

# CHAPTER 5

## From pension reforms to welfare: A unifying framework

**Camille Landais and Johannes Spinnewijn**

London School of Economics and CEPR

The evaluation of pension reforms has been the focus of considerable academic and policy discussion over recent decades. However, much of this discourse has concentrated on improving fiscal balance and examining labour supply responses to the pension system. Pension reforms have been primarily framed as responses to fiscal challenges, such as the increasing cost pressures from ageing populations (discussed in Chapter 1) and the goal of encouraging later retirement (as explored in Chapters 3 and 4). While these considerations are crucial, they overlook a fundamental aspect of pension systems, namely, their role in providing insurance against the inability to generate income in old age and their potential to redistribute based on lifetime income. Despite a rich literature on pension systems and retirement behaviour, a coherent and simple framework to evaluate the welfare effects of reforms, particularly their redistributive and insurance consequences, remains largely absent.

In this chapter, we present a new approach, based on the conceptual framework and empirical work of Kolsrud et al. (2024), that allows for a simple and transparent evaluation of the welfare implications of pension reforms. The central tenet of this approach is to treat pensions similarly to other social insurance or tax/benefit programmes. How are these programmes comparable to pension systems? Social insurance programmes balance providing insurance against adverse events (e.g., unemployment or health shocks) with maintaining incentives to avoid these events or mitigate their impact. Tax/benefit programmes aim to redistribute from high-income to low-income individuals while preserving the incentive to increase income. In a similar spirit, pension systems serve a crucial role of insurance: they protect individuals against uncertainties regarding work capacity, career length and success, as well as longevity. Pension systems often also serve an important redistributive role, redistributing from individuals with high life-time earnings to individuals with low life-time earnings. But providing insurance and redistribution through the pension system comes at a cost: generous insurance against early retirement can diminish incentives to remain in the workforce, and redistribution based on individuals' life-time earnings reduces the incentive to generate these earnings. In other words, pension systems, like any other social insurance or transfer programmes, need to balance the cost of distorting incentives with the benefits of providing insurance

and redistribution. The core novelty of the approach is to show that both sides of the trade-off – costs and benefits – can be measured precisely and can thus be closely compared.

The framework we propose allows policymakers to analyse how reforms redistribute resources and impact individuals' ability to smooth consumption. By incorporating standard welfare economics principles, the framework captures the nuanced trade-offs between fiscal sustainability, work incentives, and individual welfare. Its transparency and direct applicability to data make it particularly appealing: simple empirical data, such as consumption measures, can be used to evaluate the real-world implications of pension reforms. This not only grounds the analysis in observable outcomes but also enables policymakers to understand the redistributive and insurance implications of their reforms. The framework's transparency makes it accessible and intuitive, allowing for clearer communication of policy trade-offs and outcomes.

Taking the framework to the data, in this chapter we demonstrate that reforms in many European countries have been regressive, shifting resources from individuals with fewer resources and poorer health to those who are better off. These regressive outcomes underscore the crucial trade-offs between fiscal sustainability and equity, which should be central in any pension reform debate.

## A UNIFIED FRAMEWORK

The proposed framework, developed in Kolsrud et al. (2024), considers the pension system as any other social insurance or transfer programme, and its direct goal is to characterise the welfare consequences of reforms, like those presented in Chapter 2.<sup>1</sup> The framework shows transparently that these reforms entail a trade-off between fiscal implications and redistributive and insurance effects. We can study the two sides of the trade-off separately – fiscal effects on the one hand and redistributive and insurance effects on the other hand – and then translate them into magnitudes that can be compared.

A practical advantage of the proposed approach is that it focuses on reforms to the actual pension profile. Doing so, we move away from the widespread but often implicit vision that the ideal pension system is one that guarantees actuarial fairness, or even stronger, that provides a savings vehicle where individuals get back in pension benefits when retired what they have contributed to the system throughout their lifetime. In practice, pension systems are far from that, and the reason for this is that most systems do provide a lot of insurance and redistribution. Knowing how to evaluate the benefits provided by such insurance and redistribution is at the heart of our framework.

1 Note that Haller (2022) develops a similar framework, closely related to ours.

## Fiscal effects

Pension reforms alter the incentives for individuals to work, save, and retire. These behavioural changes, in turn, impact government fiscal outcomes by affecting tax revenues and pension outlays. These so-called fiscal externalities can be easily measured: they are fully determined by behavioural responses on the one hand, and parameters of the tax and social insurance system on the other.

Take, for instance, a reform that increases the steepness of the pension profile (Figure 1), an element typical of most reforms undertaken by European countries over the past 20 years as discussed in Chapter 2.<sup>2</sup> Such a reform enhances incentives to work longer: the returns, in terms of future pension benefits, of retiring later have increased. If individuals respond to these incentives by retiring later, this will tend to have positive fiscal externalities, coming both from increased tax contributions and a reduced time in retirement receiving pension benefits. This is the case illustrated in Figure 1, showing how both early retirees and late retirees respond to the reform by retiring at a later age. In general, the effects of a particular pension reform on saving, labour supply and retirement timing are a priori theoretically ambiguous: these behaviours result from a complex decision-making process, involving time, uncertainty and many (unknown) parameters such as one's health status and life expectancy, work ability, preferences, career opportunities, family situation, asset level, and so on. As a result, it is not always obvious to predict ex ante which reforms will have the largest effects on, for example, labour supply or retirement savings. Still, one can always try to evaluate ex post what the behavioural responses to a specific reform have been in order to quantify its fiscal externalities.<sup>3</sup> We can summarise the overall fiscal externality of a pension reform as the additional euros of fiscal resources generated because of all these behavioural responses, i.e., above and beyond each euro mechanically generated by the reform absent any behavioural response. When retirement age increases strongly in response to a reform, the fiscal externality will be large. When retirement age is quite inelastic to a reform, the fiscal externality will be small, even if the reform cuts pension benefits drastically and mechanically reduces fiscal outlays by a large amount.

A considerable body of literature has explored how changes in incentives affect labour supply and retirement decisions. Generally, reforms that raise the retirement age or encourage longer work have succeeded, with individuals delaying retirement and extending their labour supply. But, as discussed in Chapters 3 and 4, reforms that cut pension benefits the most, or provide the largest financial incentives to retire later, are not always the ones that generate the largest fiscal externality.

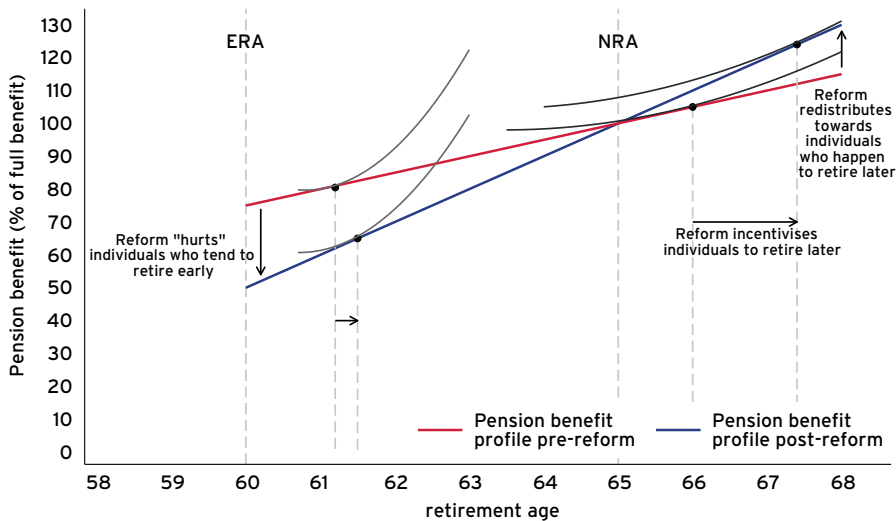
2 The figure also exemplifies the change in pension profile that followed the large 1998 Swedish pension reform analysed in Kolsrud et al. (2024).

3 As discussed in Chapter 7, when examining behavioural responses, it is essential to consider the spillover effects on other social insurance programs too. For instance, raising the minimum age at which individuals can access pension benefits tends to increase applications for disability and unemployment insurance. These spillovers can have a significant impact on the overall fiscal effects of a pension reform.

## Insurance and redistribution

While fiscal considerations are important, pension reforms can also have significant insurance and redistributive effects. In our example of a reform steepening the profile (Figure 1), we see that the reform will hurt individuals who happen to retire early, while it will redistribute towards individuals who retire later. This redistribution would be innocuous from a social welfare perspective if the value of one euro were the same for these two groups of individuals. But it is hardly the case in practice. People who retire early may do so because they want more leisure, but they may also do so due to deteriorating health or reduced earnings capacity. The former may not suffer much from the benefits cuts, but the latter will. They are also often the least able to adjust to reduced pension benefits, forcing them to cut their expenditures the most in response to a pension cut. Conversely, individuals who can continue working and remain productive at older ages may value the extra pension benefits less.<sup>4</sup>

**FIGURE 1 FISCAL EFFECTS VERSUS REDISTRIBUTIVE EFFECTS OF A PENSION REFORM: A GRAPHICAL REPRESENTATION**



Notes: The figure shows the pension benefit profile as a function age at retirement before and after a hypothetical reform strengthening incentives to retire later (i.e. steepening of the profile). It serves to exemplify the fiscal and redistributive effects of a reform, following the framework developed in Kolsrud et al. (2024). To this effect, the graph depicts indifference curves representing graphically the implicit welfare of two individuals, one retiring early, and one retiring late. The stronger curvature of the indifference curves for the early retiree reflects her preference for retiring earlier. This might be due to the fact that her work capacity is lower, due to invalidity or health, but it could also reflect a higher value of leisure, for example due to a shorter life expectancy. Retirement age is determined for each individual by their indifference curves and by the steepness of the pension profile. The figure shows that the reform induces both individuals to work longer and retire later. But it shows that the reform also redistributes away from the early retirees (who end up on a lower indifference curve), and towards later retirees (who end up on a higher indifference curve).

4 Individuals can indeed protect themselves against adverse implications of pension reforms by adjusting their savings and labour supply, as illustrated in Figure 1. However, a key insight from the social insurance literature (e.g., Chetty, 2006) is that to evaluate pension reforms the impact of individuals' behavioural responses on their own welfare can be expected to be small relative to the direct effects of the changes in pension benefits on their welfare. The intuition is that if individuals highly valued saving more or retiring later, they would have done so already, irrespective of the pension reform. This difference in welfare impacts is also apparent in Figure 1, where the direct effect dominates the behavioural response.

To fully assess the welfare effects of pension reforms, it is thus essential to consider who benefits and who loses from the reforms, and to quantify their respective gains and losses. In the hypothetical reform of Figure 1 for instance, we need to evaluate the value of the loss to earlier retirees relative to the gains for late retirees. This ratio can then be compared to the fiscal externality to determine whether steepening the pension profile further is worthwhile. That is, we ultimately want to compare the fiscal gain against the potential welfare loss when insurance and redistribution are reduced.<sup>5</sup>

## USING CONSUMPTION TO EVALUATE WELFARE

How can we precisely measure the welfare gains for those who benefit from a reform and the welfare losses for those who are hurt by the reform? This is where consumption data come in. Consumption is a particularly effective metric for evaluating welfare: it reflects the resources individuals have available, including their income from labour and capital, and not only at the individual but also at the household level. Economic theory suggests that consumption is a key indicator of welfare because it is directly related to individuals' marginal utility. Lower levels of consumption are associated with higher marginal utility, meaning that a given reduction in consumption has a more significant impact on well-being for individuals with lower incomes.

Consumption also allows us to assess individuals' ability to maintain their standard of living over time. For example, if individuals face a health shock and must withdraw prematurely from the labour market, they may need to reduce their consumption more when entering retirement. Therefore, the excess sensitivity of consumption to income shocks (like, for example, entry into retirement) is also a good indicator of the welfare value of pension benefits. The better workers are insured, the smoother their consumption path into and during retirement.

To make meaningful welfare statements about the effects of pension reforms, we can compare the consumption patterns of those who lose from the reform to those who benefit. If those who lose from the reform have fewer resources or are less able to smooth their consumption into retirement, the welfare costs may outweigh the fiscal externality.

Note that while consumption is in itself a powerful measure for welfare evaluation, a comprehensive evaluation of pension reforms requires potentially more information. In particular, it requires paying attention to the differences between winners and losers above and beyond their differences in consumption levels, such as differences in health and life expectancy. Some characteristics can affect how much individuals value the changes in benefits, but other characteristics can affect the social value we attribute to transferring resources between groups. We may find it socially undesirable to transfer

<sup>5</sup> Note that here we are focusing on redistribution *within* a cohort. But the framework can be extended to think similarly about the redistribution and insurance provided across cohorts by the pension system. These intergenerational effects are in particular important as they impact the political feasibility of reforms.

money away from early retirees and towards later retirees, even when they enjoy similar levels of consumption, if the former have worse health and lower life expectancy than the latter. Such heterogeneity in individual characteristics of winners and losers can mitigate or exacerbate the welfare effects of reforms, making it important to account for these factors when analysing consumption data.<sup>6</sup>

## DIFFERENCES IN CONSUMPTION PATTERNS BY RETIREMENT AGE

Kolsrud et al. (2024) illustrate the potential of this framework for evaluating pension systems by studying the 1998 pension reform in Sweden. This hallmark reform had various features, introducing notional adjustments to demographic factors and an overall reduction in benefits despite more generous minimum pensions. Importantly, it also strengthened incentives to retire later by steepening the pension benefit profile. This strengthening occurred at all ages after the early retirement age of 61, and in particular after the normal retirement age of 65, given that the profile was essentially flat beyond this age before the reform.

The reform induced a significant increase in retirement age, therefore creating positive fiscal externalities that, following the logic of our approach, we were able to precisely measure. We calculated that by reallocating one euro from early retirees to late retirees, the government saved an additional 15 cents due to delayed retirement.

So, this reform induced workers to retire later, but we should be concerned about its regressive nature too. Are early retirees hurt more by the decrease in pension benefits than late retirees gain from an increase in pension benefits? To investigate this question, we used detailed administrative data from Sweden and constructed a registry-based measure of household consumption expenditures, available for every Swedish resident over multiple years (see Kolsrud et al., 2020 for details). We compare how consumption patterns differ across workers who retire at different ages.

Panel A of Figure 2 considers individuals' consumption levels at the same age post-retirement and shows how these vary with retirement age. The consumption for each retirement age group is expressed relative to those retiring at age 65, corresponding to the normal retirement age at that time. The overall gradient of consumption with respect to retirement age is steep, with late retirees enjoying over 20% more consumption and premature retirees consuming up to 10% less than normal retirees. This suggests that

6 Differences in preferences over consumption will by definition change individuals' valuation of the changes in pension benefits. Understanding how consumption preferences differ between individuals or change over time is an important challenge for researchers, but one that perhaps has received excessive attention in the literature. A prominent literature starting in the 1990s documented large drops in consumption expenditures around retirement (e.g., Banks et al., 1998; Bernheim et al., 2001; Aguiar and Hurst, 2005; Stephens and Toohey, 2018). People's preferences for consumption may indeed differ when retired, and their expenditure patterns will vary due to the simple fact that some work-related expenditures are no longer due, while having more time for leisure and other activities. However, it seems too simplistic to always attribute differences in consumption patterns to differences in preferences. The literature has developed new methods to explore potential differences in preferences (e.g., Chetty, 2008; Landais and Spinnewijn, 2021), but simply studying other observable characteristics that may correlate with different consumption patterns can also be insightful.

rewarding later retirement by giving later retirees more generous pensions redistributes from low-consumption to high-consumption households. If we believe that the corresponding welfare effects are inversely related to the difference in consumption, our welfare estimates suggest a loss of just above 15 cents per euro taken from early retirees to late retirees. However, standard practice in economics is to scale this by a multiple of that, depending on how risk- or inequity-averse one is. This thus outweighs the fiscal gains from the reforms, and suggests that the overall increase in the slope has been welfare-decreasing.

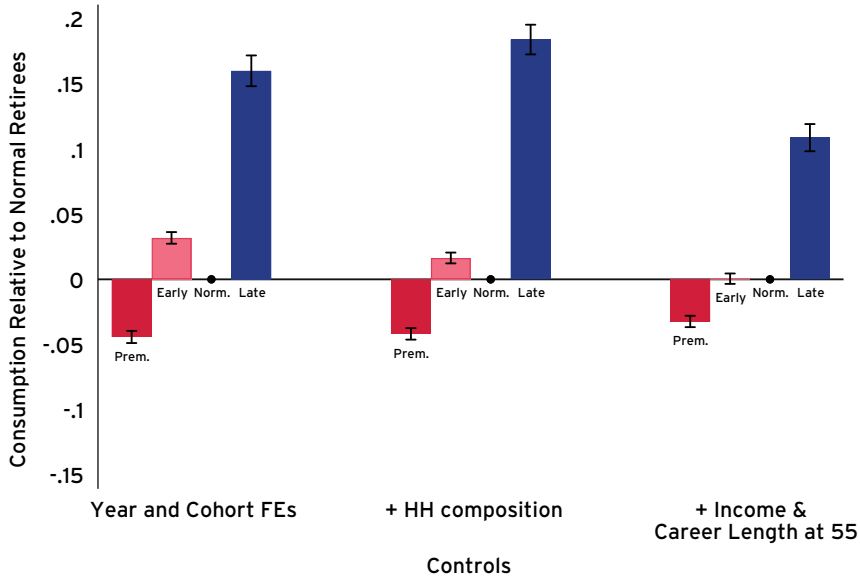
While the overall consumption gradient between retirees at ages 55 to 70 is large and positive, a closer look at the data reveals a non-monotonic relationship between retirement age and consumption. Specifically, individuals retiring between the ages of 60 and 63 have similar or higher consumption on average compared to those retiring near the normal retirement age of 65. This suggests that in this small age range, rewarding later retirement comes at no regressive cost.

The differences in consumption by retirement age become most pronounced in the years immediately surrounding retirement. Leveraging the longitudinal data, we can study how household consumption changes in the years around retirement, as shown in Panel B of Figure 2. All consumption levels are expressed relative to the levels two years before retirement. We estimate that the consumption drop from two years before retirement to two to five years after retirement is much larger for very early retirees compared to later retirees. These findings suggest that individuals who retire early are less able to maintain their standard of living throughout retirement.<sup>7</sup> In the paper, we also examine how marginal propensities to consume out of wealth shocks differ across retirement age groups (Landaïs and Spinnewijn, 2021). These results further confirm our overall finding that incentivising later retirement entails a substantial cost because it takes resources away when the marginal utility of consumption is high and provides more resources when it is low.

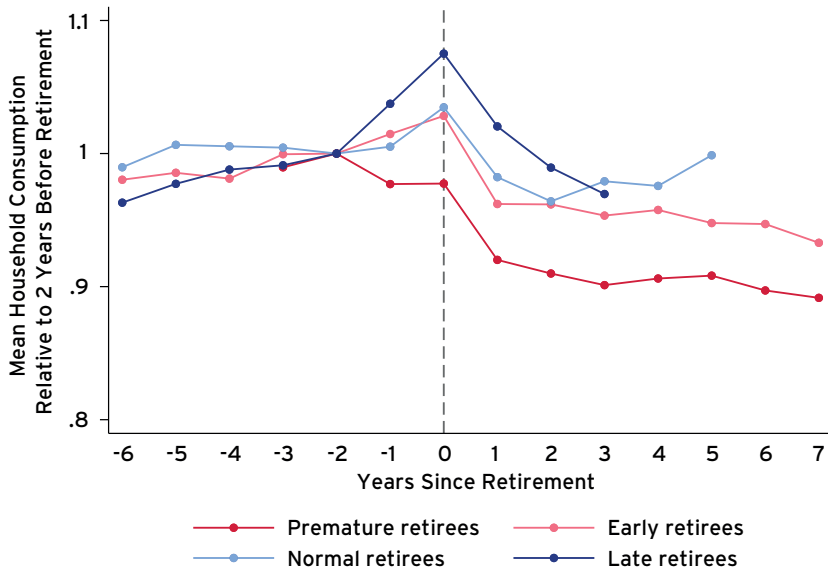
7 Interestingly though, we find a very similar drop in consumption expenditures exactly in the year of retirement for all groups, arguably reflecting differences in consumption expenditures specific to employment versus retirement.

**FIGURE 2 CONSUMPTION LEVELS AND CONSUMPTION DYNAMICS AROUND RETIREMENT, BY RETIREMENT AGE GROUP: SWEDEN**

a) Consumption levels in retirement



b) Consumption profiles: Event studies around retirement



Notes: The figure documents how consumption in retirement differs across individuals who retire at different ages. 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. In panel A, the graph reports, for all retirement age groups, the estimated average level of consumption at age 68 of individuals in that group, controlling for cohort, age, family composition, income decile and career length at 55. Panel B documents consumption dynamics around retirement. The graph plots average residualised 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.



## GOING BEYOND CONSUMPTION

The impact of pension reforms is not uniform across individuals, and there is significant heterogeneity in how different groups are affected. One key source of heterogeneity is health. Individuals who retire earlier tend to be in worse health, and there is limited evidence to suggest that they value pension benefits less than their healthier counterparts. This is illustrated in Figure 3. In fact, the data show that individuals who leave the labour market prematurely do so more often following a deterioration in health. These results further raise concerns about the equity of pension reforms that disproportionately affect early retirees. Those retiring prematurely often face a double burden of reduced income and worse health. Conversely, those retiring late are in better health and have substantially longer life expectancy too. From a lifetime perspective, the pension system is already more generous to them, and rewards for late retirement further increase this imbalance – an issue discussed in greater detail in Chapter 6.<sup>8</sup>

Wealth is another crucial factor that influences the impact of pension reforms. Individuals with higher wealth tend to have more stable consumption levels throughout retirement, as they can draw on their assets to smooth consumption. In contrast, those with lower wealth experience larger drops in consumption upon retiring, making them more vulnerable to the negative effects of pension reforms. Turning back to the differences by retirement age, we find that those retiring between the ages of 61 and 65, where the non-monotonicity appears in Figure 2, have higher household assets on average and tend to be in households where another member earns a significant income. Hence, this group is a complex mix of individuals who are forced to retire early – due to worsening health – and individuals who can afford to retire early – due to accumulated wealth and household resources. This highlights the importance of considering the wealth distribution, also at the household level, when evaluating the equity of pension reforms. It also indicates that asset tests could be effective instruments to complement the penalties that discourage early retirement in recent pension reforms. They would allow these penalties to be avoided when their welfare cost is highest.

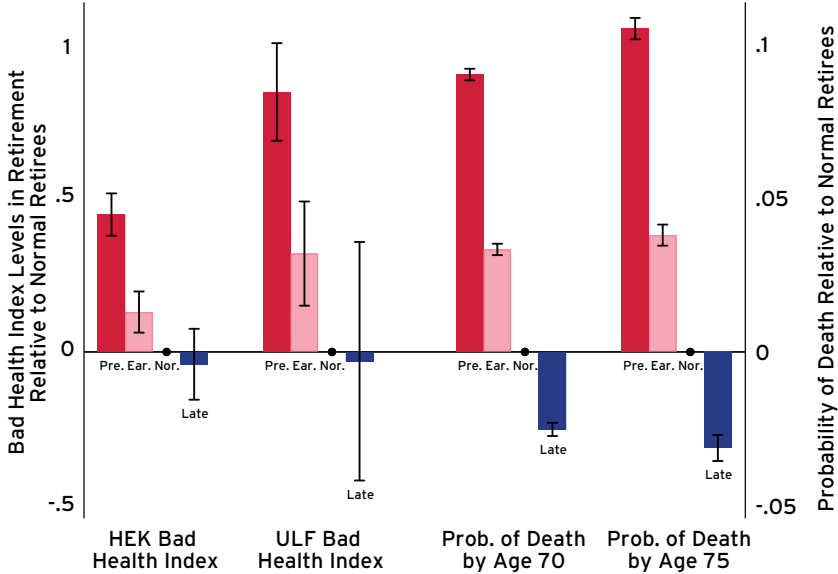
## SUMMARY OF POLICY IMPLICATIONS

The analysis presented in this chapter underscores the importance of considering both the fiscal and welfare implications of pension reforms. While ensuring the fiscal sustainability of pension systems is undoubtedly important, policymakers must also be mindful of the

<sup>8</sup> As discussed, the level of consumption is not the sole determinant of individual's valuation of changes in pension benefits either. However, by studying other observable characteristics in the Swedish context, we broadly reinforced our finding that a steeper pension benefit profile redistributes from those with a high value of pension benefits (earlier retirees) to those with a smaller value (later retirees). In particular, later retirees tend to have more education, more productive careers, and also more financial resources than those retiring very early. At the same time, we also find that the expenditure patterns were broadly similar across these groups and also change in similar ways at retirement, indicating that differences in consumption preferences are not important.

distributional consequences of these reforms. Reforms that disproportionately affect low-income or unhealthy individuals raise serious concerns about equity, as these individuals are often the least able to adjust to reduced benefits or longer working lives.

**FIGURE 3 HEALTH AND LIFE EXPECTANCY BY RETIREMENT AGE: SWEDEN**



Notes: The figure shows that individuals who retire early have much worse health and life expectancy than individuals who retire later. The evidence is drawn from Sweden and compares the average health outcomes and death probability of individuals by age at which they retire. 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 ( $66 \leq r \leq 69$ ). Results are expressed relative to the level of normal retirees. The graph shows first two indices for bad health (i.e. standardized principal components extracted from all health outcomes in the HEK and ULF surveys) and two measures of "life expectancy" (dummies for being dead by age 70, or by age 75). See Kolsrud et al. (2024) for details.

The findings suggest that reforms aimed at incentivising later retirement can have significant welfare costs for early retirees. To sum up the evidence from Kolsrud et al. (2024) in simple terms, people who retire early have lower consumption in retirement, and these lower consumption levels are in large part determined by what happens just around retirement. Early retirees also have worse health and life expectancy, and early retirement strongly correlates with the incidence of negative health shocks. As a consequence, pension reforms that transfer resources from early to late retirees have significant redistributive and insurance costs.

The patterns observed in Sweden are not unique. Data from other countries, particularly from the Survey of Health, Ageing and Retirement in Europe (SHARE) and Health and Retirement Study (HRS) datasets, reveal similar patterns in consumption across retirement ages. In both the United States and Europe, individuals who retire later tend to maintain higher levels of consumption, while those who retire earlier experience

sharper declines. In many countries, we also find a non-monotonicity between early and normal retirement ages. This consistency across different contexts suggests that the issues identified in Sweden are likely to be relevant in other advanced economies as well.

In conclusion, a more comprehensive and data-driven approach to pension reform is needed – one that takes into account not just fiscal costs and benefits, but also redistributive and insurance effects. We have shown how consumption can be a key metric for evaluating these welfare effects. By adopting a transparent, welfare-based framework, policymakers can design reforms that are both equitable and sustainable, ensuring that pension systems provide adequate support for all individuals, regardless of their health, wealth, or retirement timing.

Following this logic, Kolsrud et al. (2024) suggest that some reforms may actually have both positive fiscal externalities and positive redistributive gains. We have already shown this for strengthening incentives to retire later between the early and normal retirement age - the age range where the consumption gradient with retirement age reverses. Second, we also find that at any given retirement age, individuals who retire with shorter career durations are significantly wealthier and enjoy higher levels of consumption than individuals who have longer careers (having started earlier or experienced fewer interruptions). Reforms that strengthen labour supply incentives by penalising shorter careers and/or rewarding longer careers, conditional on retirement age, are therefore unambiguously good from a social welfare perspective. They provide stronger incentives and redistribute in the desired direction. With pension reforms, surprisingly enough, you may sometimes have your cake and eat it!

## REFERENCES

- Aguiar, M. and E. Hurst (2005), “Consumption versus Expenditure”, *Journal of Political Economy* 113(5): 919–948.
- Banks, J., R. Blundell, and S. Tanner (1998), “Is There a Retirement-Savings Puzzle?”, *The American Economic Review* 88(4): 769–788.
- Bernheim, B.B., J. Skinner, and S. Weinberg (2001), “What Accounts for the Variation in Retirement Wealth Among U.S. Households?”, *The American Economic Review* 91(4): 832–857.
- Chetty, R. (2006), “A general formula for the optimal level of social insurance”, *Journal of Public Economics* 90(10-11): 1879–1901.
- Chetty, R. (2008), “Moral hazard versus liquidity and optimal unemployment insurance”, *Journal of Political Economy* 116(2): 173–234.
- Haller, A. (2022), “Welfare Effects of Pension Reforms”, working paper.

Kolsrud, J., C. Landais, and J. Spinnewijn (2020), “The Value of Registry Data for Consumption Analysis: An Application to Health Shocks”, *Journal of Public Economics* 189: 1040–1088.

Kolsrud, J., C. Landais, D. Reck, and J. Spinnewijn (2024), “Retirement Consumption and Pension Design”, *American Economic Review* 114(1): 89 – 133.

Landais, C. and J. Spinnewijn (2021), “The Value of Unemployment Insurance”, *Review of Economic Studies* 6(88): 3041–3085.

Stephens, M. and D. Toohey (2018), “Changes in Nutrient Intake at Retirement”, NBER Working Paper No. 24621.

### ABOUT THE AUTHORS

**Camille Landais** is a Professor of Economics, Director of the Suntory and Toyota International Centres for Economics and Related Disciplines (STICERD), and Co-Director of the Hub For Equal Representation in the Economy (H.E.R) at the London School of Economics. He is also the President-Delegue at the French Council of Economic Advisers (CAE). He received his PhD from the Paris School of Economics and was a Research Fellow at Stanford University (Stanford Institute for Economic Policy Research). His research interests cover topics such as taxation, social insurance, and pro-social behaviours.

**Johannes Spinnewijn** is Professor in Economics at the London School of Economics. His research focuses on various topics in public economics, including the design of social insurance and healthcare and how this is affected by behavioral biases. Johannes is Director of the Public Economics programme at CEPR and Research Fellow at the Institute for Fiscal Studies. He has been awarded the British Academy’s Wiley Prize for outstanding early career achievement in research. Prior to joining LSE, he completed his PhD at the Massachusetts Institute of Technology.