human capital, and welfare reform," *Econometrica*, 84, 1705–1753.

Blundell, R., E. French, and G. Tetlow (2016b): "Retirement incentives and labor supply," in *Handbook of the economics of population aging*, Elsevier, vol. 1, 457–566.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016c): "Consumption inequality and family labor supply," *American Economic Review*, 106, 387–435.

Bottazzi, R., T. Jappelli, and M. Padula (2006): "Retirement expectations, pension reforms, and their impact on private wealth accumulation," *Journal of Public Economics*, 90, 2187–2212.

——— (2011): "The portfolio effect of pension reforms: evidence from Italy," *Journal of Pension Economics and Finance*, 10, 75–97.

Brugiavini, A. (1999): "Social Security and Retirement in Italy," in *Social Security and Retirement around the World*, National Bureau of Economic Research, Inc, NBER Chapters, 181–237.

Carroll, C. D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 112, 1–55.

——— (2006): "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics Letters*, 91, 312–320.

Chetty, R. (2012): "Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 80, 969–1018.

De Nardi, M. (2004): "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71, 743–768.

De Nardi, M., E. French, and J. B. Jones (2010): "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118, 39–75.

Deaton, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1221–48.

Druedahl, J. (2019): "A Guide On Solving Non-Convex Consumption-Saving Models," .

Druedahl, J. and T. H. Jørgensen (2017): "A general endogenous grid method for multi-dimensional models with non-convexities and constraints," *Journal of Economic Dynamics and Control*, 74, 87–107.

FAGERENG, A., C. GOTTLIEB, AND L. GUISO (2017): "Asset market participation and portfolio choice over the life-cycle," *The Journal of Finance*, 72, 705–750.

FELDSTEIN, M. S. (1974): "Social Security, Induced Retirement, and Aggregate Capital Accumulation," *Journal of Political Economy*, 82, 905–26.

FELLA, G. (2014): "A generalized endogenous grid method for non-smooth and non-concave problems," *Review of Economic Dynamics*, 17, 329–344.

FRENCH, E. (2005): "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour," *Review of Economic Studies*, 72, 395–427.

FRENCH, E. AND J. B. JONES (2011): "The effects of health insurance and self-insurance on retirement behavior," *Econometrica*, 79, 693–732.

GALE, W. G. (1998): "The effects of pensions on household wealth: A reevaluation of theory and evidence," *Journal of Political Economy*, 106, 706–723.

GOMES, F. AND A. MICHAELIDES (2005): "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *Journal of Finance*, 60, 869–904.

GOURIEROUX, C. S., A. MONTFORT, AND E. RENAULT (1993): "Indirect Inference," *Journal of Applied Econometrics*, 8, S85–S118.

GOURINCHAS, P.-O. AND J. A. PARKER (2002): "Consumption Over the Life Cycle," *Econometrica*, 70, 47–89.

GUISO, L. AND T. JAPPELLI (1991): "Intergenerational transfers and capital market imperfections: Evidence from a cross-section of Italian households," *European Economic Review*, 35, 103–120.

HAAN, P. AND V. PROWSE (2014): "Longevity, life-cycle behavior and pension reform," *Journal of Econometrics*, 178, 582–601.

HECKMAN, J. J. (2010): "Building bridges between structural and program evaluation approaches to evaluating policy," *Journal of Economic literature*, 48, 356–98.

ISKHAKOV, F., T. H. JØRGENSEN, J. RUST, AND B. SCHJERNING (2017): "The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks," *Quantitative Economics*, 8, 317–365.

JAPPELLI, T. AND M. PADULA (2015): "Investment in financial literacy, social security, and portfolio choice," *Journal of Pension Economics and Finance*, 14, 369–411.

Jappelli, T. and L. Pistaferri (2010): "Does consumption inequality track income inequality in Italy?" *Review of Economic Dynamics*, 13, 133–153.

Khorunzhina, N. (2013): "Structural estimation of stock market participation costs," *Journal of Economic Dynamics and Control*, 37, 2928–2942.

Kirpatrick, S., C. D. Gelet, and M. P. Vecchi (1983): "Optimization by Simulated Annealing," *Science*, 220, 621–630.

Krueger, A. B. and J.-S. Pischke (1992): "The effect of social security on labor supply: A cohort analysis of the notch generation," *Journal of labor economics*, 10, 412–437.

Lachowska, M. and M. Myck (2018): "The effect of public pension wealth on saving and expenditure," *American Economic Journal: Economic Policy*, 10, 284–308.

Lockwood, L. M. (2018): "Incidental bequests and the choice to self-insure late-life risks," *American Economic Review*, 108, 2513–50.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage risk and employment risk over the life cycle," *American Economic Review*, 100, 1432–67.

Low, H. and L. Pistaferri (2015): "Disability insurance and the dynamics of the incentive insurance trade-off," *American Economic Review*, 105, 2986–3029.

Lusardi, A. and O. S. Mitchell (2014): "The economic importance of financial literacy: Theory and evidence," *Journal of economic literature*, 52, 5–44.

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

Miniaci, R. and G. Weber (1999): "The Italian Recession Of 1993: Aggregate Implications Of Microeconomic Evidence," *The Review of Economics and Statistics*, 81, 237–249.

O'Dea, C. (2018): "Private pensions and public pension design," Tech. rep., Working Paper, Yale University.[1407, 1434].

Pischke, J.-S. (1995): "Measurement Error and Earnings Dynamics: Some Estimates from the PSID Validation Study," *Journal of Business Economic Statistics*, 13, 305–314.

Scholz, J. K., A. Seshadri, and S. Khitatrakun (2006): "Are Americans Saving "Optimally" for Retirement?" *Journal of Political Economy*, 114, 607–643.

Snyder, S. E. and W. N. Evans (2006): "The effect of income on mortality: evidence from the social security notch," *The review of economics and statistics*, 88, 482–495.

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

Todd, P. E. and K. I. Wolpin (2006a): "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 96, 1384–1417.

——— (2006b): "Ex Ante Evaluation of Social Programs," PIER Working Paper Archive 06-022, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.

Vissing-Jorgensen, A. (2004): "Perspectives on Behavioral Finance: Does Irrationality Disappear with Wealth? Evidence from Expectations and Actions," in *NBER Macroeconomics Annual 2003, Volume 18*, National Bureau of Economic Research, Inc, NBER Chapters, 139–208.

Appendix A Additional Figures and Tables

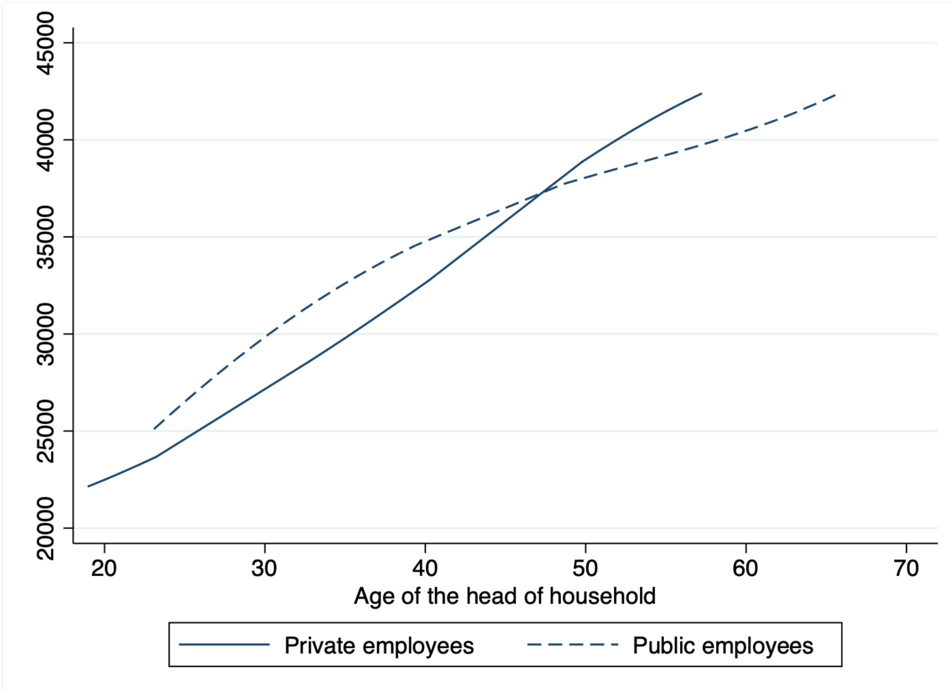

Figure 9: Estimated profiles of earnings over the working life

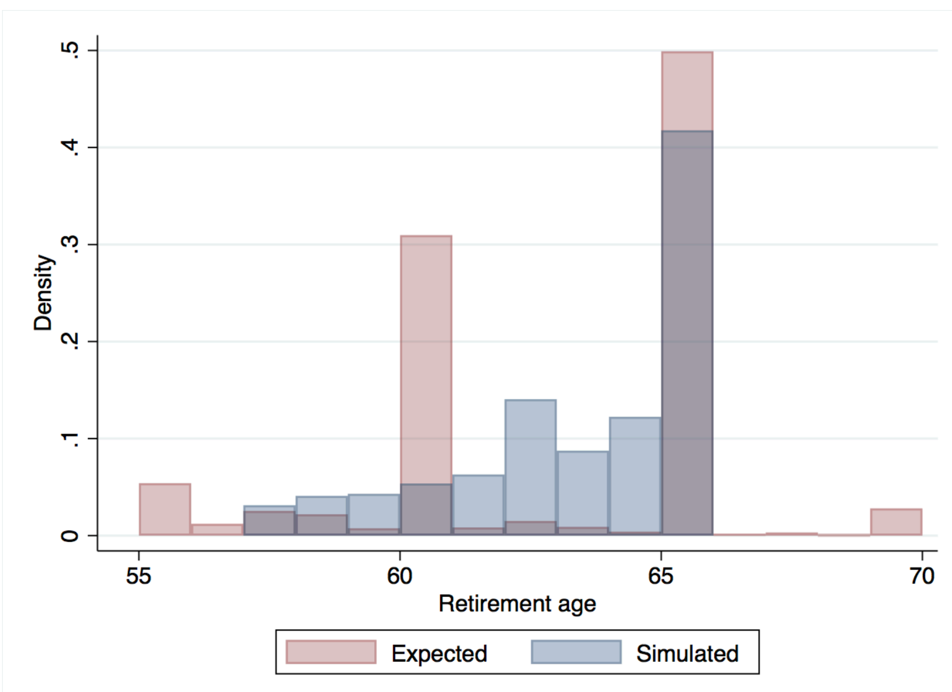

Figure 10: Expected retirement in the SHIW data and predicted by the model

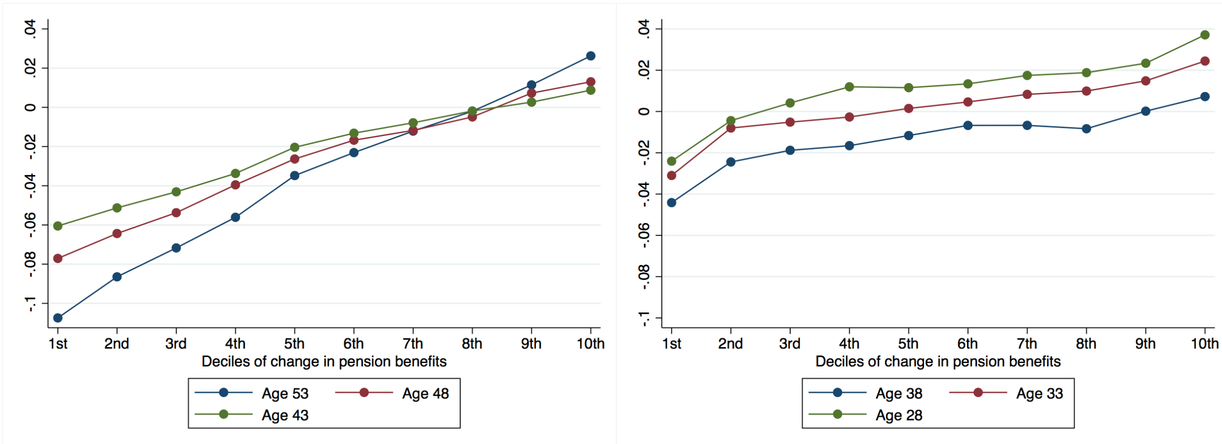

(a) Older workers (b) Middle-aged workers

Figure 11: Effects on consumption by changes in pension benefits and age at the time of reform

Table 7: Estimated parameters of the earnings process

		Private	Public
Panel A: Age coefficients			
	Age	-0.001022	0.080613**
		(0.029975)	(0.0323525)
	Age^2	0.000613	-0.0013848*
		(0.000736)	(0.000767)
	Age^3	-0.000006	0.000008
		(0.000006)	(0.000006)
		Private	Public
Panel B: Variances of shocks to earnings			
	Transitory	0.023609***	0.026754***
		(0.002828)	(0.003207)
	Permanent	0.015156***	0.010900***
		(0.002545)	(0.002632)

Notes: The estimation of the parameters of the income process uses data from SHIW for the period 1986-2008. Sample selection as described in section 2.2 is applied. The group-specific coefficients of the polynomial in age for earnings have been estimated using OLS with the inclusion of dummies for household size, cohort and year fixed effects. Earning process variances estimated using GMM. Clustered standard error in parentheses. Three stars indicate statistical significance at the 1 percent confidence level, two stars at the 5 percent level and one star at the 10 percent level.

Table 8: Other exogenous and pension parameters

Parameter			Value
Risk free rate		r_f	1.0302
Excess risky assets return		μ_S	0.0194
Std. deviation of risky assets returns		σ_S	0.2620
Return in the tail event		r_{tail}	-0.50
Starting age		t_0	25
Terminal age		T	90
Retirement age			
	Before the reform (all)		60
	After the reform (older)		61
Evolution average earnings		h_2	
	Before the reform		
	Private-employees		0.2
	Public-employees		1.0
	After the reform		
	Private-employees		0.1
	Public-employees		0.1
GDP growth rate		g	0.015
Accrual rate		ρ	
	Private-employees		0.02
	Public-employees		0.023
Contribution rate		τ	0.33
Transformation coefficient		α	
	Retirement age		
	57		0.04720
	58		0.04860
	59		0.05006
	60		0.05163
	61		0.05334
	62		0.05541
	63		0.05706
	64		0.05911
	65		0.06136

Note: r_f and μ_S are computed as described in the main text, g is the average real GDP growth rate from Istat National Account data. The after-reform retirement age apply to older workers only. For each group and pension regime, h_1 is obtained as $1 - h_2$.

Table 9: Mean simulated replacement rate before and after the pension reforms

	Pre-reform	Post-reform								
		Retirement age								
	60	57	58	59	60	61	62	63	64	65
Cohort 45-50										
Private employees	0.74					0.74				
Public employees	0.86					0.74				
Cohort 60-65										
Private employees	0.74	0.76	0.79	0.79	0.89	0.84	0.88	0.90	0.88	0.91
Public employees	0.86	0.90	0.88	0.87	0.89	0.89	0.90	0.88	0.80	0.83

Note: Replacement rates for selected cohorts of households we consider in the analysis. The replacement rate is simulated considering pension rules and parameters in Table 8 and outlined in Section 2.1. Moreover, we use growth rate of earnings for each group as estimated in Section 4.1. We assume the year of birth of each household belonging to a particular cohort to be median year in the year of birth group. Households born in 1945-1950 are older workers, and then subject to a retributive pension system. The 1960-1965 cohort are middle-aged workers in our definition and the expected pension benefit is simulated using a pro-rata model, setting the number of years of contribution in 1995 to 10.

Appendix B Solution algorithm

The dynamic programming problem that we consider has no analytical solution, and it is solved numerically by backward induction from the end of life (age 90). At each age, we compute the optimal policy rules and the value function, given the current state variables, the optimal policy rules and value function in the next period, building on the Endogenous Grid Method (EGM) developed by Carroll (2006). In his seminal work, the author introduces the method to solve a standard dynamic model for consumption and asset accumulation as in Gourinchas and Parker (2002). The complication in the solution of our model compared to the standard dynamic model for consumption à la Gourinchas and Parker arises from the inclusion of two discrete choices (retirement age and whether to participate to the financial market or not), the presence of a bequest motive and multiple state variables. Fella (2014), Blundell et al. (2016a), Iskhakov et al. (2017), Druedahl and Jørgensen (2017) and Druedahl (2019) have proposed algorithms to deal with the combination of discrete and continuous choices. As shown by Iskhakov et al. (2017), the addition of discrete choices in the standard problem of consumption introduces kinks in the value function at the points of the state space where the household is indifferent between alternatives in discrete choice space that propagate backwards through the expected marginal utility of consumption. This might imply that the uniqueness of the Euler equation is not ensured. As in Iskhakov et al. (2017) and Blundell et al. (2016a), in our problem the kinks in the value function appear at the level of assets where the agent is indifferent between alternative labor supply decisions at the extensive margin. In addition, in our model kinks might emerge at the points where agents are indifferent between participating to the financial markets, or not. To deal with these potential issues, we follow Iskhakov et al. (2017), Druedahl and Jørgensen (2017) and Druedahl (2019) and employ an upper envelop algorithm.

The solution of the model starts at the end of life (age 90). Because the probability of being alive at age 91, d_{90}, is equal to zero, the household is retired, and assuming that the household does not allocate any share of wealth to risky assets in the terminal period, the problem in (3.2) can be written as:

$$V_T(X_T) = \max_{c_T}[u(c_T) + \beta E_t b(a_T)] \tag{10}$$

Following Carroll (2006), we define an exogenous grid for end-of-period wealth $a_T = m_T - c_T$. As the form of the utility $u(C_T)$ and bequest $b(a_T)$ functions is known, we can compute the optimal consumption decision for each point of the state space X:[52]

$$c_T = [\theta(a_T + k)^{-\gamma}]^{-1/\gamma} \tag{11}$$

We can then obtain the endogenous grid for cash-on-hand as $m_T = a_T + c_T$.

Starting from $T-1$, the solution uses a combination of EGM and an upper envelope algorithm to compute policy functions and value functions until the beginning of households' working life, proceeding in two steps. The first step can be seen as an EGM conditional on a given level of permanent income, average earnings and defined contribution wealth, along the discrete choice dimensions (labor supply at retirement and allocation of wealth). In the second step, we use the conditional consumption functions obtained in step 1 and disregard non-optimal points by constructing the upper envelope associated to the different alternatives of the discrete choices on an *exogenous* grid for cash-on-hand (this avoids costly interpolations of the continuation value) as suggested by Druedahl and Jørgensen (2017) and Druedahl (2019). We finally compute the optimal policy functions and the associated value function at time t by comparing the conditional value-of-choice at time t.

We start describing the solution of the model during retirement (starting from period $T-1$), where households choose only consumption and the share of wealth to allocate to risky assets (retirement is modeled as an absorbing state) in the presence of a bequest

[52]Notice that in the last period permanent income, average earnings before retirement and defined contribution benefits are known.

motive. Consider that, after retirement, permanent income, average earnings and defined contribution benefits are known and remain constant over time. Denote with $\bar{X}$ the set of states permanent income Y, average earnings H and defined contribution benefits Γ, and with m_t^f the exogenous grid for cash-on-hand. The problem after retirement can then be written conditional on $\bar{X}$, along alternatives in the portfolio choice dimension (ω) as:

$$V_t^R(m_t^f \mid \bar{X}, \omega) = \max_{c_t}[u(c_t) + \beta(d_{t+1} E_t V_{t+1}^R(X_{t+1}) + (1 - d_{t+1}) b(a_t))] \tag{12}$$

subject to (8).

The foc for (12) with respect to c_t is

$$u'(c_t^R \mid \bar{X}, \omega) = \beta E_t[d_{t+1} R_{t+1}^p {V^{R}}'_{t+1}(X_{t+1}) + (1 - d_{t+1}) b'(a_t)] \tag{13}$$

Using the notation in Carroll (2006), we define a function:

$$\mathfrak{v}_t(a_t \mid \bar{X}, \omega) = \beta E_t \left[d_{t+1} V_{t+1}^R \left((\omega R_{t+1} + (1-\omega) R_f) a_t + PB, \bar{X}\right) + (1 - d_{t+1}) b(a_t)\right]$$

Using $u'(c_{t+1}^R) = {V^{R}}'_{t+1}(X_{t+1})$, this definition implies:[53]

$$\mathfrak{v}'_t(a_t \mid \bar{X}, \omega) = \beta E_t[R_{t+1}^p d_{t+1} u'(c_{t+1}^R(R_{t+1}^p a_t + PB, \bar{X})) + (1 - d_{t+1}) b'(a_t)]$$

From the latter we can rewrite (13) as:

$$u'(c_t^R \mid \bar{X}, \omega) = \mathfrak{v}'_t(a_t \mid \bar{X}, \omega) \tag{14}$$

The application of the EGM allows then to find the value of consumption that yields the same marginal valuation as that associated to end-of-period wealth a_t:

$$c_t^R(a_t \mid \bar{X}, \omega) = u'^{-1}\left(\beta E_t[R_{t+1}^p d_{t+1} u'(c_{t+1}^R(R_{t+1}^p a_t + PB, \bar{X})) + (1 - d_{t+1}) b'(a_t)]\right) \tag{15}$$

where $u'^{-1}(.) = q(z_t)(.)^{-1/\gamma}$ when the household is retired. The expectation of the marginal utility at time $t+1$ is computed given the information available at time t over the possible realization of returns from risky assets.

[53]The application of the Envelope Theorem with bequests tells us that:

$$V'_t(X_t) = \beta E_t[(R_{t+1}^p d_{t+1} V'_{t+1}(X_{t+1}) + (1 - d_{t+1}) b'(a_t))]$$

allowing to write:

$$u'(c_t) = V'_t(X_t)$$

Combining the consumption function $c_t(a_t \mid \bar{X}, \omega)$ with the the exogenous grid of end-of-period wealth a_t, and using the assumed dynamics of current period wealth with the presence of costs ψ to participate to the financial markets ($\omega > 0$), we can get the endogenous grid for cash-on-hand m^e:

$$\begin{aligned} m_t^e(a_t \mid \bar{X}, \omega = 0) &= c_t^R(a_t \mid \bar{X}, \omega = 0) + a_t \\ m_t^e(a_t \mid \bar{X}, \omega > 0) &= c_t^R(a_t \mid \bar{X}, \omega > 0) + a_t + \psi \end{aligned} \tag{16}$$

As discussed in Iskhakov et al. (2017), the presence of the discrete choice (portfolio choices, in this case) may imply that the value function is not globally concave. In this case, the optimality conditions for consumption would not be sufficient.[54] To solve this issue, we implement an upper envelope algorithm that disregards non-optimal consumption points similar to Druedahl and Jørgensen (2017) and Druedahl (2019). For each combination of $\bar{X}$ and discrete portfolio choices, the upper envelope algorithm returns consumption and value functions evaluated on the *exogenous* grid for cash-on-hand m_t^f: $c_t^R(m_t^f \mid \bar{X}, \omega)$ and $V_t^R(m_t^f \mid \bar{X}, \omega)$.

We can then obtain the optimal portfolio choice and the value function at time t by simply comparing the conditional value functions associated with different alternatives along the portfolio dimension:

$$V_t^R(m_t^f \mid \bar{X}) = \max_{\omega} \left(V_t^R(m_t^f \mid \bar{X}, \omega) \right)$$

Once the discrete choice is computed, it is straightforward to derive the corresponding consumption functions $c_t(m_t^f \mid \bar{X})$ for each element of the state space.

Between ages 57 and 65, households also decide whether to retire or keep working, and the problem of the household can now be written as in (3.2).

The solution of the problem of the household that decides to retire ($R = 1$) is identical to that of the retired household described above (this allows to obtain the choice-specific value function $v_t(X_t, R = 1)$ and consumption function $c_t(X_t, R = 1)$). We consider now the problem of the household that decides to work one year longer ($R = 0$). Compared to the solution of the problem of the household at retirement described above, the differences are: (i) the computation of the continuation value needs to consider the non-separability in utility between consumption and leisure; (ii) the usage of working state-specific continuation values V_t, as opposed to the retirement state-specific ones V_t^R; (iii) the additional uncertainty with respect to labor income. For each combination of permanent income Y, average earnings H,

[54] Blundell et al. (2016a) and Iskhakov et al. (2017) noticed however that uncertainty may suffice to "concavify" the expected continuation value.

defined contribution wealth Ξ and portfolio choice ω (and considering the dynamics in (3), (4), (5) and (7)), the application of the EGM allows to compute the consumption function for the worker as:

$$c_t(a_t, Y_t, H_t, \Xi_t \mid \omega, R=0) = u'^{-1}\bigg(\beta E_t[R^p_{t+1} d_{t+1} u'(c_{t+1}(X_{t+1}) + (1-d_{t+1})b'(a_t)]\bigg) \quad (17)$$

where $u'^{-1}(.) = q(z_t)(\frac{.}{\exp(\phi_1)})^{-1/\gamma}$. The expectation of the marginal utility at time $t+1$ is computed given the information available at time t over the possible realizations of permanent labor income shocks and returns from risky assets.

With the conditional consumption functions $c_t(a_t, Y_t, H_t, \Xi_t \mid \omega, R=0)$ at hand, an application of the upper envelope algorithm as in Druedahl and Jørgensen (2017) and Druedahl (2019) allows us to disregard non-optimal consumption points and obtain the consumption and value functions evaluated on the *exogenous* grid for cash-on-hand m^f_t: $c_t(m^f_t, Y_t, H_t, \Xi_t \mid, \omega, R=0)$ and $v_t(m^f_t, Y_t, H_t, \Xi_t \mid, \omega, R=0)$.

We can finally determine the portfolio choice of the worker in period t, her consumption and corresponding value function by comparing the conditional value functions associated with different alternatives along the portfolio dimension as described above for the retired household.

The comparison of the value functions for the household deciding to retire $v_t(m^f_t, Y_t, H_t, \Xi_t \mid, R=1)$ and to work at least one year longer $v_t(m^f_t, Y_t, H_t, \Xi_t \mid, R=1)$ allows to obtain to obtain the optimal retirement choice at time t.

Discretization The problem has five continuous state variables (cash-on-hand, labor income, average earnings, defined contribution wealth and defined contribution benefits) that need to be discretized. Further, we need to discretize the set of admissable values for the exogenous end-of period wealth. The set of admissible values are chosen using an exponential grid, such that there are smaller gaps between successive entries on the grid at lower levels (to better capture the concavity of the value function). We place end-of-period wealth and exogenous cash-on-hand on a grid with 50 elements. Labor income, average earnings, defined contribution wealth and defined contribution benefits are placed on grids with 10 elements.

The model has three control variables: retirement choice, portfolio choice and consumption. Consumption is obtained with the EGM and then the admissible values are not chosen ex-ante. We use a grid for the admissible values for the share of wealth allocated to risky assets with 6 elements: 0%, 7.5%, 17.2%, 37.5%, 68.7%, 100%. To perform the numerical integration, the density functions for permanent shocks to earnings and the returns from risky assets were approximated, following Tauchen (1986), using a 5-points gaussian quadrature method.

In order to evaluate the next-period consumption (and value) function associated to values of the states that do not lie on the discrete set of points of the state space, we use linear interpolation in multiple dimensions.

Appendix C Simulation and estimation

C.1 The simulated dataset with pension reforms

We use the following strategy to generate the simulated dataset that we use to match certain statistics of the actual data and the effects of the reform estimated using the diff-in-diff approach:

1. we solve the dynamic programming problem for each cohort-sector of employment group under a retributive pension system (the pre-reform pension regime), setting a group-specific exogenous retirement age, to derive policy functions for the pre-treatment period;

2. we solve the dynamic programming problem for each cohort-sector of employment group under a pro-rata model (the post-reform pension regime), introducing the choice of the age of retirement, to derive policy functions for the post-treatment period;

3. provided with the policy functions obtained in (1), we simulate the life cycle profiles for each group of households, under the assumption that the pension eligibility rules and the pension award formula remain constant over the life cycle as determined by the legislator in the pre-treatment period. Specifically, we simulate the model for a share of households in each employment group that is the same as that in the SHIW data, conditional on the year of birth. These shares are reported in Table 10. The resulting life cycle profiles represent the age profiles for outcomes that would have been observed if the pension reform was not introduced, that is the counterfactual profiles;

4. starting from the life cycle profiles obtained from (3), we simulate the introduction of the reform at cohort-specific ages corresponding to year 1995. The introduction of the reform is simulated by employing the policy functions for the post-treatment period (pro-rata model and endogenous retirement ages). The resulting profiles represent the actual life cycle profiles for the treated households.

5. Finally, we pool the simulated actual profiles for each group obtained from (3) and (4) to generate a simulated dataset that we will use to compute statistics to be matched with those estimated from the SHIW data.

Table 10: Share of each cohort of households in the SHIW data

Cohort	Private	Public
40-45	0.0996	0.0811
45-50	0.0958	0.0806
50-55	0.1018	0.0898
55-60	0.1058	0.0815
60-65	0.0947	0.0447
65-70	0.0941	0.0302

Notes: The Table reports the share of households in each cohort-sector group as computed in the SHIW data. The simulated dataset is constructed to match this composition.

C.2 Estimation

The estimation of the model is in two steps, as in Gourinchas and Parker (2002). In the first step, we set values for the exogenous parameters of the model as outlined in Section 4 of the main text. In the second step, we estimate the remaining preference and costs parameters as well as the probability of disastrous events using an indirect inference approach (Gourieroux et al., 1993).

Denote $\kappa = (\beta,\, \gamma,\, \theta,\, \psi,\, \phi_1,\, \phi_2,\, p_{tail})$ as the collection of model parameters, $\hat{\lambda^d}$ as the vector of auxiliary moments/parameters estimated in the data and $\hat{\lambda^s}(\kappa)$ as the corresponding simulated moments/parameters obtained for a given set of structural parameter values κ. We estimate κ minimizing the following indirect inference statistical criterion:

$$\hat{\kappa} = \arg\min_{\kappa} \left(\widehat{\lambda^d} - \widehat{\lambda^s}(\kappa)\right)' W \left(\widehat{\lambda^d} - \widehat{\lambda^s}(\kappa)\right) \tag{18}$$

where W is the inverse of a matrix with the bootstrapped variance of each sample moment in the diagonal. We choose this weighting matrix to overcome the small sample bias issues from the adoption of the optimal weighting matrix shown by Altonji and Segal (1996). We search for the global minimizer of (18) by employing a simulated annealing algorithm (Kirpatrick et al., 1983). Given initial values and bounds for the parameter values, the algorithm chooses probabilistically where to move in the parameter space. To avoid local minima, the "temperature" of the system, which controls the neighboring function, is reduced by 0.05 percent at each iteration. We report the converged estimates after 2000 iterations.

Standard errors We compute the standard errors of the structural parameters as the square roots of the diagonal elements of:

$$var(\hat{\kappa}) = \left(1 + \frac{1}{J}\right) \frac{\partial \widehat{\lambda^s}(\kappa)}{\partial \kappa}' W \frac{\partial \widehat{\lambda^s}(\kappa)}{\partial \kappa}$$

where $(1 + \frac{1}{J})$ is the adjustment for simulation error, with J the ratio of the number of observations in the simulated dataset to the number of observation in the sample of SHIW data that we use to estimate the auxiliary moments/parameters ($J \simeq 5$). We compute the gradient $\frac{\partial \widehat{\lambda^s}(\kappa)}{\partial \kappa}$ by finite difference.

Sensitivity To provide insights into what sources of variation help the most in pinning down the model parameters, we compute the sensitivity measure proposed by Andrews et al. (2017). We then obtain the sensitivity matrix $\widehat{\Lambda}$ as:

$$\widehat{\Lambda} = \left(\widehat{G}(\hat{\kappa})'\widehat{W}\widehat{G}(\hat{\kappa})^{-1}\right) \widehat{G}(\hat{\kappa})'\widehat{W}$$

where $\widehat{G}(\hat{\kappa}) = \frac{\partial \widehat{\lambda^s}(\kappa)}{\partial \kappa}$. Since the units of the elements of λ are not naturally comparable, we multiply the absolute value of each element of $\widehat{\Lambda}$ by $\widehat{\lambda^s}(\kappa)/100$, so that the sensitivity measure can be interpreted as the effect on the model parameter of a one percent increase in the auxiliary moment. The elements of the sensitivity matrix are plotted in Figure 12.

Confidence intervals for the difference between data and simulations We construct the 95 percent confidence interval for the difference between data and simulations reported in Table 3 in the main text similarly to Low and Pistaferri (2015). We compute the standard error of the difference $(\widehat{\lambda^d} - \widehat{\lambda^s}(\hat{\kappa}))$ as: $\sqrt{\left(1 + \frac{1}{J}\right) se^2_{\widehat{\lambda^d}}}$, where $se_{\widehat{\lambda^d}}$ are the bootstrapped standard errors of the auxiliary moments estimated using the SHIW data.

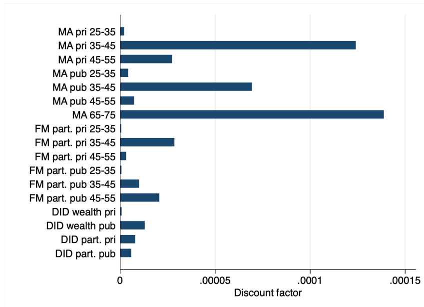

(a) β

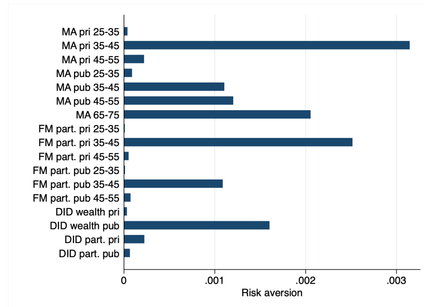

(b) γ

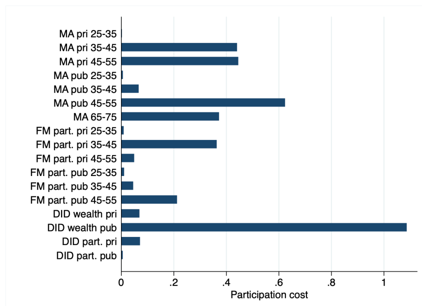

(c) ψ

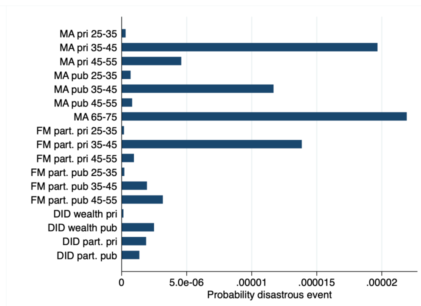

(d) p_{tail}

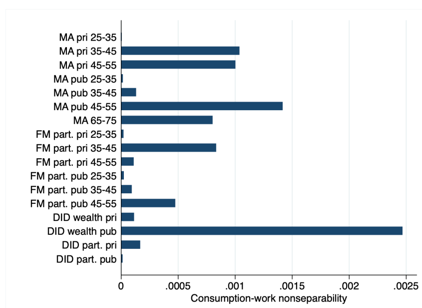

(e) Non-separability

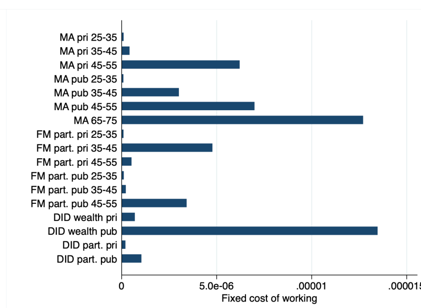

(f) ϕ_2

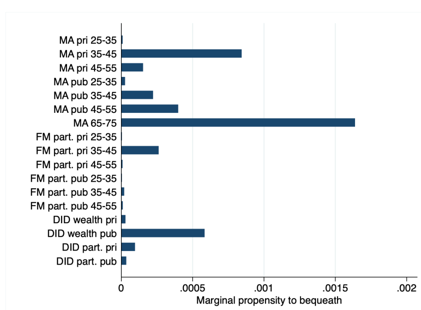

(g) Marginal propensity to bequeath

Figure 12: Sensitivity measure

Notes: This Figure reports the absolute value of the scaled sensitivity matrix as defined in Andrews et al. (2017). The sensitivity measure has been rescaled to indicate the effect of a one percent increase in the moments on the parameters.