Appendix A   Additional Institutional Details

Appendix A.1   Review of the Swedish Pension System

Appendix A.1.1   Details on Public pensions and pension reform

The public pension system in Sweden has undergone large reforms the last two decades and is in the process of going from a defined benefit (DB) system to a system based on notional defined contributions (NDC). The NDC system is expected to be fully phased-in around year 2040. Cohorts born before 1938 receive their pension benefits from the old ATP system, which is a DB scheme. Here we will review both the old and the new system. We also describe the treatment of couples and how the pension system interacts with other parts of the social insurance systems, mainly disability insurance (DI) and unemployment insurance (UI). Pension benefits in both the ATP system and the NDC system are financed by payroll taxes.

Implementing the Reform   Cohorts born between 1938 and 1953 receive their pension benefits from both the DB and the NDC schemes, with the weight on the NDC scheme increasing gradually over time. Cohorts born in 1954 onwards will receive all pension benefits from the NDC scheme. The cohorts at or near retirement age during the period spanned by our consumption data are those for whom the ATP system was the main determinant of benefits and the NDC was just beginning to be phased in. Individuals born in 1938 receive 80% of their pension benefits from the ATP system and 20% from the NDC system. Each cohort then gets another 5-percentage point from the NDC scheme. For example, individuals born in 1939 get 25% of their pension benefits from the NDC system while the 1953 cohort gets 95% of their pension benefits from the NDC system.

The ATP system.   ATP stands for *Allmän tilläggspension* in Swedish, which means "General supplementary pension." The word "supplementary" refers to the fact that there is also a minimum basic pension, which we call simply the minimum pension in the main text. We refer to combined public old-age pension benefits system prior to the reform as the ATP system, which is common terminology.

The ATP system is a DB scheme. Pension benefits are based on 1) the 15 years in an individual's career where pensionable earnings were the highest, 2) the total number of years in which an individual earns pension rights (career length, with a maximum of 30 years), and 3) the claiming age. Pensionable earnings are labor income and income from social insurance benefits that in turn are based on labor income, such as unemployment insurance, sickness insurance, parental leave benefits, workers' compensation and disability insurance. Capital income is not considered to be pensionable earnings nor are transfers that are not based on previous labor earnings, for instance social aid.

Pension rights can be earned between ages 16 and 64 - earnings at age 65 or beyond have no effect on pension rights. Annual earnings are converted to pension rights by dividing earnings in a year by a *base amount* (BA) for that year, which produces the *ATP points* used to calculate pension benefits. The BA serves to index pension rights and benefits to prices, with some discretion by the government.[38] Annual ATP points are capped at 6.5 BAs, which corresponds

---

[38]The BA is used to calculate benefits throughout the Swedish social insurance system. It is set each year by the

empirically to the median of the earnings distribution for 55 year olds in 2000.

For a worker claiming their public pension at age 65, the annual ATP pension benefit received by an individual i in year t is given by the following formula:

$$(12) \qquad b_{it} = 0.6 \cdot AP_i \cdot \min\left(\frac{N_i}{30}, 1\right) \cdot BA_t,$$

where 0.6 is the replacement rate for a worker with 30 years of contribution, AP_i are the average number of ATP pension points accrued by the individual during the highest earning 15 years, N_i are the number of contributing years and BA is the base amount in year t. The highest attainable pension benefit from the ATP system in year t is $0.6 \cdot 6.5 \cdot BA_t$.

The normal retirement age in the ATP system is 65, but pension benefits can be claimed from age 61. Claiming early reduces pension benefits by 0.5 percentage points for each month of early withdrawal relative to the month an individual turn 65. For example, individuals who claim pension benefits a year before turning 65 get their pension benefits reduced by $12 \cdot 0.5 = 6$ percentage points. Individuals who claim after 65 receive an extra 0.7 percentage point increase in pension benefits for every additional month that claiming is postponed. There is no earnings test whereby working while claiming reduces benefits, though the progressivity of the income tax schedule disincentivizes working while claiming to some degree.

For individuals with short careers or low lifetime labor earnings there is a basic pension which serves as a floor for pension benefits. The basic pension is a function of the BA and the number of years the individual has resided in Sweden. Thirty years of residence is required for full basic pension. Married individuals receive lower basic pension benefits than singles.[39] Our data shows that a quarter of all 66 year olds received basic pension in 2007.

The new NDC system In the NDC system, income-related pension benefits are calculated as the sum of wage-indexed lifetime pensionable earnings and the sum is divided by life expectancy. Unlike with the ATP, there is no upper age limit for accumulation of pension rights: as long as an individual works, their pension rights grow. The income base amount replaces the old base amount (BA) and is indexed to average wage growth instead of prices.[40] Pensionable earnings are capped at 7.5 income base amounts. Pensions in the NDC system can be claimed from age 61. However, retiring and claiming pensions earlier means that a smaller sum of pensionable earnings is divided by longer life expectancy. This decreases the net present value of the individual's pension and results in smaller pension benefits.

Just as in the ATP system there is a minimum pension for individuals with short careers and low accumulated pensionable earnings, which is now called the *guaranteed pension*. To avoid

---

Swedish government and tracks the CPI closely. However, the government can make discretionary decisions not to raise the BA or raise it more or less than the annual inflation rate. The BA also defines the minimum earnings governing whether the individual earns any ATP pension rights in a year, which was 1 BA. The BA for 2000 was 36.600 kronor or 18% of the median labor earnings among 55 year olds (see Appendix Figure A-1).

[39]Formally, the basic pension for singles is calculated as $1.529 \cdot \min\left(H_i/30, 1\right) \cdot BA_t$ where H is the number of residential years in Sweden. For married 1.529 is replaced by 1.349 which means that the basic pension is $1.529/1.349 - 1 \approx 12\%$ lower for married pensioners.

[40]The income base amount is determined by the Swedish government, just like the BA.

confusion, we continue to call the guaranteed pension the minimum pension for NDC in the main text. The guaranteed pension in the NDC system is a function of the *enhanced base amount*. This amount tracks the CPI, like the BA, but is slightly larger. Retirees with income-related pension benefits below 2.13 base amounts for singles and 1.93 enhanced base amounts for couples, receive the guaranteed pension (see Appendix Figure A-2). About 30% of all individuals receiving pension benefits are expected to receive basic pensions in 2040 when the NDC system is phased in.[41]

Treatment of singles and couples The Swedish pension system is highly individualized. Household composition is mainly used when minimum pensions are determined. As mentioned above, married individuals receive lower minimum pensions in the NDC system and in the ATP system. The minimum pension benefit in both systems is about 10% lower for married individuals, relative to singles.

The Swedish pension system also contains a survivor's benefit, which is paid out for a year after one's spouse has passed. Both widows and widowers are eligible to this benefit. Before 1990 the survivor's benefit was considerably more generous and was paid out for the rest of the survivor's life, but, unlike the current survivor's benefit, only widows and not widowers were eligible. Women who had married before 1989 and had a joint child with their husband born before December 31, 1989 and women who had been married since 1984 receive a survivor's pension based on the passed husband's ATP pension. Otherwise, widows aged below 65 and widows born before 1930 receive 40% of the husband's ATP pension while widows born 1930 and later and who are 65 years or older receive a lower survivor's pension which depends negatively on the widow's own pension and her year of birth. These more generous survivor's benefits are still paid out for those fulfilling the listed requirements above.

Interaction with other social insurance programs Social insurance benefits that are based on previous labor income counts as pensionable income in both the ATP and the NDC system. Individuals who are unemployed, receive sickness benefits or disability insurance also collect pension rights. Individuals can receive social insurance benefits until they become 65 years old.

Before 2003, disability insurance (DI) was integrated with the pension system. DI benefits were calculated as ATP pension benefits but with actual earnings being replaced by an assumed earnings profile in the calculation of pension rights (Jönsson, Palme and Svensson [2012]). Workers who were DI claimants when they reached 65 became public pension claimants and received pension benefits at the same level as DI benefits. In 2003, DI became part of the sickness insurance system. Since then, DI benefits are 64 percent of labor income from the best three years from a five-to-eight-year period leading up to disability claiming. In the new DI system benefits are slightly higher than in the old system, but the pension rights earned from receiving DI is lower (Laun and Wallenius [2015]).

---

[41]Scenarios can be found in this government report (in Swedish): http://www.sou.gov.se/wp-content/uploads/2013/05/d99edc83.pdf. The number referred to in the text is taken from figure 13 and assumes future price indexation of basic pensions.

**Appendix A.1.2   Other Pensions**

Nine out of ten workers in Sweden are covered by collective bargaining agreements negotiated between trade unions and employer organizations. The terms of *occupational pensions* are a component of these collective bargaining agreements. There are four different occupational pension schemes: one for private sector blue collar workers, one for private sector white collar workers, one for local government employees and one for central government employees. Contributions to occupational pensions, which are are mandatory for workers covered by collective bargaining agreements, are paid in by employers to pension funds that are jointly owned and administered by trade unions and employer organizations. Like the 401(k) pension plans in the US, contributions receive deferred income tax treatment. In most schemes, pension benefits can be claimed at age 55 but the recipient is not allowed to work after claiming them. Claiming earlier results in an actuarial downward adjustment of the pension benefits. It is also possible to claim occupational pensions without claiming public pensions.

Individuals can also contribute voluntarily to *private pensions*.[42] Like occupational pensions, private pensions can be claimed from age 55 onward without incurring penalties. For example, individuals who claim their private pension can continue to work and the income earned from private pension does not affect social insurance eligibility.

---

[42]In Sweden individuals can save in so-called pension insurance policies. These are savings vehicles that invest in both risky assets, such as stocks, and low-risk assets like short-term bonds. While working, the individual saves money and after retirement or at a specified age, such as 55 or 60 years old, the individual receives an annuity each month from the policy for either a specified time, often 5-20 years, or for life. Hagen [2015] reports that 25-30 percent of all individuals claim their occupation pensions for a specified number of years. Surveys done by private pension providers indicate similar figures for private pension payments. The individual is typically guaranteed a certain minimum monthly payment by the issuer, hence the wording pension insurance. Until 2016 saving in private pension policies was tax deductible.

Figure A-1: THE BASE AMOUNT (BA) AND THE ENHANCED BASE AMOUNT (EBA), 1991-2011
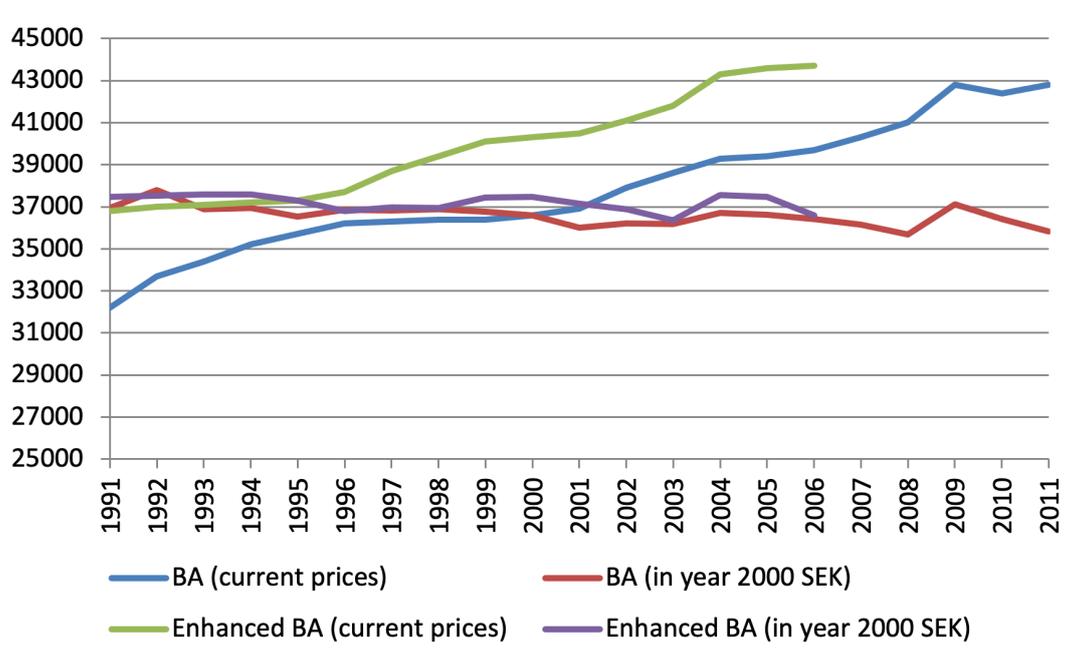
Notes: This figure shows the Base Amount (BA) and Enhanced Base Amount (EBA) over time. Both the BA and EBA are indexed against inflation.

Figure A-2: RELATIONSHIP BETWEEN INCOME DEPENDENT PENSION AND MINIMUM GUARANTEE
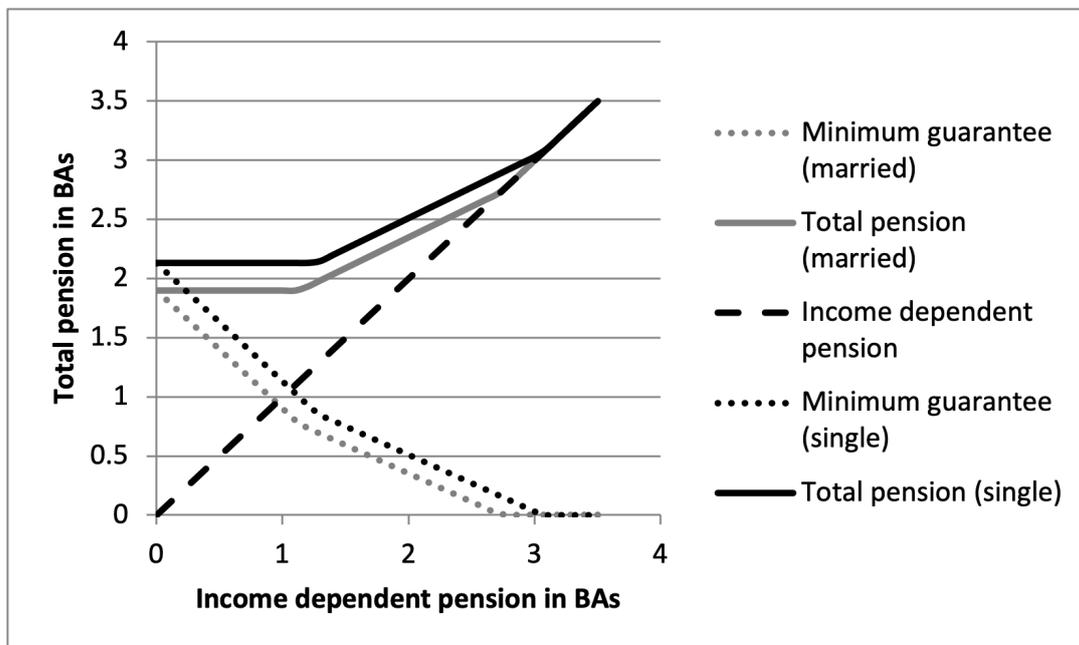
Notes: The income related pension is the same for singles and married. Total pension is the sum of the income dependent pension and the minimum guarantee.

Figure A-3: Distributions of Job Exit and Pension Claiming Ages
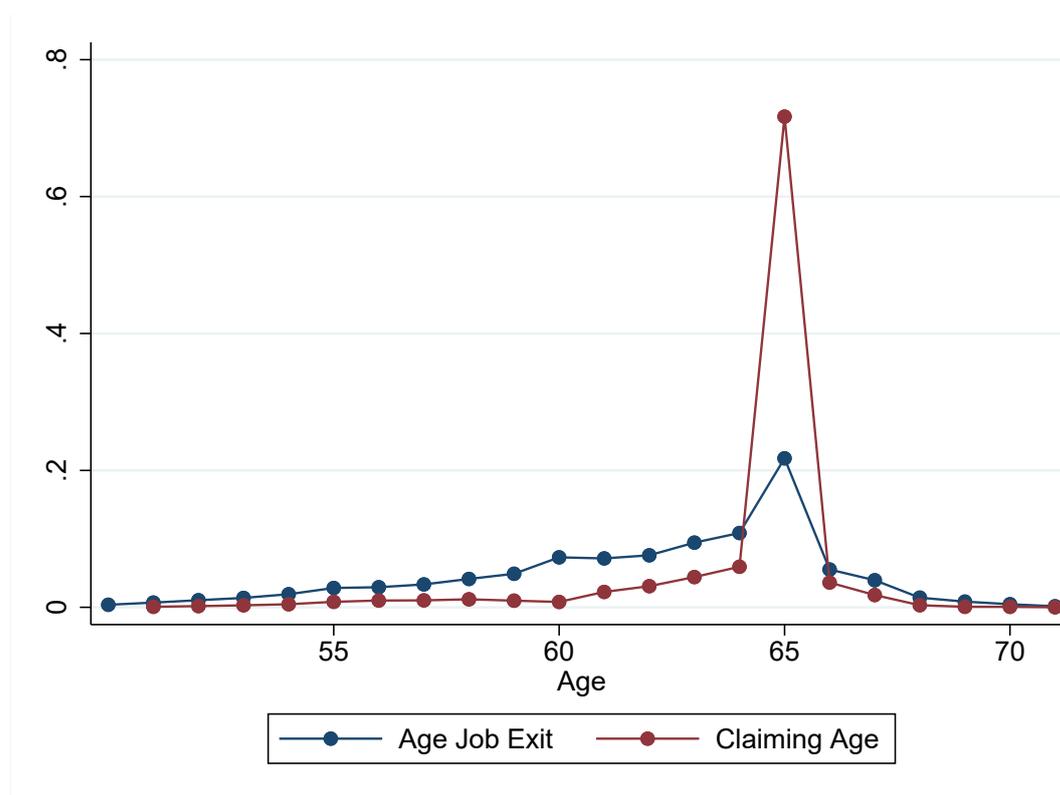
Notes: This figure shows the density distributions of job exit age and pension claiming age for workers born between 1938 and 1943.

Appendix A.2 Pension Simulation Details

Here we provide further details on our simulations of pension benefits. We use these simulations in the main text to characterize the effects of the Swedish pension reform on the profile of benefits over the retirement age, and to derive benchmark values for participation tax rates to quantify the fiscal externality.

Appendix A.2.1 Constructing Simulations

To guide our simulations, we imagine a hypothetical worker, aged 55, who is planning their retirement at some age between 55 and 70. The worker needs to know the effect that retiring at different ages will have on their pension benefits and overall income. The worker characteristics that are inputs for the simulation are:[43]

- The worker's birth cohort. We assume the worker is born in 1941 throughout, which is the midpoint of the birth cohorts we study in our empirical analysis.

- The worker's lifespan. Using mortality data, we estimate the expected lifespan of an individual from the 1941 cohort who reaches age 65. Based on this, we assume the worker lives until age 84.

---

[43]We do not consider aspects of the pension system like survivor benefits, under which pension benefits may also depend on marital status and gender.

- The workers marital status. This only matters for the minimum pension in either system; we assume the individual is single.

- The number of years worked before age 55. We calibrate this based on empirical data, see below.

- The workers' annual (pre-tax) earnings at 55. We calibrate this based on empirical data as well, see below.

- Whether the individual claims non-pension social insurance benefits (UI or DI) after retiring, and the duration and generosity of social insurance benefits. We calibrate these based on empirical data, and we present results with and without non-pension social insurance benefits.

- The age at which the individual claims their pension. We mainly assume the individual claims at 65, which as seen in Figure A-3 is the modal case.[44] We vary this in a sensitivity check.

- The age at which the individual retires (permanently stops working). This is the x-axis of the figures derived from this calculator. We vary this from age 55 to 70 in one-year increments for each specification of the above characteristics.

Given these inputs, we first simulate a complete earnings path for our individual. For years before the worker turns 55, the earnings history is based on empirical earnings growth rates, given the number of years worked and earnings at 55.[45] For years after age 55, we use a constant growth rate based on average earnings growth from 1996 to 2011. This ensures that idiosyncracies in earnings growth do not generate noise in our simulated NPVs and tax rates, and it is consistent with the intuition that a worker contemplating retirement knows their earnings history before age 55 but only knows their expected earnings after age 55.

Given the earnings history and other characteristics, we then calculate the workers' lifetime pension benefits in either the ATP and NDC system, as a function of the exit age, given the assumed claiming age and longevity. The worker will receive pension benefits from claiming age until death. As we did with earnings histories, for both the ATP and NDC systems we use actual, empirical basic amounts ("income base amounts" in NDC) up to age 55, and after age 55 we use the average growth rate of the base amounts from 1996-2011. Once again this ensures that idiosyncracies in base amounts do not generate noise in the NPVs and participation tax rates. By design, the average growth rates of base amounts are very similar to that of the price index for ATP and the wage index for NDC.

We then calculate the NPV at age 55 of lifetime pension benefits at each possible retirement age from 55 to 70. We include non-pension social insurance benefits and the pension rights they

---

[44]Of individuals retiring between 60 and 63, 76% claim their pension at age 65, and only 13% claim at job exit or one year later. Of individuals retiring between 55 and 59, 52% claim their pension at age 65, and only 4% claim at age 61, the earliest age possible.

[45]For simplicity we assume the worker worked continuously from some starting year until age 55. For example, a worker with 30 years of experience at age 55 would be assumed to start working at age 25.

provide in this NPV, but as we shall see this has a small effect. We use a discount rate of 0.98 to calculate NPVs, under which the adjustments to benefits in the NDC system that should be actuarially fair on average are in fact actuarially fair. Thus we obtain the slope of the pension benefit profile over retirement ages for a worker with the specified characteristics.

Next, we simulate participation tax rates at each possible retirement age. For a given age a, these are defined as

$$(13) \qquad \text{Participation Tax Rate}_a = \frac{\text{income tax}_a + \text{payroll tax}_a - [NPV_a - NPV_{a-1}]}{\text{Gross earnings}_a},$$

where NPV_a is the net present value at 55 of pension benefits for a worker retiring at age a, and both payroll tax and gross earnings include employer payroll tax contributions.[46] Finally, conceptually it is useful to separate out the component of the participation tax rate that is directly attributable to the pension system, i.e. payroll taxes that fund pensions (a flat tax rate of 18.5% of gross earnings in both systems) and the change in the NPV of pensions. This is calculated similarly to the above, as:

$$(14) \qquad \text{Implicit Tax Rate}_a = \frac{\text{pension payroll tax}_a - [NPV_a - NPV_{a-1}]}{\text{Gross earnings}_a}.$$

The difference between implicit and total participation tax rates therefore represents non-pension payroll taxes and income taxes.

Appendix A.2.2   Accounting for Heterogeneity by Lifetime Earnings

Our simulator performs all of the above for any specified set of worker characteristics. Our goal is to use these simulations to paint a reasonably complete picture of how the reform affected the slope of the pension benefit profile on average, accounting for differences across workers. The main form of heterogeneity we should account for in doing so is heterogeneity in earnings histories. We use some empirical moments to calibrate our simulations along these lines.

Specifically, we divide the sample of individuals born from 1938-1943 – the main cohorts of interest for our analysis – into 20 vigintiles based on individuals' accrued ATP pension rights as of age 55. Accrued pension rights are an attractive proxy for earnings histories; we do not observe full earnings histories, but this proxy mechanically captures the features of the earnings history that matter for pension benefits. Some complications arise from the cap on ATP pension rights: all individuals in the 20th vigintile have the maximum possible ATP pension at 55. In the 19th vigintile, 63% of individuals have the maximum possible ATP pension at 55. Individuals reaching the cap are split randomly between the 19th and 20th vigintiles.
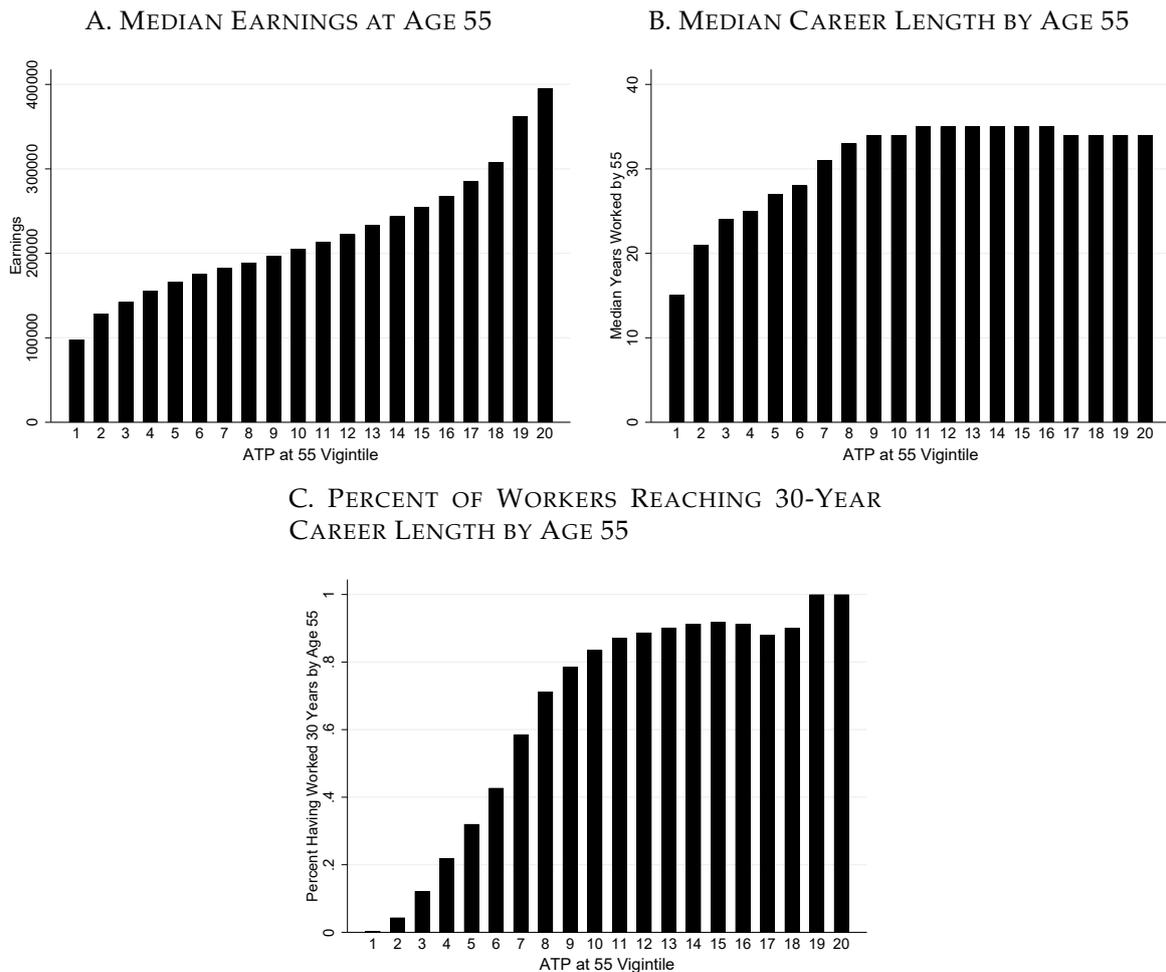
We then think of these 20 vigintiles of accrued ATP rights at age 55 as 20 different workers, each of whom represents 5% of the population of interest. We run the simulator described above 20 different times, where the worker characteristics are based on the characteristics of a

---

[46]This gross earnings concept is sometimes called "super-gross" earnings, to distinguish it from earnings gross of income and employee payroll taxes but not employer payroll taxes.

typical worker in the given vigintile of accrued ATP rights at 55. We use one set of moments to discipline labor earnings and public pension benefits, and another to account for non-pension social insurance benefits.

Labor Earnings and Pension Benefits We estimate the median earnings and median years worked within each vigintile, and use plug these into the simulator for each of the 20 hypothetical workers. These medians are plotted in Figure A-4 below, along with the fraction of workers who have worked beyond 30 years by age 55, which is important for the ATP system.

Figure A-4: EARNINGS AND CAREER LENGTH BY VIGINTILE OF ATP AT 55

A. MEDIAN EARNINGS AT AGE 55   B. MEDIAN CAREER LENGTH BY AGE 55

C. PERCENT OF WORKERS REACHING 30-YEAR CAREER LENGTH BY AGE 55

Notes: Panel A of this figure shows the median earnings at 55 of workers born between 1938 and 1943 for each ATP at 55 vigintile. Panel B shows the median years worked by age 55 for each vigintile. Panel C shows the percent who reach a career length of 30 years by the age of 55 for each ATP at 55 vigintile.

To validate this basic approach and the way we construct earnings histories, the most important thing to verify is that the earnings history we construct implies a reasonable level of accrued pension rights as of age 55. Although we divided individuals into vigintiles based on observed pension rights at age 55 in the data, our simulator constructs ATP pension rights at 55 based on the simulated earnings history, i.e. based on earnings at 55, career length, and average earnings growth in the full population. In Figure A-5, we verify that simulated ATP rights accrued as of age 55 closely match actual, empirical ATP rights accrued as of age 55,

implying that the simulation constructs realistic earnings histories throughout the distribution, and therefore that it will provide an accurate picture of the pension benefits profile and participation tax rates through the distribution.

Figure A-5: VALIDATION OF SIMULATED EARNINGS HISTORIES
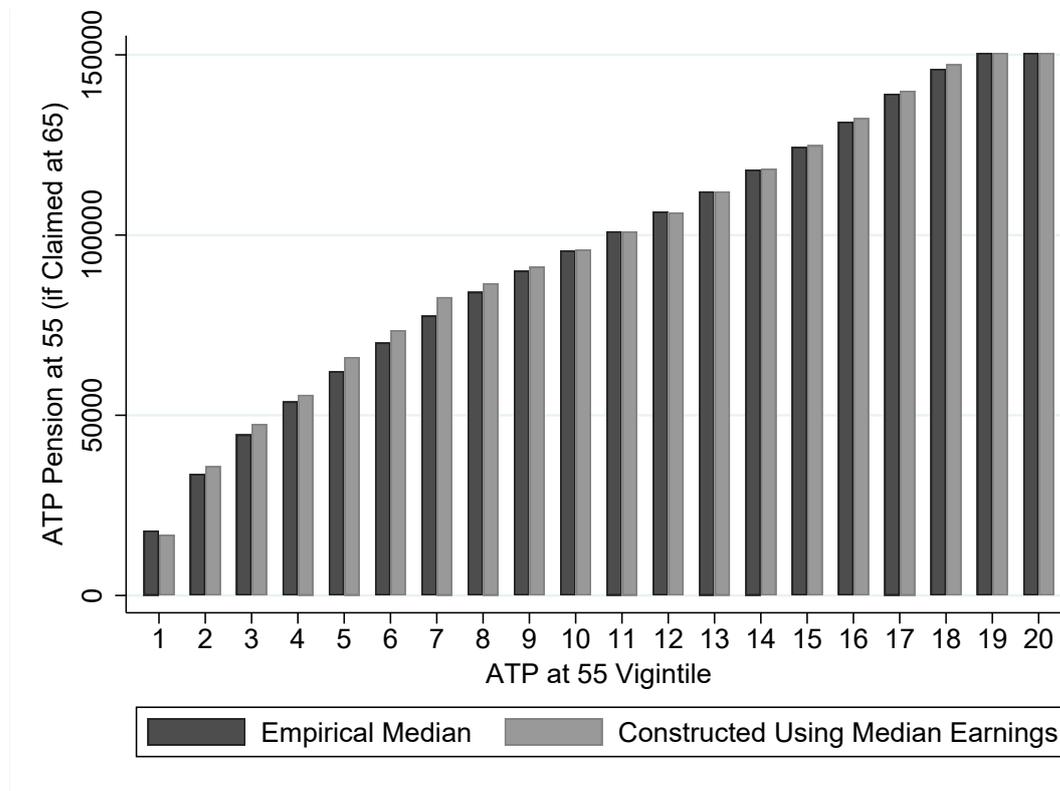
Notes: This figure shows the median ATP pension that workers born between 1938 and 1943 were eligible for at age 55 (assuming the pension is claimed at age 65) plotted alongside the simulated ATP pension. These are shown for each ATP at 55 vigintile.

Non-Pension Social Insurance   We next consider how we should empirically discipline non-pension social insurance benefits received after job exit (and before pension claiming).

These programs turn out to matter little for the shape of the pension profile, but we should account for them because early retirees do claim these benefits with some regularity. Figure A-6 plots the empirical proportion of individuals receiving UI or DI after they retire by ATP vigintile at 55. Panel (b) focuses on premature retirees, those retiring before 61. We observe that low-income, premature retirees in particular are likely to claim UI or DI after exiting and before claiming, which makes sense given our other findings (e.g. on health shocks) and the fact that these workers exit the labor market before they can claim their public pension (at 61).

Figure A-6: Proportion of Individuals Receiving UI or DI After Retirement by ATP Vigintile at 55

A. All (Retirement Ages 56-69)   B. Premature Retirees (Ret. Ages 56-60)
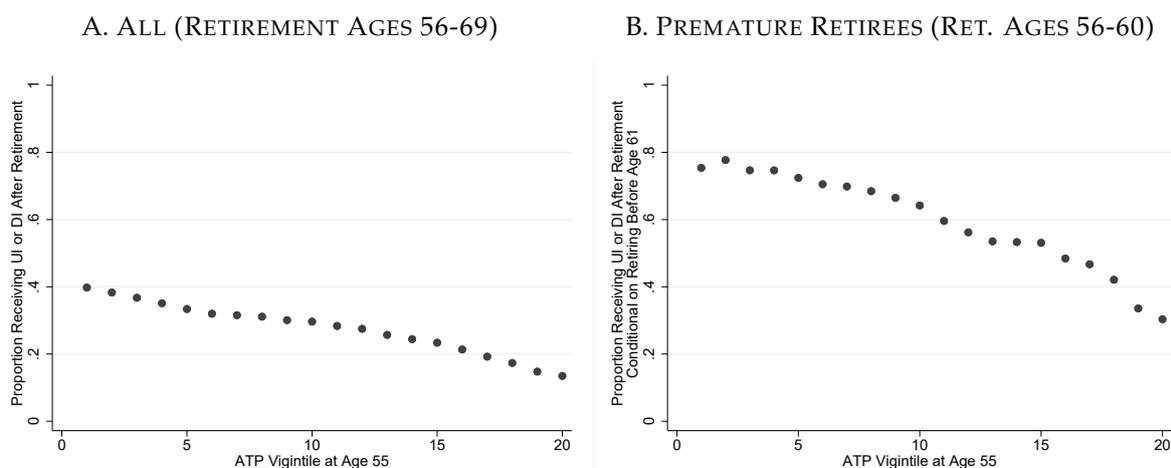
Notes: These graphs show the proportion of individuals in each ATP at 55 vigintile who receive UI or DI after retiring. Panels B restricts to individuals retiring before age 61.

To account for the effect of these benefits on the pension profile and incentives, we suppose that our hypothetical age-55 worker knows that in the event they retire before 65 there is some probability that they will claim UI or DI afterwards, and they have some expectations of how much and how long they would receive these benefits. The worker factors these possibilities into their expected NPV of benefits (pension benefits plus other social insurance benefits claimed during these ages). We therefore estimate the following three parameters by vigintile of ATP at 55 and (exact) exit age: (1) the probability of claiming UI or DI after exit, 2) the median annual benefit amount (conditinal on receiving benefits), and 3) the median benefit duration (in years). We estimate both benefit amounts and durations conditional on claiming UI or DI after exiting. From the last of these we find that assuming the individual claims for one year if exiting at age 63 or earlier, and the individual does not claim if exiting at age 64 or later, provides a reasonable approximation to reality.

We specify the NPV of pension benefits for a given age and vigintile of ATP at 55 as the weighted mean of the NPV of pension benefits without any non-pension social insurance claims and the NPV of benefits if the individual claims non-pension social insurance benefits for one year after exiting. The weights are given by the probability of claiming from (1) above and the levels of non-pension SI benefits are the median generosity of benefits from (2) above. The NPV in the case where the individual claims non-pension SI benefits accounts for adjustments to pension benefits from social insurance receipt, and to the value of these benefits themselves.

We also present results for the simpler case where individuals do not claim any non-pension SI benefits, to show how much this matters.

Appendix A.2.3   Results

Given these calibrations, we then simulate the NPV of pension benefits and participation tax rates for each of the 20 hypothetical workers. To arrive at Panel A of Figure 1 in the main

text, we average the resulting NPVs across individuals and subtract the level shift in overall benefits from the NDC system. The latter step is quite straightforward and we describe how this is done at the very end of this Appendix. Until then, in order to provide a complete and transparent characterization of the NDC reform and address some conceptual issues that are unrelated to the levels issue, we plot the NPV of benefits in the actual NDC system rather than the illustrative, budget-neutral version of NDC used in Figure 1. As a result, the NPV of benefits in the NDC system in the next few figures is lower than what we plot in Figure 1, because the NDC system decreased benefits for most workers.

Figure A-7 shows the NPV of benefits for different retirement ages compare in the ATP and NDC systems. We observe the same change in the steepness of the pension benefits profile as Figure 1, along with a level decrease in benefits in the NDC system. We also show how assuming individuals never claim non-pension SI benefits ("Without Benefits") affects our picture of the pension profile. We observe that our treatment of non-pension SI benefits matters very little, even for premature retirees. Intuitively, the main reason these benefits matter little is that the typical non-pension benefit duration is relatively short compared to the duration of receipt of public pension benefits.

Figure A-7: NPVs With and Without Post-Retirement Benefits
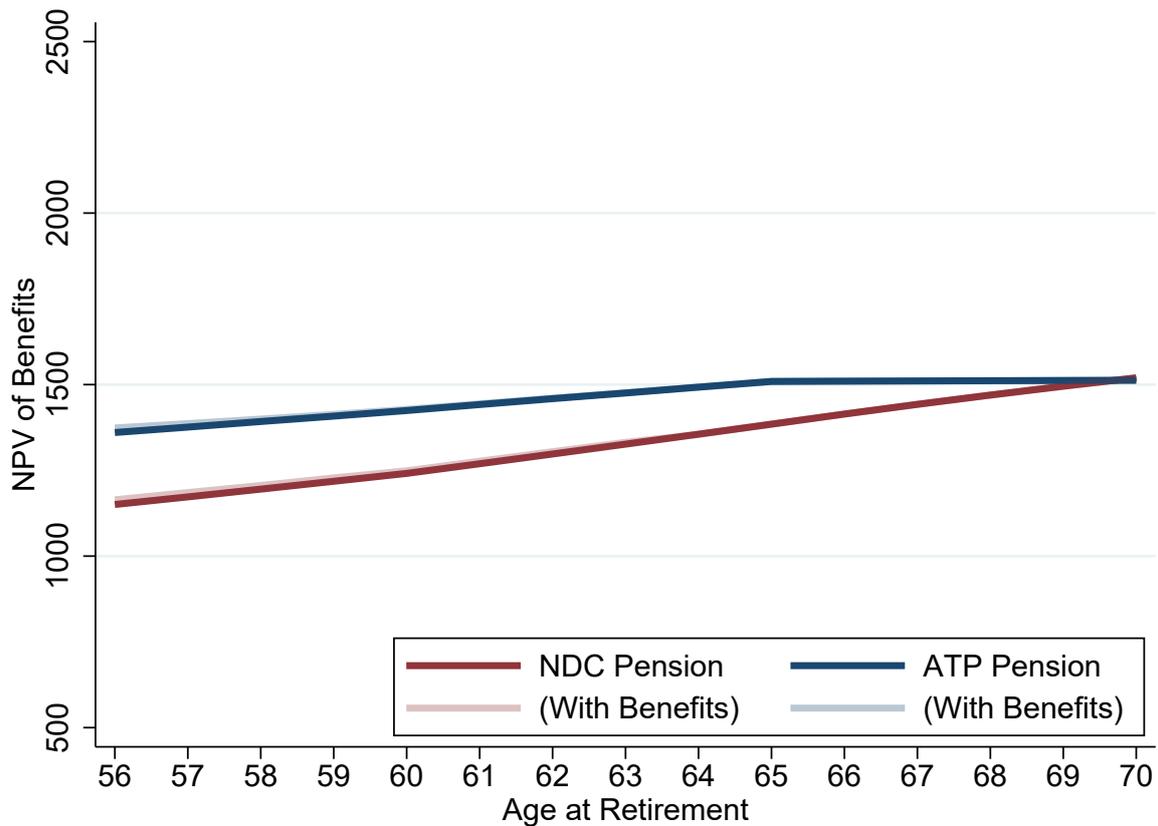
Notes: This graph shows the mean net present value (NPV) of pension benefits across the 20 ATP at 55 vigintiles for each retirement age. The opaque lines show this for an individual who does not receive post-retirement UI or DI benefits. The transparent lines show the weighted mean of the NPVs without post-retirement benefits and the NPVs with 1 year of post-retirement benefits, with the size of the benefit equal to x. x is equal to the median post-retirement benefits received for each retirement age and ATP vigintile. The weights are the probabilities of receiving post-retirement benefits for each retirement age and ATP vigintile but are set to zero for ages 64 and greater.

To get some sense of how the reform affected the steepness of the pension profile heterogeneously through the distribution of lifetime earnings, we also plot the NPVs of the ATP and NDC system in the top and bottom decile of the lifetime earnings (averaging across the top and bottom two vigintiles). These results are in Figure A-8. We observe that in the bottom decile, the higher minimum pension benefit in the NDC system resulted in a level increase in benefits for some workers, along with a flatter profile in the NDC system than in the ATP system (in contrast to most of the distribution). In the top decile, meanwhile, the cap on ATP pension benefits – attained by maxing out pensionable income for at least 15 years and contributing for at least 30 years – was binding for nearly all workers, while the higher cap on the NDC system is not. This results in a steepening of the pension profile and, at later retirement ages, higher benefits after the reform than before. For all other parts of the distribution, where the minimum and maximum on benefits are seldom binding, the qualitative effects of the reform on the pension benefits is similar to that of Figure A-7.

Figure A-8: Net Present Value of Pension Benefits by Age at Retirement

A. 1st Decile

B. 10th Decile
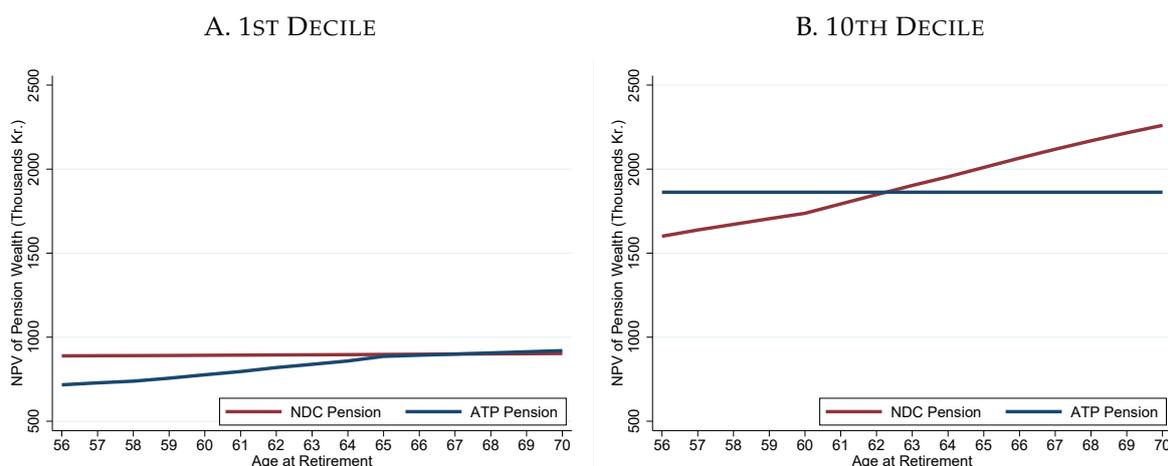
Notes: These graphs show the net present values (NPVs) of pension wealth by age at retirement for the top and bottom deciles of the distribution of ATP at age 55. The graph for each decile is created using the average NPVs of both vigintiles within that decile. Calculations are for individuals born in 1941 with a discount factor of 0.98.

We next turn to the participation and implicit tax rates, which, as discussed in Section VI, are an essential determinant of the fiscal externality from a change in steepness. Figure A-9 plots these tax rates, averaging once again over our 20 hypothetical workers. Most importantly for the welfare calculations in Section VI, we observe that a participation tax rate of 0.45 for each retirement age provides a reasonable approximation to reality in either system. The most prominent effect of the NDC reform was to decrease both the implicit and participation tax rates after age 65. This occurs because working past 65 did not accumulate pension rights in the ATP system, which acts as an implicit tax on earnings, while the NDC system allows individuals to accumulate pension rights.

We note that the tax rates in Figure A-9 vary slightly and somewhat arbitrarily across retirement ages before 65. This occurs because the empirical moments underlying our specification of non-pension social insurance benefits vary somewhat with retirement ages, which introduces some noise into the simulated tax rates. To show this and understand how these benefits contribute to the tax rates overall, we also simulate participation and implicit tax rates in both systems for a scenario in which individuals never claim non-pension SI benefits. In this case, the tax rates flatten out and are virtually constant across ages, and the participation tax rate before 65 is slightly lower at about 0.4 in both the ATP and NDC system. Using a participation tax rate of 0.4 rather than 0.45 would have a negligible impact on the benchmark fiscal externality we use in the main text, changing it from about .15 to .13.

Figure A-9: Participation and Implicit Tax Rates, With and Without Non-Pension Social Insurance Benefits
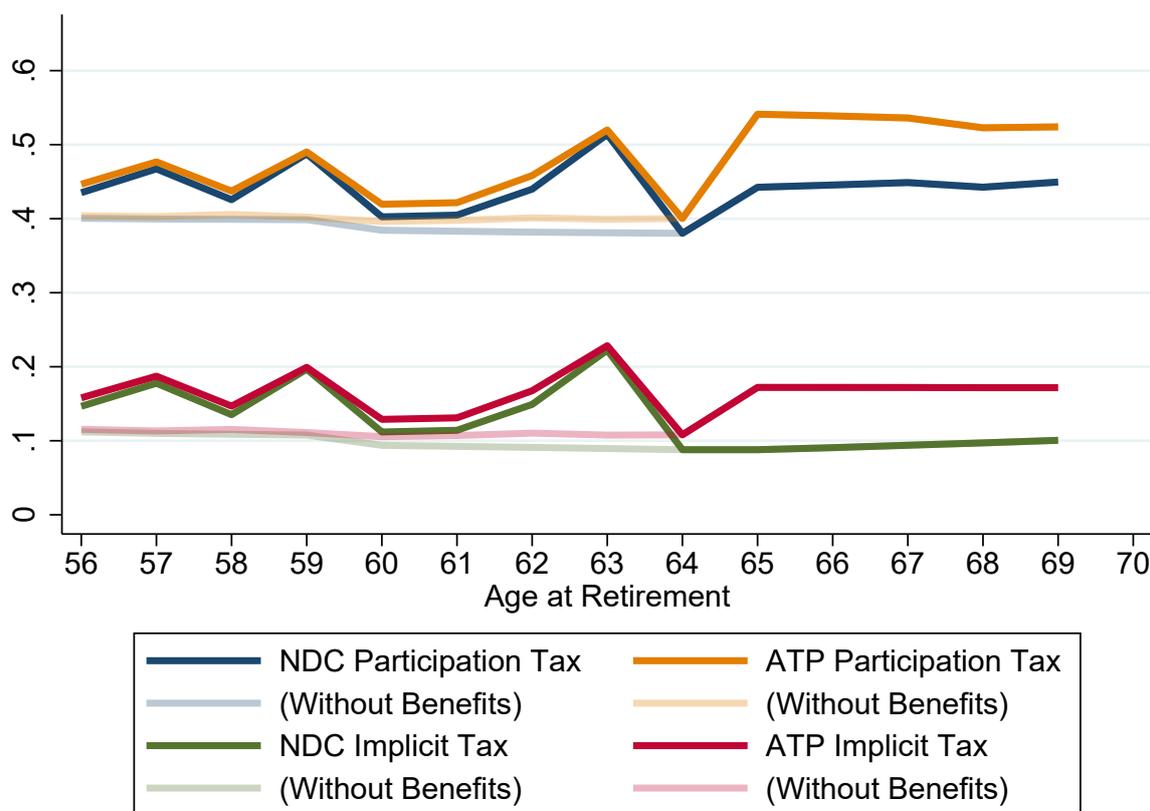
Notes: This graph shows the mean implicit and participation tax rates of pension benefits across the 20 ATP at 55 vigintiles for each retirement age. The opaque lines show this for an individual who does not receive post-retirement UI or DI benefits. The transparent lines show the participation and implicit tax rates when the NPVs are equal to the weighted mean of the NPVs without post-retirement benefits and the NPVs with 1 year of post-retirement benefits with the size of the benefit equal to x. x is equal to the median post-retirement benefits received for each retirement age and ATP vigintile. The weights are the probabilities of receiving post-retirement benefits for each retirement age and ATP vigintile but are set to zero for ages 64 and greater.

## Appendix A.2.4 Alternative Claiming Age Specification

As discussed in the main text, we primarily focus on incentives to retire at other ages, setting aside the question of the claiming age. In the simulations, we held the claiming age fixed at its modal value, age 65. While justified based on the Swedish case (see Figure A-3 and the discussion in Section A), this choice creates some difficulties in interpreting the incentives for retiring after age 65. Here we discuss these complications and simulate an alternative scenario for illustrative purposes.

Most importantly, the participation tax rates in Figure A-9 increase modestly after age 65 even in the NDC system. We observe that this does not derive from the implicit tax rate, which captures everything to do with the pension system, but rather the residual component of the participation tax rate. Rather, this derives from progressive income tax rates. If an individual claims at 65 and works at some age beyond 65, the individual would face a higher average

income tax rate on their labor and pension income combined than on their labor income alone. To show that this does in fact drive the increase in the participation tax rates in the NDC system, and get some idea of how a later claiming date would affect the relevant tax rates for late retirees, we plot the average pension profile and tax rates for a scenario in which individuals always claim at 70. We continue to average across 20 simulations, each representing 5% of the lifetime earnings distribution. For simplicity, we focus on the case where individuals never claim non-pension social insurance benefits after retiring, which are not material for the main point of this exercise.

We observe that the pension profile in Figure A-10 is very similar to Figure 1/A-7. The main difference is that the level difference between ATP and NDC profiles is slightly larger, which occurs because ATP system incorporates slightly more generous adjustments for those claiming after 65. In Figure A-11, we observe that the implicit and participation tax rates before age 65 are very similar to Figure A-9 (without non-pension benefits), suggesting a flat participation tax rate of about 0.4. After age 65, NDC participation tax rate remains constant at around 0.4 or just below in the claim at 70 specification. This confirms that the increase in this tax rate at 65 in Figure A-9 is driven by the progressivity of the income tax schedule. As such, this increase in participation tax rates is spurious for the purpose of understanding the incentives faced by late retirees retiring after 65 – such workers typically also claim after 65. The most important implication of all this is that using a constant participation tax rate at different retirement ages in our benchmark for the fiscal externality provides a good approximation to reality, even after age 65.

Figure A-10: AVERAGE NET PRESENT VALUE OF PENSION BENEFITS - CLAIM AT AGE 70
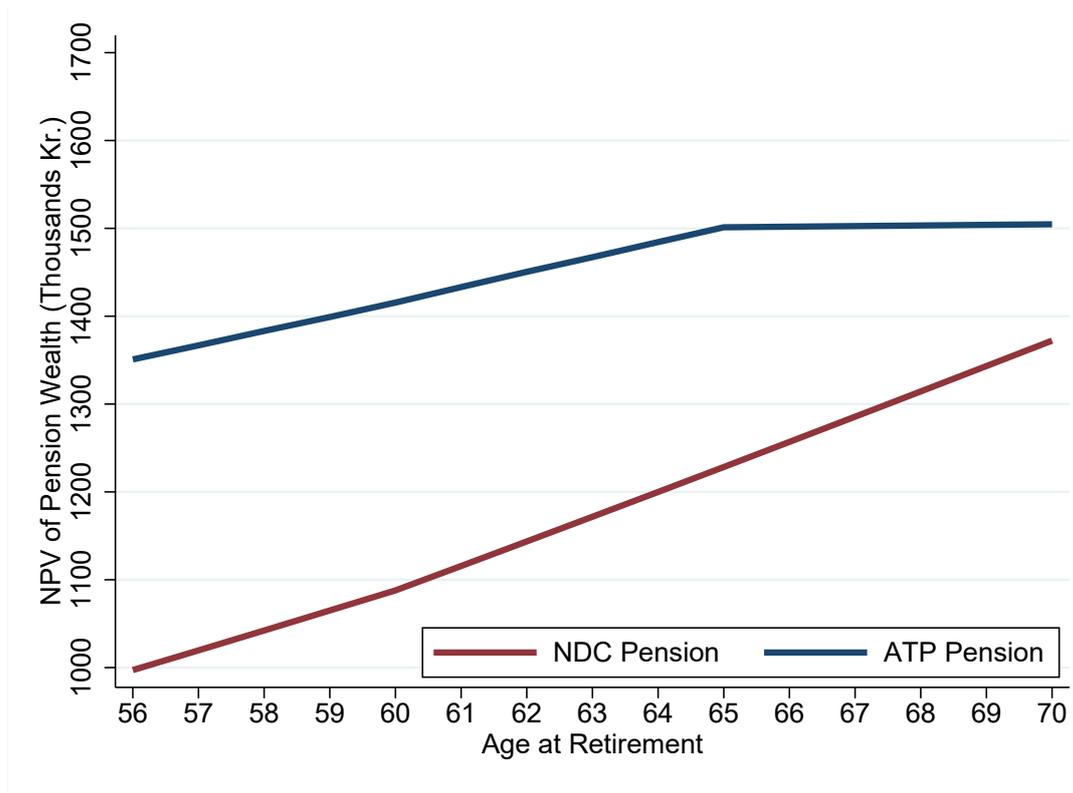
Notes: This graph shows the net present value of pension wealth by age at retirement averaged across all ATP at 55 vigintiles. Calculations are for individuals born in 1941 with a discount factor of 0.98.

Figure A-11: AVERAGE IMPLICIT TAX RATE - CLAIM AT AGE 70
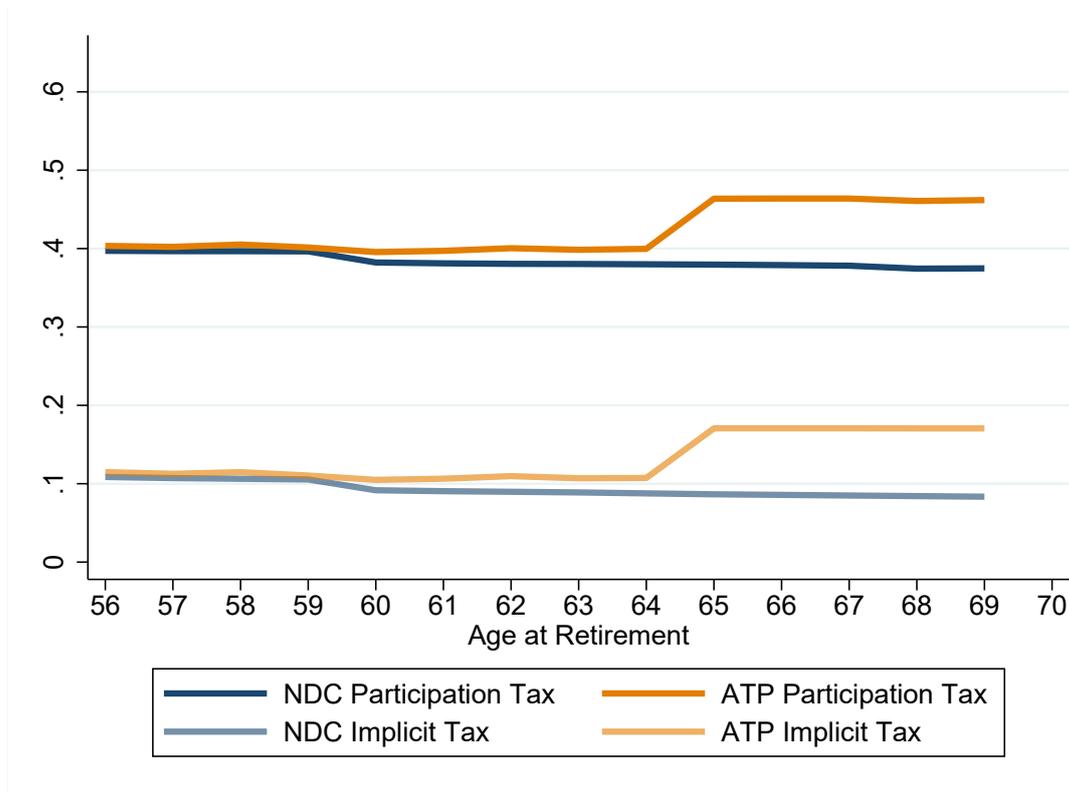
Notes: This figure shows the average participation tax rate and implicit tax rate across all 20 ATP at 55 vigintiles by age at retirement .

Appendix A.2.5   A Balanced-Budget NDC Reform

The above simulates pension benefits profiles for the actual ATP and NDC pension schemes. As one of the goals of the reform was to promote fiscal sustainability, the reform was not budget-neutral. In our theoretical framework, we characterize the effects that this reform would have had if it were budget neutral. As such, we calculate a profile that has the same budget as the ATP scheme but the same slope as the NDC scheme. We call this "budget-neutral" NDC in Figure 1.

Let $f(r)$ denote the fraction of individuals with retirement age r. Denoting the NPV of benefits at age r in the ATP and NDC schemes by ATP_r and NDC_r, respectively, our goal is to find a profile $\widehat{NDC}_r$ with the desired properties.
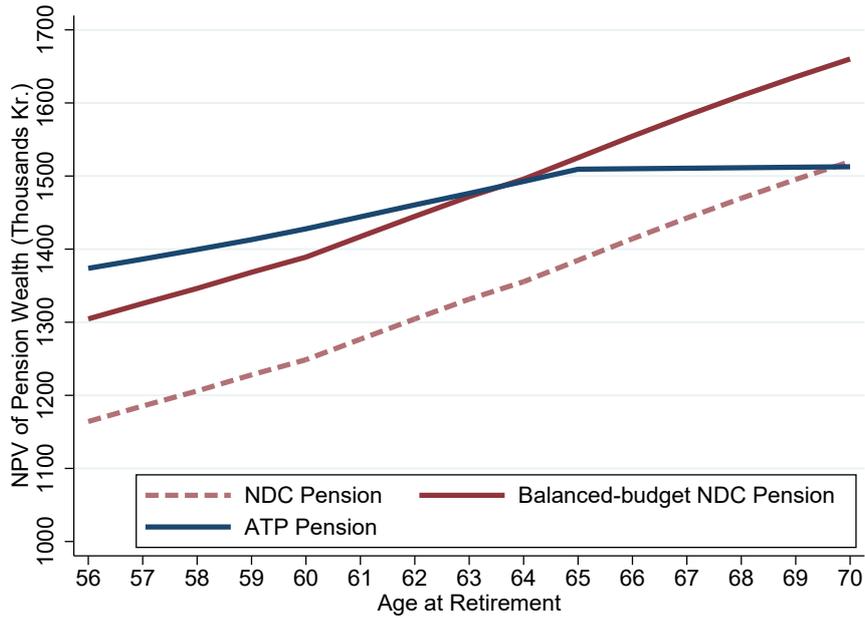
Keeping the budget fixed at the ATP level requires:

$$(15) \qquad \sum_{r=56}^{69} [ATP_r f(r)] = \sum_{r=56}^{69} [\widehat{NDC}_r f(r)]$$

Keeping the slope of the profile the same as the NDC throughout requires that for any r,

$$(16) \qquad \widehat{NDC}_r = \Delta + NDC_r.$$

67

Figure A-12: NET PRESENT VALUE OF PENSION BENEFITS - ACTUAL NDC PROFILE

Notes: This figure shows the net present value (NPV) of pension wealth by age at retirement averaged across all ATP at 55 vigintiles. Calculations are for individuals born in 1941 with a discount factor of 0.98. NPVs are shown for both the actual NDC pension and the balanced-budget version of the NDC pension.

Plugging this into equation (15) and solving for Δ we obtain:

$$(17) \qquad \Delta = \frac{\sum_{r=56}^{69}[ATP_r f(r)] - \sum_{r=56}^{69}[NDC_r f(r)]}{\sum_{r=56}^{69}[f(r)]}.$$

Figure 1 in the main text draws on the budget-neutral version of the NDC reform, $\widehat{NDC_r}$. Figure A-12 compares the ATP profile (ATP_r), the actual NDC profile (NDC_r) and the budget-neutral NDC profile $\widehat{NDC_r}$. The implementation results that characterize the change in slope in the Swedish reform are also based on a comparison of $\widehat{NDC_r}$ and ATP_r (see Appendix H for further details on the implementation).

Appendix B Data - Additional Details

Residual Measure of Consumption Expenditures

Our third registry data source is granular data on wealth from the wealth registry. These data were collected by Statistics Sweden 1999-2007, years when Sweden was taxing wealth.[47] The data contains information on real estate, stocks, bonds, other securities, debt, and bank account holdings. With this data we construct a residual consumption measure using the budget identity:

$$(18) \qquad\qquad Consumption = Income - Saving.$$

The consumption measure is one of consumption *expenditure* and records consumption on all goods paid for by taxed and recorded income. A number of recent papers have use such consumption measures, based on Scandinavian population registers with detailed information on income and assets.

All details on the data and programs used to create our measure of consumption can also be found at: http://sticerd.lse.ac.uk/_new/research/pep/consumption/default.asp.

To construct our consumption measure we follow the same method as Kolsrud, Landais and Spinnewijn [2020]. The income measure used is disposable income which is constructed by Statistics Sweden, and is included in the LISA panel. It contains the net-of-tax value of labor earnings, capital earnings, flow value of student loans (received and amortized) and social insurance benefits. Saving is defined as the change in asset holdings and debt after we have accounted for passive capital gains. Capital income and student loans are removed from the disposable income measure to prevent double counting. For stocks and bonds we use the number of securities each person holds on December 31st each year and valuate them according to the end-of-day price on December 31st. For real estate we use data from the property register which covers real estate transactions which are then linked to buyers and sellers. When individuals have no transactions consumption from real estate is zero. Debt is the sum of all types debt; mortgages, consumer credits and student loans. We cannot separate mortgages or consumer credits from the stock of debt an individual holds.

Specifically, consumption expenditures C_{it} by household i in period t is written as

$$(19) \qquad\qquad C_{it} = Z_{it} - \sum_k p_{kt} \left[A_{ikt} - A_{ikt-1} \right],$$

where Z_{it} captures all sources of income and transfers, $\mathbf{A}_{it} = A_{i1t}, .., A_{iKt}$ denotes the portfolio of assets and $\mathbf{p}_t = p_{i1t}, .., p_{iKt}$ the corresponding vector of prices at which they are traded. With wealth data spanning 1999-2007 we can estimate consumption expenditure 2000-2007.

---

[47]The wealth tax was installed in 1947 and repealed in 2006. Data was also collected in 2007. Before 1999 only data on total wealth is available and, mostly, only for individuals or households subject to wealth tax, about 5 percent of the population.

The wealth data are annual and financial assets are recorded on December 31st each year. This means that we cannot detect intra-year trading. Baker et al. [forthcoming] find that the error this creates is small on average though it may be important for some households. We also do not account for trading fees. However, these can be seen as a consumption expenditure; individuals purchase a service − investment counseling − which they pay for and this cost is included in the consumption expenditure measurement. See Kolsrud, Landais and Spinnewijn [2020] for further detail on the consumption expenditure measure.

Career Length at 55

Figure B-1: CAREER LENGTH AT 55

A. Distribution of Career Length at 55
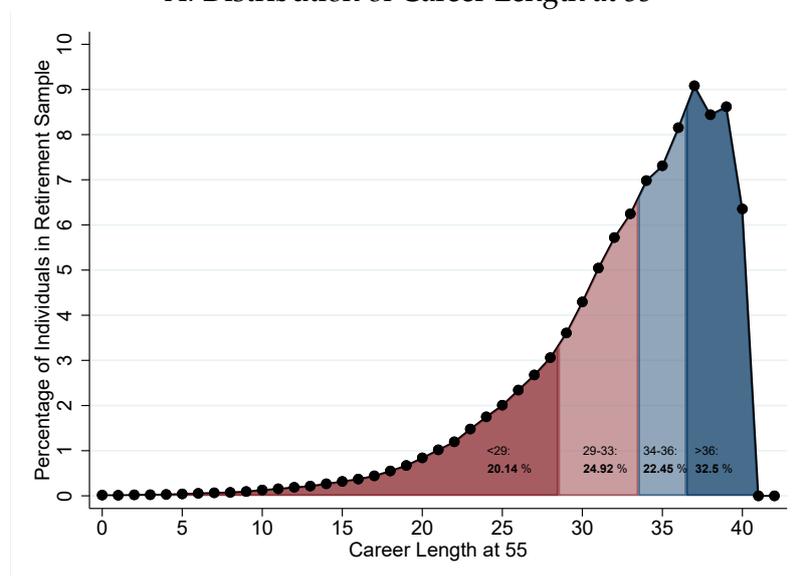
B. Distribution of Retirement Age by Career Length at 55 Groups

Notes: The figure provides information about the distribution of career length at age 55 in our sample. Career length at 55 is defined as the sum of years during which an individual made positive contr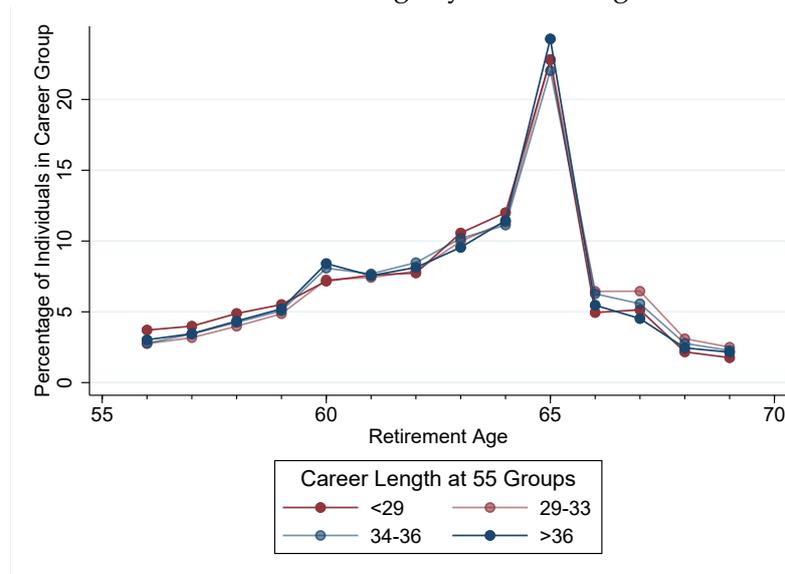ibutions to the ATP pension system before turning 55. We observe actual contributions before 55 years of age for individuals born after 1938 subject to one limitation: the data only spans as far back as 1960, which implies that we are missing contributions from ages 17-22 and before, depending on the cohort. To overcome this limitation, we impute additional years of contributions as follows. First, we leverage data from the 1990 wave of LISA on formal education to infer the number of years of schooling for individuals in our sample. Next, we assume that individuals are schooled continuously from age 7, and that they do not work and study at the same time. Finally, we assume that they start working the year after leaving formal education. If this year falls before 1960, we impute the difference as additional years of contribution to the ATP system. Panel A reports the distribution of our measure of career length at 55 among individuals from the 1938 to 1950 cohorts in Sweden. Compared to Figure 2, we also include cohorts 1944 to 1950, for which we observe actual contributions from age 16, to reduce reliance on the imputation procedure. The groups of career length at 55 used in the analysis are roughly based on quartiles of this distribution. Panel B reports the distribution of retirement age in our baseline sample, as in Figure 2, but splitting the sample into groups based on career length at 55.

Figure B-2: HETEROGENEITY & SELECTION INTO CAREER LENGTH AT 55

A. Socio-Demographic Characteristics
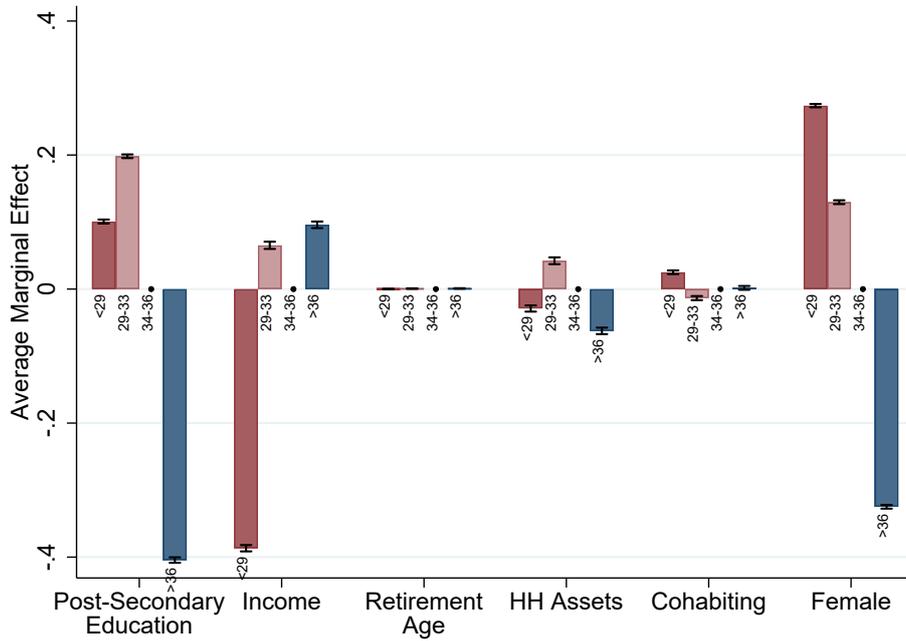
B. Health and Life Expectancy
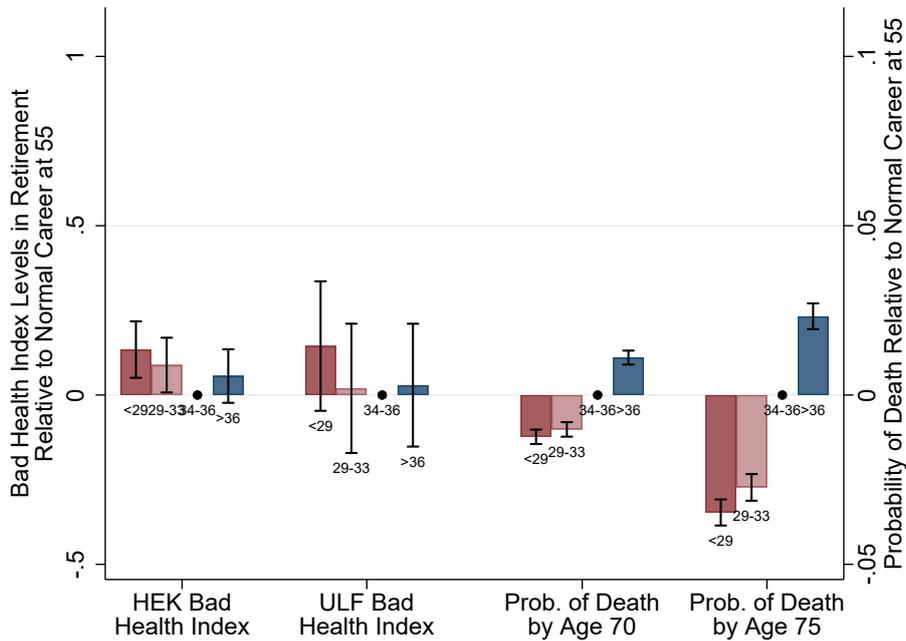
Notes: The figure documents patterns of heterogeneity across groups based on career length at age 55. The construction of the figure is exactly analogous to Figure 7. Panel A displays estimates from a multinomial logit prediction model for belonging to one of the 4 different career length groups. The model includes cohort fixed effects, a dummy for having post-secondary education, the within-cohort rank of average income between 52 and 55, retirement age, the within-cohort rank of average household assets between 1999 and 2007, a dummy for being married or cohabitating and a gender dummy. We report for each regressor the estimated average marginal effects on the relative probability to select into each of the group, using normal retirees as reference category. Panel B explores selection on health and life expectancy. The graph reports estimates from the analogue of specification (8) (with cohort and age fixed effects and controls for family structure). We replace consumption by our two indices for bad health (i.e. standardized principal components extracted from all health outcomes in the HEK and ULF surveys) and two measures of "life expectancy" (dummies for being dead by age 70, or by age 75). For the latter outcomes, we have one observation per individual and drop age fixed effects in the regression.

72

Appendix C Consumption Levels & Heterogeneity

Consumption Differences By Retirement Age: Robustness

Figure C-1: CONSUMPTION DIFFERENCES BY RETIREMENT AGE
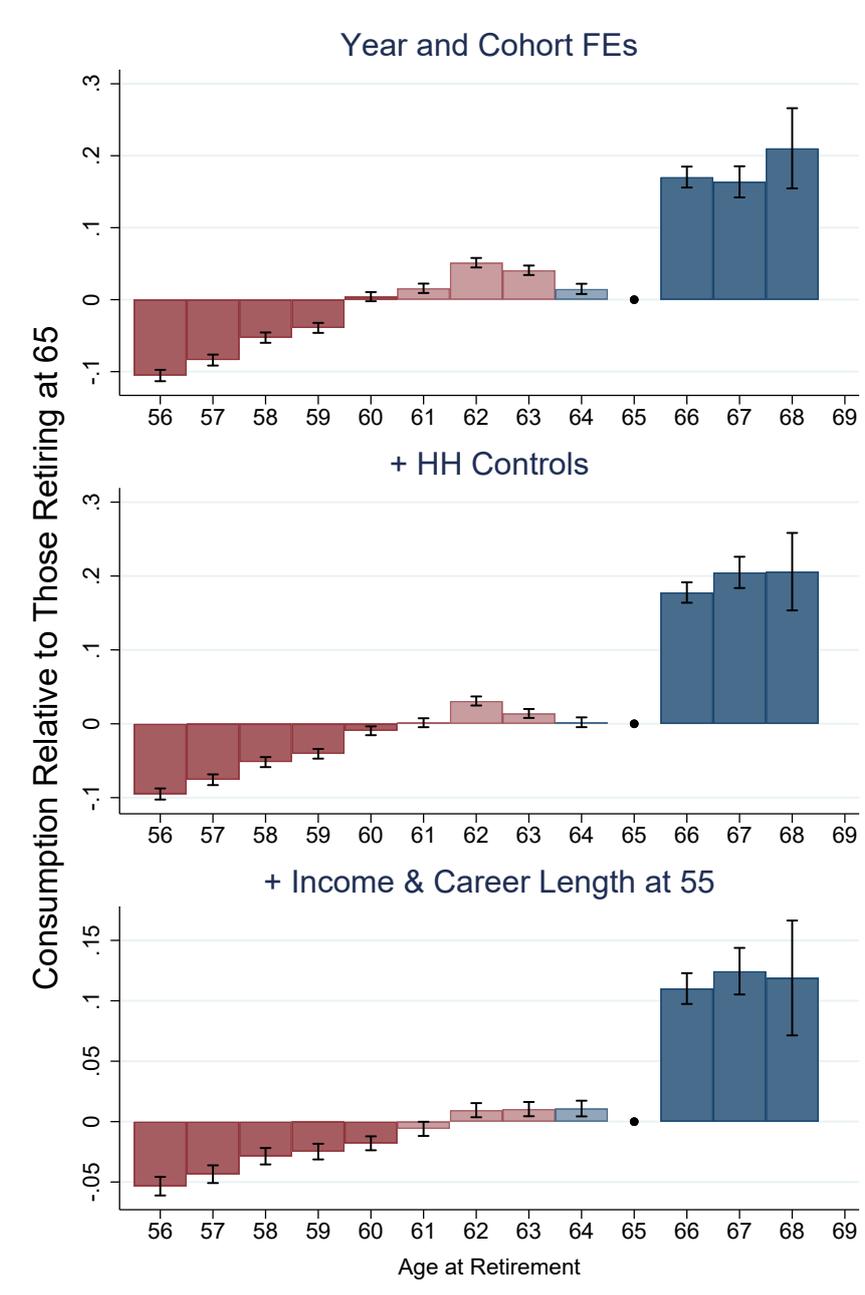
Notes: The figure report estimates of a fully non-parametric version of specification (8) where we compare consumption levels across all retirement ages (rather than aggregating retirement ages into four groups). The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed. Individuals who retire at 65 are the reference category. The graph reports for all retirement age, the estimated coefficients α_j from specification (8), scaled by $\mathbf{E}_j[\tilde{C}_{it}]$, the average level of consumption of individuals who retire at 65 from the same cohort, age, family composition, income decile and career length at 55 group as the average individual retiring in age group j. The top panel starts with results from model (8) where only year and age fixed effects are included. The middle and bottom panels show the same estimated coefficients when sequentially adding controls for family composition, within-cohort deciles of average income between ages 52 and 55 and group of career length at 55 in the vector of controls $\mathbf{X}$.

Figure C-2: CONSUMPTION DIFFERENCES BY RETIREMENT AGE GROUPS: BY AGE AT WHICH CONSUMPTION IS OBSERVED

A. Consumption at Age 66

B. Consumption at Age 67

C. Consumption at Age 68

D. Consumption at Age 69
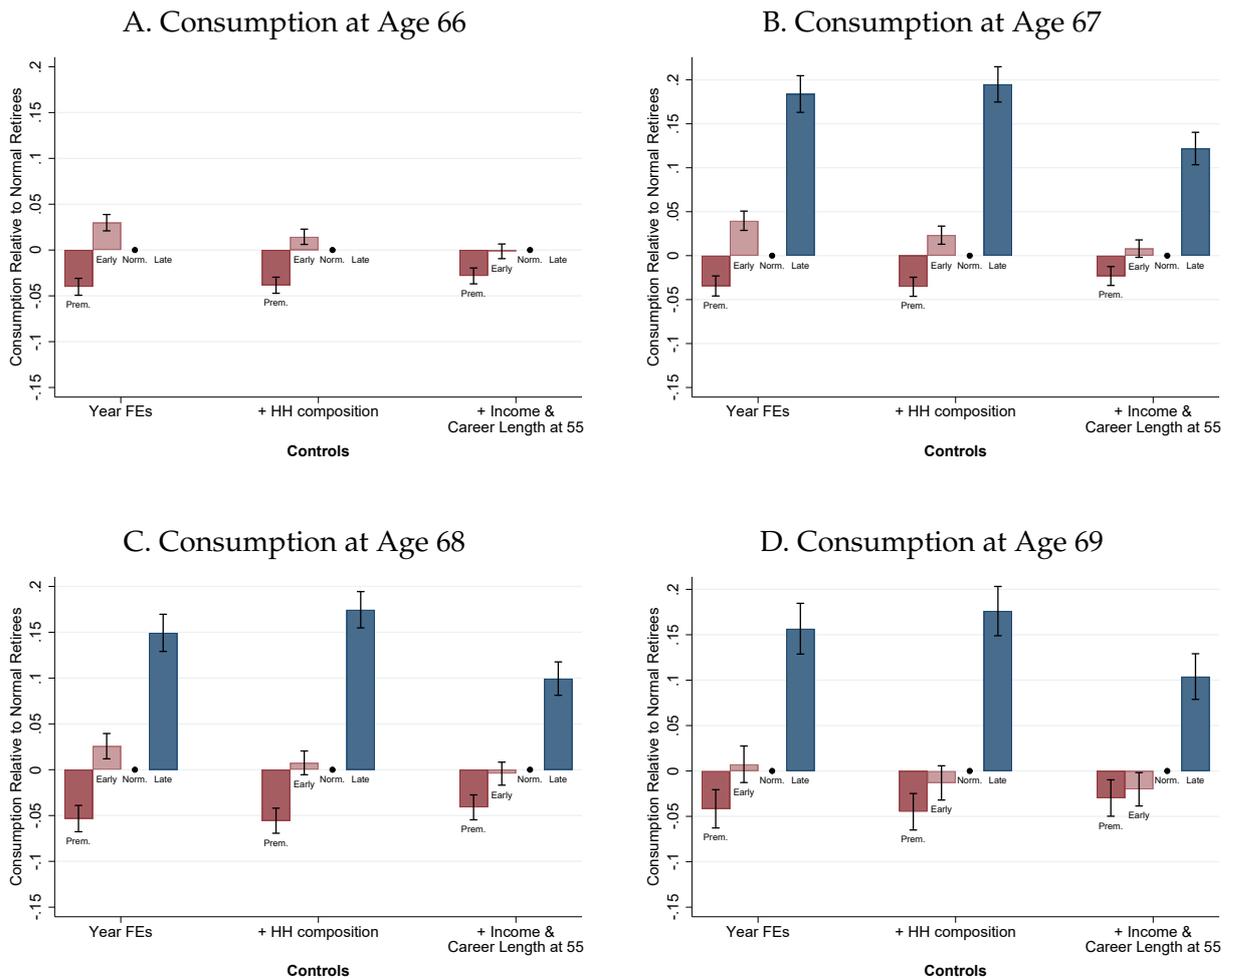
Notes: The figure shows that the consumption patterns hold irrespective of the age at which consumption is observed during retirement. We run regressions similar to specification (8), but separately for each age t. Because t is now fixed, we remove age fixed effects from the specification and control for year fixed effects γ_y. In effect, we compare consumption at age t of individuals retiring in different age groups *within the same cohort*. The graph confirms the very strong positive gradient of consumption with retirement age, at all ages at which consumption is observed.

Figure C-3: Consumption Differences by Retirement Age: Split By Household Structure

A. Couples (Married or Cohabiting)
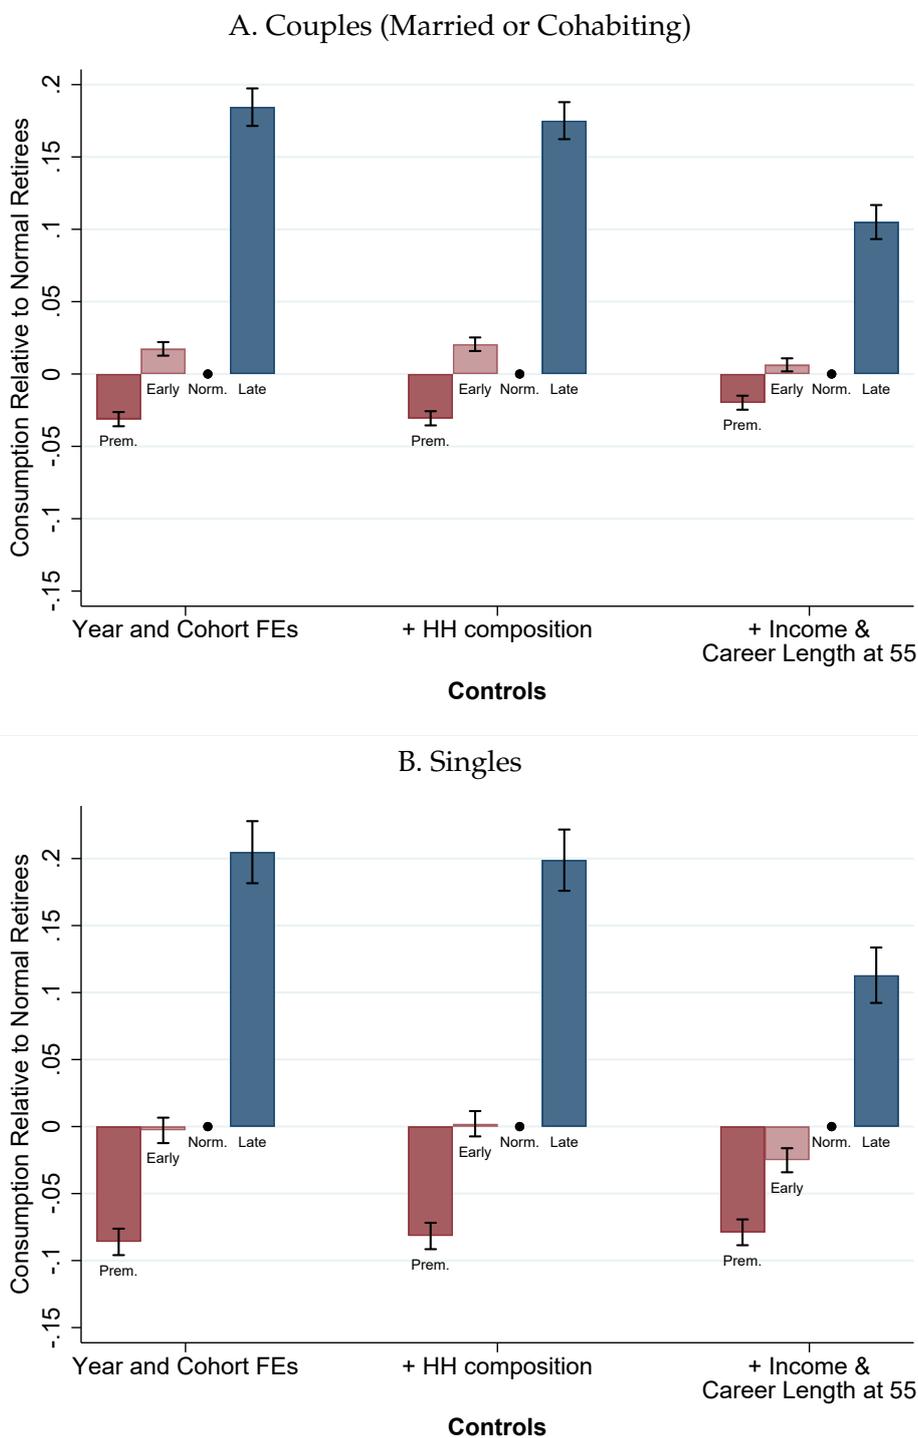
B. Singles

Notes: The figure reproduces estimates of consumption differences in retirement by retirement age group, similar to Figure 3 but splitting the sample between individuals who are single vs married/cohabiting at the time of retirement. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed.

Figure C-4: CONSUMPTION DIFFERENCES IN RETIREMENT ACROSS RETIREMENT AGE GROUPS: COMPARING RETIREES IN THE TOP AND BOTTOM DECILE OF ATP AT 55
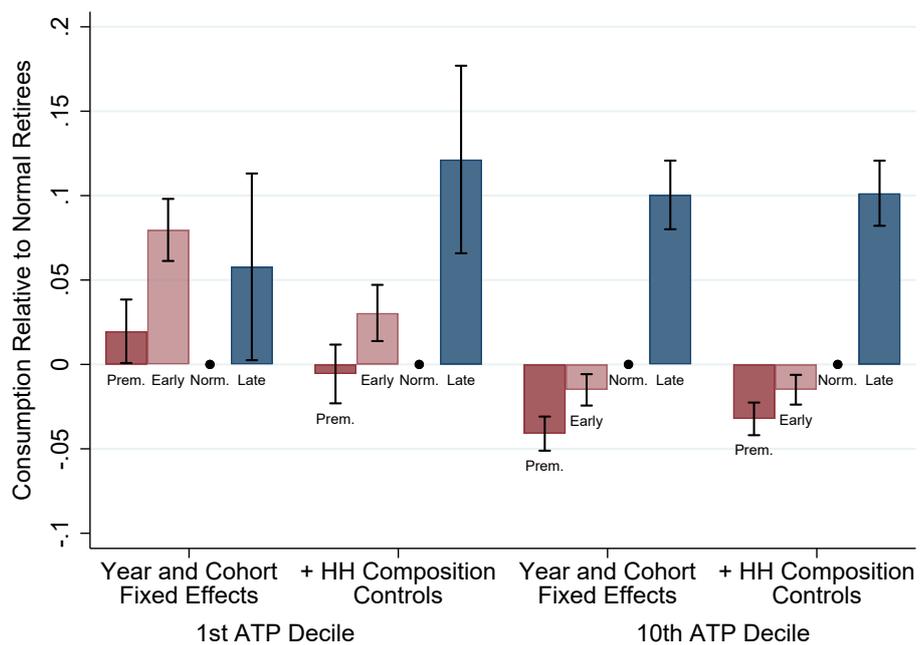
Notes: The figure reports consumption in retirement across individuals who retire at different ages relative to normal retirees. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed. Individuals are grouped into four retirement age categories: premature retirees ($56 \leq r \leq 60$), early retirees ($61 \leq r \leq 63$), normal retirees ($64 \leq r \leq 65$) and late retirees ($66 \leq r \leq 69$). Results are shown for individuals in the 1st and 10th ATP at 55 deciles, with only year and age fixed effects as well as with added controls for family composition.

Figure C-5: CONSUMPTION DIFFERENCES BY RETIREMENT AGE: ALTERNATIVE DEFINITION OF RETIREMENT

Notes: The figure documents consumption differences across retirement age groups using an alternative measure of retirement age that accounts for the time spent in UI or DI after an individual stops working. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed. Individuals are grouped into four retirement age categories using this alternative measure of retirement age: premature retirees ($56 \leq r \leq 59$), early retirees ($60 \leq r \leq 63$), normal retirees ($64 \leq r \leq 65$) and late retirees ($66 \leq r \leq 69$). Normal retirees are the reference category. The graph reports for all retirement age groups, the estimated coefficients α_j from specification (8), scaled by $\mathbf{E}_j[\tilde{C}_{it}]$, the average level of consumption of individuals who retire between 64 and 65 from the same cohort, age, family composition, income decile and career length at 55 group as the average individual retiring in age group j. We start, on the left hand side of the graph, with results from model (8) where only year and age fixed effects are included. The rest of figure shows the same estimated coefficients when sequentially adding controls for family composition, within-cohort deciles of average income between ages 52 and 55 and group of career length at 55 in the vector of controls **X**.

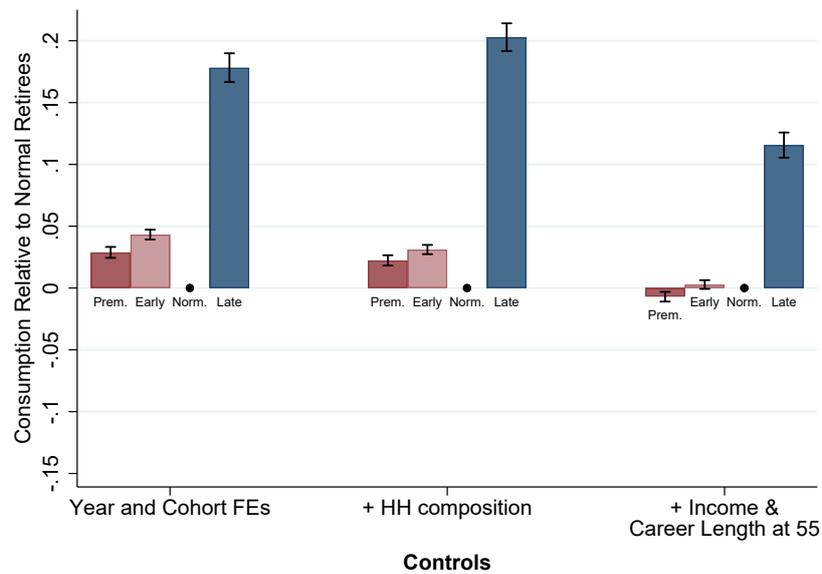Figure C-6: CONSUMPTION DIFFERENCES BY RETIREMENT AGE: CONTROLLING FOR PRE-RETIREMENT CONSUMPTION
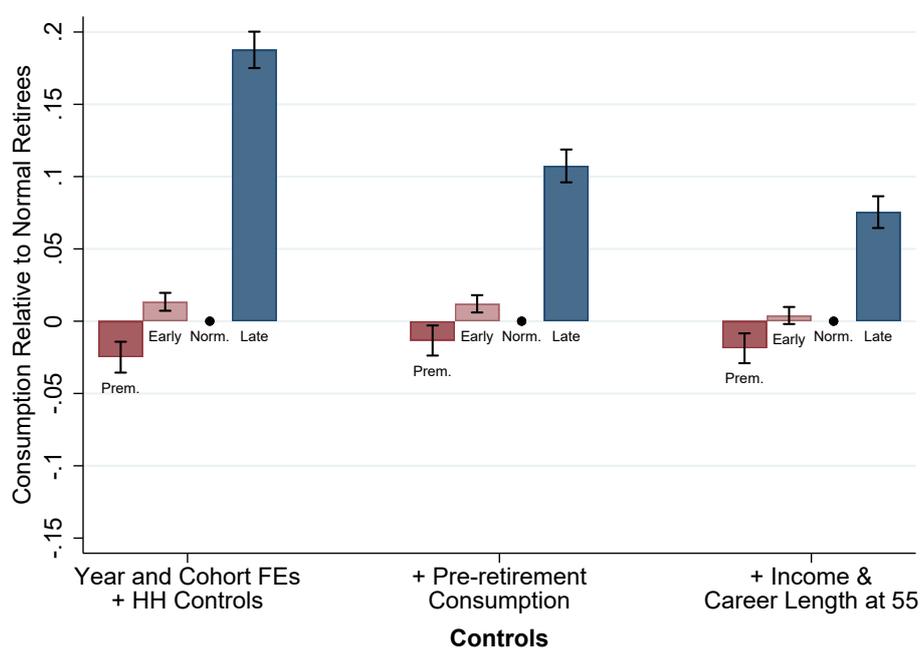
Notes: The figure depicts estimates of consumption differences in retirement by retirement age group, similar to Figure 3, but adding non-parametric controls for consumption before retirement. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed and for whom we also observe consumption two years before retirement. We start, on the left hand side of the graph, with results from the second specification from Figure 3, where year and age fixed effects and household structure controls are included. The rest of figure depicts these estimated post-retirement consumption differences when sequentially adding controls for within-cohort deciles of consumption two years before retirement, within-cohort deciles of average income between ages 52 and 55 and group of career length at 55 in the vector of controls **X**.

Table C-1: CONSUMPTION DIFFERENCES IN RETIREMENT BY RETIREMENT AGE: OB DECOMPOSITION

	Premature	Early	Normal	Late
	A. Mean residual consumption in retirement (SEK 2003)			
	201,886	211,054	204,713	260,670
	B. Difference with Late Retirees			
	-58,785	-49,617	-55,957	-
Fraction explained by:				
Consumption at r-2	0.18	0.19	0.23	-
	(0.01)	(0.01)	(0.01)	
Career Length at 55	0.06	0.03	0.03	-
	(0.01)	(0.00)	(0.00)	
Income at 52-55	0.19	0.22	0.24	-
	(0.01)	(0.01)	(0.01)	
Implied fraction explained by late career dynamics	0.56	0.57	0.50	

Notes: The table reports the results from a Oaxaca-Blinder (OB) decomposition of the differences in mean residual consumption among retirees. As in Figure C-6, the sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed and for whom we also observe consumption two years before retirement. First, consumption is residualized on a set of year and cohort fixed effects, a dummy for being married or cohabiting and a dummy for having children at home. Panel A reports average residual consumption in retirement by retirement age group. Then, the sample is split by retirement groups and residual consumption is regressed on dummies for within-cohort deciles of consumption 2 years before retirement, groups of career length at 55 and within-cohort deciles of average income between 52 and 55 years of age. Panel B reports the fraction of the difference in mean residual consumption that is accounted for by differences in endowments in the three control variables. That is to say, for each control variable $X_{it}^k = \left(d_{it}^{k,1} \quad d_{it}^{k,2} \quad \ldots \right)'$, where $d_{it}^{k,g}$ are dummies for the groups of X_{it}^k, we compute:

$$\frac{\left[\mathbf{E}_j(X_{it}^k) - \mathbf{E}_{Late}(X_{it}^k) \right]' \hat{\beta}_k^{Late}}{\mathbf{E}_j(\check{C}_{it}) - \mathbf{E}_{Late}(\check{C}_{it})}$$

where $\mathbf{E}_j$ denotes an expectation over individuals in retirement group j, $\check{C}_{it}$ is residualized consumption and $\hat{\beta}_k^{Late}$ is the vector of coefficient estimates for the base category (late retirees, in this case).

Decomposition of Consumption Expenditures at Age 60

Figure C-7: DECOMPOSITION OF CONSUMPTION EXPENDITURES AT AGE 60 BY RETIREMENT AGE

A. Premature Retirees

B. Early Retirees

C. Late Retirees
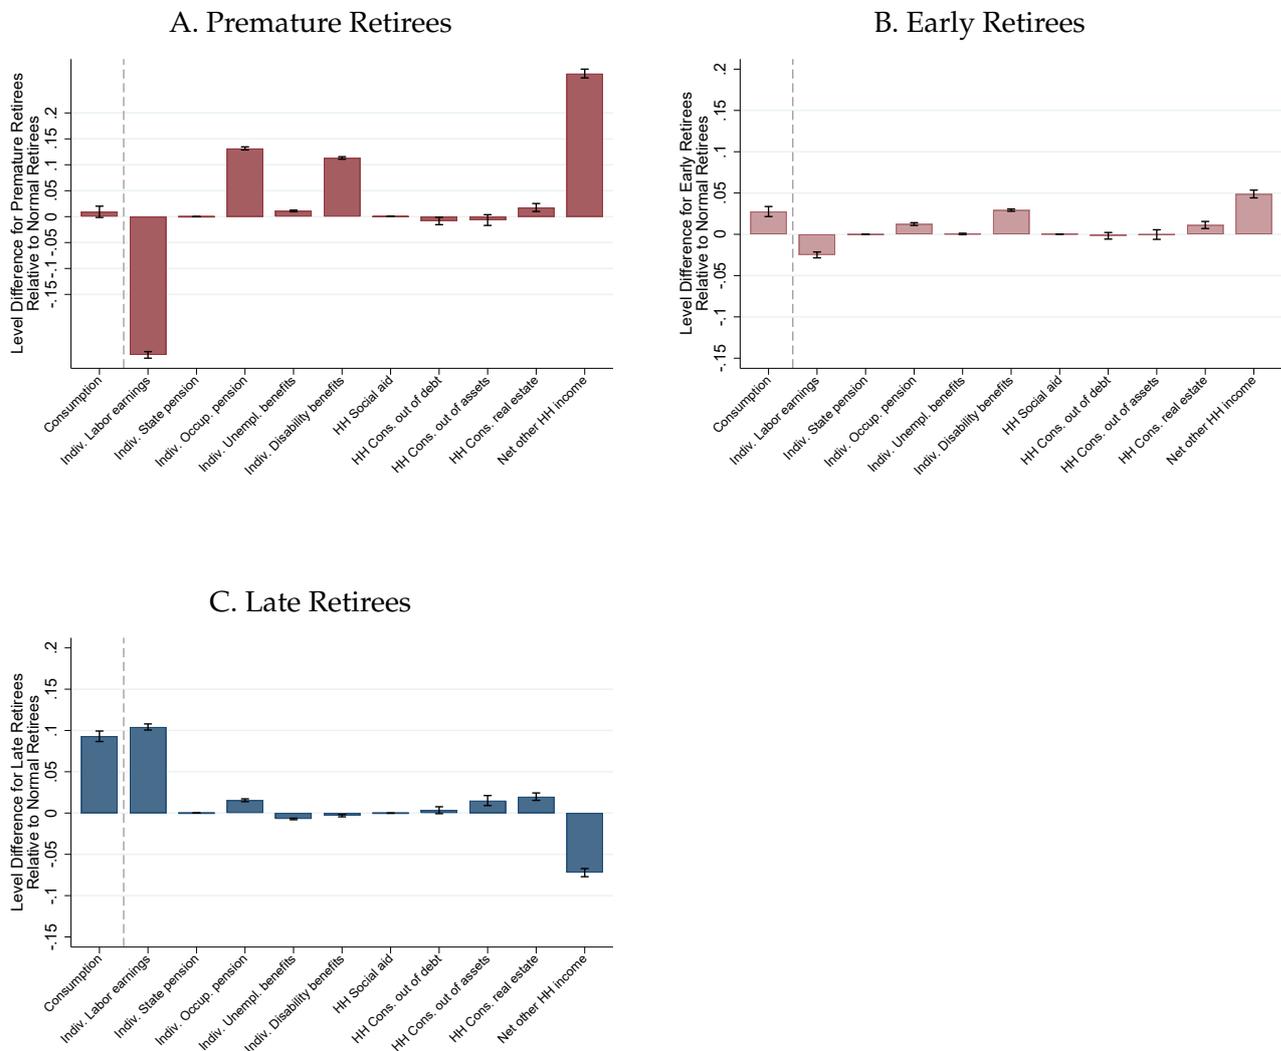
Notes: The figure decomposes consumption differences at age 60 across individuals who retire at different ages. The sample comprises all individuals from cohorts 1938 to 1943. Individuals are grouped into four retirement age categories: premature retirees ($56 \leq r \leq 60$), early retirees ($61 \leq r \leq 63$), normal retirees ($64 \leq r \leq 65$) and late retirees ($66 \leq r \leq 69$). We decompose our measure of household expenditures into a set of components that shed light on the consumption means available to individuals. These components include own income, (which we break down into own earnings, pensions, and other transfers such as UI, or DI), consumption out of debt, consumption out of assets, consumption out of real estate, and other household income (e.g. earnings from other members of the household, etc). We run specification (8) separately for each component evaluated at age 60, and report for all retirement age groups, the estimated coefficients α_j, using normal retirees as the reference category. As in Figure 3, the coefficients α_j are scaled by $\mathbf{E}_j[\tilde{C}_{it}]$, the average level of consumption of individuals who retire between 64 and 65 from the same cohort, age and family composition as the average individual retiring in age group j. All regressions include year and age fixed effects as well as controls for family composition.

Consumption Shares

Figure C-8: Differences in Consumption Shares during Retirement
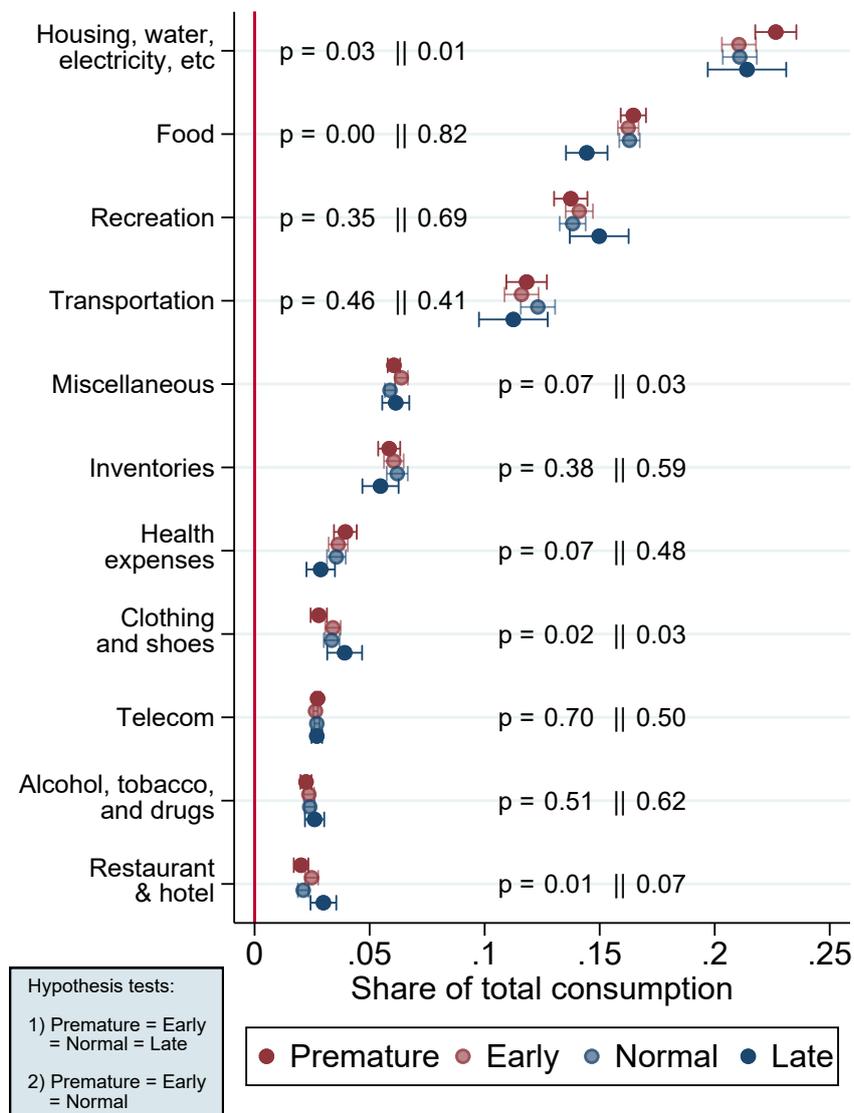
Notes: The figure examines consumption patterns in retirement. Using 3,373 observations from the HUT survey, total consumption is divided into 11 spending categories. The share of total consumption represented by each category is then regressed on a set of year and cohort fixed effects, a dummy for being married or cohabiting, a dummy for having children at home, dummies representing the retirement age groups and a constant. We plot the estimated conditional mean of the predicted consumption share by retirement age group. We also report p-values testing whether these conditional means are constant across retirement age groups. The first p-value corresponds to a joint test for the equality of the conditional means for the four retirement age groups, i.e. a joint test of equality of the three non-omitted retirement age groups and 0. The second p-value corresponds to a joint test for the equality of the conditional means for the first three retirement age groups (Premature, Early and Normal), i.e. a joint test for the equality of the first two non-omitted retirement age groups and 0.

Figure C-9: Change in Consumption Shares around Retirement
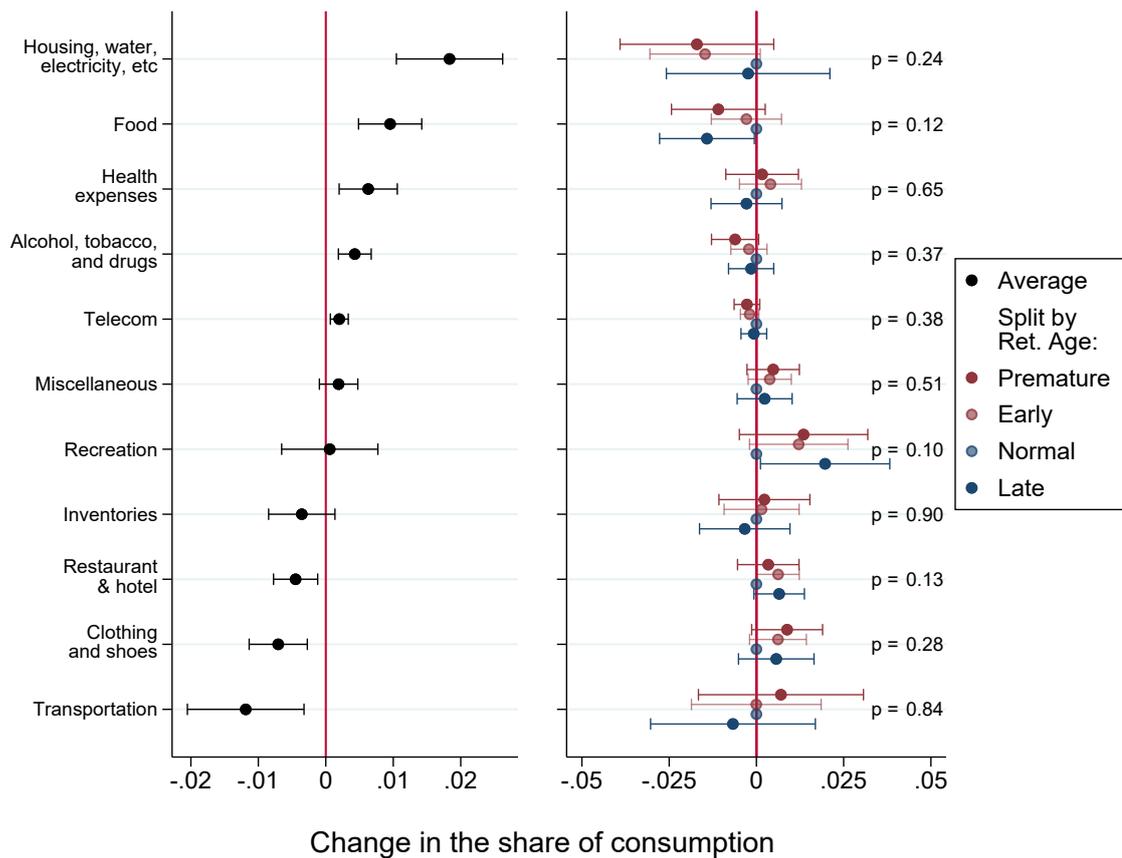
Notes: This figure examines changes in consumption patterns around retirement. Using 5,205 observations from the HUT survey, total consumption is divided into 11 spending categories. The share of total consumption represented by each category is then regressed on a set of year and cohort fixed effects, a dummy for being married or cohabiting, a dummy for having children at home retirement age group dummies, a retirement dummy, and, crucially, interactions of the retirement dummy with the the retirement age group dummies. On the right, we report the coefficients of the interaction between retirement and retirement age group, which can be interpreted as differences in the change of the shares relative to the change for Normal retirees (the baseline group). The p-values correspond to a test of equality of the partial effect of retirement for the four retirement age groups, i.e. a joint test of equality of the three non-omitted interaction terms and 0. On the left, we constrain the regression so that the partial effect of retirement is equal across retirement age groups (removing the interaction terms between retirement and the age groups) and report the coefficient on the retirement dummy.

Appendix D   Robustness of Consumption Patterns by Retirement Age Across Contexts

In this appendix, we explore the external validity, across contexts and data sources, of the consumption patterns by retirement age we documented in Sweden. We note of course that the consumption patterns across retirement age groups will depend on the policy environment (e.g. the steepness of the pension profile, the availability of other insurance mechanisms against consumption risk in old age, etc.) which differ across countries and over time. Most countries share very similar institutions (see OECD [2015, 2017, 2019]), with pension profiles that penalize early retirement and it is therefore interesting to investigate whether the broad patterns of consumption hold in these contexts as well.

One of the difficulty is of course the limited availability of data with both detailed consumption and retirement information. We use two surveys that contain such information: the Survey of Health, Ageing and Retirement in Europe (SHARE) for a large set of European Countries, which contains information on food consumption, and the Health and Retirement Study (HRS) for the US, which contains a broader measure of consumption.

Appendix D.1   Evidence from SHARE

The SHARE is a multidisciplinary and cross-national panel database of micro data on health, socio-economic status and family networks of about 140,000 individuals aged 50 and older. The survey took place in 2004, 2007, 2011, 2013, 2015 and 2017; it has a small panel structure and covers the 27 EU countries. It is harmonised with the US Health and Retirement Study (HRS). However consumption in the SHARE survey is only available for food items.

To make the analysis comparable to the analysis we conducted in Sweden, we restrict the SHARE sample to the cohorts born between 1938 and 1958, and to individuals aged between 50 and 75. We only keep for analysis countries that are repeatedly sampled since 2004, which leaves us with 11 countries: Austria, Belgium, Denmark, France, Germany, Greece, Italy, the Netherlands, Spain, Sweden, Switzerland and the United States.

We define retirement as the year an individual reports having stopped working for pay. In terms of methodology, we follow a similar approach as in our baseline analysis and regress consumption of individual i at age t living in country l on a series of dummies for retirement age, and we control for country fixed effects, year fixed effects and age fixed effects:

$$(20) \qquad C_{it}^l = \sum_j \alpha_j \cdot \mathbb{1}[r = j] + \gamma_y + \gamma_t + \gamma_l$$

In practice, we follow the same grouping of retirement age as in Sweden: we define as premature retirees individuals who retire at or before age 60, early retirees as individuals retiring between age 61 and 63, normal retirees as people retiring between 64 and 65, and late retirees for people who retire after 65. All results are expressed relative to the consumption level of the normal retirees.

In terms of aggregating results across countries, we run all regressions at the individual level with country fixed effects and report results for 3 weighting options: (i) the no weight option in which we do not include any weight in the regression (so all individual observations in the SHARE sample are given equal weight irrespective of the country population size or sampling frame); (ii) the population weight option uses weight corresponding to the sampling frame of each country in the survey, and reweights each individual weight so that the sum of weight in each country reflects a country's relative population size; (iii) finally the equal weight option (our preferred option) uses weight corresponding to the sampling frame of each country in the survey and reweights each individual weight so that the sum of weights in each country is the same (in other words, all countries are given equal weight in the regression).

Results In Figure D-1 below, we report estimates of the α_j coefficients for each retirement age group, scaled by $\mathbf{E}_j[\tilde{C}_{it}^l]$, the predicted consumption level from specification (21) when omitting the contribution of the retirement age group dummies. Specifically, $\mathbf{E}_j[\tilde{C}_{it}^l]$ corresponds to the average level of consumption of individuals who retire between 64 and 65 from the same country, cohort and age as the average individual retiring in age group j.

Results show that the overall patterns of food consumption by retirement age are very similar on average in the SHARE sample as the consumption patterns found in Sweden: there is a strong positive gradient, with the level of food consumption of premature retirees being significantly lower than that of late retirees. We also find evidence of non-monotonicity, with the level of food consumption of early retirees being slightly larger than that of normal retirees on average across the 12 countries in our sample.

We note however that the differences in consumption levels across retirement age groups are smaller overall in the SHARE survey than what we found in Sweden. We believe that this may be because the SHARE survey can only focus on food consumption, for which there is generally much less variance than for other types of expenditures. We also note that the small sample size within each country makes these estimates imprecise. And we turn for further investigations to the HRS data that has more information on consumption, and the largest sample size within the countries sampled in the SHARE/HRS data.

Figure D-1: FOOD CONSUMPTION LEVELS BY RETIREMENT AGE: SHARE DATA
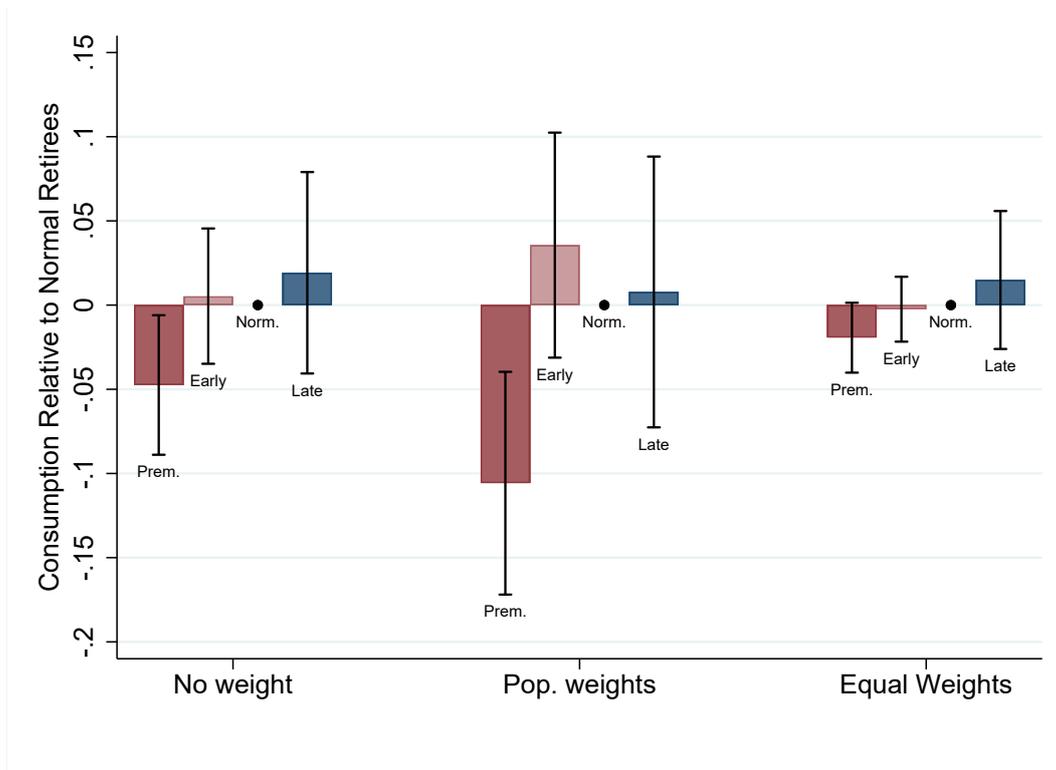
Notes: The figure documents differences in food consumption across retirement age groups. The sample comprises all individuals aged 50 to 75 from cohorts 1938 to 1958 who are observed in the SHARE data from Austria, Belgium, Denmark, France, Germany, Greece, Italy, the Netherlands, Spain, Sweden, Switzerland or the United States (HRS data). Individuals are grouped into four retirement age categories using this alternative measure of retirement age: premature retirees ($56 \leq r \leq 60$), early retirees ($61 \leq r \leq 63$), normal retirees ($64 \leq r \leq 65$) and late retirees ($66 \leq r \leq 69$). Normal retirees are the reference category. The graph reports for all retirement age groups, the estimated coefficients α_j from specification (20), where we control for age, year and country fixed effects.

Appendix D.2   Evidence from the US Using HRS Data

The HRS data has slightly richer information on consumption than the SHARE data, and a slightly larger sample size. This allows us to provide more detailed results for the US to investigate the external validity of the consumption patterns by retirement age found in the Swedish context.

The sample is composed of all individuals interviewed for the consumption module (CAMS) of the HRS. While the HRS takes place every two years since 1992, the CAMS modules happen every two years since 2001, making up 9 waves in total, and are composed of randomly selected HRS participants. In the final sample, we drop individuals for which consumption, age or the date of retirement are not observed. We are left with 13,498 observations, corresponding to 3,808 individuals and distributed across waves in the following way:

Wave	Nb of observations
2001	1,755
2003	1,524
2005	1,581
2007	1,738
2009	1,601
2011	1,534
2013	1,414
2015	1,278
2017	1,073

Consumption Measure in the HRS   The HRS special modules contain rich information about consumption. The following expenditure items are available:

- Automobiles: automobile or truck purchase, payments related to car (referred to as finance charges or interest/principal), vehicle insurance, gasoline, vehicule maintenance (parts, repairs and servicing);

- Household appliances: refrigerator, washer-dryer, dishwasher, television, computer, mortgage;

- Home cost: rent, property tax, homeowner's or renter's insurance, electricity, water, heating, telephone, cable and internet, housekeeping supplies, home repairs and maintenance, gardening and yard supplies, household furnishings and equipment;

- Food: food and beverages inside the home, dining and drinking out;

- Clothing and apparel;

- Personal care products and services;

- Health: health insurance, out-of-pocket cost of prescription and non-prescription medications, out-of-pocket cost of healthcare services, out-of-pocket cost of medical supplies;

- Hobbies/holidays: trips and vacations, tickets to movies/events, hobbies

- Other: contributions (to religious, educational, charitable or political organisations), gifts.

We focus on expenditure items that are reported in every wave. Excluded categories that do not appear in every wave are usually rather small: sport equipments, personal care products and services, gardening and yard supplies, home furnishings and equipment.

Consumption variables were originally expressed in nominal terms. We use CPI data and express all consumption in 2003 USD.

Retirement Age: Definition   The HRS survey allows to infer the date of retirement in several ways:

- It asks individuals to directly report the month and year in which they retire.

- In the HRS waves (every two years since 1992), respondents are asked to report their occupation, namely whether they are currently working for pay, unemployed, temporarily laid-off/sick, disabled, retired, or homemaker. Those option choices are not mutually exclusive and individuals are given the possibility to select themselves into several categories.

- In the CAMS waves (every two years since 2001), respondents are asked whether they are currently retired.

In order to be consistent with our definition of retirement in the Swedish context, we define retirement as a permanent switch to reporting one's occupation status as not working for pay. And retirement age is defined as the first year in which the individual does not report his occupation status as working for pay.

Methodology   We follow a similar methodology as in the Swedish context and regress household consumption C_{it} of individual i at age t in year y

$$(21) \qquad C_{it} = \sum_j \alpha_j \cdot \mathbb{1}[\mathrm{r} = \mathrm{j}] + \gamma_y + \gamma_t$$

In practice, we group retirement ages into two-years bins, and use individuals retiring between 64 and 65 as the reference category. We control for year fixed effects γ_y and age fixed effects γ_t, so that in effect, we compare consumption of individuals retiring in different age groups *within the same cohort, at the same age*, as in the main analysis. Figure D-2 below reports the estimated coefficients α_j for all retirement age groups, scaled by $\mathbf{E}_j[\tilde{C}_{it}]$, the predicted consumption level from specification (21) when omitting the contribution of the retirement age group dummies. As before, $\mathbf{E}_j[\tilde{C}_{it}^l]$ corresponds to the average level of consumption of individuals who retire between 64 and 65 from the same cohort and age as the average individual retiring in age group j.

Results   The patterns of consumption by retirement age revealed in Figure D-2 are similar to those found in the Swedish context (see for instance Figure C-1). First, we see a strong overall gradient of consumption with retirement age: "Premature" retirement (i.e. before age 60) is

associated with significantly lower consumption, while individuals who retire late (i.e. after 65) experience much larger levels of consumption, at the same age, than other individuals from the same cohort. Interestingly, we also detect the presence of non-monotonicity in the relationship between consumption and retirement age: this relationship is locally decreasing in the retirement age range 60 to 65.

The measure of expenditures used in the HRS is clearly not perfectly comparable to the measure we use in our main analysis: it is less comprehensive than the one we have in Sweden. But the comparison of results across these contexts and data sources is nevertheless very informative. Overall these results confirm that the large gradient in consumption level between individuals who retire very late vs very prematurely is a robust finding across contexts and data sources. Second, it also confirms that the non-monotonicity in the relationship between retirement age and consumption is also robust across contexts and data: for most people retiring between 60 and 65, there is no gradient, or if anything a negative gradient between consumption level and retirement age.

We should stress that the overall gradient found in the HRS data is bigger than the one we document in Sweden. There is more than a 40% difference in consumption levels at the same age between the premature and late retirees in the US (compared to a 15 to 20% difference in Sweden). This could be due to the presence of a steeper pension profile in the US compared to Sweden and the fact that insurance against shocks in late career (such as UI, and DI) is generally much less generous in the US than in Sweden. These results in turn suggest that the social marginal utility cost of increasing the steepness of the pension profile is much larger in the US than in Sweden.

Figure D-2: CONSUMPTION LEVELS BY RETIREMENT AGE IN THE US: HRS DATA
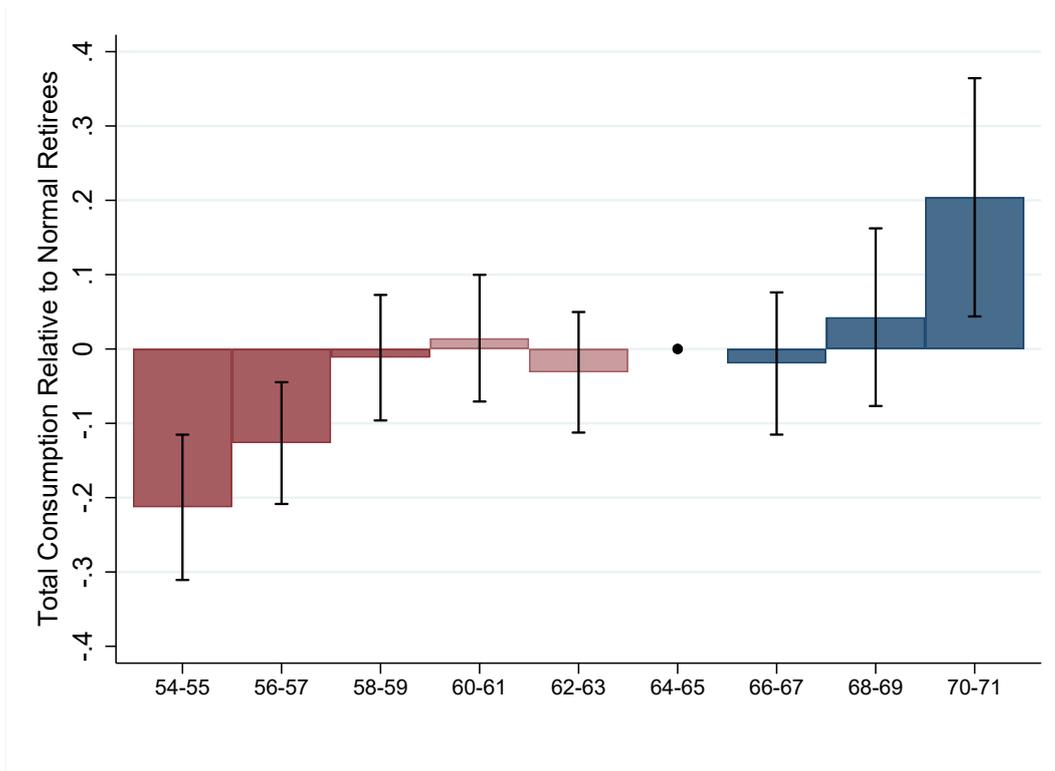
Notes: This figure documents how consumption differs across individuals who retire at different ages. The sample is composed of all individuals born between 1938 and 1958 interviewed for the consumption module (CAMS) of the Health and Retirement Study (HRS). The CAMS modules happen every two years since 2001, making up 9 waves in total, and are composed of randomly selected HRS participants. We drop individuals for which consumption, age or the date of retirement are not observed. Individuals are grouped into nine retirement age categories from 54 to 71. Retirement ages $64 - 65$ are the reference category. The graph reports for all retirement age groups, the estimated coefficients α_j from specification (21), scaled by $\mathbf{E}_j[\tilde{C}_{it}]$, the average level of consumption of individuals who retire between 64 and 65 from the same cohort and age as the average individual retiring in age group j.

Appendix E   Health & Work Longevity Risk

In this Appendix section, we further explore the role of health shocks and work longevity risk in shaping retirement consumption. In practice, we provide details on (i) the health data we are using to measure health status, (ii) the differences in health status in retirement across retirement age groups, and (iii) the dynamics of health around retirement.

Health Data

We use two surveys with detailed information on health outcomes and health expenditures. The first is the living condition survey (ULF) which contains various health measures for a representative sample of approximately seven thousand households, every year from 1997 to 2011. These measures include both subjective, such as self-reported illnesses, pain or reduced work capacity, as well as objective outcomes (number of visits to a physician in the last 12 months, body mass index, etc).[48] The second survey is the household finance survey (HEK), which samples an average of 30k individuals every year, and is also available from 1997 to 2011. The survey contains very precise information on health-related expenditures (number of visits to a doctor, to a physiotherapist, expenditures on pharmaceuticals, on outpatient care, etc).[49]

Both surveys are repeated cross-sections, but can be matched at the individual level with the administrative registers. In practice, this means that we observe for each individual surveyed in ULF and HEK their full (i.e. past and future) labor market and pension histories, consumption, etc. This allows us to investigate health dynamics around retirement using pseudo-panel techniques.

The literature on the impact of health on retirement has long recognized the potential measurement issues, leading to attenuation bias, in using only a specific subset of objective measures of health, as they may only partially capture the overall health status of an individual (Bound [1991], Stern [1989]). And while subjective measures may address these measurement issues, they can also be prone to justification bias (Butler et al. [1987]). To deal with these concerns, we follow Blundell et al. [2021]. We build, for each survey, a composite index of health by extracting the principal component of all objective and subjective measures available in the survey.

Table E-1 provides descriptive statistics on the samples from the ULF and HEK surveys that we match to our administrative data. To maximize power, we focus on cohorts 1938 to 1950. The table compares individuals matched in the ULF and HEK samples, to all individuals from our baseline sample of retirees. The table shows that the distribution of age at retirement is very similar across samples, and so are demographic and pension characteristics.

The table also reports descriptive statistics for the various health proxies that we combine into two health indices, by extracting their first principal components. Measures from the HEK (which is a household finance survey) are mostly objective measures of health expenditures. We use the following variables:

---

[48]This study is similar to the SILC survey conducted within the European Union.

[49]Importantly, the survey does not only report out-of-pocket expenditures, but also all expenditures that are directly taken care of by private and public health insurance.

- BANTGYM: Number of visits to a physiotherapist in last 12 months

- BANTLAK: Number of visits to a doctor in last 12 months

- BFRIMED: dummy for having access to free pharmaceuticals. When expenditures on pharmaceuticals reach a certain threshold (around 2000SEK per year) individuals become eligible to free pharmaceuticals.

- UMED: Pharmaceutical expenditures (under the cap).

- BFRISJU: a dummy for having access to free outpatient care. Similarly, when expenditures on outpatient care reach a certain threshold (around 1200SEK per year) individuals become eligible to free outpatient care.

- USJUKA: Total out-of-pocket expenditures for healthcare (excl. rehab) in last 12 months.

- UFORBR: Expenditures for assistive technology (e.g. motorized wheelchair, etc.) UHJALP: Expenditures for renting of assistive technology

In the ULF data, we have both subjective and objective measures of health. We extract the principal component from a Principal Components Analysis (PCA) on the following variables: Number of visits to a physician in the last 12 months, a dummy for individuals reporting having a long term / chronic illness, the number of long term illnesses reported, a dummy for reporting having serious health difficulties and/or pain, a dummy for having reduced work capacity, and the body mass index.

We create two health indices corresponding to the first component extracted from a PCA on these two sets of variables, and we then standardize both indices.

Table E-1: Descriptive Statistics: Health Information From HEK & ULF Surveys

	Retirement Sample		Retirement x HEK Sample		Retirement x ULF Sample	
	Mean	(s.d.)	Mean	(s.d.)	Mean	(s.d.)
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
I. Retirement						
Premature Retiree	23.81 %		18.09%		18.53%	
Early Retiree	25.68 %		26.56%		26.09%	
Normal Retiree	34.60 %		38.20%		37.92%	
Late Retiree	15.91 %		17.16%		17.46%	
Age at Retirement	62.91	(3.10)	63.27	(2.87)	63.25	(2.93)
II. Demographics						
Cohort	1940.67	(1.73)	1944.06	(3.54)	1943.91	(3.46)
Fraction Men	49.29 %	(50.00)	48.77%	(49.99)	48.62%	(49.98)
Married at 59	66.86 %	(47.07)	73.71%	(44.02)	66.85%	(47.08)
Kid at Home at 59	17.65 %	(38.12)	21.41%	(41.02)	18.78%	(39.06)
And Kid < 18	3.48 %	(18.33)	4.48%	(20.69)	3.84%	(19.22)
Post-Secondary Edu.	24.67%	(43.11)	30.21%	(45.92)	28.71%	(45.25)
III. Income and Wealth at 59, SEK 2003(K)						
Total Earnings	209	(160)	240	(173)	233	(156)
Net Wealth	777	(2339)	955	(1819)	876	(1529)
Bank Holdings	84	(312)	105	(264)	95	(210)
Portfolio Value	248	(1648)	289	(1252)	256	(1059)
Consumption	201	(534)	239	(844)	225	(529)
IV. Health (HEK)						
Visited Physio.			15.89%	(36.56)		
No. Physio. Visits			1.68	(5.28)		
Visited Doctor			68.38%	(46.50)		
No. Doctor's Visits			2.89	(3.83)		
Free Pharmaceuticals			25.83%	(43.77)		
Pharm. Expenses			746.2	(762.30)		
Free Outpatient Care			23.27%	(42.26)		
Healthcare Expenditure			366.40	(553.50)		

Table E-1: Descriptive Statistics: Health Information From HEK & ULF Surveys

	Retirement Sample		Retirement x HEK Sample		Retirement x ULF Sample	
	Mean	(s.d.)	Mean	(s.d.)	Mean	(s.d.)
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Assistive Tech. Exp.			5.50	(95.70)		
Ass. Tech. Rent Exp.			6.40	(203.50)		
V. Health (ULF)						
Visited Physician					38.61%	(48.69)
Has Long-Term Illness					54.75%	(49.78)
No. of LT Illnesses					.93	(1.13)
Difficulties/Pain					16.21%	(36.86)
Reduced Work Cap.					10.15%	(30.19)
Body Mass Index					256.87	(36.46)
N (Unique Individuals)	418,033		19,416		7,022	
Cohorts	[1938,1943]		[1938,1950]		[1938,1950]	

Selection on Health Across Retirement Age Groups

We first document how health and life expectancy differs across retirement age groups.

We start by running specification (8), replacing consumption on the left-hand side by two indices for bad health: the first corresponds to the standardized principal component extracted from all health outcomes available in the HEK survey, and the second index is similarly constructed based on all health variables from the ULF survey (see above). We include the same cohort and age fixed effects and controls as specification (8) above, so we effectively compare the health in retirement of individuals of the same cohort, and at the same age, who retired at different ages. On the right of Panel B, we focus on differences in "life expectancy" by using as an outcome a dummy for having died by age 70, or by age 75.[50]

For all outcomes, we document a very steep negative health gradient over retirement ages. That is, earlier retirement is strongly associated with having significantly worse health. This effect appears particularly strong for premature retirees: their health, measured by our bad health indices, is between .5 and .75 standard deviations worse than that of late retirees. Premature retirees are also almost 14 percentage points more likely to have died by age 75 than late retirees. Interestingly, we do not find any significant non-monotonicity for health outcomes: early retirees do not enjoy better health status or longer life expectancy than normal retirees

---

[50]These results rely on a specification similar to (8), although we now only have one observation per individual: as a result, we drop age fixed effects, and only include cohort fixed effects, as well as controls for family structure.

despite being wealthier and more likely to be female. In Figure E-1, we report estimates separately for each available health outcomes composing our two health indices. Results confirm the existence of the same strong negative gradient for all health measures, irrespective of their subjective or objective nature.

Health Dynamics Across Retirement Age Groups

Are these differences in health during retirement due to late career health shocks? A large literature has indeed argued that health shocks are a significant determinant of retirement and represent an important part of work longevity risk. Or are these differences persistent health heterogeneity that preexisted retirement? Could there even be reverse causality, i.e. earlier retirement causing a worsening of health status?

To shed light on these questions, we investigate how health dynamics around retirement varies across retirement age groups. We adopt a similar methodology as in our analysis of consumption dynamics in Figure 4, and compare health dynamics across retirement age groups. This allows us to check whether the differences in health outcomes when retired documented in Figure 7 pre-date retirement, and whether they are caused by early retirement (Kuhn et al. [2018], Fitzpatrick and Moore [2018], Bozio, Garrouste and Perdrix [2021]) We regress health outcomes of individual i at age t on dummies for belonging to each of the four retirement age groups interacted with dummies for being at event time $e = t - r$ relative to retirement:

$$(22) \qquad H_{it} = \sum_{j}\sum_{k} \alpha_{jk} \cdot \mathbb{1}[\mathrm{r} = \mathrm{j}] \cdot \mathbb{1}[\mathrm{e} = \mathrm{k}] + \gamma_y + \gamma_t + \mathbf{X}'\beta + v_{it}.$$

Due to the limited sample size of the health surveys, we group event times e by bins of 2 years, from 6 years before to 5 years after retirement and we report for each retirement age group the sequence of estimated coefficients $\hat{\alpha}_{re}$ around the event of retirement. We control in the regression for a series of cohort and age fixed effects, to account for the cohort and age profiles of health outcomes, as well on the same vector $\mathbf{X}$ of baseline controls (i.e. controls for household structure).

Figure E-2 Panel A reports the results from specification (22) where we use our bad health indices as an outcome, pooling both HEK and ULF surveys together.[51] The graph indicates the existence, in the pre-retirement period, of a significant gradient in health across retirement age groups. Premature retirees have a bad health index around .25 standard deviations higher than other retirees already five years prior to retirement. But we also see a clear fanning out of health outcomes just around retirement, driven by a significant worsening of the health of premature retirees. As a result, the post-retirement differences in health between premature retirees and the other three groups are twice as large (around .5 standard deviation in our bad health index) as their pre-retirement level. Interestingly, there is no significant variation in

---

[51]In practice, this means that we run specification (22) on the combined ULF and HEK samples, with H_{it} being the standardized first principal component from the ULF health outcomes if individual i is observed in the ULF sample, and H_{it} being the standardized first principal component from the HEK health outcomes if individual i is observed in the HEK sample. In Appendix Figure E-3, we report the results where instead of pooling the data, we run separate regressions on the ULF and HEK samples.

health around retirement for early, normal and late retirees (once controlling for the age profile of health).

Panel B confirms these dynamic patterns, using as an outcome the fraction of individuals reporting that they are experiencing pain. The graph shows that premature retirees have a 5 percentage points higher probability of experiencing pain 5 years prior to retirement compared to other retirees. But this probability increases steadily up to retirement, at which point it is 15 percentage points larger than for the other three groups, and persists at this high level after retirement. Again, we find no significant evolution of the probability to report pain around retirement for early, normal and late retirees. Appendix Figures E-3 and E-4 show that these dynamic health patterns replicate across various health outcomes, such as the fraction experiencing reduced work capacity, or the fraction reporting retiring due to health reasons. Overall, these results provide evidence that premature retirees (and to a smaller extent early retirees) experience significant negative health shocks around retirement, with persistent effects throughout retirement. This in turn implies that a large fraction of health differences when retired between premature (and to a lesser extent early) retirees and other retirees are due to negative health shocks experienced just around retirement. We also note that the evidence displayed here alleviates concerns about potential reverse causality in the relationship between retirement age and health during retirement. If reverse causality was at play, that is if health differences in retirement across groups were driven by the absence of work in old age being detrimental for health, we would expect to observe a (potentially gradual) decrease in health, similar for all groups, after retirement. To the contrary, we observe that the degradation of health happens entirely prior to retirement, and is highly heterogeneous across groups.

To summarize, premature retirees seem to experience negative consumption shocks just prior to retirement and these correlate strongly with proxies for the incidence of work longevity risk such as health shocks. This suggests that flatter pension profiles offer particularly valuable insurance against work longevity risk.

Furthermore, our results suggest that health shocks affect the timing of retirement primarily for premature retirees, and not so much for the rest of the population. This reconciles the results from Blundell et al. [2016b] that health dynamics explain only a limited part of the overall distribution of the timing of retirement and from Gustman and Steinmeier [2018]) that they are particularly strong for the people who stop working prematurely. Indeed, if there is no significant correlation between health dynamics and the timing of retirement for most retirees (i.e. for the large fraction of the population that retires after 63), the sensitivity of labor supply to health in old age is also highly heterogeneous.

Figure E-1: Differences in Health Status by Retirement Age: Separate Estimates For Each Component of HEK and ULF Bad Health Indices

A. ULF Survey Outcomes
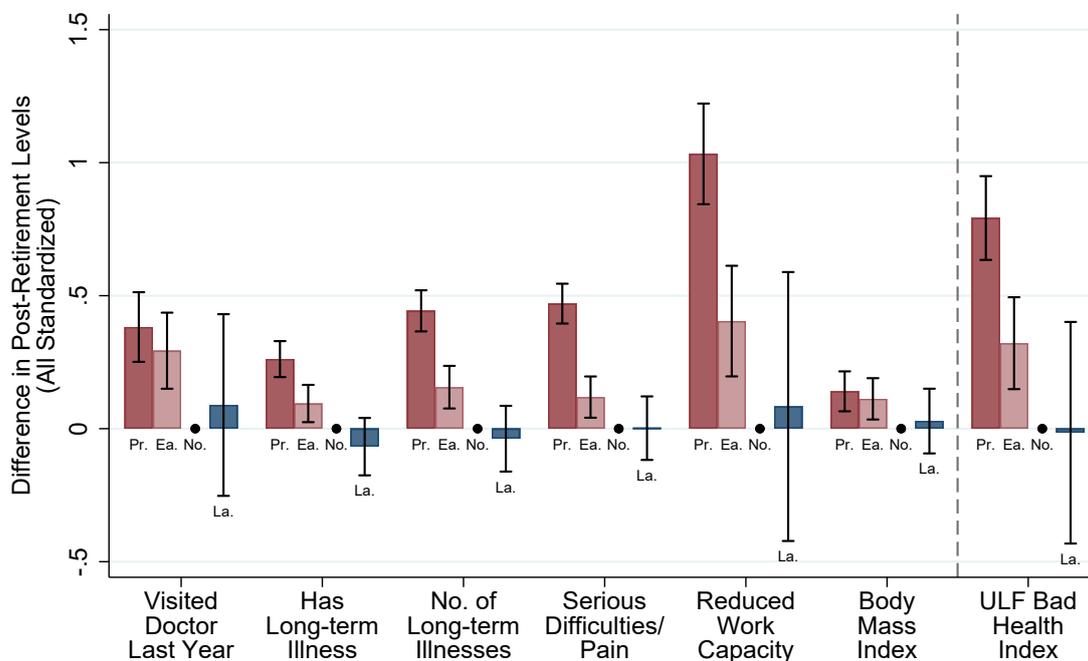
B. HEK Survey Outcomes
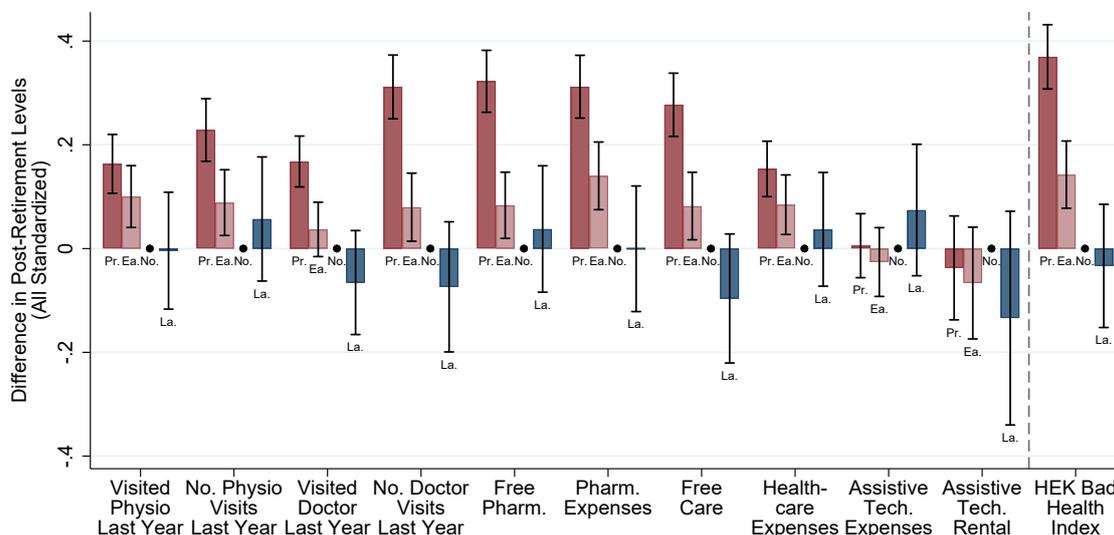
Notes: The figure documents differences in health outcomes across retirement age groups. The sample comprises all individuals from cohorts 1938 to 1943 who are observed either in the ULF or HEK surveys, and who are retired at the time of the survey. Individuals are grouped into four retirement age categories using our measure of retirement age: premature retirees ($56 \leq r \leq 60$), early retirees ($61 \leq r \leq 63$), normal retirees ($64 \leq r \leq 65$) and late retirees ($66 \leq r \leq 69$). Normal retirees are the reference category. The graph reports for all retirement age groups, the estimated coefficients α_j from specification (8), where we control for age and cohort fixed effects, as well as controls for family composition in the vector of controls **X**. All outcomes are standardized.

Figure E-2: HEALTH DYNAMICS AROUND RETIREMENT, BY RETIREMENT AGE GROUP
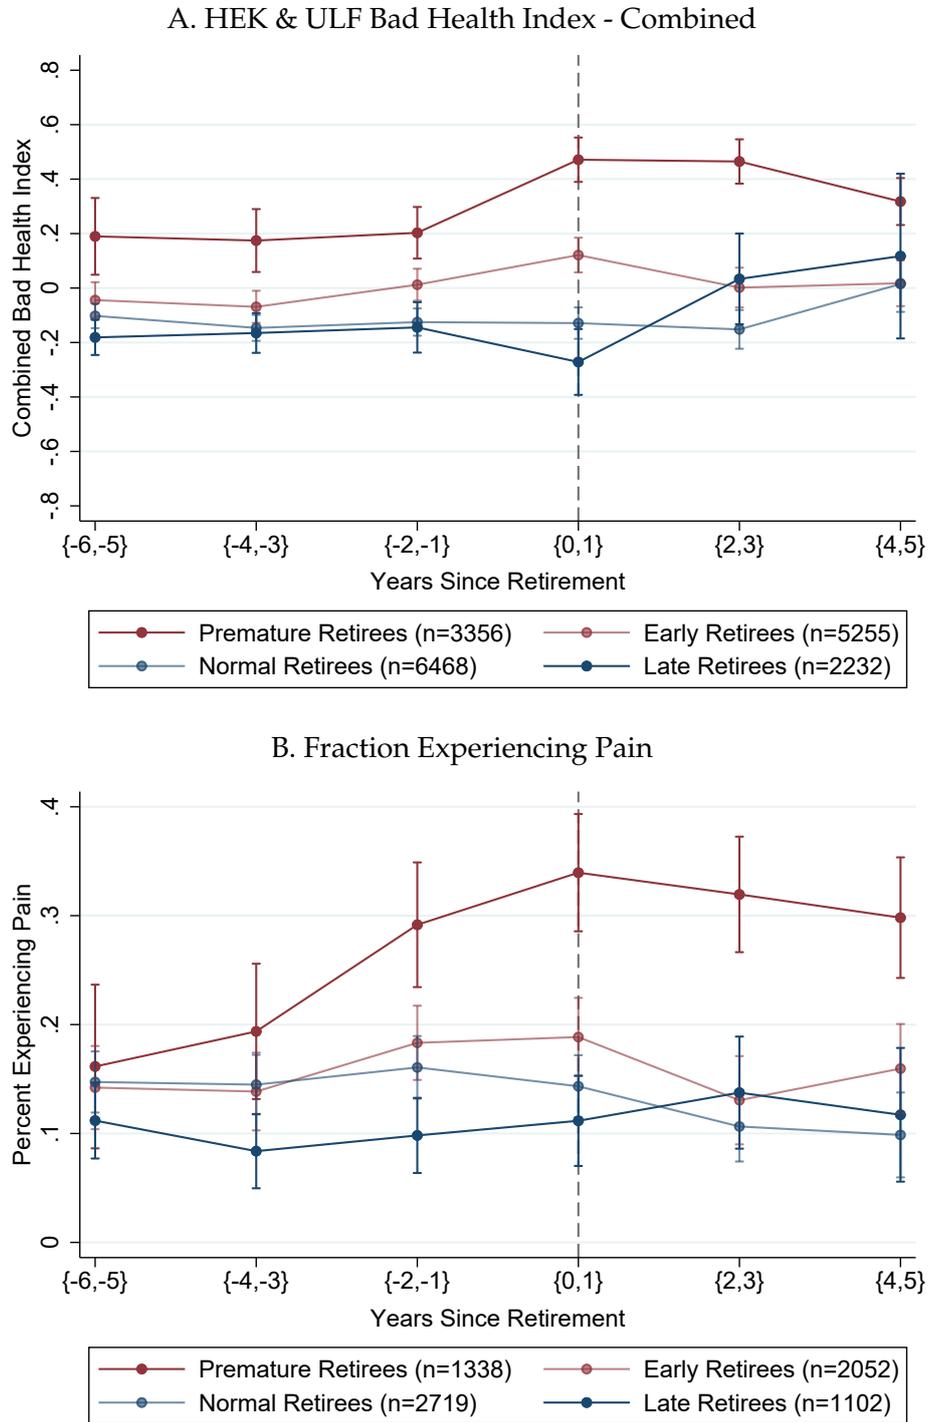
A. HEK & ULF Bad Health Index - Combined

B. Fraction Experiencing Pain

Notes: The figure documents heterogeneity in health dynamics around retirement, by retirement age group. Both panels report, for each retirement age group, the sequence of estimated coefficients $\hat{\alpha}_{re}$ from specification (22). where we control for cohort and age fixed effects and on the usual vector **X** of our baseline controls for household structure. Panel A uses our bad health indices as an outcome, pooling both HEK and ULF surveys together. Panel B uses as an outcome the fraction of individuals reporting that they are experiencing pain.

Figure E-3: HEALTH DYNAMICS AROUND RETIREMENT BY RETIREMENT AGE GROUP: HEK AND ULF BAD HEALTH INDICES
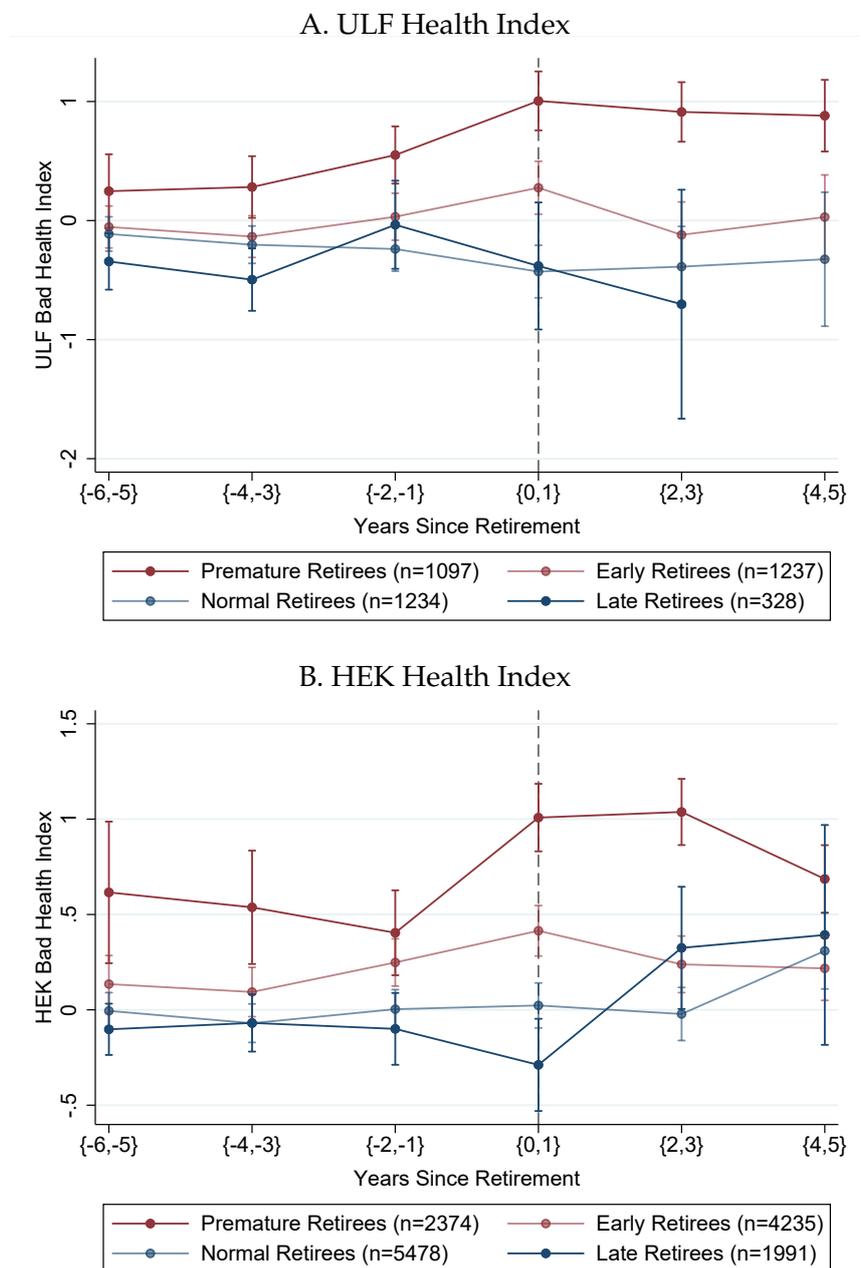
A. ULF Health Index

B. HEK Health Index

Notes: The figure documents heterogeneity in health dynamics around retirement, by retirement age group. Both panels report, for each retirement age group, the sequence of estimated coefficients $\hat{\alpha}_{re}$ from specification (22). where we control for cohort and age fixed effects and on the usual vector **X** of our baseline controls for household structure. Panel A uses the ULF bad health index as an outcome. Panel B uses the HEK bad health index as an outcome.

Figure E-4: HEALTH DYNAMICS AROUND RETIREMENT BY RETIREMENT AGE GROUP: HEK AND ULF BAD HEALTH INDICES
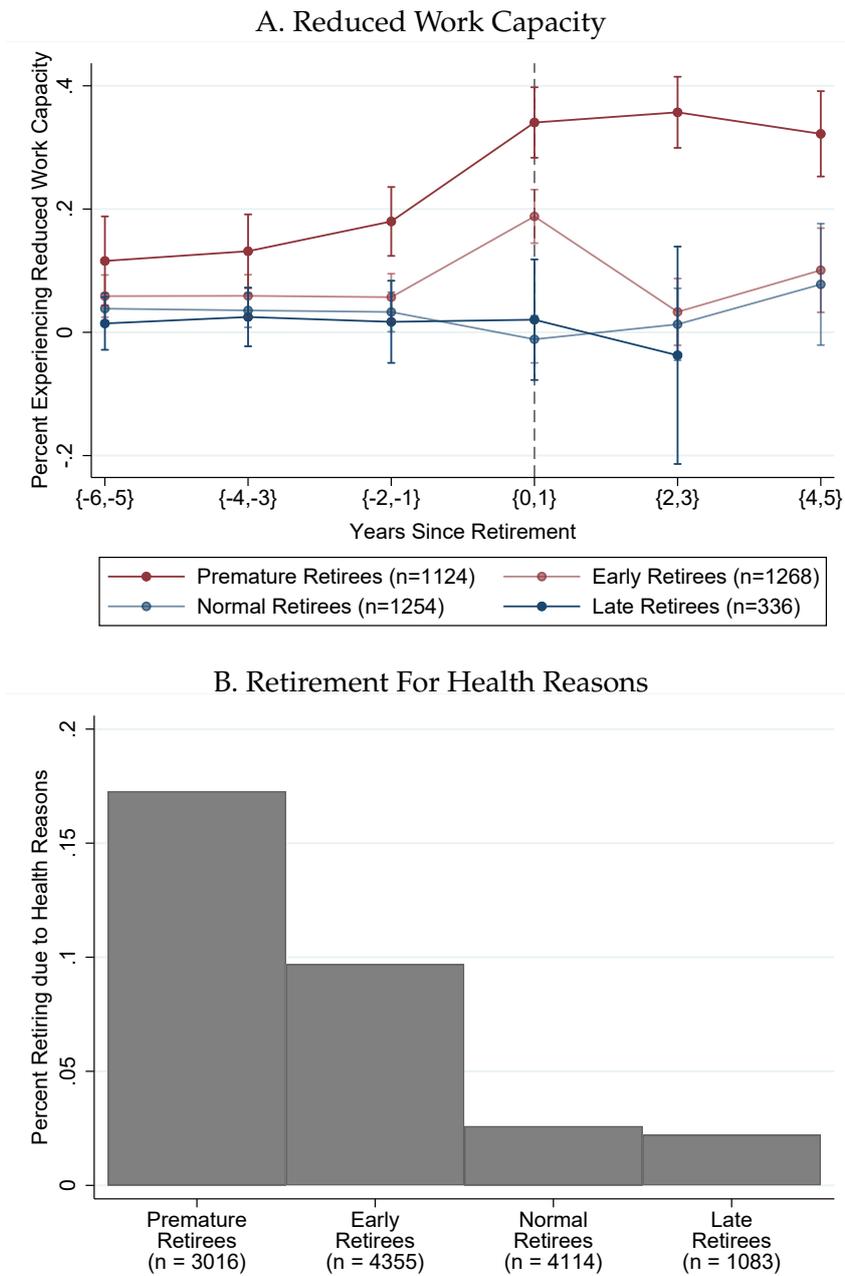
A. Reduced Work Capacity

B. Retirement For Health Reasons

Notes: The figure documents heterogeneity in health dynamics around retirement, by retirement age group. Panel A reports, for each retirement age group, the sequence of estimated coefficients $\hat{\alpha}_{re}$ from specification (22) similar to Figure E-2 where we use the fraction reporting reduced work capacity in the ULF survey as an outcome. In Panel B, we report the fraction of individuals reporting that they retired due to health reasons in the ULF survey, by retirement age groups.

Appendix F   Marginal Propensities to Consume

The marginal propensity to consume (MPC) out of a wealth shock when retired is another informative empirical moment that can be used to infer the social marginal cost of changes to retirement pension design. As shown in Landais and Spinnewijn [2021], differences in MPCs across retirement groups can capture differences in the liquidity value that pension benefits bring.

Appendix F.1   Mapping from MPC's to SMU's

Following Landais and Spinnewijn [2021], we approximate the ratio of social marginal utilities relying on the difference in marginal propensities to consume when retired. The higher the marginal propensity to consume for individuals with features x relative to those with features x', the higher the cost of making the pension profile more generous towards individuals with features x'.

MPC Implementation. *Assuming $c(\pi_{i,t}) = c_{x,t}$, $\zeta(\pi_{i,t}) = \zeta_{x,t}$ and $\omega_i = \omega_x$ for $x(\pi_{i,t}) = x$ and this for any i, x, t, and, in addition, assuming both $\frac{\partial u(c_{x,t}, \zeta_{x,t})}{\partial \zeta_{x,t}}$ and the relative curvature in preferences over c and ζ to be similar across retirement ages, we can approximate*

$$(23) \qquad \frac{SMU_{x,t}}{SMU_{x',t}} \cong \frac{\omega_x}{\omega_{x'}} \times \frac{\frac{mpc_{x,t}}{1 - mpc_{x,t}}}{\frac{mpc_{x',t}}{1 - mpc_{x',t}}},$$

where $mpc_{x,t} = \frac{dc_{x,t}}{dy_{x,t}}$.

This approximation relies on the relative curvature of utility over consumption and the resources used to smooth consumption at the margin (e.g., future consumption, household earnings) being similar across individuals with different features x.[52] More details are provided in Appendix G and in Landais and Spinnewijn [2021].

The value of this alternative implementation is again twofold. First, differences in the marginal propensity to consume reflect differences in the shadow price of consumption: the higher this price, the higher the propensity to consume out of an exogenous increase in income. By considering the MPCs, we thus narrow our welfare focus further on the liquidity value that pensions provide.[53] Second, by using yet another alternative consumption moment we again rely on different implementation assumptions. The main advantage of this MPC approach is that it does not require knowledge about the curvature in consumption preferences γ itself, but only on how preference curvatures differ across beneficiaries. The preference parameter γ is crucial for translating consumption differences into differences in marginal utilities in the first two implementations, but generally hard to estimate empirically (see Chetty and Finkelstein [2013]).

---

[52]This property holds when individuals have CARA preferences over consumption and use future consumption to smooth current consumption at the margin. However, it can be violated when individuals with different features x are more or less likely to use bequests at the margin and preferences over bequests are less curved than preferences over future consumption as in Lockwood [2018].

[53]We can expect this to provide a lower bound on the consumption smoothing gains as individuals who face a higher shadow price of consumption may do so because they already need to rely more on alternative resources to smooth their retirement consumption.

Appendix F.2 Empirical Strategy: Quasi-Random Wealth Shocks

We now turn to estimating MPCs across retirement age groups to provide another measure of the social cost of incentivizing later retirement. The challenge in measuring heterogeneity in MPCs lies in finding a credibly exogenous source of variation in income or wealth that applies similarly across the population of retirees.

We use variation in individuals' financial wealth coming from quasi-random shocks to the price of stocks that individuals hold in their portfolio, as in Di Maggio, Kermani and Majlesi [2020] and Andersen, Johannesen and Sheridan [2021].

Sample Construction & Validation of Empirical Strategy

We start from the KURU register, which has disaggregated information over the period 1999 to 2007 on all quantities of stocks, by ISIN number, held by individuals outside of mutual funds. We then match this data with information from the financial company SIX on prices of all listed stocks at the Stockholm stock exchange for each ISIN over the entire period 1990-2015. For each individual i, we define the passive capital gains on her portfolio in year $t + 1$ as:

$$\text{KG}_{i,t+1} = \sum_j (p_{j,t+1} - p_{j,t}) \cdot a_{ijt} = \sum_j \Delta p_{j,t+1} \cdot a_{ijt}$$

where a_{ijt} is number of stocks of company j held by individual i on 31st of December of year t and Δp_{jt+1} is the change in the price of stock j between 31st of December of year $t + 1$ and 31st of December of year t. Note that we consider passive capital gains at annual frequency, between 31st of December of each year, as this is the frequency at which we can also observe consumption. Throughout the analysis, we also exclude the top and bottom 1% of passive capital gains in the sample. We show below that our results are robust to various other approaches to dealing with outliers. We then match this data with our baseline retirement sample. Table F-1 provides descriptive statistics on this matched sample, and evaluates its representativeness, compared to our baseline sample. We observe financial portfolios in the KURU data for almost half of the individuals from our baseline retirement sample. The fraction of premature, early, normal and late retirees is remarkably similar in both samples. Other observable characteristics such as cohort, gender, education, labor market history at 55, earnings prior to retirement, or pensions received, are also well balanced across the two samples. As could be expected, the main difference is that individuals observed in the KURU data are somewhat wealthier on average. We therefore re-estimate the consumption differences across retirement age groups for the matched KURU sample: reassuringly, the consumption patterns, shown in Figure F-2, are virtually identical to those of our baseline sample (see Figure 3).

With this data in hand, we now show that conditional on a rich set of portfolio characteristics, innovations to stock prices generate persistent and plausibly exogenous wealth shocks (see also Andersen, Johannesen and Sheridan [2021]). For this purpose, we examine the serial correlation of passive capital gains, by regressing leads and lags of passive capital gains on current passive capital gains. For all years $k \in \{-6, ..., 6\}$, we estimate the following specification:

(24) $$\text{KG}_{i,t+k} = \alpha_k \text{KG}_{i,t+1} + \mathbf{X}'\beta$$

where $KG_{i,t+k} = \sum_j \Delta p_{j,t+k} \cdot a_{ijt}$ represents the passive capital gains that an individual would have accrued between $t + k - 1$ and $t + k$, still assuming the same portfolio as in year t. To account for the fact that portfolios of different value and of different risk structure face different stock price trends, the vector $\mathbf{X}$ controls non parametrically for the value of the portfolio in year t, as well as for the average returns and variance of the portfolio in the 6 years prior to year t.[54]

Appendix Figure F-3 plots the estimated coefficients $\hat{\alpha}_k$ for all time horizons $k \in \{-6, ..., 6\}$, revealing that current passive capital gains display no correlation with either past or future passive capital gains, conditional on portfolio value and structure. In other words, residual passive capital gains on listed stocks are as good as random, which implies that they generate random and persistent shifts in financial wealth. To provide a visual representation of the dynamic impact of passive capital gains on the value of an individual's portfolio, we correlate leads and lags of one's portfolio value $V_{i,t} = \sum_j p_{j,t} \cdot a_{ijt}$ with passive capital gains in year $t + 1$. More precisely, we regress the change in portfolio value $\Delta_{t,t+k} V_i = V_{i,t+k} - V_{i,t}$ of individual i between t and $t + k$ on her passive capital gains in $t + 1$, conditioning on a rich vector of portfolio characteristics $\mathbf{X}$:[55]

$$(25) \qquad \Delta_{t,t+k} V_i = \alpha_k^V KG_{i,t+1} + \mathbf{X}'\beta + \nu_{itk} \ , \forall k \in \{-3, ..., 3\}$$

Figure F-4 Panel A plots the estimated coefficients $\hat{\alpha}_{t+k}^V$ for all time horizons $k \in \{-3, ..., 3\}$. The graph shows that a passive capital gain of one krona is associated with a sharp, immediate, and permanent increase in portfolio value of about .6 krona.[56] These sharp dynamic patterns in portfolio values, driven by the randomness of stock price shocks, lend support to our strategy, which consists in treating passive capital gains conditional on portfolio characteristics $\mathbf{X}$, as an instrument for wealth. In Appendix Figure F-1, we show the variation in residualized capital gains that is key to our identification strategy. More than 31 percent of the passive capital gains/losses we exploit have absolute value over 10,000 SEK, which represent sizeable shocks.[57] Furthermore, the graph highlights that the distribution of our instrument is similar across retirement age groups.

---

[54]In practice, we use 50-tiles of portfolio value interacted with vigintiles of average returns in the past six years, and 50-tiles of portfolio value interacted with vigintiles of average variance in the past six years.

[55]The vector $\mathbf{X}$ controls non parametrically for the value of the portfolio in year t, as well as for the average returns and variance of the portfolio in the 6 years prior to year t. In practice, we use 50-tiles of portfolio value interacted with vigintiles of average returns in the past six years, and 50-tiles of portfolio value interacted with vigintiles of average variance in the past six years.

[56]Two related factors explain why $\hat{\alpha}_1^V$ is lower than 1, as one would have anticipated. First, because of the yearly frequency (between December and December) at which we observe stock price movements, and because of the presence of within-year trading, many portfolios change structure over the course of a year. For instance, an individual may have sold in January of $t + 1$ all her stocks a_j she held in December of year t. If all the price appreciation $\Delta p_{j,t+1}$ of stock j between December of year t and December of year $t + 1$ actually happened after January, e.g. between February and December of year $t + 1$, then KG_{t+1} will overstate the true capital gains experienced in $t + 1$. To the extent that intra-year trading is uncorrelated with the evolution of prices between these two dates, this will simply introduce measurement error. But, and this is the second factor, individuals may also *endogenously* realize their capital gains, thus decreasing portfolio value V_{t+1} by the share of passive capital gains that is realized. To deal with both issues, our approach consists in treating passive capital gains KG_{t+1} as an instrument for the change in financial wealth ΔV_{t+1}.

[57]These shocks are large compared to the variation exploited in the existing literature on wealth shocks. For instance, only 9% of the lottery shocks in Cesarini et al. [2016] are larger than 10,000 SEK.

Appendix F.3 MPC: Results

Our strategy relies on identifying the effect of wealth shocks on consumption by instrumenting wealth shocks by passive capital gains. We start by representing graphically the evolution of consumption around the time of a passive capital gain shock, which corresponds to the reduced-form dynamics of our IV. More precisely, we regress the change in consumption $\Delta_{t,t+k}C_i$ between year t and $t+k$ on the passive capital gains experienced in year $t+1$, conditioning on the same vector of portfolio characteristics as in (25):

$$(26) \qquad \Delta_{t,t+k}C_i = \alpha_k^C KG_{i,t+1} + \mathbf{X}'\beta + \varepsilon_{itk}$$

Panel B of Figure F-4 plots the estimated coefficients $\hat{\alpha}_k^C$ from the above specification, for all year horizons $k \in \{-3,...,3\}$. The graph conveys two important insights. First, in support of our identification strategy, we observe no sign of correlation between an individual's current passive capital gains and her consumption path in previous years. The absence of pre-trend in consumption indeed lends credibility to the validity of our instrument. Second, the figure shows that, in response to a passive capital gain of 1 krona, consumption increases immediately, significantly and persistently by about .03 krona. The sharpness of these consumption patterns, which closely mimic the dynamics of portfolio value in Panel A, suggests that our strategy truly captures the causal effect of the induced wealth shock on consumption.

To estimate the marginal propensity to consume, the increase in consumption estimated in Panel B needs to be scaled by the corresponding increase in wealth estimated from the first stage. In Panel A we get that the value of financial wealth increases by about .6 krona in response to a passive capital gain of 1 krona. Therefore, the estimated increase in yearly consumption of .03 krona translates into a marginal propensity to consume of .03/.6=.05 after a year, and of .15 after three years.

In Table F-2, we report 2SLS estimates of MPCs corresponding to the evidence presented in Figure F-4. We focus on average yearly consumption in the three years following a wealth shock $\overline{C}_{i,t,t+3}$, and estimate the following 2SLS model:

$$(27) \qquad \overline{C}_{i,t,t+3} - C_{i,t} = \alpha_{IV}^C \Delta_{t,t+1}V_i + \mathbf{X}'\beta + \eta_{it}$$
$$(28) \qquad \Delta_{t,t+1}V_i = \alpha_1^V KG_{i,t+1} + \mathbf{X}'\gamma + \varepsilon_{it}$$

Note that the vector $\mathbf{X}$ conditions on the same rich set of portfolio characteristics as in (25) and also includes controls for year and cohort fixed effects, as well as household structure, as in our analysis of consumption level and of consumption dynamics in the previous sections. It finally includes a dummy for being retired in t. So in effect, we allow the dynamics of consumption to flexibly differ across individuals with different retirement status, household structure or portfolio allocations. The coefficient α_{IV}^C identifies the average yearly marginal propensity to consume in years $t+1$ to $t+3$, out of an increase in financial wealth V_i generated by a random passive capital gains incurred between t and $t+1$. We obtain an MPC estimate *over a three years horizon*, by multiplying the coefficient α_{IV}^C by three. Standard errors are clustered at the

individual level, and we explore the robustness of our results to alternative inference strategies in Appendix Table F-5.

Results reported in Panel A confirm the graphical evidence from Figure F-4. We find an average marginal propensity to consume of .17 (.01) over three years. This estimate lies at the lower end of the distribution of MPC estimates found in the literature, but can be rationalized by the fact that our population of interest is on average older and wealthier than in other similar studies. Furthermore, our results are in line with estimated MPCs in Di Maggio, Kermani and Majlesi [2020] who also rely on passive capital gains shocks as instruments for wealth shocks. We also report in the last column of Table F-2 the estimates from a placebo test where we replicate specification (27) using as an outcome the change in consumption in the three years prior (rather than after) the wealth shock. The lack of any significant pre-trend is an important validation of the credibility of our identification strategy.

In Panel B, we split the sample according to retirement status at the time of the passive capital gain shock to explore how MPCs differ before vs after retirement. We find that marginal propensities to consume increase significantly after retirement. The MPC of retired individuals is .30 (.04), compared to .13 (.01) for individuals who are still working. Because we are comparing retired and non-retired individuals conditional on age and cohort fixed effects, these results are not simply capturing the fact that older, retired individuals have a shorter horizon over which to smooth consumption, driving their MPCs up. Rather, it suggests that retirement is associated with an increase in the value of liquidity conditional on age.

In Panel C, we then split the sample and estimate MPCs by retirement age groups, to see how the value of liquidity varies with retirement age. The results show significant heterogeneity in MPCs across retirement age groups with a strong overall negative gradient of MPCs with retirement age. MPCs for premature and early retirees are around .34 over three years, and markedly larger than for normal retirees (.09). Interestingly we find that the MPC of late retirees is small, and not significantly different from zero. In other words, while the value of additional liquidity seems to be high for individuals who retire early or prematurely, it seems negligible for late retirees. These results accord with the earlier evidence, indicating that individuals who retire earlier are less resourceful and more likely to be subject to negative, uninsured shocks, and as a consequence, value additional liquidity to increase their consumption more relative to individuals who retire late.

The results from Panel C of Table F-2 control for age and retirement status and therefore compare MPCs of individuals who retire at different ages while in the same retirement state. Yet, the estimates may capture different LATEs across retirement age groups, as they will place more weights on the MPCs of retired people among the premature retirees, and more weights on non-retired individuals among the late retirees. Ideally, we would therefore like to compare the MPCs of the different retirement age groups *only while retired*. Having enough power to do so however, requires adding more cohorts to our original sample, in order to observe a long panel of consumption while retired for all retirement age groups. We do so in Table F-3 where we enlarge our sample to include all cohorts from 1932 to 1943, and restrict the sample to individuals who are retired at the moment of experiencing a capital gain shock. Panel A reports

estimates for all retirement age groups together: the estimated MPC over a three-year horizon is .28 (.04). Reassuringly, this is almost identical to the estimated MPC for retired individuals in Panel B of Table F-2, which focused on cohorts 1938 to 1943. In Panel B, we report the estimated MPC when splitting the sample into our four retirement age groups. The estimates are sensitive to the specification: the MPC estimate for the premature retirees is somewhat lower, but the standard errors are relatively large. The results do confirm that the MPCs in retirement are low and insignificant for late retirees, while they are generally higher for individuals retiring earlier. Table F-4 provides further sensitivity analysis: it shows the presence of significantly larger MPCs when focusing on smaller capital gains by excluding the top and bottom 5% of capital gains. But the gradient with retirement age remains robustly negative.

Appendix F.4   Quantifying Welfare Costs

While our MPC estimates for the different retirement-age groups are more imprecise and more sensitive to the specification, we briefly illustrate again how to quantify the welfare costs of providing late-career incentives. Column (6) in Appendix Table H-1 reports the consumption smoothing costs based on the MPC implementation in equation (23). We use the MPC estimates for retired individuals, but this requires using the extended sample as reported in Table F-3. On average the estimated consumption smoothing costs are similar as for the other implementations - we come back to this in Section VI when evaluating the overall change in late-career incentives due to the Swedish reform - but the pattern in retirement age-specific costs is somewhat different compared to the other implementations. On the one hand, we do not find that the welfare cost is higher when steepening incentives for the premature retirees than for the early retirees. This reflects the low MPC estimate for the premature retirees relative to the early and normal retirees in the extended sample. This, however, does not appear for the baseline sample in Table F-2. On the other hand, we do find a high cost (.76) of steepening the profile for late retirees, coinciding with the estimate from the consumption-level implementation. This is driven by their near-zero MPC estimates, which indeed imply that the value of providing extra liquidity to late retirees is very small.

Table F-1: Descriptive Statistics on MPC Sample (i.e. Retirement Sample Matched to KURU Data on Financial Portfolios)

	Retirement Sample		Retirement x Stock Sample	
	Mean	(s.d.)	Mean	(s.d.)
I. Retirement				
Premature Retirement Probability	23.81 %		22.39%	
Early Retirement Probability	25.68 %		30.84%	
Normal Retirement Probability	34.60 %		35.35%	
Late Retirement Probability	15.91 %		11.41%	
Age at Retirement	62.91	(3.1)	62.80	(2.77)
II. Demographics				
Cohort	1940.67	(1.73)	1940.56	(1.68)
Fraction Men	49.29 %	(50.00)	52.27%	(49.95)
Fraction Married	66.86 %	(47.07)	72.59%	(44.6)
Kid at Home (≥ 1)	17.65 %	(38.12)	17.14%	(37.68)
Kid at Home Under 18 (≥ 1)	3.48 %	(18.33)	2.98%	(17.00)
Post-Secondary Education	24.67%	(43.11)	30.29%	(45.95)
III. Pension Information, SEK 2003				
State Pension	78.50	(52.90)	83.80	(55.60)
Occupational Pension	62.10	(92.60)	81.70	(112.10)
ATP Pension at 55	95.60	(38.80)	103.30	(38.30)
IV. Income and Wealth at 59, SEK 2003(K)				
Total Earnings	209	(160)	229	(178)
Net Wealth	777	(2339)	1136	(2438)
Bank Holdings	84	(312)	113	(422)
Portfolio Value	248	(1648)	248	(1648)
Consumption	201	(534)	217	(657)
N (Unique Individuals)	418,033		182,544	
Cohorts	[1938,1943]		[1938,1943]	

Notes: The table reports descriptive statistics from our baseline sample of retirees and for the baseline sample matched with portfolio information on stock ownership (KURU). Both samples are restricted to cohorts 1938 to 1943 who retire between age 56 and 69. The matched sample comprises 182,544 unique individuals. Retirement is defined as labor earnings dropping permanently below one Base Amount. Panel I reports statistics on the distribution of retirememt age. Premature retirement is defined as individuals retiring between age 56 and 60; early retirement, between age 61 and 63; normal retirement, between age 64 and 65; and late retirement, between age 66 and 69. Panel II reports various demographic information. Panel III reports the average state and occupational pension benefits received. Total ATP points correspond to the total number of ATP points accumulated in the state pension system at age 55. Panel IV focuses on income and wealth measured at age 59. Wealth and consumption is aggregated at the household level. Note that based on the average exchange rate between 2000 and 2007, 1SEK $\approx$ 0.11USD.

Figure F-1: DISTRIBUTION OF RESIDUALIZED PASSIVE CAPITAL GAINS BY RETIREMENT AGE GROUP
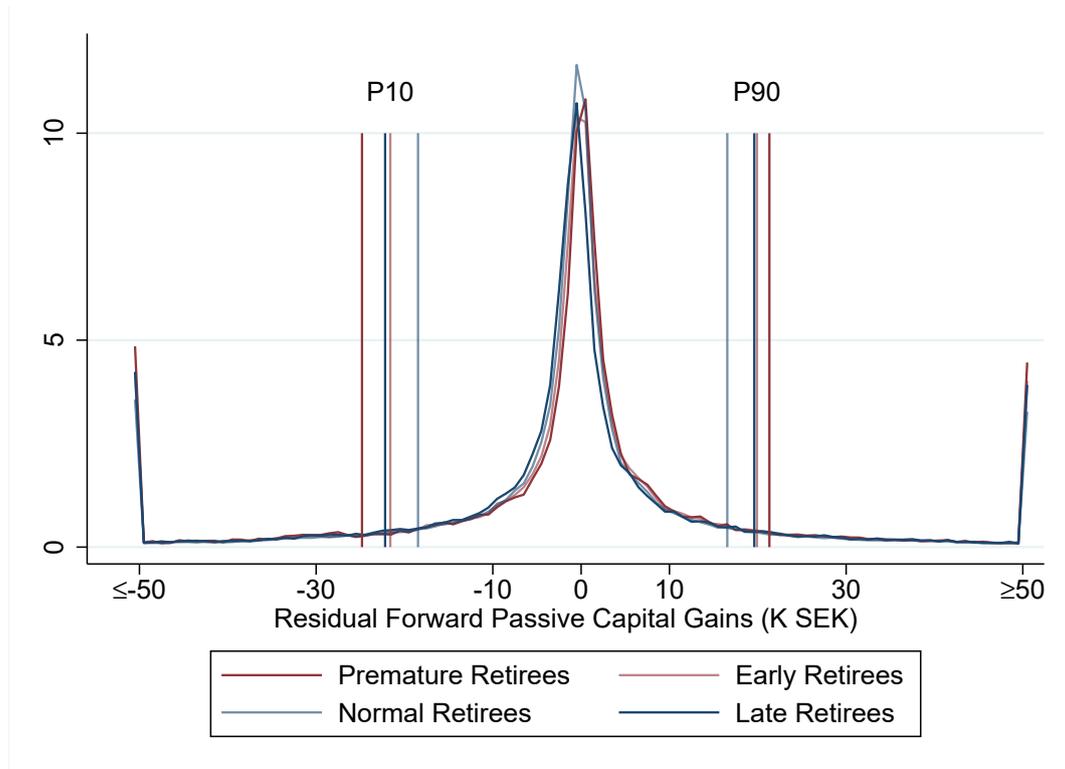
Notes: The figure plots the distribution of residualized passive capital gains. The sample is the baseline retirement sample merged with the KURU register, which has disaggregated information over the period 1999 to 2007 on all quantities of stocks, by ISIN number, held by individuals outside of mutual funds. The sample is described in Table F-1 above. For each individual i, passive capital gains on her portfolio in year $t+1$ are defined as $KG_{i,t+1} = \sum_j (p_{j,t+1} - p_{j,t}) \cdot a_{ijt} = \sum_j \Delta p_{j,t+1} \cdot a_{ijt}$ where a_{ijt} is number of stocks of company j held by individual i on 31st of December of year t and Δp_{jt+1} is the change in the price of stock j between 31st of December of year $t+1$ and 31st of December of year t. The passive KG are then residualized on a set of portfolio characteristics, capturing the value of the portfolio in year t, as well as the average returns and variance of the portfolio in the 6 years prior to year t. In practice, we use 50-tiles of portfolio value interacted with vigintiles of average returns in the past six years, and 50-tiles of portfolio value interacted with vigintiles of average variance in the past six years. In Figure F-3, we show that these residualized passive KG follow a random walk. The Figure plots the distribution of residualized $KG_{i,t+1}$, and also indicates the 10th and 90th percentile of the distribution, for each retirement age group. More than 31% percent of the residual passive capital gains/losses we exploit have absolute value over 10,000 SEK, which represent sizeable shocks. These shocks are large compared to the variation exploited in the existing literature on wealth shocks. For instance, only 9% of the lottery shocks in Cesarini et al. [2016] are larger than 10,000 SEK. Furthermore, the graph highlights that the distribution of our instrument is similar across retirement age groups.

Figure F-2: Consumption Differences by Retirement Age Group in Baseline Sample and MPC Sample
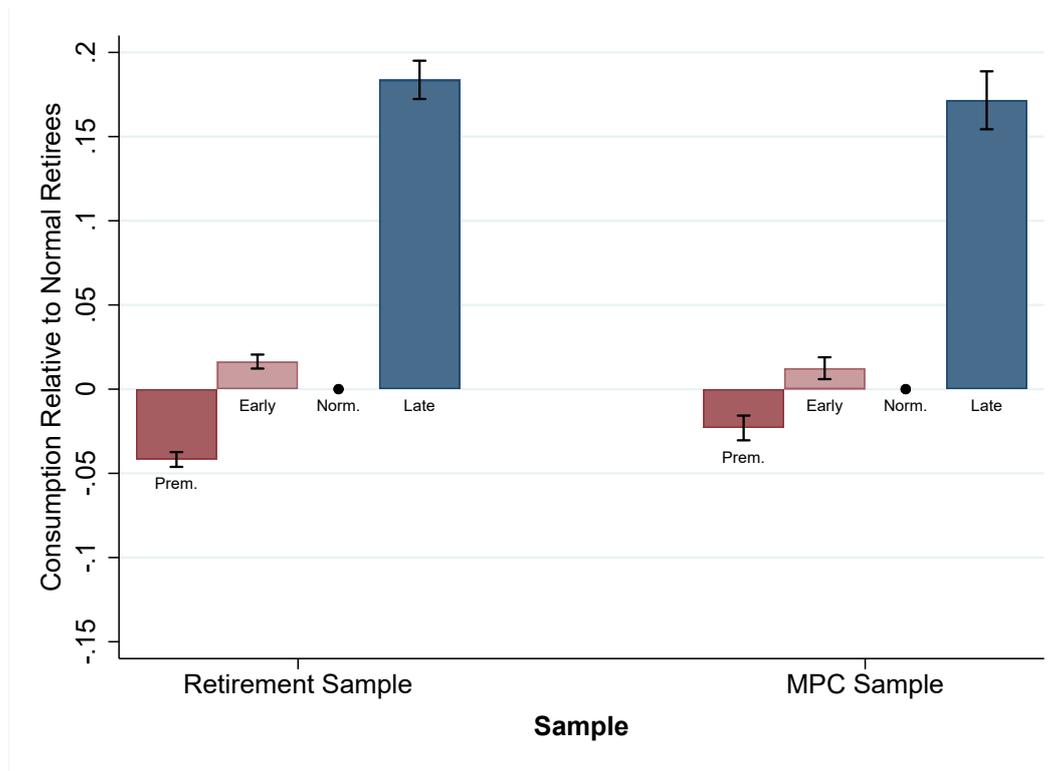
Notes: The figure replicates in the MPC sample our baseline estimates of consumption differences in retirement from Figure 3. Both samples comprise individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed. Individuals are grouped into four retirement age categories: premature retirees ($56 \leq r \leq 60$), early retirees ($61 \leq r \leq 63$), normal retirees ($64 \leq r \leq 65$) and late retirees ($66 \leq r \leq 69$). Normal retirees are the reference category. The graph reports for all retirement age groups, the estimated coefficients α_j from specification (8), scaled by $\mathbf{E}_j[\tilde{C}_{it}]$, the average level of consumption of individuals who retire between 64 and 65 from the same cohort, age and family composition as the average individual retiring in age group j. On the left hand side of the graph, we reproduce results from Figure 3 for the model (8) with year and age fixed effects and controls for family composition. On the right hand side of the graph, we plot the estimates obtained from the same model run on the MPC sample.

Figure F-3: Serial Correlation In Residual Passive K Gains & Passive Value of Portfolio

A. Serial Correlation In Residual Passive K Gains

B. Predicted Passive Value of Portfolio
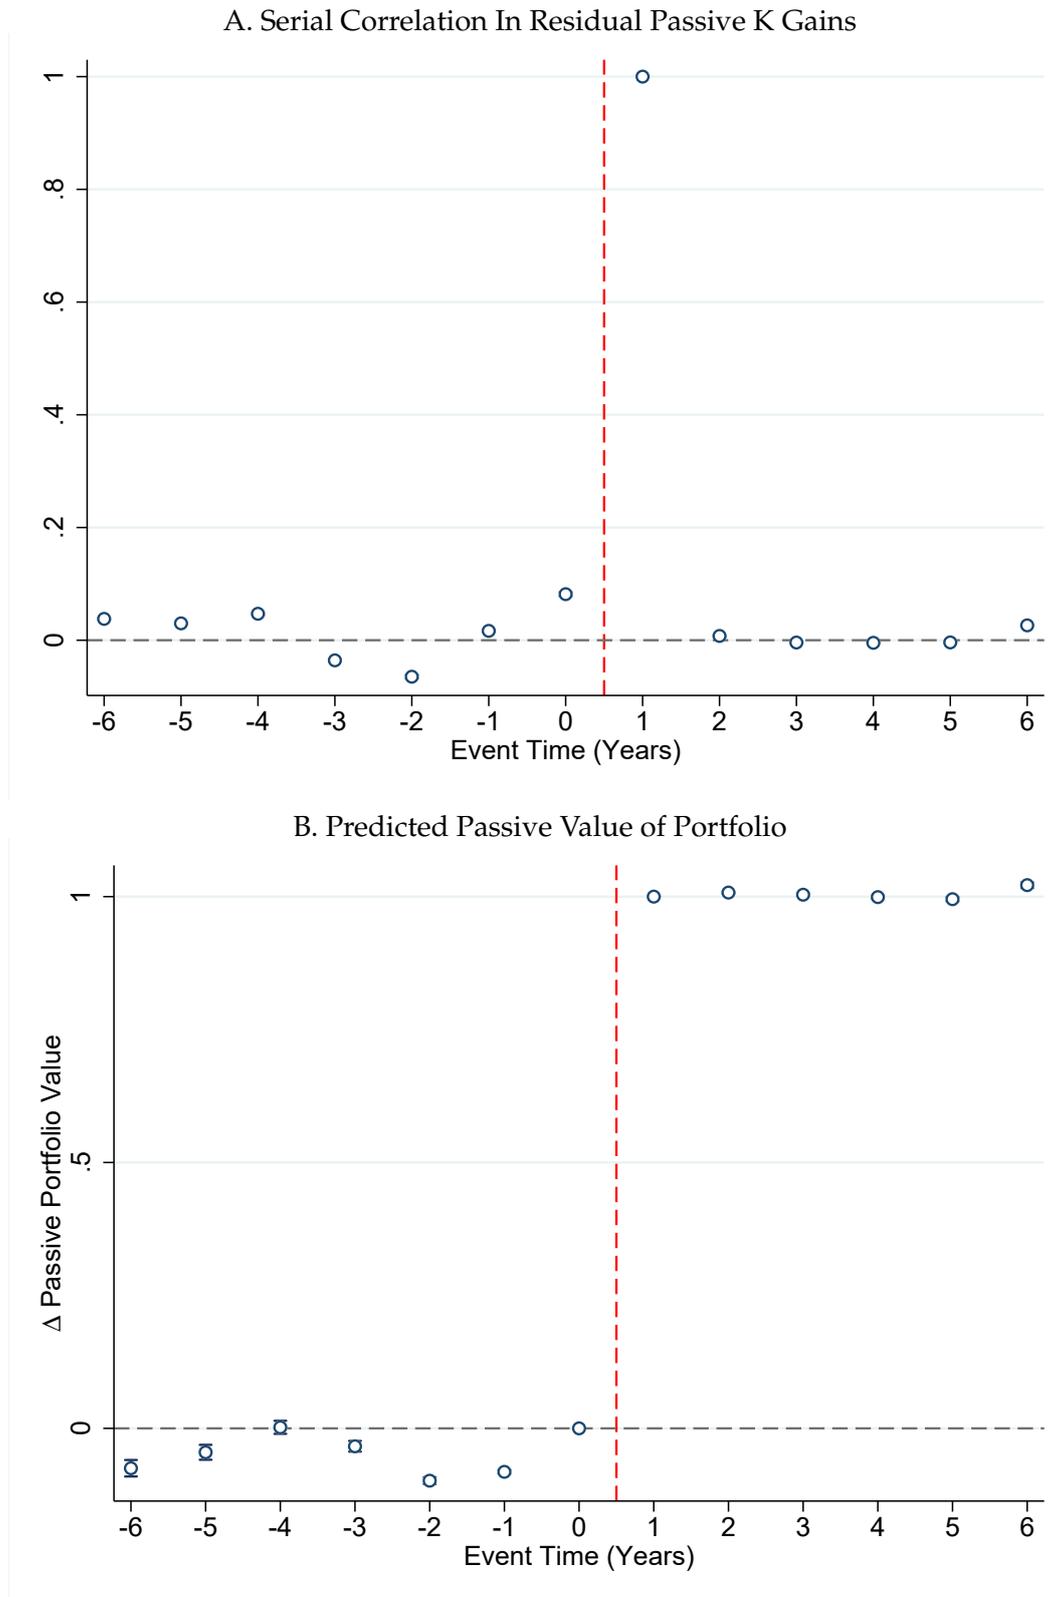
Notes: Panel A plots for each time horizon $k \in \{-6, ..., 6\}$ the serial correlation of the residual passive capital gain at k and the current residual passive capital gain, that is the coefficient α_k from regression (25). We control for the value of portfolio in year t, the average returns and variance of the portfolio in the 6 years prior to year t. Panel B examines the predictive effect of the residual on the change in passive portfolio value for each time horizon $k \in \{-6, ..., 6\}$. The passive portfolio value in year $t + k$ is defined as $\sum_j p_{j,t+k} \cdot a_{ijt}$ where a_{ijt} is number of stocks of company j held by individual i on 31st of December of year t and p_{jt+k} is the price of stock j in 31st of December of year $t + k$. It is therefore the value that the portfolio held in year t would be worth in year $t + k$ if the owner of the portfolio had not rebalanced it.

Figure F-4: MARGINAL PROPENSITIES TO CONSUME OUT OF WEALTH SHOCKS

A. First Stage: Portfolio Value
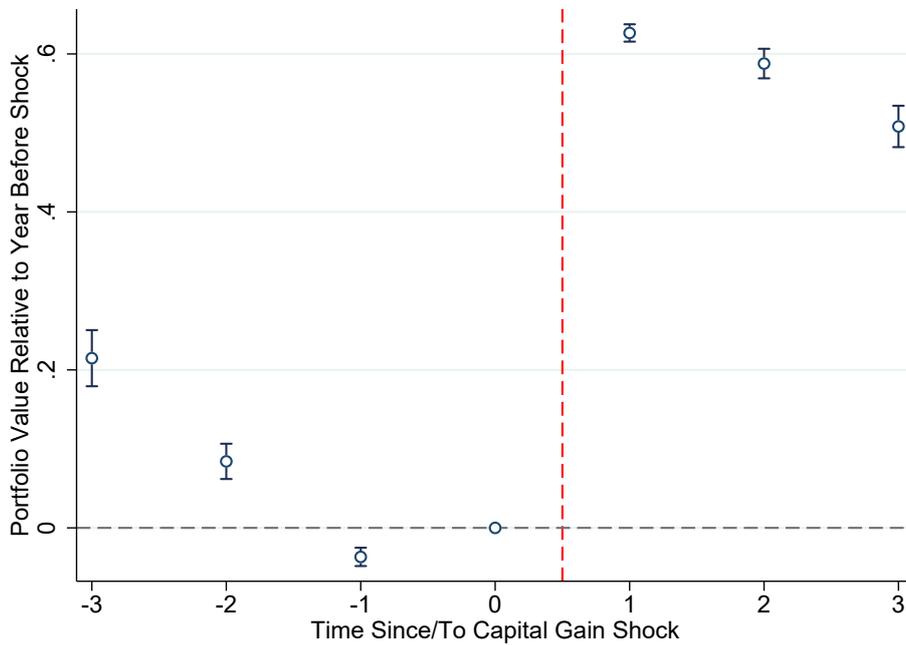
B. Reduced-Form: Consumption
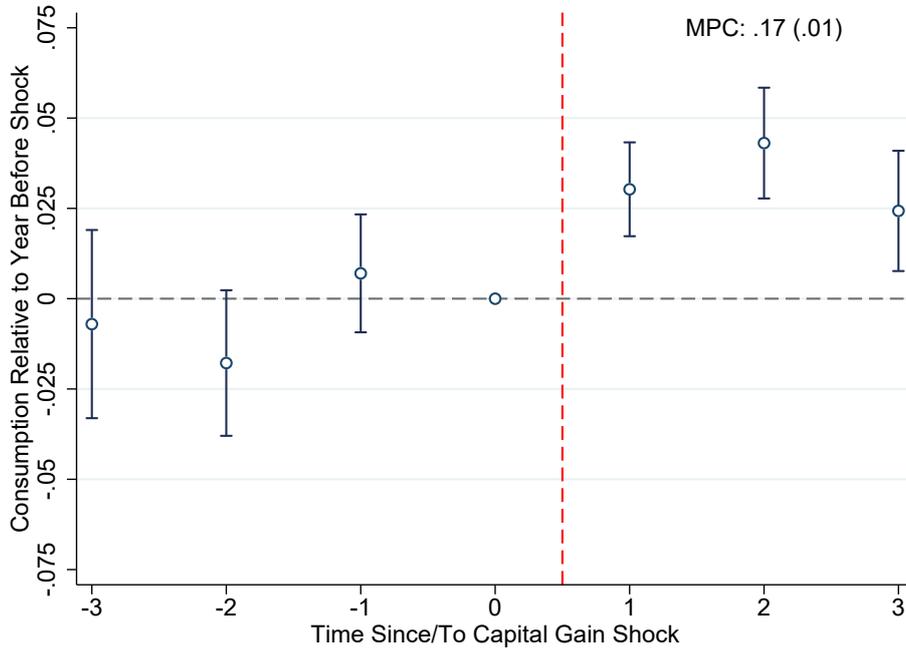
Notes: Panel A reports the estimates of the first stage regression, that is the regression of the change in portfolio value between t and $t + k$ at year t on the passive capital gains at $t + 1$, that is a year after the wealth shock, controlling for the value of the portfolio in year t, as well as for the average returns and variance of the portfolio in the 6 years prior to year t (see equation (25)). Panel B reports the estimates of the reduced form regression, that is, for each year k, the regression of the change in consumption between t and $t + k$ on the forward passive capital gain at $t + 1$, controlling for the value of the portfolio in year t, as well as for the average returns and variance of the portfolio in the 6 years prior to year t (see equation (27)). It also reports the implied marginal propensity to consume, which is the ratio of the reduced form and the first stage over the three years.

Table F-2: 2SLS Estimates of Marginal Propensity to Consume Out of Wealth Shocks

	First Stage α_1^V	Reduced Form $3 \times \alpha_{rf}^C$	IV Result $3 \times \alpha_{IV}^C$	Placebo Test α_1^P
	A. Whole Sample			
	0.66	0.11	0.17	0.04
	(0.01)	(0.03)	(0.04)	(0.01)
N	546,836	546,836	546,836	546,836
# of Indiv. Clusters	133,133	133,133	133,133	133,133
	B. By Retirement Status			
Non Retired in t	0.66	0.09	0.13	-0.01
	(0.01)	(0.03)	(0.04)	(0.02)
N	472,980	472,980	472,980	472,980
# of Indiv. Clusters	129,099	129,099	129,099	129,099
Retired in t	0.71	0.21	0.30	0.07
	(0.03)	(0.09)	(0.13)	(0.05)
N	74,022	74,022	74,022	74,022
# of Indiv. Clusters	42,735	42,735	42,735	42,735
	C. By Retirement Age Group			
Premature Retirees	0.69	0.23	0.34	-0.01
	(0.04)	(0.08)	(0.12)	(0.07)
N	48,425	48,425	48,425	48,425
# of Indiv. Clusters	16,667	16,667	16,667	16,667
Early Retirees	0.63	0.22	0.34	0.03
	(0.02)	(0.05)	(0.08)	(0.03)
N	172,747	172,747	172,747	172,747
# of Indiv. Clusters	46,972	46,972	46,972	46,972
Normal Retirees	0.68	0.06	0.09	0.03
	(0.01)	(0.04)	(0.06)	(0.02)
N	261,351	261,351	261,351	261,351
# of Indiv. Clusters	56,705	56,705	56,705	56,705
Late Retirees	0.70	0.00	0.01	0.06
	(0.03)	(0.08)	(0.12)	(0.05)
N	63,707	63,707	63,707	63,707
# of Indiv. Clusters	12,845	12,845	12,845	12,845

Notes: The table reports the 2SLS results from equations (27) and (28). Column (1) reports the estimates of the first stage, obtained by regressing the change in portfolio value of the individual between t and $t + k$ on the passive capital gains in $t + 1$, controlling for the value of portfolio in year t, the average returns and variance of the portfolio in the 6 years prior to t, but also adding a dummy for the retirement status and controlling for year, cohort fixed effects and household structure. We cluster the standard errors at the individual level. Column (2) reports the estimates of the reduced form, obtained by regressing the average yearly consumption in the three years following the wealth shock on the change in the value of the portfolio in year t instrumented by the passive capital gains. We add the same controls as in the first stage. The estimates are multiplied by three in order to obtain the MPC over a three years horizon. Column (3) reports the instrumental variable results, obtained by taking the ratio of the reduced form to the first stage, over a three years horizon. Column (4) presents the results of the placebo test, which is a replication of equation (26) where the outcome is the change in yearly consumption in the three years before the shock. The results are presented for three panels. Panel A consists of the observations considered in the baseline sample from regression (8) matched with KURU data. Panel B considers this same sample split according to the retirement status at the time of the passive capital gain shock. Panel C is a split of this same sample by retirement age group. For each sample, we trim the change in portfolio value at the 1% level and the passive capital gain each year at the 1% level.

Table F-3: 2SLS Estimates of MPCs: Sample restricted to Individuals Who Are Retired At Time of KG Shock

	First Stage α_1^V	Reduced Form $3 \times \alpha_{rf}^C$	IV Result $3 \times \alpha_{IV}^C$	Placebo Test α_1^P
A. All Retired Individuals (Cohorts 1932-1943)				
	0.70	0.19	0.28	0.06
	(0.02)	(0.07)	(0.11)	(0.02)
N	110,153	110,153	110,153	110,153
# of Indiv. Clusters	59,419	59,419	59,419	59,419
B. By Retirement Age Group (Cohorts 1932-1943)				
Premature Retirees	0.77	0.18	0.24	0.04
	(0.05)	(0.14)	(0.18)	(0.09)
N	20,064	20,064	20,064	20,064
# of Indiv. Clusters	9,595	9,595	9,595	9,595
Early Retirees	0.69	0.26	0.37	-0.09
	(0.03)	(0.13)	(0.19)	(0.13)
N	38,301	38,301	38,301	38,301
# of Indiv. Clusters	21,118	21,118	21,118	21,118
Normal Retirees	0.69	0.25	0.36	0.03
	(0.03)	(0.12)	(0.18)	(0.05)
N	37,647	37,647	37,647	37,647
# of Indiv. Clusters	21,036	21,036	21,036	21,036
Late Retirees	0.78	0.05	0.07	0.08
	(0.05)	(0.20)	(0.24)	(0.08)
N	14,084	14,084	14,084	14,084
# of Indiv. Clusters	7,589	7,589	7,589	7,589

Notes: This table follows the same approach as for Table F-2 where we enlarge the set of cohort to 1932 - 1943 and restrict to individuals retired at the moment of the capital gain shock.

Table F-4: 2SLS Estimates of MPCs: Robustness to Size of KG Shocks

	First Stage α_1^V	Reduced Form $3 \times \alpha_{rf}^C$	IV Result $3 \times \alpha_{IV}^C$
	A. Without Top/Bottom 5% of KG Shocks		
All Retirees	0.34	0.17	0.49
	(0.01)	(0.04)	(0.13)
Premature Retirees	0.29	0.37	1.26
	(0.02)	(0.12)	(0.44)
Early Retirees	0.32	0.26	0.81
	(0.01)	(0.08)	(0.23)
Normal Retirees	0.38	0.07	0.20
	(0.01)	(0.07)	(0.17)
Late Retirees	0.36	0.05	0.14
	(0.02)	(0.14)	(0.38)

Notes: This table shows the estimates of the 2SLS approach presented in equation (27). Column (1) reports the estimates of the first stage, column (2) the estimates of the reduced form, multiplied by three to obtain the MPC over a three years horizon. The IV result is presented in column (3). This sample is composed of the observations from the baseline analysis matched with the KURU information, trimming the value of portfolio at the 5% level. We also drop all values of passive capital gain above the 99-th percentile each year. We cluster the standard errors at the individual level.

Table F-5: 2SLS Estimates of MPCs: Robustness to Alternative Clustering

	First Stage α_1^V	Reduced Form $3 \times \alpha_{rf}^C$	IV Result $3 \times \alpha_{IV}^C$
Cluster by 50-tile of PF Value PF x 20-tile of Average PF Past Returns			
Baseline	0.66	0.11	0.17
	(0.04)	(0.03)	(0.04)
Number of Observations	546,836	546,836	546,836
Number of clusters	972	972	972

Notes: This table shows the results of the MPC analysis on the baseline sample, this time clustering at the cinquantile of portfolio value times vigintile of average portfolio past returns.

Appendix G   Conceptual Framework

This appendix provides more detail underlying the model setup and the derivations of the welfare impact of a pension reform and its various empirical implementations.

Model Setup   As stated in the main text, the individual's expected lifetime utility is given by

$$(29) \qquad \mathcal{U}_i(c, \zeta, \pi) = \sum_{t=0}^{T} \beta^t \int u\left(c\left(\pi_{i,t}\right), \zeta\left(\pi_{i,t}\right)\right) dF\left(\pi_{i,t}\right),$$

where $c\left(\pi_{i,t}\right)$ is the individual's consumption choice and $\zeta\left(\pi_{i,t}\right)$ represents all other choices and characteristics, either affecting an individual's utility or his or her or the government's budget constraint. This includes decisions regarding labor supply, home production choices, financial investments, bequests, etc, but also shocks to health, financial or human capital. We often use short-hand notation $c_{i,t}$ and $\zeta_{i,t}$ for these. Without loss of generality, we assume that all individual heterogeneity is captured through realizations of the state variable over the lifetime, including the starting values $\pi_{i,0}$.

Implicitly our analysis here considers a single cohort, so that age and time are equivalent. Inter-cohort/inter-generational concerns may affect optimal benefit levels, but we focus here on the within-cohort distribution of pension benefits. Despite our use of a deterministic final period T, we can capture life expectancy concerns affecting the marginal utility of consumption through the reduced-form ζ parameter.

The model is set up in reduced-form, but the various exogenous and endogenous factors in standard retirement models (see Blundell, French and Tetlow [2016]) can be captured through ζ and how it affects the utility of consumption c. Like in all structural models of retirement, $\zeta\left(\pi_{i,t}\right)$ includes the extensive labor supply choice, which is denoted by $s\left(\pi_{i,t}\right)$ and takes value 1 if an individual is employed and value 0 if an individual is retired. We assume that an individual retires only once, denoting the retirement age choice once someone has decided to retire by $r\left(\pi_{i,t}\right)$. We thus have $s\left(\pi_{i,t}\right) = 0$ for $t \geq r\left(\pi_{i,t}\right)$ and $s\left(\pi_{i,t}\right) = 1$ otherwise. Hence, the number of individuals retiring at each r equals $S(r-1) - S(r)$, where $S(r) = \int \int s\left(\pi_{i,r}\right) dF\left(\pi_{i,r}\right) di$ is the survival rate into employment.

We note that $\zeta\left(\pi_{i,t}\right)$ can also include exogenous factors to either capture relevant heterogeneity across workers from the start $\pi_{i,0}$ (e.g., in preferences, health or ability) or risks that individuals face (e.g., health or ability shocks) and realize over time, represented by the CDF $F\left(\pi_{i,t}\right)$ (see French and Jones [2011]). The general set up can also accommodate mortality risks and preferences over bequests as in French [2005]:

$$u\left(c\left(\pi_{i,t}\right), \zeta\left(\pi_{i,t}\right)\right) = \zeta_M\left(\pi_{i,t}\right) \tilde{u}\left(c\left(\pi_{i,t}\right), \tilde{\zeta}\left(\pi_{i,t}\right)\right) + \left(1 - \zeta_M\left(\pi_{i,t}\right)\right) \tilde{v}\left(\zeta_B\left(\pi_{i,t}\right)\right),$$

where ζ_M denotes the survival probability and ζ_B denotes any bequeathed wealth. The setup can also accommodate health shocks affecting required medical expenditures and/or the utility of consumption net of these medical expenditures:

$$u\left(c\left(\pi_{i,t}\right), \zeta\left(\pi_{i,t}\right)\right) = \zeta_{X_1}\left(\pi_{i,t}\right) \times \tilde{u}\left(c\left(\pi_{i,t}\right) - \zeta_{X_2}\left(\pi_{i,t}\right), \tilde{\zeta}\left(\pi_{i,t}\right)\right),$$

where ζ_{X_2} denotes the medical expenditures and ζ_{X_1} scales the utility of non-medical expenditures (e.g., Blundell, Borella, Commault, De Nardi, 2021 no 2020).

We focus on workers' extensive labor supply and the age at which they retire. If $s(\pi_{i,t}) = 0$ (retirement), the individual receives pension benefits $b(\pi_{i,t})$. If $s(\pi_{i,t}) = 1$ (employment), the individual earns wages $w(\pi_{i,t})$ and pays taxes $\tau(\pi_{i,t})$. In either case ($s \in 0,1$) after-tax income is denoted by $y(\pi_{i,t})$. Assets $a_{i,t+1}(\pi_{i,t})$ evolve in the usual fashion, based on previously accumulated assets and saving in year t, with a gross rate of return $R(\pi_{i,t})$. The individual's optimization problem is therefore to maximize $\mathcal{U}_i$ subject to the following resource constraints for each history $\pi_{i,t}$:

$$(30) \qquad a_{t+1}(\pi_{i,t}) = R(\pi_{i,t})\left[a_t(\pi_{i,t-1}) + y(\pi_{i,t}) - c(\pi_{i,t})\right],$$

$$(31) \qquad y(\pi_{i,t}) = \begin{cases} w(\pi_{i,t}) - \tau(\pi_{i,t}) \text{ if } s(\pi_{i,t}) = 1 \\ b(\pi_{i,t}) \text{ if } s(\pi_{i,t}) = 0. \end{cases}$$

Taxes $\tau(\pi_{i,t})$ and pension benefits $b(\pi_{i,t})$ can depend in a flexible way on a worker's employment history, including the number of years worked and the corresponding earnings. The government's objective is

$$(32) \qquad \max \ \mathcal{W}(b,\tau) = \int_i \omega_i \mathcal{U}_i(b,\tau) + \lambda GBC(b,\tau)\,di,$$

where

$$(33) \qquad GBC(b,\tau) = \Sigma_t \frac{1}{R^t}\int\int \left[s(\pi_{i,t})\tau(\pi_{i,t}) - (1 - s(\pi_{i,t}))b(\pi_{i,t})\right]f(\pi_{i,t})\,d\pi_{i,t}di - G_0.$$

Note that we can simplify this further re-writing the budget constrains as a function of the average tax paid by workers at age r, τ_r, and the net present value of the pension benefits received for workers retiring at age r:

$$NPV_r \equiv \frac{1}{R^r}\Sigma_{t=r}^T \frac{1}{R^{t-r}}\int\int b(\pi_{i,t})\frac{1[r(\pi_{i,t}) = r]}{S(r-1) - S(r)}dF(\pi_{i,t})\,di.$$

The government's budget constraint becomes

$$(34) \qquad GBC(b,\tau) = \Sigma_r \left[S(r)\frac{\tau_r}{R^r} - [S(r-1) - S(r)]NPV_r\right] - G_0,$$

clearly illustrating how government revenues and expenditures change with the age at which workers decide to retire. The model can in principle be extended with claiming decisions, as well as pathways to retirement through DI or UI, which then should be accounted for in the NPV_r.

Characterization The policy variation we consider is a uniform change in the benefits received by all retired individuals with the same feature x. That is, $db(\pi_{i,t}) = db_{x,t}$ for $x(\pi_{i,t}) = x$ and $s(\pi_{i,t}) = 0$. To characterize the welfare impact, we can invoke the envelope theorem, implying that the only first-order effect on workers' welfare comes from the direct effect of the

benefit receipt. We write:

$$
\begin{aligned}
SMU_{x,t} &= E\left(\omega_i \beta^t \frac{\partial u\left(c_{i,t}, \zeta_{i,t}\right)}{\partial c} \middle| x_{i,t} = x\right) \\
&= \int\int \omega_i \beta^t \frac{\partial u\left(c\left(\pi_{i,t}\right), \zeta\left(\pi_{i,t}\right)\right)}{\partial c} \frac{1\left[x\left(\pi_{i,t}\right) = x, s\left(\pi_{i,t}\right) = 0\right]}{G\left(x,t\right)} dF\left(\pi_{i,t}\right) di,
\end{aligned}
$$

where $G(x,t) = \int\int 1[x\left(\pi_{i,t}\right) = x, s\left(\pi_{i,t}\right) = 0]dF\left(\pi_{i,t}\right)di$. To compare this value to its fiscal cost, we should account for the fiscal externality of any response in $\zeta_{i,t'}$ throughout the individual's lifetime (for any t') and the implications this change in behavior has on the distribution of future states $F\left(\pi_{i,t'+k}\right)$ (for any k). In principle, individuals can change their earnings throughout their lifetime - with further consequences on the tax revenues and expected pension payments, captured through the history $\pi_{i,t}$. The change in benefits and retirement behavior can also change individuals' health and life expectancy and the labor supply of other individuals in the household with corresponding fiscal consequences (see Blundell, French and Tetlow [2016]). If we consider only the behavioral response at the extensive labor supply margin by directly affected workers, the impact on the budget constraint of the pension change $db_{x,t}$ can then be written as:

$$
1 + FE_{x,t} \equiv \frac{1}{R^t} + \Sigma_{r'}\left\{\left[\frac{\tau_{r'}}{R^{r'}} - \left(NPV_{r'+1} - NPV_{r'}\right)\right]\frac{\frac{\partial(1-S(r'))}{\partial b_{x,t}}}{S\left(r-1\right) - S\left(r\right)}\right\}.
$$

Putting the two effects together, the welfare impact per dollar spent on $b_{r,t}$ equals for $\beta R = 1$:

$$
SMU_{x,t} - \lambda\left[1 + FE_{x,t}\right].
$$

Note that this relates directly to the marginal value of public funds (MVPF) of spending on specific pension beneficiaries (Hendren and Sprung-Keyser [2020]). When the social value per dollar spent, accounting for the fiscal externality, is larger for individuals with features x vs. x',

(35)
$$
\frac{SMU_{x,t}}{1 + FE_{x,t}} > \frac{SMU_{x',t}}{1 + FE_{x',t}}
$$

we can increase welfare from spending that extra dollar on pension benefits for the former and spending a dollar less on the latter.

Consumption Implementations for Social Marginal Utility Terms   We now consider the use of consumption moments to evaluate the social marginal utility of consumption for subgroups retiring with different features. We illustrate these consumption-based implementations for individuals retiring at different ages r, but this naturally holds for any other feature x.

We assume that the only relevant heterogeneity occurs across workers retiring at different ages, so that $c\left(\pi_{i,t}\right) = c_{r,t}$ and $\zeta\left(\pi_{i,t}\right) = \zeta_{r,t}$ for $r\left(\pi_{i,t}\right) = r$. The consumption-level implementation

then immediately follows from a Taylor approximation or $\frac{\partial u(c_{r,t},\zeta_{r,t})}{\partial c}$ around $(c_{r',t},\zeta_{r,t})$,

$$\frac{\partial u\left(c_{r,t},\zeta_{r,t}\right)}{\partial c} \cong \frac{\partial u\left(c_{r',t},\zeta_{r,t}\right)}{\partial c}\left[1 + \frac{-\frac{\partial^2 u\left(c_{r',t},\zeta_{r,t}\right)}{\partial c^2}c_{r',t}}{\frac{\partial u\left(c_{r',t},\zeta_{r,t}\right)}{\partial c}}\frac{c_{r',t}-c_{r,t}}{c_{r',t}}\right].$$

Denoting the relative risk aversion parameter by $\gamma\left(c_{r',t},\zeta_{r,t}\right) = \frac{-\frac{\partial^2 u\left(c_{r',t},\zeta_{r,t}\right)}{\partial c^2}c_{r',t}}{\frac{\partial u\left(c_{r',t},\zeta_{r,t}\right)}{\partial c}}$, we have

$$\frac{E\left(\omega_i\beta^t\frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial c}\Big|\, r_i = r\right)}{E\left(\omega_i\beta^t\frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial c}\Big|\, r_i = r'\right)} = \frac{\omega_r \times \frac{\partial u\left(c_{r',t},\zeta_{r,t}\right)}{\partial c}}{\omega_{r'} \times \frac{\partial u\left(c_{r',t},\zeta_{r',t}\right)}{\partial c}}\left[1 + \gamma\left(c_{r',t},\zeta_{r,t}\right)\frac{c_{r,t}-c_{r',t}}{c_{r',t}}\right].$$

When there is heterogeneity within a group of individuals retiring at the same age, we need to correct for the covariances between the welfare weights ω_i, marginal utility of consumption $\frac{\partial u\left(c_{r',t},\zeta_{i,t}\right)}{\partial c}$, the curvature $\gamma\left(c_{r',t},\zeta_{i,t}\right)$ and the consumption drop $\frac{c_{i,t}-c_{r',t}}{c_{r',t}}$ when expressing the average of the product of these terms as a function of the product of the average of these terms (see Andrews and Miller [2013]).

The consumption-drop implementation follows from a Taylor approximation for $\frac{\partial u(c_{r,t},\zeta_{r,t})}{\partial c}$ around $\left(c_{r,pre},\zeta_{r,pre}\right)$,

$$\frac{\partial u\left(c_{r,t},\zeta_{r,t}\right)}{\partial c} \cong \frac{\partial u\left(c_{r,pre},\zeta_{r,t}\right)}{\partial c}\left[1 + \frac{-\frac{\partial^2 u\left(c_{r,pre},\zeta_{r,t}\right)}{\partial c^2}c_{r,pre}}{\frac{\partial u\left(c_{r,pre},\zeta_{r,t}\right)}{\partial c}}\frac{c_{r,pre}-c_{r,t}}{c_{r,pre}}\right],$$

where we denote the relative risk aversion parameter again by $\gamma\left(c_{r,pre},\zeta_{r,t}\right)$. We again assume that the only relevant heterogeneity occurs across retirement ages. Hence, we now have

$$\frac{E\left(\omega_i\beta^t\frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial c}\Big|\, r_i = r\right)}{E\left(\omega_i\beta^t\frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial c}\Big|\, r_i = r'\right)} = \frac{\omega_r \times \frac{\partial u\left(c_{r,pre},\zeta_{r,t}\right)}{\partial c}\times\left[1+\gamma\left(c_{r,pre},\zeta_{r,t}\right)\frac{c_{r,pre}-c_{r,t}}{c_{r,pre}}\right]}{\omega_{r'}\times\frac{\partial u\left(c_{r',pre},\zeta_{r',t}\right)}{\partial c}\times\left[1+\gamma\left(c_{r',pre},\zeta_{r',t}\right)\frac{c_{r',pre}-c_{r',t}}{c_{r',pre}}\right]}$$

$$= \frac{\omega_r \times \frac{\partial u\left(c_{r,pre},\zeta_{r,pre}\right)}{\partial c}\times \phi_{r,t}/\phi_{r,pre}\times\left[1+\gamma\left(c_{r,pre},\zeta_{r,t}\right)\frac{c_{r,pre}-c_{r,t}}{c_{r,pre}}\right]}{\omega_{r'}\times\frac{\partial u\left(c_{r',pre},\zeta_{r',pre}\right)}{\partial c}\times \phi_{r',t}/\phi_{r',pre}\times\left[1+\gamma\left(c_{r',pre},\zeta_{r',t}\right)\frac{c_{r',pre}-c_{r',t}}{c_{r',pre}}\right]}.$$

To compare these two implementations, we could separate out the difference in marginal utility due to pre-retirement consumption differences too. Applying another Taylor expansion for the pre-retirement consumption levels gives

$$\frac{E\left(\omega_i\beta^t\frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial c}\Big|\, r_i = r\right)}{E\left(\omega_i\beta^t\frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial c}\Big|\, r_i = r'\right)} = \frac{\omega_r}{\omega_{r'}}\frac{\frac{\partial u\left(c_{r',pre},\zeta_{r,pre}\right)}{\partial c}}{\frac{\partial u\left(c_{r',pre},\zeta_{r',pre}\right)}{\partial c}}\times\frac{\phi_{r,t}/\phi_{r,pre}}{\phi_{r',t}/\phi_{r',pre}}$$

$$\times\left[1+\gamma\left(c_{r',pre},\zeta_{r,pre}\right)\frac{c_{r',pre}-c_{r,pre}}{c_{r',pre}}\right]\times\frac{1+\gamma\left(c_{r,pre},\zeta_{r,t}\right)\frac{c_{r,pre}-c_{r,t}}{c_{r,pre}}}{1+\gamma\left(c_{r',pre},\zeta_{r',t}\right)}\frac{c_{r',pre}-c_{r',t}}{c_{r',pre}}.$$

MPC Implementations For the MPC implementation, we rely on the MPC approach proposed by Landais and Spinnewijn [2021]. To illustrate their approach, we denote by $\tilde{\zeta}_{i,t}(\in \zeta_{i,t})$ the resource used at the margin to increase consumption $c_{i,t}$. This could for example be future consumption or other earnings in the household. $p_{i,t}$ units of $\tilde{\zeta}_{i,t}$ translate into one unit of consumption. The price $p_{i,t}$ can thus be interpreted as the shadow price of consumption and is allowed to differ across individuals. The optimizing behavior of a worker implies

$$(36) \qquad \frac{\partial u\left(c_{i,t}, \zeta_{i,t}\right)}{\partial c} + p_{i,t} \times \frac{\partial u\left(c_{i,t}, \zeta_{i,t}\right)}{\partial \tilde{\zeta}_{i,t}} = 0.$$

From the implicit differentiation of this optimality condition, we can derive the marginal propensity of consumption smoothing with respect to state-specific income $y(\pi_{i,t})$ for any $\pi_{i,t}$. Assuming separable preferences as in Landais and Spinnewijn [2021], we can obtain:

$$(37) \qquad \frac{\frac{dc_{i,t}}{dy_{i,t}}}{1 - \frac{dc_{i,t}}{dy_{i,t}}} = p_{i,t} \times \frac{\frac{\partial^2 u(c_{i,t}, \zeta_{i,t})}{\partial \tilde{\zeta}^2} \Big/ \frac{\partial u(c_{i,t}, \zeta_{i,t})}{\partial \tilde{\zeta}}}{\frac{\partial^2 u(c_{i,t}, \zeta_{i,t})}{\partial c^2} \Big/ \frac{\partial u(c_{i,t}, \zeta_{i,t})}{\partial c}} \simeq p_{i,t}.$$

Hence, the marginal propensity to consume is increasing in the price $p_{i,t}$ and decreasing in the curvature of preferences over consumption relative to the curvature of preferences over the resources used to smooth consumption at the margin. Our implementation assumes that the only relevant heterogeneity occurs across retirement ages, but the approximation also relies on the relative curvature in preferences to be similar across individuals with different retirement ages. Landais and Spinnewijn [2021] consider the MPC approach in the context of unemployment insurance. They show how this approximation is exact when individuals have CARA preferences and use future consumption (through their savings or by taking out loans) to smooth current consumption at the margin. In this case, the price depends on the interest rate the individual faces. In the context of retirement, the use of bequests become relevant as studied by Lockwood [2018] and showing that preferences over bequests are less curved. The price of using bequests, however, is again the interest rate. Hence, if some individuals/groups are more likely to use bequests at the margin, this depresses their MPC and we would wrongly attribute this to their price of consumption being lower.

Combining equations (36) and (37), we then obtain:

$$\frac{\frac{\partial u(c_{r,t}, \zeta_{r,t})}{\partial c}}{\frac{\partial u\left(c_{r',t}, \zeta_{r',t}\right)}{\partial c}} = \frac{\frac{\frac{dc_{r,t}}{dy_{r,t}}}{1 - \frac{dc_{r,t}}{dy_{r,t}}}}{\frac{\frac{dc_{r',t}}{dy_{r',t}}}{1 - \frac{dc_{r',t}}{dy_{r',t}}}} \times \frac{\frac{\partial u(c_{r,t}, \zeta_{r,t})}{\partial \tilde{\zeta}_{r,t}}}{\frac{\partial u\left(c_{r',t}, \zeta_{r',t}\right)}{\partial \tilde{\zeta}_{r',t}}} \simeq \frac{\frac{\frac{dc_{r,t}}{dy_{r,t}}}{1 - \frac{dc_{r,t}}{dy_{r,t}}}}{\frac{\frac{dc_{r',t}}{dy_{r',t}}}{1 - \frac{dc_{r',t}}{dy_{r',t}}}}.$$

The approximation in the MPC implementation relies on the marginal cost of using resources to increase consumption to be similar across retirement age groups, i.e., $\frac{\partial u(c_{r,t}, \zeta_{r,t})}{\partial \tilde{\zeta}_{r,t}} \simeq \frac{\partial u\left(c_{r',t}, \zeta_{r',t}\right)}{\partial \tilde{\zeta}_{r',t}}$. Landais and Spinnewijn [2021] propose this MPC implementation to compare within-individual

differences in marginal utility when empoyed vs. unemployed and argue that it is likely to have $\frac{\partial u(c_{u,t},\zeta_{u,t})}{\partial \tilde{\zeta}_{u,t}} > \frac{\partial u(c_{e,t},\zeta_{e,t})}{\partial \tilde{\zeta}_{e,t}}$ if $p_{u,t} > p_{e,t}$. Indeed, when hit by unemployment, an individual faces lower income and is more reliant on other resources to increase her income. Unemployment is therefore likely to increase the shadow price of consumption, but also the disutility of using more resources to smooth consumption. When comparing the MPC's across individuals instead, we also need to factor in a substitution effect, implying that individuals facing higher $p_{r,t}$ may reduce their use of this resource to smooth consumption. The approximation in the MPC implementation will thus depend on how big these potentially offsetting effects are. We refer the interested reader for more discussion on robustness and extensions to Landais and Spinnewijn [2021].

Fiscal Externality of Steeper Incentives Consider a budget-balanced reform at retirement age $\tilde{r}$ with $db_{r,t} = db_{r>\tilde{r},t}$ for $r > \tilde{r}$ and $db_{r,t} = db_{r\leq\tilde{r},t}$ for $r \leq \tilde{r}$ with $db_{r>\tilde{r},t} = -\frac{1-S(\tilde{r})}{S(\tilde{r})}db_{r\leq\tilde{r},t}$. For simplicity, we drop the age subscript t. Using $T_r = \frac{\tau_{r'}}{R^{r'}} - (NPV_{r'+1} - NPV_{r'})$, we can express the impact on social welfare as:

$$
\begin{aligned}
dW &= (1 - S(\tilde{r}))\, SMU_{r\leq\tilde{r}}db_{r\leq\tilde{r}} + S(\tilde{r})\, SMU_{r>\tilde{r}}db_{r>\tilde{r}} \\
&\quad - \lambda(1 - S(\tilde{r}))\left[1 - \Sigma_{r'}T_{r'}\frac{\partial S(r')}{\partial b_{r\leq\tilde{r}}}\frac{1}{1 - S(\tilde{r})}\right]db_{r\leq\tilde{r}} \\
&\quad - \lambda S(\tilde{r})\left[1 - \Sigma_{r'}T_{r'}\frac{\partial S(r')}{\partial b_{r>\tilde{r}}}\frac{1}{S(\tilde{r})}\right]db_{r>\tilde{r}} \\
&= S(\tilde{r})\,db_{r>\tilde{r}}\left[SMU_{r>\tilde{r}} - SMU_{r\leq\tilde{r}}\right] + \lambda\left[\Sigma_{r'}T_{r'}\left[\frac{\partial S(r')}{\partial b_{r\leq\tilde{r}}}db_{r\leq\tilde{r}} + \frac{\partial S(r')}{\partial b_{r>\tilde{r}}}db_{r>\tilde{r}}\right]\right].
\end{aligned}
$$

The second equality uses the budget-neutrality of the reform. We now make the following assumptions regarding the response of the survival rates to changes in the benefit policy.

- *Assumption 1:* for any $\tilde{r}$, $\frac{\partial S(r)}{\partial b_{r\leq\tilde{r}}} \cong 0$ for $r > \tilde{r}$; $\frac{\partial S(r)}{\partial b_{r>\tilde{r}}}$ for $r \leq \tilde{r}$

- *Assumption 2:* for any $\tilde{r}$, $\Sigma_{r'\leq\tilde{r}}\frac{\partial S(r')}{\partial b_{r\leq\tilde{r}}} \cong \Sigma_{r'>\tilde{r}}\frac{\partial S(r')}{\partial b_{r>\tilde{r}}}\frac{1-S(\tilde{r})}{S(\tilde{r})}$ and $T_{\tilde{r}} \cong T$

- *Assumption 3:* for any $\tilde{r}$, $-\frac{\partial S(\tilde{r})}{\partial b_{r\leq\tilde{r}}} = \frac{\partial S(\tilde{r})}{\partial b_{r>\tilde{r}}} \cong \frac{\partial S(\tilde{r})}{\partial w_{\tilde{r}}}$

Assumption 1 follows from small changes in the policy for given retirement ages only affecting individuals who are at the margin of retiring at those ages. Assumption 2 is weaker than the assumption that income effects do not matter. Instead it assumes that for a budget-balanced change in the profile, the negative income effect on the retirement of early retirees is equal to the positive income effect on the retirement of late retirees. Assumption 3 relies on the fact that the change in the survival rate at $\tilde{r}$ only depends on the change in local slope of the pension profile $d[b_{r>\tilde{r}} - b_{r\leq\tilde{r}}]$ and thus that locally income effects are small relative to substitution effects.

We can approximate the welfare impact in the following three steps using Assumptions 1-3

respectively:

$$
\begin{aligned}
dW \;\cong\; & S\left(\tilde{r}\right)db_{r>\tilde{r}}\left[SMU_{r>\tilde{r}}-SMU_{r\le\tilde{r}}\right]+\lambda\left[\Sigma_{r'\le\tilde{r}}T_{r'}\frac{\partial S\left(r'\right)}{\partial b_{r\le\tilde{r}}}db_{r\le\tilde{r}}+\Sigma_{r'>\tilde{r}}T_{r'}\frac{\partial S\left(r'\right)}{\partial b_{r>\tilde{r}}}db_{r>\tilde{r}}\right] \\
\cong\; & S\left(\tilde{r}\right)db_{r>\tilde{r}}\left[SMU_{r>\tilde{r}}-SMU_{r\le\tilde{r}}\right]+\lambda T_{\tilde{r}}\left[\frac{\partial S\left(\tilde{r}\right)}{\partial b_{r>\tilde{r}}}db_{r>\tilde{r}}+\frac{\partial S\left(\tilde{r}\right)}{\partial b_{r\le\tilde{r}}}db_{r\le\tilde{r}}\right] \\
\cong\; & S\left(\tilde{r}\right)db_{r>\tilde{r}}\left[SMU_{r>\tilde{r}}-SMU_{r\le\tilde{r}}\right]+\lambda T_{\tilde{r}}\frac{\partial S\left(\tilde{r}\right)}{\partial w_{\tilde{r}}}\left[db_{r>\tilde{r}}-db_{r\le\tilde{r}}\right].
\end{aligned}
$$

We can finally rewrite and re-express the welfare impact in terms of elasticities:

$$
\begin{aligned}
dW \;\cong\; & \left(\tilde{r}\right)db_{r>\tilde{r}}\left\{\left[SMU_{r>\tilde{r}}-SMU_{r\le\tilde{r}}\right]+\lambda T_{\tilde{r}}\frac{\partial S\left(\tilde{r}\right)}{\partial w_{\tilde{r}}}\frac{1}{S\left(\tilde{r}\right)}\left[1-\frac{db_{r\le\tilde{r}}}{db_{r>\tilde{r}}}\right]\right\} \\
=\; & S\left(\tilde{r}\right)db_{r>\tilde{r}}\left\{\left[SMU_{r>\tilde{r}}-SMU_{r\le\tilde{r}}\right]+\lambda\frac{T_{\tilde{r}}}{w_{\tilde{r}}}\left[\varepsilon_{S(\tilde{r}),w_{\tilde{r}}}-\varepsilon_{1-S(\tilde{r}),w_{\tilde{r}}}\right]\right\} \\
=\; & S\left(\tilde{r}\right)db_{r>\tilde{r}}\left\{\left[SMU_{r>\tilde{r}}-SMU_{r\le\tilde{r}}\right]+\lambda\frac{T_{\tilde{r}}}{w_{\tilde{r}}}\varepsilon_{\frac{S(\tilde{r})}{1-S(\tilde{r})},w_{\tilde{r}}}\right\}
\end{aligned}
$$

Normalizing with respect to the social marginal utility of individuals retiring at the normal retirement age and assuming $SMU_{NRA}\cong\lambda$, we have that the net welfare return, expressed in monetary terms, of a dollar of pension benefits taken from early retirees ($r\le\tilde{r}$) and given to late retirees ($r>\tilde{r}$) is equal to:

$$
\Delta W_{\tilde{r}}=\frac{dW/\left[S\left(\tilde{r}\right)db_{r>\tilde{r}}\right]}{SMU_{NRA}}\cong\frac{T_{\tilde{r}}}{w_{\tilde{r}}}\times\varepsilon_{\frac{S(\tilde{r})}{1-S(\tilde{r})},w_{\tilde{r}}}-\frac{SMU_{r\le\tilde{r}}-SMU_{r>\tilde{r}}}{SMU_{NRA}}.
$$

Empirical Inputs	Economic Interpretation	Assumptions	Challenges
		Consumption-Level Implementation – Equation (6)	
$E_{r>\tilde{r}}(c)$, $E_{r\le\tilde{r}}(c)$: Average consumption levels of individuals retiring before vs after $\tilde{r}$	Captures both the redistributive and insurance value of profile reform	Homogeneous relative risk aversion γ	Measuring γ
		$\omega_r \frac{\partial u(\bar{c},\zeta_{r,t})}{\partial c}$ constant across retirement ages r	Gauging selection into retirement ages based on *SMU* of consumption, driven by ω_r or $\zeta_{r,t}$
		Taylor approximation (Chetty [2006])	
		Heterogeneity within retirement age group negligible (Andrews and Miller [2013])	
		Consumption-Drop Implementation – Equation (7)	
$\Delta c_{r>\tilde{r}}$, $\Delta c_{r\le\tilde{r}}$: Average drop in consumption around retirement of individuals retiring before vs after $\tilde{r}$	Captures only the insurance value of profile reform	Homogeneous relative risk aversion γ	Measuring γ
		$\omega_r \frac{\partial u(c_{r,pre},\zeta_{r,t})}{\partial c}$ constant across retirement ages r	Gauging selection into retirement ages based on *changes* in *SMU* of consumption around retirement, driven by $\frac{\zeta_{r,t}}{\zeta_{r,pre}}$
		Taylor approximation (Chetty [2006])	
		Heterogeneity within retirement age group negligible (Andrews and Miller [2013])	
		MPC Implementation – Equation (23)	
$mpc_{r>\tilde{r}}$, $mpc_{r\le\tilde{r}}$: Average marginal propensity to consume in retirement of individuals retiring before vs after $\tilde{r}$	Captures the liquidity value of profile reform	Constant relative curvature of u over consumption c and resources in ζ across retirement ages (Landais and Spinnewijn [2021])	Finding exogenous unanticipated income shocks to identify MPCs across retirement ages
		Heterogeneity within retirement age group negligible (Andrews and Miller [2013])	

Notes: The table summarizes our three proposed empirical implementations for the measurement of the social marginal value $\frac{SMU_{r\le\tilde{r}}}{SMU_{r>\tilde{r}}}$ of steepening the pension profile at age $\tilde{r}$. We consider a marginal and budget-balanced steepening of the pension profile at a given retirement age $\tilde{r}$ by reducing pensions for individuals retiring before age $\tilde{r}$ by some small amount $db_{r\le\tilde{r}}$, and increasing them for individuals retiring after age $\tilde{r}$ by $db_{r>\tilde{r}}$. For each implementation, we provide the empirical inputs necessary to measure the social marginal value of the reform, and the assumptions and challenges involved. See sections I and VI for details.

Appendix H Welfare Implementation Details

This appendix provides further detail on the welfare implementation described in Section VI and illustrated in Figure 10 and Tables 2 and H-1 (below). We estimate the consumption smoothing costs for budget-neutral reforms that steepen the pension profile. The terms correspond to the welfare effects of transferring a dollar for individuals retiring *before* a specific age to individuals retiring *after* that age. The values we obtain can then be compared with the fiscal externality to compute the net welfare effect. Below we also provide a back-of-the envelope calculation showing that a fiscal externality of .15 is a reasonable benchmark to evaluate the net welfare gain.

Table H-1: CONSUMPTION SMOOTHING COST OF INCENTIVIZING LATER RETIREMENT

	Baseline		Sensitivity			Alternative
	Cons. levels	ΔC	Risk aversion	Health Dep.	Welfare Wgts	MPC
	$\gamma = 4, \omega = 1$	$\gamma = 4$	$\gamma = 2, \omega = 1$	$SMU \sim$ Health	$\gamma = 4, \omega \sim$ Life Exp.	
	(1)	(2)	(3)	(4)	(5)	(6)
A. Age-Specific Profile Change						
$\tilde{r} = 60$	0.34	0.21	0.17	0.20	0.41	-0.27
$\tilde{r} = 63$	0.28	0.12	0.14	0.18	0.33	0.08
$\tilde{r} = 65$	0.76	0.14	0.38	0.70	0.79	0.76
B. Swedish Pension Reform						
	0.37	0.14	0.18	0.30	0.40	0.21

Notes: This table presents estimates of the consumption smoothing cost of incentivizing later retirement. Panel A considers three age-specific reforms, while Panel B considers the Swedish pension reform as described in Appendix H.1 below. Column (1) repeats the results from the consumption levels implementation, using the difference in consumption levels to approximate the the difference in SMU's (see equation (6)), while (2) shows uses the difference in consumption drops (see equation (7)). Columns (3), (4) and (5) explore the sensitivity of the consumption levels results: (3) considers a change of the curvature in preferences, (4) allows for health-dependent marginal utility, and (5) assigns welfare weights that depend on life expectancy. Finally, column (6) shows the results for the alternative implementation using the difference in MPC's (see equation (23)).

Appendix H.1 Consumption Smoothing Cost

We first describe in detail how we approximate the consumption smoothing cost for the age-specific policies. This reform involves a steepening of the pension profile at a given retirement age $\tilde{r}$ by reducing pensions for individuals retiring before age $\tilde{r}$ by some small amount $db_{r \leq \tilde{r}}$, and increasing them for individuals retiring after age $\tilde{r}$ by $db_{r > \tilde{r}}$. Budget balance requires that $db_{r > \tilde{r}} = -\frac{1 - S(\tilde{r})}{S(\tilde{r})} db_{r \leq \tilde{r}}$, where $1 - S(\tilde{r})$ is the share of individuals who retired before age $\tilde{r}$. The consumption smoothing cost per dollar transferred then equals

$$\frac{SMU_{r \leq \tilde{r}} - SMU_{r > \tilde{r}}}{SMU_{NRA}}.$$

We can in principle implement this change in benefits for individuals at any given age t, but in our implementation we only use the consumption years we observe after retirement in our baseline sample. For brevity, we drop the age subindices.

Baseline Implementations Figure 10 and columns (1) of Table 2 and Table H-1 follow the baseline implementation using the difference in consumption levels across retirement age groups, relative to the normal-retirement age group, scaled by the relative risk aversion,

$$(38) \qquad \frac{SMU_{r \leq \tilde{r}} - SMU_{r > \tilde{r}}}{SMU_{NRA}} \approx \gamma \times \left[\frac{E_{r > \tilde{r}}(c)}{E_{r \in NRA}(c)} - \frac{E_{r \leq \tilde{r}}(c)}{E_{r \in NRA}(c)} \right].$$

We obtain the estimates of the consumption levels for people retiring at age r relative to normal retirees, using regression (8). For each age r, we approximate the consumption smoothing cost of steepening the profile at age r as the difference in the weighted average of consumption levels for people above age r and this same difference for people below age r. The weights used are the fraction of people at each retirement age. The consumption smoothing cost is obtained by multiplying the value obtained by γ, for which we set the baseline value at 4 (see Landais and Spinnewijn [2021]). The grey bars in Figure 10 show these values for each age. In Table 2, we are only reporting the results for implementing the reform at the cut-off ages between the different retirement-age groups (i.e., ages 60, 63 and 65). This is repeated in Table H-1.

Column (2) shows the results applying the baseline consumption drops implementation in equation (11),

$$(39) \qquad \frac{SMU_r}{SMU_{NRA}} \cong \frac{1 + \gamma \frac{c_{r,pre} - c_{r,t}}{c_{r,pre}}}{1 + \gamma \frac{c_{NRA,pre} - c_{NRA,t}}{c_{NRA,pre}}},$$

where we assumed a $\gamma = 4$ and that the welfare weights multiplied by the marginal utility of consumption before retirement are equal across retirement ages. The numbers we use for the consumption drops come from Figure 4. Following the equation above, for each retirement-age group, we normalise the SMU by the value for normal retirees. Then, for each age cut-off age $\tilde{r}$, we obtain the consumption smoothing cost of steepening the profile at age $\tilde{r}$ by taking the difference between the weighted average of these rescaled values for the retirement-age groups above $\tilde{r}$, where the weights are the fraction of people in each group and this same weighted average for people below $\tilde{r}$. We assume again $\gamma = 4$. The results are presented in Table 2 and repeated in Table H-1 column (2).

Sensitivity Analysis Columns (3), (4) and (5) of Table H-1 present results when making alternative assumptions on the curvature in consumption preferences, on the sensitivity of the SMU to health, and on the welfare weights, respectively.

Column (3) is obtained by applying the same method as for the baseline levels implementation, but reducing the curvature in consumption preferences to $\gamma = 2$.

Column (4) explores the robustness of the consumption smoothing cost estimates from the levels implementation to health-dependence in the marginal utility of consumption. Following the results in Finkelstein, Luttmer and Notowidigdo [2013], we assume that the semi-elasticity of the marginal utility of consumption to changes in a bad health measure, measured in stan-

dard deviations, is -0.25. To account for this, we adjust the SMU terms as follows:

$$(40) \qquad \frac{SMU_r}{SMU_{NRA}} \cong \left(1 + \gamma \frac{c_{NRA} - c_r}{c_{NRA}}\right)(1 - 0.25(H_r - H_{NRA})),$$

where H_r is the estimated bad health measure for individuals retiring at age r, measured in standard deviations. We use as our health measure the pooled health index described in Appendix E.[58] The differences $H_r - H_{NRA}$ are estimated for the retirement age groups as in Figure 7, i.e. using specification (8) with the health measure as the outcome variable. The reported consumption smoothing cost is simply the difference of the weighted averages of the SMUs for individuals below and above the cutoff age.

Column (5) presents a sensitivity analysis in which we assign welfare weights to each retirement age r that depend on life-expectancy. We follow Chetty et al. [2016] to estimate the life expectancy and Becker, Philipson and Soares [2005] to adjust the welfare weights.
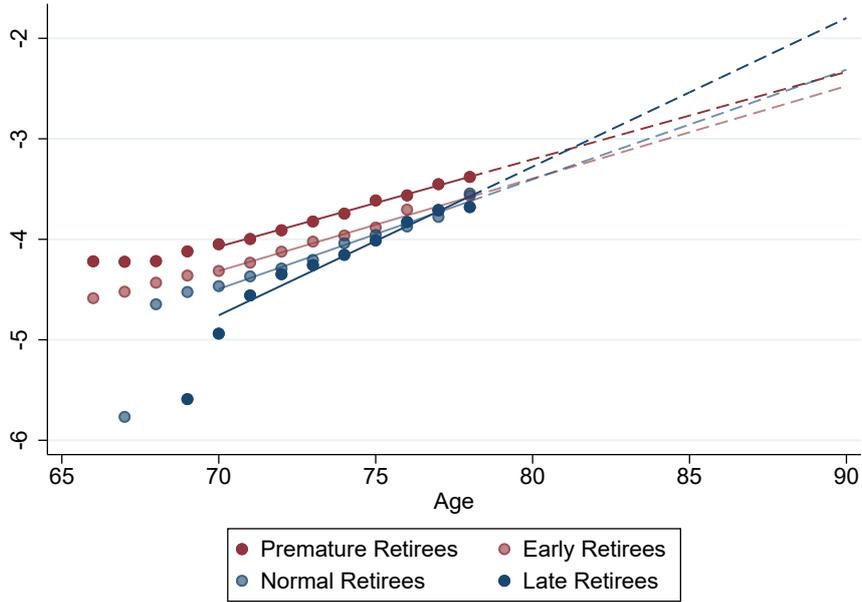
For each retirement age group, we can compute the mortality rate at each age t, defined as the number of people who were alive at $t - 1$ but died at age t divided by the number of people who are alive at age t. Since the mortality register provides death years up until 2017, we will assume that all the people who have a missing death year are alive in 2017.

For the ages $[66; 78]$, we simply calculate the empirical mortality rates in the different retirement age groups, as illustrated in Figure H-1. To obtain mortality rates at higher ages, we implement a Gompertz extrapolation for each retirement age group. Specifically, we run the regression: $ln(mortality) = \alpha + \beta age + \epsilon$. We restrict the regression sample to the mortality values for ages $[70; 78]$ given that up to 69 the mortality rates are mechanically different for the different retirement age groups by definition. This is shown in Figure H-1. We then compute the expected life expectancy at 65 using the true mortality rates in the range $[65; 78]$ and the estimated ones in the range $[79; 90]$.

---

[58]We note that Finkelstein, Luttmer and Notowidigdo [2013] used the number of chronic diseases as their preferred health measure. We observe a similar health measure in the ULF survey but not the HEK survey. Panel A of Figure E-1 suggests that the health gradient across groups would be less steep if we used the number of Long-term illnesses instead of the pooled bad health index. Using this alternative health measure would therefore bring the consumption smoothing effects in column (4) of Table H-1 closer to the original estimate in column (1), effectively making our estimated consumption smoothing effects less sensitive to accounting for health.

Figure H-1: TRUE AND INTERPOLATED MORTALITY VALUES FOR EACH RETIREMENT AGE
GROUP

Notes: This figure plots the true mortality rates (dots) and the imputed mortality rates (line) using a Gompertz extrapolation, for each retirement age group. For the extrapolation, we consider only the computed mortality rates in the range $[70; 78]$ (solid line). The mortality rates from the dashed line are then used to compute the expected discounted lifetime by retirement age group.

The goal of this sensitivity analysis is to compute compensating consumption differentials that would equalize the expected lifetime utility for individuals with different retirement ages and use these compensating differentials to adjust the SMU's. This is done by computing Δx_j, for each retirement age group j in the formula below:

$$(41) \qquad \sum_{k=65}^{90} S_{k,NRA}\beta^k u(\bar{c}) = \sum_{k=65}^{90} S_{k,j}\beta^k u(\bar{c} + \Delta x_j),$$

where $S_{k,j}$ is the survival rate at k for retirement age group j. Formally, $S_k = \prod_{i=0}^{k}(1 - m_i)$, where m_i is the mortality rate at age k we computed above. Assuming CRRA preferences, we can approximate:

$$(42) \qquad \sum_{k=65}^{90} S_{k,NRA}\beta^k = \sum_{k=65}^{90} S_{k,j}\beta^k (1 + \gamma \Delta x_j)$$

which simplifies to:

$$(43) \qquad \gamma \Delta x_j = \frac{\sum_{k=65}^{90} S_{k,NRA}\beta^k - \sum_{k=65}^{90} S_{k,j}\beta^k}{\sum_{k=65}^{90} S_{k,j}\beta^k},$$

which corresponds to the relative difference in expected discounted lifetimes.

We then obtain a value for the consumption smoothing cost by applying the same method

126

Figure H-2: LINEAR SPLINES FOR THE CONSUMPTION DROPS AND MPC IMPLEMENTATIONS

A. CONSUMPTION DROPS

B. MPC
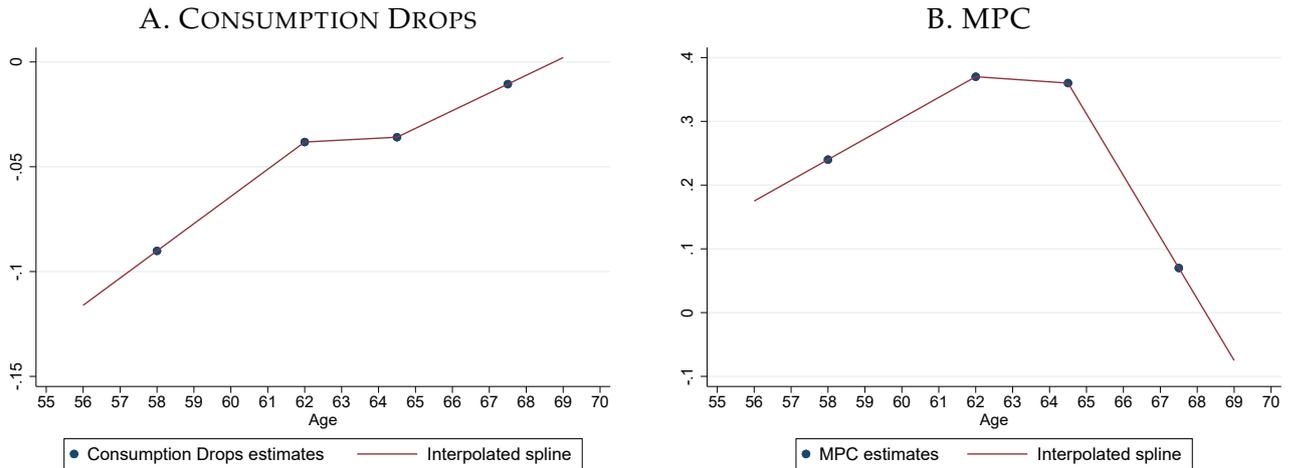
Notes: Panel A presents the consumption drops estimates from Figure 4 for the four retirement age groups (dots) and the interpolated linear spline between each of them. The consumption drop estimate for each retirement age group is assumed to lie at the midpoint of the interval. For instance, for the premature retirees (age range [56;60]) we assign the consumption drop to 58. We obtain age-specific values by interpolating a linear spline between each point (solid line).
Panel B replicates this same approach using the MPC values from Table F-3.

as for the baseline levels implementation, except that we now subtract from the consumption level the Δx_t term for each age t. Intuitively, if retirement-age group j has lower life expectancy, then Δx_j represents how much we need to increase consumption for that group to compensate them for the lower expected lifetime. We then subtract this value from their actual consumption level to obtain a corresponding increase in the SMU. The results are shown in Table H-1 column (4).

Alternative MPC Implementation   Column (6) in Table H-1 present the results for the MPC implementation described in Appendix Appendix F. Following equation (23),

$$(44) \qquad \frac{SMU_r}{SMU_{NRA}} \cong \frac{\frac{mpc_r}{1-mpc_r}}{\frac{mpc_{NRA}}{1-mpc_{NRA}}},$$

assuming now that welfare weights are similar across retirement ages. For the marginal propensities to consume, we take values from Table F-3. We then compute the odds ratio of the marginal propensities to consume, rescaled by the odds ratio of the normal retirees, following equation (11). Similar as above, we obtain the consumption smoothing cost of steepening the profile at age $\tilde{r}$ by taking the difference between the weighted average of these rescaled values for the retirement-age groups above $\tilde{r}$, where the weights are the fraction of people in each group and this same weighted average for people below $\tilde{r}$.

Swedish Pension Reform   Panel B of Table H-1 shows the welfare effects of the change in slope of the pension profile due to the Swedish pension reform. That is, we compute a profile that has the same slope as the NDC scheme but the same budget as the ATP scheme, denoted by $N\hat{D}C$, as described in Appendix A.2.5. The consumption smoothing cost of this reform, per

127

dollar transferred from early to late retirees, equals:

$$\Sigma_r \mu_r \frac{SMU_r}{SMU_{NRA}},$$

where $\mu_r = \frac{f(r)\left(\widehat{NDC}_r - ATP_r\right)}{\Sigma_{r \leq \bar{r}} f(r)\left(\widehat{NDC}_r - ATP_r\right)} / S(\bar{r})$, where $\bar{r}$ is the retirement age at which the pension profiles intersect (i.e., $\widehat{NDC}_{\bar{r}} = ATP_r$). The weights are thus composed of the product of (i) the relative frequency of the retirement age group (ii) the difference between the new pension profile $N\hat{D}C$ and the ATP one. This sum is then rescaled by the total value of the pension dollars taken away from the early retirees and given to the late retirees. Note that this formulation corresponds to $\frac{SMU_{r \leq \bar{r}} - SMU_{r > \bar{r}}}{SMU_{NRA}}$ when using the age-specific pension reforms considered above.

The age-specific estimates we take for the consumption-level implementations in columns (1), (3), (4), and (5) are again the ones reported in Figure 10. For instance, for the baseline levels implementation in column (1), we will take the consumption levels from regression (8). Since we have only estimated health differences, consumption drops and MPCs at the retirement age group level, we obtain age-specific values by interpolating a linear spline, as shown in Figure H-2 Panels A and B.

Heterogeneity Analysis Section VI briefly reported on some heterogeneity analysis, to account for the fact that the reform had a differential impact on different categories of people. Table H-2 presents these results, which all follow our baseline consumption levels implementation, but using the estimates from regression (8) restricted to the relevant sample.

Column (1) reproduces our baseline results, i.e., using the baseline consumption levels implementation and baseline sample. Columns (2) and (3) restricts the sample to to the bottom and top decile of ATP points accumulated at 55, which corresponds to low income/short-career vs. high income/long-career individuals respectively.For the implementation of the Swedish pension reform, we also calculate the corresponding change in pension benefits, following Appendix Figure C-4. Columns (4) and (5) consider single people and cohabiting people respectively. Lastly, column (6) uses the consumption analysis for the baseline sample, but changes the definition of retirement, as described in footnote 37.

Table H-2: HETEROGENEITY ANALYSIS: CONSUMPTION SMOOTHING COST OF STEEPER PRO-FILE

	Levels Baseline (1)	Bottom 10% (2)	Top 10% (3)	Couples (4)	Singles (5)	UI/DI (6)
			A. Age-Specific Profile Change			
$\tilde{r} = 60$	0.34	0.16	0.19	0.30	0.49	0.12
$\tilde{r} = 63$	0.28	0.10	0.22	0.24	0.40	0.15
$\tilde{r} = 65$	0.76	0.45	0.46	0.71	0.88	0.75
			B. Swedish Pension Reform			
	0.37	0.21	0.21	0.20	0.30	0.18

Notes: This table shows the results of the heterogeneity analysis of the baseline consumption levels implementation. Column (1) replicates the estimates for the baseline levels implementation. Column (2) and column (3) produce estimates for the sample restricted to the bottom decile of ATP points accrued at age 55 and top decile respectively. Column (4) and (5) present the analysis restricting to couples and singles respectively, while column (6) replicates the baseline analysis redefining retirement for those who exit the labor market through UI/DI.

Appendix H.2 Fiscal Externality Benchmark

For the implementation of the fiscal externality, we use the approximation in equation (10) and assume that both $\varepsilon_{\frac{S(\tilde{r})}{1-S(\tilde{r})},w_{\tilde{r}}}$ and $\frac{T_{\tilde{r}}}{w_{\tilde{r}}}$ are age-independent. We then use $\varepsilon_{S(\tilde{r}),w_{\tilde{r}}} = .22$, which corresponds to the extensive labor supply elasticity estimated in Laun [2017] based on the labor supply responses of workers over 65 to age-specific earned income tax credits. Using $\frac{S(NRA)}{1-S(NRA)} = 0.53$, corresponding to the share of individuals retiring at 65 or later vs. before in our baseline sample, we then obtain

$$\varepsilon_{\frac{S(NRA)}{1-S(NRA)},w_{NRA}} = \varepsilon_{S(NRA),w_{NRA}} \left[1 + \frac{S(NRA)}{1-S(NRA)} \right] \approx 0.35.$$

We also take $\frac{T_{\tilde{r}}}{w_{\tilde{r}}} \cong \frac{T}{w} = 0.45$. This participation tax rate relies on the pension calculator from Appendix A. See in particular Figure A-9 and the supplementary discussion around Figure A-11. Hence, putting the two terms together we obtain a fiscal externality of 0.15. That is, we would gain 15 cents per dollar transferred from individuals retiring before $\tilde{r}$ to individuals retiring after $\tilde{r}$. Without non-pension social insurance benefits, we would obtain a participation tax rate of about 0.4 rather than 0.45 (see Figure A-9). A participation tax rate of .4 would reduce the fiscal externality to .13, which is a negligible difference for our purposes.

As briefly discussed in the main text, it seems a reasonable assumption that the fiscal externality is similar across retirement ages. Figure A-9 shows indeed that the participation tax is indeed stable across retirement ages. Regarding the labor supply elasticity and how this varies between early and late retirees, the literature rightly points out that this elasticity is not a structural parameter and depends on what portion of workers are near the margin of retirement at a given age (French [2005], French and Jones [2012], Blundell, French and Tetlow [2016]). Existing studies mostly point out how this matters for labor supply elasticities

at prime working age versus around retirement, rather than at early versus late retirement ages. As discussed, one would also wish to account for life-cycle dynamics and compositional effects, as later retirees are different from earlier retirees in ways that could matter for their labor supply elasticity (e.g. they are less subject to negative health shocks and have longer life expectancies). An attractive way to identify how the Frisch labor supply elasticity varies with $\tilde{r}$ is therefore to compare similar local variations in the profile of the net-present-value of pensions (such as kinks) at different retirement ages $\tilde{r}$, in the exact same context. This is exactly the type of variation leveraged in Seibold [2021] with more than 400 such local variations in the same German context. His results (cf. his Figure 5) indicate that the responses to similar local changes in the pension profile appear remarkably constant across retirement age groups. Furthermore, they suggest the absence of any systematic and significant heterogeneity in responsiveness across other observable characteristics (such as education, birth cohorts, lifetime earnings, unionization or health) that may correlate with retirement age.