

NBER WORKING PAPER SERIES

RETIREMENT CONSUMPTION AND PENSION DESIGN

Jonas Kolsrud
Camille Landais
Daniel Reck
Johannes Spinnewijn

Working Paper 31628
<http://www.nber.org/papers/w31628>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2023

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 FajardoSteinhauser, Diego Ferreras Garrucho, Alice Lapeyre, and Quirin Von Blomberg for outstanding research assistance. We gratefully acknowledge support from ERC grants #679704 and #716485. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Jonas Kolsrud, Camille Landais, Daniel Reck, and Johannes Spinnewijn. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Retirement Consumption and Pension Design
Jonas Kolsrud, Camille Landais, Daniel Reck, and Johannes Spinnewijn
NBER Working Paper No. 31628
August 2023
JEL No. H55

ABSTRACT

This paper analyzes consumption to evaluate the distributional effects of pension reforms. Using Swedish administrative data, we show that on average workers who retire earlier consume less while retired and experience larger drops in consumption around retirement. Interpreted via a theoretical model, these findings imply that reforms incentivizing later retirement incur a substantial consumption-smoothing cost. Turning to other features of pension policy, we find that reforms that redistribute based on 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.

Jonas Kolsrud
Department of Economics and Statistics
Linnaeus University
351 95 Växjö
Sweden
and National Institute of Economic Research,
Stockholm.
jonas.kolsrud@lnu.se

Daniel Reck
University of Maryland
Department of Economics
Tydings Hall
7343 Preinkert Dr
College Park, MD 20742
and NBER
dreck@umd.edu

Camille Landais
Department of Economics
London School of Economics
Houghton Street
London WC2A 2AE
United Kingdom
and CEPR
c.landais@lse.ac.uk

Johannes Spinnewijn
Department of Economics
London School of Economics, Houghton Street
London, WC2A 2AE
J.Spinnewijn@lse.ac.uk

Online appendices are available at <http://www.nber.org/data-appendix/w31628>

Retirement Consumption and Pension Design

Jonas Kolsrud Camille Landais Daniel Reck Johannes Spinnewijn*

May 2, 2023

Abstract

This paper analyzes consumption to evaluate the distributional effects of pension reforms. Using Swedish administrative data, we show that on average workers who retire earlier consume less while retired and experience larger drops in consumption around retirement. Interpreted via a theoretical model, these findings imply that reforms incentivizing later retirement incur a substantial consumption-smoothing cost. Turning to other features of pension policy, we find that reforms that redistribute based on 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.

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 is 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

*Kolsrud: Dept of Economics and Statistics, Linnaeus University, 351 95 Växjö, Sweden; National Institute of Economic Research, jonas.kolsrud@lnu.se. Landais: Dept of Economics, London School of Economics, Houghton Street London, WC2A 2AE, c.landais@lse.ac.uk. Reck: Dept of Economics, University of Maryland, 3114 Tydings Hall, 7343 Preinkert Dr, College Park, MD 20742; dreck@umd.edu. Spinnewijn: 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.

¹The 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, Kanninen and Ravaska \[2022\]](#)). For further details, see [OECD \[2015, 2017, 2019\]](#).

fiscal effects as 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 complexities of 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 effect of such a reform. As the direct welfare effect depends on differences in the marginal utility of a dollar of pension benefits across groups of individuals, we label this a *consumption smoothing* effect.² We map empirical differences in consumption 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 Zweimüller [2013]; Manoli and Weber [2016]; Laun [2017]; Manoli and Weber [2018]; Gruber, Kanninen and Ravaska [2022]; Seibold [2021]; Lalive, Magesan and Staubli [2022]; Haller [2022]), and, potentially, behavioral internality effects [Mullainathan, Schwartzstein and Congdon, 2012; Spinnewijn, 2015; Reck and Seibold, 2021].

We use administrative registry data from Sweden, and registry-based measures of household consumption (see Kolsrud, Landais and Spinnewijn [2020]), to inform this trade-off between consumption smoothing and incentives. We start by closely examining the consumption smoothing effects of late-career incentives, due to the global policy focus on these incentives;

²While 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.

we subsequently examine other dimensions of pension benefits. We first attempt to understand therefore whether and to what extent early retirees have higher social marginal utilities of consumption than late retirees. We study how two consumption moments vary with the retirement age: the level of consumption during retirement, and the change in consumption around retirement. The consumption levels approach captures both the redistributive and insurance value of pension reforms, under the assumption that the non-consumption determinants of marginal utility in retirement do not systematically vary across groups. The consumption dynamics approach captures only the insurance value of pension reforms, under the weaker assumption that these non-consumption determinants of marginal utility evolve similarly around retirement across groups. We assess the assumptions underlying each approach in multiple ways.

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. Likewise, 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 at retirement between those two groups.³ We formally map these facts about consumption to the differential value of pension benefits to different groups.

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. Hence, within this set of ages, incentivizing later retirement is arguably less costly than at other ages. 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.

We supplement our main empirical results on consumption with data on a rich set of observable characteristics and with more granular data related to consumption structure. Doing so helps us assess the validity of the assumptions necessary to map estimated consumption patterns into welfare measures. More broadly, it informs the mechanisms underlying the differences in consumption levels and in consumption dynamics. For example, a number of factors suggest 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

³The drop in consumption at retirement has been widely studied and debated (e.g., [Banks, Blundell and Tanner \[1998\]](#), [Bernheim, Skinner and Weinberg \[2001\]](#), [Aguiar and Hurst \[2005\]](#), [Battistin et al. \[2009\]](#), [Stephens and Toohey \[2018\]](#)), but without considering how this drop varies by the retirement age.

retirement age. Overall, however, individuals retiring later are more well-resourced, in better health, and subject to 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. Furthermore, our analysis of survey data reveals that the composition of consumption is remarkably similar across retirement age groups both before and after retirement, which is consistent with the assumptions we use to map these consumption moments to welfare. 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. These findings generally reinforce our finding that incentivizing later retirement incurs a substantial welfare cost.

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. However, 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 internalities 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 tax 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 may 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.

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 (Banks, Blundell and Tanner [1998], Bernheim, Skinner and Weinberg [2001], Aguiar and Hurst [2005], Aguiar and Hurst [2013]; see De Nardi, French and Jones [2016], Jappelli and Pistaferri [2010] for reviews). Several papers have also aimed to uncover the importance of different determinants of retirement (see Blundell, French and Tetlow [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] and Ndiaye [2020] examine this trade-off using more structural approaches than ours. Structural approaches admit a fuller characterization of optimal policy, but the characterization may be model-dependent; we show how the welfare impact of pension reforms can be connected, for a large class of models, to moments that we can directly estimate. Complementary to our work is Haller [2022], who takes a similar sufficient-statistics approach to optimal pension design but focuses on the fiscal externality side of the trade-off. His work relates to a large empirical literature studying incentives and retirement behavior (e.g., Staubli and Zweimüller [2013]; Manoli and Weber [2016, 2018]; Gruber, Kanninen and Ravaska [2022]; Seibold [2021]; Lalive, Mage-

san and Staubli [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 I develops the conceptual framework that guides our empirical analysis. Section II describes the Swedish institutional setting and our data. Section III analyzes consumption levels and consumption dynamics by the retirement age. Section IV maps these consumption patterns to the welfare cost of late-career incentives. Section V considers consumption along other dimensions that are relevant for pension design: early-career labor supply, income, and wealth. Section VI 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.

I Conceptual Framework

This section presents a general framework to evaluate the design of pension policies. The framework guides our empirical analysis and motivates our focus on specific consumption moments in the data. Pension benefits are often a complex function of individuals' employment history, including their retirement age and past contributions. The value of pension benefits conditional on having specific features depends on the social marginal utility of these benefits for individuals with the specified features. To evaluate a pension reform, we first show that it suffices to compare the relevant social marginal utilities to the fiscal externality due to the behavioral responses triggered by the reform. We then show that differences in social marginal utilities across beneficiaries can be related to differences in their consumption.⁵

Setup Our model encompasses the rich heterogeneity and non-separabilities that are standard in retirement models (e.g., French [2005], French and Jones [2011]). We assume that 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 their utility and choices at that time. This can include an individual's past earnings and savings, shocks to their health, human capital or financial capital, etc. An individual chooses $c(\pi_{i,t})$ and $\zeta(\pi_{i,t})$ determining their flow utility $u(c(\pi_{i,t}), \zeta(\pi_{i,t}))$ 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. In particular, the marginal utility of consumption may be different under employment versus retirement, in accordance with a large literature on non-separabilities in consumption-leisure (see Jappelli and Pistaferri [2017]).

⁴Additionally, 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.

⁵In Appendix G we discuss further details regarding the setup and provide the full derivation of all equations and approximations.

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

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

In addition to the consumption decision $c(\pi_{i,t})$, we zoom in on the key decision to stay in the labor force or to retire, denoted by $s(\pi_{i,t}) \in \{1, 0\}$ and included in $\zeta(\pi_{i,t})$. If $s(\pi_{i,t}) = 0$ (retirement), the individual receives pension benefits $b(\pi_{i,t})$, which can depend on the individual's employment history in a general way. If $s(\pi_{i,t}) = 1$ (employment), the individual earns wages $w(\pi_{i,t})$ and pays taxes $\tau(\pi_{i,t})$. Assets $a_{t+1}(\pi_{i,t})$ evolve in the usual fashion, based on previously accumulated assets and saving in year t , with a gross rate of return $R(\pi_{i,t})$. The individual's optimization problem is to maximize \mathcal{U}_i subject to the corresponding budget constraint for each history $\pi_{i,t}$. We denote the resulting indirect utility by $U_i(b, \tau)$.

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

$$(2) \quad \max \mathcal{W}(b, \tau) = \int_i \omega_i U_i(b, \tau) di + \lambda GBC(b, \tau).$$

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). This welfare effect depends on two terms.

The first term captures the marginal value of the pension benefit:

$$(3) \quad \frac{\partial \int_i \omega_i U_i(b, \tau) di}{\partial b_{x,t}} = G(x, t) \times \underbrace{E \left(\omega_i \frac{\partial u(c_{i,t}, \zeta_{i,t})}{\partial c} \Big|_{x_{i,t} = x} \right)}_{\equiv SMU_{x,t}},$$

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 individuals' welfare, due to the envelope theorem.

The second term captures the fiscal cost of the marginal change in pension benefits:

$$(4) \quad \lambda \frac{\partial GBC(b, \tau)}{\partial b_{x,t}} = \lambda G(x, t) \times [1 + FE_{x,t}].$$

This expression makes explicit that the fiscal cost includes the fiscal externality due to the agents' behavioral responses, denoted by $FE_{x,t}$. That is, the same changes in behavior imply that the fiscal cost of increasing expected pension expenditures by one dollar may differ from one dollar. 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 invoked envelope theorem relies on individuals' optimizing their behavior. The presence of behavioral biases would add a third term to the welfare effect, consisting of marginal internalities and the corresponding behavioral responses to the reform [Mullainathan, Schwartzstein and Congdon, 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.⁶

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. Putting together equations (3) and (4), optimality of the relative benefits of early and late retirees, or more generally any two groups with different characteristics x and x' , requires that

$$(5) \quad \frac{SMU_{x,t}}{SMU_{x',t}} = \frac{1 + FE_{x,t}}{1 + FE_{x',t}}.$$

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

Equation (5) 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.

⁶An important concern highlighted in prior work is that individuals may not be adequately prepared for retirement. 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 (3). Moreover, in the empirically dominant case of so-called 'passive savers' documented in [Chetty et al., 2014], passive individuals would not change their savings behavior in response to the reforms (and the active savers who do reoptimize would not experience first-order changes in their welfare due to envelope conditions). 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].

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 (5) can be used to evaluate reforms to pension benefits that incentivize later retirement, as Sweden and many other countries have recently done.

Consumption Smoothing The focus in our empirical analysis is on the consumption smoothing aspect of pension reforms, i.e. the left-hand side of equation (5). We defer further analysis of the fiscal externality due to later retirement to the welfare illustration in Section VI. 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? We follow a standard approach in the social insurance literature that relates differences in marginal utilities to differences in consumption (e.g., Gruber [1997], Chetty and Finkelstein [2013]). Using Taylor series approximations around different consumption benchmarks, we can provide two alternative characterizations of the SMU for retirees with features x relative to the SMU of those retiring with features x' . Denoting the welfare weight and the relative risk aversion by ω_x and γ_x respectively, we find:

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

$$(6) \quad \frac{SMU_{x,t}}{SMU_{x',t}} \cong \frac{\omega_x}{\omega_{x'}} \times \frac{\phi_{x,t}}{\phi_{x',t}} \times \left[1 + \gamma_x \frac{c_{x',t} - c_{x,t}}{c_{x',t}} \right],$$

where t refers to an age or time period after retirement and $\phi_{x,t} = \frac{\partial u(\cdot, \zeta_{x,t})}{\partial c}$ captures the impact of ζ_x when retired on the marginal utility of consumption (evaluated at $c_{x',t}$ for both subgroups).

Consumption-Drop Implementation. Assuming $c(\pi_{i,t}) = c_{x,t}$, $\zeta(\pi_{i,t}) = \zeta_{x,t}$ for $x(\pi_{i,t}) = x$ and this for any i, t, x , we can approximate

$$(7) \quad \frac{SMU_{x,t}}{SMU_{x',t}} \cong \frac{\tilde{\omega}_{x,pre}}{\tilde{\omega}_{x',pre}} \times \frac{\phi_{x,t}/\phi_{x,pre}}{\phi_{x',t}/\phi_{x',pre}} \times \frac{1 + \gamma_x \frac{c_{x,pre} - c_{x,t}}{c_{x,pre}}}{1 + \gamma_{x'} \frac{c_{x',pre} - c_{x',t}}{c_{x',pre}}},$$

where t refers to an age or time period after retirement, while “pre” refers to an age or time before retirement. $\tilde{\omega}_{x,pre} = \omega_x \cdot \frac{\partial u(c_{x,pre}, \zeta_{x,pre})}{\partial c}$ represents the (Pareto-weighted) marginal utility of pre-retirement consumption of group x . And $\phi_{x,t}/\phi_{x,pre} = \frac{\partial u(\cdot, \zeta_{x,t})}{\partial c} / \frac{\partial u(\cdot, \zeta_{x,pre})}{\partial c}$ captures the impact of changes in ζ_x around retirement on the marginal utility of consumption (evaluated at $c_{x,pre}$ for subgroup x).

The first implementation highlights that the difference in $SMUs$ crucially depends on the difference in consumption levels across retirement groups receiving different pension benefits. For example - everything else equal - the lower is consumption by early retirees relative to

late retirees, the higher is the cost of incentivizing later retirement by steepening the pension profile. The second implementation instead highlights the relation between the difference in *SMUs* and the difference in consumption drops around retirement. For the same example, it indicates that the cost of providing late-career incentives is higher the larger the drop in consumption around retirement for early retirees relative to late retirees.

Redistribution vs. Insurance The two implementations capture different aspects of consumption smoothing, which may be important to separately identify from a policy perspective. In particular, the consumption-level implementation accounts for all differences in consumption when retired to evaluate welfare, including differences that pre-date retirement. This is desirable when planners wish to capture the overall *redistributive* effects of pension reforms.⁷ The second implementation highlights that only the differences in consumption drops around retirement remain relevant when planners wish to take the differences in social marginal utilities that exist before retirement as given (e.g., $\omega_x \frac{\partial u(c_{x,pre}, \zeta_{x,pre})}{\partial c} = \omega_{x'} \frac{\partial u(c_{x',pre}, \zeta_{x',pre})}{\partial c}$). This more narrow *insurance* perspective is often used in the social insurance literature. Pension policies may provide valuable insurance as workers face work-longevity risk due to shocks to productivity, health or ability.⁸ The focus on this insurance value is often motivated by the availability of other policy tools for redistribution or insurance of earnings differences during an individual's working life (e.g., progressive income taxes, unemployment insurance).

Heterogeneity and Non-Separability The two alternative implementations show that the mapping from consumption into *SMUs* requires different information and assumptions on preferences depending on the consumption moments used. A fundamental implementation challenge may arise when the marginal utility of consumption for individuals retiring with different features x differs, *conditional on consumption*. As shown in equation (6), if the conditional marginal utility of consumption $\phi_x = \frac{\partial u(\cdot, \zeta_x)}{\partial c}$ differs for different retirement groups, their differences in consumption levels would no longer be sufficient to evaluate welfare.⁹ Differences in the non-consumption determinants of marginal utilities ζ_x – e.g., a smaller value of consumption relative to leisure – may be particularly concerning as they can drive selection into retirement. Still, any *permanent* differences in ζ_x across groups become inconsequential for the consumption-drop implementation, when the focus is on the insurance value of the reform. A potential concern remains that different *within-individual* changes in ζ_x around retirement confound the mapping from consumption drops for each group and their *SMU*, as captured by $\phi_{x,t} / \phi_{x,pre}$ in equation (7).

The issue of non-separabilities is abundantly highlighted in the retirement-consumption literature, in particular in explaining the puzzle why consumption drops between employment and

⁷Other welfare concerns, for example about differences in health or life expectancy, may imply different weights (ω_x) across retirement groups. We study such differences in Section III.

⁸In their classic paper on pension design, [Diamond and Mirrlees \[1978\]](#) consider a social planner who uses public pensions to provide insurance against work-longevity risk due to disability. In their framework, consumption drops are expected to be higher for people who retire earlier as this is induced by disability shocks.

⁹Another 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 (see [Andrews and Miller \[2013\]](#)). The within-group heterogeneity, however, is only relevant to the extent that this differs across these groups.

retirement (e.g., [Aguiar and Hurst \[2005\]](#), [Stephens and Toohey \[2018\]](#)). For example, observed consumption expenditures could translate differently into real consumption while employed vs. retired (e.g. due to differential reliance on home production), or actual preferences over consumption could differ (e.g., due to differential complementarities with leisure). However, it is important to note that our welfare evaluation compares the marginal utility of individuals *when retired*. So we circumvent the common concerns raised in relation to the so-called retirement-consumption puzzle to the extent that they affect the relevant retirement groups in the same way.

To gauge this further, we perform several empirical exercises in which we leverage the richness of the Swedish administrative data in Sections III and IV. In particular, we examine the composition of consumption and a rich set of additional observable characteristics to gain insight into the sources of consumption differences across retirees. We also use consumption surveys to study more directly the potential importance of heterogeneity in consumption preferences across retirement groups. Finally, following [Landais and Spinnewijn \[2021\]](#), we study yet another consumption moment in [Appendix F](#) – i.e., the marginal propensity to consume – to assess the sensitivity of our results to potential differences in marginal utilities.¹⁰

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 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 section III 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 V, 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.¹¹ Public pensions are complex and vary across countries, but together, income, early- and late-career labor supply capture the most important dimensions of benefit variation observed across all pension systems. We discuss how Swedish pension benefits map onto these conceptual dimensions in the next section.

¹⁰The different 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.

¹¹Wealth 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.

II Institutional Background & Data

A Institutional features of the Swedish pension system

The Swedish Pension system comprises three primary pillars: public pensions, occupational pensions, and private pensions. We focus on public pensions, but account for the presence of the other pillars. Sweden is transitioning from a defined benefit system, called the “ATP” scheme, to a “Notional Defined Contribution” (NDC) scheme. The goal of this paper is not to evaluate this reform, but rather to evaluate conceptual reforms that inform pension design in general. As such, we describe here how the NDC reform relates to conceptual reforms considered above; [Appendix A](#) presents a more comprehensive review of Swedish pensions.

Public Pension Reform The NDC reform was passed in 1994 and has been phasing in gradually across cohorts since 1998.¹² Pension benefits are financed by payroll taxes before and after the reform. In the pre-reform ATP system, pension benefits are determined by average earnings in the 15 highest-earning years of one’s career, and career length up to a 30-year cap.¹³ Pension rights can be earned between ages 16 and 64, but not beyond 65. Annual pensionable earnings are capped at around the empirical median of the earnings distribution for 55-year-olds in 2000. There is a minimum benefit level for retirees with short careers and low lifetime earnings. About one quarter of 66 year olds received the minimum pension benefit in 2007.

The post-reform NDC system resembles a DC system from a worker’s perspective. Benefits are an annuity closely linked to a worker’s lifetime contributions through payroll taxes. Unlike a typical private DC scheme, 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. Labor income after age 65 is pensionable. The maximum annual pensionable earnings was increased by about 25%, and the minimum pension benefit level was increased by about 40%. About 30% of all individuals receiving pension benefits are expected to receive basic pensions in 2040 when the NDC system is fully phased in.

Conceptualizing the Reform How does the NDC reform differently affect workers along the three key dimensions of pension benefits discussed above – late career labor supply, early career labor supply, and permanent income? Like recent reforms in many countries, incentivizing later retirement was an explicit goal of the NDC reform. Two features of the Swedish reform directly contributed to this aim: allowing workers to earn pension rights beyond age 64, and removing the 30-year contribution cap.¹⁴

Panel B of [Figure 1](#) illustrates how the Swedish reform increased late-career incentives, mirroring the stylized reform shown in [Panel A](#).¹⁵ We use simulations of lifetime income and pension

¹²Cohorts born before 1938 receive their pension benefits from the ATP system. Those born between 1938 and 1953 receive a weighted mixture of ATP and NDC benefits, with increasing weight on the NDC benefits over time.

¹³Pensionable earnings includes labor income and income from social insurance benefits based on labor income and excludes capital income and transfers that are not based on previous labor earnings.

¹⁴For 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.

¹⁵[Figure 1B](#) abstracts away from the overall level effect of the reform on pension benefits. However, to promote

benefits for a representative set of workers born in 1941. Holding earnings history fixed before 55, we calculate the net present value at age 55 of workers' pension benefits at different retirement ages (see [Appendix A](#)). The removal of the cap on pension rights after age 65 has a salient effect: the ATP schedule mechanically flattens out after 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.

While the late-career labor supply dimension was a focal point of the reform, some provisions affected early-career labor supply incentives and the link between benefits and 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. Turning to 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. Unlike many countries, Sweden has no earnings test whereby pension benefits are reduced for those continuing to work after claiming. 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 65. In the NDC system, the adjustments are on average actuarially fair by design, because pension benefits are scaled by life expectancy at claiming age. Empirically, we observe more variation in retirement ages than in claiming ages – see [Appendix Figure A-3](#).¹⁶

Other Social Insurance Programs Motivated by our conceptual framework, we focus on retirement defined as the moment individuals stop working permanently. Although they are not labelled pension benefits, components of the Swedish social insurance system like disability insurance and unemployment insurance can cushion the shock of losing employment for the elderly. Such benefits affect workers' labor supply incentives in old age, both directly and because these benefits generate "pensionable earnings." We include these benefits as part of the overall pension system in [Figure 1B](#).¹⁷

Households and Individuals Like most of the Swedish tax and transfer system, the pension system is entirely individualized, with two exceptions. First, the minimum benefit in both systems is about 10% lower for married individuals. Second, there is a survivor's benefit that is paid out for a year after one's spouse has passed – see [Appendix A](#) for details.

fiscal sustainability, the NDC reform enacted a reduction in pension benefits for most workers. We illustrate the level effects in [Figure A-12](#).

¹⁶In the cohorts we study, 69% of workers claim their pension at age 65, but only about 22% retire at age 65; far more workers retire before 65 than claim before 65. In [Figure 1B](#), we assumed workers claimed at 65 for simplicity; we discuss and conduct further analysis on this point in [Appendix A](#).

¹⁷Doing so has little impact on the NPV of pension benefits or the implied fiscal externalities from incentivizing later retirement (see [Appendix A](#) for details and sensitivity analysis).

B Data

We rely on uniquely rich data from several Swedish population registries, 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 covering all 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, and transfers, including sickness, disability and unemployment benefits. From LISA, we construct a measure of retirement, defined as the moment individuals stop working permanently. We follow [Karlstrom, Palme and Svensson \[2004\]](#) and categorize an individual as retired when her annual labor earnings permanently fall below one “Base Amount” – about 18% of median labor earnings.¹⁸

Our second main registry dataset covers pensions. Data from the ATP system contains pension contributions from 1960 onwards for all individuals born 1938 and later. Contributions under the NDC system are available from the late 1990s, when it was created. In addition, the data covers all pension benefits that individuals accrue and receive: old age state pension benefits, occupational pension and private pension savings.

Consumption Measures We use the registry-based measure of annual household consumption expenditures for all Swedish households created for the years 2000 to 2007 by [Kolsrud, Landais and Spinnewijn \[2020\]](#) (see also [Appendix B](#)). The construction of this measure relies on the identity, 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, taxes, and transfers, and annual capital income. Second, we have precise data on wealth from the population wealth tax register (*Förmögenhetsregistret*), which contains detailed individual 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 wealth tax register data with data 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.

We complement our registry-based consumption measure with survey-based measures of consumption expenditure from the Swedish consumption expenditure survey (HUT). This allows us to investigate the structure of consumption expenditures across categories. [Appendix B](#) and [Kolsrud, Landais and Spinnewijn \[2020\]](#) present details on the construction of our consumption measures and a thorough assessment of the robustness and consistency of registry-based vs survey measures of expenditures.

¹⁸The one base amount (BA) threshold is widely used to define labor force participation in administrative data. In the ATP system, individuals must earn at least 1 BA in order to earn pension rights in a year ([Appendix A](#)). Specifically, we assume retirement takes place during the last year in which an individual’s earnings exceed 1 BA.

Health and Mortality We further complement our data with the death register, as well as with two large surveys containing detailed information on health and health expenditure: the living condition survey (ULF) and the household finance survey (HEK). 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 retirement ages for these cohorts. The vast majority of individuals retire between 55 and 70, with a peak at 65. To analyze 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 (inclusive); early retirement, ages 61-63; normal retirement, ages 64-65; and late retirement, ages 66-69. 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 because we observe the full ATP contribution history for cohorts born from 1938 onwards, and our consumption data spans 2000 to 2007. Our sample selection allows us to observe consumption before and during retirement for each retirement age group in each cohort. This selection of cohorts therefore allows us to control for both age and cohort effects in consumption. For these cohorts, the ATP system was the main determinant of pension benefits and the NDC was just beginning to be phased in.

[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.

III Consumption Patterns By Retirement Age

This section studies the differences in consumption patterns across individuals retiring at different ages. We estimate both differences in consumption levels when retired and differences in consumption dynamics around retirement, following the conceptual framework in [Section I](#). We also analyse potential sources of heterogeneity underlying different consumption patterns. The estimated differences are used to evaluate the welfare costs of pension reforms that strengthen late-career labor supply incentives, which we do in the next section.

A Consumption Levels During Retirement By Retirement Age

We first document how consumption levels differ 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 \(6\)](#).

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 :

$$(8) \quad C_{it} = \sum_j \alpha_j \cdot \mathbb{1}[r = j] + \gamma_y + \gamma_t + \mathbf{X}'\beta + \varepsilon_{it}.$$

We estimate model [\(8\)](#) 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 (6), 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 (8) 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 $E_j[\tilde{C}_{it}]$, the average predicted consumption level in retirement age group j from specification (8) when omitting the contribution of the retirement age group dummies.¹⁹ We start, on the left hand side of the graph, with results from model (8) where only year and age fixed effects are included. The rest of figure shows the same estimated coefficients when sequentially adding controls for family composition 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 10 to 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

¹⁹ $E_j[\tilde{C}_{it}]$ 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 .

household composition and other pension determinants, but there remains a clear flattening of the consumption gradient around this retirement age range.²⁰

Robustness The two main consumption patterns are robust across different specifications. 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 (8), but separately for each age t between 66 and 69.²¹ We document a strong positive gradient of consumption with retirement age for each age t . The consumption of late retirees when retired 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 consumption patterns are similar across household structures. We replicate specification (8), 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-monotonicity between early and normal retirees appears to be driven virtually entirely by couples.

The consumption patterns also appear when using survey data, which allows us to consider different countries too. We use survey data from SHARE and HRS which contain information on retirement and survey measures of consumption expenditures for similar cohorts in 11 European countries and the US. While the policy environment obviously differs across countries, many countries share institutions similar to those described in Section A, which penalize early retirement. We report the results in Appendix D. They confirm the large gradient in consumption levels between individuals who retire very late versus very prematurely, suggesting that this is a robust finding, not only across data sources, but also across contexts. Also the non-monotonicity is strikingly robust: 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.

B Consumption Dynamics Around Retirement by Retirement Age

We now turn to the consumption dynamics around retirement and revisit the consumption drops at retirement which have drawn much 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.

We start by residualizing household consumption on a set of cohort fixed effects and age fixed

²⁰In Appendix Figure C-1, we also report estimates of a fully non-parametric version of specification (8) where we compare consumption levels across all retirement ages (rather than aggregating retirement ages into four groups). 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.

²¹Because 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*.

effects. Figure 4 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. Note that the graph scales residual consumption of each group by its level two years prior to retirement.²²

During the initial period up until two years before retirement, all retirement age groups apparently experience similar trends in consumption. But when focusing 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, French and Tetlow \[2016\]](#), that earnings ability shocks are 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 remarkably similar drop in consumption, of about 5%, right at retirement. As mentioned, a large literature has focused on this drop in consumption at retirement, sometimes called the “retirement-consumption puzzle” (e.g., [Aguilar and Hurst \[2005\]](#), [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. But importantly, we find that consumption drops right 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 4 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,

²²Because 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 4.

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 above highlights the presence of significant heterogeneity in consumption dynamics across retirement age groups around the time of retirement. How important are these late-career divergences in explaining differences in consumption during retirement? In other words, are the large consumption differences during retirement documented in Figure 3 already present before retirement, or does late-career consumption dynamics play a meaningful role in determining post-retirement consumption?

The first simple way to address this question is to replicate the type of analysis of Figure 3 using pre-retirement consumption levels instead, to see how much of these consumption differences across retirement age groups pre-date retirement. In Figure 5, we look at consumption at age 55 by retirement age groups. The results show the presence of significant consumption differences between the premature and the late retirees already at age 55. But the overall consumption gradient by retirement age is much more muted at age 55 than in post-retirement consumption. It is about twice as small compared to the gradient observed in Figure 3. Interestingly, this gradient also entirely disappears when controlling for income and career length at 55. This suggests that a significant fraction of the differences in retirement consumption across retirement age groups emerges in the last few years prior to retirement.

To further gauge the role of late career vs early career dynamics in determining retirement consumption, we replicate our estimates of consumption differences across retirement age groups from specification (8), 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 finally 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.²³ 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, revealing that at least half of the consumption differences across retirement age groups emerges in the very last stages of workers' careers.

²³Note 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.

C Heterogeneity and Selection into Retirement Age

Our results show large differences across retirement age groups in both the consumption levels when retired and the consumption dynamics around retirement. We leverage the rich data environment in Sweden to shed further light on potential sources of heterogeneity and on the nature of the late-career shocks underlying these differences.

A first natural step is to decompose our measure of household consumption expenditures into the different components we used to construct it, to shed light on the resources available to individuals around retirement. 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 (8) 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 6 reports the estimates $\hat{\alpha}_j$ for each component, scaled by $E_j[\tilde{C}_{it}]$, 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.²⁴ We note that Figure 6 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.²⁵ This evidence suggests that many individuals in this group retire earlier in part because they have the means to do so.

We obtain very similar insights from studying the selection on observables into one of the four different retirement age groups. We estimate a multinomial logit prediction model including a

²⁴As explained in section II, 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.

²⁵Note 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.

large set of socio-economic characteristics as well as cohort fixed effects. In Panel A of Figure 7, 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 are significantly more highly educated than all other retirees, earned much higher incomes and accumulated more assets. Premature retirees find themselves at the opposite end of the spectrum. 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. This is suggestive of significant wealth effects on labor supply around retirement (Giupponi [2019], French et al. [2022]). Early retirees 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.²⁶ Perhaps surprisingly, we find no significant differences in career lengths at age 55 across the different retirement age groups.

We finally consider differences in health and life expectancy across retirement age groups. Panel B of Figure 7 documents the presence of a steep negative health gradient over retirement ages. Earlier retirement is strongly associated with having significantly worse health, using two different health indices constructed using health surveys.²⁷ 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, already in the pre-retirement period, of a significant gradient in health across retirement age groups. 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.²⁸ Negative health shocks thus seem to be an important dimension of work longevity risk, especially for those leaving the labor market prematurely (see also Blundell et al. [2021]). The worsening health dynamics correlate strongly with the negative consumption dynamics experienced by premature retirees just prior to retirement.

Overall, our analysis thus implies that steeper pension profiles tend to redistribute to individ-

²⁶The average age difference between couples in our data is 3.8 years while the difference in age of retirement 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.

²⁷The 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.

²⁸Appendix 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 retiring due to health reasons.

uals who are already well-off across multiple observable dimensions, including having better health and life prospects. Steeper pension profiles also reduces insurance against work longevity risk for people having to retire earlier, for example due to disability or bad health.

IV Welfare Cost of Late-Career Incentives

This section maps the estimated consumption patterns into estimates of the consumption smoothing costs from strengthening incentives for later retirement. We follow the consumption implementations presented earlier in the conceptual framework (see Section I). We consider the welfare effect of a stylized pension reform, simply steepening 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} .²⁹ This type of reform is illustrated in Panel A of Figure 1 for $\tilde{r} = 65$. An overall change in the pension profile by retirement age, as shown in Panel B for the Swedish reform, can be evaluated as a combination of stylized reforms at different ages.

A Consumption-based Implementations

We first consider the consumption-level implementation of welfare cost. Building on equation (6), we can approximate the consumption smoothing cost of the stylized pension reform by

$$(9) \quad \frac{SMU_{r \leq \tilde{r}} - SMU_{r > \tilde{r}}}{SMU_{NRA}} \approx \gamma \times \left[\frac{E_{r > \tilde{r}}(c)}{E_{r \in NRA}(c)} - \frac{E_{r \leq \tilde{r}}(c)}{E_{r \in NRA}(c)} \right],$$

where the differences are expressed relative to the normal retirement age group as estimated in regression (8). We use a CRRA risk aversion parameter γ of 4 (see Landais and Spinnewijn [2021]) and thus assume for our baseline implementations no differences in welfare weights across retirement ages, nor in marginal utilities conditional on consumption. We consider alternative assumptions below.

Figure 10 plots the resulting consumption smoothing cost estimates of the stylized reform around retirement age $\tilde{r} \in [56, 67]$. For this consumption levels implementation, the consumption smoothing costs range between .27 and .80. Hence, per dollar(/krona) transferred from individuals 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 age 61 or after the age 65. Note that the stylized reform redistributes resources from *everyone below* a certain retirement age to *everyone above*. If we consider instead a more flexible reform like 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

²⁹To 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.

retirees, suggesting that there is locally no trade-off between incentives and redistribution for those retiring between 61 and 65.

We can also quantify the welfare costs of providing late-career incentives using the consumption-drop implementation of the *SMU*'s. As discussed, this focuses on the welfare cost of reducing insurance by steepening the pension-profile.³⁰ Column (1) in Table 2 repeats the estimation of consumption smoothing costs using the consumption-level implementation, but now when transferring resources between the four retirement-age groups considered earlier. Column (2) shows the corresponding estimates using the consumption-drop implementation, following equation (7). We use the consumption drops from two years before to two years after retirement (see Figure 4), scaled by $\gamma = 4$. The estimated welfare 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 retirees 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). We further note the same non-monotone pattern arises again, implying the welfare cost is lowest when steepening incentives more locally between the early and normal retirement age, rather than before and after some age.

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 shocks. However, beyond reducing insurance, providing more incentives also reduces redistribution towards individuals with lower pre-retirement consumption.

B Robustness of Welfare Estimates

The validity of our mapping from consumption moments to *SMUs* depends crucially on the assumption that preferences do not differ across retirement age groups, as discussed in Section I. The welfare estimates also rely on a specific estimate for the curvature in consumption preferences.

Preference Heterogeneity The consumption-level implementation depends on the absence of differences in non-consumption determinants of the marginal utility of consumption across retirement age groups, as shown in equation (6). The consumption-drop implementation is robust to any persistent difference, but still relies on the absence of differences in within-individual changes in the non-consumption determinants of marginal utilities around retirement, as shown in equation (7).

To gauge this further, we first 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 C-8 ranks the 11 consumption categories by their importance for the

³⁰Recall that 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.

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.

We can also use the survey data to gauge the importance of within-individual changes in preferences around retirement. As discussed, retired individuals may have more time, increase their home production of goods, search for better prices, spend less on work-related expenditures, etc. and this may change the marginal utility of consumption beyond the differences in consumption. Appendix Figure C-9 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 consumption-drop implementation is that these changes in non-consumption determinants of marginal utilities are similar across retirement age groups. On this front, Appendix Figure C-9 is reassuring, as the changes in consumption structure are very similar across all retirement age groups.

Taken together, our results seem to suggest that the differences in consumption across retirement age groups are driven more by differences in the means for consumption and differences in shocks than by differences in preferences for consumption. Our results further confirm that the mechanism behind the retirement-consumption puzzle is *prima facie* inconsequential to evaluate the relative value of pension benefits across retirement age groups.

Marginal Propensity to Consume The preference parameter γ is crucial for translating consumption differences into differences in marginal utilities in the two consumption-based implementations, but generally hard to estimate empirically (see Chetty and Finkelstein [2013]).³¹ As shown in Landais and Spinnewijn [2021], differences in marginal propensity to consume (MPC) can be used to capture differences in SMUs, without having to make specific assumptions on the value of the curvature parameter γ . The main intuition is that 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. Using MPCs thus narrows our welfare focus further on the liquidity value that pensions provide.

Appendix F describes the method in detail, and offers an empirical implementation using variation in individuals' financial wealth, from quasi-random shocks to the price of stocks that individuals hold in their portfolio. We find that MPCs out of wealth shocks increase markedly

³¹Reducing the curvature 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)). 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 implementations ($\gamma \geq 4$).

after retirement. Furthermore, the estimated MPCs after retirement exhibit a strong negative gradient over the retirement age: the value of additional liquidity appears to be high for individuals who retire early or prematurely, but it appears negligible for late retirees. Overall, we find that the welfare costs of increasing late-career incentives inferred from these MPC estimates generally accord with those obtained from our consumption-level and consumption-drop implementations.

V Welfare Cost of Alternative Dimensions of Pension Reforms

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 I: 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 VI.

A 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.³² 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 NDC system, treating contributions in all years of the career equally. More generally, when strengthening the link between pension contributions and benefits – e.g. when switching 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.³³

Panel A of Figure 8 illustrates how consumption varies by career length at 55, using the specification in model (8) 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 notable. 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 Figure 8: those with very short careers also have low consumption. Comparing across specifications, we

³²We define career length as of age 55 as the number of years prior to age 55 in which an individual had pensionable income.

³³In principle, governments could even reform pension benefits to specifically target early-career labor supply, but they seldom do.

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 7 and revealing striking heterogeneity behind the consumption patterns in Figure 8. Focusing first on those with long careers versus those with medium-length careers, they tend to 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 9 reveals that the consumption differentials in Figure 8 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 8. 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 rather than insurance against late-career shocks: if we condition on pre-retirement consumption, the estimated consumption smoothing gains would drop down to basically zero.

B 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).³⁴ 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 8 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 9, 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 stabilizes or declines 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, Skinner and Weinberg \[2001\]](#) of a substantial wealth gradient in the drop in consumption at retirement. The larger drop in consumption at retirement for low-wealth individuals explains about 30% of the overall wealth gradient in consumption at retirement from Figure 8.

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-

³⁴Note 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.

VI 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.

A Late-Career Incentives

To evaluate the social welfare effect of strengthening late-career incentives, we should compare the consumption smoothing cost estimates from Section IV to the fiscal externality associated with behavioral responses to the changed incentives. To calibrate the fiscal externality in simple terms, we focus on retirement responses and abstract from the fiscal implications of changes in saving and other labor supply responses. We can then approximate the net welfare gain per dollar(/krona) transferred from individuals retiring before \tilde{r} to those retiring after by

$$(10) \quad \Delta W_{\tilde{r}} \cong \frac{\tau_{\tilde{r}} - [NPV_{\tilde{r}+1} - NPV_{\tilde{r}}]}{w_{\tilde{r}}} \times \varepsilon_{\frac{S(\tilde{r})}{1-S(\tilde{r})}, w_{\tilde{r}}} - \frac{SMU_{r \leq \tilde{r}} - SMU_{r > \tilde{r}}}{SMU_{NRA}}.$$

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} , governing how much the survival rate into employment at age \tilde{r} increases due to stronger incentives to continue working at age \tilde{r} .³⁵ 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.

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 on labor earnings and thus positive. Our simulation results suggest a participation tax rate of about .45 (see Appendix A). For the locally relevant behavioral elasticity, we use the labor supply elasticity of .22 estimated by Laun [2017] for Swedish individuals over age 65 circa 2007.³⁶ Altogether, we obtain a fiscal

³⁵Here 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} . 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, so we also assume that the fiscal externalities of the opposing income effects cancel out for a budget-balanced reform. Finally, we express the welfare effect relative to the value of a dollar given to our reference group, those retiring at the normal retirement age (SMU_{NRA}), which we assume to be approximately equal to marginal cost of public funds λ .

³⁶The identification of this elasticity comes from variation in age-specific tax credits rather than a pension reform. Nevertheless, this estimated elasticity is in line with estimates from pension reforms in other European countries (e.g. .25 in Manoli and Weber [2016] and .33 in French et al. [2022]), and somewhat higher than the elasticities found in Seibold [2021] ($\approx .1$) in Germany using bunching methods.

externality from inducing later retirement of .15: the government collects 15 cents per dollar transferred from individuals retiring before \tilde{r} to individuals retiring after \tilde{r} . [Appendix G](#) and [Appendix H](#) provide details on the derivation and the implementation of the fiscal externality respectively.

Assuming a *constant* fiscal externality of 0.15, we find that the consumption smoothing costs of reforming late career incentives are larger than or comparable to the fiscal benefits. In [Appendix Table H-1](#), we evaluate the overall change in late-career incentives from the Swedish reform depicted in [Figure 1B](#), aggregating across stylized reforms. We find a consumption smoothing cost of 0.36 using the consumption levels implementation capturing the redistributive and insurance values of pension benefits, and a consumption smoothing cost of 0.14 using the consumption drops implementation capturing only the insurance value. The consumption levels estimate would suggest that the net welfare effect of strengthening incentives is negative, especially below the early entitlement age (61) and above the normal retirement age (65). The non-monotonicity in consumption smoothing costs 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. In contrast, the Swedish reform strengthened incentives especially for late retirees (see [Figure 1B](#)).

We attach some caveats to these welfare and policy conclusions. First, our discussion presumed that the fiscal externality is similar across retirement ages. The participation tax is indeed 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. We have little evidence on how labor supply elasticities vary in the age range of interest (see [Appendix H.1](#) for further discussion). The most direct evidence comes from [Seibold \[2021\]](#), who finds similar labor supply responses to financial incentives across 400 kinks in the German pension profile and no significant heterogeneity in responsiveness across observable characteristics (e.g. education, birth cohort, lifetime earnings, unionization or health) that correlate with retirement age.

Second, our implementations map consumption moments into *SMU*'s under specific assumptions on preferences and welfare weights (see [Table G-1](#)). We gauge the sensitivity of our conclusions to different implementation assumptions in [Appendix Table H-1](#). Adjusting the *SMU*'s for estimated differences in health could reduce the consumption smoothing cost by 20%, assuming that earlier retirees having worse health also have lower marginal utility of consumption, following [Finkelstein, Luttmer and Notowidigdo \[2013\]](#) (see column (3)). However, because late retirees live longer, assigning welfare weights to retirement age groups based on their differential life-expectancy, following [Becker, Philipson and Soares \[2005\]](#), would increase the consumption smoothing costs by 10% (see column (4)). [Appendix H.1](#) provides details on the methodology underlying these calculations.

Third, our analysis demonstrates the value of providing insurance against work-longevity risk. As discussed above, DI and UI provide complementary insurance to the pension system, especially for premature retirees, whose pathway into retirement is often through DI or UI (see [Appendix Figure A-6](#)). Our consumption-based estimates of the marginal value of extra trans-

fers already account for all resources retirees can rely on, including social insurance transfers. Accounting for the availability of UI/DI modestly increases the fiscal externality from inducing individuals to work longer, from .13 to .15 (see Appendix Figure A-9).³⁷

B Other Policy Implications

What do the results on the other dimensions of pension benefits examined in Section V imply about pension design? First, while we find 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 potential responses in early-career labor supply (e.g., French et al. [2022]), educational attainment, or family-related career interruptions.

Second, our results suggest that valuable redistribution could be accomplished through the income and/or wealth dimensions. Any gains from redistributing along these dimensions should obviously be compared to the relevant fiscal externality due to behavioral responses. Another natural question is whether income taxes could accomplish the same redistribution more effectively (Atkinson and Stiglitz [1976]). Our results nevertheless suggest that conditioning on wealth (or perhaps lifetime income) would allocate pension benefits to those who most value them. Wealth predicts not only consumption in retirement but also the drop in consumption at retirement, suggesting that low-wealth retirees particularly value the insurance provided by public pensions.

Naturally, under-saving due to behavioral biases could be one reason why low wealth is associated with a larger drop in consumption at retirement. As we discussed in Section I, our consumption smoothing estimates are valid in the presence of behavioral biases. However, fully accounting for biases would require incorporating an additional internality effect into the first-order welfare effects of a pension reform (equation (3)). Passive saving tends to mute behavioral responses to incentives [Chetty et al., 2014], which in turn would mute internality effects for the case of under-saving. Other behavioral frictions – like inattention to financial incentives or reliance on statutory retirement ages [Seibold, 2021] – may also distort retirement decisions and mute labor supply responses to reforms to the pension benefit schedule. This could dampen the fiscal externality effect, as in Chetty, Looney and Kroft [2009] and/or introduce internality effects relevant for the evaluation of pension reforms [Reck and Seibold, 2021].

We focused on evaluating reforms along a single dimension, but our overall approach can inform pension design across multiple dimensions. First, along each dimension we consider,

³⁷ 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 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.

conditioning on other determinants of pension benefits only modestly reduces consumption differences, suggesting that redistributing along one dimension with a pension reform cannot eliminate the consumption smoothing effects of redistributing along other dimensions. Second, the average change in late-career incentives from Sweden's 1998 reform (Figure 1B) masks significant heterogeneity, especially for low-income, short-career workers affected by the increase in the minimum pension and high-income, long-career workers affected by the increase in the effective maximum pension (see Appendix Figure A-8). In Table H-2, we quantify the consumption smoothing cost of steeper profiles for different subgroups of individuals. Consumption differences by retirement age are less pronounced for individuals in the bottom decile of pension rights accumulated by age 55 (see Appendix Figure C-4), where the pension profile is very flat. Hence, strengthening retirement incentives for this group would have been less costly, but the Swedish reform did the opposite by increasing the minimum pension. Additionally, the flatter consumption gradient for couples compared to singles implies lower consumption smoothing costs of steeper incentives for couples, presumably because couples can rely on intra-household insurance.

VII 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 find 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 and drops in consumption around retirement, as well as supplementary analysis of patterns of selection into early retirement and the composition of consumption. 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, future work could examine retirement consumption in other countries. Our first pass at doing so suggests that consumption differences by retirement age are larger in the US, perhaps because the relative generosity of the social insurance system in Sweden reduces individuals' exposure to work longevity risk. Relatedly, one could study the optimal design of pension and other social insurance programs jointly, accounting for the sometimes fuzzy boundaries between programs (Inderbitzin, Staubli and Zweimüller [2016]). Second, as we briefly 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. 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. Finally, 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 directly (e.g., [Gruber, Kanninen and Ravaska \[2022\]](#), [Seibold \[2021\]](#), [Reck and Seibold \[2021\]](#)).

References

- Aguiar, Mark, and Erik Hurst.** 2005. "Consumption versus Expenditure." *Journal of Political Economy*, 113(5): pp. 919–948.
- Aguiar, Mark, and Erik Hurst.** 2013. "Deconstructing Life Cycle Expenditure." *Journal of Political Economy*, 121(3): 437–492.
- Andersen, Asger L., Niels Johannesen, and Adam Sheridan.** 2021. "Dynamic Spending Responses to Wealth Shocks: Evidence from Quasi-lotteries on the Stock Market." Center for Economic Policy Research Working Paper.
- Andrews, Isaiah, and Conrad Miller.** 2013. "Optimal social insurance with heterogeneity." *Working Paper*.
- Atkinson, A.B., and J.E. Stiglitz.** 1976. "The design of tax structure: Direct versus indirect taxation." *Journal of Public Economics*, 6(1): 55–75.
- Baily, Martin Neil.** 1978. "Some Aspects of Optimal Unemployment Insurance." *Journal of Public Economics*, 10(3): 379–402.
- Baker, Scott, Lorenz Kueng, Steffen Meyer, and Michaela Pagel.** forthcoming. "Consumption Imputation Errors in Administrative Data." *Review of Financial Studies*.
- Banks, James, Richard Blundell, and Sarah Tanner.** 1998. "Is There a Retirement-Savings Puzzle?" *The American Economic Review*, 88(4): 769–788.
- Barr, Nicholas, and Peter Diamond.** 2009. "Reforming Pensions: Principles, Analytical Errors and Policy Directions." *International Social Security Review*, 62(2): 5–29.
- Battistin, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber.** 2009. "The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach." *The American Economic Review*, 99(5): 2209–2226.
- Becker, Gary S., Tomas J. Philipson, and Rodrigo R. Soares.** 2005. "The Quantity and Quality of Life and the Evolution of World Inequality." *American Economic Review*, 95(1): 277–291.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg.** 2001. "What Accounts for the Variation in Retirement Wealth Among U.S. Households?" *The American Economic Review*, 91(4): 832–857.
- Blundell, Richard, Eric French, and Gemma Charlotte Tetlow.** 2016. "Retirement Incentives and Labor Supply." In *Handbook of the Economics of Population Aging*. Vol. 1, , ed. John Piggot and Alan Woodland, Chapter 0, 457–566. Elsevier.
- Blundell, Richard, Jack Britton, Monica Costa Dias, and Eric French.** 2016b. "Cognition, Health, Employment and Wages near Retirement in England." Mimeo.
- Blundell, Richard, Jack Britton, Monica Costa Dias, and Eric French.** 2021. "The Impact of Health on Labor Supply near Retirement." *Journal of Human Resources*.
- Bound, John.** 1991. "Self-Reported Versus Objective Measures of Health in Retirement Models." *The Journal of Human Resources*, 26(1): 106–138.
- Bozio, Antoine, Clémentine Garrouste, and Elsa Perdrix.** 2021. "Impact of Later Retirement on Mortality: Evidence from France." *Health Economics*, 30: 1178–1199.

- Butler, J. S., Richard V. Burkhauser, Jean M. Mitchell, and Theodore P. Pincus.** 1987. "Measurement Error in Self-Reported Health Variables." *The Review of Economics and Statistics*, 69(4): 644–650.
- Cesarini, David, Erik Lindqvist, Robert Östling, and Björn Brügemann.** 2016. "Wealth, Health, and Child Development: Evidence from Administrative Data on Swedish Lottery Players." *The Quarterly Journal of Economics*, 131(2): 687–738.
- Chetty, Raj.** 2006. "A general formula for the optimal level of social insurance." *Journal of Public Economics*, 90(10-11): 1879–1901.
- Chetty, Raj, Adam Looney, and Kory Kroft.** 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review*, 99(4): 1145–77.
- Chetty, Raj, and Amy Finkelstein.** 2013. "Social insurance: Connecting theory to data." In *Handbook of Public Economics*. Vol. 5, 111–193. Elsevier.
- Chetty, Raj, John N Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen.** 2014. "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark." *Quarterly Journal of Economics*, 129(3): 1141–1219.
- Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler.** 2016. "The Association Between Income and Life Expectancy in the United States, 2001–2014." *JAMA*, 315(16): 1750–1766.
- De Nardi, Mariacristina, Eric French, and John Bailey Jones.** 2016. "Savings After Retirement: A Survey." *Annual Review of Economics*, 8: 177–204.
- Deshpande, Manasi, and Lee M. Lockwood.** 2022. "Beyond Health: Nonhealth Risk and the Value of Disability Insurance." *Econometrica*, 90(4): 1781–1810.
- Diamond, Peter, and James Mirrlees.** 1978. "A Model of Social Insurance with Variable Retirement." *Journal of Public Economics*, 10(3): 295–336.
- Diamond, Peter, and James Mirrlees.** 1982. "Social Insurance with Variable Retirement and Private Saving." MIT Working Papers 296.
- Diamond, Peter, and James Mirrlees.** 1986. "Payroll-Tax Financed Social Security with Variable Retirement." *Scandinavian Journal of Economics*, 88: 25–30.
- Di Maggio, Marco, Amir Kermani, and Kaveh Majlesi.** 2020. "Stock Market Returns and Consumption." *Journal of Finance*, 75(6): 3175–3219.
- Fadlon, Itzik, and Torben Heien Nielsen.** 2019. "Household labor supply and the gains from social insurance." *Journal of Public Economics*, 171: 18–28. Trans-Atlantic Public Economics Seminar 2016.
- Finkelstein, Amy, Erzo F. P. Luttmer, and Matthew J. Notowidigdo.** 2013. "What Good is Wealth Without Health? The Effect of Health on the Marginal Utility of Consumption." *Journal of the European Economic Association*, 11(1): 221–258.
- Fitzpatrick, Maria, and Timothy J. Moore.** 2018. "The Mortality Effects of Retirement: Evidence from Social Security Eligibility at age 62." *Journal of Public Economics*, 157(C): 121–137.
- French, Eric.** 2005. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement." *Journal of Public Economics*, 89(1): 1–28.

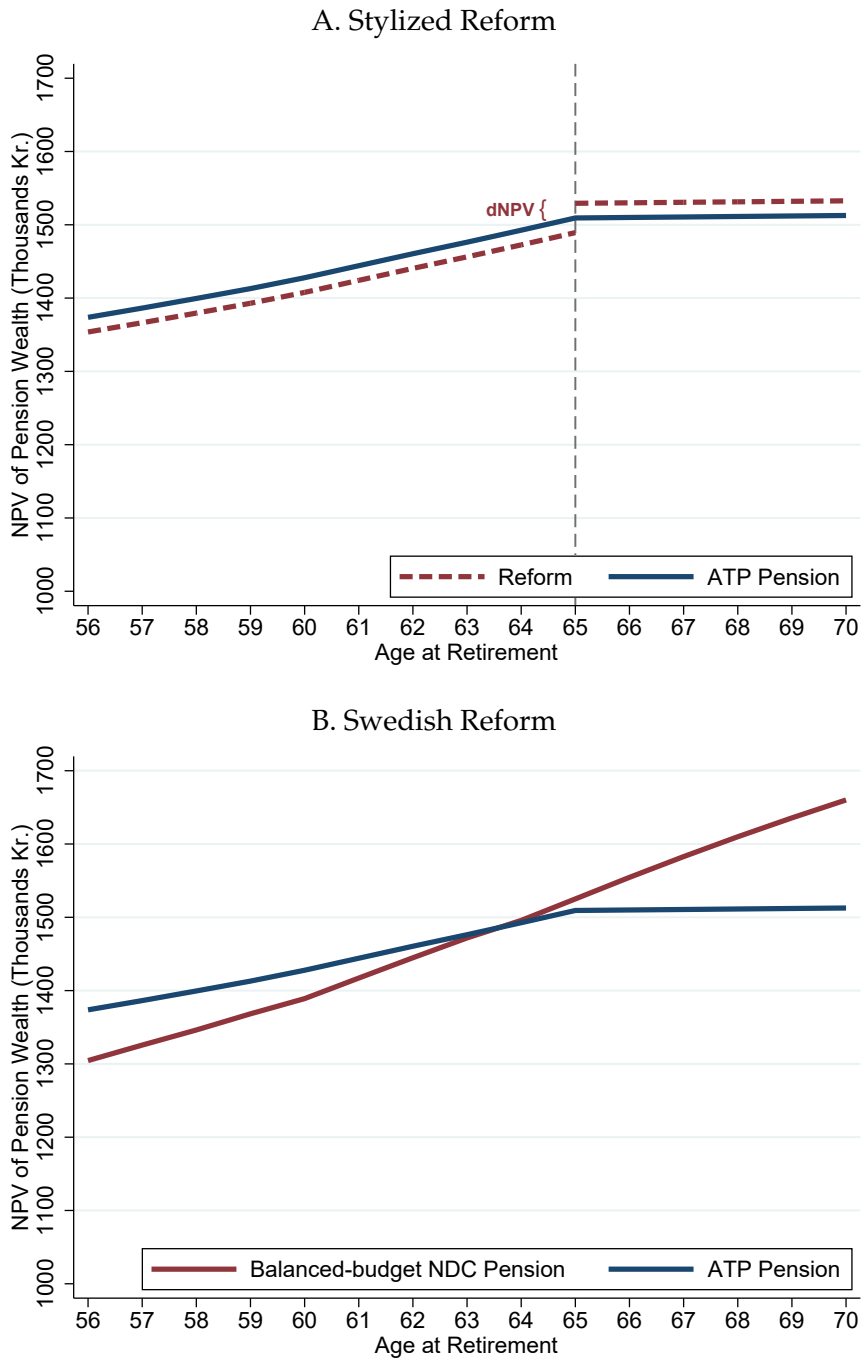
- ment Behaviour." *The Review of Economic Studies*, 72(2): 395–427.
- French, Eric, and John Bailey Jones.** 2011. "The Effects of Health Insurance and Self-Insurance on Retirement Behavior." *Econometrica*, 79(3): 693–732.
- French, Eric, and John Bailey Jones.** 2012. "Public Pensions and Labor Supply Over the Life Cycle." *International Tax and Public Finance*, 19: 268–287.
- French, Eric, and John Bailey Jones.** 2017. "Health, Health Insurance, and Retirement: A Survey." *Annual Review of Economics*, 9: 383–409.
- French, Eric, Attila Lindner, Cormac O’Dea, and Tom Zawisza.** 2022. "Labor Supply and the Pension Contribution-Benefit Link." Working Paper.
- Giupponi, Giulia.** 2019. "When Income Effects are Large: Labor Supply Responses and the Value of Welfare Transfers." CEP Discussion Paper No. 1651.
- Gruber, Jonathan.** 1997. "The consumption smoothing benefits of unemployment insurance." *The American Economic Review*, 87(1): 192–205.
- Gruber, Jonathan, and David A. Wise.** 1999. "Introduction to "Social Security and Retirement around the World"." In *Social Security and Retirement around the World. International Social Security*, , ed. Jonathan Gruber and David A. Wise. University of Chicago Press.
- Gruber, Jonathan, Ohto Kanninen, and Terhi Ravaska.** 2022. "Relabeling, retirement and regret." *Journal of Public Economics*, 211: 104677.
- Gustman, Alan, and Thomas L Steinmeier.** 2000. "Retirement in Dual-Career Families: A Structural Model." *Journal of Labor Economics*, 18(3): 503–45.
- Gustman, Alan, and Thomas L Steinmeier.** 2018. "The Role of Health in Retirement." *Research in Labor Economics*, 46: 229–297.
- Hagen, Johannes.** 2015. "The Determinants of Annuitization: Evidence from Sweden." *International Tax and Public Finance*, 22(4): 549–578.
- Haller, Andreas.** 2022. "Welfare Effects of Pension Reforms." Working Paper.
- Hendren, Nathaniel.** 2017. "Knowledge of future job loss and implications for unemployment insurance." *American Economic Review*, 107(7): 1778–1823.
- Hendren, Nathaniel.** 2020. "Measuring Ex Ante Welfare in Insurance Markets." *The Review of Economic Studies*, 88(3): 1193–1223.
- Hendren, Nathaniel, and Ben Sprung-Keyser.** 2020. "A Unified Welfare Analysis of Government Policies." *Quarterly Journal of Economics*, 135(3): 1209–1318.
- Hospido, Laura, and Gema Zamarro.** 2014. "Retirement Patterns of Couples in Europe." *IZA Journal of European Labor Studies*, 3(1): 1–18.
- Inderbitzin, Lukas, Stefan Staubli, and Josef Zweimüller.** 2016. "Extended Unemployment Benefits and Early Retirement: Program Complementarity and Program Substitution." *American Economic Journal: Economic Policy*, 8(1): 253–288.
- Jappelli, Tullio, and Luigi Pistaferri.** 2010. "The Consumption Response to Income Changes." *Annual Review of Economics*, 2(1): 479–506.
- Jappelli, Tullio, and Luigi Pistaferri.** 2017. *The Economics of Consumption: Theory and Evidence.*

OUP USA.

- Jönsson, Lisa, Mårten Palme, and Ingemar Svensson.** 2012. "Disability Insurance, Population Health, and Employment in Sweden." In *Social Security Programs and Retirement around the World.*, ed. David Wise, 79–217. Chicago:University of Chicago Press.
- Karlstrom, Anders, Marten Palme, and Ingemar Svensson.** 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, Jonas, Camille Landais, and Johannes Spinnewijn.** 2020. "The Value of Registry Data for Consumption Analysis: An Application to Health Shocks." *Journal of Public Economics*, 189: 1040–1088.
- Kolsrud, Jonas, Camille Landais, Peter Nilsson, and Johannes Spinnewijn.** 2018. "The optimal timing of unemployment benefits: Theory and evidence from Sweden." *American Economic Review*, 108(4-5): 985–1033.
- Kuhn, Andreas, Stefan Staubli, Jean Philippe Wuellrich, and Josef Zweimüller.** 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, Rafael, Arvind Magesan, and Stefan Staubli.** 2022. "How Social Security Reform Affects Retirement and Pension Claiming." *American Economic Journal: Economic Policy.* Forthcoming.
- Landais, Camille, and Johannes Spinnewijn.** 2021. "The Value of Unemployment Insurance." *Review of Economic Studies*, 6(88): 3041–3085.
- Laun, Lisa.** 2017. "The Effect of Age-Targeted Tax Credits on Labor Force Participation of Older Workers." *Journal of Public Economics*, 152(C)(102-118).
- Laun, Tobias, and Johanna Wallenius.** 2015. "A Life Cycle Model of Health and Retirement: The Case of Swedish Pension Reform." *Journal of Public Economics*, 127(C): 127–136.
- Lockwood, Lee M.** 2018. "Incidental Bequests and the Choice to Self-Insure Late-Life Risks." *American Economic Review*, 108(9): 2513–50.
- Manoli, Dayanand S., and Andrea Weber.** 2016. "Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions." *American Economic Journal: Economic Policy*, 8(4): 160–182.
- Manoli, Dayanand S., and Andrea Weber.** 2018. "The Effects of the Early Retirement Age on Retirement Decisions." NBER Working Paper No. 22561.
- Mullainathan, Sendhil, Joshua Schwartzstein, and William J Congdon.** 2012. "A Reduced-Form Approach to Behavioral Public Finance." *Annual Review of Economics*, 4(1): 511–540.
- Ndiaye, Abdoulaye.** 2020. "Flexible Retirement and Optimal Taxation." Working Paper.
- O’Dea, Cormac.** 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.* Paris:OECD Publishing.


- OECD.** 2017. "Recent Pensions Reforms." In *Pensions at a Glance 2017: OECD and G20 Indicators*. Paris:OECD Publishing.
- OECD.** 2019. *Pensions at a Glance 2019: OECD and G20 indicators*. Paris:OECD Publishing.
- Reck, Daniel, and Arthur Seibold.** 2021. "The Welfare Economics of Reference Dependence." Working Paper.
- Seibold, Arthur.** 2021. "Reference Points for Retirement Behavior: Evidence from German Pension Discontinuities." *American Economic Review*, 111(4): 1126–25.
- Spinnewijn, Johannes.** 2015. "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs." *Journal of the European Economic Association*, 13(1): 130–167.
- Staubli, Stefan, and Josef Zweimüller.** 2013. "Does Raising the Early Retirement Age Increase Employment of Older Workers?" *Journal of Public Economics*, 108(C): 17–32.
- Stephens, Melvin, and Desmond Toohey.** 2018. "Changes in Nutrient Intake at Retirement." NBER Working Paper No. w24621.
- Stern, Steven.** 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 ($56 \leq r \leq 59$), early retirees ($60 \leq r \leq 63$), normal retirees ($64 \leq r \leq 65$) and late retirees ($66 \leq r \leq 69$). Normal retirees are the reference category. The graph reports, for all retirement age groups, the estimated coefficients α_j from specification (8), scaled by $E_j[\tilde{C}_{it}]$, the average level of consumption of individuals who retire between 64 and 65 from the same cohort, age, family composition, income decile and career length at 55 group as the average individual retiring in age group j . We start, on the left hand side of the graph, with results from model (8) where only year and age fixed effects are included. The rest of figure shows the same estimated coefficients when sequentially adding controls for family composition, within-cohort deciles of average income between ages 52 and 55 and group of career length at 55 in the vector of controls \mathbf{X} .

Figure 4: CONSUMPTION DYNAMICS AROUND RETIREMENT, BY RETIREMENT AGE GROUP



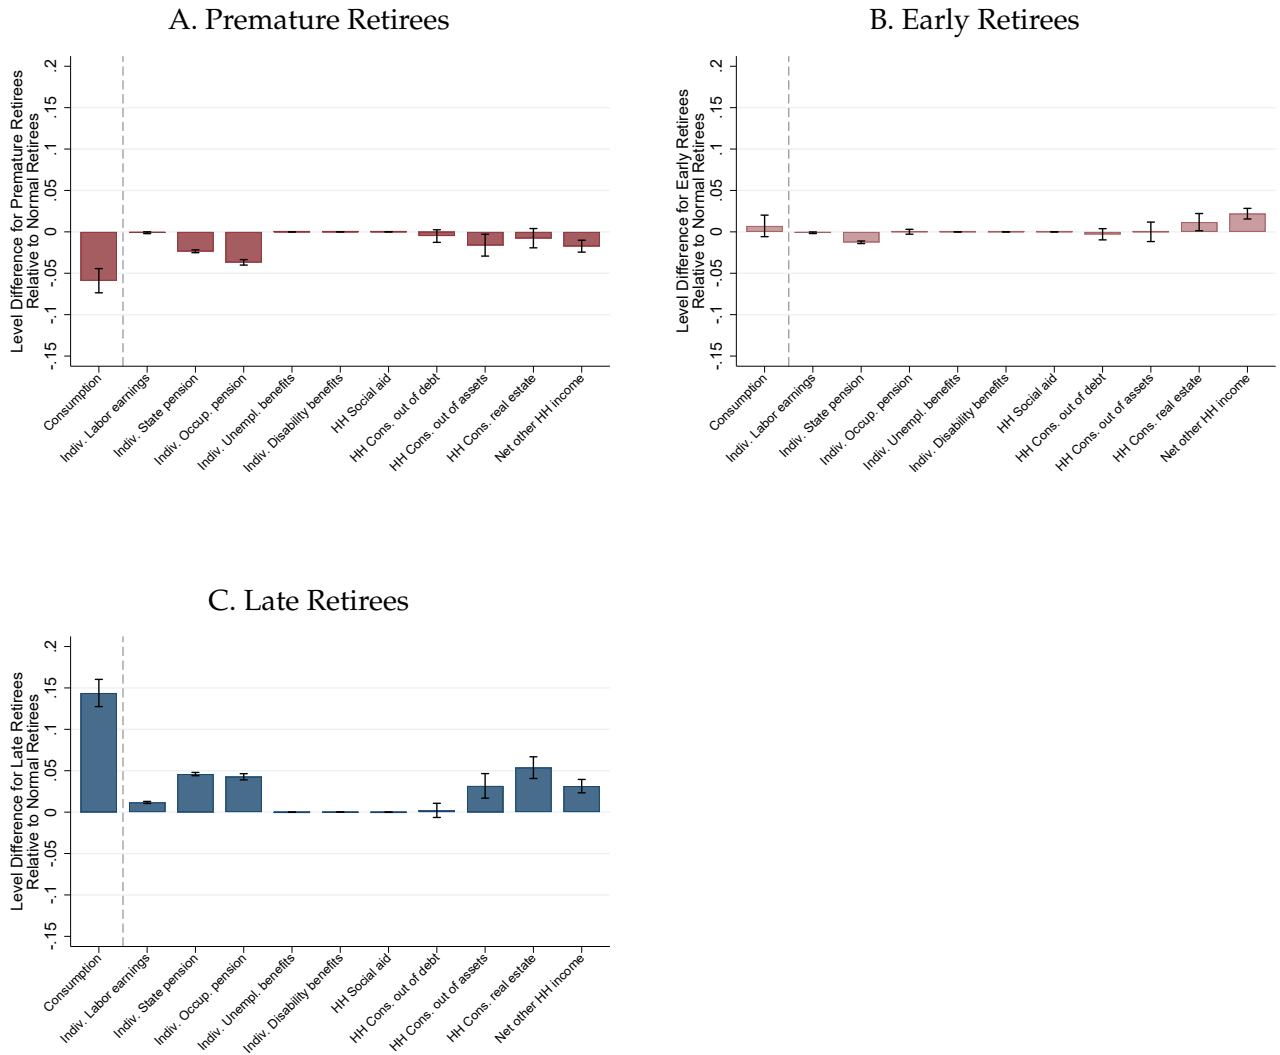
Notes: The figure documents consumption dynamics around retirement. In both panels, household consumption is first residualized on a set of cohort fixed effects and age fixed effects and household structure controls, as in specification (8). Panel A plots average residualized consumption as a function 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 5: CONSUMPTION DIFFERENCES AT AGE 55 ACROSS RETIREMENT AGE GROUPS



Notes: The figure documents how consumption at age 55 differs across individuals who retire at different ages. The sample comprises all individuals from cohorts 1945 to 1949 who are 55 and still working at the time their consumption is observed. Note that we use different cohorts relative to 3 because we only see consumption at age 55 for individuals born after 1945. Individuals are grouped into four retirement age categories: premature retirees ($56 \leq r \leq 59$), early retirees ($60 \leq r \leq 63$), normal retirees ($64 \leq r \leq 65$) and late retirees ($66 \leq r \leq 69$). Normal retirees are the reference category. The graph reports, for all retirement age groups, the estimated coefficients α_j from specification (8), scaled by $E_j[\hat{C}_{it}]$, the average level of consumption of individuals who retire between 64 and 65 from the same cohort, age, family composition, income decile and career length at 55 group as the average individual retiring in age group j . We start, on the left hand side of the graph, with results from model (8) where only year and age fixed effects are included. The rest of figure shows the same estimated coefficients when sequentially adding controls for family composition, within-cohort deciles of average income between ages 52 and 55 and group of career length at 55 in the vector of controls X .


Figure 6: 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 (8) 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 $E_j[\tilde{C}_{it}]$, 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 7: HETEROGENEITY & SELECTION INTO RETIREMENT AGE

A. Socio-Demographic Characteristics

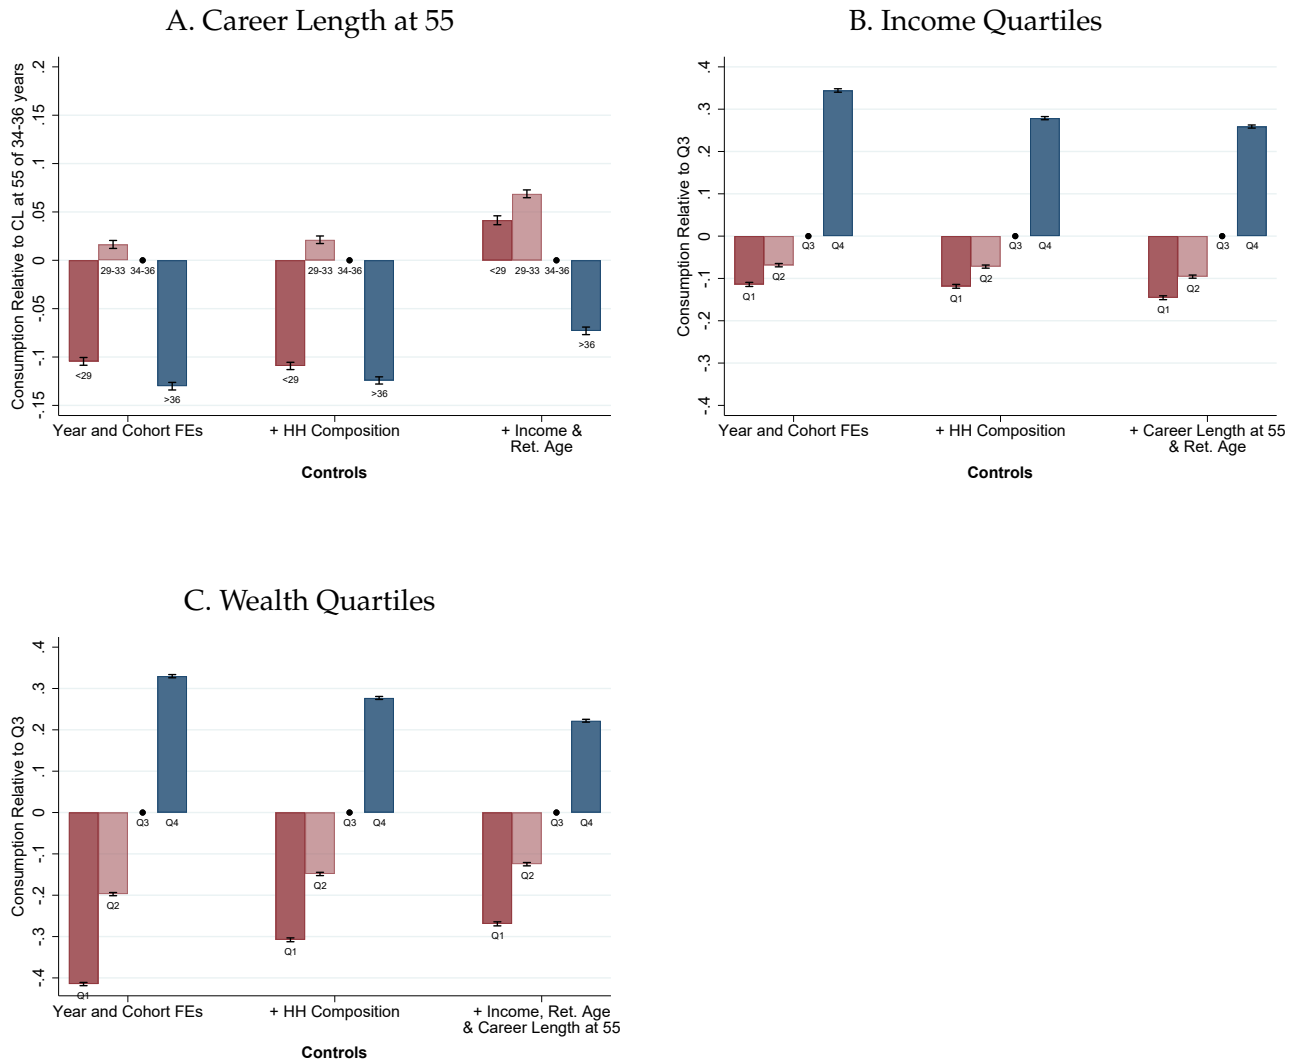


B. Health and Life Expectancy



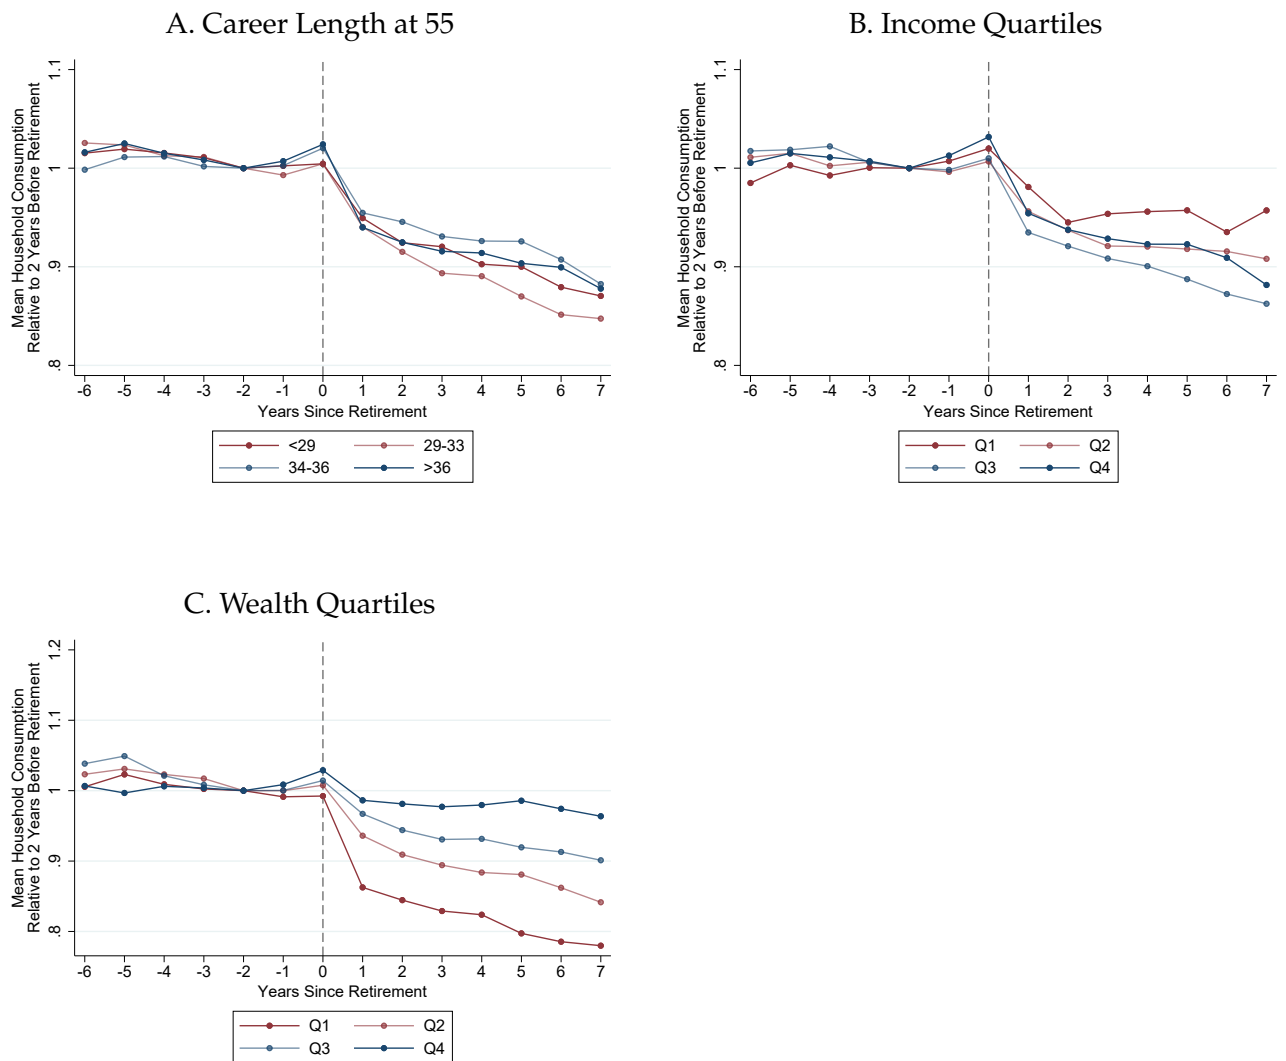
Notes: The figure documents patterns of heterogeneity across retirement age groups. Panel A displays estimates from a multinomial logit prediction model for retiring in one of the 4 different age groups. The regression sample includes one observation for each of the 418,033 unique individuals of our 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 (8) (with cohort and age fixed effects and controls for family structure). We replace consumption by our two indices for bad health (i.e. standardized principal components extracted from all health outcomes in the HEK and ULF surveys; 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 8: 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 (8), scaled by $E_j[\bar{C}_{it}]$, 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 8, we begin on the left-hand side with estimates from model (8) 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 9: CONSUMPTION DYNAMICS AROUND RETIREMENT ACROSS ALTERNATIVE POLICY DIMENSIONS



Notes: The figure documents consumption dynamics around retirement across other policy dimensions. As in Figure 4, household consumption is first residualized on a set of cohort fixed effects, age fixed effects and household structure controls, following specification (8). 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 10: 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 (10). 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 consumption levels implementation,

$$(11) \quad \frac{SMU_{r \leq \bar{r}} - SMU_{r > \bar{r}}}{SMU_{NRA}} \approx \gamma \times \left[\frac{E_{r > \bar{r}}(c)}{E_{r \in NRA}(c)} - \frac{E_{r \leq \bar{r}}(c)}{E_{r \in NRA}(c)} \right],$$

where the differences in consumption levels are based on estimates in regression (8) 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 (≥ 1)	17.65 %	(38.12)
Kid at Home Under 18 (≥ 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)
<hr/>		
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

	Consumption levels	Consumption drops
	C	ΔC
	$\gamma = 4, \omega = 1$	$\gamma = 4$
	(1)	(2)
A. Retirement Age		
Premature \rightarrow Early-Late	0.34	0.21
Prem.-Early \rightarrow Normal-Late	0.28	0.12
Prem.-Normal \rightarrow Late	0.76	0.14
B. Career Length at 55		
Q1 \rightarrow Q2-Q4	0.27	0.01
Q1-Q2 \rightarrow Q3-Q4	-0.11	0.04
Q1-Q3 \rightarrow Q4	-0.37	-0.01
C. Income		
Q1 \rightarrow Q2-Q4	0.74	-0.04
Q1-Q2 \rightarrow Q3-Q4	0.89	-0.03
Q1-Q3 \rightarrow Q4	1.32	0.02
D. Wealth		
Q1 \rightarrow Q2-Q4	1.40	0.33
Q1-Q2 \rightarrow Q3-Q4	1.45	0.27
Q1-Q3 \rightarrow Q4	1.69	0.26

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 10. 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 implementation using the difference in consumption levels to approximate the difference in SMU 's (see equation (6)). 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 (7)). [Appendix H](#) provides more details underlying the welfare calculations.