Retirement Consumption and Pension Design

Jonas Kolsrud* Camille Landais† Daniel Reck‡ Johannes Spinnewijn§¶

October 29, 2022

Abstract

This paper leverages consumption data to evaluate the distributional effects of pension reforms. Using administrative data from Sweden, we show that on average workers who retire earlier have lower consumption while retired, experience larger drops in consumption around retirement, and show higher marginal propensities to consume out of wealth shocks. These findings imply that reforms incentivizing later retirement, as many countries have recently done, incur a substantial consumption smoothing cost. Accounting for selection on health and life expectancy further increases the redistributive cost of such reforms. Turning to other features of pension policy, we find that reforms that redistribute on the basis of early career labor supply would have opposite-signed redistributive effects, while differentiating on wealth may help to target pension benefits toward those who are vulnerable to larger drops in consumption around retirement.

*National Institute of Economic Research and Stockholm University, Box 12090, 10223 Stockholm, Sweden +46(0)8-4535932, jonas.kolsrud@konj.se.

†Dept of Economics, London School of Economics, Houghton Street London, WC2A 2AE, c.landais@lse.ac.uk

‡Dept of Economics, University of Maryland, Tydings Hall, 3114 Preinkert Dr, College Park, MD 20742; dreck@umd.edu

§Dept of Economics, London School of Economics, Houghton Street London, WC2A 2AE, j.spinnewijn@lse.ac.uk

¶We thank Richard Blundell, Peter Diamond, Eric French, Francois Gerard, Jon Gruber, Johannes Hagen, Xavier Jaravel, Atilla Lindner, Petra Persson and seminar participants at Bonn University, CEPR, Gothenburg University, IFS, LSE, Mannheim, NYU, Stanford University, UC Berkeley, UPenn, and Yale University for helpful discussions and suggestions. We also thank Francesco Armillei, Marion Brouard, Juliette Caucheteux, Henriette Druba, Sebastian Ernst, Miguel Fajardo-Steinhauser, Diego Ferreras Garrucho, Alice Lapeyre, and Quirin Von Blomberg for outstanding research assistance. We gratefully acknowledge support from ERC grants #679704 and #716485.

1 Introduction

Many countries have undertaken large reforms to their public pension systems over the past two decades, and more seem likely to follow suit in the near future. These reforms are perhaps the most substantial reforms to social insurance policy in the developed world over the last 20 years. Public discussion of pension reforms largely focused on restoring fiscal sustainability because of ageing populations. In particular, a common theme of the reforms taken in most countries – including Austria, Belgium, Canada, Denmark, France, Germany, Spain, Sweden, and the UK – has been to incentivize workers to retire later in life (see e.g., Gruber and Wise [1999], OECD [2019], Barr and Diamond [2009]).[1] Incentives to work longer have desirable fiscal effects: workers who retire later pay more tax. But a coincident feature of these reforms is that the burden of making the pension system fiscally sustainable falls more heavily on some workers (e.g. early retirees) than on others (e.g. late retirees). The welfare costs due to this aspect of pension reforms are not well understood.

In this paper, we evaluate the redistributive costs of pension reforms using consumption measures constructed from Swedish administrative data. We focus on questions involving the optimal within-cohort distribution of pension benefits, such as the relative amount of benefits provided to early versus late retirees. Doing so allows us to separate thorny questions about the overall generosity of pensions and whether they are funded or pay-as-you-go, about which much has been written, from questions about how pension benefits vary with the timing of retirement or other individual characteristics, about which comparatively little has been written. Empirically, we use rich data on consumption and other information covering the population of Sweden to examine whether and to what extent potential reforms to pension benefits schedules are progressive or regressive.

We begin by developing a theoretical framework, which accommodates the complex environment that comprises real-world pension policy and can be applied to characterize the welfare effects of virtually any change in pension benefits. We use this framework to guide our analysis of consumption data. Specifically, we characterize the welfare effects of budget-neutral changes to the pension benefits schedule that redistribute on the basis of some characteristic, such as the retirement age. Our main focus is on the direct consumption smoothing (redistributive) effect of such a reform, which depends on differences in the marginal value of a dollar of pension benefits across groups of individuals.[2] We map empirical differences in con-

[1] The precise manner in which countries changed their pension benefits schedules to incentivize later retirement varies. The most common characteristic of reforms is to tighten the link between lifetime earnings and benefit amounts, as in the change from a defined benefit to defined contribution pension scheme. We describe the components of the Swedish reform along these lines below. In countries where, unlike Sweden, pension claiming and job exit are closely linked, reforms sometimes incentivize later retirement by rewarding delays in claiming public pension benefits. Another common feature of recent reforms is to increase the minimum age at which one can claim public pension benefits, which typically incentivizes workers who would otherwise retire early to work longer. A final feature of recent reforms that has ambiguous effects on incentives to work, but may nevertheless induce later retirement, concerns changes to statutory retirement ages like the "Normal Retirement Age" (see Seibold [2021]; Gruber et al. [2022]). For further details, see OECD [2015, 2017, 2019].

[2] While earlier work focused on insurance against work longevity risk the pension system provides (Diamond and Mirrlees [1978]), individuals may choose to work longer or retire earlier for various reasons. Moreover, the pension system not only provides insurance against end-of-career shocks, but redistributes between individuals with different employment histories more generally.

sumption patterns across groups to differences in the marginal value of pension benefits across these same groups, building on prior literature relating patterns in consumption to the value of social insurance (Gruber [1997], Hendren [2017], Landais and Spinnewijn [2021]). Doing so allows us to estimate the direction and magnitude of the consumption smoothing effects of reforms like those incentivizing later retirement. As in other theory on social insurance (Baily [1978], Chetty [2006]), the optimal policy would trade off these consumption smoothing costs against the fiscal benefits of incentivizing later retirement, which have been the focus of comparatively more research (e.g., Staubli and Zweimüler [2013]; Manoli and Weber [2016]; Laun [2017]; Manoli and Weber [2018]; Gruber et al. [2022]; Seibold [2021]; Lalive et al. [2022]; Haller [2022]), and, potentially, behavioral internality effects [Mullainathan et al., 2012; Spinnewijn, 2015; Reck and Seibold, 2021].

We use administrative registry data from Sweden, and registry-based measures of household consumption (see Kolsrud et al. [2020]), to inform this trade-off between consumption smoothing and incentives. Due to the apparent global policy focus on incentives to work late in workers' careers, we begin by closely examining the consumption smoothing effects of these late-career incentives, turning later on to potential reforms to other dimensions of pension benefits. We therefore begin by attempting to understand whether earlier retirees have higher social marginal utilities of consumption than late retirees, and if so, just how much higher. Our first approach to answering this question is to compare consumption levels in retirement between workers who retired at different ages. We assess the assumptions underlying this approach with additional data, including consumption surveys, and we implement two additional approaches to identify the relevant consumption smoothing effect, based on consumption dynamics around retirement and marginal propensities to consume when retired, respectively.

Our empirical findings suggest that strengthening late-career incentives to work entails a substantial and potentially pivotal consumption smoothing cost. We estimate a steep gradient of consumption over the retirement age, with those retiring after 65 enjoying about 20% higher consumption than those retiring before age 60, evaluated at the same age. The estimated steepness of this gradient is robust to a number of measurement concerns. Other patterns in consumption data also suggest a large consumption smoothing cost of strengthening late-career work incentives. Those retiring before 60 experience nearly a 10 percent decline in consumption when they retire, while those retiring after 65 experience virtually no decline in consumption; indeed this differential drop in consumption explains a substantial portion of the difference in consumption in retirement between those two groups.[3] In addition, we find a substantial marginal propensity to consume out of wealth shocks for those retiring before 65, but we estimate a near-zero marginal propensity to consume for late retirees, suggesting that the liquidity value of pension benefits to late retirees is relatively low. We formally map each of these facts about consumption to the differential value of pension benefits to different groups and also gauge the potential for differential consumption preferences, either permanent or changing around retirement, using linked consumption surveys. We find that despite

[3]The drop in consumption at retirement has been widely studied and debated (e.g., Bernheim et al. [2001], Aguiar and Hurst [2005a], Battistin et al. [2009], Stephens and Toohey [2018]), but without considering how this drop varies by the retirement age.

relying on different premises these approaches tell a broadly consistent story.

While the overall consumption gradient between retirees at ages 55 to 70 is clearly positive, we also document a notable non-monotonicity between the early and normal retirement age (resp. 61 and 65). The consumption gradient is much flatter in this range and, in some specifications, negative. That is, individuals retiring between those ages have similar or higher consumption on average compared to individuals retiring at the normal retirement age. We conduct some supplementary analysis of consumption by retirement age using data from the US Health and Retirement Study (HRS) and the Survey of Health, Aging and Retirement in Europe (SHARE). The patterns in measures of consumption we estimate with these data are strikingly similar to our findings based on the Swedish population register data, including the non-monotonicity for individuals retiring in the years just before the normal retirement age. Hence, within this set of ages, incentivizing later retirement is arguably less costly than at other ages.

Despite the attention paid to this dimension of pension benefits in public debate, the age at retirement is just one input to public pension benefits. If making pensions more fiscally sustainable by adjusting these incentives is costly, one wonders if adjusting benefits along other margins might have different consumption smoothing effects. We examine two other inputs to pension benefits: early-career labor supply and income before retirement. Together with late-career labor supply (i.e. retirement behavior), early-career labor supply and income while working capture most of the within-cohort variation in public pension benefits. We also examine consumption by wealth, which one can view as a proxy for lifetime income or as a prospective evaluation of an asset test for pension benefits.

In contrast to workers with long careers late in life (i.e., late retirees), workers with long careers as of age 55 have about 12% lower consumption than workers with medium length careers as of age 55. In other words, while incentivizing work late in life specifically reallocates resources from relatively needy to relatively less needy workers, incentivizing work early in life does the opposite. Furthermore, an analysis of consumption dynamics reveals that, unlike with the retirement age dimension, the differences in consumption in retirement by career length as of age 55 are entirely driven by longer term differences in consumption rather than differences that emerge around retirement. We also estimate large positive consumption gradients with pre-retirement income and household wealth, with those in the top quartile of income or wealth enjoying over 40% more consumption during retirement than those in the bottom quartile. In the case of income, these differences in consumption pre-date retirement and we observe no differential drop in consumption at retirement across income groups. For wealth, however, we find a much larger drop in consumption at retirement in the bottom wealth quartile than in other groups. This result suggests that conditioning pension benefits on wealth in particular may help allocate pension benefits to those who value them most, though of course these benefits should be traded off against potential fiscal and/or internality effects.

In general, to evaluate reforms, we should compare the consumption smoothing welfare effects we estimate with the relevant fiscal externalities (and potentially welfare effects due to behavioral internalities). Our results suggest, for instance, that the consumption smoothing effect of incentivizing later retirement is negative, but how would they compare quantitatively to the relevant fiscal externality? To answer this question, we compare our estimated consumption smoothing costs to plausible values for the relevant fiscal externality, based on our analysis of the size of the relevant fiscal incentives in Sweden and prior estimates of the response of Swedes' retirement timing to those incentives (Laun [2017]). The size of the relevant behavioral elasticity and how it might vary with workers' age are uncertain, but our estimates suggest that the consumption smoothing costs exceed the fiscal benefits of incentivizing later retirement. Owing to the non-monotonicity in consumption over retirement ages, incentivizing later retirement at very early or very late ages is especially costly, while a doing so for ages 60 to 65 specifically can be desirable. The results therefore suggest the desirability of an S-shaped reform: flat incentives below age 60 and above age 65 and steep incentives between these ages. This contrasts with the recent Swedish reform that provided stronger incentives at all ages and in particular after 65. Naturally, some caution is warranted - when extrapolating these results to the optimal profile or beyond the Swedish context - as our analysis is local and conditional on the tax and transfer system in place.

In our empirical analysis, we supplement the facts about consumption with data on a rich set of observable characteristics and with more granular data related to consumption structure. This helps us to better understand the mechanisms that underlie the differences in consumption levels and in consumption dynamics. It also helps us assess the validity of the assumptions necessary to map estimated consumption patterns into welfare measures. For example, a number of factors suggests that the non-monotonicity in the consumption gradient over retirement ages is driven by relatively well-off married couples in which the secondary earner retires before the normal retirement age. Overall, however, individuals retiring later are not only more well-resourced, they are also in better health and face lower mortality risk. Studying the evolution of health around retirement, we find that health shocks in the years just before retirement are more prevalent for workers retiring very early. The strong health/mortality gradient over retirement ages contrasts sharply with the small differences in health status that we find across individuals with long vs short career length as of age 55. These findings generally reinforce our finding that incentivizing later retirement in particular incurs a substantial welfare cost.

Our work contributes to a sizable recent literature using the calculus of variations to characterize the welfare effects of reforms in terms of reduced-form sufficient statistics. This approach has proven useful for the analysis of other social insurance programs, especially unemployment insurance (Baily [1978], Chetty [2006]). Our framework builds on Kolsrud et al. [2018] who incorporated heterogeneity and dynamic considerations in this approach for the analysis of unemployment insurance. We extend this to the context of retirement, which proves particularly useful because of the dynamics inherent to the life-cycle and the important selection effects into retirement. A large literature has studied consumption smoothing over the life-cycle and into retirement in particular (Bernheim et al. [2001], Aguiar and Hurst [2005b]; see De Nardi et al. [2016], Jappelli and Pistaferri [2010] for reviews). Several papers have also aimed to uncover the importance of different determinants of retirement (see Blundell et al. [2016], French and Jones [2017] for reviews). Our conceptual framework allows one to connect virtually any feature of public pension policy to consumption moments and patterns of dynamic selection to be able to evaluate its value, which we illustrate by considering reforms along a number of dimensions of pension benefits. We also rely on recent advances in the estimation of the value of social transfers (e.g., Hendren [2017, 2020]; Fadlon and Nielsen [2019]; Deshpande and Lockwood [2022]; Landais and Spinnewijn [2021]), following up on the seminal work by Gruber [1997], but here applied to public pensions.

Our work also contributes to a small but recently expanding literature on the trade-off between incentives and insurance in pension design specifically. The theoretical foundations of this approach were laid by Diamond and Mirrlees [1978, 1982, 1986]. Some recent papers have re-examined this basic trade-off, using both theory and empirical analysis. O'Dea [2018] takes a structural approach to this trade-off. He contrasts the value of lifetime-earnings-based pensions with policies like minimum pensions that provide an income floor; his results also suggest that current policy under-values the insurance benefits of pension provision. In contrast, we use a sufficient statistics framework to characterize welfare effects and consider local reforms to the pension benefits schedule, such as a change in the return to additional years of work. Ndiaye [2020]'s approach to this trade-off is in the spirit of the macro public finance literature, characterizing the optimal retirement wedge and how this wedge changes with the age of retirement. In his model it is the fixed cost of work and how it correlates with productivity that determines whether inducing later retirement generates positive redistributive value. While our paper does not attempt to provide a full characterization of the optimal policy, we show how the welfare impact of pension reforms can be connected for a large class of models to moments that are directly estimable in the data. Complementary to our work is Haller [2022], who takes a similar sufficient-statistics approach to optimal pension design as in our paper, but focuses on the fiscal externality component of the trade-off. His work relates to a large empirical literature studying incentives and retirement behavior (e.g., Staubli and Zweimüler [2013]; Manoli and Weber [2016, 2018]; Gruber et al. [2022]; Seibold [2021]; Lalive et al. [2022]) and exploits Austrian pension reforms in the benefit generosity and early entitlement age to compare the corresponding average fiscal externalities. In contrast, our main empirical contribution is to estimate the consumption smoothing effects of pension reforms.[4]

The rest of the paper proceeds as follows. Section 2 develops the conceptual framework that guides our empirical analysis. Section 3 describes the Swedish institutional setting and our data. Sections 4 and 5 consider the consumption smoothing effects of incentivizing later retirement, using consumption levels and other moments of consumption. Section 6 considers consumption along other dimensions that are relevant for pension policy: early-career labor supply, income, and wealth. Section 7 describes how our consumption smoothing estimates enter into an overall welfare analysis of pension reforms in order to draw out the policy implications of our results. The final section concludes.

2 Conceptual Framework

In this section we briefly present a conceptual framework to evaluate pension design. This framework guides our empirical analysis and motivates our focus on specific consumption

[4] Additionally, our evaluation of the slope of the benefit profile also requires us to unpack retirement dynamics beyond looking at the average fiscal impact of a reform.

moments in the data. Pension benefits are an often complex function of individuals' employment history, including their retirement age and past contributions. The value of pension benefits conditional on specific features depends on the social marginal utility of these benefits to its beneficiaries. To evaluate a pension reform, we show that it suffices to compare the relevant social marginal utilities to the fiscal externality due to the behavioral responses triggered by the reform. In Appendix G we discuss further details regarding the setup and provide the full derivation of all equations and approximations.

Setup Our model can encompass the rich heterogeneity and non-separabilities in standard retirement models (e.g., French [2005], French and Jones [2011]). At any point in time t, the state variable $\pi_{i,t} \in \Pi_t$ captures all aspects of individual i's history and characteristics relevant for determining her utility and choices at that time. This can include an individual's past earnings and savings, shocks to her health, human capital or financial capital, etc. We assume that an individual chooses $c\left(\pi_{i,t}\right)$ and $\zeta\left(\pi_{i,t}\right)$ determining her flow utility $u\left(c\left(\pi_{i,t}\right), \zeta\left(\pi_{i,t}\right)\right)$ at time t given history $\pi_{i,t}$. The key innovation here is to capture all individual features - both exogenous and endogenous - that affect utility, other than consumption c, by the reduced-form variable ζ. This can include labor supply, home production, bequests and other choices, but also health status, preferences, and other characteristics. What matters for the value of (public) pensions is how the factors embedded in ζ modify the marginal utility of consumption, regardless of whether these factors are exogenous or endogenous.

Individual expected utility is the present discounted value of expected flow utility integrating over possible future states:

$$\mathcal{U}_i\left(c, \zeta, \pi\right) = \sum_{t=0}^{T} \beta^t \int u\left(c\left(\pi_{i,t}\right), \zeta\left(\pi_{i,t}\right)\right) dF\left(\pi_{i,t}\right). \tag{1}$$

We zoom in on the key decision to stay in the labor force or to retire, denoted by $s\left(\pi_{i,t}\right) \in \{1, 0\}$ and included in $\zeta\left(\pi_{i,t}\right)$. Obviously the marginal utility of consumption may be different under employment ($s = 1$) versus retirement ($s = 0$), in accordance with a large literature on non-separabilities in consumption-leisure (see Jappelli and Pistaferri [2017]). If $s\left(\pi_{i,t}\right) = 0$ (retirement), the individual receives pension benefits $b\left(\pi_{i,t}\right)$, which can depend on the individual's employment history in a general way. If $s\left(\pi_{i,t}\right) = 1$ (employment), the individual earns wages $w\left(\pi_{i,t}\right)$ and pays taxes $\tau\left(\pi_{i,t}\right)$. In either case ($s \in 0, 1$) after-tax income is denoted by $y\left(\pi_{i,t}\right)$. 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 $R\left(\pi_{i,t}\right)$. The individual's optimization problem is therefore to maximize $\mathcal{U}_i$ subject to the following constraints for each history $\pi_{i,t}$:

$$a_{i,t+1}(\pi_{i,t}) = R\left(\pi_{i,t}\right)\left[a_{i,t}(\pi_{i,t-1}) + y\left(\pi_{i,t}\right) - c\left(\pi_{i,t}\right)\right], \tag{2}$$

$$y\left(\pi_{i,t}\right) = \begin{cases} w\left(\pi_{i,t}\right) - \tau(\pi_{i,t}) \text{ if } s\left(\pi_{i,t}\right) = 1 \\ b\left(\pi_{i,t}\right) \text{ if } s\left(\pi_{i,t}\right) = 0. \end{cases} \tag{3}$$

We denote the resulting indirect utility by $U_i\left(b, \tau\right)$.

The government's problem is to maximize a generalized utilitarian social welfare function with welfare weights ω_i, subject to a government budget constraint,

$$\max \ \mathcal{W}(b,\tau) = \int_i \omega_i U_i(b,\tau)\,di + \lambda GBC(b,\tau). \tag{4}$$

The government budget constraint requires that the net present value of taxes collected while working equals the net present value of pensions paid out while retired.

Pension Policy Pension benefits $b(\pi_{i,t})$ can depend in a flexible way on a worker's employment history, including the retirement age, the number of years worked and the corresponding earnings. Recent pension reforms have changed how these features map into pension benefits. To evaluate the welfare effect of these reforms, we can group retired individuals by the features x determining the pension benefits (e.g., their retirement age) and consider the welfare effect of a change in pension benefits $b_{x,t}$ received at age t by individuals who retire with features x (e.g., a retirement age above 65):

$$\frac{\partial \mathcal{W}(b,\tau)}{\partial b_{x,t}} = \frac{\partial \int_i \omega_i U_i(b,\tau)\,di}{\partial b_{x,t}} + \lambda \frac{\partial GBC(b,\tau)}{\partial b_{x,t}}. \tag{5}$$

The first term of the welfare impact - capturing the marginal value of the pension benefit - can be re-written as

$$\frac{\partial \int_i \omega_i U_i(b,\tau)\,di}{\partial b_{x,t}} = G(x,t) \times \underbrace{E\left(\omega_i \frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial c} \middle| x_{i,t} = x\right)}_{\equiv SMU_{x,t}}, \tag{6}$$

denoting the share of individuals with features x and retired at age t by $G(x,t)$ and assuming $\beta = R = 1$. This term equals the average social marginal utility of transferring a dollar to individuals at age t, having retired with features x, which we denote going forward by $SMU_{x,t}$. Importantly, the value thus only depends on the social marginal utility of consumption for the beneficiaries of the increased pension benefits. Behavioral responses, including changes in labor supply, retirement, and/or savings, only have a second-order effect on agents' welfare, due to the envelope theorem. The same changes in behavior, however, imply that the fiscal cost of increasing expected pension expenditures by one dollar may differ from one dollar. We re-write the second term in equation (5) as

$$\frac{\partial GBC(b,\tau)}{\partial b_{x,t}} = G(x,t) \times [1 + FE_{x,t}], \tag{7}$$

making explicit that the fiscal cost includes the fiscal externality due to the behavioral responses, $FE_{x,t}$. For example, when increasing benefits for individuals retiring at later age, the later retirement age increases the taxes received and reduces the pension benefits paid.

The envelope condition relies on agents optimizing their behavior. The presence of behavioral biases would require including a third term in equation (5), consisting of marginal internalities and the corresponding behavioral responses to the reform [Mullainathan et al., 2012; Spinnewijn, 2015]. We note, however, that the first two terms would still be present if we incorporated biases, so it remains valuable to characterize the welfare effect occurring through the SMU, which is our main focus. Moreover, under-saving due to behavioral biases would mainly act to increase individuals' marginal utility in retirement. In other words, an individual's marginal utility might be higher in retirement because they saved too little and thus are forced to consume less; this is implicitly already captured by the SMU term in equation (6).[5]

Pension Reform We can now compare the effect a marginal change in benefit level $b_{x,t}$ for individuals who retired with features x, relative to a marginal change in the benefit level $b_{x',t}$ for individuals who retired with features x'. For example, we can think of a pension reform that incentivizes later retirement as one that increases benefits for those retiring after some age r and decreases them for those retiring before that age, as illustrated in Panel A of Figure 1. Optimality of the relative benefits of early and late retirees, or more generally any two groups with different characteristics x and x', requires, based on equation (6) that

$$\frac{SMU_{x,t}}{SMU_{x',t}} = \frac{1+FE_{x,t}}{1+FE_{x',t}}. \tag{8}$$

Otherwise, we can find a *budget neutral* reform of the profile that increases social welfare.[6]

Equation (8) resembles the classic insurance-incentives trade-off often studied for other social insurance policies (Baily [1978], Chetty [2006]). The left-hand side reflects the consumption-smoothing value of re-allocating transfers across groups, accounting for potential differences in welfare weights and the marginal utility of consumption. Importantly, this does not require a comparison of individuals who are working vs. retired, but only of retired individuals who are or could be treated differently by the pension system. The right-hand side reflects the relative fiscal externality caused by the changing incentives when reforming pension benefits. We note that a number of concerns affecting the optimal level of pension benefits, such as fiscal sustainability or inter-generational redistribution, are immaterial for the evaluation of a budget neutral within-cohort reform such as this. Formally, this is captured by the fact that we can characterize the welfare effect of such a reform without reference to the marginal cost of public funds, λ. Thus, if we consider the feature x to be a retirement age group, equation (8) can be used to evaluate reforms to pension benefits that incentivize later retirement, as Sweden and many other countries have recently done.

[5]While equation (6) does not capture any welfare effects due to internalities and behavioral responses to reforms, we note that in the empirically dominant case of so-called 'passive savers' documented in [Chetty et al., 2014], individuals would not change their savings behavior in response to the reforms. The absence of behavioral responses would imply that the additional welfare impact due to the potential bias correction is still only of second-order importance. Of course, behavioral frictions can play at other margins too. One example is the large impact that statutory retirement ages have relative to financial incentives on individuals' retirement behavior [Seibold, 2021]. Our focus here is on the welfare effect through the SMU channel and, briefly, the fiscal effect, and we defer consideration of internalities to other work [see e.g. Reck and Seibold, 2021].

[6]This relates 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',

$$\frac{SMU_{x,t}}{1+FE_{x,t}} > \frac{SMU_{x',t}}{1+FE_{x',t}} \tag{9}$$

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

Differences in Consumption The focus in our empirical analysis is on the consumption smoothing aspect of pension reforms, i.e. the left-hand side of equation (8). We defer further analysis of the fiscal externality due to later retirement to the welfare illustration in Section 7. How can we shed empirical light on the difference in social marginal utilities between groups of pension beneficiaries, like for example individuals retiring at different ages? A standard approach in the social insurance literature (see Chetty and Finkelstein [2013]) is to study differences in consumption and to use a Taylor series approximation to map these into differences in marginal utilities *:*

$$\frac{\partial u\left(c_{i,t},\zeta_{it}\right)}{\partial c} \cong \frac{\partial u\left(c_0,\zeta_{i,t}\right)}{\partial c}\left[1 - \frac{-\frac{\partial^2 u(c_0,\zeta_{i,t})}{\partial c^2}c_0}{\frac{\partial u(c_0,\zeta_{i,t})}{\partial c}}\frac{c_{i,t}-c_0}{c_0}.\right] \tag{10}$$

We can now use this approximation to compare the SMU for retirees with features x relative to the SMU of those retiring with features x'. Setting $c_0 = c_{x',t}$ in equation (10), we find:

Consumption-Level Implementation. *Assuming that for any i, t, x, $c\left(\pi_{i,t}\right) = c_{x,t}$, $\zeta\left(\pi_{i,t}\right) = \zeta_{x,t}$ and $\omega_i = \omega_x$ for $x\left(\pi_{i,t}\right) = x$, we can approximate*

$$\frac{SMU_{x,t}}{SMU_{x',t}} \cong \frac{\omega_x \frac{\partial u\left(c_{x',t},\zeta_{x,t}\right)}{\partial c}}{\omega_{x'}\frac{\partial u\left(c_{x',t},\zeta_{x',t}\right)}{\partial c}} \times \left[1 + \gamma_{x,t}\frac{c_{x',t}-c_{x,t}}{c_{x',t}}\right], \tag{11}$$

where γ equals the relative risk aversion in consumption preferences.

The implementation highlights how the difference in SMU's depends crucially on the difference in consumption across retirement groups receiving different pension benefits. For example, the lower is consumption by early retirees relative to late retirees, the higher is the cost of incentivizing later retirement by adjusting the pension benefit schedule.

Limitations of the Consumption-Level Implementation. A first practical implementation assumption is that preference heterogeneity occurs across individuals retiring with different features x rather than across individuals retiring with the same features. Otherwise, the aggregation would need to account for the within-group covariance between preferences and consumption.[7] A second practical implementation assumption relates to how the social marginal utility of consumption for individuals retiring with different features x compare, *conditional on consumption*. Concerns about this assumption motivate several additional empirical exercises. The main concern is that the marginal utility of consumption itself, conditional on observed consumption, could be different for different groups. More specifically, observed consumption expenditures could translate differently into real consumption across groups (e.g. due to differential reliance on home production), or actual preferences over consumption could differ across groups (e.g., due to differential complementarities with leisure). It is important to note here that we compare the consumption of individuals *when retired,* so the common concerns raised in relation to the so-called retirement consumption puzzle – drawing inference from

[7] Andrews and Miller [2013] generalize the characterization of the SMU's for when within-group heterogeneity is present. Since we are interested in the relative SMU's across groups with different features, the within-group heterogeneity is only relevant to the extent that it differs across these groups.

differences in consumption when employed vs. retired – do not apply here. An additional concern is that one may wish to use different welfare weights (ω_x) across different groups, for example due to differences in health or life expectancy.

We leverage the richness of the Swedish administrative data to address these potential limitations of the consumption-level implementation empirically. We follow recent work in the social insurance literature by using other consumption moments to learn about differences in social marginal utilities (e.g., Hendren [2017], Landais and Spinnewijn [2021]): we complement the consumption-level implementation with two other implementations that focus respectively on consumption dynamics and on marginal propensity to consume, each of which relies on different assumptions regarding the mapping of consumption patterns into SMUs. We also examine the composition of consumption and a rich set of additional observable characteristics to gain insight into the drivers of differences in consumption, allowing us to critically assess our assumptions about the map from consumption to marginal utility.

Dimensions of Pension Policy As we turn toward connecting this conceptual approach to consumption data, we must specify which dimensions of pension benefits to consider, i.e. which feature we specify as the variable x. Because the most important recent reforms to retirement pension designs across developed nations have focused on strengthening incentives to supply labor late in life (Gruber and Wise [1999], OECD [2019], Barr and Diamond [2009]), and because changes in late-career labor supply incentives are essentially equivalent to changing the benefits one receives as a function of their retirement age, the primary characteristic of interest x that we focus on in our empirical implementation is age at retirement r. That is, we compare in sections 4 and 5 the consumption patterns of individuals who retire at different retirement ages r in order to measure the social marginal cost of incentivizing later retirement, as captured by the ratio $\frac{SMU_{r,t}}{SMU_{r',t}}$. In section 6, we further investigate consumption patterns across other important characteristics for pension benefits, namely early-career labor supply, income, and wealth, to provide further insights on the welfare effects of alternative reforms to the design of retirement pensions. Public pensions are complex and vary across countries, but together, income, early- and late-career labor supply capture the important sources of variation in pension benefits in most countries' pension benefits. Wealth is not typically a direct determinant of pension benefits, but introducing an asset test into pension benefits has been debated in the US and elsewhere. We discuss how Swedish pension benefits map onto these conceptual dimensions in the next section.

3 Institutional Background & Data

This section provides an overview of Swedish pension institutions and the administrative data from Sweden that we use. We describe key features of the Swedish pensions system and explain how the major pension reform in Sweden maps into the dimensions of pension benefits that are the focus of our empirical analysis. A more comprehensive review of Swedish pensions is in Appendix A.

3.1 Institutional features of the Swedish pension system

The Swedish Pension system consists of three primary components: public pensions, occupational pensions, and private pensions. We focus on the first of the three, public pensions, but account for the presence of the other two components. Sweden has undertaken a large reform to its public pension system and is currently transitioning from a defined benefit system, called the "ATP" scheme, to a new, "Notional Defined Contribution" (NDC) scheme.[8]

Public Pension Reform In the pre-reform ATP system, pension benefits are determined by earnings averaged over the 15 highest-earning years in an individual's career and the total number of years in which an individual earns pension rights, up to a maximum of 30 years.[9] 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. Annual ATP points are capped at 6.5 BAs, which corresponds empirically to the median of the earnings distribution for 55 year olds in 2000; earnings in a given year beyond this level do not increase pension rights. For individuals with short careers and 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. Our data shows that a quarter of all 66 year olds received the basic pension in 2007.

The post-reform NDC system resembles a DC system from a worker's perspective. A given worker's benefits are an annuity closely linked to that worker's lifetime contributions through payroll taxes. Unlike a typical private DC scheme, however, the system retains its Pay-As-You-Go structure, as pension points are only notional. Pension benefits in the NDC system are calculated from 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, pensionable earnings grow.[10] Pensionable earnings are capped at a higher level (at 7.5 income base amounts) than the ATP system. 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*. The new minimum benefit is about 40% higher than the benefit under the ATP system. About 30% of all individuals receiving pension benefits are expected to receive basic pensions in 2040 when the NDC system is fully phased in.

The reform was passed in the Parliament in 1994 and phased in gradually across cohorts starting in 1998. Cohorts born before 1938 receive their pension benefits from the ATP system.

[8]ATP stands for *Allmän tilläggspension* in Swedish, which means "General supplementary pension." The word "supplementary" in the title refers to the fact that there is also a basic old-age pension benefit whose amount does not depend on a person's earnings history. We refer to combined public old-age pension benefits system prior to the reform as the ATP system, which is common terminology.

[9]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, like social aid or housing aid.

[10]The BA has been replaced by an "income base amount" which is indexed to average wage growth instead of prices. However, just like with the BA, the government keeps some discretion over how to set the income base amount.

Those born between 1938 and 1953 receive a weighted mixture of ATP and NDC benefits, with increasing weight on the NDC benefits over time. 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. Pension benefits in both the ATP system and the NDC system are financed by payroll taxes.

Conceptualizing the Reform The goal of this paper is not to provide a full evaluation of the Swedish reform, but rather to evaluate conceptual reforms that inform pension design in general. Nevertheless, it is instructive to map the Swedish NDC reform to the types of conceptual reforms we evaluate below, as it is a great example of how such reforms can affect the design of pension systems across our three key policy dimensions of interest: the provision of late-career incentives, the provision of early career incentives, and redistribution by permanent income level.

As mentioned, our primary focus is on reforms to pension benefit schedules that incentivize late-career labor supply, i.e. incentives to retire later in life. Incentivizing later retirement was an explicit goal of the Swedish reform; indeed this appears to be the most consistent feature of recent public pension reforms around the world. Two main features of the Swedish reform explicitly contributed to this aim: allowing workers to earn pension rights beyond age 64, and removing the 30-year contribution cap.[11]

We illustrate how the Swedish reform increased late-career incentives in Panel B of Figure 1, which quantifies how pensions benefits change as a function of workers' retirement age, mirroring the stylized reform shown in Panel A.[12] We construct the Figure using simulations of lifetime income and pension benefits for a representative set of workers born in 1941. Holding earnings history and income fixed, we calculate the net present value at age 55 of workers' pension benefits at different retirement ages (see Appendix A).[13] The removal of the cap on pension rights after age 65 has the most salient effect on these pension benefits schedules: the ATP schedule mechanically flattens out after age 65 and the NDC schedule does not. More subtly, the removal of the 30-year contribution cap increases the slope of the pension benefits profile over all retirement ages.[14]

While the late-career labor supply dimension was a key focus of the reform in public discussions, some provisions also affected two other key dimensions of pension benefits: early-career

[11]For a small subset of workers, a third feature, the increase in the cap on pensionable earnings, increased the return to work at later ages as well.

[12]Figure 1B abstracts away from the overall level effect of the reform on pension benefits. However, to promote fiscal sustainability, the NDC reform enacted a reduction in pension benefits for most workers. We illustrate the level effects in Figure A-12 and further discuss them in Appendix A.

[13]To account for how the reform affected workers differently depending on their lifetime earnings, we consider 20 hypothetical earnings histories, calibrated based on vigintiles of the distribution of accrued ATP points at age 55, median earnings and years worked at 55 for workers in each vigintile, and historical earnings growth. Further analysis in Appendix A suggests that this approach provides a reasonably accurate account of how pension benefits change as a function of the retirement age through the distribution of earnings histories. Averaging the NPV of pension benefits in the NDC and ATP system across the 20 hypothetical workers, each of whom roughly represents 5% of the lifetime earnings distribution, we arrive at Figure 1 Panel B.

[14]Beyond 30 contribution years, earning income and contributing payroll taxes in a given year in the ATP system could still increase the value of one's pension benefits by increasing the average income on which pension benefits were based – average income in the top 15 contribution years. Workers earning above this average late in life could increase the average marginally for each additional year they worked.

labor supply and the link between pension and permanent income. Even holding late-career labor supply fixed, removing the 30-year career length cap implicitly redistributes from workers with short early careers to workers with long early careers. Separately, the reform had a complex effect along the income dimension. The reform increased the cap on pensionable income by roughly 25%, but it also increased the minimum pension benefit amount by about 40%. The first of these provisions redistributes toward workers near the top of the *annual* income distribution, while the second provision redistributed toward workers near the bottom of the *lifetime* income distribution, which comprises workers with low annual income and/or short careers.

Retirement vs. Claiming A dimension of pension design that we ignore in this paper, due to the specifics of our context, are claiming incentives. Pensions can be claimed from age 61, which we refer to as the early retirement age. Unlike many other countries, Sweden has no earnings test whereby pension benefits are reduced for those continuing to work after claiming the pension benefits. In the ATP system, claiming before age 65 resulted in a nearly actuarially fair reduction in benefits, while benefits are adjusted slightly more for those claiming after age 65. In the NDC system, the adjustments are on average actuarially fair by design: claiming pensions earlier means that the sum of pensionable earnings is divided by the longer life expectancy. Consistent with the idea that retirement and claiming are de-coupled in Sweden, we observe much more variation in retirement ages compared to claiming ages, as illustrated in Appendix Figure A-3.[15] In quantifying the effects of the reform on incentives above we focused for simplicity on the case where individuals claim at 65, which empirically is the most relevant case; we discuss and conduct further analysis on this point in Appendix A.

Other Social Insurance Motivated by our conceptual framework, we focus on retirement defined as the moment individuals stop working permanently. On top of the three pillars of the pension system, other components of the social insurance system in Sweden, such as disability insurance or unemployment insurance, provide transfers to cushion the shock of losing employment for the elderly. Although these are not explicitly called "pension" benefits, such benefits received by workers at the moment they stop working do affect the profile of their labor supply incentives in old age. Non-pension social insurance benefits also contribute to "pensionable earnings" in determining a workers pension benefits, in both the ATP and NDC system. Because of all this, we integrate these benefits as part of the overall pension system when computing the NPVs of benefits related to stopping to work at different ages. We provide details on these computations in Appendix A, and explore the robustness of the pension benefits schedules from Figure 1B to alternative treatments of UI and DI benefits.[16]

We finally note that with two exceptions, the pension system, like most of the Swedish tax and

[15]In the cohorts we study, 69% of workers claim their pension at age 65, but only about 22% retire at age 65, with far more workers retiring before 65 than claiming before 65. 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.

[16]In Appendix A, we show that although it is especially common for low-income workers retiring early to claim social insurance after exiting the labor force, accounting for insurance benefits and their induced additional pension rights has a very small effect on the average NPV of benefits from pensions and other social insurance combined (see Figure A-7). As a result, this has a small effect on the fiscal externality from incentivizing individuals to work longer (see Figure A-9).

transfer system, is entirely individualized. The first exception is that the minimum pension benefit in both systems is about 10% lower for married individuals than for singles. The second is that there is a survivor's benefit that is paid out for a year after one's spouse has passed, see Appendix A for details.

3.2 Data

We rely on uniquely rich data on consumption, employment, pensions, and health. The data comes from several Swedish registries covering the universe of the population, as well as additional surveys, which can all be linked using a unique personal identifier (*personnummer*).

Labor Market History and Pensions Our first source of information on labor supply history in old age is LISA, a panel containing the universe of individuals residing in Sweden aged 16 years or above, between 1990 and 2017. LISA includes socio-demographic variables such as age, education, marital status, household composition and place of residence. It also contains information on labor market status, labor earnings, various types of transfers such as sickness benefits, disability benefits and unemployment benefits.

From LISA, we construct a registry-data measure of retirement, defined as the moment individuals stop working permanently. To do this we follow Karlstrom et al. [2004] and categorize an individual as retired when her labor earnings permanently fall below one Base Amount − about 18% of median labor earnings.[17]

Our second main registry data source is data on pension contributions in both the old ATP system and the new NDC system. Data from the ATP system contains contributions from 1960 onwards for all individuals born 1938 and later. The NDC contributions are available from the late 1990s when the NDC system was initiated. In addition, the data also contains information on all pension benefits that individuals accrue and receive: old age state pension benefits, occupational pension and private pension savings.

Consumption Measures To measure consumption, we use the registry-based measure of annual household consumption expenditures for the universe of Swedish households created for all years 2000 to 2007 by Kolsrud et al. [2020]. The construction of this measure relies on the identity coming from the household's budget constraint between consumption expenditures and income net of changes in assets. We aggregate consumption at the household level using administrative identifiers of household structure created by Statistics Sweden. The quality of our consumption expenditure measure owes to the comprehensiveness of income and asset data in Sweden. First, LISA contains exhaustive disaggregated information on all earnings, all taxes and transfers and capital income on an annual basis. Second, we have precise data on wealth coming from the wealth tax register (*Förmögenhetsregistret*), which covers the asset portfolios for the universe of Swedish individuals with detailed information on the stock of all financial assets (including pension wealth and different types of debt) and real assets as of December of each year. We complement the information from the wealth tax register with data

[17] The one base amount (BA) threshold is widely used to define labor force participation in the administrative data. One BA also corresponds to the minimum earnings threshold allowing individuals to earn pension rights in the ATP system. See further details in Appendix A. Note that we define the year in which the event of retirement takes place as the last year in which the individual earns more than one BA.

on financial asset transactions (KURU), and data on real estate transactions from the housing registries (*Fastighetsprisregistret*), which enable us to disentangle the contribution of savings from that of price changes in the evolution of asset balances. The KURU register also allows us to construct measures of wealth shocks using random variation in asset prices that we exploit in section 5.2.

We complement our registry-based measure of consumption with survey-based measures from the Swedish consumption expenditure survey (HUT) that we match to our administrative data. This allows us in particular to investigate the structure of consumption expenditures across categories. We refer the reader to Appendix B and to Kolsrud et al. [2020] for further details on the construction of our consumption measures, and for a thorough assessment of the robustness and consistency of registry-based vs survey measures of expenditures. We note that our imputed measure of consumption is capturing, like most survey measures, expenditures rather than consumption. We discuss at length how this affects the mapping between consumption and welfare when we present our empirical results below.

Health and Mortality We also complement our data with the death register, as well as with two large surveys, the living condition survey (ULF) and the household finance survey (HEK), containing detailed information on health and health expenditures. We provide all details on data construction and on the computation of our composite health indices in Appendix E.

Sample and Descriptive Statistics Our main sample focuses on all individuals from cohorts 1938 to 1943. Figure 2 displays the distribution of age at retirement among individuals belonging to these cohorts. It shows that the vast majority of individuals retire between 55 and 70, with a peak at 65. For our analysis of late-career incentives, we define four retirement age groups based on this empirical distribution. 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. We drop from our sample the small group of individuals whom we observe retiring before 55, or after 70.

We chose these cohorts and retirement age groups for two practical reasons. First, we only observe the full ATP contribution history for cohorts born from 1938 onwards. Second, given our consumption data spans years 2000 to 2007, this sample selection allows us to observe, for each cohort, consumption during retirement, as well as before retirement, for all retirement age groups. This specification of cohorts therefore allows us to control for both age and cohort effects in consumption.

Table 1 provides summary statistics for this baseline sample, with information on retirement patterns, demographics, income, wealth and pensions. The sample comprises 418,033 unique individuals, with an average age at retirement of 62.9.

4 Retirement Consumption & the Welfare Cost of Late-Career Incentives

This section studies consumption differences across individuals retiring at different ages, and uses these estimated differences to evaluate the costs of reforms that strengthen late-career labor supply incentives, i.e. the steepening of the pension benefit profile along the retirement-age dimension.

4.1 Consumption Levels By Retirement Age

We start by documenting how consumption differs across individuals who retire at different ages. We measure these differences at the same age, and in the same state, i.e. when individuals are retired, in order to be consistent with the welfare implementation of equation (11).

Empirically, we simply regress household consumption C_{it} of individual i at age t in year y, on a series of dummies that capture an individual's retirement age r:

$$C_{it} = \sum_{j} \alpha_j \cdot \mathbb{1}[\mathrm{r} = \mathrm{j}] + \gamma_y + \gamma_t + \mathbf{X}'\beta + \varepsilon_{it}. \tag{12}$$

We estimate model (12) including consumption at all ages $t > r$, that is we restrict the sample to individual X year observations for which individuals are observed as being retired. To control for business cycle fluctuations and for the life-cycle profile of consumption, we include both year fixed effects γ_y and age fixed effects γ_t. In effect, we compare consumption of individuals *from the same cohort, at the same age*, who are currently retired, but who have retired at different ages. In practice, we group retirement ages into four groups, as explained above: 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 systematically use normal retirees as the reference category.

The vector of controls $\mathbf{X}$ comprises two sets of variables. First, we include a series of dummies capturing household composition because we measure total consumption at the household level. By including them, we control for any mechanical relationship between consumption and retirement age, in case the latter correlates with family composition. Second, we include dummies corresponding to the main determinants of pension benefits besides the age of retirement: income and career length early in life. Specifically, we control for deciles of individuals' average income between 52 and 55 and for the number of years individuals' have been employed before the age of 55. This allows us to compare the consumption levels of individuals who retire at different ages, but would have otherwise received the same pension benefits. Adding this second set of controls is not necessary to inform the consumption smoothing effect of a reform according to equation (11), but it reveals the extent to which differences in the value of pension benefits across retirement age groups are attributable to differences in other determinants of public pension benefits.

Figure 3 reports the estimated coefficients α_j from specification (12) for all retirement age groups. We estimate the regression using consumption levels (rather than logs) but to facilitate interpretation, we scale the estimates α_j for all retirement age groups by $\mathbf{E}_j[\tilde{C}_{it}]$, the average predicted consumption level in retirement age group j from specification (12) when omitting the contribution of the retirement age group dummies.[18] We start, on the left hand side of

[18] $\mathbf{E}_j[\tilde{C}_{it}]$ therefore corresponds to the average level of consumption of individuals who retire between 64 and 65 from the same cohort, age, family composition and ATP quartile at age 55 as the average individual retiring in age group j.

the graph, with results from model (12) 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 and the other determinants of pension benefits.

Two important insights emerge. First, the estimates reveal the presence of a very strong positive gradient of consumption with retirement age. When retired, the level of consumption of premature retirees is 5% lower than consumption of normal retirees from the same cohort, at the same age. Late retirees, to the contrary, enjoy a level of consumption that is between 10 and 20% larger than normal retirees at the same age. Importantly, the magnitude of the overall gradient remains large when controlling for family structure and other pension determinants: this suggests that the large differences in consumption between individuals who retire very early and those who retire very late is not primarily driven by differences in household composition or labor market history. The second insight is that, while the overall gradient is positive, the relationship between consumption and retirement age also exhibits a clear non-monotonicity. Indeed, consumption is actually larger for early retirees compared to normal retirees, with a significant difference of about 3%. The non-monotonicity is dampened when controlling for household composition and other pension determinants, but there remains a clear flattening of the consumption gradient around this retirement age range. [19]

4.2 Robustness and External Validity

The two main consumption patterns are robust across different specifications in the Swedish context and also appear when using survey data and considering different countries.

We first provide additional evidence of robustness in the Swedish context. In Appendix Figure C-2, we show that the consumption patterns hold irrespective of the age at which consumption is observed during retirement. We run regressions similar to specification (12), but separately for each age t.[20] We document a very strong positive gradient of consumption with retirement age at all ages at which consumption is observed. The consumption of late retirees is systematically 15 to 20% larger than that of premature retirees. The non-monotonicity also obtains for any age at which consumption is observed. In Appendix Figure C-3, we further show that the the consumption patterns are similar across household structures. We replicate specification (12), splitting the sample between single vs couples, where family structure is defined as of the year of retirement. We observe a large negative gradient between early and late retirees for couples and singles; the non-mononotonicity between early and normal retirees appears to be driven virtually entirely by couples.

We briefly turn to the question how robust these patterns of consumption by retirement age are across contexts and data sources. The consumption gradient across retirement age groups likely depends 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 differs

[19] In Appendix Figure C-1, we also report estimates of a fully non-parametric version of specification (12) where we compare consumption levels across all retirement ages (rather than aggregating retirement ages into four groups). One additional insight that emerges is the sharp difference in consumption levels between individuals who retire before age 65 and individuals who retire just after 65.

[20] 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*.

across countries and over time. But many countries share similar institutions to those described in Section 3.1, with pension profiles that penalize early retirement and it is therefore important to investigate whether the broad patterns of consumption documented in Sweden hold across institutional contexts as well. Another challenge when exploring robustness across contexts is the limited availability of data with detailed information on both consumption and retirement behavior.

We examine this using survey data from SHARE and HRS which contain, for similar cohorts, information on retirement and survey measures of consumption expenditures for 11 European countries and the US. We report the results in Appendix D, which overall confirm that the large gradient in consumption levels between individuals who retire very late versus very prematurely is a robust finding across all contexts and data sources. Our non-monotonicity is also strikingly robust across contexts and data: for most people retiring between 61 and 65, there is no gradient, or if anything a negative gradient between consumption level and retirement age. Interestingly, the overall gradient found in the HRS data for the US is larger than the one we find in Sweden. There is 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 an even steeper pension profile in the US, and the fact that insurance against shocks in late career (such as UI and DI) is 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.

4.3 Sources of Heterogeneity

Our results so far show large differences across retirement age groups in consumption levels measured at the same age, when retired. To shed further light on the consumption heterogeneity we leverage the rich Swedish administrative and survey data to explore differences in both resources used to fund the consumption expenditures and categories on which the resources are spent. Doing so helps us not only to understand the potential mechanisms driving such consumption differences, but also to map these consumption differences into welfare.

We first decompose our measure of household expenditures into a rich set of components that shed light on the consumption means available to individuals. These components include own income (which we break down into own labor earnings, public and occupational pensions, and other government 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 (12) separately for each component evaluated at age 68 on the sample of all individuals from cohorts 1938 to 1943 who are retired by age 68. Figure 4 reports the estimates $\hat{\alpha}_j$ for each component, scaled by $\mathbf{E}_j[\tilde{C}_{it}]$, with one panel for each retirement age group.

Results reveal that the main reason why late retirees enjoy much larger consumption than other retirees is their significantly larger flow of consumption out of wealth, i.e., financial assets and real estate, including imputed rents. Together, these flows account for more than half of the difference in overall consumption between late and normal retirees. In addition, the late retirees enjoy higher pension benefits, both from the public system and from occupational pensions. The opposite is true for premature retirees. The figure shows that the lower levels of consumption of premature retirees are driven by a combination of lower flows across all available means of consumption. They have lower pension benefits, including occupational pensions. They also have significantly lower consumption out of wealth and lower consumption out of the income of other household members. Interestingly, in Appendix Figure C-7, we replicate the same exercise at age 60, which reveals that premature retirees have a much higher incidence of unemployment insurance and disability insurance receipt. This evidence suggests that individuals who retire prematurely not only have limited means to smooth consumption, but may also be more likely to have experienced negative earnings shocks due to unemployment or disability in their late career.[21] We provide further evidence on the dynamics of consumption across retirement age groups in Section 5.1. We note that Figure 4 also illuminates the drivers of the non-monotonicity highlighted above. Panel B shows that early retirees enjoy higher consumption despite having lower pensions, because they have both higher consumption flows from wealth, and also, significantly larger consumption flows out of the income of other household members.[22] This evidence suggests that many individuals in this group retire earlier in part because they have the means to do so.

We confirm these insights when studying the selection on observables into one of the four different retirement age groups. We estimate a multinomial logit prediction model including a large set of socio-economic characteristics as well as cohort fixed effects. In Figure 5, we report for each regressor the estimated average marginal effects on the relative probability to select into each of the groups, using normal retirees again as reference category. Late retirees appear relatively well-off across multiple observable dimensions. They are significantly more highly educated than all other retirees, healthier, earned much higher incomes and accumulated more assets. They do not have longer careers as of age of 55. At the opposite end of the spectrum, premature retirees exhibit the lowest educational attainment, the worst health and mortality, and the lowest levels of wealth. Like late retirees, premature retirees do not have significantly different career lengths at age 55 and are more likely to be male.

Patterns of selection related to early vs normal retirees point to the same mechanisms underlying the non-monotonicity in the consumption gradient. Early retirees have higher income than normal retirees and even as high levels of average household assets as the late retirees. They are also more likely to be cohabitating or married, and to be female, compared to normal retirees. A possible explanation for the specific household patterns lies in complementarities in labor supply decisions around retirement: early retirees, who are more often women and more often enjoy an above-average consumption, may time their retirement with that of their older partner.[23] The strong correlation between wealth and retirement age also hints at the presence

[21] As explained in section 3, we consider retirement as the age an individual stops working. And because UI and DI may provide financial support until pension claiming for premature and early retirees, we explicitly account for UI and DI when computing the incentives provided by the pension profile to stop working at different ages. In Appendix Figure C-5, we show that the consumption differences across retirement age groups are robust to using an alternative measure of retirement that accounts for the time spent in UI or DI after an individual stops working.

[22] Note that these estimates control for household structure. Differences in the contribution of income from other household members therefore does not reflect differences in household structure, but differences in the magnitude of income flows generated by household members for a given household structure.

[23] The average age difference between couples in our data is 3.8 years while the difference in age of retirement

of significant wealth effects on labor supply around retirement (Giupponi [2019], French et al. [2020]).

We finally consider differences in the expenditure shares for different consumption categories using the data from consumption surveys (HUT). Differences in consumption structure by retirement age would indicate the presence of significant preference heterogeneity. Figure 6 ranks the 11 consumption categories by their importance for the sub-sample of retired individuals surveyed in the HUT. Quite strikingly, the differences in expenditure shares across retirement age groups are small and insignificant. The one potential exception is the group of late retirees, who seem to spend for example less on food at home and more on restaurants and hotels, as well on recreation, but the differences remain small and mostly insignificant. This implies that preference heterogeneity across retirement-age groups, if any, can only exist to the extent that it does not translate into different consumption expenditure patterns.

Taken together, our results seem to suggest that the differences in consumption across retirement age groups we found initially in Figure 3 are driven more by differences in the means for consumption than by differences in preferences for consumption. This is in line with the necessary assumptions for mapping consumption differences to differences in the value of pension benefits via equation (11).

4.4 Quantifying Welfare Costs

We now map the estimated consumption patterns into estimates of the consumption smoothing costs from strengthening incentives for later retirement. In particular, we consider a simple 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 using this to increase pensions for individuals retiring after age $\tilde{r}$ by $db_{r>\tilde{r}} = -\frac{1-S(\tilde{r})}{S(\tilde{r})}db_{r\leq\tilde{r}}$, where $1-S(\tilde{r})$ is the share of individuals who retired before age $\tilde{r}$.[24] This type of reform is illustrated in Panel B of Figure 1 for $\tilde{r}=65$. Building on equation (11), we can approximate its consumption smoothing cost by

$$\frac{SMU_{r\leq\tilde{r}}-SMU_{r>\tilde{r}}}{SMU_{NRA}} \approx \gamma \times \left[\frac{E_{r>\tilde{r}}(c)}{E_{r\in NRA}(c)} - \frac{E_{r\leq\tilde{r}}(c)}{E_{r\in NRA}(c)}\right], \tag{13}$$

where the differences are expressed relative to the normal retirement age group as estimated in regression (12). This baseline implementation assumes no differences in welfare weights across retirement ages, nor in marginal utilities conditional on consumption, but we consider alternative assumptions below.

The blue bars in Figure 11 depict the resulting consumption smoothing costs of steepening the profile for each retirement age $\tilde{r} \in [56, 67]$, using a CRRA risk aversion parameter γ of 4 (see Landais and Spinnewijn [2021]). For this baseline implementation, the consumption smoothing costs range between .27 and .80. Hence, per dollar(/krona) transferred from individuals

is 3.3 years, suggestive of a joint retirement decision for couples. Gustman and Steinmeier [2000], using US data from the National Longitudinal Survey of Mature Women for the US, and Hospido and Zamarro [2014], using the European SHARE dataset, report similar findings on the average age differences and a joint retirement decision for couples.

[24]To be precise, we can implement this change in benefits for individuals at any given age t, but would need to scale by the share of individuals retiring before vs. after age $\tilde{r}$ among the individuals still alive at that age t. For brevity, we drop the age subindices.

retiring early to individuals retiring late, social welfare decreases by between 27 and 80 cents due to the loss of consumption smoothing. The figure also shows a clear non-monotonicity in the consumption smoothing costs, reflecting the non-monotonicity in the consumption levels. The consumption smoothing cost of inducing later retirement is lower at ages between the early and normal retirement age compared to the age before the early retirement age or after the normal retirement age. The stylized reforms we consider in Figure 11 redistribute resources from *everyone below* a certain retirement age to *everyone above*. If we consider instead more flexible reforms reducing pensions only for individuals retiring early and increasing pensions only for individuals at the normal retirement age, we would improve the consumption smoothing across retirees, suggesting that there is locally no trade-off between incentives and redistribution for those retiring between 61 and 65.[25]

5 Using Alternative Consumption Moments to Evaluate Late-Career Incentives

This section complements the comparison of consumption levels with the comparison of other consumption moments. We consider the change in consumption around retirement and the marginal propensity to consume when retired. The advantage of studying additional consumption moments is twofold. First, the mapping from these alternative consumption moments into *SMU*s relies on different assumptions on preferences, and thus allows to relax specific assumptions underlying the baseline implementation. Second, these alternative consumption moments also capture different aspects of consumption smoothing, which may be important to separately identify from a policy perspective.[26]

5.1 Consumption Dynamics Around Retirement

We first study the consumption dynamics around retirement and revisit the consumption drops at retirement which have drawn large attention in the literature. In contrast with prior work, the focus of our analysis is on how these consumption drops are different across individuals retiring at different ages.

5.1.1 Mapping between Consumption Changes around Retirement and SMUs

We first illustrate how the difference in consumption changes around retirement can be used as an alternative to capture differences in SMU's relative to the baseline implementation using the difference in consumption levels. Using a Taylor expansion of the ratios of social marginal utilities around pre-retirement consumption, we find:

[25] Reducing the curvature in consumption preferences γ of course reduces the consumption smoothing cost linearly. Column (2) in Appendix Table H-1 shows the consumption smoothing cost for $\gamma = 2$ instead of $\gamma = 4$ (repeated in column (1)). However, recent work in the context of unemployment (e.g. Hendren [2017], Landais and Spinnewijn [2021]) suggests that, if anything, the consumption-based approach we employ here requires more curvature than in our baseline implementation ($\gamma \geq 4$). Note also that as long as the risk aversion preferences are constant across retirement-age groups, the qualitative pattern in our estimates of the consumption smoothing cost remains the same.

[26] Our baseline and alternative implementations to measure the social marginal value of pension reforms have important similarities, but beyond differing in their underlying assumptions and in the interpretation of the social marginal value they capture, they also differ in the empirical inputs they require and thus in the challenges they entail. To facilitate comparison, Table G-1 in Appendix G summarizes the strengths and weaknesses of each approach.

Consumption-Drop 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, t, x, *we can approximate*

$$\frac{SMU_{x,t}}{SMU_{x',t}} \cong \frac{\omega_x \frac{\partial u\left(c_{x,pre},\zeta_{x,pre}\right)}{\partial c}}{\omega_{x'} \frac{\partial u\left(c_{x',pre},\zeta_{x',pre}\right)}{\partial c}} \times \frac{\theta_{x,pre,t}}{\theta_{x',pre,t}} \times \frac{1+\gamma_x \frac{c_{x,pre}-c_{r,t}}{c_{x,pre}}}{1+\gamma_{x'} \frac{c_{x',pre}-c_{x',t}}{c_{x',pre}}}, \tag{14}$$

where γ_x *denotes the relative risk aversion in consumption preferences,* t *refers to an age or time period after retirement, while "pre" refers to an age or time before retirement, and* $\theta_{x,pre,t} = \frac{\frac{\partial u(c_{x,pre},\zeta_{x,t})}{\partial c}}{\frac{\partial u(c_{x,pre},\zeta_{x,pre})}{\partial c}}$ *captures the impact of changes in* ζ_x *around retirement on the marginal utility of consumption.*

This implementation highlights that - everything else equal - the larger the drop in consumption around retirement for individuals retiring with features x rather than x', the larger the relative value of pension benefits given to these respective groups.[27] While it again assumes that the relevant heterogeneity occurs across the different groups, the value of this alternative implementation is twofold.

First, by relying on differences in *within-individual changes* in consumption across individuals (rather than differences in post-retirement *consumption levels*), we do not need to be concerned about *permanent* differences in non-consumption determinants of marginal utilities ζ_x – e.g., a smaller value of consumption relative to leisure – driving selection into retirement. Such differences would confound the translation from the difference in consumption levels to SMU in the consumption-level implementation, but they would not confound the within-individual differences in the consumption-drop implementation.

That being said, while permanent differences in ζ_x across groups are now inconsequential, changes in ζ_x around retirement, to the extent that they affect marginal utility of consumption, as captured by the term $\theta_{x,pre,t}$, could still affect the mapping between the consumption drop of each group and their SMU. This is a standard issue, abundantly highlighted in the retirement-consumption literature. Note however that if we are interested in pension design, and therefore in the ratios of the SMU across different groups, these effects are only relevant to the extent that they affect retirement groups differentially ($\theta_x/\theta_{x'} \neq 1$).

The second argument for focusing on this implementation is that planners might wish to evaluate the welfare of the pension system taking the differences that arise before retirement as given (i.e., $\omega_x \frac{\partial u\left(c_{x,pre},\zeta_{x,pre}\right)}{\partial c} = \omega_{x'} \frac{\partial u\left(c_{x',pre},\zeta_{x',pre}\right)}{\partial c}$). This can be motivated by the availability of other policy tools for redistribution and insurance of earnings differences during an individual's working life (progressive income taxes, unemployment and disability insurance, etc.) This welfare perspective relates to Diamond and Mirrlees [1978], where the social planner

[27] To compare this implementation to the first one based on differences in post-retirement consumption, 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

$$\frac{SMU_{x,t}}{SMU_{x',t}} \cong \frac{\omega_x}{\omega_{x'}} \frac{\frac{\partial u\left(c_{x',pre},\zeta_{x,pre}\right)}{\partial c}}{\frac{\partial u\left(c_{x',pre},\zeta_{x',pre}\right)}{\partial c}} \times \frac{\theta_{x,pre,t}}{\theta_{x',pre,t}} \times \left[1+\gamma_{x'} \frac{c_{x',pre}-c_{x,pre}}{c_{x',pre}}\right] \times \frac{1+\gamma_x \frac{c_{x,pre}-c_{x,t}}{c_{x,pre}}}{1+\gamma_{x'} \frac{c_{x',pre}-c_{x',t}}{c_{x',pre}}}.$$

uses public pensions to provide insurance against work longevity risk, coming from disability shocks later in life. Under this more narrow social insurance perspective, the differential consumption *drops* around retirement for retirees with different characteristics x are more informative than the differential consumption *levels* across these retirees. In particular, in the Diamond and Mirrlees [1978] framework, consumption drops are expected to be higher for people who retire earlier as this is induced by disability shocks.

5.1.2 Consumption Dynamics by Retirement Age

To understand the social marginal cost of providing steeper work incentives in late career, we apply our analysis of consumption dynamics to the case where the feature of interest x is the retirement age r. That is, we examine the dynamics of consumption around retirement, contrasting these dynamic patterns across individuals who retire at different ages. We start by residualizing household consumption on a set of cohort fixed effects and age fixed effects. Figure 7 Panel A plots residualized consumption as a function of time to retirement. We do this separately for premature, early, normal and late retirees. By residualizing, we effectively compare the dynamics of consumption of individuals from the same cohort, and at the same age, but who retire at different ages.[28]

During the initial period up until two years before retirement, all retirement age groups apparently experience similar trends in consumption. We report on the graph the yearly consumption levels of the four retirement age groups two years before retirement. It reveals that the premature, early and normal retirees not only experience similar trends, but also remarkably similar consumption *levels* at this point. This suggests that for these three retirement age groups, there are no sizeable differences in consumption patterns up to two years before retirement. Late retirees, however, clearly stand out in terms of consumption levels in the pre-retirement period - although they experience quite similar trends at this point. Their consumption level is already 15% larger than the other three groups two years before retirement. In other words, part of the difference in consumption *when* retired between late retirees and the other retirees originated already before retirement.

We focus next on the period just before retirement. The graph highlights significant divergence in consumption across retirement age groups in the two years leading to retirement. Premature retirees experience a clear decline in consumption just before retirement, compared to all other groups. This decline amounts to a drop of 2.5% in two years relative to their prior consumption level. And it represents a drop of almost 5% compared to the consumption trend of early and normal retirees, the latter two groups sharing extremely similar dynamics just before retirement. In contrast, the consumption of late retirees increases sharply, by about 8%, in the two years just before retirement. This finding suggests that premature retirees experience negative shocks just prior to retirement, while late retirees are hit by positive shocks. This is in line with the evidence, reviewed in Blundell et al. [2016], that earnings ability shocks are

[28]Note that the graph scales residual consumption of each group by its level two years prior to retirement. Because of the year and cohort coverage of our consumption and retirement pension data, the earliest we can observe consumption among all premature retirees is 3 years prior to retirement. And the latest we can observe consumption among all the late retirees is three years after retirement. This explains the differential coverage of the residualized consumption series in terms of event time in Figure 7.

important determinants of labor supply decisions in old age.

Following a clear fanning out of consumption levels across groups in the period just before retirement, all groups experience a strikingly similar drop in consumption, of about 5%, right at retirement. A large literature has focused on this drop in consumption at retirement, sometimes called the "retirement-consumption puzzle" (Aguiar and Hurst [2005b], Aguiar and Hurst [2013],Banks et al. [2019], Stephens and Toohey [2018]). Whether an individual's consumption drop is driven by lack of insurance on the one hand or by work-related expenditures or other complementarities between consumption expenditures and leisure on the other hand, has indeed critical implications for the mapping between consumption dynamics around retirement and the insurance value of pensions for this individual, implications to which we come back below. But importantly, we already note that consumption drops at retirement are almost identical across all groups. In other words, whatever drives the retirement consumption puzzle cannot account for the large differences in consumption when retired between individuals who retire earlier vs later.

Finally, after retirement, consumption patterns follow similar trends across all groups. The differences in consumption that emerge just prior to retirement seem to persist, more or less unaltered, well past retirement.

Panel B of Figure 7 summarizes the evidence on consumption dynamics into two moments: the estimated consumption drop in the year of retirement (i.e., between the age of retirement r and $r+1$), and the estimated consumption drop in a larger time window around retirement (i.e., between $r-2$ and $r+2$), encompassing dynamics of consumption prior to retirement. The graph confirms that while consumption drops *at retirement* are virtually identical for all groups, consumption drops *around retirement* are significantly different across retirement age groups, and exhibit a stark overall gradient by retirement age. The percentage drop in consumption around retirement of premature retirees is 6 percentage points larger than that of late retirees. But interestingly, there is once again evidence of some non-monotonicity, similarly to what we found for consumption levels in retirement: consumption drops around retirement are weakly decreasing with retirement age for the early and normal retirement age groups.

Late-career shocks The evidence on consumption dynamics from Figure 7 suggests that a significant fraction of the differences in retirement consumption across retirement age groups documented in Section 4 emerges in the last few years prior to retirement, which points to the importance of late career shocks in determining retirement consumption.

To further gauge the role of late career vs early career dynamics in retirement consumption, we replicate our estimates of consumption differences across retirement age groups from specification (12), but adding controls that capture career history and consumption history up to two years before retirement. That is, we include non-parametric controls for consumption levels two years prior to retirement, and then, we also include non-parametric controls for income levels between age 52 and 55, as well as for career length at age 55. Results displayed in Appendix Figure C-6 confirm that a significant gradient in retirement consumption remains even after controlling for these rich set of controls. We can properly measure the contribution of early vs late career dynamics to consumption differences in retirement through an Oaxaca-

Blinder decomposition of consumption differences between retirement groups in Table C-1. To this effect, we regress consumption while in retirement for each retirement age group on non-parametric controls for income levels between 52 to 55, career length at 55, and consumption levels two years prior to retirement.[29] We find that these variables explain an important part of consumption differences in retirement across all retirement age groups, but that more than 50% of consumption differences across groups remain unexplained after controlling for them, implying that more than half of the consumption differences across retirement age groups emerges in the very last stages of workers' careers.

Health shocks & work longevity risk The above evidence shows that late career dynamics play an important role in shaping retirement consumption differences across retirement age groups. We can relate these dynamic patterns to observable work longevity risk. We focus specifically on health, which is one important dimension of work longevity risk.

We first document the presence of a steep negative health gradient over retirement ages in Panel B of Figure 5. Earlier retirement is strongly associated with having significantly worse health, using two different health indices constructed using health surveys.[30] The difference in health appears to be 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. The panel also shows mortality gradients that are as pronounced. For example, premature retirees are also almost 14 percentage points more likely to have died by age 75 than late retirees.

We also examine the dynamics of health outcomes around retirement in Appendix Figure E-2. We find 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 document 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 deviations in our bad health index) as their pre-retirement level.[31]

To summarize, we find that the negative consumption shocks experienced by premature retirees just prior to retirement correlate strongly with proxies for the incidence of work longevity risk such as health shocks. This suggests that flatter pension profiles offer valuable insurance against work longevity risk for people having to retire prematurely.

Change in consumption shares around retirement As mentioned above, looking at consumption changes around retirement allows to control for persistent differences in non-consumption determinants of marginal utilities ζ, which may affect the translation of retirement consumption differences across groups into differences in social marginal utilities. But focusing on

[29] Note that we residualize first consumption on year fixed-effects, cohort fixed-effects and household structure fixed effects, to be consistent with our baseline analysis of consumption differences across retirement age groups.

[30] The ULF health index is based on both subjective and objective health measures, while the HEK health index measure is based on health expenses. More detail on the survey and the construction of the health indices is in Appendix Appendix E.

[31] Appendix Figures E-3 and E-4 show that these dynamic health patterns replicate across various health outcomes, such as the fraction reporting pain, the fraction experiencing reduced work capacity, or the fraction reporting retiring due to health reasons.

consumption drops around retirement (rather than consumption levels in retirement) comes with its own challenges. The main difficulty in translating consumption dynamics around retirement into welfare is that retirement may be associated, at the individual level, with fundamental changes to non-consumption determinants of marginal utilities ζ. This is captured in the term θ in our implementation formula (14). Such changes have been abundantly discussed in the literature on the "retirement-consumption puzzle": retired individuals have more time, they may increase their home production of goods, they may search for better prices, they spend less on work-related expenditures, etc.

Appendix Figure C-8 shows that the structure of consumption changes indeed at retirement, in line with existing evidence in the literature. For example, retirement is associated with a decline in the expenditure share of clothing, transportation and restaurants, and an increase in the share spent on housing, food and health. But what matters for the implementation of formula (14) is not whether $\theta \neq 1$ but whether θ differs across groups, as it is the ratio of θ's across retirement-age groups that determines the social marginal cost of late-career incentives. On this front, Appendix Figure C-8 is reassuring: it shows that changes in consumption structure are extremely similar across all retirement age groups.

This result further confirms that the mechanism behind the retirement-consumption puzzle is *prima facie* inconsequential to evaluate the relative insurance value of pension benefits across retirement age groups, and ultimately, the optimal pension profile by retirement age.

5.1.3 Quantifying Welfare Costs

We can now quantify the welfare costs of providing late-career incentives using the alternative implementation of the SMU's based on the consumption drops. Column (1) in Table 2 repeats the estimates of the consumption smoothing costs using the baseline implementation, but now when transferring resources between the four retirement-age groups. Column (2) shows the corresponding estimates using the alternative implementation based on consumption drops, following equation (14). We use the consumption drops from two years before to two years after retirement (see Figure 7), scaled by $\gamma = 4$.[32]

The resulting consumption smoothing costs are smaller when using the consumption drops compared to the consumption levels. Still, as shown in the empirical analysis, the differences in consumption drops around retirement capture a substantial share of the differences in consumption levels post-retirement. The cost from transferring resources from the premature to all later retires equals .21 rather than .34, and from the premature and early retirees to normal and late retirees equals .12 rather than .28. It is only when transferring resources to the late retirees that the welfare cost is substantially smaller when using the consumption drops (.14 instead of .76). Besides, the same non-monotone pattern arises with the welfare cost being lowest when steepening the incentives between the early and normal retirement age, rather than before or after. Overall, our evidence suggests that much of the loss in consumption smoothing when providing more incentives is driven by the loss of insurance against work longevity

[32] The implementation assumption here is that the welfare weights multiplied by the marginal utility of consumption before retirement are the same across the retirement-age groups.

shocks, rather than by reducing redistribution towards individuals with lower pre-retirement consumption.

5.2 Marginal Propensity 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.

5.2.1 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*

$$\frac{SMU_{x,t}}{SMU_{x',t}} \cong \frac{\omega_x}{\omega_{x'}} \times \frac{\frac{mpc_{x,t}}{1-mpc_{x,t}}}{\frac{mpc_{x',t}}{1-mpc_{x',t}}}, \tag{15}$$

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.[33] 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.[34] 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

[33] 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].

[34] 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.

two implementations, but generally hard to estimate empirically (see Chetty and Finkelstein [2013]).

5.2.2 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 et al. [2020] and Andersen et al. [2021]. 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:

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

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.[35]

Appendix F provides all the details on our sample construction and gauges sample selection issues. It also shows that conditional on a rich set of portfolio characteristics, innovations to stock prices generate persistent and plausibly exogenous wealth shocks (see also Andersen et al. [2021]). 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 $\Delta_{t,t+k}\mathrm{V}_i = \mathrm{V}_{i,t+k} - \mathrm{V}_{i,t}$ of individual i between t and $t+k$ on her passive capital gains in $t+1$, conditioning on a rich vector of portfolio characteristics $\mathbf{X}$:[36]

$$\Delta_{t,t+k}\mathrm{V}_i = \alpha_k^V \mathrm{KG}_{i,t+1} + \mathbf{X}'\beta + \nu_{itk} \quad ,\forall k \in \{-3, ..., 3\} \tag{16}$$

Figure 8 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

[35] 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 in Appendix F that our results are robust to various other approaches to dealing with outliers.

[36] 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.

permanent increase in portfolio value of about .6 krona.[37] 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.[38] Furthermore, the graph highlights that the distribution of our instrument is similar across retirement age groups.

5.2.3 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 (16):

$$\Delta_{t,t+k}C_i = \alpha_k^C \mathrm{KG}_{i,t+1} + \mathbf{X}'\beta + \varepsilon_{itk} \tag{17}$$

Panel B of Figure 8 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

[37] 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}.

[38] 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.

year, and of .15 after three years.

In Table 3, we report 2SLS estimates of MPCs corresponding to the evidence presented in Figure 8. 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:

$$\overline{C}_{i,t,t+3} - C_{i,t} = \alpha^C_{IV} \Delta_{t,t+1} V_i + \mathbf{X}'\beta + \eta_{it} \quad (18)$$

$$\Delta_{t,t+1} V_i = \alpha^V_1 \text{KG}_{i,t+1} + \mathbf{X}'\gamma + \varepsilon_{it} \quad (19)$$

Note that the vector $\mathbf{X}$ conditions on the same rich set of portfolio characteristics as in (16) 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^C_{IV} 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^C_{IV} 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-3.

Results reported in Panel A confirm the graphical evidence from Figure 8. 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 et al. [2020] who also rely on passive capital gains shocks as instruments for wealth shocks. We also report in the last column of Table 3 the estimates from a placebo test where we replicate specification (18) 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 3 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 4 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 3, 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. Appendix F provides further sensitivity analysis. In particular, we find 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.

5.2.4 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 (15). We use the MPC estimates for retired individuals, but this requires using the extended sample as reported in Table 4. On average the estimated consumption smoothing costs are similar as for the other implementations - we come back to this in Section 7 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 3. 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.

6 Evaluating the Welfare Cost of Alternative Pension Policy Dimensions

While late-career incentives have been a key focus in public discussions of pension reforms globally, pension benefits vary greatly along other dimensions as well. In this section, we deepen our analysis by examining the other dimensions of pension benefits we discussed in Section 2: early-career labor supply, income, and wealth. We analyze these other dimensions of pension benefits using similar methods to the previous two sections here, and we turn to the policy implications of these findings in Section 7.

6.1 Early-Career Labor Supply: Career Length at Age 55

We focus on career length *as of age 55* as the main feature of interest, in order to assess the consumption smoothing effects of reforms that incentivize early career labor supply, holding all else fixed.[39] While by retiring at a later age individuals lengthen their careers later in life, they tend to be rewarded by the pension system for the total number of years they have contributed, whether those come early or later in their careers. In Sweden for example, the number of contribution years was capped at 30 in the pre-reform ATP system, but this cap was lifted in the NPV system, treating contributions in all years of the career equally. More generally, when strengthening the link between pension contributions and benefits like when moving from defined benefits to defined contribution plans, one increases the rewards for work not only later in the career, but also earlier. Examining consumption by career length at age 55 sheds light on the corresponding distributional consequences.[40]

Panel A of Figure 9 illustrates how consumption varies by career length at 55, using the specification in model (12) and replacing retirement age with career length. As above, we split workers into groups based on their career length at 55. The distribution of career length is shown in Panel A of Appendix Figure B-1. We construct four roughly equal-sized groups based on quartiles of career length at age 55, with cutoffs at 29, 34, and 36 years of work experience by age 55. In the top three quartiles, we observe a negative gradient between consumption and career length. The contrast with the gradient when we considered late career labor supply in Figure 3 is striking. There we found that workers who retired later - and thus have the longest careers counting from age 55 - enjoyed significantly higher consumption than other groups. Here, the pattern is the opposite: those with long careers before 55 have 13 to 15 percent lower consumption than those with medium-length careers. We also observe a non-monotonicity in

[39]We define career length as of age 55 as the number of years prior to age 55 in which an individual had pensionable income.

[40]In principle, governments could even reform pension benefits to specifically target early-career labor supply, but they seldom do.

Figure 9: those with very short careers also have low consumption. Comparing across specifications, we observe that the negative gradient from medium-length to long-career individuals continues to hold with controls for household composition and even when controlling for income and retirement age. In contrast, controlling for income (using average income between 52 and 55 as before) and the retirement age significantly increases the relative consumption of short-career individuals versus other groups.

We briefly examine further what can explain the difference in gradients along the two dimensions. First, surprisingly, we find virtually no correlation between career length at 55 and retirement age, as shown by the retirement age distributions for the different career length quartiles in Panel B of Appendix Figure B-1. Second, we relate observables to career length at age 55 in Appendix Figure B-2, mirroring Figure 5 and revealing striking heterogeneity behind the consumption patterns in Figure 9. Focusing first on those with long careers versus those with medium-length careers, they tend be less highly educated and male, and they have somewhat higher mortality. Their income at 55 is modestly higher than those with medium-length careers but their assets are slightly lower. Note that working more than 36 years by age 55 essentially requires starting work as soon as one becomes an adult, and then working nonstop until 55. Focusing on those with short careers instead, we note that working fewer than 29 years prior to age 55 requires spending significant time outside the labor force as an adult. Those with short careers by age 55 are more likely to be female, low-income, and somewhat highly educated, and they have somewhat lower mortality. In other words, the data suggest that gender and family dynamics play a role in explaining why this group has lower consumption.

Third, we ask again whether these differences in consumption across groups emerge around retirement, or if they are more permanent. Panel A of Figure 10 reveals that the consumption differentials in Figure 9 primarily reflect longer-term consumption differentials. In every career length group, consumption is roughly level before retirement and then it falls by about 7% after retirement, and continues to fall modestly after that. We observe a modest divergence after retirement, with short-career workers having larger declines in consumption; the size of this divergence is very small compared to the 10% differences in consumption in Figure 9. In summary, all of these workers experience drops in consumption at retirement, but the differences in retirement consumption are virtually entirely due to longer-term differences in consumption across groups.

Our evidence thus indicates that, in contrast with providing late-career incentives, providing early-career incentives tends to be a progressive intervention, especially in the presence of a minimum pension to protect those with the shortest careers. As above, we can quantify the consumption smoothing cost of providing stronger incentives early in the career (holding late-career incentives fixed), which we report in Table 2. For example, transferring resources from the lower career-length quartiles to the higher career length quartiles provides gains of up to .37 cents per dollar, again for a risk aversion of $\gamma = 4$. However, such gains reflect redistribution across individuals with different pre-retirement histories: if we would condition on pre-retirement consumption, the estimated consumption smoothing gains drop down to basically zero.

6.2 Income History and Wealth

We next analyze the income dimension of the pension benefit schedule. Doing so allows us to assess the redistributive value of reforming existing minimum and maximum pensionable income thresholds, or in changing the map from annual income to pension benefits generally. As discussed above, the changes in minimum and maximum pensions in the Swedish reform disproportionately rewarded those near the bottom of the lifetime income distribution, and those near the top of the annual income distribution. To examine this dimension practically, we examine the consumption gradient over annual income at specific ages (averaged over ages 52 to 55) and over wealth (averaged over our sample period to account for volatility in asset prices).[41] Examining the gradient over wealth helps us to understand the effect of the redistribution on the basis of lifetime income embedded in many pension systems, and it informs the redistributive value of introducing explicit asset-testing in pension benefits.

Figure 9 presents also estimates of the gradient of consumption in retirement by quartiles of income (Panel B) and of wealth (Panel C). We observe large positive gradients, even larger than what we found for the retirement age or career length at 55. Those in the top income quartile enjoy 40% to 45% more consumption in retirement than those in the lowest income quartile. Adding either set of controls makes relatively little difference. Those in the top wealth quartile enjoy 45% to 70% more consumption in retirement than those in the lowest wealth quartile. When adding controls for income between 52 and 55, retirement age, and career length at 55 - other determinants of pension benefits - a substantial consumption gradient remains.

Turning to the consumption dynamics in Figure 10, we find that the consumption gradient by income in Panel B is mostly driven by longer term differences in consumption rather than differences that emerge around retirement, similar to the consumption gradient by career length in Panel A. Consumption falls by about 7% in each group at retirement and then stabilizing or declining very slightly thereafter. For wealth, however, in Panel C, we observe that the consumption drop at retirement is concentrated among individuals in the bottom quartile of the wealth distribution. Consumption in the bottom wealth quartiles moves roughly in parallel with other wealth groups until retirement, where it drops by about 12%. Other groups experience significantly smaller declines, and the estimated size of the decline is monotonic in wealth throughout. This result matches the finding in Bernheim et al. [2001] of a substantial wealth gradient in the drop in consumption at retirement. We note, however, that the larger drop in consumption at retirement for low-wealth individuals explains only about 30% of the overall wealth gradient in consumption at retirement from Figure 9.

Table 2 again translates these consumption differences in welfare cost estimates. Using the consumption-level implementation, we find that transferring resources from low-income or low-wealth retirees to high-income or high-wealth retirees can entail welfare costs of more than one dollar per dollar transferred. These estimates are substantially higher than when transferring resources across the late- or early-career dimension. However, if we disregard pre-retirement differences in consumption by employing the consumption drops implemen-

[41] Note that our use of cohort and age fixed effects accounts for the fact that wealth is measured at different ages for individuals in different cohorts in our data.

tation, the welfare cost of redistributing along the income dimension disappears, while it remains substantial along the wealth dimension, ranging between 26 and 33 cents per dollar transferred. These figures for the wealth dimension are still larger than our estimates for the retirement-age dimension. This suggests that the value of smoothing the consumption drop around retirement for lower wealth individuals relative to higher wealth individuals is higher than for earlier retirees relative to later retirees..

7 Policy Implications

This section discusses the implications of our estimated consumption smoothing costs for the design of pension policy. We draw some welfare conclusions regarding the provision of late-career incentives, but also consider the other dimensions of pension benefits.

7.1 Late-Career Incentives

Our estimates from Sections 4 and 5 indicated that incentivizing late-career labor supply through pension policy comes at a sizeable welfare cost, as it transfers resources from individuals who value them more to individuals who value them less. This is especially true when strengthening incentives at premature and late retirement ages. In other words, stronger penalties for premature retirement or rewards for late retirement are highly regressive policies. The exception in Sweden occurs when providing stronger incentives to continue working between 61 and 65.

To evaluate the net welfare gain from providing stronger late-career incentives, we should compare the consumption smoothing cost to the fiscal externality associated with the behavioral responses to the changed incentives. To obtain a simple calibration of the fiscal externality, we focus on the retirement responses and make abstraction from the fiscal implications of changes in savings behavior and other labor supply responses. We can then approximate the net welfare gain per dollar(/krona) transferred from individuals retiring before $\tilde{r}$ to individuals retiring after by

$$\Delta W_{\tilde{r}} \cong \frac{\tau_{\tilde{r}} - [NPV_{\tilde{r}+1} - NPV_{\tilde{r}}]}{w_{\tilde{r}}} \times \varepsilon_{\frac{S(\tilde{r})}{1-S(\tilde{r})}, w_{\tilde{r}}} - \frac{SMU_{r \leq \tilde{r}} - SMU_{r > \tilde{r}}}{SMU_{NRA}}. \tag{20}$$

The fiscal externality depends on the retirement response multiplied by the fiscal return to later retirement. The retirement response is mainly determined by the Frisch elasticity $\varepsilon_{S(\tilde{r}),w_{\tilde{r}}}$ at age $\tilde{r}$, underlying how much stronger incentives to continue working at age $\tilde{r}$ increase the survival rate into employment at age $\tilde{r}$.[42] The fiscal return to later retirement depends on the participation tax rate, determined by both the income tax $\tau_{\tilde{r}}$ and the implicit tax embedded in the pension benefits formula $NPV_{\tilde{r}+1} - NPV_{\tilde{r}}$. The latter accounts for the changes in the net present value of pension benefits received and payroll taxes paid when retiring one year later.

[42]Here we have assumed $\frac{\partial S(\tilde{r})}{\partial b_{r>\tilde{r}}} \cong -\frac{\partial S(\tilde{r})}{\partial b_{r\leq\tilde{r}}} \cong \varepsilon_{S(\tilde{r}),w_{\tilde{r}}} \times \frac{S(\tilde{r})}{w_{\tilde{r}}}$, where $w_{\tilde{r}}$ is the wage at age $\tilde{r}$. Since a reduction in pensions for those retiring before age $\tilde{r}$ increases their survival in employment, while an increase in pensions for those retiring after age $\tilde{r}$ reduces their survival in employment, we have also assumed that the fiscal externalities of the opposing income effects cancel out for a budget-balanced change in the pension profile. Finally, we are expressing the welfare effect relative to the value of a dollar given to our reference group of individuals retiring at the normal retirement age (SMU_{NRA}), which we assume to be approximately equal to marginal cost of public funds λ.

While prior work (e.g. Gruber and Wise [1999]) has focused on calculating the implicit tax rate due to public pension incentives alone, the fiscal externality from inducing individuals to work longer is in general dominated by the income tax paid on labor earnings and thus positive. Our simulation results suggest a participation tax rate of about .45.[43] To calibrate the elasticity, we need an estimate of the locally relevant elasticity corresponding to deviations from the same pension profile around which we have estimated consumption smoothing gains. To this effect, we use the labor supply elasticity of .22 estimated by Laun [2017] using the Swedish NDC reform discussed above.[44] Taken together, this results in a fiscal externality from inducing later retirement of .15, which means that we would gain about 15 cents per dollar transferred from individuals retiring before $\tilde{r}$ to individuals retiring after $\tilde{r}$. Appendix G and Appendix H provide more details on the derivation and the implementation of the fiscal externality respectively.

Now if the fiscal externality were constant across retirement ages, our estimates would imply that the net welfare effect of providing steeper incentives has been negative, especially below the early entitlement age and above the normal retirement age. This suggests the optimality of making the retirement incentives more S-shaped, with more muted incentives at both premature and late retirement ages, and plausibly stronger local incentives for continuing to work between early and normal retirement ages. Or, in other words, given our consumption smoothing estimates, there should be strong incentive effects to be able to rationalize higher penalties for workers leaving the labor market before the early entitlement age and higher bonuses for workers continuing to work after the normal retirement age. We note that such penalties and bonuses are central in pension reforms, including in the Swedish pension reform, which increased incentives throughout the late career and the most for individuals retiring after the normal retirement age. To obtain an estimate of the consumption smoothing cost for the overall change in late-career incentives due to the Swedish reform, we can simply weight the SMU's at each retirement age with the corresponding change in pension benefits at that age from Figure 1. Using the consumption-level implementation, we find an estimate of .37, losing 37 cents per dollar transferred from earlier to later retirees, relative to a fiscal gain of 15 cents, implying that the net welfare impact of the Swedish profile change has been negative.

We next provide some further caveats to these welfare and policy conclusions. First, our discussion so far presumed that the fiscal externality is similar across retirement ages. We note that the participation tax is indeed remarkably stable across retirement ages (see Appendix Figure A-9), so the key unknown is how the labor supply elasticity varies between early and late retirees. Existing studies have shown the importance of workers near the participation margin for labor supply elasticities around retirement vs. at prime working age (French [2005], French and Jones [2012], Blundell et al. [2016]). However, to evaluate how labor supply elasticities change with retirement age, it is also important to account for potential compositional effects

[43] See Appendix A for details. We gauge the robustness of this value when accounting for non-pension social insurance benefits and changing the claiming age in the Appendix. Such benefits increase the participation tax rate by about .05.

[44] While this elasticity is arguably the relevant one for the our welfare evaluation, we note that this elasticity is in line with estimates from reforms in other European countries (e.g. .25 in Manoli and Weber [2016] and .33 in French et al. [2020]), and somewhat higher than the elasticities found in Seibold [2021]($\approx$ .1) in Germany using bunching approaches.

and life-cycle dynamics (e.g., in health and life expectancy). The ideal experiment would leverage similar local variations in the profile of the net-present-value of pensions at different retirement ages $\tilde{r}$, in the exact same context. Seibold [2021] comes closes to this ideal setup, finding similar responses across 400 kinks in the German pension profile and no significant heterogeneity in responsiveness across observable characteristics (such as education, birth cohorts, lifetime earnings, unionization or health) that may correlate with retirement age.

Second, as we emphasized before, our implementations map consumption moments into SMU's relying on specific assumptions on preferences and welfare weights (see Table G-1). However, we can gauge the sensitivity of our conclusions to different implementation assumptions (see Appendix Table H-1). The estimated consumption smoothing cost of the change in late-career incentives is reduced when using the consumption-drop implementation (.14) or the MPC implementation (.21). In other words, the fiscal externality generated by the steepening of pension profile is thus roughly offset by the estimated consumption smoothing cost when focusing only on the corresponding loss of insurance or liquidity. Similarly, using our baseline implementation, but reducing the curvature in consumption preferences to $\gamma = 2$ decreases the consumption smoothing cost by 50% and reduces the net welfare gain close to zero (see column (2)). Adjusting the SMU's for the estimated differences in health can also reduce the consumption smoothing cost (.30), when assuming that earlier retirees having worse health also have lower marginal utility of consumption following Finkelstein et al. [2013] (see column (3)).[45] However, assigning welfare weights to retirement age groups based on their differential life-expectancy would increase the welfare costs (.40, see column (4)). Indeed, individuals retiring earlier expect a shorter lifetime, so their welfare-adjusted SMU's are higher.[46]

Third, our analysis has demonstrated the value of providing insurance against work-longevity risk. DI and UI provide complementary insurance to the pension system by covering individuals who leave the labor market early, not only through the transfers received, but also through the accumulation of pension points while on DI/UI. As discussed above, especially for the premature exits, the pathway into retirement is often through DI or UI (see also Appendix Figure A-6). We note that our consumption-based estimates of the marginal value of extra transfers account for all resources retirees can rely on, including transfers received through the DI/UI system. We also note that this pathway changes the fiscal externality from inducing individuals to work longer, but these effects are very modest, increasing the fiscal externality from .13 to .15 only (see Appendix Figure A-9).[47]

[45] We take the estimated differences in health index across retirement age groups from Figure 5 and use the upper-bound estimate of Finkelstein et al. [2013] of a 25% increase in the marginal utility of consumption per standard deviation of health. We note that the evidence in the literature on even the sign of the relation between MUC and health is mixed, either using similar methods in other contexts (e.g., Kools and Knoef [2019]) or using alternative methods (e.g., Evans and Viscusi [1991]). We provide more detail in Appendix H.

[46] We obtain differences in life expectancy using a Gompertz extrapolation of mortality rates estimated by retirement age group [Chetty et al., 2016]. We then calculate compensating consumption differentials that would equalize the expected lifetime utility for individuals with different retirement ages, following Becker et al. [2005], and use these compensating differentials to adjust the SMU's. Individuals who expect a shorter lifetime are assigned a higher SMU and vice versa. We provide more detail in Appendix H.

[47] One could alternatively define retirement not as when people stop working, but as when people stop accumulating pension points. Column (6) in Appendix Table H-2 shows how with this definition change - i.e., using the consumption estimates for the alternative retirement age definition reported in Appendix Figure C-5 - the consumption smoothing cost is smaller than in the baseline case (repeated in column 1). Indeed, redistributing

7.2 Other Policy Implications

We now broaden our perspective beyond the question of late-career incentives, informed by the results in Section 6.

First, while our results above suggest that strengthening late-career incentives is costly, our results suggest that the opposite is true for incentivizing work early in life. Early-career incentives appear to be relatively effective for redistributing between high- and low-resource individuals. However, to completely evaluate a reform strengthening early-career incentives we would need to compare the consumption smoothing effects from Table 2 to the fiscal externality from behavioral responses to these incentives, including the response of labor supply early in workers careers (e.g., French et al. [2020]). Early career-incentives could also change educational attainment and family-related career interruptions, effects which would need to be accounted for in a comprehensive welfare analysis.

Second, our results suggest that valuable redistribution could be accomplished through the income and/or wealth dimensions of pension benefits. Any benefits from redistributing along these dimensions should obviously be compared to the relevant fiscal externality due to behavioral responses in labor supply or wealth accumulation. Moreover, a natural question here is whether income taxes could accomplish the same redistribution more effectively (Atkinson and Stiglitz [1976]). Our results on the wealth dimension nevertheless suggest that wealth (or perhaps lifetime income) may be a useful tag to allocate pension benefits to those who value them the most. Not only does wealth have a strong relationship with consumption in retirement, wealth also predicts the drop in consumption at retirement, suggesting that the insurance value of pension benefits against such drop in consumption is particularly high for individuals with low wealth.

Naturally, behavioral biases that underlie under-saving could be one reason why low wealth is associated with a relatively large drop in consumption at retirement, and indeed several other differences we found in consumption across groups could be due to heterogeneity in the extent of under-saving. As we discussed in Section 2, our consumption smoothing estimates remain valid in the presence of behavioral biases like those that cause under-saving. However, fully accounting for behavioral biases would require adding a third component to the first-order welfare effect of a pension reform (equation (6)), i.e. an internality effect. Internalities arise when biases break the envelope condition, so that an individual's behavioral responses to a reform has first-order implications for his or her own welfare [Mullainathan et al., 2012; Spinnewijn, 2015]. For under-saving specifically, it also seems crucial to account for heterogeneity in biases. If some individuals are active and save optimally while others are passive and prone to under-saving like in Chetty et al. [2014], then the behavioral response to a pension reform is likely driven by the active types, and the envelope theorem logic eliminates any welfare effect on active savers. Moreover, the behavior of passive savers does not respond to the reform in the first place because they are passive, so the internality effect of a pension reform is null in this type of model, as mentioned in Section 2. Behavioral frictions - like inattention to financial

resources away from people who stop working early, including those who go on UI and DI, is costlier than from people who leave the labor market early, excluding those who go on UI and DI.

incentives or reliance on statutory retirement ages [Seibold, 2021] - may also distort retirement decisions and mute labor supply responses to pension reforms. This would dampen the size of the welfare effect occurring through the fiscal externality channel, as in Chetty et al. [2009], but could also introduce additional internality effects relevant for the evaluation of pension reforms.[48]

Finally, we note that the dimensions of pension benefits may interact. As noted above, increasing the minimum pension redistributes toward individuals with low incomes and/or short careers, whom our consumption estimates suggest are relatively needy indeed. However, a more generous minimum pension also weakens the late-career incentives for these individuals. In contrast, the effective increase in the maximum pension benefit had the opposite effect, providing stronger late-career incentives for workers with long-career and/or high-income. The average change in late-career incentives implied by Sweden's 1998 reform thus masks significant heterogeneity (see Appendix Figure A-8). We can evaluate the heterogeneity in consumption smoothing costs from increasing incentives for different subgroups. For example, the consumption differences are less pronounced for individuals in the bottom decile of pension rights accumulated by age 55 (see Appendix Figure C-4 and Table H-2), partly because of the flatter pension profile. Hence, strengthening retirement incentives for this group would have been less costly, but the Swedish reform did exactly the opposite by increasing the minimum pension.[49]

8 Conclusion

As many countries endeavor to make their pensions fiscally sustainable, they face difficult questions about which individuals should bear the burden of doing so. We found that pension reforms that incentivize later retirement specifically have a substantial and potentially pivotal redistributive cost. We reach this conclusion from an analysis of the gradient of consumption over the retirement age, drops in consumption around retirement, marginal propensities to consume, and patterns of selection into early retirement. A number of findings further suggest that work longevity risk is an important driver of the redistributive cost of incentivizing later retirement. We also find that the redistributive cost of a steeper benefits profile is largest for very early and very late retirement ages, and significantly smaller between ages 61 and 65. A similar empirical approach suggests that reforms targeting labor supply early in life have better redistributive properties than reforms targeting the retirement age. Such reforms would not help address differential exposure to work longevity risk, however, and their fiscal effects are not well understood. We also find very large redistributive effects of adjusting pension benefits along the income or wealth dimension; along the wealth dimension in particular our results suggest a sizable part of these effects is due to differences in the insurance value of pensions rather than solely an across-individual redistributive effect.

Our analysis could be extended in a number of directions in future work. First, as we briefly

[48]See Reck and Seibold [2021] for an analysis of the potential internalities that arise due to reliance on statutory retirement ages in retirement decisions.

[49]We also note that the overall gradient in consumption is somewhat smaller for couples, presumably because they can rely on intra-household insurance. This translates into substantially lower consumption smoothing costs from steeper incentives for couples than for single households (columns (4) and (5) of Appendix Table H-2).

discussed in the last section, one could delve more deeply into heterogeneity in incentives to retire later for workers with different income or earnings history. Doing so would be useful for further evaluating, for instance, minimum and maximum pension benefits. Second, reflecting our discussion of UI and DI above, one could study the optimal design of pension and other social insurance programs jointly, accounting for the sometimes fuzzy boundaries between programs (Inderbitzin et al. [2016]). Third, a caveat to our finding of a potentially optimal S-shaped pension profile is that we assume that the fiscal return to incentivizing later retirement does not vary significantly over various retirement ages. Future work could speak to this question empirically by examining how the elasticity of retirement with respect to pension incentives varies between early and late retirees. Fourth, future research could seek to explicitly estimate the size of the fiscal effects of adjusting the dimensions of pension benefits besides the retirement age. Doing so would quantify another key aspect of optimal pension benefits along these other dimensions. Fifth, future research could incorporate behavioral frictions into the analysis of the optimal steepness of pension profiles. The types of behavioral frictions that seem the most likely to matter for the evaluation of steeper retirement incentives are those affecting retirement decisions specifically (e.g., Gruber et al. [2022], Seibold [2021], Reck and Seibold [2021]).

References

Aguiar, M. and Hurst, E. (2005a). Consumption versus expenditure. *Journal of Political Economy*, 113(5):pp. 919–948.

Aguiar, M. and Hurst, E. (2005b). Consumption versus expenditure. *Journal of Political Economy*, 113(5):919–948.

Aguiar, M. and Hurst, E. (2013). Deconstructing life cycle expenditure. *Journal of Political Economy*, 121(3):437–492.

Andersen, A. L., Johannesen, N., and Sheridan, A. (2021). Dynamic spending responses to wealth shocks: Evidence from quasi-lotteries on the stock market. Working paper, Center for Economic Policy Research.

Andrews, I. and Miller, C. (2013). Optimal social insurance with heterogeneity. *Working Paper*.

Atkinson, A. and Stiglitz, J. (1976). The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics*, 6(1):55–75.

Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics*, 10(3):379–402.

Baker, S., Kueng, L., Meyer, S., and Pagel, M. (forthcoming). Consumption imputation errors in administrative data. *Review of Financial Studies*.

Banks, J., Blundell, R., Levell, P., and Smith, J. P. (2019). Life-cycle consumption patterns at older ages in the united states and the united kingdom: Can medical expenditures explain the difference? *American Economic Journal: Economic Policy*, 11(3):27–54.

Barr, N. and Diamond, P. (2009). Reforming pensions: Principles, analytical errors and policy directions. *International Social Security Review*, 62(2):5–29.

Battistin, E., Brugiavini, A., Rettore, E., and Weber, G. (2009). The retirement consumption puzzle: Evidence from a regression discontinuity approach. *The American Economic Review*, 99(5):2209–2226.

Becker, G. S., Philipson, T. J., and Soares, R. R. (2005). The quantity and quality of life and the evolution of world inequality. *American Economic Review*, 95(1):277–291.

Bernheim, B. D., Skinner, J., and Weinberg, S. (2001). What accounts for the variation in retirement wealth among u.s. households? *The American Economic Review*, 91(4):832–857.

Blundell, R., Britton, J., Dias, M. C., and French, E. (2016b). Cognition, health, employment and wages near retirement in england. Mimeo.

Blundell, R., Britton, J., Dias, M. C., and French, E. (2021). The impact of health on labor supply near retirement. *Journal of Human Resources*.

Blundell, R., French, E., and Tetlow, G. C. (2016). Retirement incentives and labor supply. In Piggot, J. and Woodland, A., editors, *Handbook of the Economics of Population Aging*, volume 1, chapter 0, pages 457–566. Elsevier.

Bound, J. (1991). Self-reported versus objective measures of health in retirement models. *The Journal of Human Resources*, 26(1):106–138.

Bozio, A., Garrouste, C., and Perdrix, E. (2021). Impact of later retirement on mortality: Evidence from france. *Health Economics*, 30:1178–1199.

Butler, J. S., Burkhauser, R. V., Mitchell, J. M., and Pincus, T. P. (1987). Measurement error in self-reported health variables. *The Review of Economics and Statistics*, 69(4):644–650.

Cesarini, D., Lindqvist, E., Östling, R., and Brügemann, B. (2016). Wealth, health, and child development: Evidence from administrative data on swedish lottery players. *The Quarterly Journal of Economics*, 131(2):687–738.

Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics*, 90(10-11):1879–1901.

Chetty, R. and Finkelstein, A. (2013). Social insurance: Connecting theory to data. In *Handbook of Public Economics*, volume 5, pages 111–193. Elsevier.

Chetty, R., Friedman, J. N., Leth-Petersen, S., Nielsen, T. H., and Olsen, T. (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark. *Quarterly Journal of Economics*, 129(3):1141–1219.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–77.

Chetty, R., Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A., and Cutler, D. (2016). The association between income and life expectancy in the united states, 2001-2014. *JAMA*, 315(16):1750–1766.

De Nardi, M., French, E., and Jones, J. B. (2016). Savings after retirement: A survey. *Annual Review of Economics*, 8:177–204.

Deshpande, M. and Lockwood, L. M. (2022). Beyond health: Nonhealth risk and the value of disability insurance. *Econometrica*, 90(4):1781–1810.

Di Maggio, M., Kermani, A., and Majlesi, K. (2020). Stock market returns and consumption. *Journal of Finance*, 75(6):3175–3219.

Diamond, P. and Mirrlees, J. (1978). A model of social insurance with variable retirement. *Journal of Public Economics*, 10(3):295–336.

Diamond, P. and Mirrlees, J. (1982). Social insurance with variable retirement and private saving. Mit working papers 296.

Diamond, P. and Mirrlees, J. (1986). Payroll-tax financed social security with variable retirement. *Scandinavian Journal of Economics*, 88:25–30.

Evans, W. N. and Viscusi, W. K. (1991). Estimation of state-dependent utility functions using survey data. *The Review of Economics and Statistics*, 73(1):94–104.

Fadlon, I. and Nielsen, T. H. (2019). Household labor supply and the gains from social insurance. *Journal of Public Economics*, 171:18–28. Trans-Atlantic Public Economics Seminar 2016.

Finkelstein, A., Luttmer, E. F. P., and Notowidigdo, M. J. (2013). What Good is Wealth Without Health? The Effect of Health on the Marginal Utility of Consumption. *Journal of the European Economic Association*, 11(suppl$_1$) : 221 − −258.

Fitzpatrick, M. and Moore, T. J. (2018). The mortality effects of retirement: Evidence from social security eligibility at age 62. *Journal of Public Economics*, 157(C):121–137.

French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *The Review of Economic Studies*, 72(2):395–427.

French, E. and Jones, J. B. (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, 79(3):693–732.

French, E. and Jones, J. B. (2012). Public pensions and labor supply over the life cycle. *International Tax and Public Finance*, 19:268–287.

French, E. and Jones, J. B. (2017). Health, health insurance, and retirement: A survey. *Annual Review of Economics*, 9:383–409.

French, E., Lindner, A., O'Dea, C., and Zawisza, T. (2020). Labor supply and the pension contribution-benefit link. Working paper.

Giupponi, G. (2019). When income effects are large: Labor supply responses and the value of welfare transfers. Cep discussion paper no. 1651.

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *The American Economic Review*, 87(1):192–205.

Gruber, J., Kanninen, O., and Ravaska, T. (2022). Relabeling, retirement and regret. *Journal of Public Economics*, 211:104677.

Gruber, J. and Wise, D. A. (1999). Introduction to "social security and retirement around the world". In Gruber, J. and Wise, D. A., editors, *Social Security and Retirement around the World*, International Social Security. University of Chicago Press.

Gustman, A. and Steinmeier, T. L. (2000). Retirement in dual-career families: A structural model. *Journal of Labor Economics*, 18(3):503–45.

Gustman, A. and Steinmeier, T. L. (2018). The role of health in retirement. *Research in Labor Economics*, 46:229–297.

Hagen, J. (2015). The determinants of annuitization: Evidence from sweden. *International Tax and Public Finance*, 22(4):549–578.

Haller, A. (2022). Welfare effects of pension reforms. Working paper.

Hendren, N. (2017). Knowledge of future job loss and implications for unemployment insurance. *American Economic Review*, 107(7):1778–1823.

Hendren, N. (2020). Measuring Ex Ante Welfare in Insurance Markets. *The Review of Economic Studies*, 88(3):1193–1223.

Hendren, N. and Sprung-Keyser, B. (2020). A unified welfare analysis of government policies. *Quarterly Journal of Economics*, 135(3):1209–1318.

Hospido, L. and Zamarro, G. (2014). Retirement patterns of couples in europe. *IZA Journal of European Labor Studies*, 3(1):1–18.

Inderbitzin, L., Staubli, S., and Zweimüler, J. (2016). Extended unemployment benefits and early retirement: Program complementarity and program substitution. *American Economic Journal: Economic Policy*, 8(1):253–288.

Jappelli, T. and Pistaferri, L. (2010). The Consumption Response to Income Changes. *Annual*

Review of Economics, 2(1):479–506.

Jappelli, T. and Pistaferri, L. (2017). *The Economics of Consumption: Theory and Evidence*. OUP USA.

Jönsson, L., Palme, M., and Svensson, I. (2012). Disability insurance, population health, and employment in sweden. In Wise, D., editor, *Social Security Programs and Retirement around the World*. University of Chicago Press, Chicago.

Karlstrom, A., Palme, M., and Svensson, I. (2004). A dynamic programming approach to model the retirement behaviour of blue-collar workers in sweden. *Journal of Applied Econometrics*, 19(6):795–807.

Kolsrud, J., Landais, C., Nilsson, P., and Spinnewijn, J. (2018). The optimal timing of unemployment benefits: Theory and evidence from sweden. *American Economic Review*, 108(4-5):985–1033.

Kolsrud, J., Landais, C., and Spinnewijn, J. (2020). The value of registry data for consumption analysis: An application to health shocks. *Journal of Public Economics*, 189:1040–1088.

Kools, L. and Knoef, M. (2019). Health and consumption preferences; estimating the health state dependence of utility using equivalence scales. *European Economic Review*, 113:46–62.

Kuhn, A., Staubli, S., Wuellrich, J. P., and Zweimüler, J. (2018). Fatal attraction? extended unemployment benefits, labor force exits, and mortality. In *Inequality and Public Policy, Trans-Atlantic Public Economics Seminar 2018*, NBER Chapters. National Bureau of Economic Research, Inc.

Lalive, R., Magesan, A., and Staubli, S. (2022). How social security reform affects retirement and pension claiming. *American Economic Journal: Economic Policy*. Forthcoming.

Landais, C. and Spinnewijn, J. (2021). The value of unemployment insurance. *Review of Economic Studies*, 6(88):3041–3085.

Laun, L. (2017). The effect of age-targeted tax credits on labor force participation of older workers. *Journal of Public Economics*, 152(C)(102-118).

Laun, T. and Wallenius, J. (2015). A life cycle model of health and retirement: The case of swedish pension reform. *Journal of Public Economics*, 127(C):127–136.

Lockwood, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, 108(9):2513–50.

Manoli, D. S. and Weber, A. (2016). Nonparametric evidence on the effects of financial incentives on retirement decisions. *American Economic Journal: Economic Policy*, 8(4):160–182.

Manoli, D. S. and Weber, A. (2018). The effects of the early retirement age on retirement decisions. Nber working paper no. 22561.

Mullainathan, S., Schwartzstein, J., and Congdon, W. J. (2012). A reduced-form approach to behavioral public finance. *Annual Review of Econonomics*, 4(1):511–540.

Ndiaye, A. (2020). Flexible retirement and optimal taxation. Working paper.

O'Dea, C. (2018). Insurance, efficiency and the design of public pensions. 2018 meeting papers 1037, society for economic dynamics.

OECD (2015). *Pensions at a Glance 2015: OECD and G20 Indicators*. OECD Publishing, Paris.

OECD (2017). Recent pensions reforms. In *Pensions at a Glance 2017: OECD and G20 Indicators*.

OECD Publishing, Paris.

OECD (2019). *Pensions at a Glance 2019: OECD and G20 indicators*. OECD Publishing, Paris.

Reck, D. and Seibold, A. (2021). The welfare economics of reference dependence. Working paper.

Seibold, A. (2021). Reference points for retirement behavior: Evidence from german pension discontinuities. *American Economic Review*, 111(4):1126–25.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, 13(1):130–167.

Staubli, S. and Zweimüler, J. (2013). Does raising the early retirement age increase employment of older workers? *Journal of Public Economics*, 108(C):17–32.

Stephens, M. and Toohey, D. (2018). Changes in nutrient intake at retirement. Nber working paper no. w24621.

Stern, S. (1989). Measuring the effect of disability on labor force participation. *Journal of Human Resources*, 24(3):361–395.

Figure 1: REFORMING THE PROFILE OF PENSION BENEFITS OVER AGE AT RETIREMENT

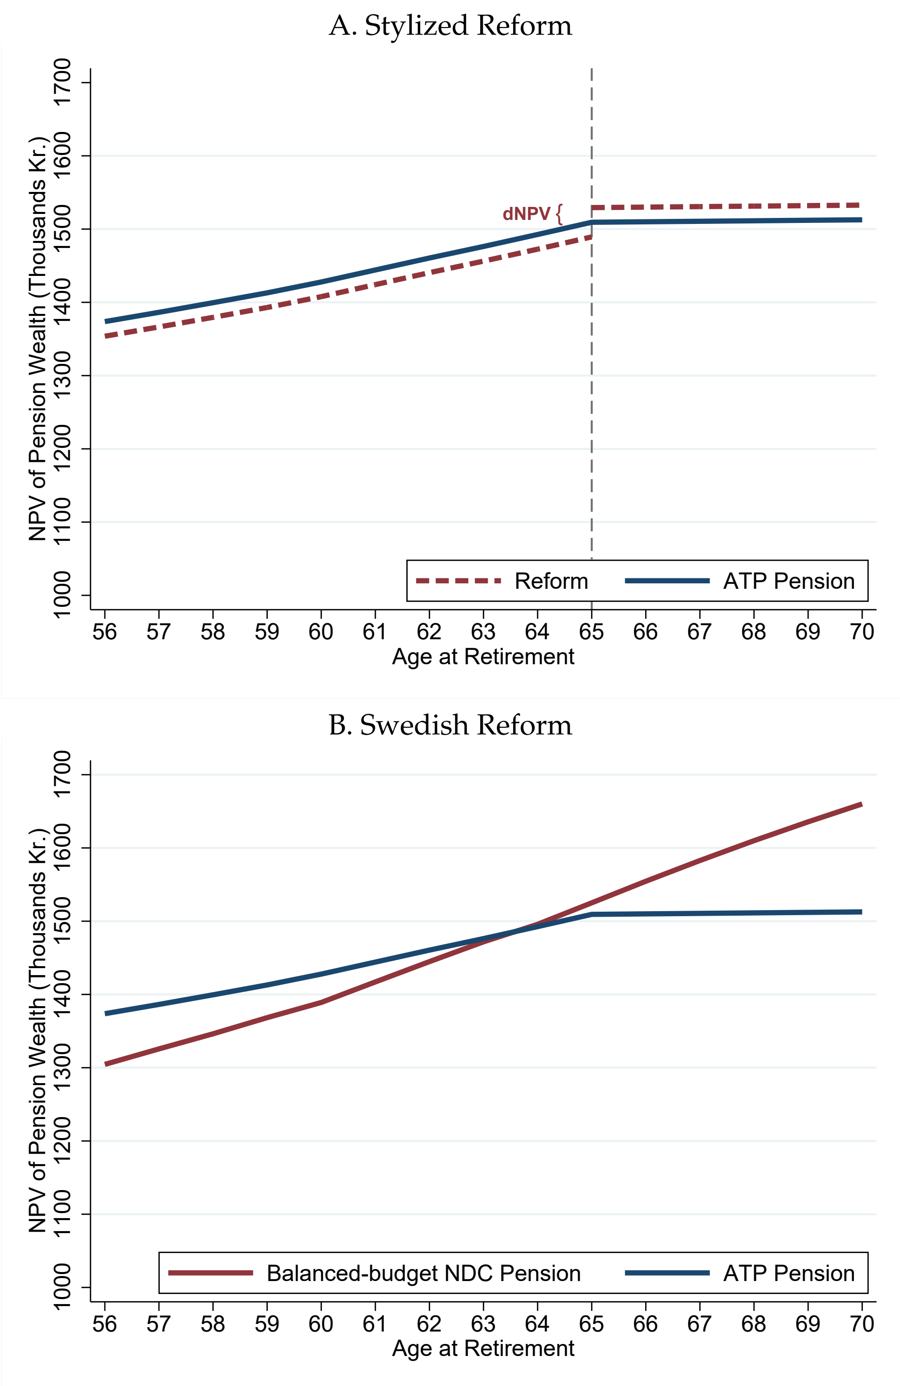

Notes: Panel A shows the effect of the Swedish pension reform on the net present value of pension wealth by age at retirement averaged across vigintiles of accrued pension rights (ATP points) at age 55. Calculations are for individuals born in 1941 with a discount factor of 0.98. To focus on the effect of the reform on the slope of the pension profile, we remove the level effect of the NDC reform on pension benefits, and call the resulting schedule "balanced budget NDC" – see also Figure A-12. Panel B illustrates a stylized balanced-budget reform in the pension profile that increases pension benefits above age 65 and decreases them below that age. Our theoretical model characterizes the welfare effects of the reform like that of Panel A, and a combination of age-specific reforms can be used to approximate the reform in Panel B.

Figure 2: DISTRIBUTION OF RETIREMENT AGE

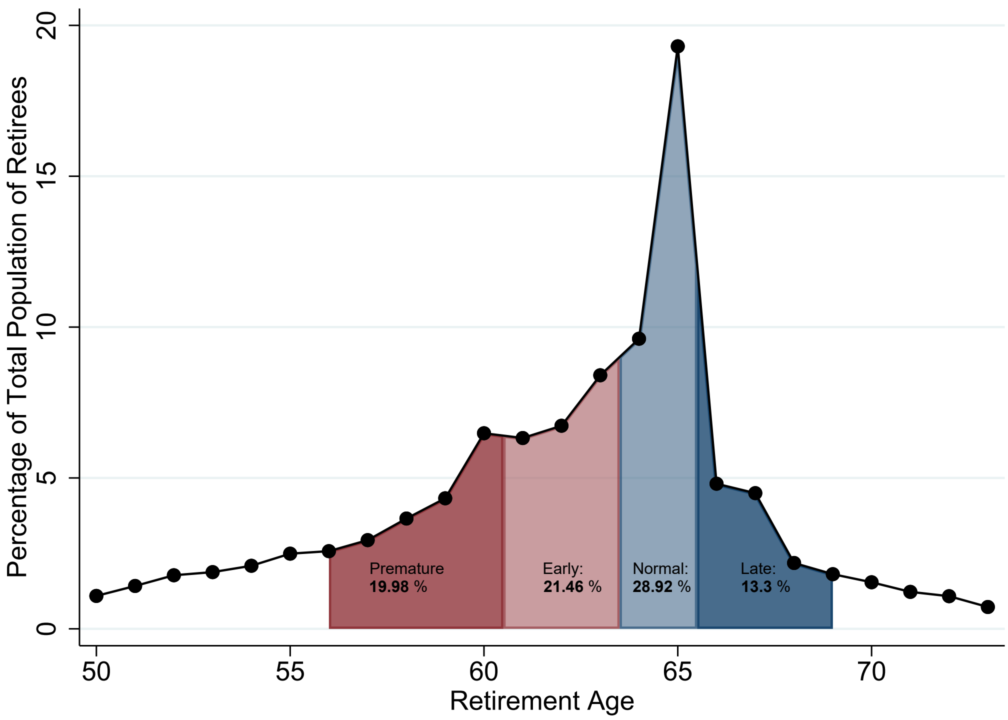

Notes: The figure reports the distribution of age at retirement among individuals from the 1938 to 1943 cohorts in Sweden. Retirement is defined as labor earnings dropping permanently below one Base Amount. In our empirical analysis, we group individuals into for categories of retirement age. Premature 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. For each group, we report the total fraction of individuals retiring in that group among the 1938 to 1943 cohorts. In the rest of the analysis, we drop from our sample the small group of individuals whom we observe retiring before 55, or after 70.

Figure 3: CONSUMPTION DIFFERENCES IN RETIREMENT ACROSS RETIREMENT AGE GROUPS

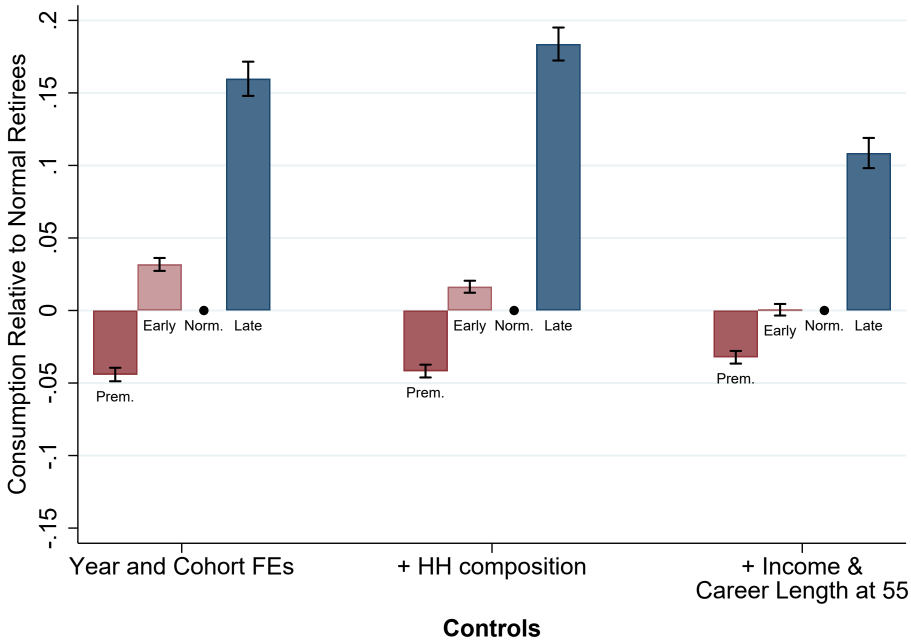

Notes: The figure documents how consumption in retirement differs across individuals who retire at different ages. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed. Individuals are grouped into four retirement age categories: premature retirees ($56 \le r \le 59$), early retirees ($60 \le r \le 63$), normal retirees ($64 \le r \le 65$) and late retirees ($66 \le r \le 69$). Normal retirees are the reference category. The graph reports, for all retirement age groups, the estimated coefficients α_j from specification (12), 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 (12) 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 $\mathbf{X}$.

Figure 4: DECOMPOSITION OF CONSUMPTION EXPENDITURES AT AGE 68 BY RETIREMENT AGE

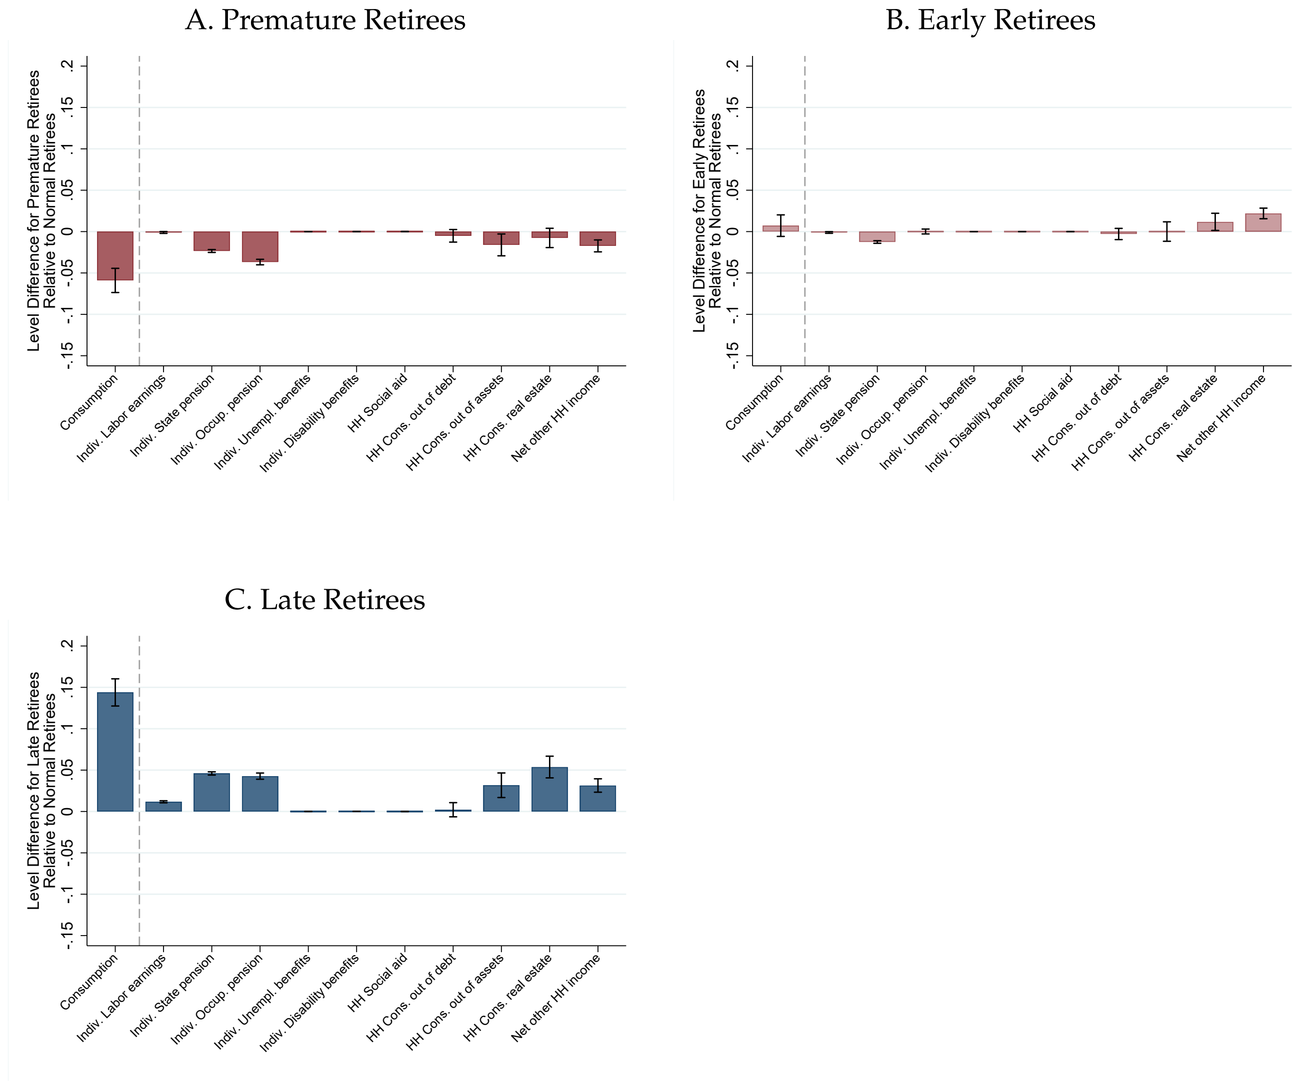

Notes: The figure decomposes consumption differences at age 68 across individuals who retire at different ages. The sample comprises all individuals from cohorts 1938 to 1943 who are retired age 68, and individuals are grouped into four retirement age categories: premature retirees ($56 \leq r \leq 60$), 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 (12) separately for each component evaluated at age 68, 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.

Figure 5: HETEROGENEITY & SELECTION INTO RETIREMENT AGE

A. Socio-Demographic Characteristics

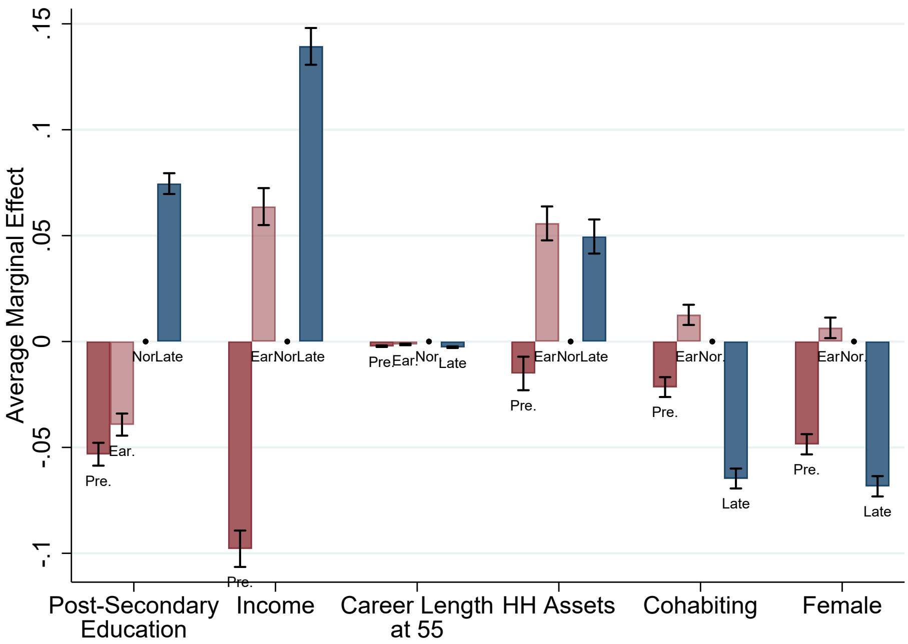

B. Health and Life Expectancy

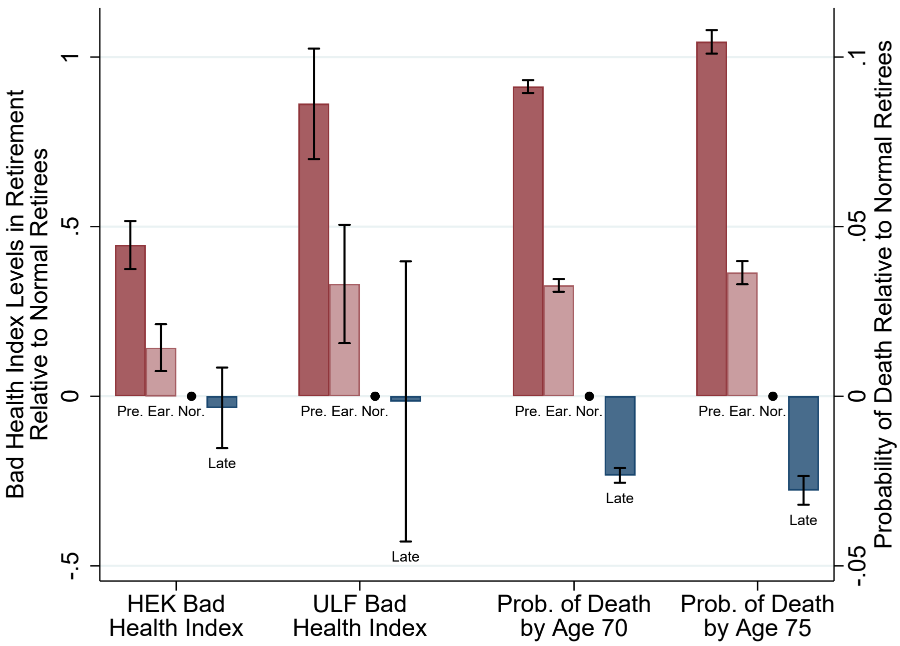

Notes: The figure documents patterns of heterogeneity across retirement age groups. Panel A displays estimates from a multinomial logit prediction model for retiring in one of the 4 different age groups. The regression sample includes one observation for each of the 418,033 unique individuals of our baseline sample. The model includes cohort fixed effects, a dummy for having post-secondary education, the within-cohort rank of average income between 52 and 55, years of career length at 55, 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 specification (12) (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; see Figure E-1 for other health outcomes) 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.

Figure 6: DIFFERENCES IN CONSUMPTION SHARES DURING RETIREMENT

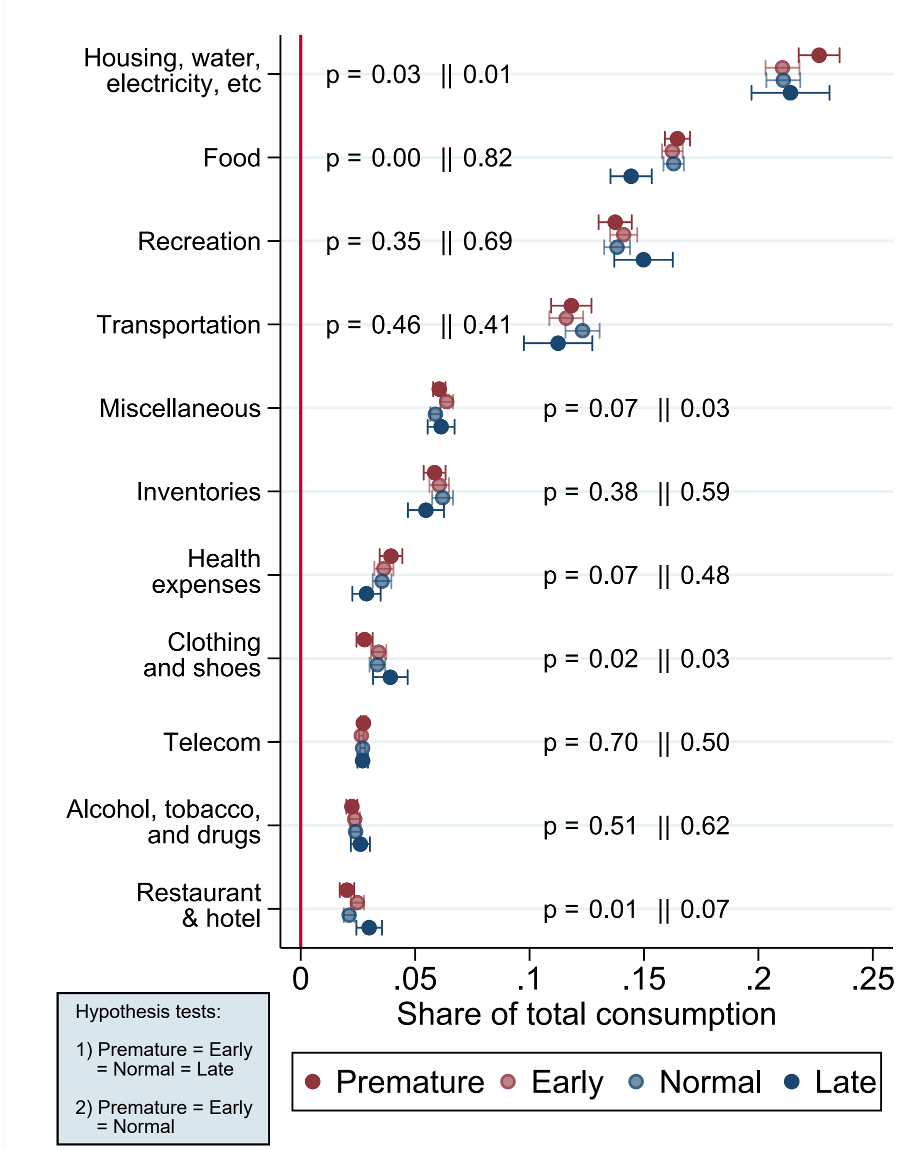

Notes: The figure examines consumption patterns in retirement. Using 3,373 observations from the HUT survey, total consumption is divided into 11 spending categories. The share of total consumption represented by each category is then regressed on a set of year and cohort fixed effects, a dummy for being 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 7: CONSUMPTION DYNAMICS AROUND RETIREMENT, BY RETIREMENT AGE GROUP

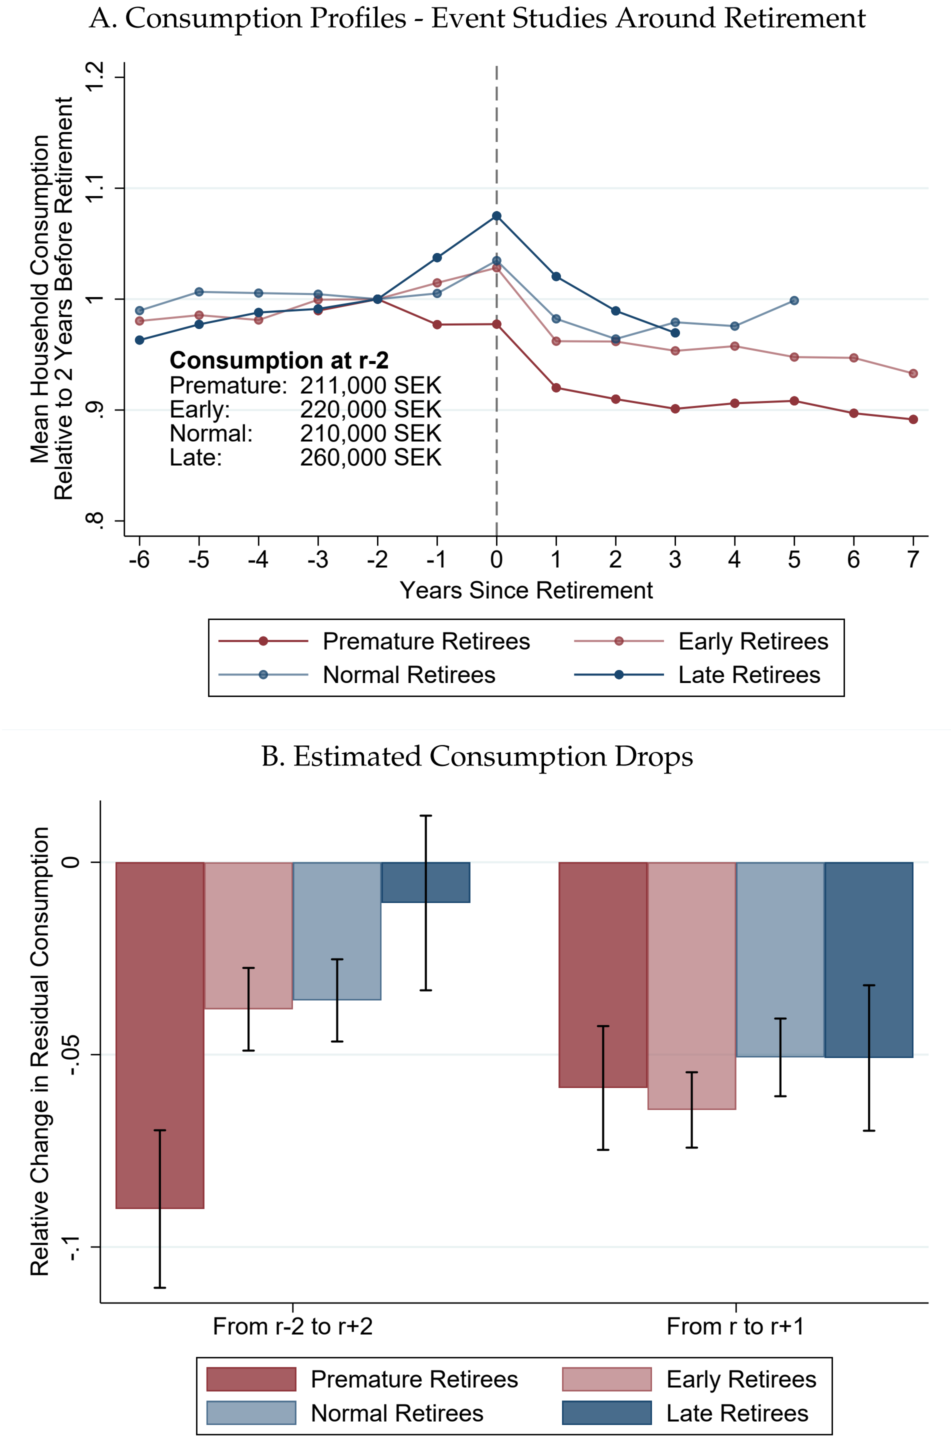

Notes: The figure documents consumption dynamics around retirement. In both panels, household consumption is first residualized on a set of cohort fixed effects and age fixed effects and household structure controls, as in specification (12). Panel A plots average residualized consumption as a function of time to retirement, separately for premature, early, normal and late retirees. The graph scales residual consumption of each group by its level two years prior to retirement (this level is also reported on the graph). Because of the year and cohort coverage of our consumption and retirement pension data, the earliest we can observe consumption among all premature retirees is 3 years prior to retirement. And the latest we can observe consumption among all the late retirees is three years after retirement. This explains the differential coverage of the residualized consumption series. Panel B reports, for each retirement age group, estimates of residual consumption changes in a 5 year period around retirement (from $r-2$ to $r+2$) and just at retirement (from r to $r+1$). The latter drop has been the focus of the "retirement-consumption puzzle" literature.

Figure 8: MARGINAL PROPENSITIES TO CONSUME OUT OF WEALTH SHOCKS

A. First Stage: Portfolio Value

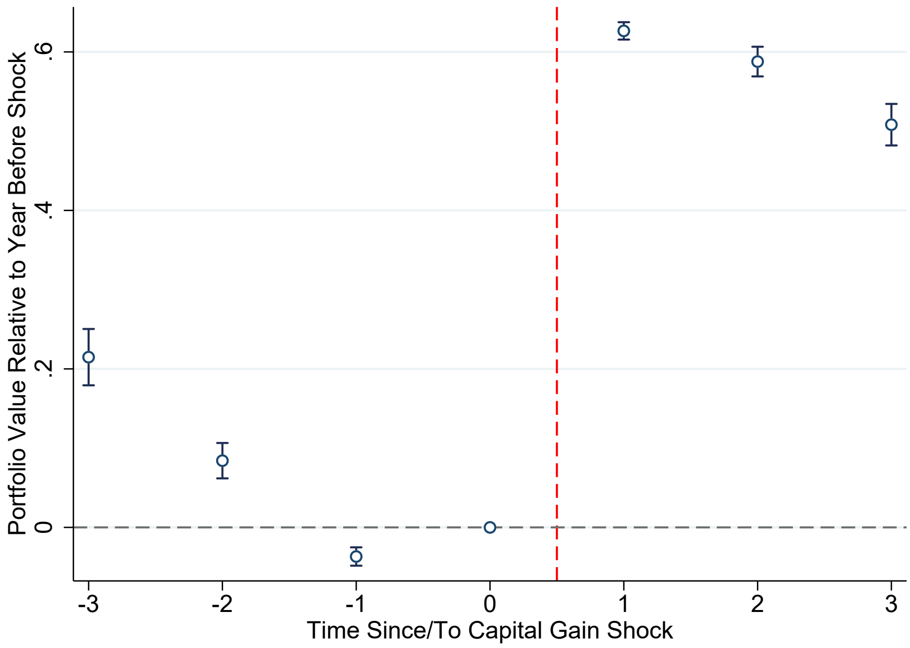

B. Reduced-Form: Consumption

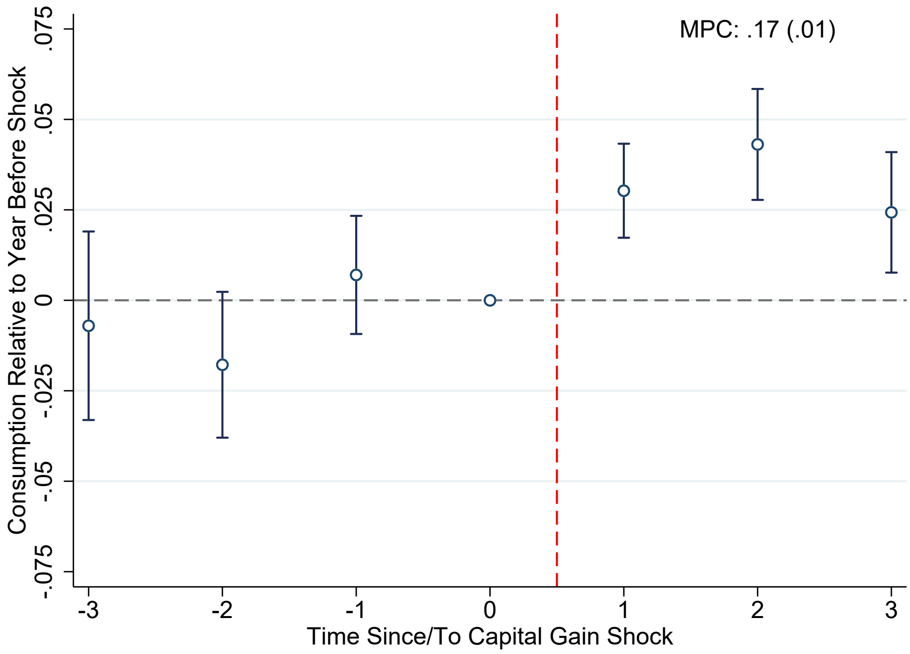

Notes: Panel A reports the estimates of the first stage regression, that is the regression of the change in portfolio value between t and $t+k$ at year t on the passive capital gains at $t+1$, that is a year after the wealth shock, controlling for the value of the portfolio in year t, as well as for the average returns and variance of the portfolio in the 6 years prior to year t (see equation (16)). 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 (18)). 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.

Figure 9: CONSUMPTION DIFFERENCES IN RETIREMENT ACROSS ALTERNATIVE POLICY DIMENSIONS

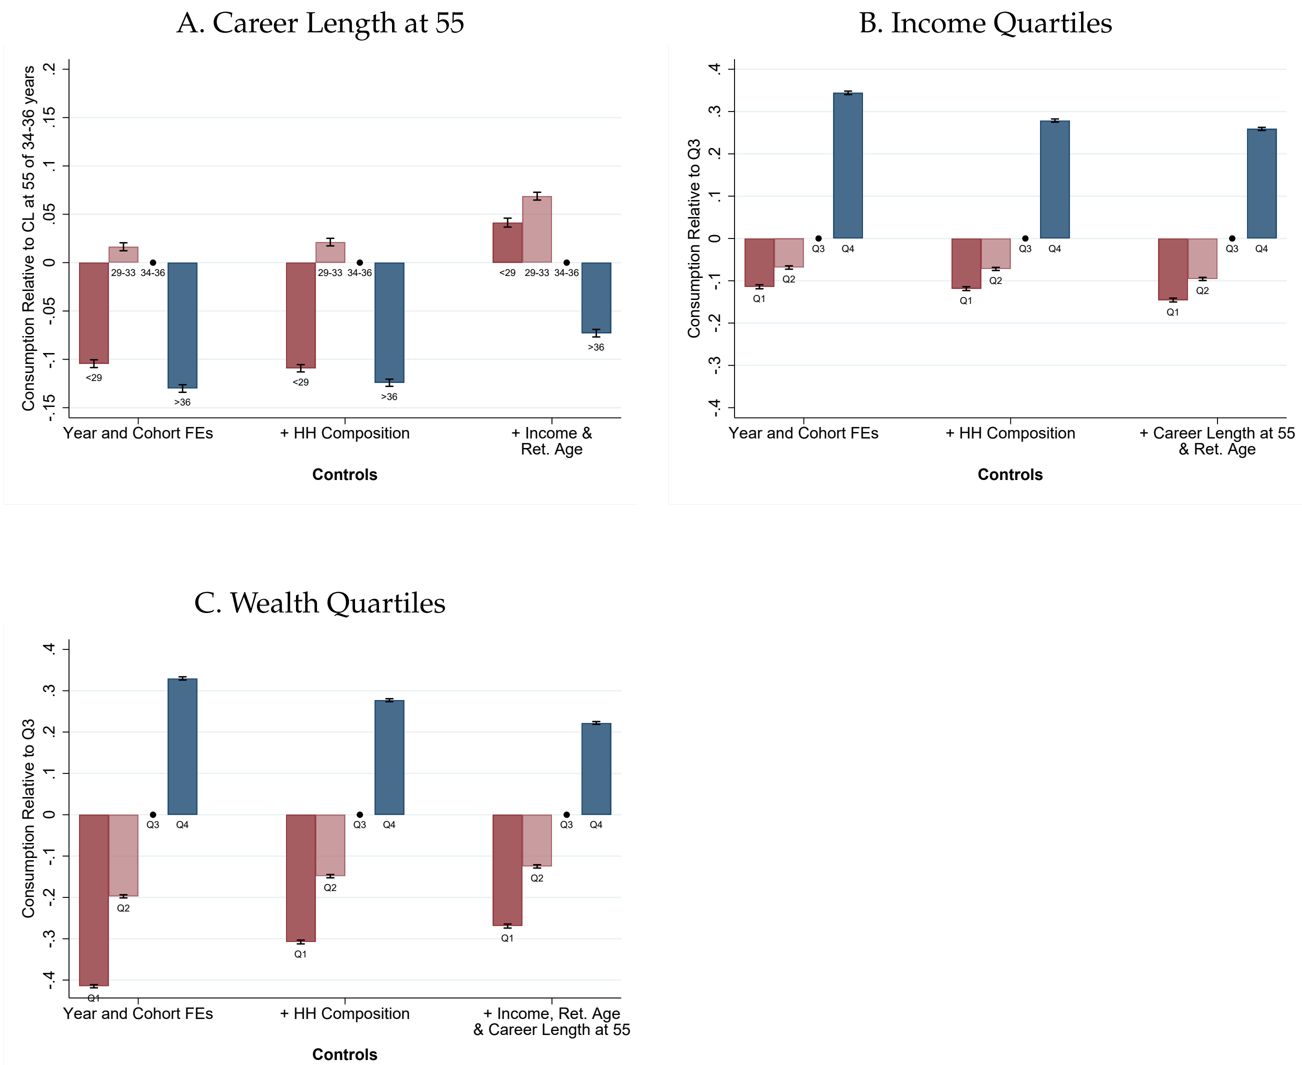

Notes: The figure documents how consumption in retirement differs across alternative dimensions of pension policy. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed. In Panel A, individuals are grouped into four career length at age 55 categories, roughly based on quartiles: fewer than 29 years, between 29 and 33 years, between 34 and 36 years, and more than 36 years of contribution. In Panels B and C, individuals are grouped into within-cohort quartiles of average income between ages 52 to 55 and average household wealth between 1999 and 2007, respectively. In all cases, the third group is the reference category. The graph reports, for all groups, the estimated coefficients α_j from the analogue of specification (12), scaled by $\mathbf{E}_j[\tilde{C}_{it}]$, the average level of consumption of individuals in the reference group from the same cohort, age, family composition and other control variables as the average individual in group j. As in Figure 9, we begin on the left-hand side with estimates from model (12) including only year and cohort fixed effects, and then we add controls for family composition and then further controls for other determinants of pension benefits.

Figure 10: Consumption Dynamics around Retirement Across Alternative Policy Dimensions

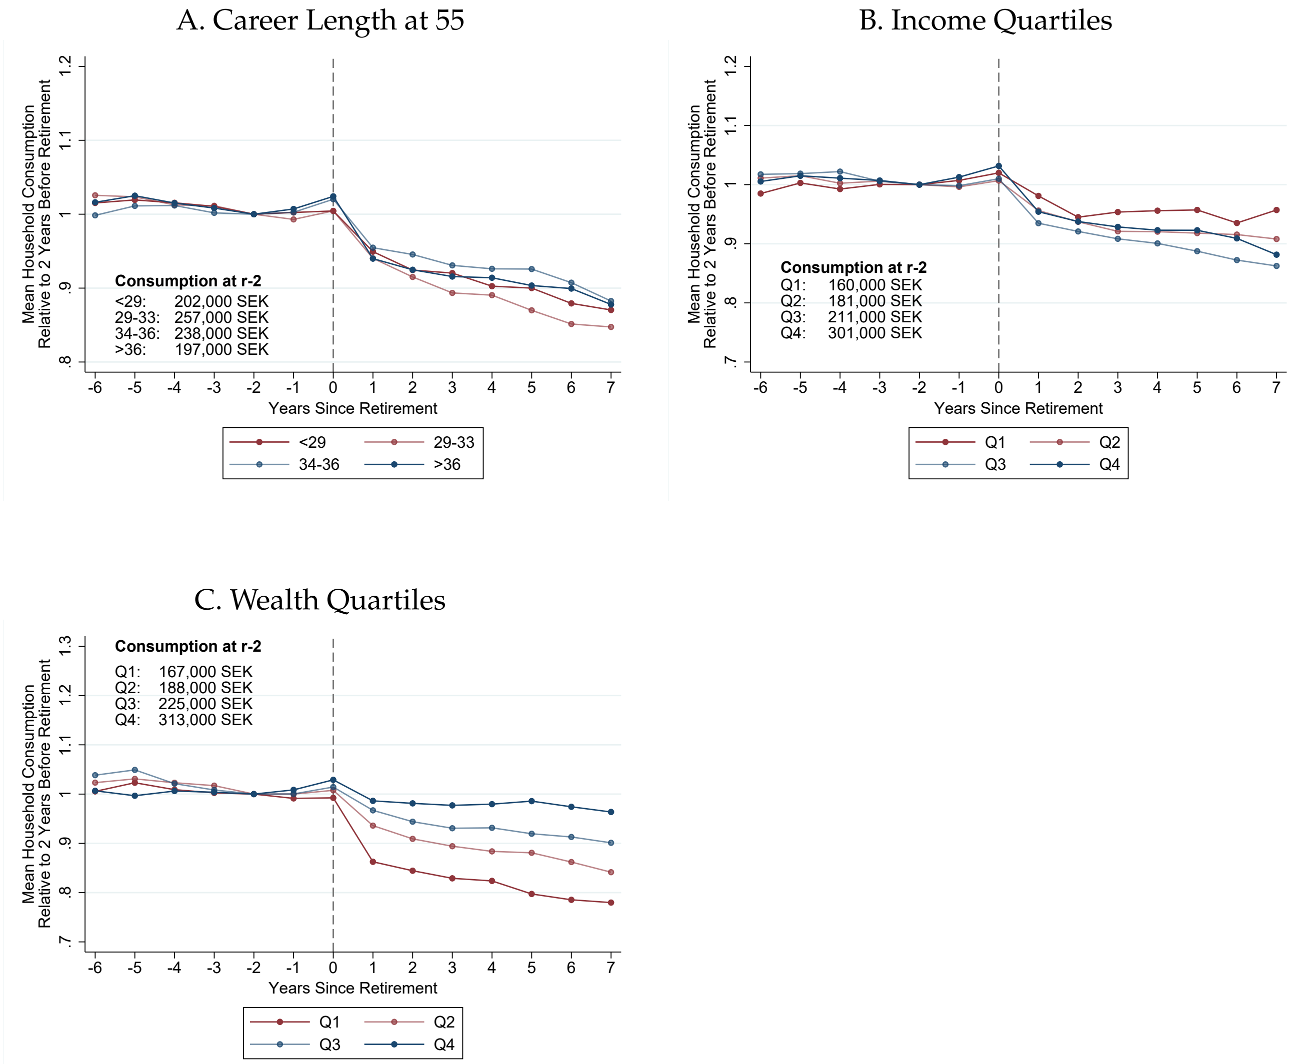

Notes: The figure documents consumption dynamics around retirement across other policy dimensions. As in Figure 7, household consumption is first residualized on a set of cohort fixed effects, age fixed effects and household structure controls, following specification (12). Panel A plots average residualized consumption as a function of time to retirement, separately for each group of career length at 55. Panels B and C separate by within-cohort deciles of average income between ages 52 to 55 and average household wealth between 1999 and 2007, respectively. Each graph scales residual consumption of each group by its level two years prior to retirement (this level is also reported on each graph).

Figure 11: Welfare Impact of Steeper Pension Profile by Retirement Age

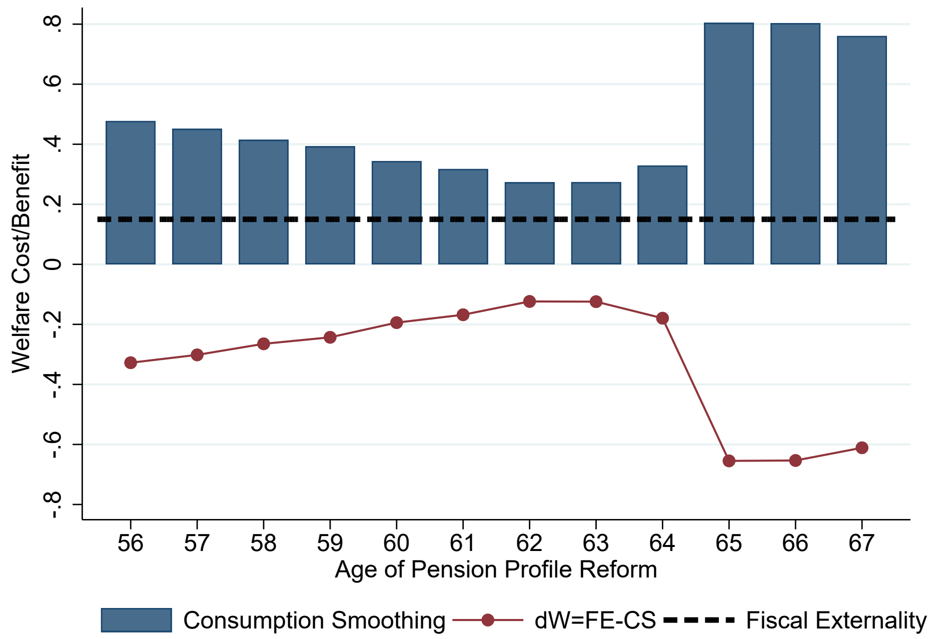

Notes: This figure reports the consumption smoothing cost of steepening the pension profile at different retirement ages (blue bars) and benchmarks them with the fiscal externality gain (dashed line), following equation (20). The difference between the two captures the net welfare impact (red line). The terms correspond to the welfare effects of transferring a dollar for individuals retiring *at or before* a specific age to individuals retiring *after* that age. The consumption smoothing costs follow our baseline implementation,

$$\frac{SMU_{r\leq\tilde{r}} - SMU_{r>\tilde{r}}}{SMU_{NRA}} \approx \gamma \times \left[\frac{E_{r>\tilde{r}}(c)}{E_{r\in NRA}(c)} - \frac{E_{r\leq\tilde{r}}(c)}{E_{r\in NRA}(c)}\right], \tag{21}$$

where the differences in consumption levels are based on estimates in regression (12) and γ is set at 4. Further details on the computation of the welfare terms are provided in Appendix H. The sensitivity of the estimates is explored in Tables 2, H-1 and H-2.

Table 1: DESCRIPTIVE STATISTICS: RETIREMENT SAMPLE

	Mean (1)	(s.d.) (2)
I. Retirement		
Fraction of Premature Retirees	23.81 %	
Fraction of Early Retirees	25.68 %	
Fraction of Normal Retirees	34.60 %	
Fraction of Late Retirees	15.91 %	
Age at Retirement	62.91	(3.1)
II. Demographics		
Cohort	1940.67	(1.73)
Fraction Men	49.29 %	(50)
Fraction Married	66.86 %	(47.07)
Kid at Home ($\geq$ 1)	17.65 %	(38.12)
Kid at Home Under 18 ($\geq$ 1)	3.48 %	(18.33)
Post-Secondary Education	24.67%	(43.11)
III. Income and Wealth at 59, SEK 2003 (K)		
Total Earnings	209	(160)
Net Wealth	777	(2339)
Bank Holdings	84	(312)
Portfolio Value	248	(1648)
Consumption	201	(534)
IV. Pensions		
State pension	78.5	(52.9)
Occupational Pension	62.1	(92.6)
ATP Pension at 55	95.6	(38.1)
N (Unique Individuals)	418,033	

Notes: The table reports descriptive statistics from our baseline sample of retirees. The sample is restricted to cohorts 1938 to 1943 who retire between age 56 and 69. The sample comprises 418,033 unique individuals. Retirement is defined as labor earnings dropping permanently below one Base Amount. Panel I reports statistics on the distribution of retirement 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 focuses on income and wealth measured at age 59. Wealth and consumption is aggregated at the household level. Panel IV 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. Note that based on the average exchange rate between 2000 and 2007, 1SEK $\approx$ 0.11USD.

Table 2: CONSUMPTION SMOOTHING COSTS OF PENSION REFORMS

	Baseline	**Cons. drops**
	Cons. levels	ΔC
	$\gamma = 4, \theta = 1$	$\gamma = 4$
	(1)	(2)
	A. Retirement Age	
Premature $\longrightarrow$ Early-Late	.34	.21
Prem.-Early $\longrightarrow$ Normal-Late	.28	.12
Prem.-Normal $\longrightarrow$ Late	.76	.14
	B. Career Length at 55	
Q1 $\longrightarrow$ Q2-Q4	.27	.01
Q1-Q2 $\longrightarrow$ Q3-Q4	-.11	.04
Q1-Q3 $\longrightarrow$ Q4	-.37	-.01
	C. Income	
Q1 $\longrightarrow$ Q2-Q4	.74	-.04
Q1-Q2 $\longrightarrow$ Q3-Q4	.89	-.03
Q1-Q3 $\longrightarrow$ Q4	1.32	.02
	D. Wealth	
Q1 $\longrightarrow$ Q2-Q4	1.40	.33
Q1-Q2 $\longrightarrow$ Q3-Q4	1.45	.27
Q1-Q3 $\longrightarrow$ Q4	1.69	.26

Notes: This table presents the estimated consumption smoothing costs of budget-neutral pension reforms that redistribute across a given policy dimension. The reforms in Panel A consist in providing steeper incentives at each retirement age $\tilde{r}$ in a specific interval, $\frac{SMU_{r<\tilde{r}} - SMU_{r>\tilde{r}}}{SMU_{NRA}}$, where $\tilde{r} \in 60, 63, 65$ coincides with the cutoffs between the retirement age groups. Panels B, C and D transfer across quartiles of the distributions of career length at 55, average income between 52 and 55 and average household wealth between 1999 and 2007, respectively. The consumption smoothing costs are expressed per dollar transferred, following equation 20. In the case of retirement age (Panel A), these costs can be compared to our benchmark fiscal externality of .15 to evaluate the net welfare gain from a reform (apart from behavioral internality effects). Column (1) reports the results for the baseline implementation, using the difference in consumption levels to approximate the difference in SMU's (see equation (11)). In column (2), we show the results for the implementation using the difference in consumption drops to approximate the difference in SMU's (see equation (14)). Appendix H provides more details underlying the welfare calculations.

Table 3: 2SLS Estimates of Marginal Propensity to Consume Out of Wealth Shocks

	First Stage α_1^V	Reduced Form $3 \text{ x } \alpha_{rf}^C$	IV Result $3 \text{ x } \alpha_{IV}^C$	Placebo Test α_1^P
	A. Whole Sample			
	.66	.11	.17	.04
	(.01)	(.01)	(.01)	(.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	.66	.09	.13	-.02
	(.01)	(.01)	(.01)	(.02)
N	472,980	472,980	472,980	472,980
# of Indiv. Clusters	129,099	129,099	129,099	129,099
Retired in t	.71	.21	.30	.22
	(.03)	(.03)	(.04)	(.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	.69	.23	.34	-.01
	(.04)	(.03)	(.04)	(.07)
N	48,425	48,425	48,425	48,425
# of Indiv. Clusters	16,667	16,667	16,667	16,667
Early Retirees	.63	.22	.34	.03
	(.02)	(.02)	(.03)	(.03)
N	172,747	172,747	172,747	172,747
# of Indiv. Clusters	46,972	46,972	46,972	46,972
Normal Retirees	.68	.06	.09	.03
	(.01)	(.01)	(.02)	(.02)
N	261,351	261,351	261,351	261,351
# of Indiv. Clusters	56,705	56,705	56,705	56,705
Late Retirees	.70	.00	.01	.06
	(.03)	(.03)	(.04)	(.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 (18) and (19). 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 (17) 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 (12) 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 4: 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)			
	.70	.19	.28	.06
	(.02)	(.02)	(.04)	(.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	.77	.18	.24	.04
	(.05)	(.05)	(.06)	(.09)
N	20,064	20,064	20,064	20,064
# of Indiv. Clusters	9,595	9,595	9,595	9,595
Early Retirees	.69	.26	.37	-. 09
	(.03)	(.04)	(.06)	(.13)
N	38,301	38,301	38,301	38,301
# of Indiv. Clusters	21,118	21,118	21,118	21,118
Normal Retirees	.69	.25	.36	.03
	(.03)	(.04)	(.06)	(.05)
N	37,647	37,647	37,647	37,647
# of Indiv. Clusters	21,036	21,036	21,036	21,036
Late Retirees	.78	.05	.07	.08
	(.05)	(.07)	(.08)	(.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 3 where we enlarge the set of cohort to 1932 - 1943 and restrict to individuals retired at the moment of the capital gain shock.

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. 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. Pension benefits in both the ATP system and the NDC system are financed by payroll taxes.

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).

The ATP system. 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,[50] 2) the total number of years in which an individual earns pension rights (with a maximum of 30 years), and 3) the claiming age.

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.[51] Annual ATP points are capped at 6.5 BAs, which corresponds 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:

[50] 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.

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

$$b_{it} = 0.6 \cdot AP_i \cdot \min\left(\frac{N_i}{30}, 1\right) \cdot BA_t, \tag{22}$$

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.[52] 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, pensionable earnings grow. The income base amount replaces the old base amount (BA) and is indexed to average wage growth instead of prices.[53] 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*. 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.[54]

[52]Formally, the basic pension for singles is calculated as $1.529 \cdot \min(H_i/30, 1) \cdot BA_t$ 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.

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

[54]Scenarios can be found in this government report (in Swedish): http://www.sou.gov.se/wp-content/

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 et al. [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]).

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

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.

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*.[55] 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.

Figure A-1: THE BASE AMOUNT (BA) AND THE ENHANCED BASE AMOUNT (EBA), 1991-2011

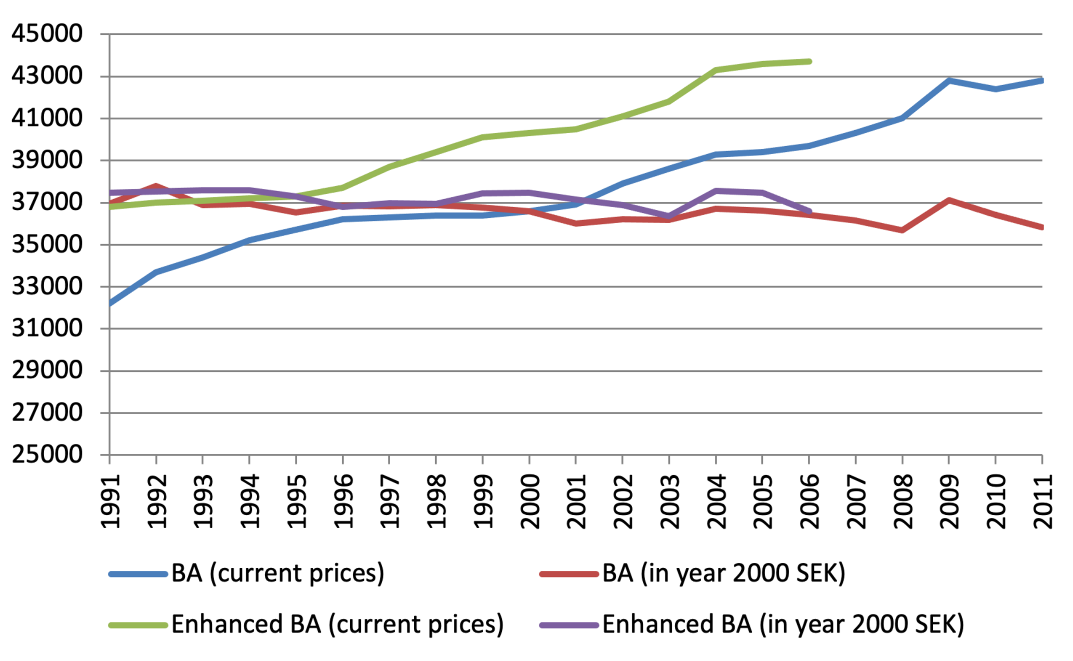

Notes: This figure shows the Base Amount (BA) and Enhanced Base Amount (EBA) over time. Both the BA and EBA are indexed against inflation.

[55] 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-2: RELATIONSHIP BETWEEN INCOME DEPENDENT PENSION AND MINIMUM GUARANTEE

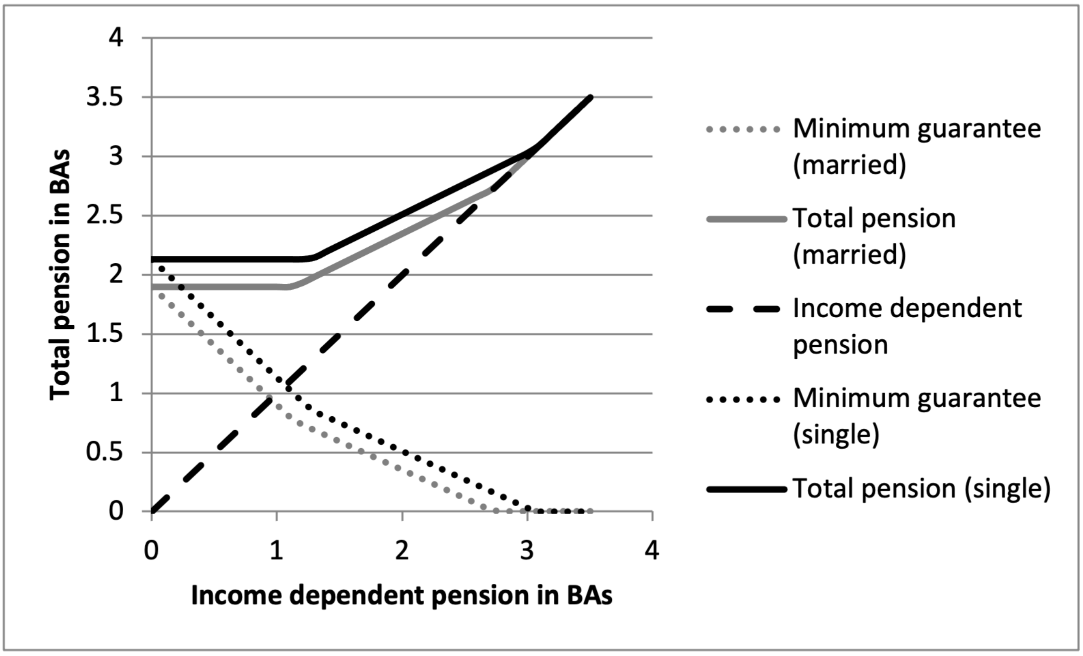

Notes: The income related pension is the same for singles and married. Total pension is the sum of the income dependent pension and the minimum guarantee.

Figure A-3: DISTRIBUTIONS OF JOB EXIT AND PENSION CLAIMING AGES

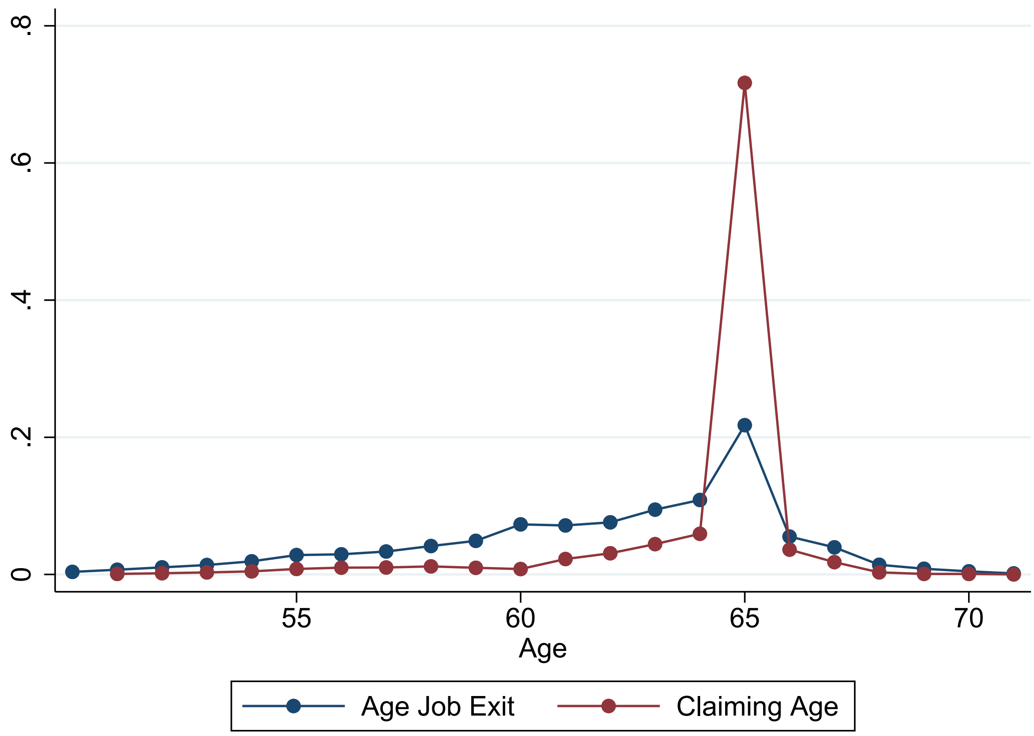

Notes: This figure shows the density distributions of job exit age and pension claiming age for workers born between 1938 and 1943.

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 wishes 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:[56]

- 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.
- 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, 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. 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,

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

given the number of years worked and earnings at 55.[57] 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 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

$$\text{Participation Tax Rate}_a = \frac{\text{income tax}_a + \text{payroll tax}_a - [NPV_a - NPV_{a-1}]}{\text{Gross earnings}_a}, \tag{23}$$

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.[58] 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:

$$\text{Implicit Tax Rate}_a = \frac{\text{pension payroll tax}_a - [NPV_a - NPV_{a-1}]}{\text{Gross earnings}_a}. \tag{24}$$

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

[57] 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.

[58] 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.

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 by lifetime income. We use some empirical moments to calibrate our simulations along these lines.

Specifically, we divide the sample of individuals born from 1988-1943 – 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 lifetime earnings; 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 full population of interest. We run the simulator described above 20 different times, where the worker characteristics are based on the characteristics of a 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

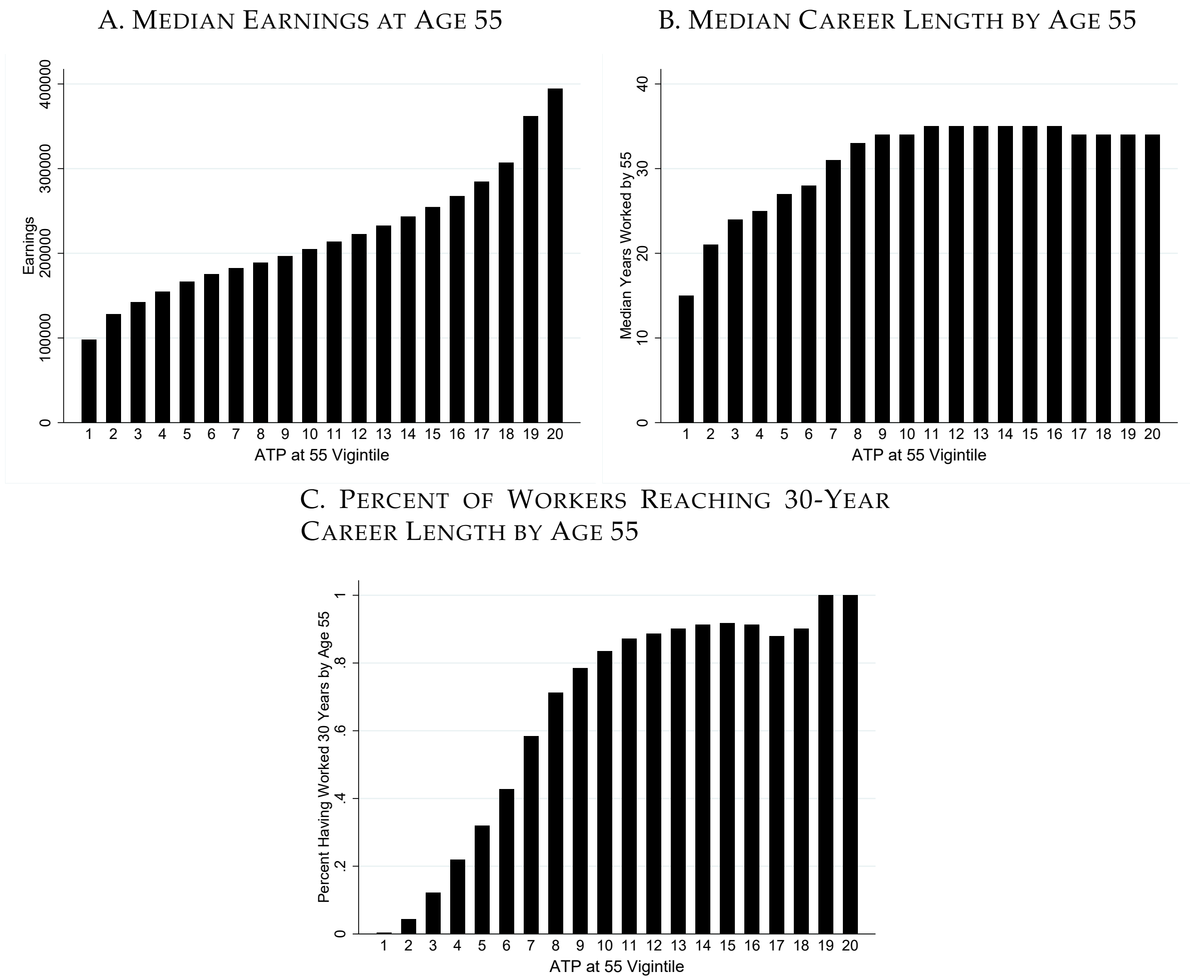

Notes: Panel A of this figure shows the median earnings at 55 of workers born between 1938 and 1943 for each ATP at 55 vigintile. Panel B shows the median years worked by age 55 for each vigintile. Panel C shows the percent who reach a career length of 30 years by the age of 55 for each ATP at 55 vigintile.

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

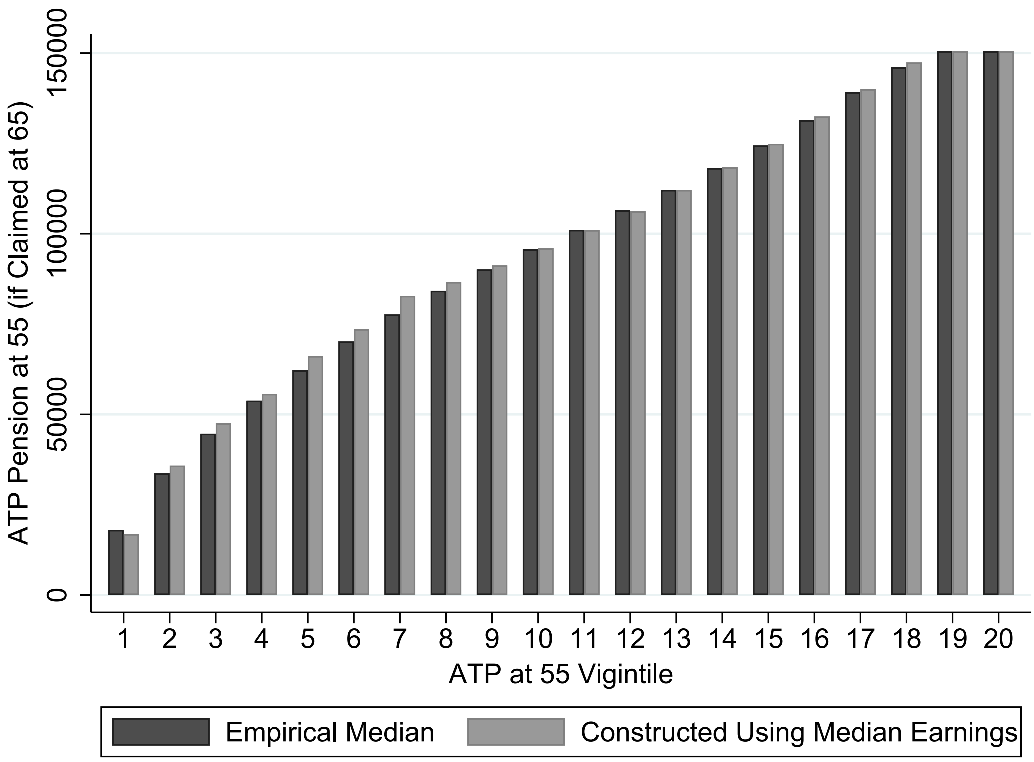

Notes: This figure shows the median ATP pension that workers born between 1938 and 1943 were eligible for at age 55 (assuming the pension is claimed at age 65) plotted alongside the simulated ATP pension. These are shown for each ATP at 55 vigintile.

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)

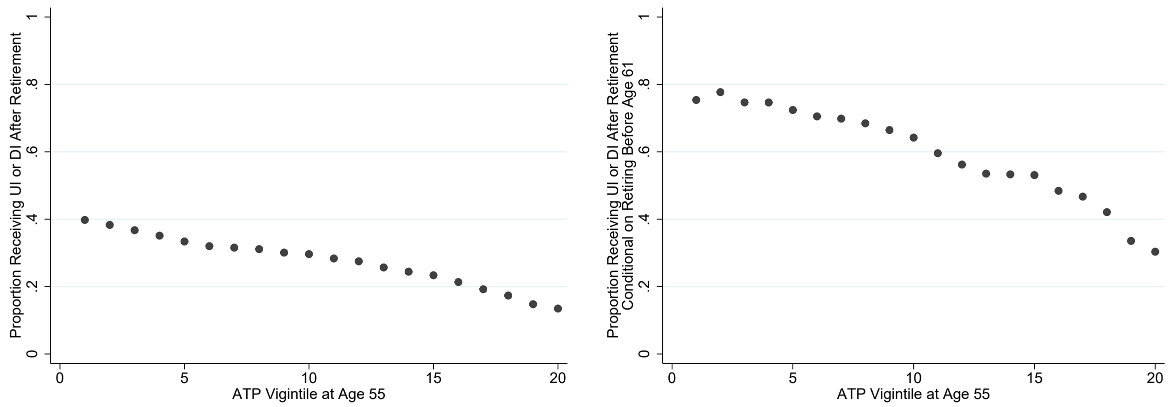

Notes: These graphs show the proportion of individuals in each ATP at 55 vigintile who receive UI or DI after retiring. Panels B restricts to individuals retiring before age 61.

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, 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

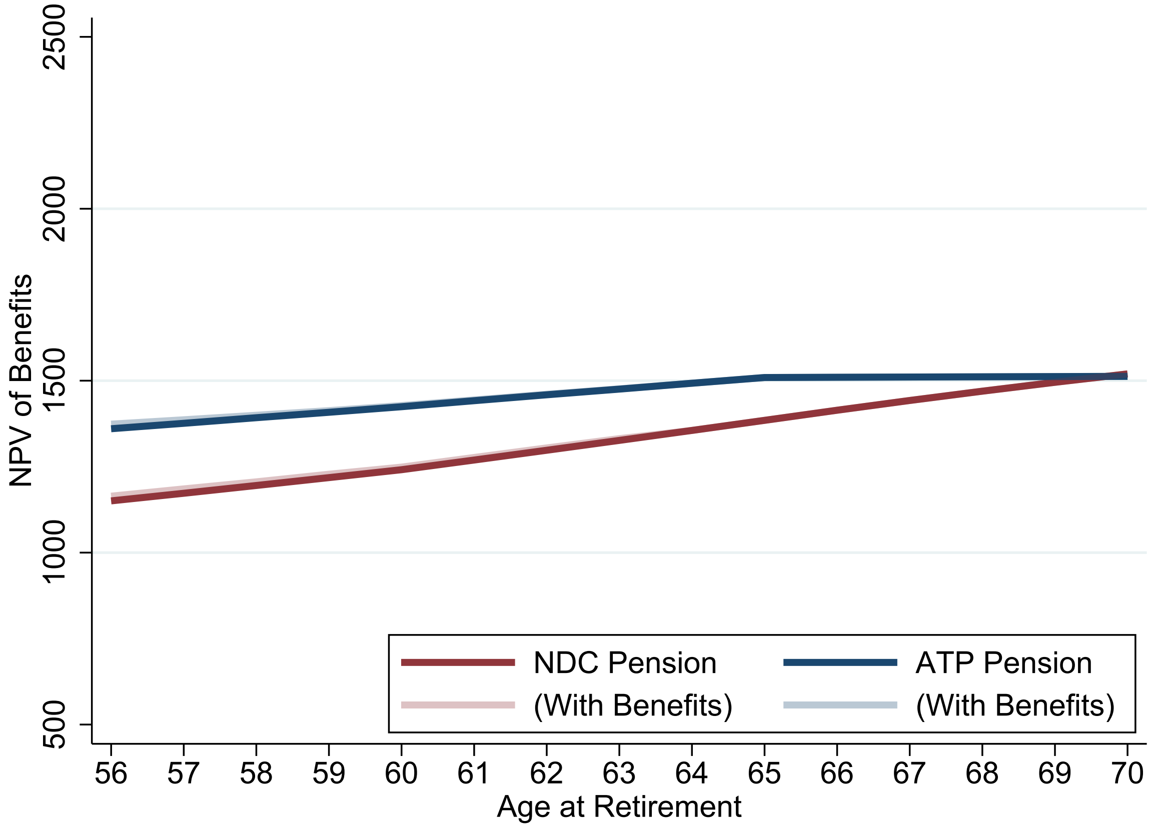

Notes: This graph shows the mean net present value (NPV) of pension benefits across the 20 ATP at 55 vigintiles for each retirement age. The opaque lines show this for an individual who does not receive post-retirement UI or DI benefits. The transparent lines show the weighted mean of the NPVs without 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

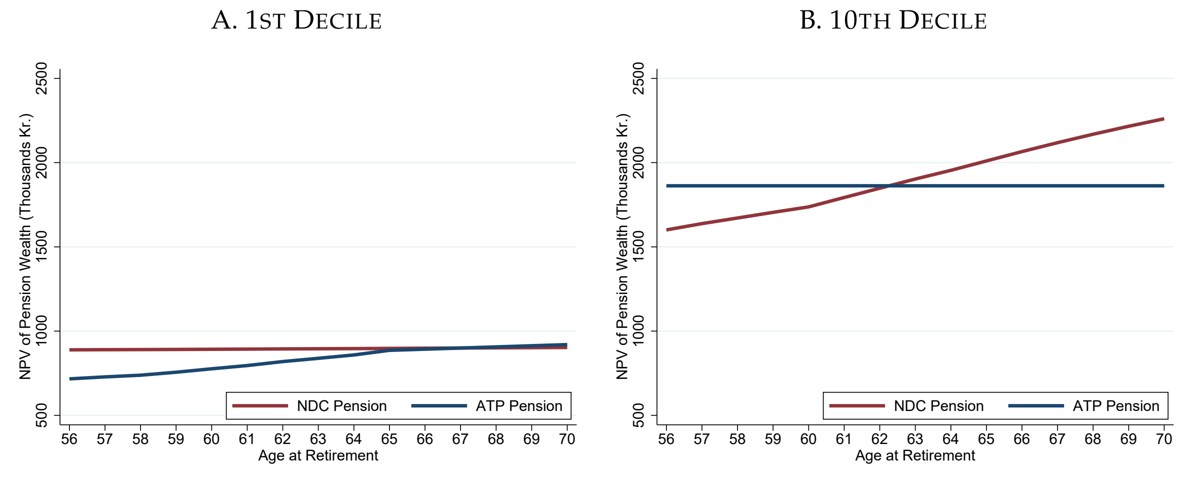

Notes: These graphs show the net present values (NPVs) of pension wealth by age at retirement for the top and bottom deciles of the distribution of ATP at age 55. The graph for each decile is created using the average NPVs of both vigintiles within that decile. Calculations are for individuals born in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

We next turn to the participation and implicit tax rates, which, as discussed in Section 7, 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 7, 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

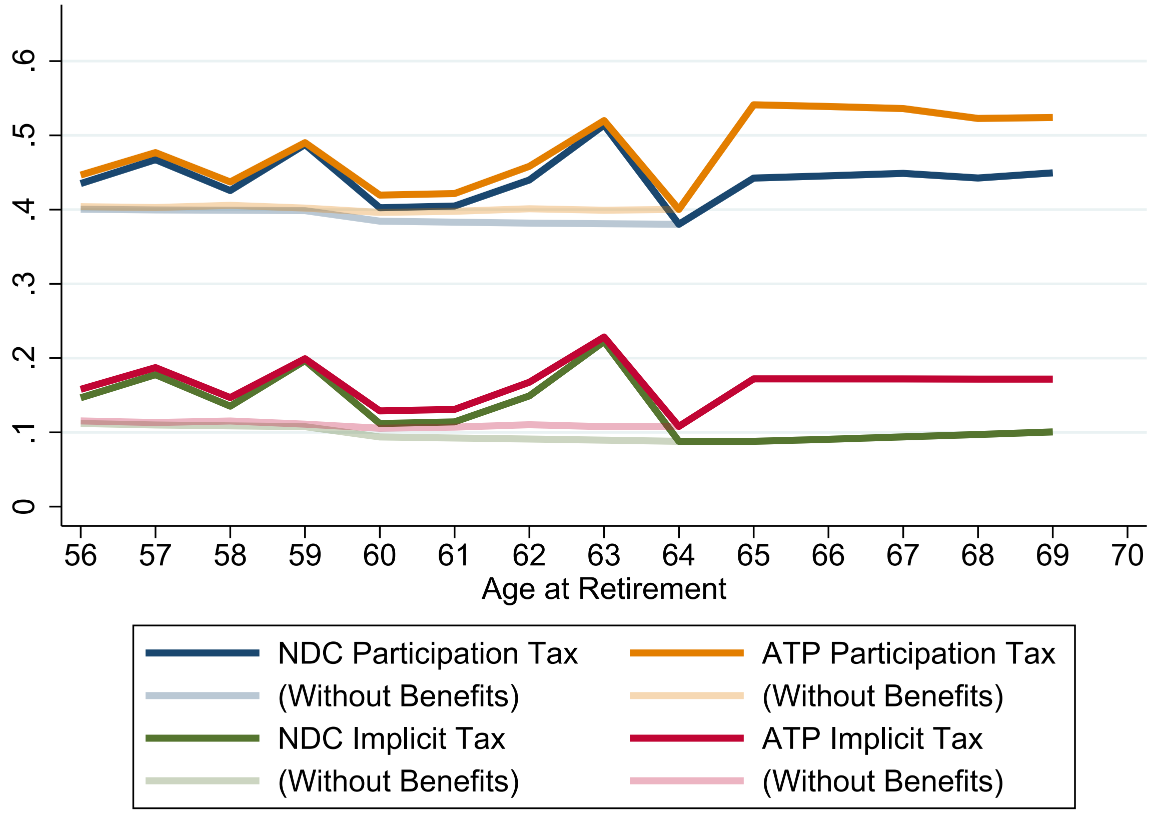

Notes: This graph shows the mean implicit and participation tax rates of pension benefits across the 20 ATP at 55 vigintiles for each retirement age. The opaque lines show this for an individual who does not receive post-retirement UI or DI benefits. The transparent lines show the participation and implicit 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 3.1), 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

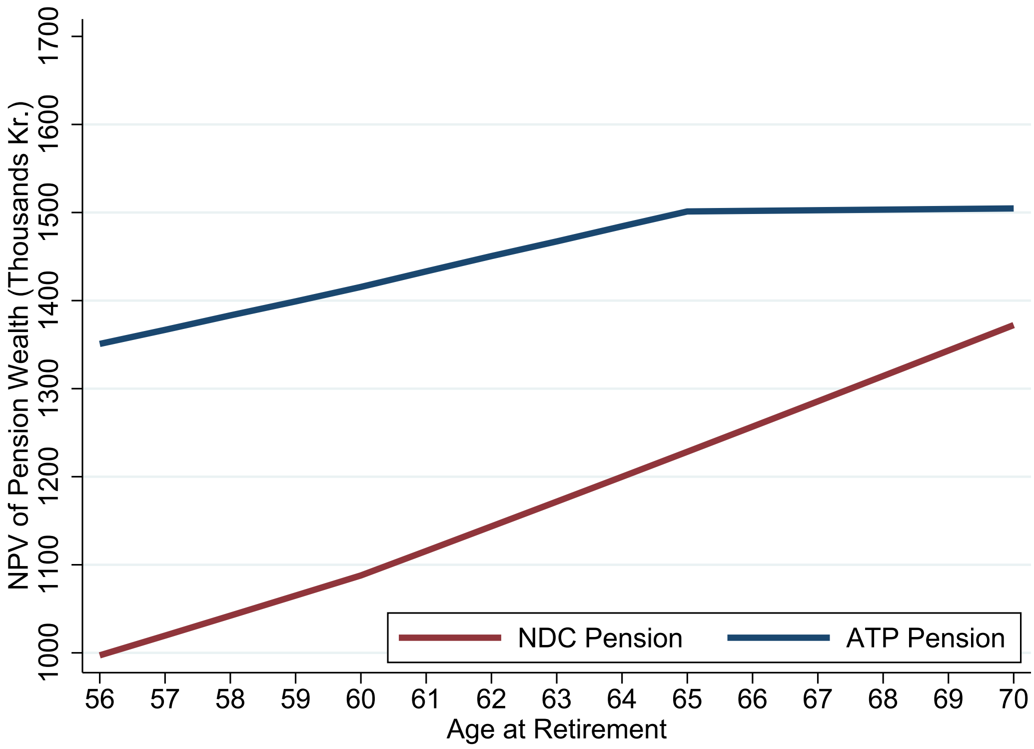

Notes: This graph shows the net present value of pension wealth by age at retirement averaged across all ATP at 55 vigintiles. Calculations are for individuals born in 1941 with a discount factor of 0.98.

Figure A-11: AVERAGE IMPLICIT TAX RATE - CLAIM AT AGE 70

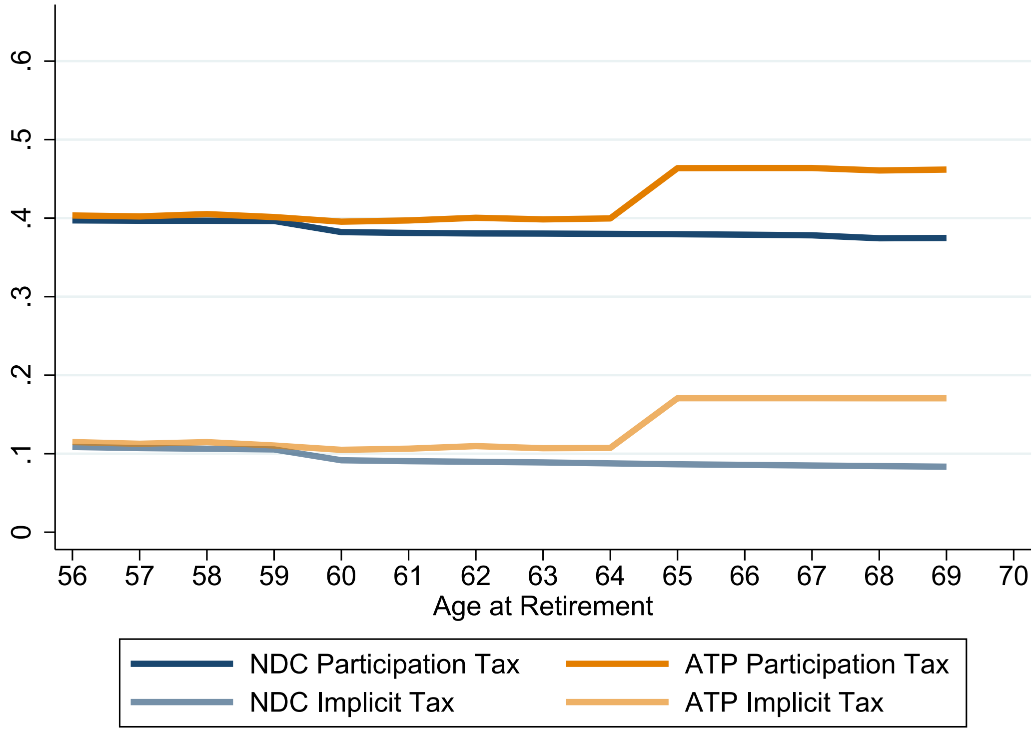

Notes: This figure shows the average participation tax rate and implicit tax rate across all 20 ATP at 55 vigintiles by age at retirement .

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:

$$\sum_{r=56}^{69} [ATP_r f(r)] = \sum_{r=56}^{69} [\widehat{NDC}_r f(r)] \tag{25}$$

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

$$\widehat{NDC}_r = \Delta + NDC_r. \tag{26}$$

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

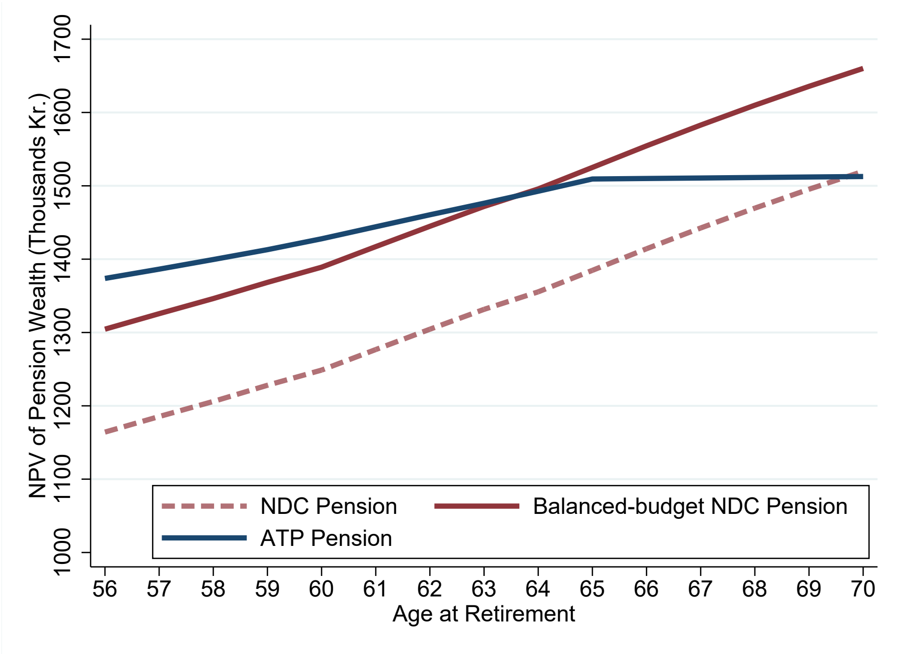

Notes: This figure shows the net present value (NPV) of pension wealth by age at retirement averaged across all ATP at 55 vigintiles. Calculations are for individuals born in 1941 with a discount factor of 0.98. NPVs are shown for both the actual NDC pension and the balanced-budget version of the NDC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

Plugging this into equation (25) and solving for Δ we obtain:

$$\Delta = \frac{\sum_{r=56}^{69}[ATP_r f(r)] - \sum_{r=56}^{69}[NDC_r f(r)]}{\sum_{r=56}^{69}[f(r)]}. \tag{27}$$

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.[59] 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:

$$Consumption = Income - Saving. \tag{28}$$

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 et al. [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

$$C_{it} = Z_{it} - \sum_k p_{kt} \left[A_{ikt} - A_{ikt-1} \right], \tag{29}$$

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.

[59] 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 et al. [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

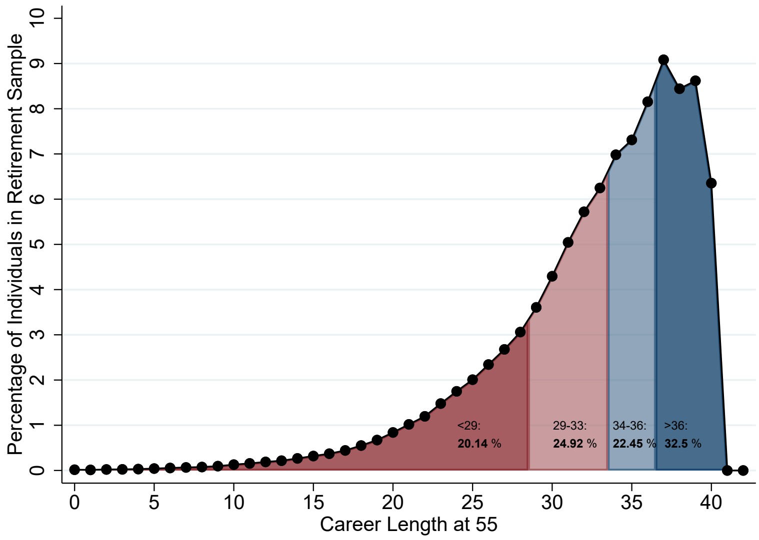

B. Distribution of Retirement Age by Career Length at 55 Groups

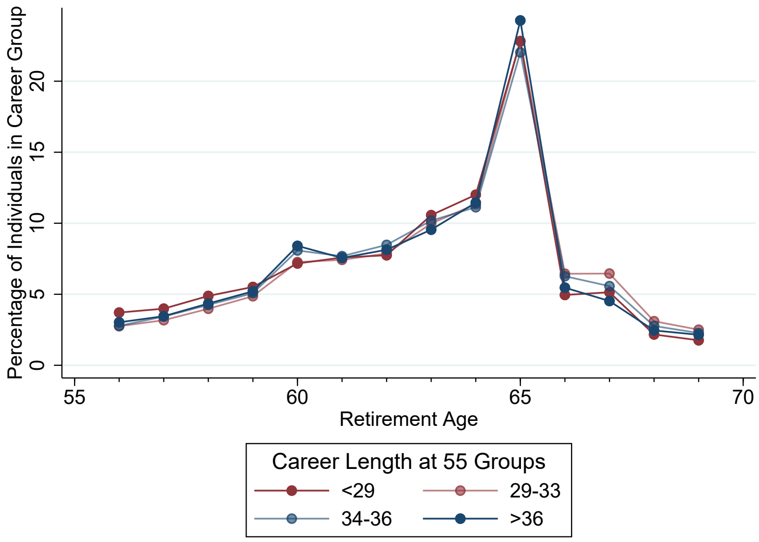

Notes: The figure provides information about the distribution of career length at age 55 in our sample. Career length at 55 is defined as the sum of years during which an individual made positive contributions to the ATP pension system before turning 55. We observe actual contributions before 55 years 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

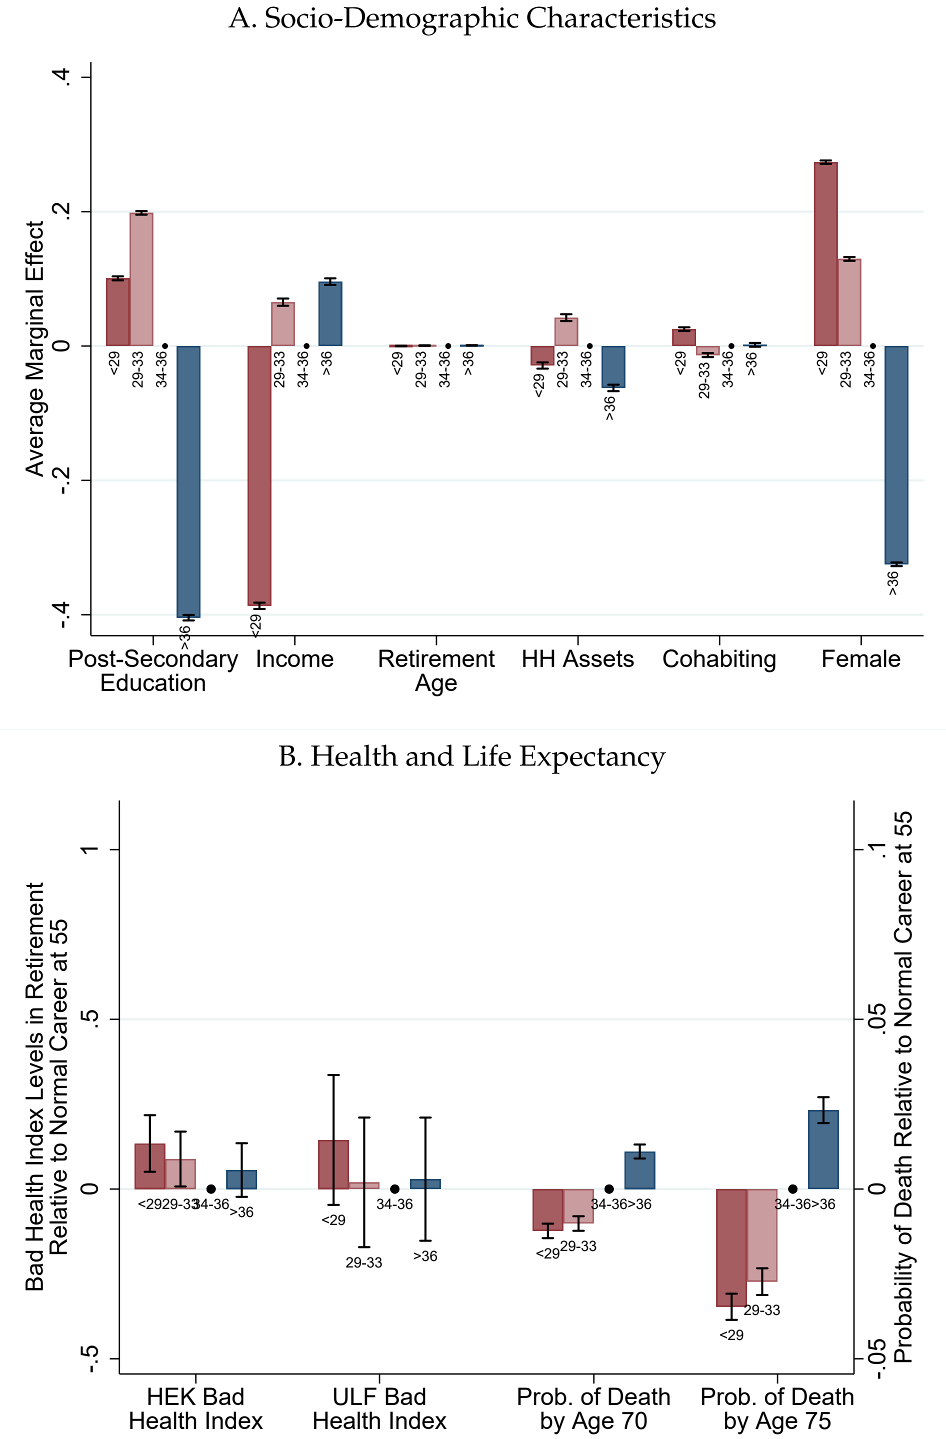

Notes: The figure documents patterns of heterogeneity across groups based on career length at age 55. The construction of the figure is exactly analogous to Figure 5. Panel A displays estimates from a multinomial logit prediction model for belonging to one of the 4 different career length groups. The 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 (12) (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.

Appendix C Consumption Levels & Heterogeneity

Consumption Differences By Retirement Age: Robustness

Figure C-1: CONSUMPTION DIFFERENCES BY RETIREMENT AGE

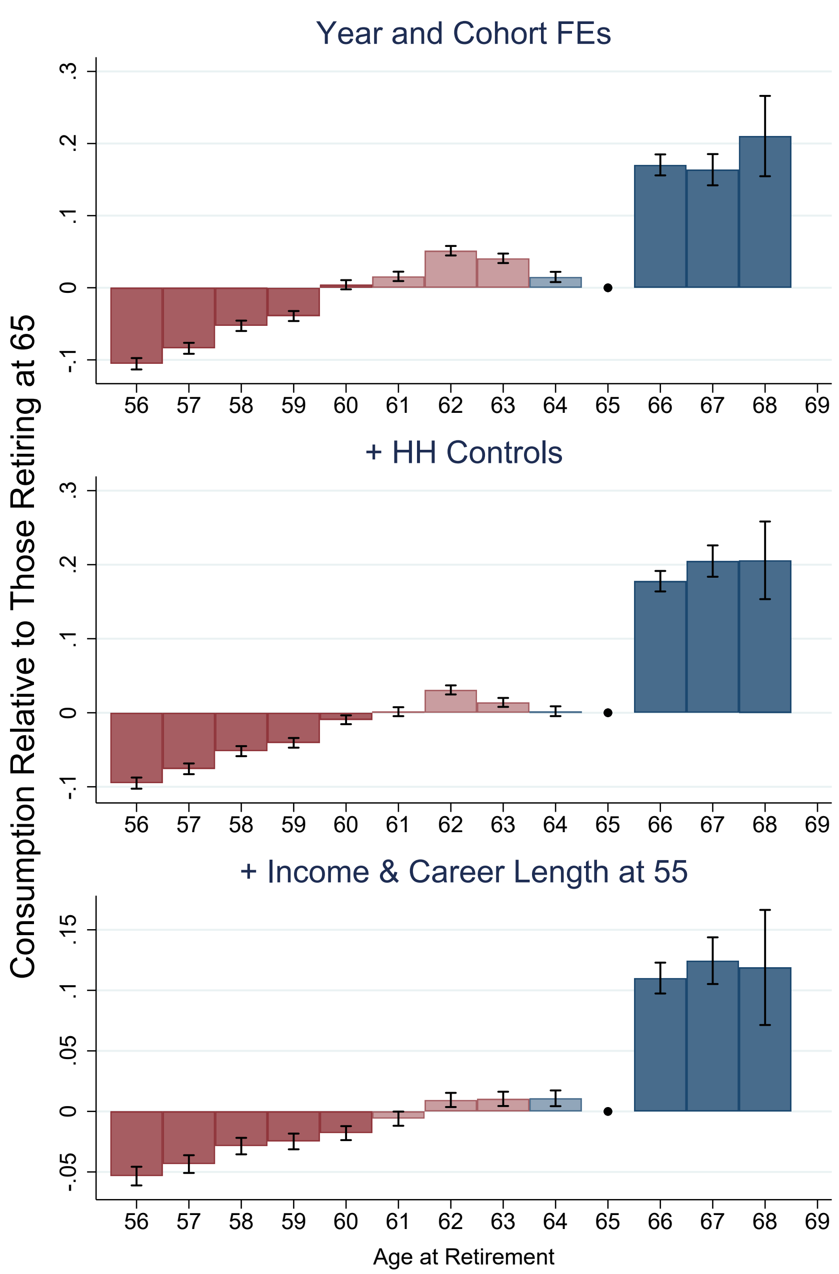

Notes: The figure report estimates of a fully non-parametric version of specification (12) where we compare consumption levels across all retirement ages (rather than aggregating retirement ages into four groups). The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time 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 (12), scaled by $\mathbf{E}_j[\bar{C}_{it}]$, 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 (12) 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

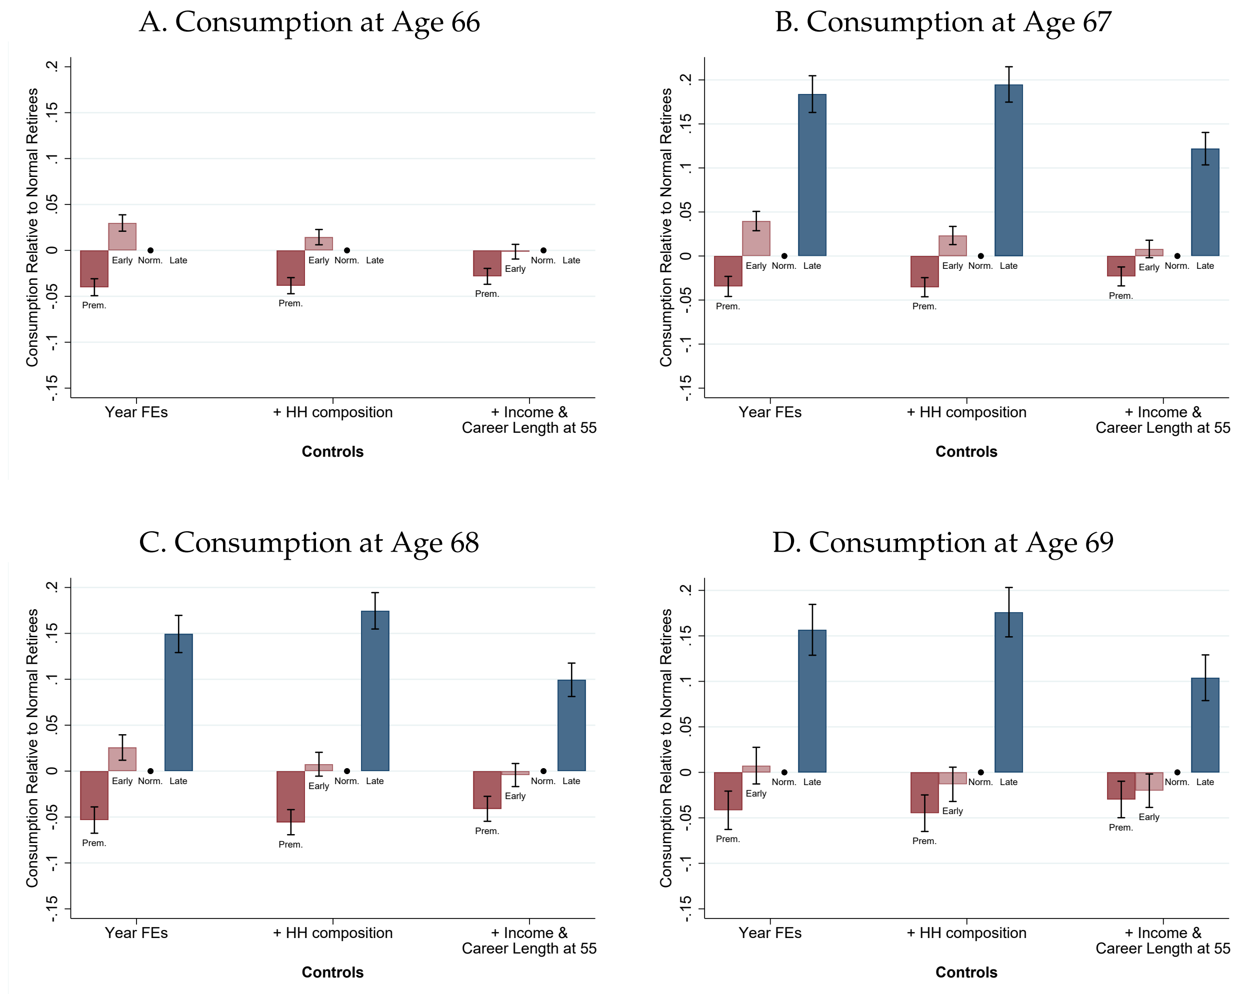

Notes: The figure shows that the consumption patterns hold irrespective of the age at which consumption is observed during retirement. We run regressions similar to specification (12), but separately for each age t. Because t is now fixed, we remove age fixed effects from the specification and control 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)

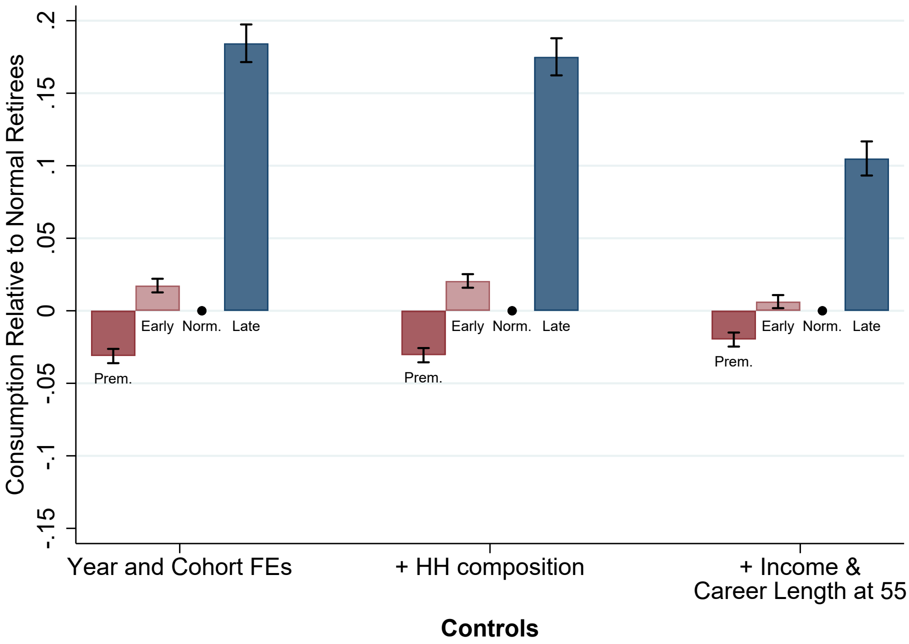

B. Singles

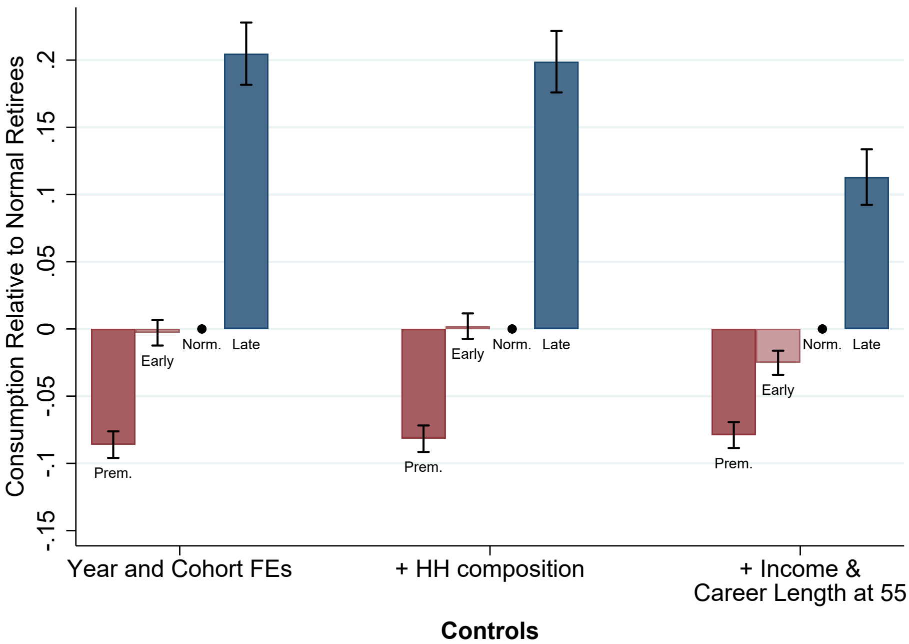

Notes: The figure reproduces estimates of consumption differences in retirement by retirement age group, similar to Figure 3 but splitting the sample between individuals who are single vs married/cohabiting at the time of retirement. The sample comprises all individuals from cohorts 1938 to 1943 who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

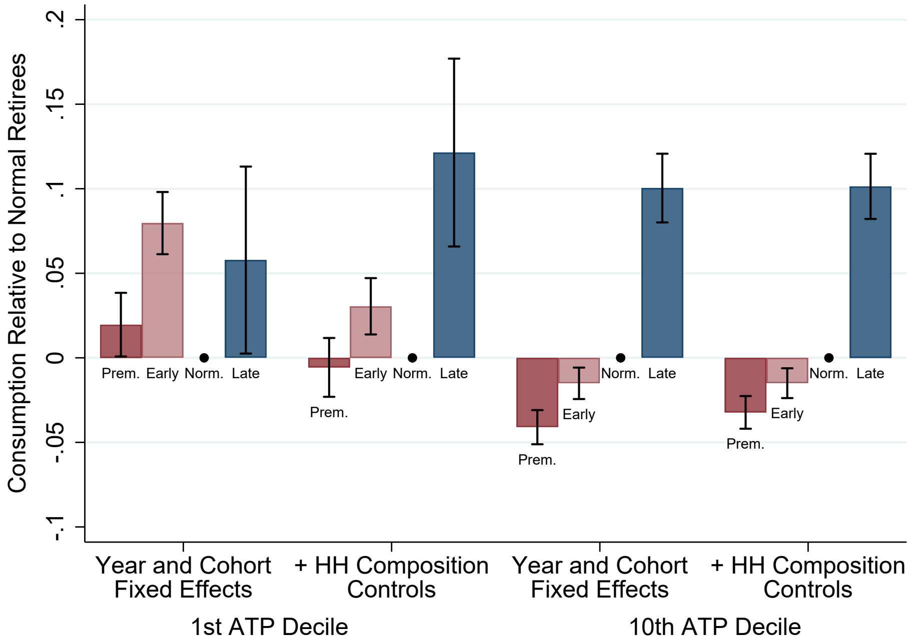

Notes: The figure reports consumption in retirement across individuals who retire at different ages relative to normal retirees. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed. Individuals are grouped into four retirement age categories: 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

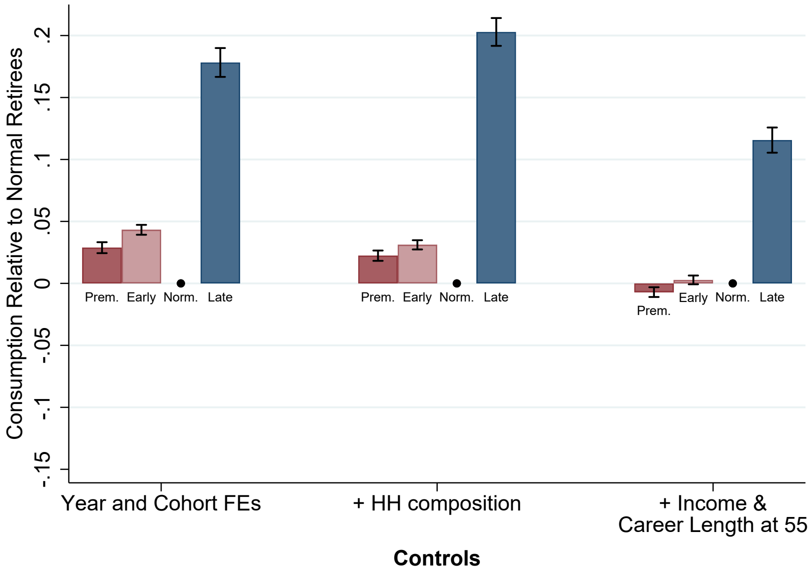

Notes: The figure documents consumption differences across retirement age groups using an alternative measure of retirement age that accounts for the time spent in UI or DI after an individual stops working. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their 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 (12), 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 (12) 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 $\mathbf{X}$.

Figure C-6: CONSUMPTION DIFFERENCES BY RETIREMENT AGE: CONTROLLING FOR PRE-RETIREMENT CONSUMPTION

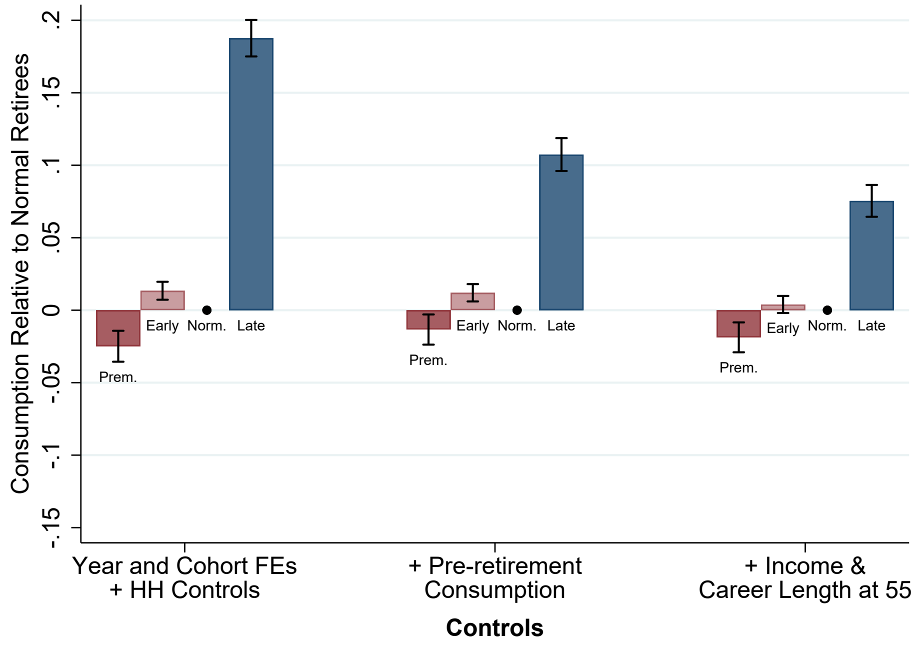

Notes: The figure depicts estimates of consumption differences in retirement by retirement age group, similar to Figure 3, but adding non-parametric controls for consumption before retirement. The sample comprises all individuals from cohorts 1938 to 1943 who are retired at the time their consumption 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 BY RETIREMENT AGE: OB DECOMPOSITION

	Premature	Early	Normal	Late
	A. Mean residual consumption (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	.18	.19	.23	-
	(.01)	(.01)	(.01)	
Career Length at 55	.06	.03	.03	-
	(.01)	(.00)	(.00)	
Income at 52-55	.19	.22	.24	-
	(.01)	(.01)	(.01)	
Implied fraction explained by late career dynamics	.56	.57	.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 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 $X_{it}^k = \begin{pmatrix} d_{it}^{k,1} & d_{it}^{k,2} & \dots \end{pmatrix}'$, 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

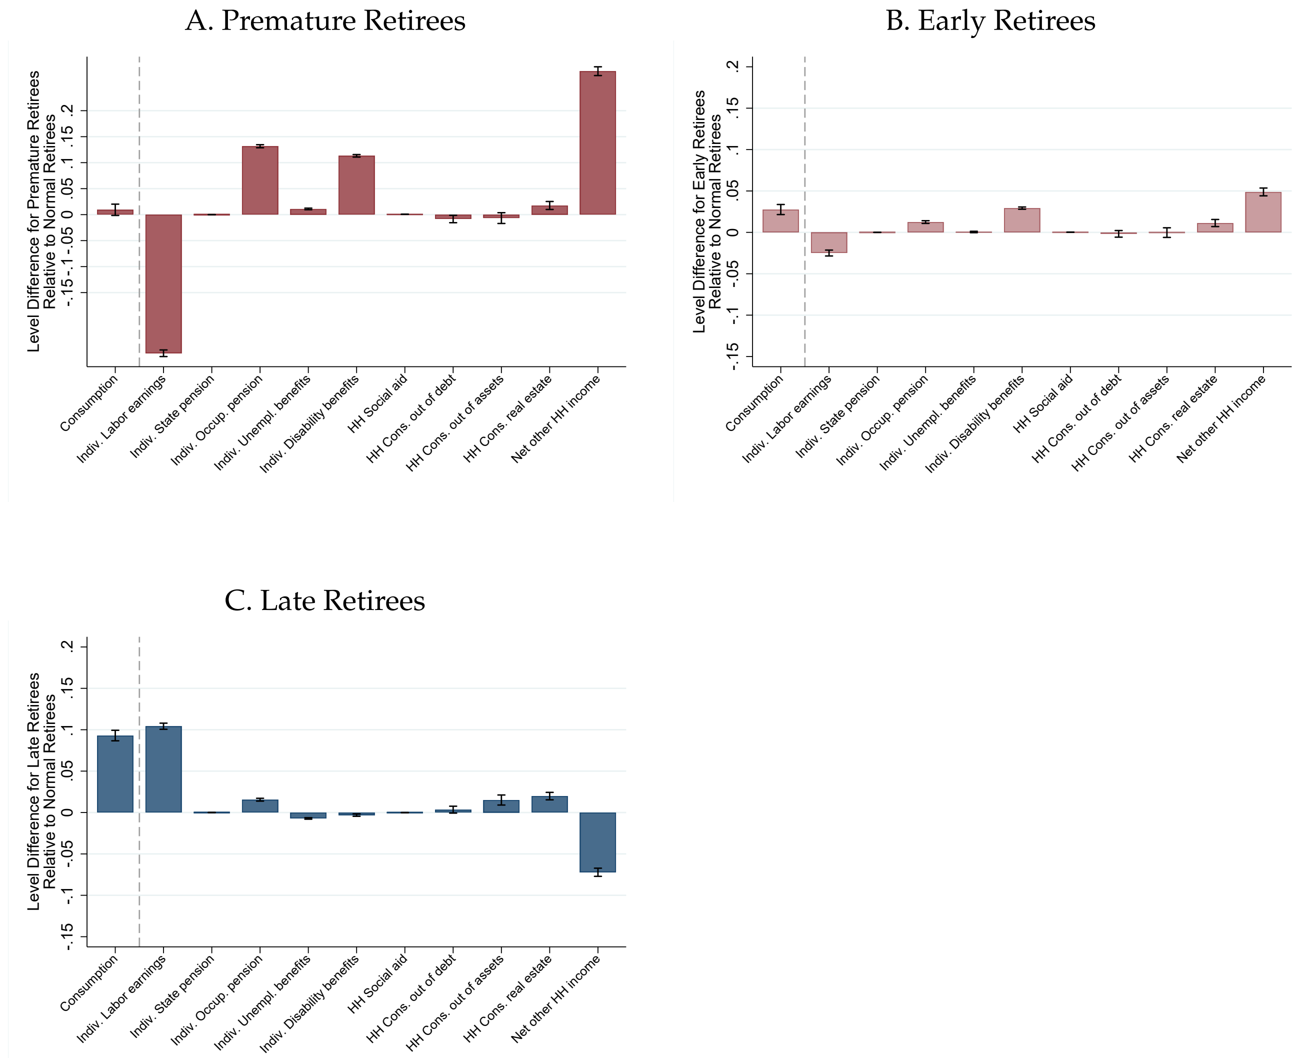

Notes: The figure decomposes consumption differences at age 60 across individuals who retire at different ages. The sample comprises all individuals from cohorts 1938 to 1943. Individuals are grouped into four retirement age categories: premature retirees ($56 \leq r \leq 60$), early retirees ($61 \leq r \leq 63$), 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 (12) 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: CHANGE IN CONSUMPTION SHARES AROUND RETIREMENT

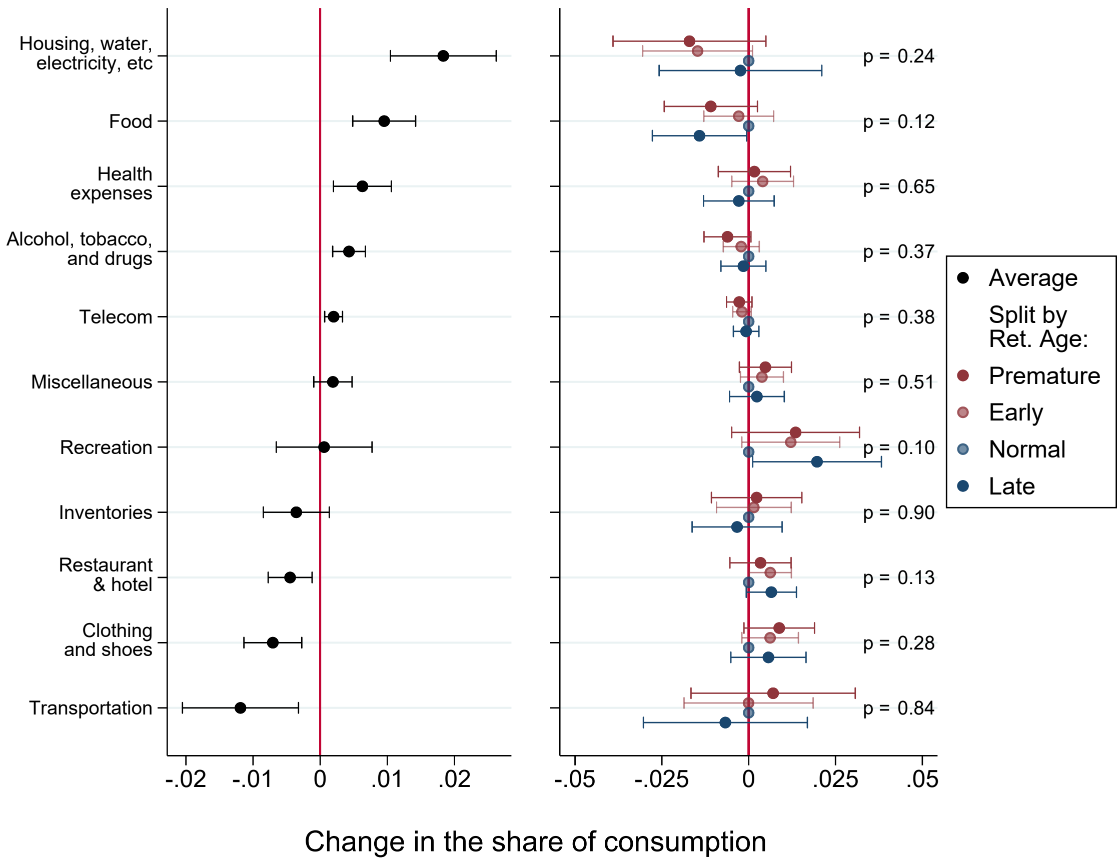

Notes: This figure examines changes in consumption patterns around retirement. Using 5,205 observations from the HUT survey, total consumption is divided into 11 spending categories. The share of total consumption represented by each category is then regressed on a set of year and cohort fixed effects, 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:

$$C_{it}^{l} = \sum_{j} \alpha_j \cdot \mathbb{1}[\mathrm{r = j}] + \gamma_y + \gamma_t + \gamma_l \tag{30}$$

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 $\mathbf{E}_j[\tilde{C}^l_{it}]$, the predicted consumption level from specification (31) when omitting the contribution of the retirement age group dummies. Specifically, $\mathbf{E}_j[\tilde{C}^l_{it}]$ 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

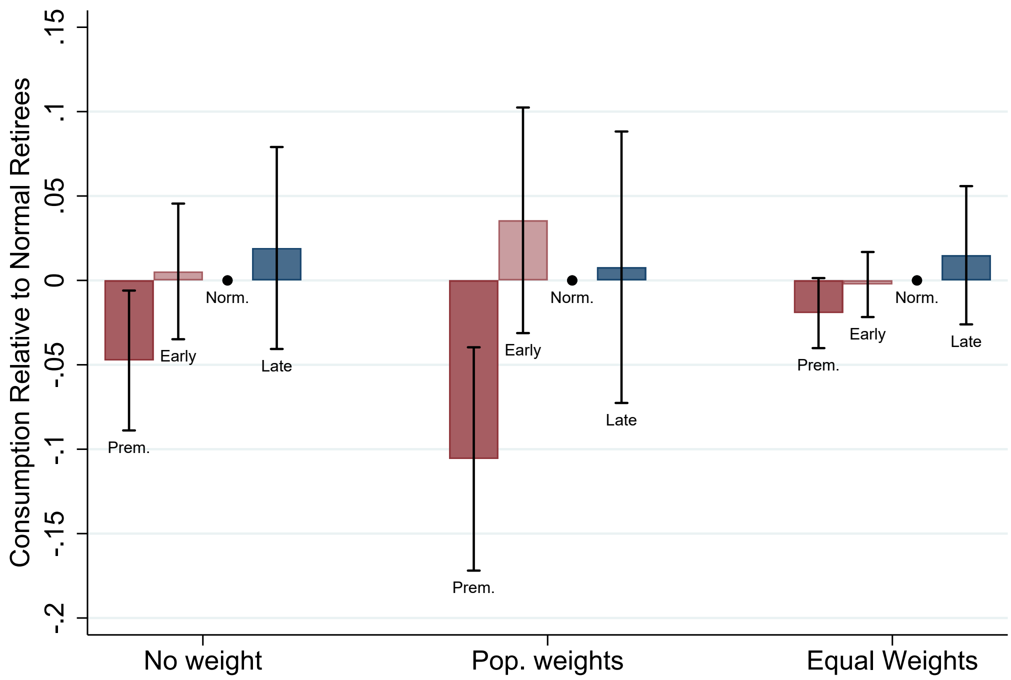

Notes: The figure documents differences in food consumption across retirement age groups. The sample comprises all individuals aged 50 to 75 from cohorts 1938 to 1958 who are observed in the SHARE data from Austria, Belgium, Denmark, France, Germany, Greece, Italy, the Netherlands, Spain, Sweden, Switzerland 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 (30), 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), vehicule 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

$$C_{it} = \sum_j \alpha_j \cdot \mathbb{1}[\mathrm{r} = \mathrm{j}] + \gamma_y + \gamma_t \tag{31}$$

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*. 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 (31) when omitting the contribution of the retirement age group dummies. As before, $\mathbf{E}_j[\tilde{C}^l_{it}]$ 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 not 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 quite 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

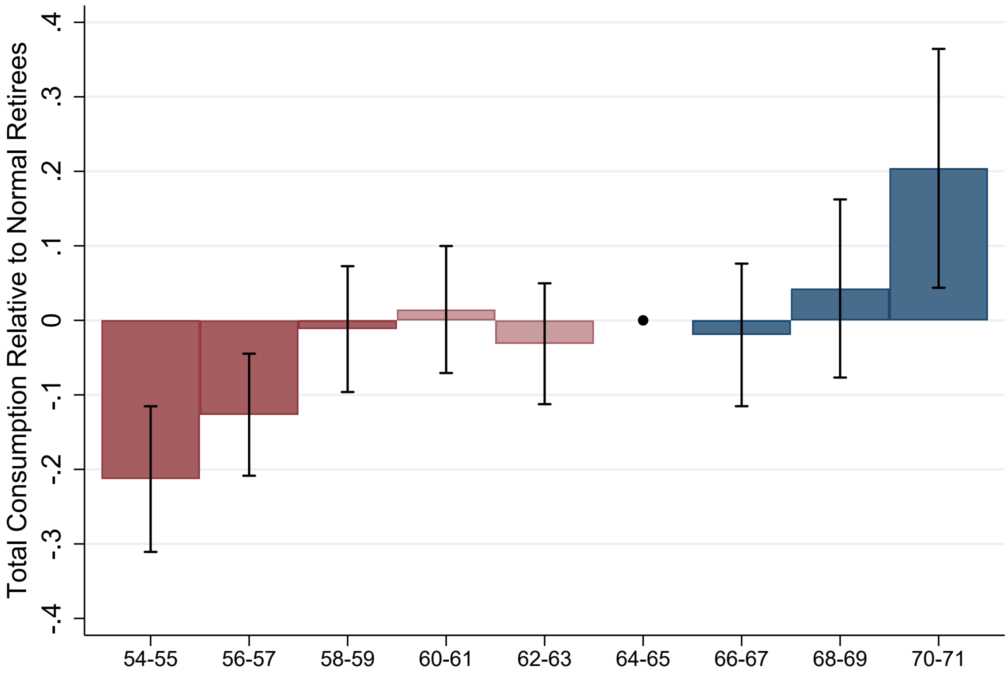

Notes: This figure documents how consumption differs across individuals who retire at different ages. The sample is composed of all individuals born between 1938 and 1958 interviewed for the consumption module (CAMS) of the Health and Retirement Study (HRS). The CAMS modules happen every two years since 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 (31), 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).[60] 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).[61]

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:

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

[61] 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)		

Continued on next page

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)
Free Outpatient Care			23.27%	(42.26)		
Healthcare Expenditure			366.40	(553.50)		
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 (12), 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 (12) 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.[62]

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

[62]These results rely on a specification similar to (12), 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.

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 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 7, and compare health dynamics across retirement age groups. This allows us to check whether the differences in health outcomes when retired documented in Figure 5 pre-date retirement, and whether they are caused by early retirement (Kuhn et al. [2018], Fitzpatrick and Moore [2018], Bozio et al. [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:

$$H_{it} = \sum_j \sum_k \alpha_{jk} \cdot \mathbb{1}[\mathrm{r} = \mathrm{j}] \cdot \mathbb{1}[\mathrm{e} = \mathrm{k}] + \gamma_y + \gamma_t + \mathbf{X}'\beta + \nu_{it}. \tag{32}$$

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 (32) where we use our bad health indices as an outcome, pooling both HEK and ULF surveys together.[63] 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

[63] In practice, this means that we run specification (32) 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 index) as their pre-retirement level. Interestingly, there is no significant variation in 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

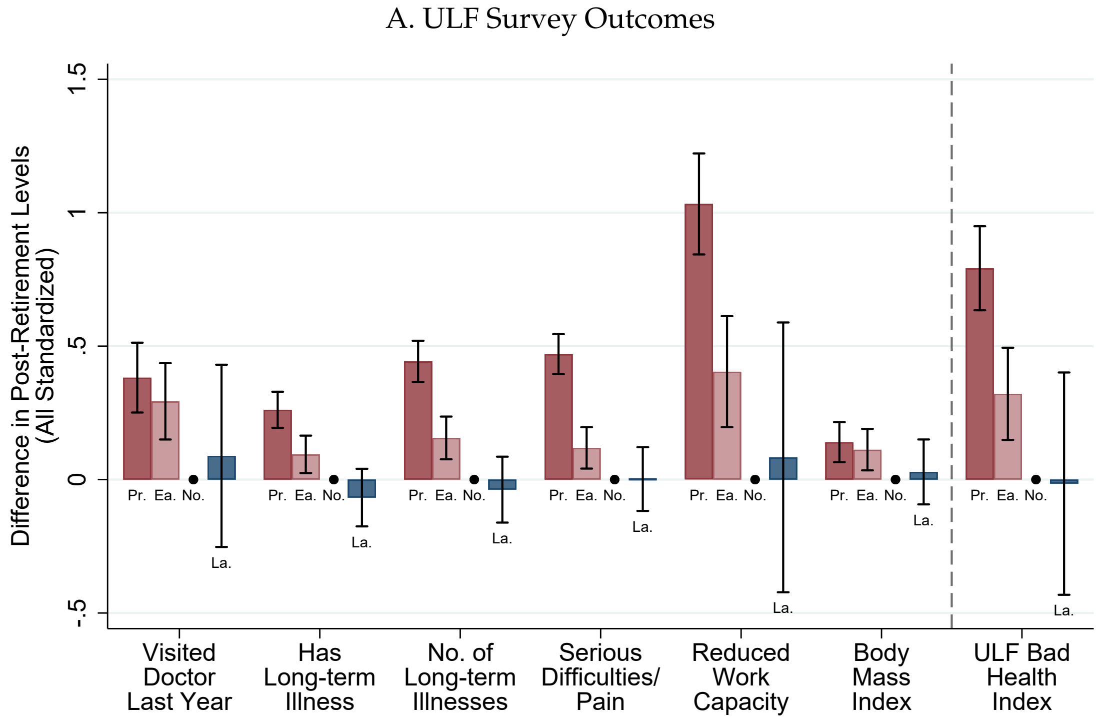

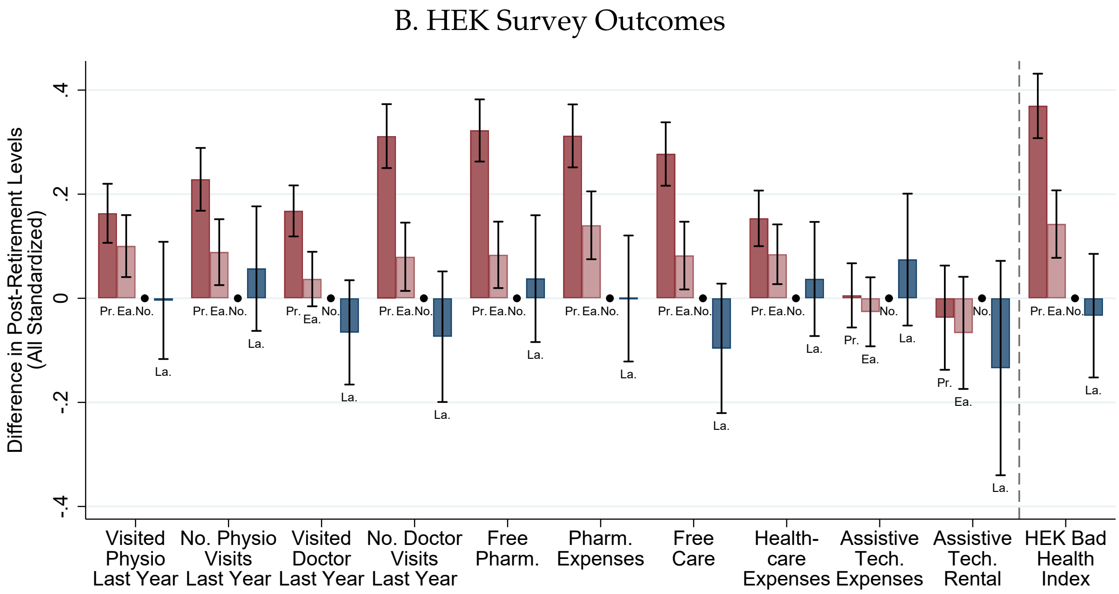

Notes: The figure documents differences in health outcomes across retirement age groups. The sample comprises all individuals from cohorts 1938 to 1943 who are observed either in the ULF or HEK surveys, and who are retired at the time of the survey. Individuals are grouped into four retirement age categories 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 (12), where we control for age and cohort fixed effects, as well as controls for family composition in the vector of controls $\mathbf{X}$. All outcomes are standardized.

Figure E-2: Health Dynamics Around Retirement, By Retirement Age Group

A. HEK & ULF Bad Health Index - Combined

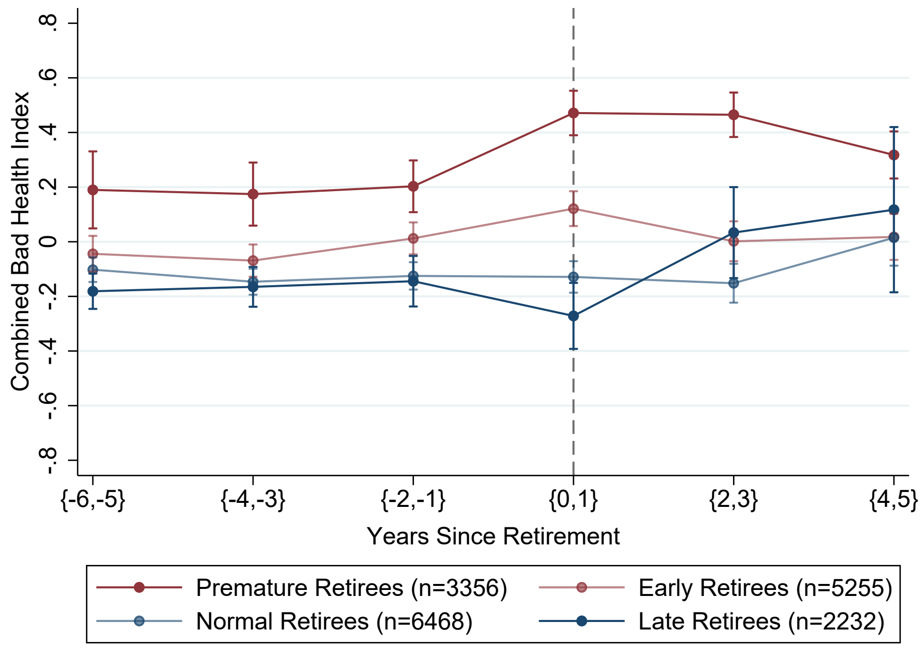

B. Fraction Experiencing Pain

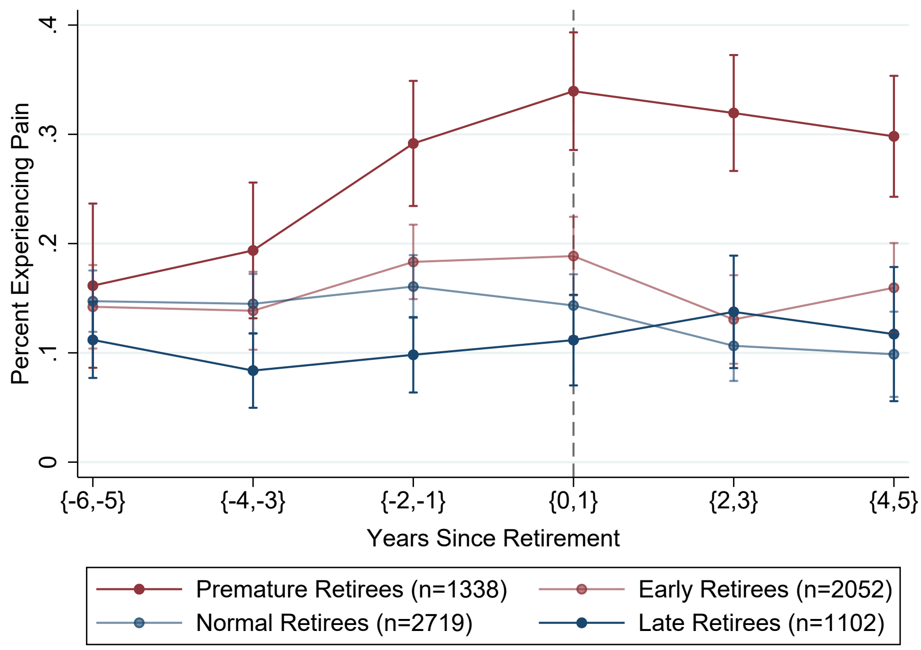

Notes: The figure documents heterogeneity in health dynamics around retirement, by retirement age group. Both panels report, for each retirement age group, the sequence of estimated coefficients $\hat{\alpha}_{re}$ from specification (32). where we control for cohort and age fixed effects and on the usual vector $\mathbf{X}$ 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

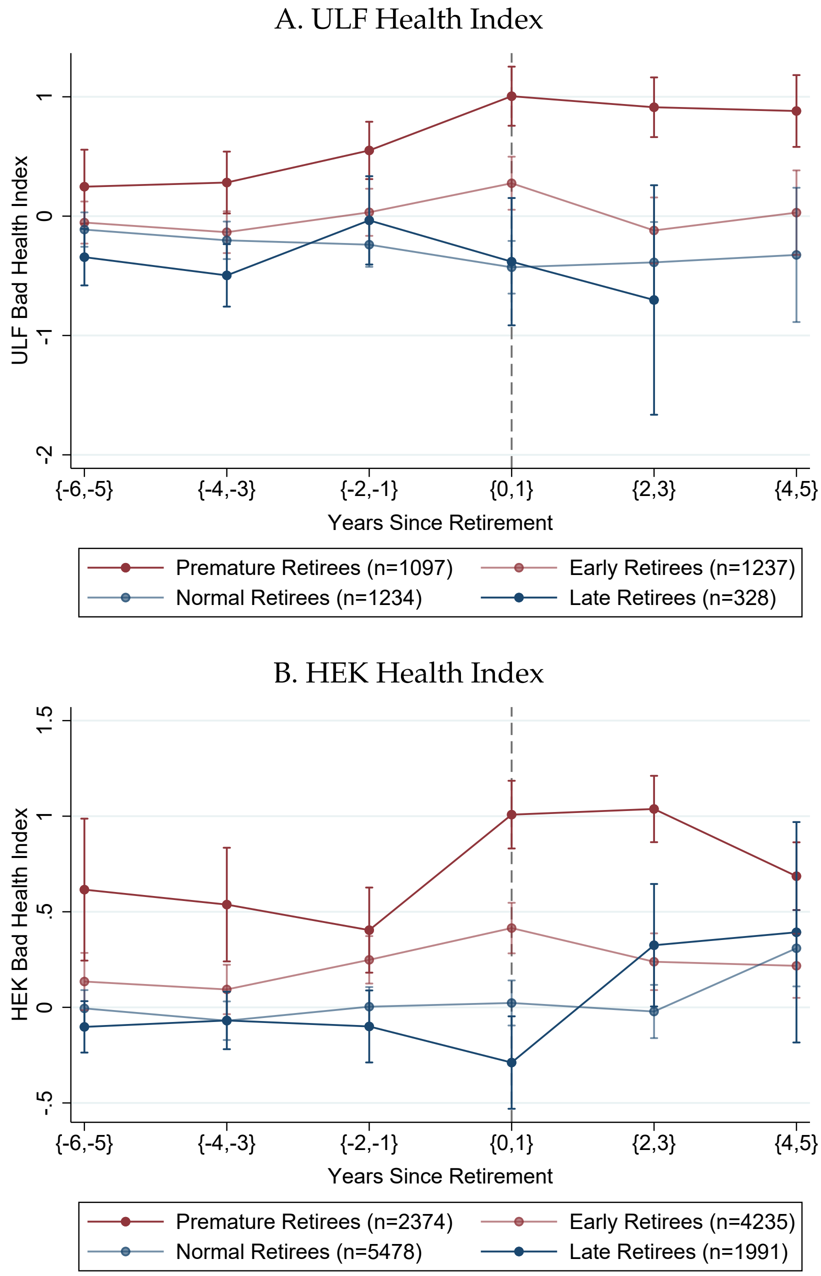

Notes: The figure documents heterogeneity in health dynamics around retirement, by retirement age group. Both panels report, for each retirement age group, the sequence of estimated coefficients $\hat{\alpha}_{re}$ from specification (32). where we control for cohort and age fixed effects and on the usual vector $\mathbf{X}$ 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

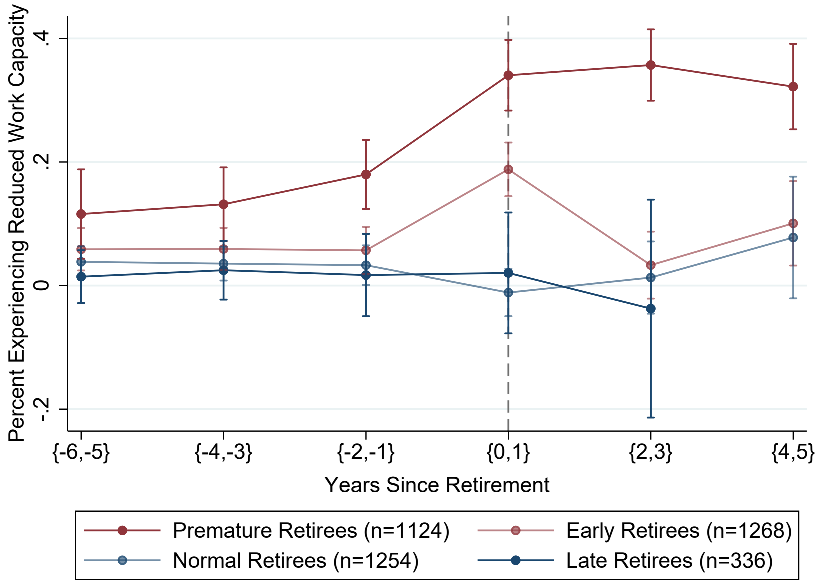

B. Retirement For Health Reasons

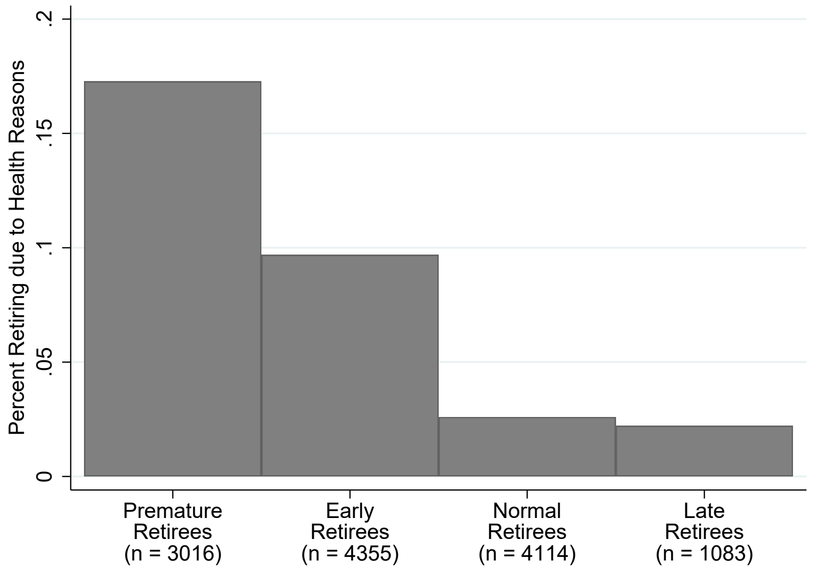

Notes: The figure documents heterogeneity in health dynamics around retirement, by retirement age group. Panel A reports, for each retirement age group, the sequence of estimated coefficients $\hat{\alpha}_{re}$ from specification (32) similar to Figure E-2 where we use the fraction reporting reduced 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

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:

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

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 et al. [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:

$$\mathrm{KG}_{i,t+k} = \alpha_k \mathrm{KG}_{i,t+1} + \mathbf{X}'\beta \tag{33}$$

where $\mathrm{KG}_{i,t+k} = \sum_j \Delta p_{j,t+k} \cdot a_{ijt}$ 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.[64]

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

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.

and 50-tiles of portfolio value interacted with vigintiles of average variance in the past six years.

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

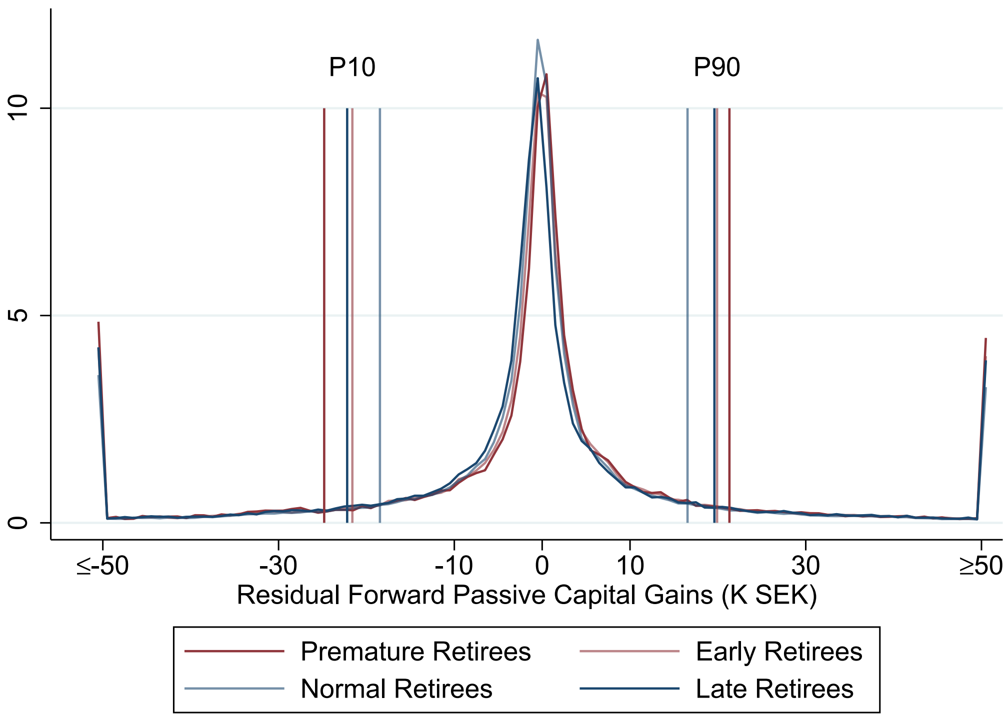

Notes: The figure plots the distribution of residualized passive capital gains. The sample is the baseline retirement sample merged with the KURU register, which has disaggregated information over the period 1999 to 2007 on all quantities of stocks, by ISIN number, held by individuals outside of mutual 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 $\text{KG}_{i,t+1} = \sum_j (p_{j,t+1} - p_{j,t}) \cdot a_{ijt} = \sum_j \Delta p_{j,t+1} \cdot a_{ijt}$ 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. 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 $\text{KG}_{i,t+1}$, 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

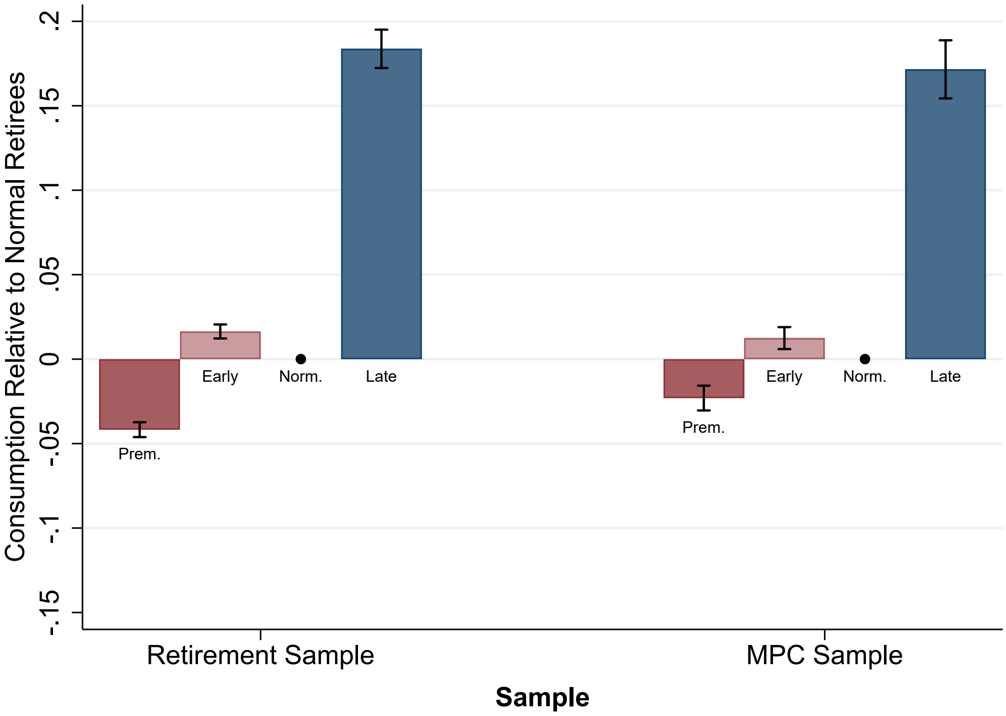

Notes: The figure replicates in the MPC sample our baseline estimates of consumption differences in retirement from Figure 3. Both samples comprise individuals from cohorts 1938 to 1943 who are retired at the time their consumption is observed. Individuals are grouped into four retirement age categories: 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 (12), 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 (12) 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

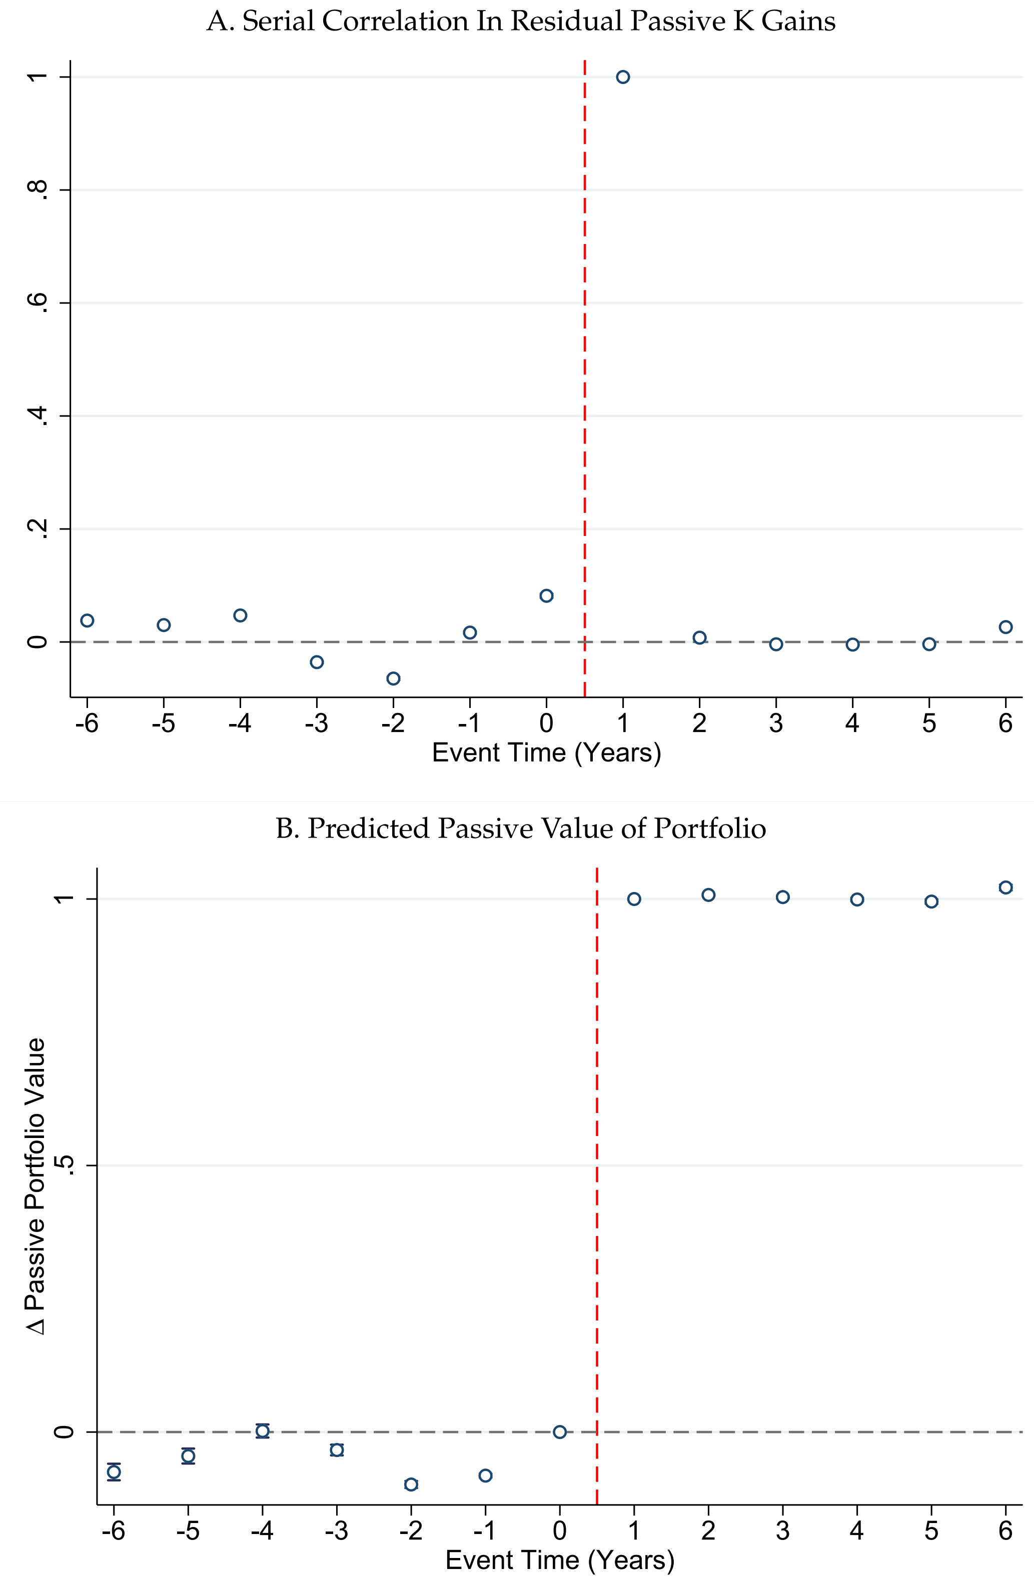

Notes: Panel A plots for each time horizon $k \in \{-6, ..., 6\}$ the serial correlation of the residual passive capital gain at k and the current residual passive capital gain, that is the coefficient α_k from regression (16). We control for the value of portfolio in year t, the average returns and 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.

Table F-2: 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	.34	.17	.49
	(.01)	(.01)	(.04)
Premature Retirees	.29	.37	1.26
	(.02)	(.04)	(.15)
Early Retirees	.32	.26	.81
	(.01)	(.03)	(.08)
Normal Retirees	.38	.07	.20
	(.01)	(.02)	(.06)
Late Retirees	.36	.05	.14
	(.02)	(.05)	(.13)

Notes: This table shows the estimates of the 2SLS approach presented in equation (18). 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-3: 2SLS Estimates of MPCs: Robustness to Alternative Clustering

	First Stage α_1^V	Reduced Form 3 x α_{rf}^C	IV Result 3 x α_{IV}^C
Cluster by 50-tile of PF Value PF x 20-tile of Average PF Past Returns			
Baseline	.66	.11	.17
	(.04)	(.01)	(.01)
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

$$\mathcal{U}_i(c, \zeta, \pi) = \sum_{t=0}^{T} \beta^t \int u\left(c\left(\pi_{i,t}\right), \zeta\left(\pi_{i,t}\right)\right) dF\left(\pi_{i,t}\right), \tag{34}$$

where $c(\pi_{i,t})$ is the individual's consumption choice and $\zeta(\pi_{i,t})$ 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 et al. [2016]) can be captured through ζ and how it affects the utility of consumption c. Like in all structural models of retirement, $\zeta(\pi_{i,t})$ includes the extensive labor supply choice, which is denoted by $s(\pi_{i,t})$ 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 $r(\pi_{i,t})$. We thus have $s(\pi_{i,t}) = 0$ for $t \geq r(\pi_{i,t})$ and $s(\pi_{i,t}) = 1$ otherwise. Hence, the number of individuals retiring at each r equals $S(r-1) - S(r)$, where $S(r) = \int \int s(\pi_{i,r})\, dF(\pi_{i,r})\, di$ is the survival rate into employment.

We note that $\zeta(\pi_{i,t})$ 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 $F(\pi_{i,t})$ (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 denote taxes by $\tau(\pi_{i,t})$ and pension benefits by $b(\pi_{i,t})$, which can depend in a flexible way on a worker's employment history, including the number of years worked and the corresponding earnings. We focus on workers' extensive labor supply and the age at which they retire. The government's objective is

$$\max\ \mathcal{W}(b,\tau) = \int_i \omega_i U_i(b,\tau) + \lambda GBC(b,\tau)\, di, \tag{35}$$

where

$$GBC(b,\tau) = \Sigma_t \frac{1}{R^t} \int\int [s(\pi_{i,t})\,\tau(\pi_{i,t}) - (1-s(\pi_{i,t}))\,b(\pi_{i,t})]\, f(\pi_{i,t})\, d\pi_{i,t} di - G_0. \tag{36}$$

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

$$GBC(b,\tau) = \Sigma_r\left[S(r)\,\frac{\tau_r}{R^r} - [S(r-1)-S(r)]\,NPV_r\right] - G_0, \tag{37}$$

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:

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

where $G(x,t) = \int\int 1[x(\pi_{i,t}) = x, s(\pi_{i,t}) = 0]dF(\pi_{i,t})\, di$. 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 $F(\pi_{i,t'+k})$ (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 et al. [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 $\beta R = 1$:

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

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 the Taylor approximation in equation (10) for $\frac{\partial u(c_{r,t},\zeta_{r,t})}{\partial c}$ around $\left(c_{r',t}, \zeta_{r,t}\right)$,

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

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

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

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 the Taylor approximation in equation (10) 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

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

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

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

$$\frac{\partial u\left(c_{i,t},\zeta_{i,t}\right)}{\partial c} + p_{i,t} \times \frac{\partial u\left(c_{i,t},\zeta_{i,t}\right)}{\partial \tilde{\zeta}_{i,t}} = 0. \tag{38}$$

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

$$\frac{\frac{dc_{i,t}}{dy_{i,t}}}{1-\frac{dc_{i,t}}{dy_{i,t}}} = p_{i,t} \times \frac{\frac{\partial^2 u(c_{i,t},\zeta_{i,t})}{\partial \tilde{\zeta}^2} / \frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial \tilde{\zeta}}}{\frac{\partial^2 u(c_{i,t},\zeta_{i,t})}{\partial c^2} / \frac{\partial u(c_{i,t},\zeta_{i,t})}{\partial c}} \cong p_{i,t}. \tag{39}$$

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 (38) and (39), we then obtain:

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

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., $\frac{\partial u(c_{r,t},\zeta_{r,t})}{\partial \tilde{\zeta}_{r,t}} \cong \frac{\partial u\left(c_{r',t},\zeta_{r',t}\right)}{\partial \tilde{\zeta}_{r',t}}$. 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 $d\left[b_{r>\tilde{r}} - b_{r\leq\tilde{r}}\right]$ 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:

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

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

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

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 ($r \leq \tilde{r}$) and given to late retirees ($r > \tilde{r}$) is equal to:

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

Table G-1: MEASURING THE SOCIAL MARGINAL VALUE OF STEEPENING THE PENSION PROFILE AT AGE $\tilde{r}$

Empirical Inputs	Economic Interpretation	Assumptions	Challenges
	Consumption-Level Implementation – Equation (11)		
$E_{r>\tilde{r}}(c), E_{r\leq\tilde{r}}(c)$: 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 $\omega_r \frac{\partial u(\bar{c},\zeta_{r,t})}{\partial c}$ constant across retirement ages r Taylor approximation (Chetty [2006]) Heterogeneity within retirement age group negligible (Andrews and Miller [2013])	Measuring γ Gauging selection into retirement ages based on *SMU* of consumption, driven by ω_r or $\zeta_{r,t}$
	Consumption-Drop Implementation – Equation (14)		
$\Delta c_{r>\tilde{r}}, \Delta c_{r\leq\tilde{r}}$: 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 $\omega_r \frac{\partial u(c_{r,pre},\zeta_{r,t})}{\partial c}$ constant across retirement ages r Taylor approximation (Chetty [2006]) Heterogeneity within retirement age group negligible (Andrews and Miller [2013])	Measuring γ Gauging selection into retirement ages based on *changes* in *SMU* of consumption around retirement, driven by $\frac{\zeta_{r,t}}{\zeta_{r,pre}}$
	MPC Implementation – Equation (15)		
$mpc_{r>\tilde{r}}, mpc_{r\leq\tilde{r}}$: 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]) Heterogeneity within retirement age group negligible (Andrews and Miller [2013])	Finding exogenous unanticipated income shocks to identify MPCs across retirement ages

Notes: The table summarizes our three proposed empirical implementations for the measurement of the social marginal value $\frac{SMU_{r<\tilde{r}}}{SMU_{r>\tilde{r}}}$ 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 $db_{r\leq\tilde{r}}$, 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 2 and 7 for details.

Appendix H Welfare Implementation Details

This appendix provides further detail on the welfare implementation described in Section 7 and illustrated in Figure 11 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	Risk aversion	Health Dep.	Welfare Wgts	ΔC	MPC
	$\gamma = 4, \theta = 1$	$\gamma = 2, \theta = 1$	$SMU \sim$ Health	$\gamma = 4, \theta \sim$ Life Exp.	$\gamma = 4$	
	(1)	(2)	(3)	(4)	(5)	(6)
A. Age-Specific Profile Change						
$\tilde{r} = 60$	.34	.17	.20	.41	.21	-.27
$\tilde{r} = 63$	.28	.14	.18	.33	.12	.08
$\tilde{r} = 65$	.76	.38	.70	.79	.14	.76
B. Swedish Pension Reform						
	.37	.18	.30	.40	.14	.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 baseline implementation, using the difference in consumption levels to approximate the the difference in SMU's (see equation (11)). Columns (2), (3) and (4) explore the sensitivity of the baseline results: (2) considers a change of the curvature in preferences, (3) allows for health-dependent marginal utility, and (4) assigns welfare weights that depend on life expectancy. Column (5) shows the results for the alternative implementation using the difference in consumption drops to approximate the difference in SMU's (see equation (14)), while column (6) uses the difference in MPC's (see equation (15)).

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 Implementation Figure 11 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,

$$\frac{SMU_{r\leq\tilde{r}} - SMU_{r>\tilde{r}}}{SMU_{NRA}} \approx \gamma \times \left[\frac{E_{r>\tilde{r}}(c)}{E_{r\in NRA}(c)} - \frac{E_{r\leq\tilde{r}}(c)}{E_{r\in NRA}(c)}\right]. \tag{40}$$

We obtain the estimates of the consumption levels for people retiring at age r relative to normal retirees, using regression (12). 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 11 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.

Sensitivity Analysis Columns (2), (3) and (4) 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 (2) is obtained by applying the same method as for the baseline implementation but reducing the curvature in consumption preferences to $\gamma = 2$.

Column (3) explores the robustness of the consumption smoothing cost estimates to health-dependence in the marginal utility of consumption. Following the results in Finkelstein et al. [2013], we assume that the semi-elasticity of the marginal utility of consumption to changes in a bad health measure, measured in standard deviations, is -0.25. To account for this, we adjust the SMU terms as follows:

$$\frac{SMU_r}{SMU_{NRA}} \cong \left(1 + \gamma\frac{c_{NRA} - c_r}{c_{NRA}}\right)\left(1 - 0.25(H_r - H_{NRA})\right), \tag{41}$$

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.[65] The differences $H_r - H_{NRA}$ are estimated for the retirement age groups as in Figure 5, i.e. using specification (12) 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 (4) presents a sensitivity analysis when assigning welfare weights to each retirement age r that depends on life-expectancy. We follow Chetty et al. [2016] to estimate the life expectancy and Becker et al. [2005] to adjust the welfare weights.

[65]We note that Finkelstein et al. [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 (3) 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.

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 $[79;90]$.

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

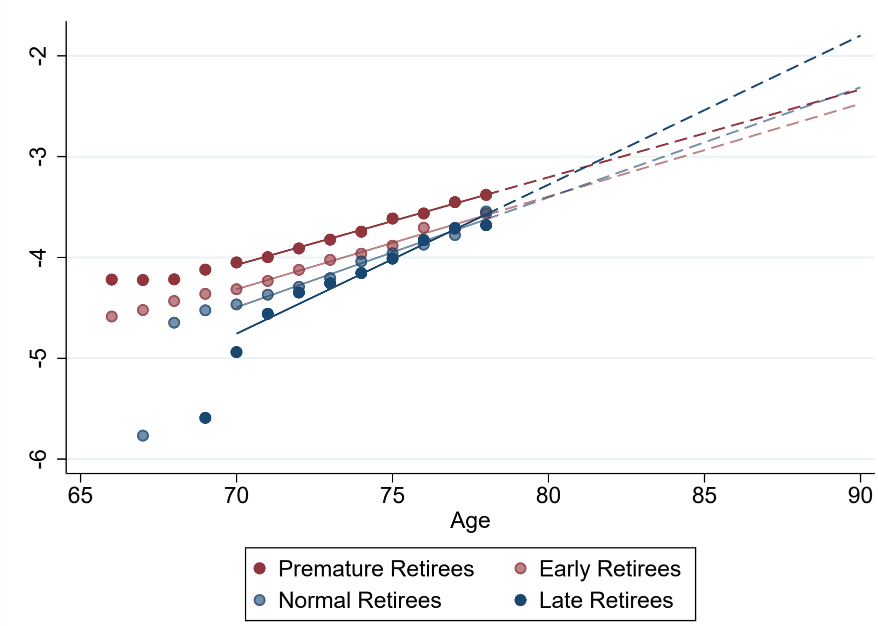

Notes: This figure plots the true mortality rates (dots) and the imputed mortality rates (line) using a Gompertz extrapolation, for each retirement age group. For the extrapolation, we consider only the computed mortality rates in the range $[70;78]$ (solid line). The mortality rates from the dashed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

$$\sum_{k=65}^{90} S_{k,NRA}\beta^k u(\bar{c}) = \sum_{k=65}^{90} S_{k,j}\beta^k u(\bar{c}+\Delta x_j), \tag{42}$$

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:

$$\sum_{k=65}^{90} S_{k,NRA}\beta^k = \sum_{k=65}^{90} S_{k,j}\beta^k(1 + \gamma\Delta x_j) \tag{43}$$

which simplifies to:

$$\gamma\Delta x_j = \frac{\sum_{k=65}^{90} S_{k,NRA}\beta^k - \sum_{k=65}^{90} S_{k,j}\beta^k}{\sum_{k=65}^{90} S_{k,j}\beta^k}, \tag{44}$$

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 as for the baseline 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 Implementations Columns (5) and (6) in Table H-1 present the results for the alternative implementations described in section 5.

Column (5) shows the results applying the alternative implementation for the consumption drops in equation (21),

$$\frac{SMU_r}{SMU_{NRA}} \cong \frac{1 + \gamma\frac{c_{r,pre} - c_{r,t}}{c_{r,pre}}}{1 + \gamma\frac{c_{NRA,pre} - c_{NRA,t}}{c_{NRA,pre}}}, \tag{45}$$

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 7. 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 (4).

Column (6) shows the results using the MPC implementation, following equation (15),

$$\frac{SMU_r}{SMU_{NRA}} \cong \frac{\frac{mpc_r}{1-mpc_r}}{\frac{mpc_{NRA}}{1-mpc_{NRA}}}, \tag{46}$$

assuming now that welfare weights are similar across retirement ages. For the marginal propensities to consume, we take values from Table 4. We then compute the odds ratio of the marginal propensities to consume, rescaled by the odds ratio of the normal retirees, following equation (21). 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

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

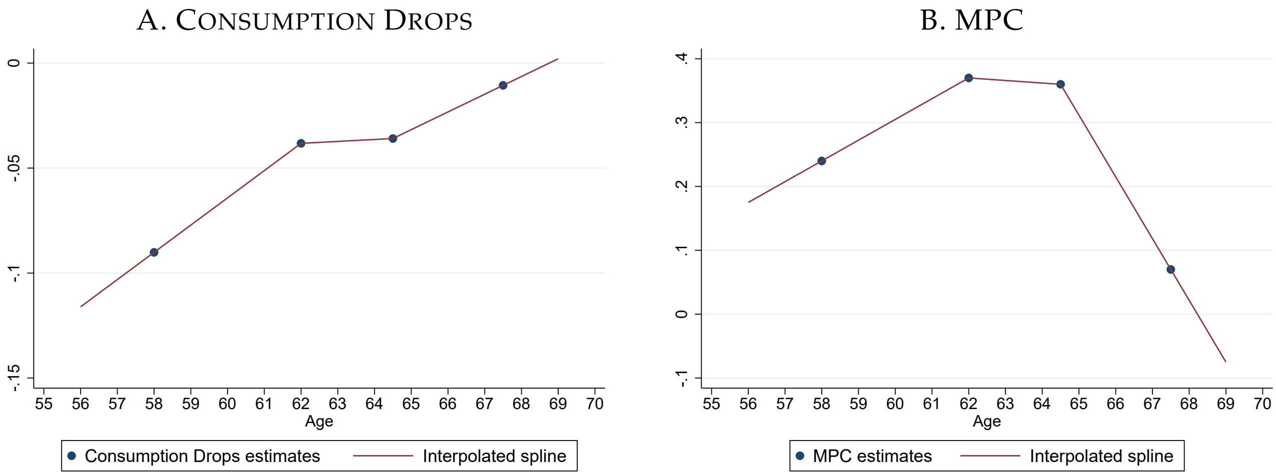

Notes: Panel A presents the consumption drops estimates from Figure 7 for the four retirement age groups (dots) and the interpolated linear spline between each of them. The consumption drop estimate for each retirement age group is assumed to lie at the midpoint of the interval. For instance, for the 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 4.

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 dollar transferred from early to late retirees, equals:

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

where $\mu_r = \frac{f(r)\left(\widehat{NDC}_r - ATP_r\right)}{\Sigma_{r \le \bar{r}} f(r)\left(\widehat{NDC}_r - ATP_r\right)} / S(\bar{r})$, 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 $\frac{SMU_{r \le \tilde{r}} - SMU_{r > \tilde{r}}}{SMU_{NRA}}$ when using the age-specific pension reforms considered above.

The age-specific estimates we take for the consumption-level implementations in columns 1-4 are again the ones reported in Figure 11. For instance, for the baseline analysis in column (1), we will take the consumption levels from regression (12). 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 7 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 the baseline implementation using the difference in consumption levels, but using the estimates from regression (12) restricted to the relevant sample.

Column (1) reproduces our baseline results, i.e., using the baseline 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 47.

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

	Baseline (1)	Bottom 10% (2)	Top 10% (3)	Couples (4)	Singles (5)	UI/DI (6)
	A. Age-Specific Profile Change					
$\tilde{r} = 60$	.34	.16	.19	.30	.49	.12
$\tilde{r} = 63$	.28	.10	.22	.24	.40	.15
$\tilde{r} = 65$	.76	.45	.46	.71	.88	.75
	B. Swedish Pension Reform					
	.37	.21	.21	.20	.30	.18

Notes: This table shows the results of the heterogeneity analysis of the baseline implementation. Column (1) replicates the estimates for the baseline analysis. Column (2) and column (3) produce the 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 (20) 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 to the Swedish pension reform. 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 et al. [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.