Jia, Zhiyang; Kruse, Herman; Vigtel, Trond C.

Working Paper

Welfare effects of pension reforms

Discussion Papers, No. 1025

Provided in Cooperation with:
Research Department, Statistics Norway, Oslo

Suggested Citation: Jia, Zhiyang; Kruse, Herman; Vigtel, Trond C. (2025) : Welfare effects of pension reforms, Discussion Papers, No. 1025, Statistics Norway, Research Department, Oslo

This Version is available at:
https://hdl.handle.net/10419/322383

Statistisk sentralbyrå
Statistics Norway

Welfare Effects of Pension Reforms

Zhiyang Jia, Herman Kruse and Trond Christian Vigtel

TALL SOM FORTELLER
DISCUSSION PAPERS
1025

Discussion Papers: comprise research papers intended for international journals or books. A preprint of a Discussion Paper may be longer and more elaborate than a standard journal article, as it may include intermediate calculations and background material etc.

The Discussion Papers series presents results from ongoing research projects and other research and analysis by Statistics Norway staff. The views and conclusions in this document are those of the authors.

Published: July 2025

Abstracts with downloadable Discussion Papers in PDF are available on the Internet:

https://www.ssb.no/discussion-papers

http://ideas.repec.org/s/ssb/dispap.html

ISSN 1892-753X (electronic)

Abstract

This paper investigates the welfare effects of two types of pension reforms aimed at addressing challenges due to aging populations. The study uses a framework by Kolsrud et al. (2024), decomposing welfare into consumption smoothing and fiscal externality effects. Norwegian administrative data is used to study the welfare effects of two reforms. The first is a hypothetical budget-neutral reform steepening pension incentives, which rewards late retirees. The second is the 2011 Norwegian old-age pension flexibility reform. We find that the first (hypothetical) reform is regressive. Based on consumption differences, such a reform incurs substantial consumption smoothing costs and results in significant overall welfare costs (0.4–0.5 NOK per 1 NOK transferred), highlighting the negative welfare impact of heavily penalizing early retirement. Conversely, the 2011 Norwegian old-age pension flexibility reform, which lowered the eligibility age (from age 67 to age 62) had a near-zero effect on total labor supply. Quasi-experimental evidence suggests this reform shifted the consumption distribution upwards and resulted in welfare gains, estimated at around NOK 138,000 per affected individual.

Keywords: Pension reform; welfare effects; consumption; retirement; pension claiming age

JEL Classification: J26

Acknowledgements: The authors graciously acknowledge the financial support from The Ministry of Labour and Social Inclusion (PENAL project). We also thank Lars J. Kirkebøen for valuable feedback and suggestions.

Addresses:
Zhiyang Jia, Statistics Norway, Research Department. E-mail: zhiyang.jia@ssb.no.
Herman Kruse, Statistics Norway, Research Department. E-mail: herman.kruse@ssb.no.
Trond C. Vigtel, Statistics Norway, Research Department. E-mail: trond.vigtel@ssb.no.

Sammendrag

Denne artikkelen undersøker velferdseffektene av to typer pensjonsreformer som tar sikte på å møte utfordringer knyttet til en aldrende befolkning. Studien benytter en modell utviklet av Kolsrud et al. (2024), som deler velferd inn i konsumutjevning og finansielle eksternaliteter. Norske registerdata benyttes for å analysere velferdseffektene av to reformer. Den første er en hypotetisk budsjettnøytral reform som skjerper pensjonsinsentivene og belønner utsatt pensjonering. Den andre er fleksibilitetsreformen i den norske alderspensjonen fra den norske pensjonsreformen i 2011. Vi finner at den første (hypotetiske) reformen er regressiv. Basert på forskjeller i konsum medfører en slik reform betydelige kostnader knyttet til konsumutjevning og resulterer i betydelige samlede velferdstap (0,4–0,5 kroner per overført krone), noe som understreker den negative velferdseffekten av å straffe tidligpensjonering. Derimot hadde fleksibilitetsreformen i 2011, som senket pensjonsalderen fra 67 til 62 år, en tilnærmet null effekt på samlet arbeidstilbud. Våre funn tyder på at reformen skiftet konsumfordelingen oppover og førte til velferdsgevinster, estimert til rundt 138 000 kroner per individ.

1. Introduction

In responses to the aging problem, public policies promoting longer working lives and delayed retirement have been implemented in many countries during last two decades. These policies mainly work through two channels to induce workers to retire later: either via shifting financial incentives for late retirees, or via directly increasing statutory retirement age. Such reforms effectively steepen the net present value (NPV) profile of pension wealth relative to retirement timing. As a result of these policies, workers stay on longer in the labor market and the government collects more tax revenue. However, 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). By nature, these types of reforms are regressive, as they tend to penalize groups with poorer health who may die before even becoming eligible. Additionally, while delayed retirement incentives – such as an increase in annual pension benefits for postponing retirement – are designed to encourage longer working lives, they disproportionately benefit individuals who are already well-off. These individuals often enjoy their work, are in better health, and would likely continue working even in the absence of such incentives. As a result, these types of policies may, to a large extent, reward behavior that would have occurred regardless, offering large transfers to those with the least need for additional financial support. In other words, while resolving the financial sustainability of the pension system, these policies could lead to an overall reduction in social welfare. Thus, understanding the welfare consequences of these different schemes is essential for designing efficient pension systems.

In this paper, we follow the framework developed by Kolsrud et al. (2024) and investigate the welfare effects of two types of pension reforms that are widely used. In the first type of reform, we increase the NPV for late retirement at the cost of a deduction of the NPV for early retirement. In other words, we change the slope of the NPV as a function of the retirement age while keeping the reform budget-neutral. Essentially, this type of reform strengthens the incentives to stay longer in the labor market.

A more direct policy tool is to increase the access age to induce workers to postpone their withdrawal from the labor force, as more than a dozen countries in the OECD have done (OECD, 2017). In contrast to this policy, the 2011 Norwegian pension reform gave individuals more flexibility in terms of when to retire, by lowering the earliest possible pension age from 67 to 62 for a major subgroup of workers (Hernæs et al., 2020), while at the same time keeping the reform budget-neutral. Thus, in the second type of reform we consider, we adjust the eligibility age of pension claiming, but with actuarial adjustments aimed at keeping the NPV unchanged.

Following Kolsrud et al. (2024), we use consumption after retirement to shed light on the welfare changes of these pension reforms. The consumption levels are imputed from the Norwegian income and wealth registry data using the method proposed by Fagereng and Halvorsen (2017).

Our contribution to the literature is twofold. Firstly, we document the empirical patterns of consumption after retirement across different groups of individuals in Norway. We find that, generally, average consumption tends to be higher for later retirees compared to earlier retirees, although there is some variation, such as early retirees potentially consuming slightly more than those retiring in the "normal" age bracket (62–66) after adjusting for income-related variables. This pattern supports the idea that those who retire earlier might be worse off and value consumption higher. Secondly, based on observed consumption patterns and the welfare framework, we evaluate the welfare implications of two pension reforms. We find that a hypothetical reform which steepens the net present value (NPV) profile of pension wealth relative to retirement timing, benefiting late retirees at the expense of early retirees entails substantial consumption-smoothing costs due to its

regressive nature, resulting in potentially large overall welfare losses. In contrast, the 2011 Norwegian reform of the pension claiming age, which increases flexibility by lowering the eligibility age while applying actuarial adjustments to keep the NPV unchanged, had a near-zero effect on total labor supply for the affected group but shifted the consumption distribution upward. Based on quasi-experimental evidence, our analysis suggests that the reform generated a sizable welfare gain: approximately NOK 138,000 per affected individual.individual.

The remainder of our paper is organized as follows: Section 2 details the welfare framework from Kolsrud et al. (2024), and explains how one can decompose the welfare effects into consumption smoothing costs and fiscal externalities. Section 3 describes the Norwegian administrative data and present the results from the consumption imputation exercise. Section 4 quantifies the welfare effects of two reforms: Section 4.1 analyses a hypothetical reform with steeper incentives, while Section 4.2 evaluates the Norwegian 2011 flexibility reform using quasi-experimental methods. Section 5 concludes.

2. Measuring welfare implications of pension policies

Evaluating the welfare effects of pension policies that alter retirement incentives is complex. Pension benefits typically are a complex function of individuals employment history, including retirement/claiming age and past contributions. Reforms involve changing how these features map into benefits. Additionally, individual welfare depends on numerous individual choices and characteristics, such as consumption, labor supply, saving, bequests, and health status. Pension policies also interact with other social insurance programs like unemployment insurance and disability insurance, which can affect pathways into retirement and require joint consideration.

In this paper, we adopt the framework developed by Kolsrud et al. (2024). They show how the framework can be used to evaluate the welfare effects of marginal reforms that change the slope of pension benefits over age at retirement. Their framework decomposes the welfare effect of such a reform into two primary components. Firstly, the consumption smoothing effect, which captures the welfare gain or loss from redistributing resources (pension benefits) across groups of individuals who have different values of consumption. It directly depends on the social marginal utility of consumption (SMU) for the affected groups. Transferring resources to individuals with a higher SMU yields a positive consumption smoothing welfare effect. Individuals with different characteristics, for example retired at different ages, may place different values on pension benefits. In other words, aggregate social welfare might be improved by transferring benefits from those who value them less to those who value them more. And secondly, the fiscal externality (FE) effect, which captures the welfare effect that arises because policy changes can trigger behavioral responses (like changes in labor supply or retirement timing) that have fiscal consequences (e.g., changes in tax revenue and other benefit payments). The welfare impact is defined as:

$$\Delta W = \text{Fiscal externality} - \text{Consumption smoothing cost} \tag{2.1}$$

This decomposition builds on a "variational" approach commonly used in social insurance literature and relies on the envelope theorem. The envelope theorem implies that, to a first order, behavioral responses only affect welfare via the fiscal externality, unless behavioral biases are significant.

In the following, we present the framework by Kolsrud et al. (2024) in more detail. Individuals are

assumed to optimize over their life time discounted expected utility:

$$\sum_{t=0}^{T} \beta^t \int u(c(\pi_{i,t}), \kappa(\pi_{i,t}))dF(\pi_{i,t}) \tag{2.2}$$

with proper budget constraints over the consumption (c) and labor supply (retirement) decisions (κ). Here the state variable $\pi_{i,t} \in \mathbf{\Pi}_t$ captures all aspects of individual i's history and characteristics relevant. We denote the life-time expected indirect utility depends on pension rules $b = b(\pi_{i,t})$ and tax rules τ:

$$U_i(b, \tau). \tag{2.3}$$

Note that function b maps the worker's employment history into the pension benefits. A pension reform is considered to be a change in this mapping.

The government tries to maximize the weighted sum of the individual indirect utility functions with the government budget constraint which is given by $GBC(b, \tau)$. Or mathematically:

$$\max W(b, \tau) = \sum_i \omega_i U_i(b, \tau) + \lambda GBC(b, \tau), \tag{2.4}$$

where ω_i is the weight for individual i, and λ is the Lagrange multiplier.

In the following, we consider the welfare effect of a small change in pension benefits $b_{x,t}$ received at age t by individuals who retire with features x. The marginal change in the first term is proportional to the marginal indirectly utility function:

$$\frac{\partial \sum_i \omega_i U_i(b, \tau)}{\partial b_{x,t}} = \sum_i \omega_i \frac{\partial U_i(b, \tau)}{\partial b_{x,t}} = G(x, t) E[\omega_i \frac{\partial u(c_{i,t}, \kappa)}{\partial c} | x = x_{i,t}] \equiv G(x, t) SMU_{x,t}, \tag{2.5}$$

where $G(x, t)$ is the group size of individuals with retirement features x. The SMU is proportional to the average marginal utility of transferring one monetary unit to that group. This captures the idea of consumption smoothing: moving resources from individuals with lower marginal utility of consumption to those with higher marginal utility generally improves welfare.

The fiscal cost of the change is given by:

$$\lambda \frac{\partial GBC(b, \tau)}{\partial b_{x,t}} = \lambda G(x, t) \times (1 + FE_{x,t}), \tag{2.6}$$

where $FE_{x,t}$ is the fiscal externality due to the individuals behavioral response. For example, assuming pension benefits are increased if one retires later, this will induce more individuals to work longer – implying higher overall tax revenue and possibly also lower overall pension expenditure.[1]

Kolsrud et al. (2024) consider a particular type of pension reform, namely a marginal reduction (-$db_{r \leq \tilde{r}}$) in pensions for individuals who retired before a given age $\tilde{r}$ and a marginal increase ($db_{r > \tilde{r}}$) for individuals who retire after age $\tilde{r}$. The relative size of the reduction is calculated such that the reform is, to some extent, budget neutral:

$$db_{r > \tilde{r}} S(\tilde{r}) - \{[1 - S(\tilde{r})]\} db_{r < \tilde{r}} = 0, \tag{2.7}$$

where $1 - S(\tilde{r})$ is the share of individuals who retired before age $\tilde{r}$.

Assuming equal weights, i.e., $\omega_i = 1$, the change in the consumption smoothing cost is given by:

$$\begin{aligned} &\sum_{t < \tilde{r}} S(t) SMU_t db_{r < \tilde{r}} + \sum_{t > \tilde{r}} S(t) SMU_t db_{r > \tilde{r}} \\ &\quad \approx (1 - S(\tilde{r})) SMU_{r < \tilde{r}} db_{r < \tilde{r}} + S(\tilde{r}) SMU_{r > \tilde{r}} db_{r > \tilde{r}} \\ &\quad = (1 - S(\tilde{r})) db_{r < \tilde{r}} (SMU_{r < \tilde{r}} - SMU_{r > \tilde{r}}). \end{aligned} \tag{2.8}$$

[1] This approach ignores the presence of behavior biases. Interested readers are referred to Kolsrud et al. (2024) for more detailed discussions.

And similarly the fiscal externality is given by:

$$\lambda(1-S(\tilde{r}))db_{r<\tilde{r}}(FTE_{r<\tilde{r}}-FTE_{r>\tilde{r}}). \tag{2.9}$$

To simplify the interpretation, the welfare effect is measured relative to the value of one NOK given to the reference group of individuals retiring at the normal retirement age $\tilde{r}$, $SMU_{\tilde{r}}$. Similar to Kolsrud et al. (2024), we assume it to be approximately equal to the marginal cost of public funds λ. This implies that we can approximate the net welfare gain per NOK transferred from individuals retiring before $\tilde{r}$ to individuals retiring after $\tilde{r}$ by:

$$\left[(FE_{r>\tilde{r}}-FE_{r<\tilde{r}})-\frac{SMU_{r<\tilde{r}}-SMU_{r>\tilde{r}}}{SMU_{\tilde{r}}}\right]. \tag{2.10}$$

2.1 Measuring consumption smoothing costs using consumption level differences

One way of quantifying the consumption smoothing cost $\frac{SMU_{r<\tilde{r}}-SMU_{r>\tilde{r}}}{SMU_{\tilde{r}}}$, suggested by Kolsrud et al. (2024), is to use the post-retirement consumption levels. In particular, we compare consumption levels among individuals who retired at different ages, measured at the same age while retired. We then can map the observed consumption differences to SMU differences using a Taylor-series approximation and a parameter for relative risk aversion (γ). The relative social marginal utility ratio between those who retired at age r and r' can be approximated by:

$$\frac{SMU_r}{SMU_{r'}} \approx 1+\gamma\frac{c_{r'}-c_r}{c_{r'}} \tag{2.11}$$

where γ is the relative risk aversion (ie. the curvature of the utility function $\gamma=-u''/u'$). $c_{r'}$ and c_r are consumption levels for different groups who retire at different ages.

2.2 Imputation of consumption levels

While we have high quality register data on many individual characteristics, such as their labor supply behavior, pension claiming behavior, earnings history and education level, the consumption levels are not directly observable in the data.

To impute consumption, we apply the method suggested by Fagereng and Halvorsen (2017):

$$C_t = Y_t+(A_{t-1}-D_{t-1})r_t-(A_t-A_{t-1}+D_{t-1}-D_t+P_t^hI_t^h)+B_t, \tag{2.12}$$

where C_t is consumption at year t, Y_t is income and r is the rate of return (interest rate). Assets accumulated the year before is represented as A_{t-1}, while D_{t-1} measures the debt level. B_t is bequest received, and $P_t^hI_t^h$ is investment in housing. Thus, $Y_t+(A_{t-1}-D_{t-1})r_t+B_t$ represent cash-on-hand and $(A_t-A_{t-1}+D_{t-1}-D_t+P_t^hI_t^h)$ measures investment. The consumption level is simply defined as the difference between cash-on-hand and investment (savings).

Following Fagereng and Halvorsen (2017), we exclude periods with housing investment because of measurement errors. Additionally, we assume no bequests. We also do not have access to detailed stock holding information, although we have the information of market values of aggregated assets classes, such as stock holdings, mutual funds, individual pension savings, and bonds from the tax return register. We use the historical annual return of Oslo Stock Exchange to calculate the returns of stocks. We make use of a weighted Oslo Stock Exchange (30%) and MSCI World index (70%) to impute the returns of mutual fund returns, and we use the Treasury bill to impute returns on bonds.

2.3 Measuring the fiscal externality

The fiscal externality component captures the welfare effect that arises from behavioral responses triggered by changes in pension policy, specifically changes in individuals' labor supply or retirement timing, and the fiscal consequences of those responses.

Following Kolsrud et al. (2024), the fiscal externality is quantified by the product of two factors. The first is the behavioral response, which is measured by the extensive labor supply elasticity. This elasticity reflects how much the likelihood of staying in employment changes in response to the financial incentives (or disincentives) provided by the pension system at a given age. The second is the fiscal return, which is measured by the participation tax rate. This rate represents the net gain to the government from an individual working an additional year instead of retiring. It is calculated by considering the income tax and payroll taxes paid on earnings, minus the change in the NPV of pension and other benefits received if the person had retired earlier.

3. Data

The analysis uses Norwegian administrative data merged with several different registers that covers the entire Norwegian population. We draw on four main administrative sources: the linked employer-employee data sets, the annual tax registers, the Central Population Register, Nav's social security register, and Statistics Norway's pension-wealth database. These data sets are linked via unique encrypted personal identification. The employer-employee register encompasses all wage workers and employers, and include contracted weekly working hours in their primary job each year, industry of the firm, and occupation of the worker. Another important source is the tax return register, which contains various sources of earnings, as well as wealth information that includes bank deposits, bonds, mutual funds, stocks, and debt. Under Norwegian law, employers, banks, brokers, insurers and other financial intermediaries must disclose individuals earnings, the value of any assets they manage on behalf of those individuals, and the income generated by those assets to both the taxpayer and the authorities. Because most entries are reported by third parties, the data are both accurate and reliable. Additionally, our dataset contains information on educational attainment, marital status, and other demographic characteristics.

The sample includes all individuals born between 1940 and 1953 and covers the period from 2000 to 2018. We impose further sample restrictions when we implement different parts of our empirical analysis below.

3.1 Consumption imputations

Based on the method described in Section 2.2, we construct a longitudinal measure of consumption for the entire Norwegian population for sample period from 2000 to 2018, applying data from administrative records collected for tax purposes.

Figure 3.1 shows the how the median consumption varies over different ages for selected years and birth cohorts. All monetary amounts in this paper are measured in NOK, deflated using the consumer price in 2015. These plots reveal a distinct life-cycle pattern, and a large increase in real consumption over time. There are real consumption growth effects over time that causes younger birth cohort to follow life cycle consumption profiles at a higher level than previous birth cohorts.

There are, unfortunately, some measurement errors in the imputed consumption. For example, there could be large swings in consumption from one year to another, even after we drop the years in which we observe large changes in housing values. This could be due to the fact we do not have

Figure 3.1 Median consumption over age: by year and by birth cohort

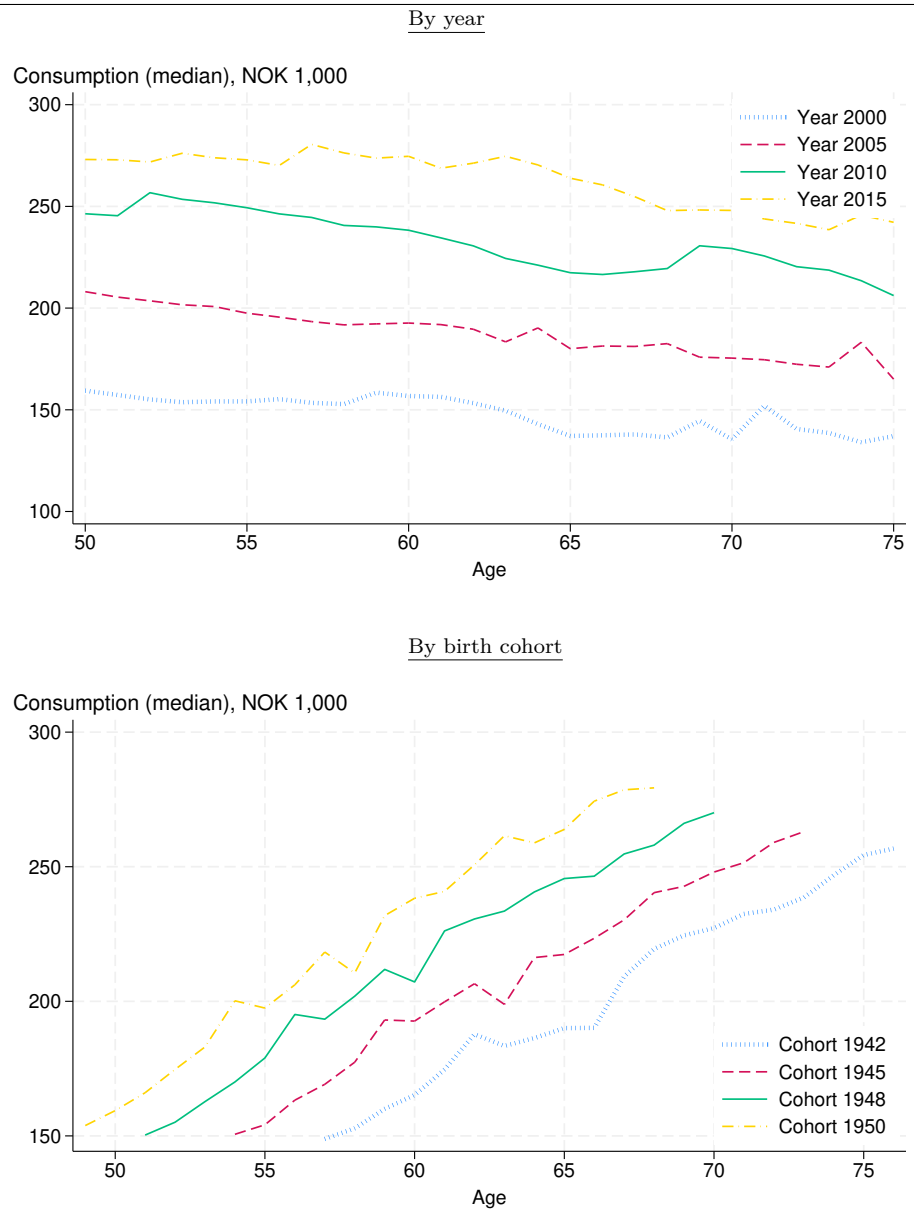

Source: Authors' own calculations using data from Statistics Norway.

the exact composition of financial assets at the individual level. Our method of using the composite stock exchange index will not fully capture actual returns on the individual level. Nevertheless, Fagereng and Halvorsen (2017) compared the imputed consumptions using the same method and those from The Norwegian Survey of Consumer Expenditures and found that the imputed values capture the main distributional patterns in terms of consumption.

4. Quantifying welfare implications

In this section, we apply the framework from Section 2 to evaluate the welfare effects of two types of pension reforms.

Figure 4.1 Distribution of retirement ages

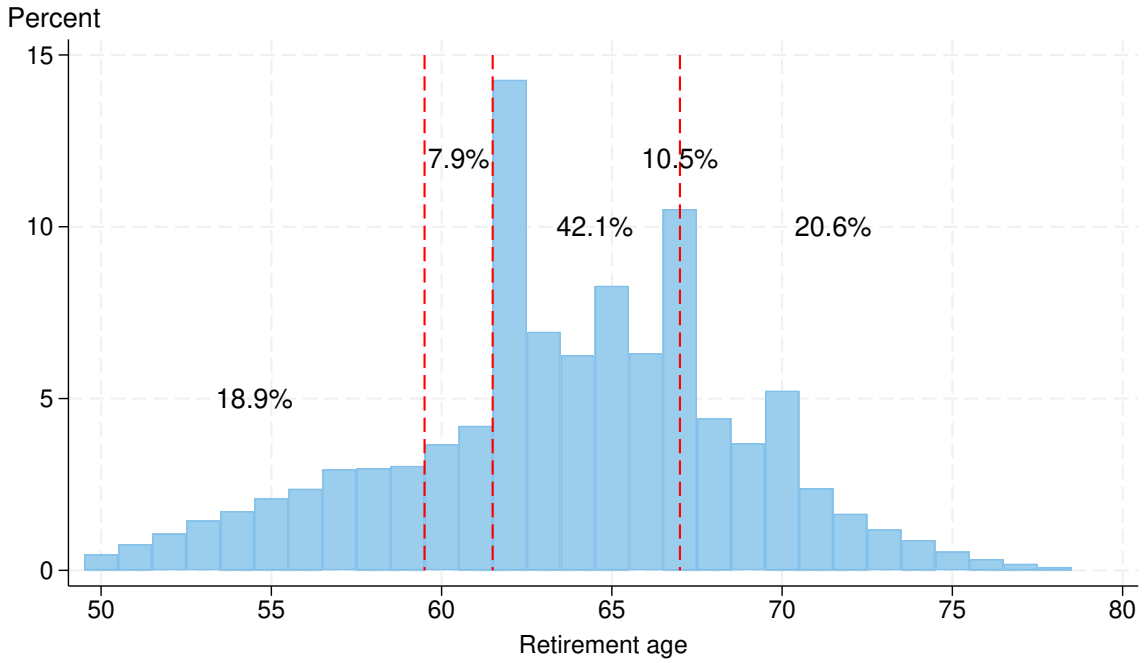

Note: The figure shows the distribution of retirement ages among individuals born between 1940 and 1953. Retirement is defined as the last year with labor income above 1 basic amount (approximately NOK 90,000 in 2015). The red dashed lines indicate the retirement age cutoffs used to define the five retirement groups: premature (below 60), early (60–61), normal (62–66), statutory (67), and late (above 67). Percentages reflect the share of individuals retiring within each group.
Source: Authors' own calculations using data from Statistics Norway.

4.1 A reform with steeper incentives

The first reform is a purely hypothetical reform. It assumes a budget-neutral reform of the old-age pension system but incorporates steeper incentives (or penalties) across retirement ages.

4.1.1 Retirement consumption patterns by retirement age

Following the method described in Section 2, we first look at the consumption patterns after retirement for different groups of individuals, which are characterized by the time when they withdraw from the labor market. We define the retirement age as the age when individuals withdraw from the labor market as the last year with labor income over 1 basic amount (approximately NOK 90,000 in 2015).

Retirees are categorized into five groups based on their age at retirement: (1) premature retirees, who exited the labor force before age 60; (2) early retirees, who retired at age 60 or 61; (3) normal retirees, who retired between ages 62 and 66; (4) statutory retirees, who retired at the statutory retirement age of 67; and (5) late retirees, defined as those who retired after age 67.

Figure 4.1 reports the distribution of age at retirement among individuals from the 1940 to 1953 birth cohorts in Norway. We see that there is considerable variation across the population. While there is a large share of individuals withdrawing from the labor force before the early retirement pension becomes available (age 62), many continue to have some labor market attachments after the official retirement age (age 70).

A regression model is used to estimate consumption differences across retirement age groups, controlling for year, birth cohort, and age. Fixed effects isolate the average consumption difference α_j

Figure 4.2 Consumption differences in retirement across retirement age groups

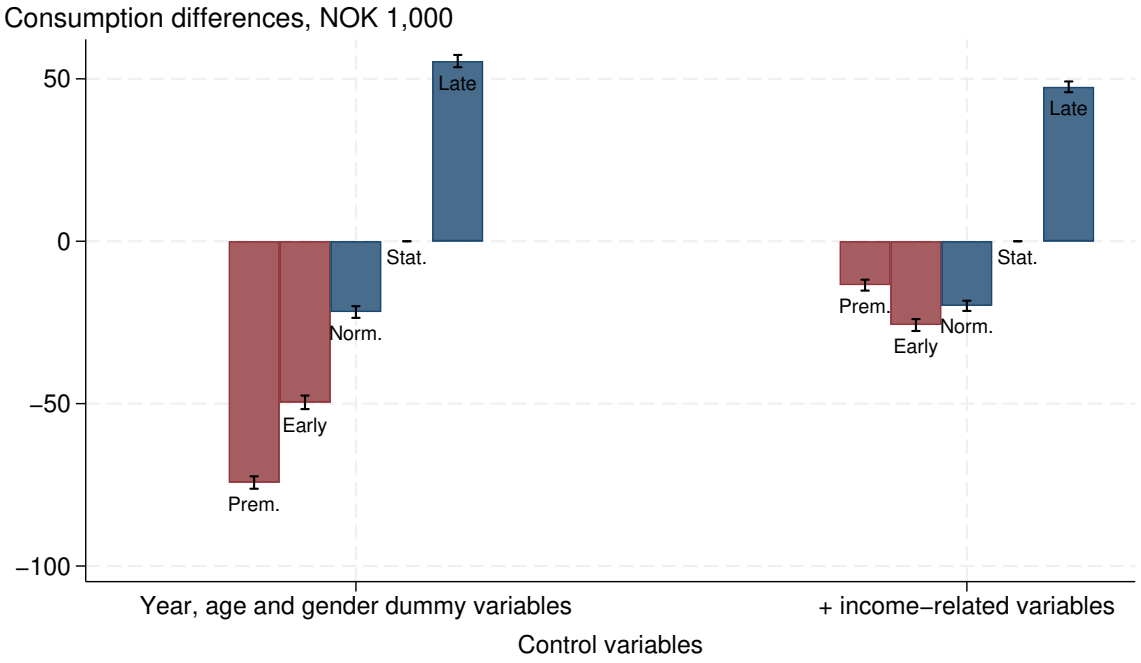

Note: The figure plots estimated coefficients α_j from Equation (4.1), capturing average consumption differences across retirement age groups relative to retiring at age 67 (the reference group). The left-hand side panel includes controls for year, age, and gender fixed effects. The right-hand side panel additionally controls for income-related variables, including lifetime earnings. Retirement groups are categorized as: premature retirees (Prem.), early retirees (Early), normal retirees (Norm.), statutory retirees (Stat.), and late retirees. Error bars represent 95% confidence intervals.
Source: Authors' own calculations using data from Statistics Norway.

between retirement groups. The model is specified as

$$C_{it} = \sum_j \alpha_j 1\{r = j\} + \gamma_y + \gamma_t + X'\beta + \epsilon_{it}, \tag{4.1}$$

where j indexes different groups by the timing of retirement (five groups), while γ_y and γ_t denotes year and age fixed effects, respectively. α_j measures the differences in consumption of individuals from the same birth cohort at the same age, but who have retired at age 67. The vector X denotes the income-related variables that represent the earnings history and that are the key determinants of the NPV.

Figure 4.2 presents the estimated coefficients $\hat{\alpha}_j$ from Equation (4.1) across all retirement age groups. The estimates are based on consumption levels with those who withdrew at age 67 as the reference group. The left-hand side of of the figure corresponds to a regression including only year and age fixed effects. The right-hand side of the figure includes the income-related variables. From the figure, we observe a robust upward slope in consumption with retirement age: average consumption tends to be higher for later retirees. After adjusting for the-income related variables (as a proxy for pension rights) the differences remain pronounced.

However, in this case, the consumption-retirement age relationship is not perfectly monotonic. Similar to what is found using Swedish data, early retirees – particularly those retiring well before the typical retirement age – tend to show a reversal, consuming slightly more than those who withdraw from labor market during ages 62 to 66. This patterns confirms our conjecture that those who retire earlier are worse off than those who retire late. Consequently, delayed retirement incentives, while implying a positive fiscal effect, could result in negative social welfare effect as it reallocates benefits from those who value consumption higher to those who value consumption lower.

4.1.2 The welfare implications for budget-neutral transfers

Using Equation (2.10), the total welfare effect can be quantified. It is computed as the difference between the *consumption smoothing cost* and the *fiscal externality*. For the consumption smoothing cost, we use the estimated consumption gaps from Equation (4.1). We also assume that the relative risk aversion parameter γ is equal to 4, a value which is used by both Kolsrud et al. (2024) and Sæverud (2024). The fiscal externality is typically found to be positive when shifting benefits from early to late retirees, as governments typically gain more from increased income tax revenue from the increased labor earnings. The values reported in the literature vary from 0.15 in Sweden to 0.31 in Denmark. We do not estimate the fiscal externality for Norway in this paper, but instead choose 0.2 as our baseline estimate. However, the main patterns of the welfare implications are robust to different choices of the fiscal externality within the range of estimates found in the literature.

Table 4.1 reports the estimated consumption smoothing costs for a transfer of 1 NOK from early retirees to late retirees. For example, if we reduce the pension benefits for those who retire earlier than 67, and distribute these benefits to those who retire at and after age 67, it induces a social cost of 0.66 NOK as the earlier retirees value the consumption (and pension) more than the late retirees. Our estimates of the consumption smoothing cost is in the higher end of the estimates from Sweden, which ranges from 0.3 to 0.8 (Kolsrud et al., 2024).

Table 4.1 Estimated consumption smoothing costs for budget-neutral transfers

Transfer type	Cost estimate
Premature (< 60) $\rightarrow$ Older groups (≥ 60)	0.62
Premature+Early (< 62) $\rightarrow$ Normal+Statutory+Late (≥ 62)	0.61
Premature+Early+Normal (< 67) $\rightarrow$ Statutory+Late (≥ 67)	0.66
Premature+Early+Normal+Statutory (≤ 67) $\rightarrow$ Late (> 67)	0.70

Source: Authors' own calculations using data from Statistics Norway.

Even after considering the positive fiscal externality, the values remain to be around 0.4–0.5 NOK for each NOK transferred, implying potentially large welfare costs of pension reforms that steepen the pension profiles to incentivize later retirement.

However, some caution is needed when interpreting these values. First of all, the reform considered is only a stylized reform which makes a marginal adjustment to the pension profile. Moreover, strong assumptions are made in the process of obtaining these estimates, in particular, preference heterogeneity within retirement age groups may bias the SMU estimates. Moreover, measurement error in consumption would also impact our estimates. Nevertheless, the main idea holds: pension designs that strongly penalize early retirement or heavily reward late retirement are particularly regressive, as they impose high welfare costs on consumption smoothing.

4.2 The Norwegian 2011 old-age pension reform

Another common type of pension reform involves changes to the eligibility age. In the following, we study both the behavioral impact and the welfare implications of such a reform using the 2011 Norwegian old-age pension reform as an example.

4.2.1 The main elements of the flexibility reform

The 2011 Norwegian pension reform introduced substantial flexibility into the public pension system, most notably within the National Insurance Scheme (NIS) and the private sector AFP ("av-

talefestet pensjon"). The main elements of the reform – such as automatic longevity adjustments and the abolition of all earnings tests in the private sector – were aimed at ensuring the long-term fiscal sustainability of the system. At the same time, the reform introduced additional flexibility in the pension access age, including a significant reduction in the earliest claiming age from 67 to 62 for a key subgroup of workers. This shift allowed individuals to begin drawing NIS old-age pensions flexibly between ages 62 and 75, with actuarial adjustments that reward later retirement and reduce benefits for earlier claims. The private sector AFP was transformed from a strictly earnings-tested early retirement scheme into a lifelong, non-tested supplement that must be claimed alongside the NIS pension. These changes effectively eliminated the implicit tax on continued work, which had previously discouraged labor force participation among older workers. By bringing uniformity to a previously fragmented set of access ages and incentives, the reform aimed to support more consistent work-retirement decisions across the population. The result was a system in which workers can adapt to increased work incentives, balance out longevity adjustments, and tailor their retirement timing more flexibly, thereby enhancing both individual choice and overall system sustainability.

4.2.2 The impact of the flexibility reform

To study the effects of the flexibility reform on consumption, we follow the framework detailed in Hernæs et al. (2020). We identify the reform effects, following Sun and Abraham (2021), by employing an event-study difference-in-difference approach. Our sample consists of individuals from birth cohorts 1943–1953 who at age 59 (i) were employed as wage earners, (ii) earned at least NOK 85,000, and (iii) were not receiving any disability benefits, whether permanent or temporary. We include only those who were ineligible for AFP, and who met the post-reform requirements for claiming the new public old-age pension, with actuarial adjustments, at age 62.

The control group consists of individuals who are 60–61 years old, while our treatment group consists of individuals who are 62–66 years old. The treatment period is 2011 and later, when workers aged 62 and older can withdraw from the labor force with no loss in pension benefits, while the 60- and 61-year-olds cannot. Thus, before 2011, both the treatment group and the control group were constrained by the eligibility age, while after 2011, only those in the control group were constrained.

We start with the following linear event-study specification to derive the average effect on consumption:

$$y_{i,a,t} = \alpha + X_i\beta + \sum_{s=60}^{66} \gamma_s DA_s + \sum_{l=2009}^{2014} \lambda_l DT_l + \sum_{s=60}^{66}\sum_{l=2009}^{2014} \eta_{s,l} DA_s DT_l + \varepsilon_{i,a,t}, \tag{4.2}$$

where $y_{i,a,t}$ is consumption of individual i at age a in year t. The vector X_i is individual characteristics, DA_s are dummy variables for age and DT_l are dummy variables for year.

Figure 4.3 shows the reform effects using a median regression on the specification in Equation (4.2). We see clearly that the reform has a positive effect on the level of consumption across both the age and year dimension. The average median effect is NOK 21,558 (with a standard error of NOK 3,409) with covariates, and NOK 21,212 (with a standard error of NOK 3,258) without covariates.

Because individuals respond differently to the flexibility reform, we expect an uneven effect across the consumption distribution and we therefore analyze the changes induced by the reform across the entire outcome distribution. We achieve this using an estimator based on the complementary cumulative distribution function (CCDF), which is defined as 1 minus the cumulative distribution function, such that $CCDF(y) \equiv 1 - F(y)$. This approach, including the construction of distribu-

Figure 4.3 Reform effects on consumption by age and year

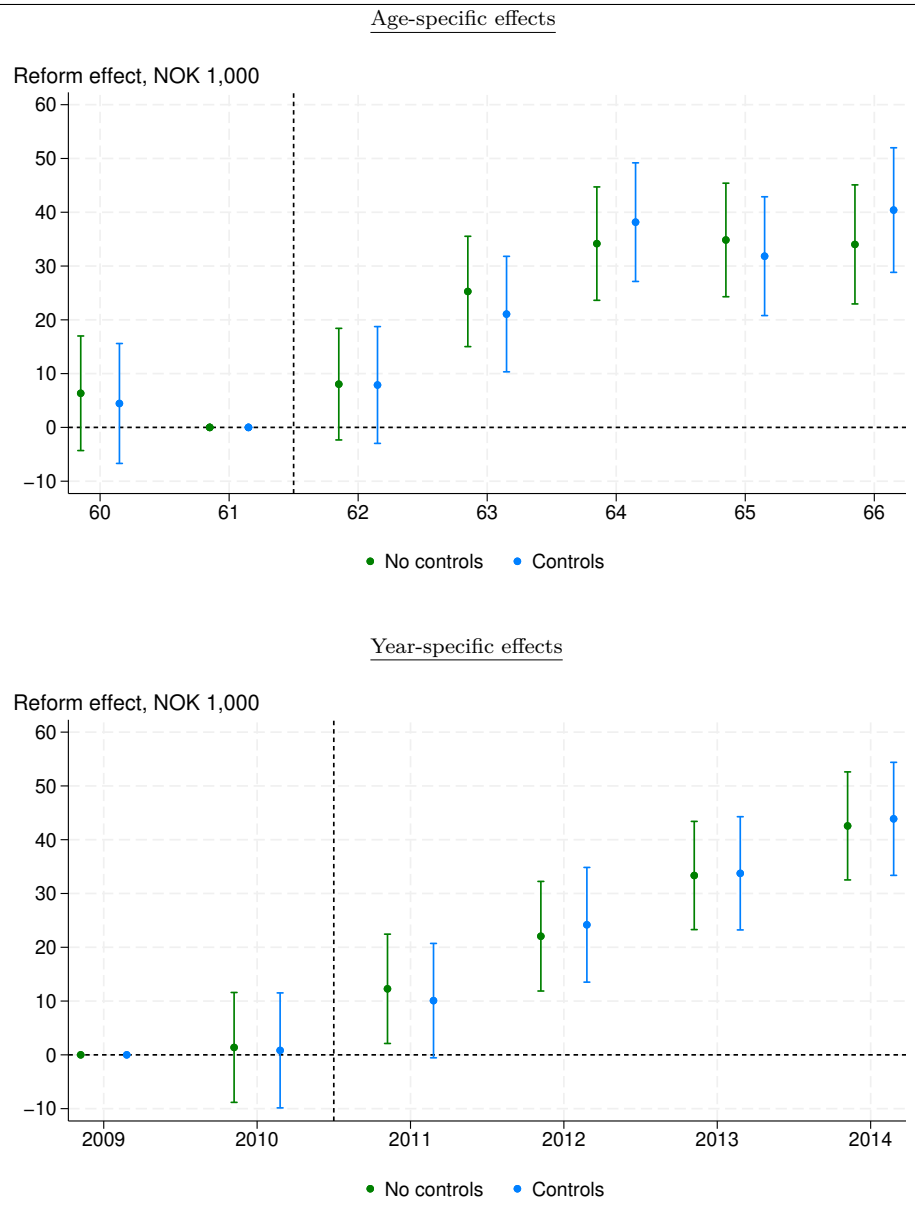

Note: Median regression results from estimating Equation (4.2), with age-specific and year-specific effects.
Source: Authors' own calculations using data from Statistics Norway

tion functions and the regression analysis of shifts and covariate impacts, is detailed in Hernæs and Jia (2013) and Brinch et al. (2017).

Following Hernæs and Jia (2013) and Brinch et al. (2017), we conduct a series of linear regressions on the probability of consumption above a series of steps, each of length NOK 75,000, up to a little over NOK 5 million. Specifically, for each level of consumption $y_{i,a,t}$ for individual i at age $a = 60, \ldots, 66$ and year t, we model the CCDF as follows:

$$\mathbf{1}\left(y_{i,a,t} > y^k\right) = \alpha^k + X_i\beta^k + \sum_{s=60}^{66} \gamma_s^k DA_s + \sum_{l=2009}^{2014} \lambda_l^k DT_l + \sum_{m=60}^{66}\sum_{l=2009}^{2014} \eta_{m,l}^k DA_m DT_l + \varepsilon_{i,a,t}^k \quad (4.3)$$

The building blocks for our estimates of the reform effects are cohort- and age-specific average effects from Equation (4.3), defined similarly to the cohort-specific average treatment effects on the treated in Sun and Abraham (2021). We aggregate these cohort- and age-specific effects using weights representing the shares of birth cohorts affected by the reform at a given age to obtain the

Figure 4.4 Effect on consumption distribution

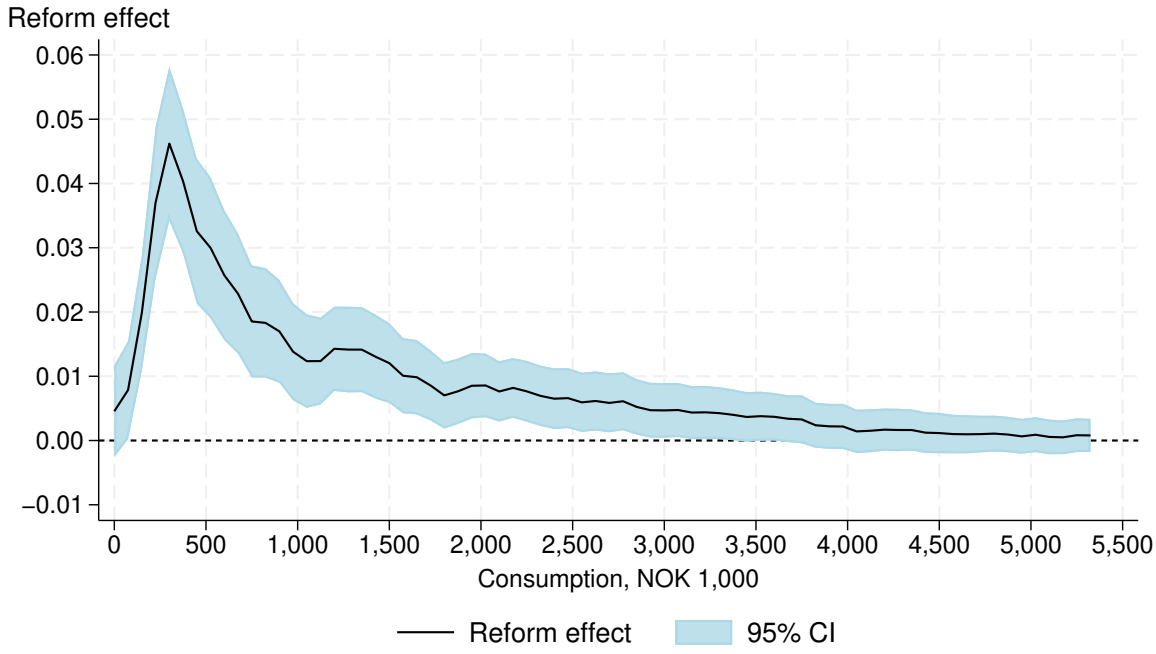

Note: Simulation results from estimation of Equation (4.3). Control variables are pre-determined and include linear controls for education length, education length squared, log average annual pre-tax earnings from age 30 to age 59, and net liquid wealth at age 59. The grey-shaded area shows the 95% confidence intervals (based on 200 non-parametric bootstrap replications for each estimation, clustered at the individual level).
Source: Authors' own calculations using data from Statistics Norway.

average treatment effects on the treated for a given age. Similarly, we can construct reform effects for a specific year or birth cohort and, eventually, the averaged reform effects over all birth cohorts from ages 62–66.

Figure 4.4 shows the average reform effect across all ages and years of the reform on the weekly working hours distribution. The main takeaway from Figure 4.4 is an increased mass of individuals with consumption above approximately NOK 300,000 and a corresponding reduction of individuals with consumption levels below NOK 300,000.

This pattern becomes even clearer when examining the corresponding shifts in the probability of being within various intervals of consumption distribution (Table 4.2). Here we see that the fraction with consumption below NOK 300,000 decreases by 4.6 percentage points after the reform, while the fraction of individuals with earnings between NOK 300,000 and NOK 900,000 increases by 2.9 percentage points.

Table 4.2 Probability of being in selected consumption intervals

Consumption interval	Probability (before reform)	Estimate	Standard error
$<$ NOK 300,000	0.455	-0.046	0.007
NOK 300,000–900,000	0.426	0.029	0.007
NOK 900,000–2,100,000	0.081	0.009	0.005
$\geq$ NOK 2,100,000	0.038	0.008	0.002

Note: Probability of being in the different consumption intervals and the effect on the probability of being in the different intervals. The effect is generated using the marginal effects from Figure 4.4, with pooled standard errors in the last column.
Source: Authors' own calculations using data from Statistics Norway.

Figure 4.5 shows the dynamics of the effect of the flexibility reform on consumption over ages 62–

Figure 4.5 Change in probability of being in selected consumption intervals, by age

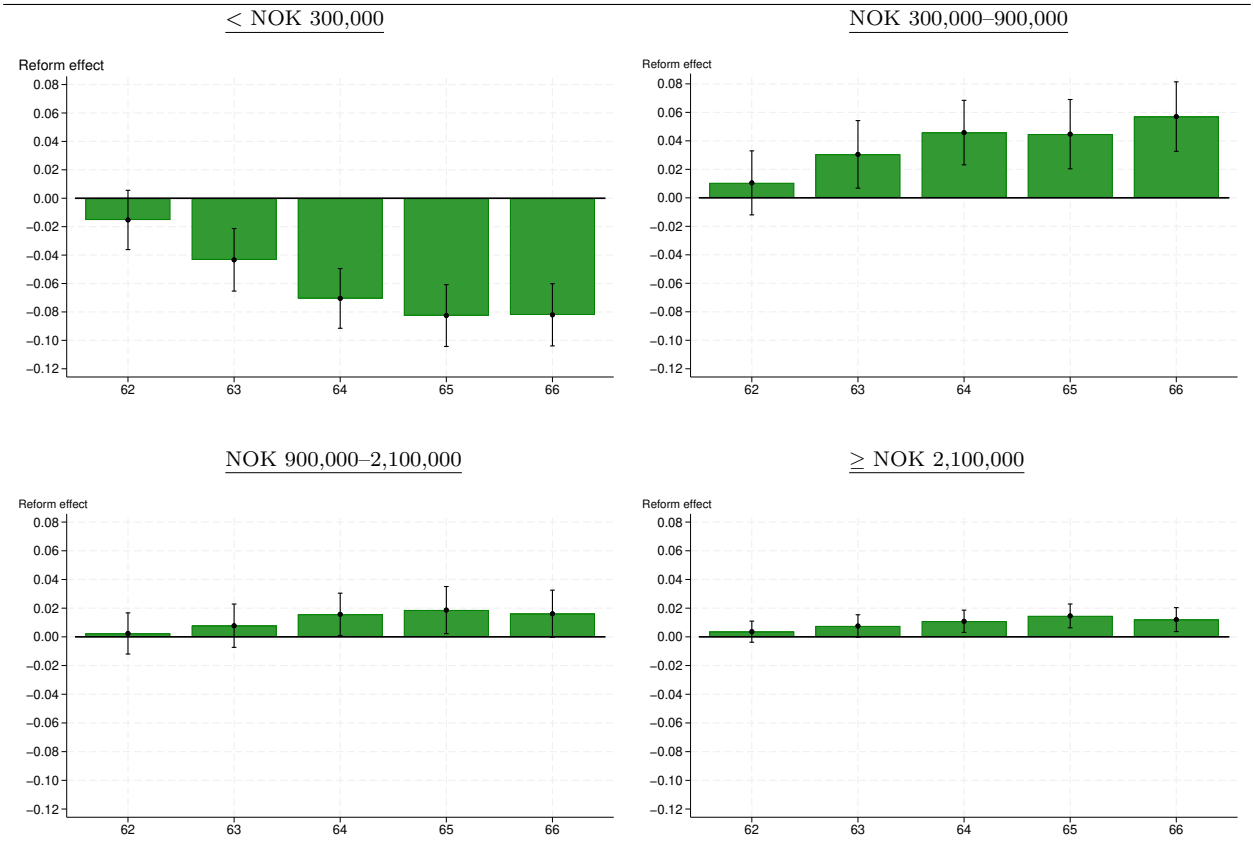

Note: The effect on the probability of being in the different consumption intervals, by age. The effects are generated using the marginal effects from Equation (4.3), with the capped lines showing the 95% confidence intervals.
Source: Authors' own calculations using data from Statistics Norway

66, and we find that the decrease in the fraction with less than NOK 300,000 in consumption and the fraction with consumption between NOK 300,000 and NOK 900,000 decreases and increases monotonically by age, respectively.

4.2.3 Measuring the welfare effects

One important result from the study by Hernæs et al. (2020) is that the flexibility reform yields a near-zero effect on total labor supply for those who were affected by the reform. This implies that we can disregard the effect due to the fiscal externality when we evaluate the welfare effect of this reform.

So far, following both Kolsrud et al. (2024) and Sæverud (2024), we have been considering the groups defined by the age when the individual actually withdraws from the labor market (retirement age). While the retirement age and claiming age are generally correlated, there is typically more variation in the retirement age than in claiming age [2]. However, we focus on the claiming age rather than the retirement age in this context, as the pension benefits are in principle defined as a function of claiming age (and not the retirement age).

4.2.4 Retirement consumption patterns by pension-claiming age

In Norway, due to the existence of different pension schemes, many individuals were able to claim pension already starting from age 62 before the 2011 reform.

[2] Note that there are no earnings tests where pension benefits are reduced for those who continue to work after claiming during our study period.

Figure 4.6 Distribution of pension-claiming age

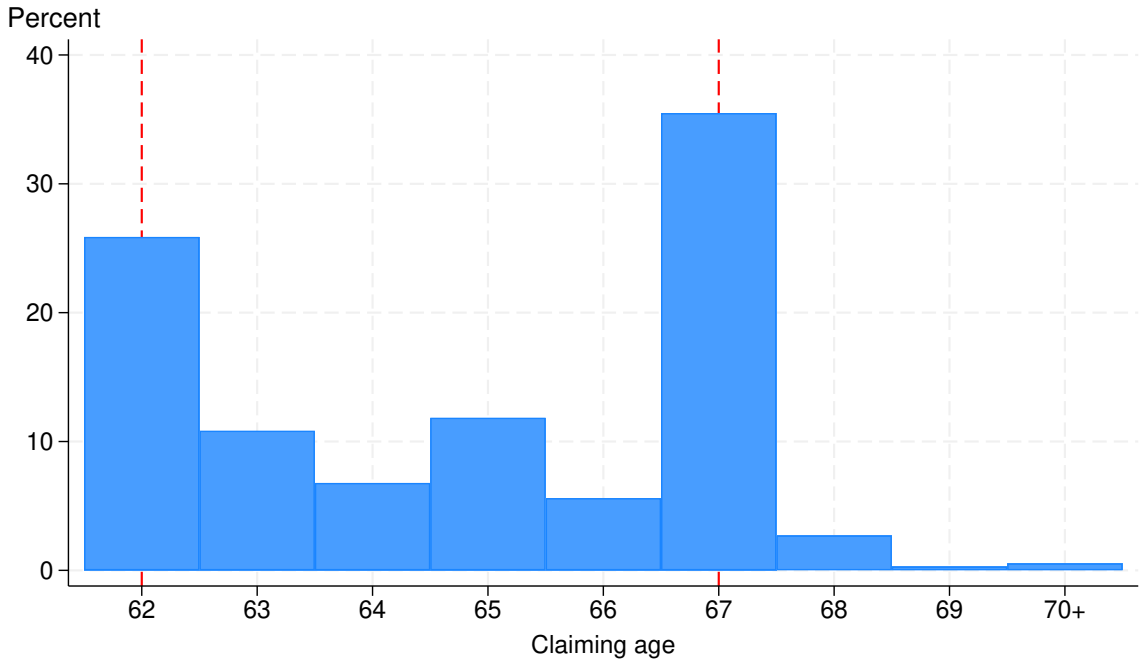

Note: The figure shows the distribution of pension-claiming ages among individuals in the sample. Claiming age is defined as the age at which old-age pension benefits were first received. The red dashed lines indicate the official pension eligibility age (67) and early claiming age (62).
Source: Authors' own calculations using data from Statistics Norway.

Figure 4.7 Consumption differences in retirement across pension-claiming age groups

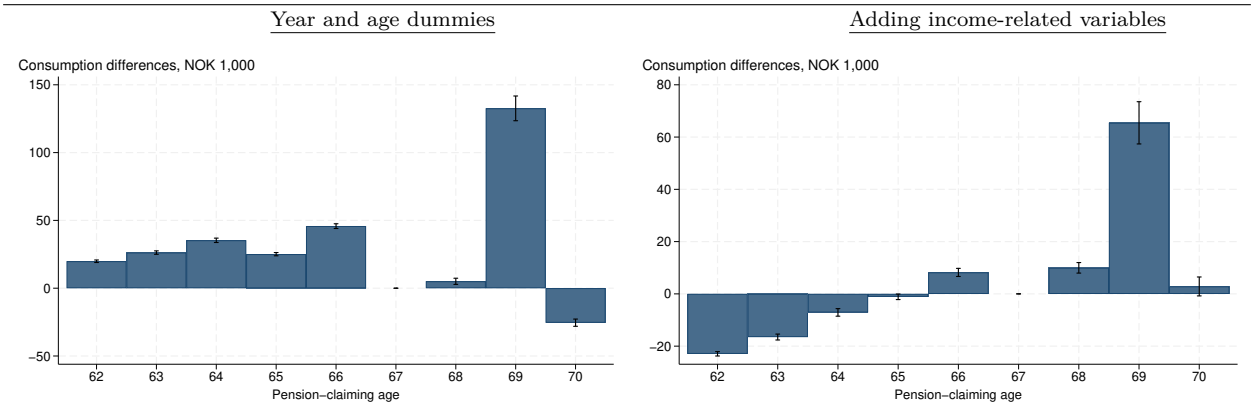

Note: The figure plots estimated coefficients α_j from equation (4.1), capturing average consumption differences across claiming age relative to claiming at age 67 (the reference group). The left panel includes controls for year, age, and gender fixed effects. The right panel additionally controls for income-related variables, including lifetime earnings. Error bars represent 95% confidence intervals.
Source: Authors' own calculations using data from Statistics Norway.

Figure 4.6 shows the distribution of pension-claiming ages. Compared with the age of retirement, there is much less variation, as we expected. The consumption differences across across pension-claiming ages are presented in Figure 4.7. In contrast to what we observed for the case of retirement age – those who claim pension at age 67 seem to have the lowest consumption of almost all groups. We suspect that this could be due to the fact that many of those who claim their pension at age 67 have already withdrawn from the labor force even if they did not have access to the pension. In other words, they are constrained by the eligibility age and lowering the eligibility age could clearly make them better off. The upward slope in consumption with age restored after we control for the pension right proxy.

4.2.5 The welfare effects

To implement the above framework of evaluating the welfare effects, we need to treat the policy change of the eligibility age as a change that steepens/flattens the profile of the NPV of pension wealth with respect to the time of claiming. It makes claiming before the new eligibility age relatively less financially attractive (either through penalties or not being able to access benefits) and can make retiring at or after the new age relatively more attractive (by shifting the point at which benefits can be received or fully accrued). In other words, reducing the eligibility age, as done by the Norwegian 2011 flexibility reform, will be treated as if it is a NPV transfer from those who claim later to those who claim earlier.

In practice, this implies that the welfare analysis here will be based on a stylized reform that redistributes resources within one birth cohort, transferring from those claiming later than age 67 to those claiming earlier. Such a reform will induce more individuals to claim earlier and adjust their labor supply accordingly, having similar behavioral impacts as the actual flexibility reform. By doing this, we can proceed as in Section 4.1. A similar approach is taken by Sæverud (2024) to study the welfare impacts of a pension reform which raised the retirement eligibility age in Denmark. However, in our case, those who claim at age 67 have substantially lower consumption after retirement than others. As an consequence, the social welfare effects of reducing the eligibility age, in terms of consumption smoothing cost, are negative. In fact, if we follow this approach and use the consumption differences reported in the left-hand side panel of Figure 4.7, transferring 1 NOK from those who claim pension at age 67 and above to those who claim pension before will induce a social cost of 0.19 NOK. Namely, increasing the flexibility actually *reduced* the social welfare.

One important issue overlooked in the implementation above is that the 2011 policy reform of lowering the access age was designed to be actuarially neutral in terms of the NPV of lifetime expected pension wealth, adjusted by the claiming age. However, for a given individual, the choice of claiming their pension benefits earlier, thereby receiving a smaller annual amount over a longer period, is indeed a decision about redistribution of their own lifetime pension wealth across time. The choice is thus about picking a different path for receiving a *given* NPV of expected lifetime benefits. In that sense, it should not considered as a redistribution across *groups*, but across different "selves". Thus, the welfare analysis presented above, particularly the "consumption smoothing" component which is fundamentally framed as evaluating redistribution across groups of individuals, cannot be readily applied to this situation. One way to get around this is to control for the "pension rights" and exploit the difference in retirement consumption between individuals who claim at different ages. In other words, to measure consumption differences across groups defined by their chosen claiming age, while ensuring these groups are otherwise comparable (e.g., same pre-retirement earnings history that determines "same total pension rights" or potential NPV before the claiming choice).

While this does not solve the fundamental problem that the flexibility reform does not redistribute across groups (individuals), one can argue that such redistribution actually happens. This is because even though the flexibility reform sets the NPV to be constant across claiming ages by design, individuals choosing different claiming ages might still exhibit differences in their actual consumption levels while retired. These differences could arise from various factors correlated with claiming choices, such as differences in health, wealth, liquidity needs, or preferences for consumption smoothing, even if the policy doesn't mechanically create wealth differences based on claiming age.

Using the above framework's consumption-level approach, we compare the consumption levels of individuals who claimed earlier versus those who claimed later. If, for example, individuals who

chose to claim earlier have lower observed consumption levels while retired, the framework would interpret this as evidence that they have a higher marginal utility relative to those who claimed later.

This would then indicate a potential consumption smoothing gain from a reform that transfers resources towards earlier claimers (or a cost from transferring away from them). For the individuals with the same expected NPV, some of them will get more while others get less, depending on the actual time of payment stop dates. If we assume a positive correlation between the claiming date and subjective life expectancy, then increased flexibility will effectively imply a transfer across groups who claim later to claim earlier. Brinch et al. (2018) studied the claiming behavior after the 2011 reform and found some individuals do claim pensions early because they gain from doing so.

Results from the right-hand side panel of Figure 4.7 can be used to evaluate the welfare impact of such a reform. This gives an estimate of -0.11. In other words, transferring 1 NOK from those who claim pension at 67 and later to those who claim earlier, given that they have the same pension rights, would increase the social welfare by 0.11 NOK. While this gives us a qualitatively "correct" welfare evaluation, some concerns remain. Firstly, it is difficult to quantify the size of transfers implied by the flexibility reform. Secondly, there is also the fact that not all individuals are impacted by the reform as many elderly were already able to claim pension from the age of 62 before the reform.

Fortunately, the above quasi-experimental evaluation of the flexibility reform presented above actually gives us an opportunity to evaluate the welfare effect by using a linear approximation, as the product of the marginal utility of consumption (MU) and the change in consumption (Δc) due to the reform, which is estimated from the quasi-experimental evidence. The welfare effect is:

$$\Delta W \approx \sum_i (MU_i \cdot \Delta c_i) \approx N \cdot \overline{MU} \cdot \overline{\Delta c}, \tag{4.4}$$

where N is the number of individuals who were affected by the reform, $\overline{MU}$ is the average marginal utility of consumption and $\overline{\Delta c}$ is the average change in consumption for those.

Similar to earlier, the welfare effect is measured in terms of the value of a 1 NOK increase in the reference group of individuals who claim the pension at age 67. This way we have:

$$\widehat{\Delta W} \approx N \cdot \frac{MU_{t<67}}{MU_{t=67}} \cdot \Delta c_{t<67} \approx N \left(1 + \gamma \frac{c_{t=67} - c_{t<67}}{c_{t=67}}\right) \Delta c_{t<67}. \tag{4.5}$$

We again assume a coefficient of relative risk aversion of $\gamma = 4$, and estimate consumption levels $c_{t=67}$ and $c_{t<67}$ using the same framework as in Equation (4.1). For the consumption change parameter Δc, we use the quasi-experimental estimate of the average change in consumption. This estimate reflects an average annual change over the age span 62–66. Since the welfare analysis requires the total change over the whole age span, we scale the estimate accordingly, yielding $\Delta c \approx$ NOK 105,000.

Substituting these parameter values into Equation (4.5), we estimate an average welfare gain of approximately NOK 138,000 per individual affected. Given that around 6,000 individuals in the 1949 birth cohort were impacted by the flexibility reform, the total welfare gain for that birth cohort amounts to approximately NOK 828 million.

5. Conclusion

This study assesses pension reforms through the trade-off between fiscal gains and consumption smoothing losses. We adopt the framework developed by Kolsrud et al. (2024) which decomposes the welfare effect into these two primary components. The consumption smoothing effect captures welfare changes from redistributing resources across groups with different social marginal utilities of consumption, while the fiscal externality reflects welfare impacts from behavioral responses affecting government finances. We apply this framework, using consumption levels imputed from Norwegian administrative data for birth cohorts born between 1940 and 1953, to quantify the consumption smoothing costs. The fiscal externality is linked to labor supply elasticity and the participation tax rate.

Our analysis first examines a hypothetical budget-neutral reform that aims to incentivize longer working lives by steepening the net present value (NPV) profile of pension wealth relative to retirement timing, increasing NPV for late retirees at the cost of early retirees. Using estimated consumption differences across retirement age groups, which reveal that average consumption tends to be higher for later retirees, we find that transferring resources from early to late retirees incurs substantial consumption smoothing costs, ranging from 0.61 to 0.70 NOK per NOK transferred. Even after considering a positive fiscal externality, potentially large welfare costs remain (around 0.4 to 0.5 NOK per NOK transferred). These findings indicate that pension designs that strongly penalize early retirement or heavily reward late retirement are regressive and impose high welfare costs on consumption smoothing. These findings highlight the importance of balanced pension designs, potentially favoring more modest incentive adjustments.

Secondly, we examine a reform involving a change in eligibility age, using the 2011 Norwegian reform as an example. This reform introduced significant flexibility by lowering the earliest claiming age from 67 to 62 for a subgroup, designed with actuarial adjustments to maintain NPV and eliminate implicit taxes on continued work. Using a quasi-experimental event-study difference-in-difference approach for affected birth cohorts, we find a near-zero effect on total labor supply for the affected group. The reform is also found to have shifted the consumption distribution upwards, increasing the share consuming above NOK 300,000 and decreasing the share below this level. While directly applying the standard welfare framework for redistribution across groups is complex for an actuarially neutral reform, a linear approximation based on the estimated consumption changes from the quasi-experimental evidence suggests that this flexibility reform provided welfare gains. For the 1949 birth cohort, the estimated welfare gain is around NOK 828 million, or roughly NOK 138,000 per affected individual.

In conclusion, we find that while reforms steepening pension wealth profiles may offer fiscal benefits, they come with substantial negative welfare effects due to consumption smoothing losses, especially when transferring resources between groups with significantly different consumption levels. In contrast, reforms like the 2011 Norwegian flexibility reform, which enhance individual choice and do not significantly impact aggregate labor supply, are found to provide welfare gains. These findings underscore the critical need to consider not only the fiscal sustainability but also the distributional and welfare implications when designing pension systems.

Bibliography

Brinch, C. N., D. Fredriksen, and O. L. Vestad (2018): "Life Expectancy and Claiming Behavior in a Flexible Pension System," *The Scandinavian Journal of Economics*, 120, 979–1010.

Brinch, C. N., E. Hernæs, and Z. Jia (2017): "Salience and Social Security Benefits," *Journal of Labor Economics*, 35, 265–297.

Fagereng, A. and E. Halvorsen (2017): "Imputing consumption from Norwegian income and wealth registry data," *Journal of Economic and Social Measurement*, 42, 67–100.

Hernæs, E. and Z. Jia (2013): "Earnings Distribution and Labour Supply after a Retirement Earnings Test Reform," *Oxford Bulletin of Economics and Statistics*, 75, 410–434.

Hernæs, E., Z. Jia, J. Piggott, and T. C. Vigtel (2020): "Work less but stay longer: Mature worker response to a flexibility reform," Tech. rep., Discussion Papers.

Kolsrud, J., C. Landais, D. Reck, and J. Spinnewijn (2024): "Retirement Consumption and Pension Design," *American Economic Review*, 114, 89133.

OECD (2017): "Pensions at a Glance 2017: OECD and G20 Indicators," OECD Publishing, Paris.

Sun, L. and S. Abraham (2021): "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 225, 175–199.

Sæverud, J. (2024): "The Impact of Social Eligibility and Pension Wealth on Retirement," Working Paper 05/24, CEBI Working Paper Series, accessed April 2025.

Appendix

Figure 5.1 Effect on consumption distribution, by year

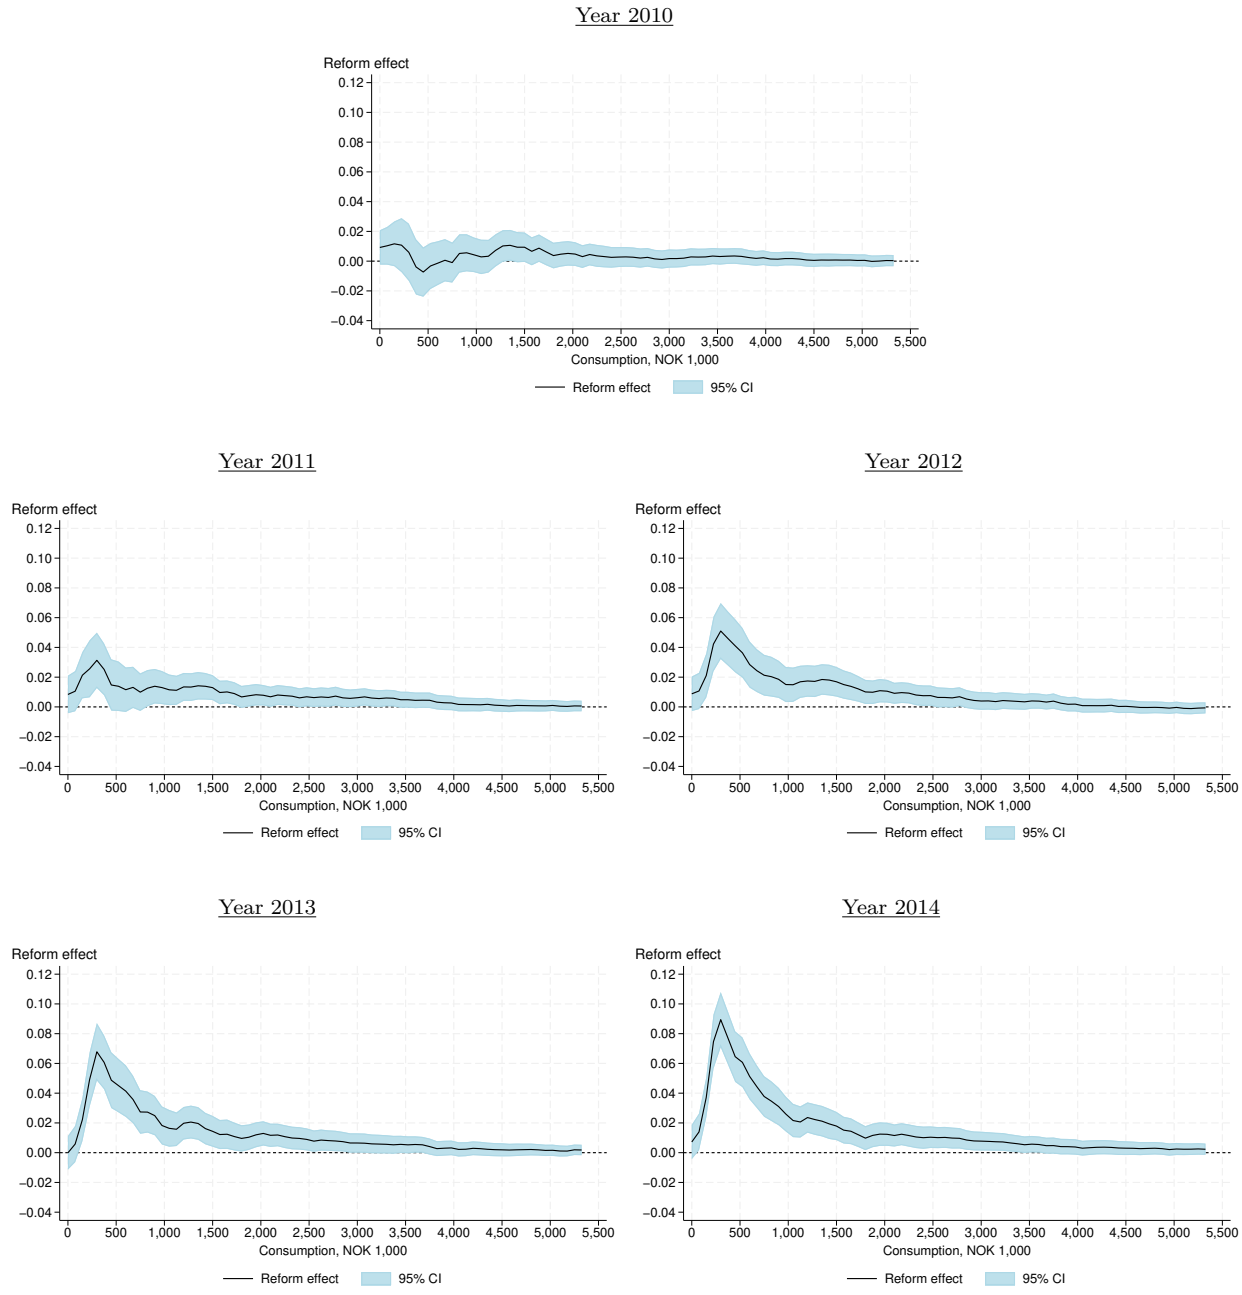

Note: The effect on the consumption distribution, by year. The effects are generated using the marginal effects from Equation (4.3), with the capped lines showing the 95% confidence intervals.
Source: Authors' own calculations using data from Statistics Norway