

Talosaga, Talosaga; Vink, Mark

Working Paper

The Effect of Public Pension Eligibility Age on Household Saving: Evidence from a New Zealand Natural Experiment

New Zealand Treasury Working Paper, No. 14/21

Provided in Cooperation with:

The Treasury, New Zealand Government

Suggested Citation: Talosaga, Talosaga; Vink, Mark (2014) : The Effect of Public Pension Eligibility Age on Household Saving: Evidence from a New Zealand Natural Experiment, New Zealand Treasury Working Paper, No. 14/21, ISBN 978-0-478-43606-8, New Zealand Government, The Treasury, Wellington

This Version is available at:

<https://hdl.handle.net/10419/205676>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

The Effect of Public Pension Eligibility Age on Household Saving: Evidence from a New Zealand Natural Experiment

Talosaga Talosaga and Mark Vink

New Zealand Treasury Working Paper 14/21

November 2014



New Zealand Government

DISCLAIMER

The views, opinions, findings, and conclusions or recommendations expressed in this Working Paper are strictly those of the author(s). They do not necessarily reflect the views of the New Zealand Treasury or the New Zealand Government. The New Zealand Treasury and the New Zealand Government take no responsibility for any errors or omissions in, or for the correctness of, the information contained in these working papers. The paper is presented not as policy, but with a view to inform and stimulate wider debate.

**NZ TREASURY WORKING
PAPER 14/21**

The Effect of Public Pension Eligibility Age on Household Saving: Evidence from a New Zealand Natural Experiment

MONTH/YEAR

November 2014

AUTHORS

Talosaga Talosaga
New Zealand Treasury
No. 1 The Terrace
Wellington
New Zealand
Email Talo.Talosaga@treasury.govt.nz
Telephone 64-4-917 6966

Mark Vink
New Zealand Treasury
No. 1 The Terrace
Wellington
New Zealand
Email Mark.Vink@treasury.govt.nz
Telephone 64-4-917 6006

ISBN (ONLINE)

978-0-478-43606-8

URL

Treasury website at November 2014:
<http://www.treasury.govt.nz/publications/research-policy/wp/2014/wp14-21>
Persistent URL: <http://purl.oclc.org/nzt/1703>

ACKNOWLEDGEMENTS

The authors wish to thank Hamish Low, and reviewers Orazio Attanasio, Adam Jaffe, Andreas Heuser and Katherine Meerman for their helpful comments and suggestions. This paper extends research originally undertaken by Mark Vink as part of the requirements for the degree of Master of Philosophy in Economics, University of Cambridge, 2013. An earlier version of this paper was presented at 55th Annual Conference of the New Zealand Association of Economists 2014, where it was awarded the Statistics New Zealand prize for the best use of statistics.

NZ TREASURY

New Zealand Treasury
PO Box 3724
Wellington 6008
NEW ZEALAND
Email information@treasury.govt.nz
Telephone 64-4-472 2733
Website www.treasury.govt.nz

Abstract

This paper examines the effect of the last increase in the eligibility age for New Zealand's public pension, New Zealand Superannuation, on household saving rates. The age of eligibility was increased progressively from 60 to 65 years old between 1992 and 2001, with little forewarning. Drawing on Household Economic Survey data, the paper uses difference-in-difference regression analysis to compare the last cohorts to receive New Zealand Superannuation at the age of 60 years old with the first to face higher eligibility ages. The policy change is found to have increased average saving rates of affected households, particularly among middle-income and older households. The increase in saving rates is associated with higher household labour supply and income, and lower expenditure. The results suggest the policy change initially lifted the aggregate household saving rate by around 2.5 percentage points with the effect declining slightly over time.

JEL CLASSIFICATION

D14: Household saving
D91: Life cycle models and saving
E21: Consumption; Saving
H55: Social security and public pensions
J26: Retirement Policies

KEYWORDS

Household Saving; Retirement Income; New Zealand Superannuation

Executive Summary

The New Zealand Treasury has argued that a reduction in the generosity of New Zealand's public pension, New Zealand Superannuation (NZS), would lead to an increase in household saving rates. This prediction is based on the simple life-cycle theoretical model of saving in which individuals smooth consumption over time. In this simple model, households respond to an unexpected decrease in public pension wealth by saving more prior to retirement, through working more and/or consuming less. However, an unexpected reduction in public pension wealth can also, in theory, lead to lower household saving rates. This may occur if individuals respond by delaying retirement from the workforce, thereby reducing the level of accumulated savings needed at retirement. An assessment of how household saving rates may respond to changes in the generosity of NZS therefore requires the support of empirical evidence. Such evidence has been lacking.

The objective of this paper is to measure the effect of the last increase in the NZS eligibility age on household saving behaviour. The age of eligibility was increased progressively from 60 to 65 years old between 1992 and 2001, with little forewarning. Because this policy change affected different birth cohorts differently, it provides a "natural experiment" that allows identification of the policy's impact on affected households. Drawing on data from the Household Economic Survey (HES) for the period 1984 to 1998, the paper uses difference-in-difference regression analysis to compare the average saving rates of the last cohorts to receive NZS at 60 years old with those of the first cohorts facing higher eligibility ages. To provide insight into the nature of the saving response, the same method is used to measure the policy's effects by year, age and income quintile, as well its effects on household income, expenditure and labour supply.

The analysis suggests the policy change caused an average increase in household saving rates of around 2 percentage points for each additional year added to the eligibility age faced by the household head. For households facing an eligibility age of 65 years old, this translates to a 10 percentage point increase in the average annual saving rate. The impact of the policy change on saving rates is evident from the first year after its announcement in 1990, and appears to have been greatest on older and middle-income households. The findings for saving rates are supported by the estimated effects of the policy change on disposable income, expenditure, and labour supply. Disposable income and labour supply are found to have increased as a result of the policy change, while the effect on expenditure appears to be negative. The positive effect on disposable income and labour supply is, again, greatest for older households.

These results suggest the lift in the eligibility age led to an initial increase in the aggregate household saving rate of around 2.5 percentage points, declining to around 2 percentage points by the year ended March 1998. To what extent would future reductions in the generosity of NZS lead to similar changes in aggregate household saving rates? The absence of a tax surcharge (an income test) is one important difference of NZS today from the scheme in the 1990s. All else being equal, this difference would likely increase the relative effects of a future reduction in the generosity of NZS. Two other important factors relate to the amount of forewarning provided and the nature of household expectations. To the extent that any future changes are signalled well in advance of implementation, households will have a longer period of time to adjust, leading to a smaller increase in annual saving rates. However, if future changes are already built into household expectations, the effects on saving will be reduced.

Table of Contents

Abstract	i
Executive Summary	ii
1 Introduction	1
2 Data	2
2.1 The data and cohort construction	2
2.2 Defining saving and the saving rate.....	2
3 The Policy Context	4
3.1 An overview of New Zealand Superannuation	4
3.2 The reforms.....	4
4 Empirical method: difference-in-differences estimation	5
4.1 The basic difference-in-differences set up.....	5
4.2 The difference-in-differences regression model	8
4.3 The common trend assumption: tests in the pre-treatment period	9
4.4 Adding interaction terms to the model	10
5 Saving Results	10
5.1 The raw data	11
5.2 Average treatment effect	11
5.3 Breakdown of effects by year, age, and income quintile	13
5.4 Summary and approximate impact on aggregate saving rates	17
6 Additional Results	17
6.1 Average income and expenditure effects	18
6.2 Breakdown of income and expenditure effects by year, age and income quintile	19
6.3 Labour supply effects.....	22
7 Sensitivity Analyses	24
7.1 An alternative measure of saving	24
7.2 Saving measured at the household level.....	25
8 Conclusion	26
References	28
Appendix A: HES Saving Measures	30
Appendix B: Additional Tables	31

List of Tables

Table 1 – Official NZS eligibility ages and the treatment variable.....	7
Table 2 – Average saving rate effects: regression results (sample period 1984 to 1998).....	12
Table 3 – Average saving rate effects: robustness checks over the pre-treatment period (1984 to 1990).....	13
Table 4 – Average saving rate effects: regression results (sample period 1984 to 1994).....	15
Table 5 – Average income & expenditure effects: regression results (sample period 1984 to 1998)	18
Table 6 – Average income & expenditure effects: robustness checks over the pre-treatment period (1984 to 1990).....	19
Table 7 – Average saving rate effects with alternative saving rate measure: regression results (sample period 1984 to 1998)	25
Table 8 – Average saving rate effects using household-level data: regression results (sample period 1984 to 1998).....	26

Appendix Table 1 – Sample size of each cohort by year.....	31
Appendix Table 2 – Average income effects: robustness checks over the pre-treatment period (1984 to 1990).....	32
Appendix Table 3 – Average expenditure effects: robustness checks over the pre-treatment period (1984 to 1990).....	33
Appendix Table 4 – Labour supply effects: robustness checks over the pre-treatment period (1984 to 1990).....	34

List of Figures

Figure 1 – NZS eligibility policies and affected cohorts	6
Figure 2 – Control and treatment group cohorts: age by survey year.....	7
Figure 3 – Mean saving rates: control group and before- and after- treated groups (sample period 1984 to 1998).....	11
Figure 4 – Saving rates: estimated treatment effects by year (sample period 1984 to 1998)	14
Figure 5 – Saving rates: estimated treatment effects by age group (sample period 1984 to 1998)	16
Figure 6 – Saving rates: estimated treatment effects by income quintile (sample period 1984 to 1998).....	17
Figure 7 – Income and Expenditure: estimated treatment effects by year (sample period 1984 to 1998).....	20
Figure 8 – Income and expenditure: estimated treatment effects by age (sample period 1984 to 1998).....	21
Figure 9 – Income and expenditure: estimated treatment effects across the income distribution (sample period 1984 to 1998).....	22
Figure 10 – Average weekly hours worked: estimated treatment effects by year (sample period 1984 to 1998).....	23
Figure 11 – Average weekly hours worked: estimated treatment effects by age (sample period 1984 to 1998).....	24

The Effect of Public Pension Eligibility Age on Household Saving: Evidence from a New Zealand Natural Experiment

1 Introduction

It has been argued that a reduction in the generosity of New Zealand's public pension, New Zealand Superannuation (NZS), would lead to an increase in household saving rates¹. This prediction is a basic result of the simple life-cycle model of saving in which individuals attempt to smooth consumption over time. In this simple model, households respond to an unexpected decrease in public pension wealth by saving more prior to retirement, through working more and/or consuming less. However, as argued by Feldstein (1974), a reduction in public pension generosity could in fact lead to lower household saving if individuals choose to delay their retirement and shorten the period over which they draw down accumulated savings. An assessment of how household saving rates actually respond to changes in the generosity of NZS therefore requires the support of empirical evidence. Such evidence has been lacking.

The objective of this paper is to measure the effect of the last increase in the NZS eligibility age on household saving behaviour. The age of eligibility was increased progressively from 60 to 65 years old between 1992 and 2001, with little forewarning. Because this policy change affected different birth cohorts differently, it provides a "natural experiment" that can be exploited to identify the policy's impact on affected households. Drawing on data from the Household Economic Survey (HES) for the period 1984 to 1998, this paper uses difference-in-difference regression analysis to compare the average saving rates of the last cohorts to receive NZS at 60 years old with those of the first cohorts facing higher eligibility ages. To provide insight into the nature of the saving response, the paper uses the same method to analyse the policy's effects by year, age and income quintile, as well its effects on household income, expenditure and labour supply.

The paper is structured as follows. Section 2 introduces the data and saving measures. Section 3 provides background on NZS and its reform. Section 4 describes the empirical methods. Sections 5 and 6 outline the main results relating to saving, and to income, expenditure and labour supply. Section 7 presents sensitivity analyses and Section 8 concludes.

¹ Recent examples include Brook (2014), Law (2013) and New Zealand Treasury (2010).

2 Data

This section provides a short summary of the data and the construction of cohorts in Subsection 2.1, and of the measures of saving used in Subsection 2.2. A more detailed discussion can be found in Vink (2014), which uses the same data and measures.

2.1 The data and cohort construction

This paper uses the 15 annual March-year HES surveys from 1983/84 to 1997/98. Each of the surveys provides a rich set of income, expenditure, labour market and demographic data from an independent and representative sample of New Zealand private households. The individual surveys are referred to henceforth by the year in which they ended.

Saving rates are calculated at the household level, rather than the individual level, reflecting an assumption that expenditure and saving decisions are made on a household basis. Calculating saving at an individual level is difficult because HES only records expenditure at the level of the household. The household head is defined as the household member with the highest gross income.² Characteristics of the household head are used to define each household's age, gender and ethnicity.

Ideally panel data, which consists of observations from the same household units over time, would be used for analysing households' saving behaviour over their lifetimes. In the absence of such data for New Zealand, "synthetic panel data" are constructed by dividing the sample into cohorts determined by the birth year of the household head, following a method described by Deaton (1985).³ Using birth cohorts (henceforth "cohorts") as the unit of analysis allows the average behaviour of these cohorts to be tracked over time.

The use of synthetic panels is valid if the membership of the population (and each cohort within it) is fixed. This may not be the case if there is a large amount of household dissolution and reformation, in which, for example, older people move in with their children so previously "old" households become "young" households in subsequent years. This problem cannot be fully overcome, but to reduce its impact, the sample is restricted to households with heads aged 70 years old or younger.

A similar problem may occur when there are common age-related changes to income, such as retirement or NZS receipt. If these age-related changes affect the relative earnings of household members on a systematic basis, they may also systematically influence the identity of the household head. To provide reassurance that the results are not being biased in this way, two alternative definitions of household head were trialled: the oldest household member and the Statistics New Zealand's "reference person".⁴ These alternatives are unlikely to have the same systematic age-related influences as highest earner. Neither materially affected the results.

2.2 Defining saving and the saving rate

The HES was not designed for the measurement of saving and there are important limitations to using it for this purpose, as discussed by Vink (2014). However, HES is the best source of household-level saving data available and it has been used in other studies

² If two individuals have equal highest income, the older individual is chosen as head. If household members have equal income and age, the individual with the lowest HES person number is chosen as head.

³ A person's year of birth is calculated as the survey year minus their age (in years) if they were surveyed in the second half of the survey year, or equal to the survey year minus their age, minus one if they were surveyed in the first half of the survey year.

⁴ The Statistics New Zealand reference person is the person who fills out the household's questionnaire. No particular criteria are used in selecting the reference person.

to analyse household saving.⁵ The preferred measure of household saving in this paper has been chosen to correspond closely to the definition used in New Zealand’s Household Income and Outlay Account (HIOA), which provides the most frequently-cited aggregate household saving measure for New Zealand. Specifically, household saving is defined as:

$$S = YD - C, \quad (1)$$

where: YD = HES “total income”
 - net taxes and transfers⁶
 C = HES “total household expenditure”
 - HES “contributions to savings”
 - HES “mortgage principal payments”
 - HES “life and health insurance payments”
 - HES “purchases of property”
 + HES “sale of property” (classified as negative expenditure in HES)

In other words, saving is defined as the difference between household disposable income and expenditure, plus mortgage principal payments and life and health insurance payments.⁷ To test whether the results are sensitive to this definition, an alternative saving measure is adopted in Subsection 7.1. More detail on both saving measures is contained in Appendix A.

Saving rates are calculated for each cohort as the ratio of that cohort’s mean saving to its mean disposable income:

$$s_t^c = \frac{\frac{1}{N_t^c} \sum_h (YD_t^{ch} - C_t^{ch})}{\frac{1}{N_t^c} \sum_h YD_t^{ch}}, \quad (2)$$

where:

N_t^c = the number of household observations in cohort group c in survey year t ;

C_t^{ch} = expenditure of household h , belonging to cohort c , in survey year t ;

YD_t^{ch} = disposable income of household h , belonging to cohort c , in survey year t .

The aggregation properties of this “ratio of means” measure (as opposed to a “mean of ratios” measure) are useful when considering the implications of results for aggregate household saving, which is similarly calculated as total household saving divided by total household disposable income. The ratio of means measure also reduces the influence of outliers and measurement error. A disadvantage of ratio of mean measures is their limited use for understanding behaviour at the household level. In Section 7.2, the analysis is repeated using the household as the unit of analysis, rather than the cohort mean. This specification allows for the control of household characteristics.

⁵ Other studies that use the HES to analyse saving behaviour include Gibson and Scobie (2001) and Coleman (2006).

⁶ The HES does not record income tax data. An adjusted data set is used, which incorporates New Zealand Treasury estimates of net taxes and transfers at the household level

⁷ These are considered financial transactions in the HIOA, and are therefore not part of consumption.

3 The Policy Context

This section describes the main features of NZS and its reform, which are relevant to the change in eligibility age considered by this paper.⁸ Much of the detail is drawn from Preston (2001), unless indicated otherwise.

3.1 An overview of New Zealand Superannuation

NZS is New Zealand's defined benefit public pension scheme, funded from general taxation on a pay-as-you-go basis. At the time the last increase in the eligibility age of NZS was announced by the Government in July 1991, virtually all New Zealanders over the age of 60 years old were eligible to receive NZS. The net rate of payment for couples was approximately 72 per cent of the net average wage, with annual adjustments ensuring that it remained within a "wage band" of 65 to 72.5 per cent. A "tax surcharge" of 20 per cent was applied to non-NZS income (with some exemptions) and affected approximately 30 per cent of NZS recipients (Boston, Dalziel *et al* 1999). Apart from the tax surcharge there were no means- or work-testing criteria for receiving NZS. Private pension schemes played a minor role in retirement income, accounting for around 7 per cent of the total income of those aged 60 years old or older in the year ended March 1991, as recorded by HES.

3.2 The reforms

There were growing concerns during the 1980s about the high and rising costs of funding NZS. In response to these concerns, the Government introduced the 65 to 72.5 per cent wage band in 1989. This wage band applied to all current and future NZS recipients and represented a significant reduction from the prevailing NZS-wage ratio of 80 per cent. At the same time the Government signalled an increase in the NZS age of eligibility from 60 to 65 years old that was to take place progressively over a 20 year period beginning in 2006. Under this policy, individuals born in 1947 or later would face higher eligibility ages, while those born earlier would be unaffected.

In June 1990, the National Party, then in opposition, announced its intention to raise the age of eligibility to 65 years old, over a 20 year period beginning in 1992 (New Zealand Herald 1990). This policy affected individuals born in 1933 or later, with those born from 1943 facing the new eligibility age of 65 years old. At the time of the announcement the National Party was polling very strongly and it went on to win the national election in November 1990 (Aimer and Vowles 1993). However, escalating concerns about the state of the Crown's financial situation led the Government to implement a more severe package of reform than it had campaigned on, including substantial welfare benefit cuts (Dalziel and Lattimore 2001). Three main changes to NZS were announced in the 1991 Budget:

- (i) An increase in the NZS eligibility age from 60 to 65 years old between 1992 and 2001. The schedule saw the eligibility age increase to 61 for those born in between 1 April 1932 and 30 June 1932, with additional increases of 3 months for each successive 3-month birth cohort (the schedule is set out in Subsection 4.1).
- (ii) The cancellation of NZS rate-of-payment adjustments in 1991 and 1992. Consumer price inflation adjustment was to be made thereafter provided the rate was above the lower bound of the wage band, which remained at 65 per cent of the net average wage for couples. These changes affected all current and future NZS recipients.

⁸ New Zealand's public pension has previously been called National Superannuation and Guaranteed Retirement Income, but is referred to consistently in this paper as NZS.

- (iii) An increase in the tax surcharge to 25 per cent and a reduction in the income exemption threshold.⁹ The proportion of NZS recipients who were subject to the surcharge continued to fluctuate at around 30 per cent until the surcharge was reduced in 1997 and abolished in 1998 (Boston, Dalziel *et al* 1999).

In 1994 the Transitional Retirement Benefit (TRB) was introduced to assist those affected by the increase in the NZS eligibility age. The age of eligibility for the TRB was 60 years old in 1994 and was increased at the same rate as the NZS eligibility age, reaching 65 years old in 2004, at which point it was phased out. Individuals were therefore potentially eligible for the TRB for three years before becoming eligible for NZS. The TRB was set at roughly the same amount as NZS and it was subject to the same tax surcharge. Although the TRB would moderate any effects of the change in NZS eligibility age on saving behaviour, its overall effects are likely to have been small because only a small proportion of potentially eligible individuals actually received the TRB.¹⁰

4 Empirical Method: Difference-in-differences Estimation

This section outlines the basic difference-in-differences framework and regression model in Subsections 4.1 and 4.2 respectively. Subsection 4.3 describes some of the robustness tests employed, and Subsection 4.4 covers extensions to the basic regression model.

4.1 The basic difference-in-differences set up

Difference-in-differences is a common technique used to measure the effects of policy changes in natural experiments.¹¹ A policy change provides a natural experiment if it is exogenous and affects some population groups (a “treatment” group) but not others (a “control” group). Because the change in the NZS eligibility age was not implemented in response to concerns about household saving rates and because it affected different cohorts differently, the policy change can be considered a natural experiment.

Difference-in-differences analysis compares changes in the outcomes for treatment and control groups before and after the policy change to identify the policy’s effect (the “treatment effect”), on the assumption that without the change the change in outcomes for the two groups would be the same. This assumption is known as the “common trends” assumption and is critical to the difference-in-differences identification strategy. For this study, the common trends assumption requires the effects of macroeconomic developments and other policy changes on the treatment and control groups to be the same.

Figure 1 shows how the policies announced in 1989, 1990 and 1991 affected the NZS eligibility age applying to different cohorts. Three different groups can be identified as follows:

- (i) the control group: cohorts born before 1932, who were not affected by any of the policies;

⁹ A larger increase in the surcharge was initially announced, but reduced later in response to public pressure.

¹⁰ The Ministry of Social Development (2002) reports that only 6 per cent of 60 to 64 year olds received the TRB between 1998 and 2001.

¹¹ Recent examples that consider the effects of public pension reform on saving include Attanasio and Brugiavini (2003), Attanasio and Rohwedder (2003) and Aquila (2011).

- (ii) the treatment group: cohorts born between 1932 and 1946, who were not affected by the 1989 policy but were affected by the policies announced in 1990 and 1991; and
- (iii) excluded group: cohorts born in 1947 or later, who were affected by all three policies.

Because cohorts in the third group were affected by the 1989 policy, they are excluded from the analysis in order to isolate the effects of the policies announced in 1990 and 1991.

Figure 1 – NZS eligibility policies and affected cohorts

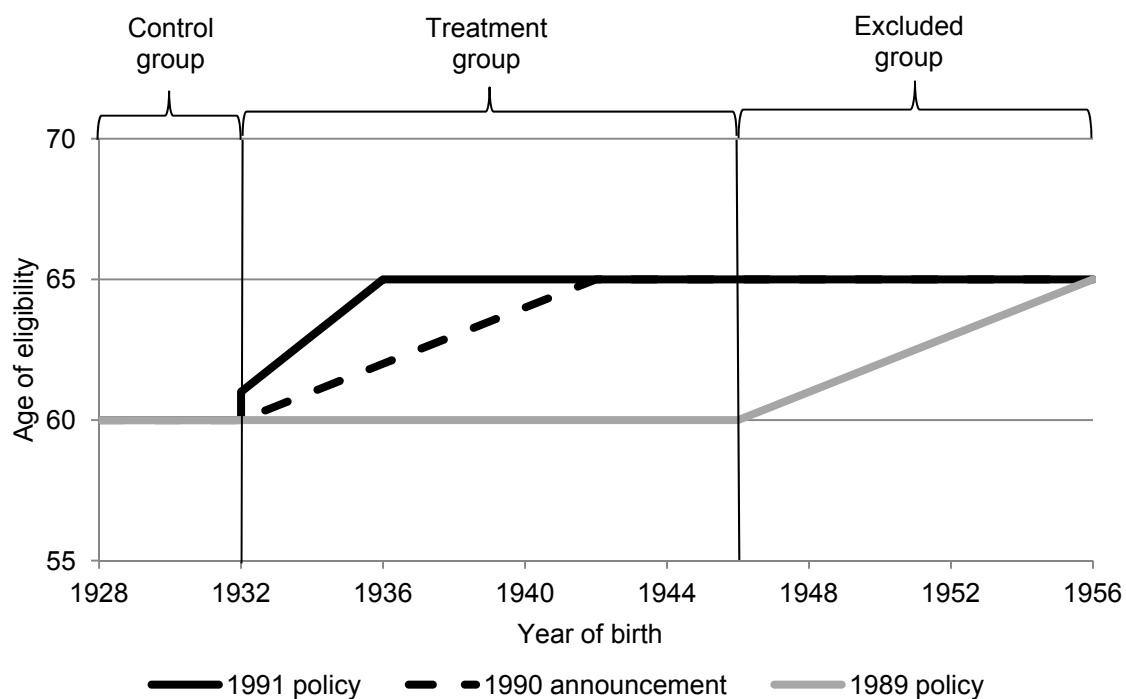
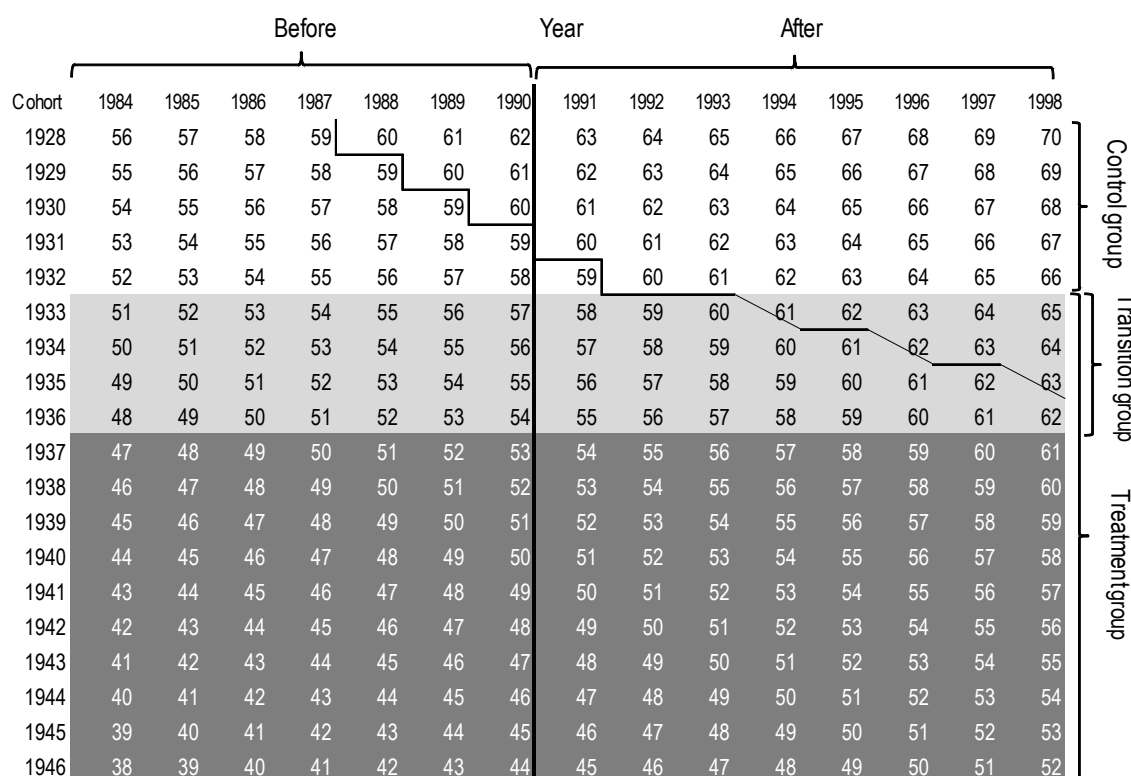


Figure 2 shows the age of each cohort in the treatment and control groups by year. Appendix Table 1 shows the number of household observations in each of these cohort-year “cells”. Cohorts born before 1928 are excluded from the control group to ensure that each cohort is represented in each survey year, given the sample restriction excluding households with heads older than 70 years old. By narrowing the range of birth years, excluding older cohorts from the control group also reduces the potential for bias relating to any nontreatment-factors that affected different cohorts differently.

The treatment is defined as taking effect from the National Party’s June 1990 announcement that, if elected, it would increase the NZS eligibility age. This announcement was close to coinciding with the commencement of the 1991 HES survey (in April 1990), so the before/after demarcation is set between the 1990 and 1991 surveys.

Figure 2 – Control and treatment group cohorts: age by survey year



Diagonal line indicates the age at which each cohort is eligible for NZS.

Cohorts born between 1933 and 1936 faced transitional NZS eligibility ages of between 61 and 65 years old, while later-born cohorts in the treatment group all became eligible for NZS at 65 years old. To capture these differences in treatment intensity, the treatment is defined as the mean number of additional years that each treated cohort had to wait before being eligible for NZS, ie, *new eligibility age – 60*.¹² The official schedule for the rise in the eligibility age and the corresponding treatment definitions for each cohort in the treatment group are shown in Table 1.¹³

Table 1 – Official NZS eligibility ages and the treatment variable

Scheduled increase in NZ eligibility age			Definition of treatment variable		
Birth date (quarter ending)	NZS age of eligibility	Year of NZS receipt	HES cohort	NZS average eligibility age	Treatment variable
Mar-1932	60	1992	1932 and earlier	60	0
Jun-1932	61	1993	1933	61.375	1.375
Jun-1933	62	1995	1934	62.375	2.375
Jun-1934	63	1997	1935	63.375	3.375
Jun-1935	64	1999	1936	64.375	4.375
Jun-1936	65	2001	1937 and later	65	5

¹² Treatment intensity could alternatively be defined as the change in the discounted value of future NZS payments. However, because of the tax surcharge at place at the time, this would require estimating each household's future non-NZS income, in addition to their life expectancies. Such estimation would likely introduce more error into the treatment intensity variable, so the simpler definition is used.

¹³ A more precise definition would set the treatment effect in 1990 differently using the eligibility age proposed in the National Party's 1990 announcement. This definition has an immaterial effect on the results and is excluded for presentational clarity.

Most of the analysis uses the full data set from 1984 to 1998, giving seven survey years in the pre-treatment period and eight years in the treatment period. The analysis is completed on this long sample period to ensure a high degree of age overlap between the control and treatment groups, which allows age to be controlled for as discussed in the next section. The disadvantage of using a longer sample period is that it increases the possibility of the common trends assumption being violated, because there is higher chance that other policy or macroeconomic changes occur within the sample period that affect the treatment and control groups differently. These potential biases are explored in Section 5.3 and are found to have little impact on the results.

4.2 The difference-in-differences regression model

Because different cohorts are at different ages in each survey period, it is necessary to control for the effect of age to account for life-cycle patterns in behaviour. This is done by implementing the difference-in-differences estimation in a regression framework with age included as an explanatory variable. Dummy variables for year are included to control for macroeconomic shocks and to improve the efficiency of the estimation. Similarly, cohort dummies are included to control for differences in behaviour between cohorts. However because the treatment itself is defined by cohort, it is possible for cohort dummies to capture part of the treatment effect and thereby bias estimates downward. The collinearity between age, cohort and year prevents year and cohort variables being included in the same regression with age, so these two sets of dummy variables are included alternately.¹⁴ Formally, the estimated difference-in-differences regression model is as follows:

$$s_t^c = \alpha + f(\text{age}_t^c) + \beta_1 \text{Treat}^c + \beta_2 \text{After}_t + \beta_3 (\text{Treat}^c \times \text{After}_t) + \underbrace{\lambda_1 D_t + \lambda_2 D^c}_{\text{Included alternately}} + u_t^c, \quad (3)$$

where:

- s_t^c = the saving rate of cohort c in year t , replaced by variables for disposable income YD_t^c , expenditure C_t^c and labour supply LS_t^c for the specifications presented in Section 6;
- $f(\text{age}_t^c)$ = a function representing a quadratic in age¹⁵;
- Treat^c = the treatment variable, as defined in Table 1;
- After_t = a dummy variable equal to one for years 1991 to 1998, and zero otherwise;
- D_t = a vector of year dummies;
- D^c = a vector of cohort dummies;

¹⁴ Methods have been developed to identify age, cohort and year effects using additional assumptions, such as those used in Vink (2014). These methods cannot be used in this context over the full sample period because treatment status varies only by cohort. The possibility of estimating age and cohort effects using pre-treatment data to apply over the full sample is also ruled out by the short pre-treatment period available.

¹⁵ For the saving, income and expenditure estimations, an age quadratic is preferred to a set of age dummy variables for the purposes of parsimony and presentational clarity. Replacing the age quadratic with dummy variables does not alter the key findings. However, age dummy variables are used for estimating labour supply effects to capture the (larger) discontinuities in labour supply at the age of retirement.

β_1	=	the estimated difference between the treatment and control groups that is constant over time;
β_2	=	the estimated post-treatment effect that is common to both the treatment and control groups;
β_3	=	the difference-in-differences estimator of the treatment effect, ie the average effect on the saving rate over the treatment period effect of each additional year added to the eligibility age;
λ_1	=	a vector of estimated year effects;
λ_2	=	a vector of estimated cohort effects; and
u_t^c	=	the residual.

By including variables to control for age, year and/or cohort, it is assumed that the effect of birth cohort does not vary with time or age, the effect of age does not vary with time or cohort, and the effect of time does not vary by age or cohort (the common trends assumption). The validity of the common trends assumption can be examined in both the pre-treatment and treatment periods as discussed in the following two subsections.

4.3 The common trend assumption: tests in the pre-treatment period

The presence of any difference in time trends in the pre-treatment period between the treatment and control group would raise significant concerns about the validity of the difference-in-differences results. Two standard robustness or “falsification” tests are used in the literature to test the common trends assumption over the pre-treatment period. The first attempts to directly estimate any difference in trends with the following regression model¹⁶:

$$s_t^c = \alpha + f(\text{age}_t^c) + \gamma_1 \text{trend} + \gamma_2 (\text{Treat} \times \text{trend}) + u_t^c; \quad t < 1991 \quad (4)$$

where *trend* is a time trend variable and (*Treat* × *trend*) is a variable that interacts the time trend with the treatment variable. The coefficient γ_1 captures any trend that is common to the treatment and control groups, while γ_2 captures any difference in time trends between the two groups. If the common trends assumption is valid, the estimated coefficient of γ_2 should be close to zero and statistically insignificant.

The second test estimates the model using a “placebo” treatment. This involves re-estimating the difference-in-differences model over the pre-treatment period, but with the assumption that the treatment took effect at an earlier date. Since this treatment precedes the announcement of the policy change investigated, the difference-in-difference estimator should be statistically insignificant and close to zero.

Two separate placebo treatments are considered. In both cases the policy announcement is assumed to occur in 1987, which is the mid-point of the pre-treatment period. In the first case, the placebo policy is the same as the actual policy, so that the only difference is the timing of the announcement. In the second case, the same increase in the eligibility age is assumed to occur, but beginning four years earlier in 1988.

¹⁶ The dependent variable in Equations 4 to 6 is shown as the cohort saving rate, however the same models are used with income, expenditure and labour supply as the independent variables.

4.4 Adding interaction terms to the model

A further check on the validity of the common trends assumption involves interacting the treatment effect with the year dummy variables to provide an estimate of the policy effect by year, as shown by Equation 5:

$$s_t^c = \alpha + f(\text{age}_t^c) + \beta_1 \text{Treat}^c + \beta_3 (\text{Treat}^c \times \text{After}_t \times D_t) + \lambda_1 D_t + u_t^c, \quad (5)$$

where β_3 is a vector of coefficients capturing the treatment effect in each of the post-treatment years.¹⁷ Given the hypothesis that the treatment effect should be broadly constant over time after controlling for age (and ignoring any adjustment lags), any changes in the treatment effect over time could be indicative of some bias in the estimation.

As an additional check, the before/after treatment demarcation (set by the After_t variable) can be set earlier than the actual policy announcement, following Autor (2003). This technique tests whether any pre-treatment effects are evident, the existence of which would suggest bias in the estimated treatment effects (in the post-treatment period). Of particular interest is whether any effects from the 1989 adjustment to the NZS-wage ratio may be captured in the treatment effect.

Interaction terms can also be added to Equation 3 to provide estimates of the treatment effect by age, as shown in Equation 6:

$$s_t^c = \alpha + f(\text{age}_t^c) + \beta_1 \text{Treat}^c + \beta_3 (\text{Treat}^c \times \text{After}_t \times \text{ageband}_a) + \lambda_1 D^c + u_t^c, \quad (6)$$

where ageband_a is a set of five-year age band dummies and β_3 is a vector of coefficients capturing the treatment effect for each age group.¹⁸ As households age over time, the households within each age band will be drawn from earlier cohorts, which may lead to bias if the time and age constant cohort effects are nonzero. A vector of cohort dummies (D^c) is included in Equation 6 to control for this potential source of bias.

5 Saving Results

This section presents the estimated effects of the increase in the NZS eligibility age on saving rates. To provide context for the subsequent regression estimates, Subsection 5.1 first presents the average saving rates of the treatment and control groups graphically. Subsection 5.2 reports difference-in-differences regression results for the average effect on saving rates and Subsection 5.3 provides a breakdown of this effect by year, age and income quintile. Subsection 5.4 summarises the findings of the preceding subsections and considers the effect on the aggregate household saving rate.

All of the regression results presented are estimated using weighted least squares, with the weights equal to the number of household observations in each year-cohort cell. This weighting method provides an efficient way of estimating parameters when using cell averages by accounting for the greater variance in cells with few observations. Additionally, Donald and Lang (2007) show standard errors can be underestimated if residuals are correlated within cohort or year. Using tests developed by Silva and Parente (2013), the residuals in all the estimated specifications reported in this and the following section were found to be correlated within years but not within cohorts. In recognition of these test results, the p-values shown in all tables are calculated to be robust to within-year correlation using the methods of Rogers (1994).

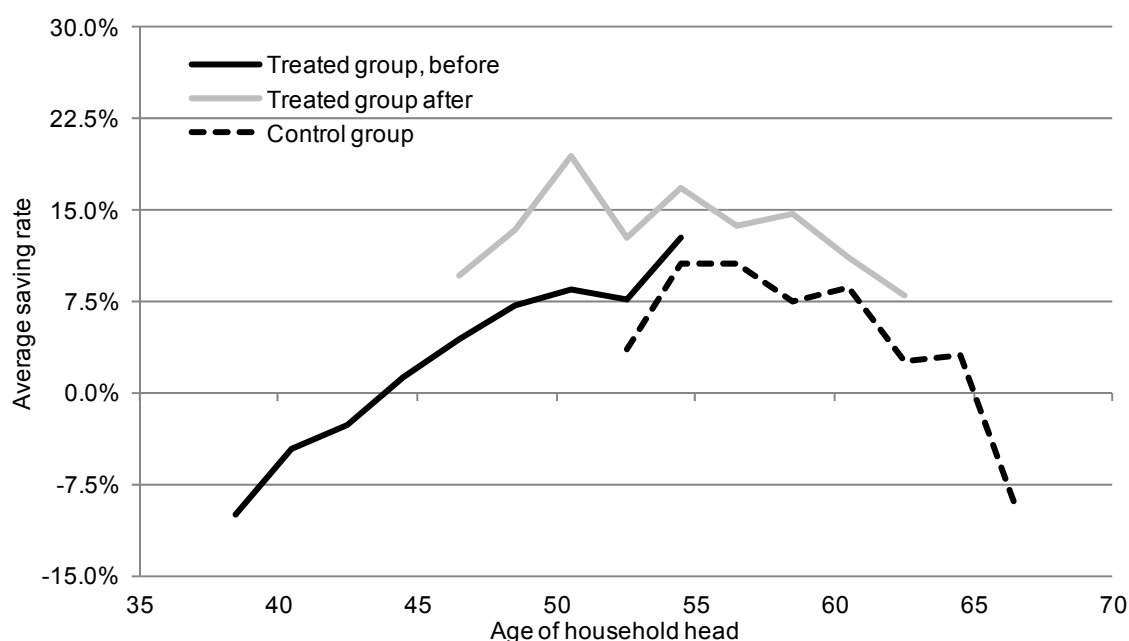
¹⁷ The After dummy variable is excluded from Equation 5 because of collinearity.

¹⁸ Age bands dummies are used instead of individual age dummies to reduce the volatility of each estimate.

5.1 The raw data

Figure 3 compares the average saving rates of the treatment group, before and after the treatment took effect, with the average saving rates of the control group. Each data point represents the mean of the saving rates for all cohorts in the relevant sample group at that particular age. To remove some of the volatility in the data, age is defined in two-year bands, so that the data point at age 60.5 years old includes cohorts at age 60 and 61 years old.¹⁹ Although there are no controls for the effects of time, cohort or treatment intensity, the figure shows a “jump” up in the saving rates of the treated group after the treatment, which is consistent with the policy having increased the saving rates of affected cohorts.

Figure 3 – Mean saving rates: control group and before- and after- treated groups (sample period 1984 to 1998)



5.2 Average treatment effect

Table 2 provides several estimates of the average treatment effect on the annual saving rates of all households in the treatment group. In the first column, Specification (a) shows estimates for Equation 3, excluding year and cohort dummies. The estimated coefficient of the difference-in-differences estimator (β_3) is positive and statistically significant, indicating that each one-year increase in the eligibility age increased the annual saving rate of affected households by an average of 2.1 percentage points over the treatment period. For households facing an eligibility age of 65 years old, this translates to a 10.5 percentage point increase in the average annual saving rate.

Turning to the other variables in Specification (a): the estimated coefficient for β_2 suggests average saving rates across the whole sample were 4.4 percentage points lower in the period following the policy change, after excluding the effects of the treatment and controlling for age. The estimated coefficient for β_1 is small and not statistically significant, suggesting no significant difference between the average saving rates of the treatment and control groups (again, excluding treatment effects and controlling for age). Finally, the variables for age are both statistically significant and their estimated coefficients are consistent with the expected “hump-shape” profile of life-cycle saving behaviour, with saving rates peaking around the age of 55.

¹⁹Ages are excluded if they are observed in two or fewer survey years to reduce noise at the ends of each line.

Table 2 – Average saving rate effects: regression results (sample period 1984 to 1998)

Specification	(a) Equation 3 excluding year/cohort dummy variables	(b) Equation 3 with year dummy variables	(c) Equation 3 with cohort dummy variables
Estimated coefficients (p-values)			
Age	0.082*** (0.000)	0.077*** (0.001)	0.110*** (0.000)
Age squared	-0.001*** (0.000)	-0.001*** (0.001)	-0.001*** (0.000)
Treated group (β_1)	0.001 (0.820)	-0.006 (0.301)	
After (β_2)	-0.044* (0.085)		-0.051* (0.085)
Difference-in-difference estimator of treatment effect (β_3)	0.021*** (0.002)	0.022*** (0.001)	0.013** (0.027)
Additional controls			
Year effects		Y	
Cohort effects			Y
R squared	0.371	0.434	0.426
N	285	285	285

*p<0.1, ** p<0.05, *** p<0.01, p-values are robust to within-year correlation.

Specifications (b) and (c) include the dummy variables in Equation 3 to account for year or cohort effects.²⁰ These specifications show similar results to Specification (a), although the magnitude of the difference-in-difference estimator is significantly smaller in Specification (c) than in the other two specifications. This difference may reflect the fact that the cohort dummies in Specification (c) are capturing some of the treatment effect (which is, by definition, cohort-related) and/or it may reflect cohort-related bias in the estimates of Specifications (a) and (b). The results from Specifications (b) and (c) are therefore best interpreted as providing upper and lower bounds of the true effect.

As robustness checks, two additional specifications were estimated (results not shown). The first excluded cohorts that faced transitional NZS eligibility ages (those born between 1933 and 1936) from the treatment group in Specification (a), as a simple way of checking whether the definition of treatment intensity affects the results. Excluding these cohorts did not significantly change the estimated average treatment effect. The second check collapsed the data into before and after periods following the suggestion of Bertrand, Duflo *et al* (2004), who show that serial correlation in the residuals can lead to underestimated standard errors and false inference when using difference-in-difference regression.²¹ Collapsing the data involves calculating the saving rate for each cohort in the pre-treatment period and the post-treatment period, rather than by year, and re-estimating the model with this reduced data set. Estimated treatment effects were still statistically significant after collapsing the data.

²⁰ The after and treatment variables are removed from Specifications (b) and (c) respectively because of collinearity.

²¹ Tests for autocorrelation in the cohort saving rates suggested no autocorrelation when age controls were included

Table 3 shows results for the three tests of the common trend assumption, outlined in Subsection 4.3. For Specification (d), the estimate of γ_1 shows that, after accounting for age, average saving rates were trending up by 1.2 percentage points each year in both the treatment and control groups. The zero (to three decimal places) estimate for γ_2 indicates no other time trend differences between the treatment and control group over the pre-treatment period. This result provides assurance that differential time trends are not leading to bias in the results presented in Table 2. Similarly, the placebo regression results in Specifications (e) and (f) provide support for the robustness of the central results, with the estimated treatment effect, β_3 , small and statistically insignificant for both specifications, as expected.²²

Table 3 – Average saving rate effects: robustness checks over the pre-treatment period (1984 to 1990)

Specification	(d) Equation 4: difference in trends	(e) Placebo treatment (assumed age increase between 1992-2001)	(f) Placebo treatment (assumed age increase between 1988-1997)
Estimated coefficients (p-values)			
Age	0.053** (0.035)	0.027 (0.291)	0.030 (0.261)
Age squared	-0.000* (0.059)	-0.000 (0.481)	-0.000 (0.420)
Time trend (γ_1)	0.012** (0.030)		
Treatment x Time trend (γ_2)	0.000 (0.975)		
Treated group (β_1)		0.022 (0.359)	0.021 (0.436)
After (β_2)		-0.004 (0.677)	-0.001 (0.550)
Difference-in-difference estimator of placebo effect (β_3)		0.009 (0.290)	0.005 (0.716)
R squared	0.345	0.399	0.399
N	133	133	133

*p<0.1, ** p<0.05, *** p<0.01, p-values are robust to within-year correlation.

5.3 Breakdown of effects by year, age, and income quintile

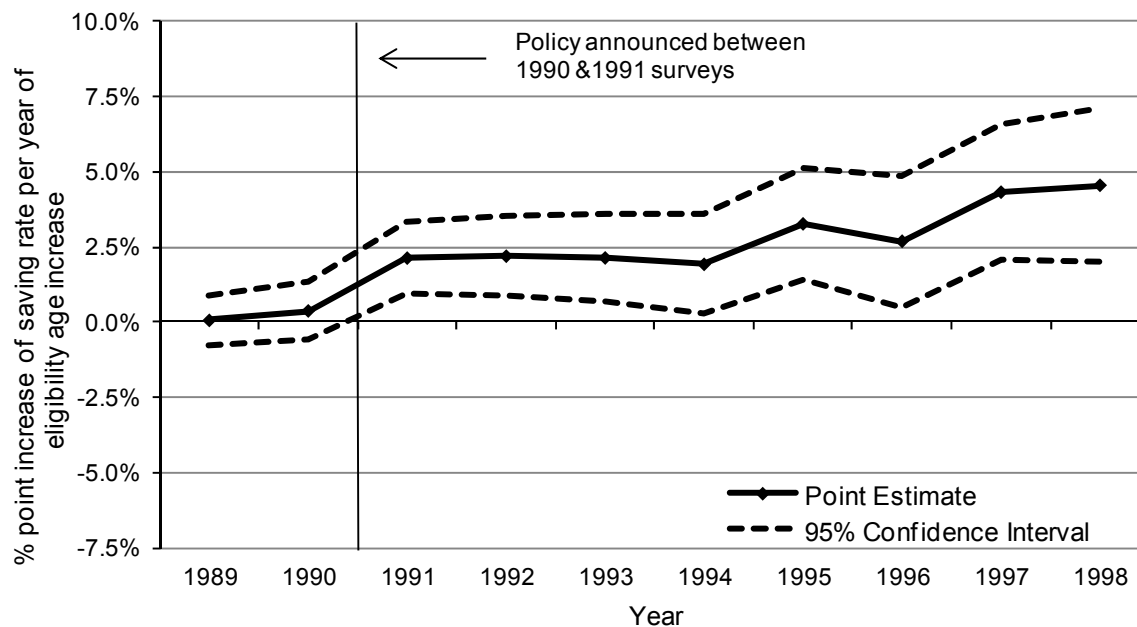
5.3.1 Treatment effects by year

Figure 4 shows estimated treatment effects by year for Equation 5, along with 95 per cent confidence intervals. Each point estimate corresponds to an estimated coefficient in the vector β_3 . The treatment is defined as taking effect from the 1989 survey, two years earlier than the policy announcement, to check for any potential sources of bias immediately prior

²² Differences in the age variable coefficients compared with Specifications (a), (b) and (c) reflect the shorter age range in the pre-treatment sample. The results (not presented) were not sensitive to the inclusion of year or cohort dummy variables.

to the treatment. The estimates in Figure 4 do not suggest significant pre-treatment effects, with the estimated coefficients close to zero in both 1989 and 1990. These results suggest the decrease in the NZS-wage ratio announced in 1989 had a similar effect on both the treatment and control groups and is therefore not biasing the estimated treatment effects.

Figure 4 – Saving rates: estimated treatment effects by year (sample period 1984 to 1998)



The treatment effect is immediately evident in first year after the policy announcement, with the estimate statistically significant and of a near-constant magnitude across the first four years of the treatment period. From 1995 there appears to be a rise in the estimated treatment effects. The rise in the effect is most apparent, and statistically significant, in the last two years of the sample, notwithstanding some widening in the confidence intervals through time (partially reflecting smaller samples in later years). The rise in treatment effects may reflect changes in other policies that occurred in that period or the downturn associated with the Asian Financial Crisis. These potential sources of bias are explored in more detail in the income and expenditure breakdowns discussed in Subsection 6.2. It is also possible that the rise in effects over time reflects an increasing treatment effect with age, as discussed in the following subsection. Unfortunately it is difficult with the data available to separately identify these age and time treatment effects.²³

Re-estimating Specifications (a) through (c) with the 1995 to 1998 period excluded from the sample results in slightly higher (but not significantly different) estimates to those presented in Table 3, suggesting no upward bias from the later years on the overall estimated average treatment effect. These estimates over the shorter time period are shown in Table 4.

²³ The standard approach would be to add additional interaction variables to the difference-in-differences model to separately identify the different effects. The linear relationship between age, time and cohort makes estimating such a model impossible.

Table 4 – Average saving rate effects: regression results (sample period 1984 to 1994)

Specification	(g) Equation 3 excluding year/cohort dummy variables	(h) Equation 3 with year dummy variables	(i) Equation 3 with cohort dummy variables
Age	0.044** (0.046)	0.039* (0.096)	0.070** (0.024)
Age Squared	-0.000 (0.126)	-0.000 (0.206)	-0.001** (0.043)
Treated group (β_1)	0.008* (0.062)	0.002 (0.729)	
After (β_2)	-0.080** (0.010)		-0.089*** (0.007)
Difference-in-difference estimator of treatment effect (β_3)	0.026*** (0.001)	0.027*** (0.001)	0.020*** (0.001)
Year effects		Y	
Cohort effects			Y
R squared	0.424	0.437	0.462
N	209	209	209


*p<0.1, ** p<0.05, *** p<0.01, p-values are robust to within-year correlation.

5.3.2 Treatment effects by age

Figure 5 shows the treatment effects by age of household head, estimated using Equation 6. The profile of effects appears to increase steadily with age, with estimates of just over 0.5 percentage point for households with heads aged between 45 and 49 years old and of around 2.5 percentage points for households with heads aged between 60 and 64 years old. This rising profile could be explained by the fact that younger cohorts could smooth their adjustment over a longer time period, whereas older individuals had a shorter time period between learning of the policy change and their retirement, and therefore their (annual) adjustment was greater.

There are two main, and possibly offsetting, potential sources of bias in these results. First, many of the treated cohorts in the 60 to 64 year old group faced the transitional eligibility ages and are therefore receiving NZS in the treatment period. If receiving NZS is associated with retirement from the workforce, this would lead to downward bias as retirees start dissaving. Second, and in the opposite direction, because most of the older age group is observed later in the sample period, the higher treatment effects may reflect time-related bias, as noted in the previous subsection. However, of those in the treatment group, the 60 to 64 year old age group is least likely to be affected by time-related bias as they are closest in age to the control group.

Figure 5 – Saving rates: estimated treatment effects by age group (sample period 1984 to 1998)



5.3.3 Treatment effects by income quintile

As a final breakdown of the results, Figure 6 shows the estimated treatment effect by disposable income quintile. These estimates are calculated by running five separate regressions analogous to Specification (a), one for each income quintile, with the quintiles calculated for each cohort-year cell. The five point estimates in the figure represent the estimated coefficient of β_3 for each of these regressions.²⁴


The estimated treatment effects for the middle three income quintiles are broadly similar at around 3 percentage points, and each is statistically significant at the 5 per cent level. The effects on these quintiles are clearly driving the overall average estimate.

The estimate for the highest income quintile is insignificantly different from zero. This estimate is likely to reflect the influence of the tax surcharge, which meant that NZS was (or was expected to be) abated away for most of this group while the surcharge was in place. The highest-income households were therefore only weakly affected, if at all, by the change in the NZS eligibility age.

The estimate for the lowest income quintile is negative at -1.8 percentage points. The confidence intervals are noticeably wider for this estimate than for the other quintiles, reflecting greater dispersion in saving rates at lower levels of disposable income. The concurrent reduction in welfare benefits is likely to be a confounding factor that is biasing the treatment estimate for this quintile. The degree to which the cuts in welfare benefits may bias the average treatment effects is limited by the lowest income quintile's proportionally smaller share of total income and expenditure. Separate income and expenditure effects by quintile are examined in Subsection 6.2.

²⁴ An important caveat to this approach is that households transition between income quintiles over time, meaning the population membership of each cohort-quintile group is not fixed. Unfortunately more stable measures of lifetime income, such as highest education qualification, are not available for the full sample.

Figure 6 – Saving rates: estimated treatment effects by income quintile (sample period 1984 to 1998)



5.4 Summary and approximate impact on aggregate saving rates

The preceding subsections have provided a range of estimates for the effects on household saving rates of an additional year in the NZS eligibility age faced by the household head. Estimates of the average treatment effect range from 1.3 to 2.7 percentage points, depending on the sample period and the control variables used. As the midpoint of this range, 2 percentage points is selected as the preferred estimate for the average effect. This estimate also represents the lowest of the shorter-period estimates, which are less likely to be biased by non-treatment-related factors.

The effect of increasing the eligibility age on the aggregate household saving rate is approximated in a simple way by summing the treatment effect for each cohort weighted by their respective treatment intensities and aggregate income shares. On the basis of this calculation, the increase in eligibility age was associated with a lift in the annual aggregate household saving rate of around 2.5 percentage points in 1991, with the effect decreasing to around 2 percentage points by 1998, as the treatment group ages and earns a smaller proportion of total household income.

There are key two caveats that should be considered when interpreting this approximation. First, the calculation only involves cohorts included in the treatment group – the effects on younger cohorts are not incorporated, but are likely to be relatively muted given the increasing effect by age. Second, Vink (2014) shows there are level and trend differences between the HES and HIOA saving rate measures, introducing measurement error in HES-based estimates of aggregate saving effects.

6 Additional Results

To help explain the responses in household saving behaviour presented in Section 5, this section considers the separate effects of the policy change on real disposable income, real expenditure and household labour supply. Subsection 6.1 considers the effect on average household income and expenditure and Subsection 6.2 breaks down this effect by year, age and income quintile. Subsection 6.3 presents the effect of the policy change on household labour supply.

The disposable income and expenditure measures used in this section are defined in Equation 1, deflated by the consumer price index measure of inflation. The natural log is taken of each cohort's average disposable income and expenditure, so that the treatment effect estimates can be interpreted approximately in percentage change terms.

6.1 Average income and expenditure effects

Table 5 provides estimates of the average effect of the treatment on the annual income and expenditure of all households in the treatment group. The specifications shown are estimated using Equation 3 and are therefore analogous to those shown in Table 3. Taken together, the magnitudes of the estimates for each income/expenditure specification pair are consistent with the magnitudes of the estimated treatment in the corresponding saving rate specifications.²⁵ For income, the estimated treatment effects range between an approximate increase of 1.3 to 2.4 per cent in annual real household disposable income for each additional year added to the NZS eligibility age.²⁶ On the other hand, the estimated effect on real expenditure is marginally negative but statistically insignificant in each specification. These results suggest that the average household saving response was generated primarily by increases in income rather than by reductions in expenditure.

Table 5 – Average income & expenditure effects: regression results (sample period 1984 to 1998)

Specification	Real disposable income			Real expenditure		
	(j) Eq 3 excl year/cohort dummy variables	(k) Eq 3 with year dummy variables	(l) Eq 3 with cohort dummy variables	(m) Eq 3 excl year/cohort dummy variables	(n) Eq 3 with year dummy variables	(o) Eq 3 with cohort dummy variables
Age	0.168*** (0.000)	0.184*** (0.000)	0.202*** (0.000)	0.086*** (0.000)	0.105*** (0.000)	0.090*** (0.001)
Age squared	-0.002*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)
Treated group (β_1)	-0.124*** (0.005)		-0.110** (0.031)	-0.073** (0.045)		-0.051 (0.217)
After (β_2)	0.001 (0.863)	-0.004 (0.581)		0.000 (0.997)	0.002 (0.726)	
Difference-in-difference estimator of treatment effect (β_3)	0.024** (0.014)	0.017* (0.067)	0.013 (0.144)	-0.001 (0.918)	-0.007 (0.389)	-0.002 (0.764)
Year effects		Y			Y	
Cohort effects			Y			Y
R squared	0.745	0.764	0.761	0.791	0.822	0.809
N	285	285	285	285	285	285

*p<0.1, ** p<0.05, *** p<0.01, p-values are robust to within-year correlation.

²⁵ The small difference between the estimated effects on saving and the residual of the estimated effects on income and expenditure result from saving being measured as a rate, while income and expenditure are measured in log terms.

²⁶ The treatment effect is not statistically significant at the 10 per cent level if cohort dummies are included.

Table 6 shows the key results for the three tests of the common trend assumption, outlined in Subsection 4.3, for real disposable income. These test results raise some concerns about the potential for biases in the results reported in Table 5. In particular, the placebo regression estimates for disposable income in Specification (q) is of a similar magnitude to those in Table 5. The results for expenditure are less concerning. The full set of results relating to Table 6 is contained in the appendix.

Table 6 – Average income & expenditure effects: robustness checks over the pre-treatment period (1984 to 1990)

Specification	Income			Expenditure		
	(p) Equation 4: difference in trends	(q) Placebo treatment (1992-2001)	(r) Placebo treatment (1988-1997)	(s) Equation 4: difference in trends	(t) Placebo treatment (1992-2001)	(u) Placebo treatment (1988-1997)
Treatment x Time trend (γ_2)	-0.000 (0.119)			-0.000* (0.075)		
Difference-in-difference estimator of placebo effect (β_3)		0.017 (0.116)	0.004* (0.078)		0.007 (0.547)	0.002 (0.473)
R squared	0.343	0.423	0.424	0.616	0.614	0.614
N	133	133	133	133	133	133

*p<0.1, ** p<0.05, *** p<0.01, p-values are robust to within-year correlation.

6.2 Breakdown of income and expenditure effects by year, age and income quintile

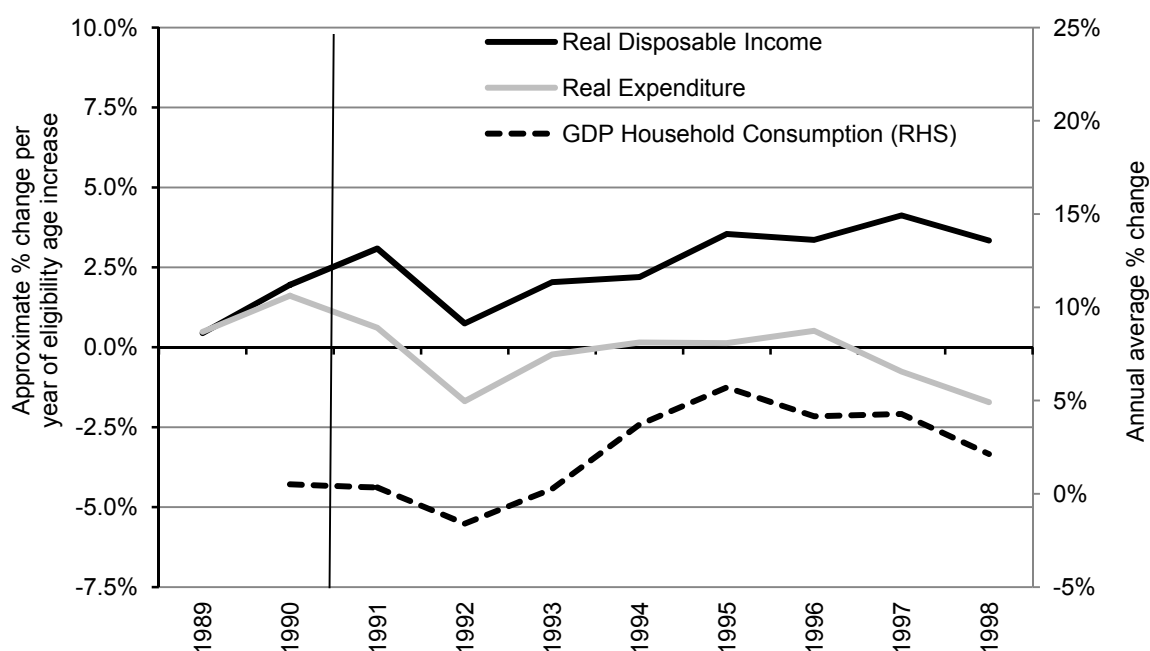
6.2.1 Treatment effects by year

A breakdown of the treatment effects on income and expenditure by year, along with growth in GDP household consumption expenditure, is shown in Figure 7. Consistent with the corresponding analysis of saving rate effects shown in Figure 4, the treatment is defined as taking effect from 1989 and estimated using the model in Equation 5, with each point estimate representing one of the estimated coefficients in the vector β_3 . Confidence intervals are excluded for presentational clarity. The GDP data is included to illustrate the possible relationship between the treatment effects and the economic cycle. Such a relationship makes sense if income and expenditure for the (younger) treatment group are more sensitive to economic developments than the (older) control group, which would be explained by higher proportions of the control group having retired from the work force.²⁷

Focussing first on the pre-treatment period, the positive treatment effect in 1990 for both income and expenditure is consistent with the concerns raised by the robustness tests examined in the previous subsection. Potential explanations for these positive pre-treatment effects include the Government's reduction in the NZS-wage ratio in 1989 and economic growth over the 1980s. An implication is that the subsequent estimated year effects and the estimated average effects shown in Table 5 are likely to be biased upwards. This bias suggests the true positive effect on income is lower than estimated, and the true negative effect on expenditure is greater than estimated.

²⁷ Although retirees would still be subject to the effects of fluctuations in wages (through the NZS-wage link) and in investment income, they would be unaffected (by definition of retirement) to changes in the economy-wide employment rate.

Figure 7 – Income and Expenditure: estimated treatment effects by year (sample period 1984 to 1998)



In 1991, the first post-treatment year, the estimated treatment effects are positive and negative respectively for income and expenditure (compared with the previous period, 1990), which is aligned with the positive treatment effect on saving shown in Figure 4. The magnitude of the change in the income and expenditure treatment effects is similar between 1990 and 1991. This suggests that the treatment effect on saving was caused by responses in both the income and expenditure.

For the following five years, 1992 to 1996, the estimated treatment effects of the two measures broadly track the economic cycle in parallel: both decline during the sharp economic recession in 1991 (1992 HES year), before rising with the subsequent economic recovery. However, in the last two years of the sample, 1997 and 1998, the estimated treatment effects diverge, consistent with the rise in the saving estimate in those years. This pattern could be explained by policy and/or economic shocks that had different effects on the treatment and control groups because of, for example, differences in the proportions of each that were receiving NZS and/or in the labour market. These differences increased over time and are greatest at the end of the sample period.²⁸

There are three shocks of note that occurred late in the sample. First, and possibly the most significant, was the removal of the tax surcharge on NZS. This policy was announced in the 1996 Budget, but implemented from April 1997. Anticipation of the permanent increase in income resulting from the surcharge removal likely boosted the control group’s expenditure in the 1997 HES survey by more than it did for the younger treatment group, because most households in the treatment group had to wait longer before they would benefit. The actual surcharge removal in survey year 1998 then directly increased the control group’s (current) disposable income compared with the treatment group. Such effects are consistent with the observed drop in the expenditure treatment effect in 1997 and the subsequent drop in the disposable income treatment effect in 1998. Second, income tax reductions implemented in 1996 (affecting HES surveys from 1997) would have boosted the disposable income (and expenditure, but to a lesser extent if

²⁸ By 1997, all of the households in the control group were eligible for NZS, and they were working an average of 7 hours per week. This, compares with only two (of fifteen) cohorts eligible for NZS in the treatment group, in which households were working an average of 50 hours per week.

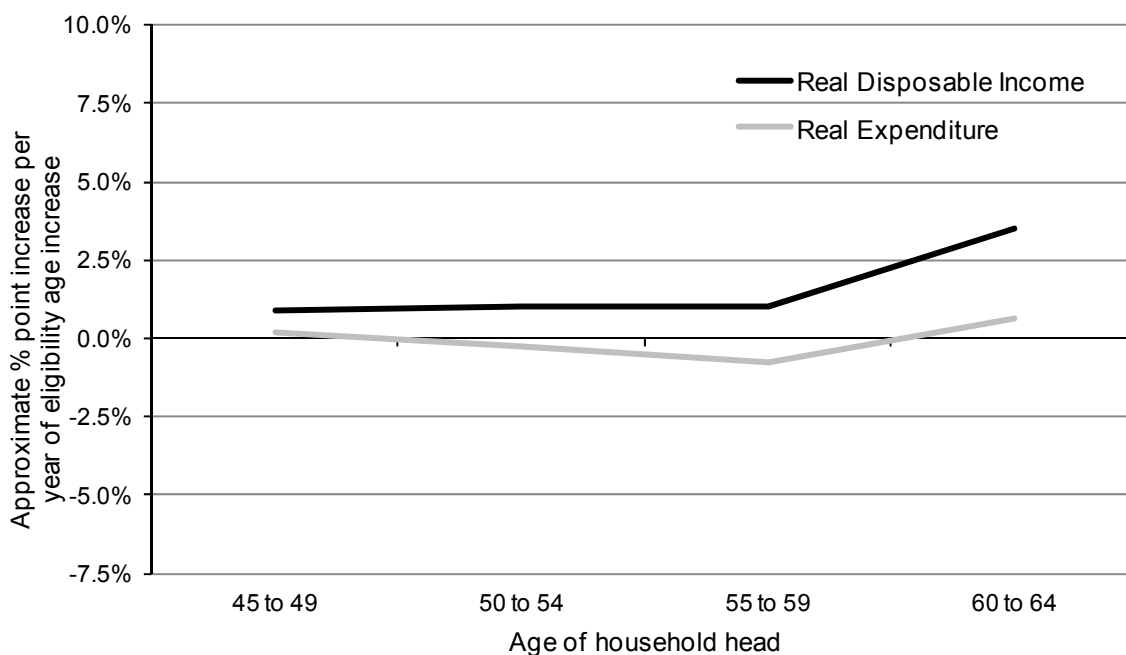
some of the tax cut was saved) of the treatment group compared with the control group. Finally, the recession of 1997/98 is also likely to have had greater effects on the treatment than control group.

An important question is whether the differing effects on income and expenditure over time undermine the reliability of the saving estimates. The divergence in income and expenditure treatment effects for 1997 and 1998 does indicate that estimates for those years should be viewed with caution. However, the parallel movement of the income and expenditure effects over the earlier treatment years suggests that the corresponding saving estimates are valid.

6.2.2 Treatment effects by age

Figure 8 shows treatment effects by age for income and expenditure, estimated using Equation 6. The pattern of effects is in line with the corresponding estimates for saving rates by age presented in Subsection 5.3.2. The estimated treatment effects on income for the oldest age group are perhaps the most striking of these results, with income increasing by approximately 17.5 per cent for households with heads facing an eligibility age of 65 years old. For most households in this group (those not affected by the tax surcharge), the increase in NZS eligibility age represented a direct loss in current income (rather than expected future income) compared with those of a similar age in the control group. The large positive treatment effect on income therefore implies these households more than offset their loss of NZS income with additional earned income. Furthermore, as noted in Subsection 5.3.2, the estimates for this age group are least likely to be affected by any positive bias from differential economic or policy effects, given they are closest in age to the control group. Subsection 6.3 considers what might be driving this result by examining changes in labour supply.

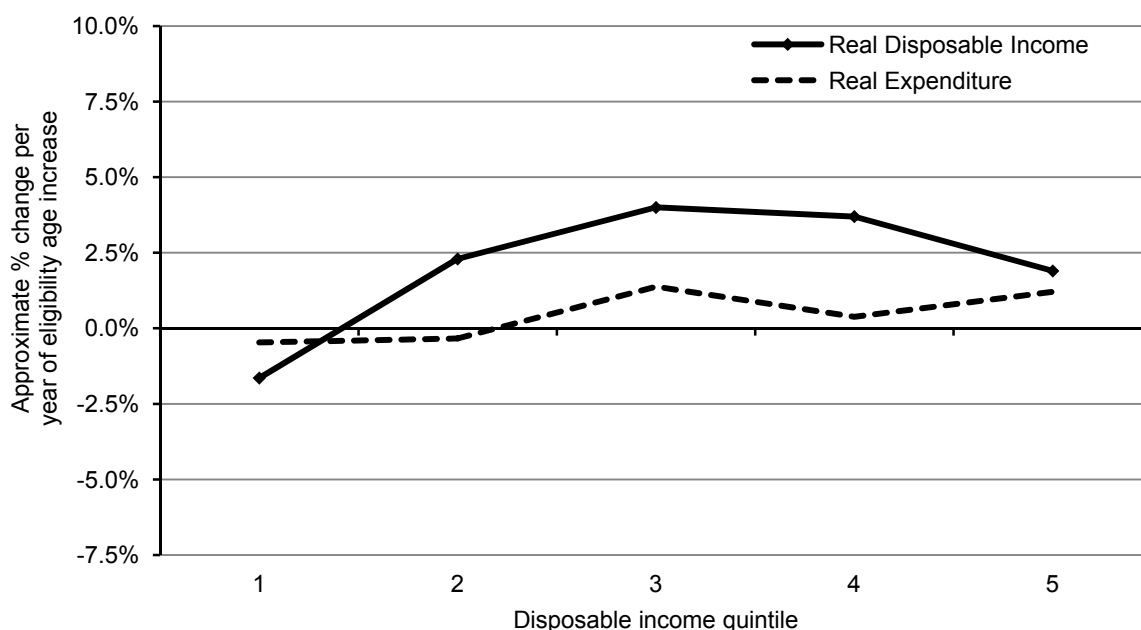
Figure 8 – Income and expenditure: estimated treatment effects by age (sample period 1984 to 1998)



6.2.3 Treatment effects across the income distribution

Estimated treatment effects for income and expenditure by disposable income quintile are shown in Figure 9. Again, the pattern of effects is in line with the corresponding estimates for saving rates by income quintile presented in Section 5.3.3. The likelihood of an upward bias in the estimates caused by economic growth helps to explain the positive point estimates for expenditure for the upper quintiles, which would otherwise appear anomalous. The negative estimates for the lowest quintile likely reflect the direct effects of the loss of current income from NZS for households in the 60 to 65 year old age group and from the 1991 benefit cuts.

Figure 9 – Income and expenditure: estimated treatment effects across the income distribution (sample period 1984 to 1998)



6.3 Labour supply effects

This subsection uses the HES measure of average weekly hours worked totalled for all members of the household to provide greater insight into the treatment effects on income. Like the other dependent variables, the average hours worked measure is calculated for each cohort in each year. The focus of the section is on the breakdown of treatment effects over time and by age, which are estimated using Equations 5 and 6 respectively.²⁹ Results for the three tests of the common trends assumption, outlined in Subsection 4.3, do not raise concerns about potential biases and are included in Appendix B.

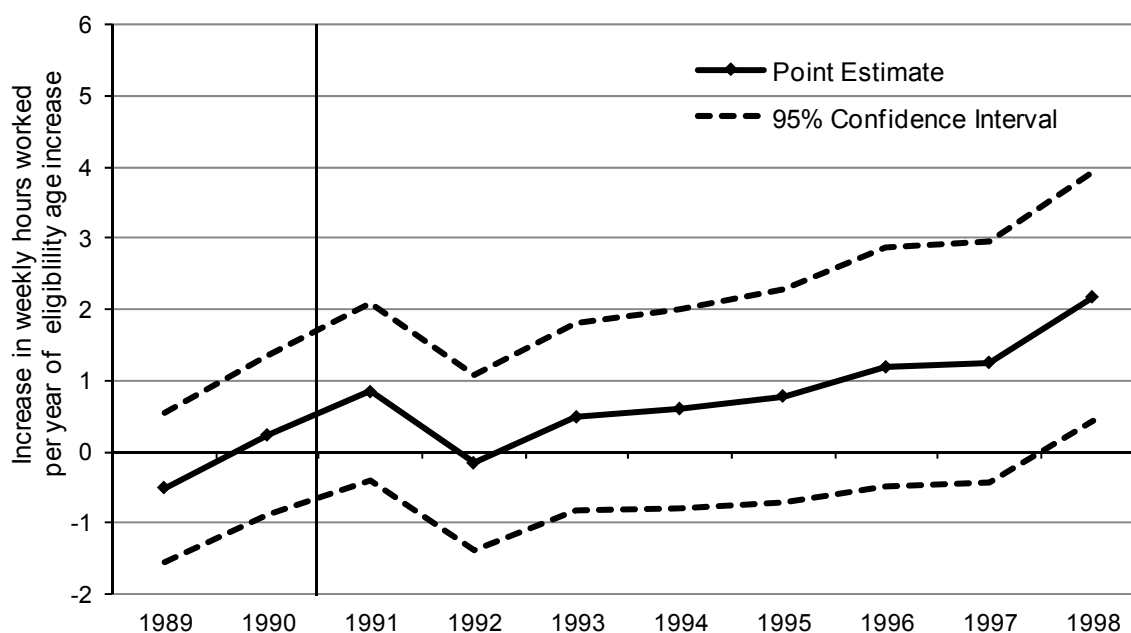
6.3.1 Treatment effects by year

Figure 10 shows treatment effects for labour supply by year, estimated using Equation 5, with the treatment defined as taking effect from the 1989 survey. The pattern of estimated effects is very similar to the corresponding pattern for income in Figure 7. This similarity suggests that a change in hours worked drove the observed income effects and supports the argument that the estimated income and expenditure effects are influenced by the economic cycle. The continuing rise in the labour supply treatment effects later in the

²⁹ One difference from these equations is the use of dummy variables to control for age, rather than an age quadratic (which are employed in the rest of this paper). Although the inclusion of dummy variables reduces degrees of freedom (and thereby the statistical significance of the estimates presented), an age quadratic is not well suited to capturing the discontinuities in the life-cycle patterns of labour force participation.

sample (1997 and 1998) is the one main difference with the pattern of income effects. This may have been caused by the removal of the NZS tax surcharge. The removal would have directly increased the control group's total current disposable income (reducing the estimated income treatment effect). However, it probably only had a minor impact on the control group's labour supply because household heads in the group were aged between 65 and 69 by 1997, and many would have already withdrawn permanently from the labour market.

Figure 10 – Average weekly hours worked: estimated treatment effects by year (sample period 1984 to 1998)



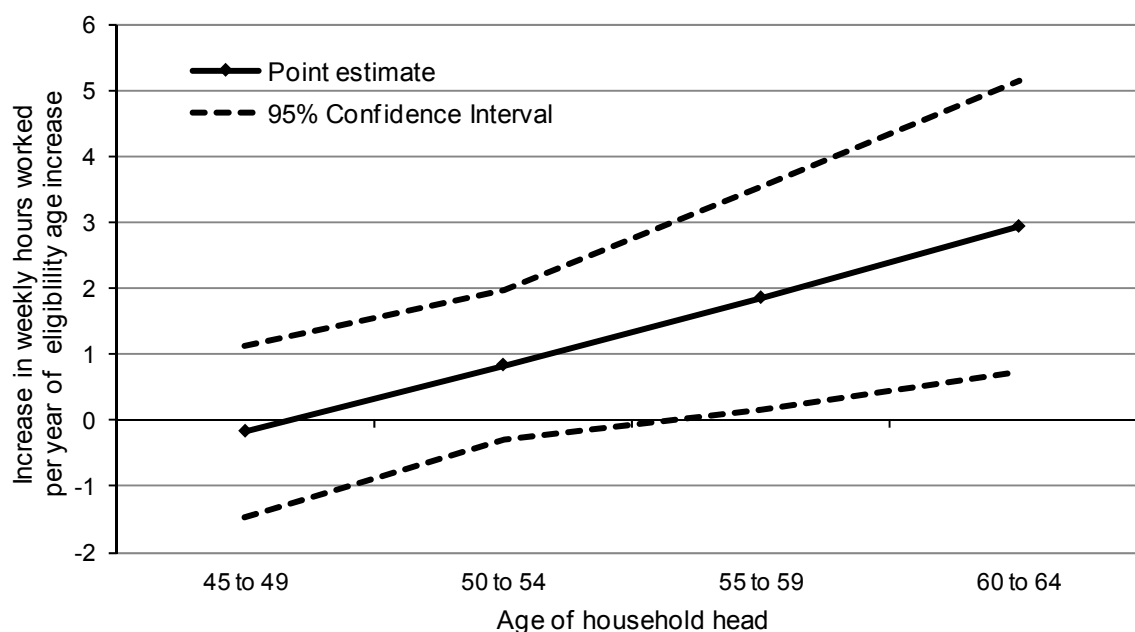
6.3.2 Treatment effects by age

The estimated treatment effects on labour supply by age of household head are shown in Figure 11. The effects show a marked increase with age. Moreover, as discussed in the previous two subsections covering to age-specific effects, any bias in the age estimates are likely to be lowest for the oldest cohorts in the treatment group.

The estimated average increase of approximately three hours per week for households headed by 60 to 64 year olds corresponds to an approximate 15 per cent increase in labour supply for each additional year of NZS eligibility. This represents a total increase in labour supply of 75 per cent for cohorts that faced an eligibility age of 65 years old. Although these estimates may appear high, their magnitude is consistent with the corresponding estimated treatment effect for income of around 17.5 percent, which reflects additional income earned over and above what was required to offset the direct losses of NZS. The labour supply estimate is also in line with the findings of Hurnard (2005) who found that the rate of labour force participation (not hours) for 60 to 64 year olds approximately doubled as a result of the increase in the NZS eligibility age.³⁰

³⁰ Hurnard (2005) does not explicitly state this result, which was derived from his separate estimates for males and females.

Figure 11 – Average weekly hours worked: estimated treatment effects by age (sample period 1984 to 1998)



The question remains as to why older households appear to have worked by more than was required to offset their loss in NZS income. Two potential explanations relate to the presence of “frictions” in the labour market. First, households may not have been able to optimally choose their preferred number of hours worked, for example because of a lack of part-time work opportunities. Second, social norms may have encouraged workers to retire at the point they became eligible for NZS, rather than at the optimal time from an economic perspective.

7 Sensitivity Analyses

This section presents sensitivity analyses relating to the average saving results presented in Subsection 5.2. Subsection 7.1 uses an alternative, economic, measure of saving. Subsection 7.2 undertakes the analysis at the household level, rather than the cohort level, which allows for the inclusion of household level demographics.

7.1 An alternative measure of saving

The main measure of saving used in this paper corresponds closely to the HIOA measure of saving as outlined in Subsection 2.2. However, other saving definitions may be preferable from an economic perspective. In this section, the difference-in-difference regression model of Equation 3 is re-estimated using an alternative measure of saving that is better aligned with economic concepts. There are two main adjustments. First, spending on items that provide consumption benefits over more than a year (such as consumer durable goods, capital goods, education and healthcare services) is classified as saving rather than expenditure. Second, the inflation components of interest income and interest expenses are removed from income and expenditure respectively. This adjustment accounts for the fact that inflation reduces the real value of borrowed principal. The inflation component of interest is compensation for this reduction and not an actual income or expense.³¹

³¹ Gorman, Scobie *et al*(2013) show that these adjustments have significant effects on the trend and level of aggregate measures of saving.

Table 7 shows the results from re-estimating the difference-in-difference regression model of Equation 3 using this economic measure of saving as the dependent variable. Specification (a) from Table 3 is also reproduced for comparison. Estimates using the two alternative measures of saving are almost identical, demonstrating that the results are not particular to the choice of saving measure. Results are similar if cohort or year dummies are included.

Table 7 – Average saving rate effects with alternative saving rate measure: regression results (sample period 1984 to 1998)

Specification	(a) Equation 3 with HIOA-comparable saving measure excluding year/cohort dummy variables	(s) Equation 3 with economic saving measure excluding year/cohort dummy variables
Age	0.082*** (0.000)	0.070*** (0.000)
Age squared	-0.001*** (0.000)	-0.001*** (0.000)
Treated group (β_1)	0.001 (0.820)	-0.001 (0.905)
After (β_2)	-0.044* (0.085)	-0.046** (0.032)
Difference-in-difference estimator of treatment effect (β_3)	0.021*** (0.002)	0.021*** (0.004)
R squared	0.321	0.383
N	285	285

*p<0.1, ** p<0.05, *** p<0.01, p-values are robust to within-year correlation.

7.2 Saving measured at the household level

The analysis in this paper has been based on saving rates measured at the cohort level. The main advantage of this method is that it minimises the influence of measurement errors and outliers. However, it does not allow for the consideration of differences in characteristics at the household level, which may have an influence on a household's saving rate. If these characteristics – such as family structure - vary across cohorts and/or across time, they may introduce bias into the treatment effect estimates.

This section repeats the analysis presented in Table 2 (with the same HIOA-comparable saving definition), but uses saving rates calculated at the household level as the dependent variable. To minimise the effect of outliers and measurement error, the estimation is undertaken using quantile regression at the median.³² The model is also estimated with the inclusion of dummy variables to control for the following household characteristics: gender of household head; whether the household head is of Maori or Pacific Island ethnicity; whether the household is renting, has a mortgage, or is rent and mortgage free; and whether the household head is a single parent or partnered with children.

Table 8 shows the results from the household-level estimations. Household characteristics are excluded from Specification (u) but included in (v). Again, Specification (a) from Table 3 is reproduced for comparison. The estimated effect of the policy change in both specifications using household-level data is not significantly different from the estimate

³² P-values are calculated taking into account error clustering by year using the methods of Machado, Parente *et al* (2012).

using cohort averages. These results provide assurance that the results are not particular to the measurement of saving at the cohort level, or biased by differences in the demographic composition of the sample. Results are similar if cohort or year dummies are included.

Table 8 – Average saving rate effects using household-level data: regression results (sample period 1984 to 1998)

Specification	(a) Equation 3 cohort level, excl year/cohort dummy variables	(t) Equation 3 household level excl demographics and year/cohort dummy variables	(u) Equation 3 household level with demographics excl year/cohort dummy variables
Age	0.070*** (0.000)	0.047*** (0.000)	0.045*** (0.000)
Age squared	-0.001*** (0.000)	-0.000*** (0.000)	-0.000*** (0.002)
Treated group (β_1)	-0.001 (0.905)	-0.061** (0.043)	-0.053* (0.059)
After (β_2)	-0.046** (0.032)	-0.001 (0.750)	-0.001 (0.797)
Difference-in-difference estimator of treatment effect (β_3)	0.021*** (0.004)	0.018*** (0.002)	0.017*** (0.005)
Household Level		Y	Y
Demographics			Y
R squared	0.321	0.003	0.006
N	285	13263	13263

*p<0.1, ** p<0.05, *** p<0.01, p-values are robust to within-year correlation.

8 Conclusion

This paper has assessed the effects of the last increase in the NZS eligibility age, from 60 to 65 years old, on household saving rates using difference-in-differences regression analysis. The analysis suggests the policy change caused an average increase in household saving rates of around 2 percentage points for each additional year added to the eligibility age faced by the household head. For households facing an eligibility age of 65 years old, this translates to a 10 percentage point increase in the average annual saving rate. The impact of the policy change on saving rates is evident from the first year after its announcement in 1990, and it appears to have been greatest on older and middle-income households. These results are generally robust to a variety of sensitivity and falsification tests.

The findings for saving rates are validated by the estimated effects of the policy change on disposable income, expenditure, and labour supply, although the findings for income and expenditure are less robust than for saving rates. Disposable income and labour supply are found to have increased as a result of the policy change, while the effect on expenditure is negative. The positive effects on disposable income and labour supply are, again, greatest for older households.

The positive income effect on households with heads in the 60 to 64 year-old age group is particularly interesting, given it reflects additional income earned over and above that

required to offset the direct loss of NZS income resulting from the policy change. This “overcompensation” in income is corroborated by a corresponding large (approximately 15 per cent) increase in these households’ labour supply for each additional year in the NZS eligibility age faced. Labour market frictions, such as a lack of part-time work opportunities, provide one plausible explanation for the apparent excess sensitivity of labour supply and income to the policy change.

The results suggest the lift in the eligibility age led to an initial increase in the aggregate household saving rate of around 2.5 percentage points, declining to around 2 percentage points by 1998. To what extent would future reductions in the generosity of NZS lead to similar changes in aggregate household saving rates? The absence of a tax surcharge, an income test, is one important difference of NZS today from the scheme in 1990s. All else being equal, this difference would likely increase the relative effects of a future reduction in the generosity of NZS. Two other important factors relate to the amount of forewarning provided and the nature of household expectations. To the extent that any future changes are signalled well in advance of implementation, households will have a longer period of time to adjust, leading to a smaller increase in annual saving rates. However, if future changes are already built into household expectations, the effects on saving will be reduced.

References

- Aguila, Emma (2011) "Personal retirement accounts and saving." *American Economic Journal: Economic Policy* 3(4): 1-24.
- Aimer, Peter and Jack Vowles (1993) *Voters' vengeance: 1990 election in New Zealand and the fate of the fourth Labour Government*. Auckland University Press.
- Attanasio, Orazio P and Agar Brugiavini (2003) "Social security and households' saving." *Quarterly Journal of Economics* 118(3): 1075-1119.
- Attanasio, Orazio P and Susann Rohwedder (2003) "Pension wealth and household saving: Evidence from pension reforms in the United Kingdom." *American Economic Review* 93(5): 1499-1521.
- Autor, David H (2003) "Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing." *Journal of Labor Economics* 21(1): 1-42.
- Bertrand, Marianne, Esther Duflo and Sendhil Mullainathan (2004) "How much should we trust differences-in-differences estimates?" *Quarterly Journal of Economics* 119(1): 249-275.
- Boston, Jonathan, Paul Dalziel and Susan St John (1999) *Redesigning the welfare state in New Zealand: Problems, policies, prospects*. Oxford University Press Auckland.
- Brook, Anne-Marie (2014) "Options to narrow New Zealand's saving-investment imbalance." Wellington, New Zealand Treasury, *Working Paper* 14/17.
- Coleman, Andrew (2006) "The life-cycle model, savings and growth." *Reserve Bank workshop on Housing, Savings, and the Household Balance Sheet*, Vol. 14
- Dalziel, Paul and Ralph Lattimore (2001) *The New Zealand macroeconomy: A briefing on the reforms and their legacy*. Oxford University Press.
- Deaton, Angus (1985) "Panel data from time series of cross-sections." *Journal of Econometrics* 30(1): 109-126.
- Donald, Stephen G and Kevin Lang (2007) "Inference with difference-in-differences and other panel data." *The Review of Economics and Statistics* 89(2): 221-233.
- Feldstein, Martin (1974) "Social security, induced retirement, and aggregate capital accumulation." *The Journal of Political Economy* 82(5): 905-926.
- Gibson, John K and Grant M Scobie (2001) "Household saving behaviour in New Zealand: A cohort analysis." Wellington, New Zealand Treasury, *Working Paper* 01/18
- Gorman, Emma, Grant M Scobie and Yongjoon Paek (2013) "Measuring saving rates in New Zealand: An update." Wellington, New Zealand Treasury, *Working Paper* 13/04
- Hurnard, Roger (2005) "The effect of New Zealand Superannuation eligibility age on the labour force participation of older people." Wellington, New Zealand Treasury, *Working Paper* 05/09
- Law, David (2013) "Retirement income policy and national savings." Wellington, New Zealand Treasury, *Working Paper* 13/28
- Machado, Jose Antonio Ferreira, Paulo MDC Parente and JMC Santos Silva (2012) "qreg2: Stata module to perform quantile regression with robust and clustered standard errors." *Statistical Software Components*.
- Ministry of Social Development (2002) "The Statistical Report for year ending June 2002."
- New Zealand Herald (1990, 22 June) "National does big u-turn on superannuations reform." *New Zealand Herald*.
- New Zealand Treasury (2010) "Saving in New Zealand - Issues and options." *Report Prepared for the Savings Working Group*.

- Preston, David A (2001) *Retirement income in New Zealand: the historical context*. Office of the Retirement Commissioner.
- Rogers, William (1994) "Regression standard errors in clustered samples." *Stata Technical Bulletin* 3(13).
- Silva, JMC and Paulo MDC Parente (2013) "Quantile regression with clustered data." University of Essex, Department of Economics.
- Vink, Mark (2014) "Intergenerational developments in household saving behaviour." Wellington, New Zealand Treasury, *Working Paper* 14/23.

Appendix A: HES Saving Measures

This appendix provides additional detail relating to HES and the calculation of household saving.

Preferred Saving Measure

Table 1 provides the expenditure categories and codes relating to the saving measure outlined in Subsection 2.2.

Table 1 – HES expenditure categories used in preferred saving measure calculation

Category	Item reference number codes
HES total expenditure	0001-7269
HES contribution to savings	6900-6909
HES mortgage principal payments	1210-1217
HES life and health insurance payments	6903
HES purchases of property	1100-1109, 1538
HES sale of property	1110-1119

Alternative saving measure – inflation adjustment

The inflation adjusted measure of saving used in Subsection 7.1 is calculated as follows. In each year the ratio of annual inflation to nominal interest rates is applied to each household's interest payments and receipts to provide "inflation components" of interest payments and receipts. The inflation component of interest payments is added to saving and deducted from expenditure. The inflation component of interest receipts is deducted from saving (and income). The Reserve Bank of New Zealand's measures of "floating first mortgage" and "6-month term deposit" rates are used to approximate the average interest rates for interest payments and receipts respectively.

Alternative saving measure – including investment items

Table 2 outlines the investment expenditures added to the alternative saving measure including investment expenditure described in Subsection 7.1.

Table 2 – Expenditure items classified as "investment" expenditure

Category	Item reference number codes
Health	5200 – 5299, 6000-6099
Education	6200-6299, 6702-6703
Durable goods	2100-2179, 2200-2339, 2400-2419, 2500-2519, 4200-4229
Building-permit fees	1300
Office Equipment	5650-5669
Sales of durable and capital goods	7000-7269
Other capital goods	1100-1199, ,5506, 5507

Appendix B: Additional Tables

Appendix Table 1 – Sample size of each cohort by year

Cohort	Year															
	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	Total
1928	40	55	55	45	75	50	40	45	45	50	30	35	45	35	25	670
1929	50	50	50	45	60	35	40	40	35	55	25	30	45	30	35	625
1930	45	55	40	60	60	45	40	40	45	50	45	35	30	35	25	650
1931	50	55	50	50	50	45	40	45	40	55	45	35	25	45	35	665
1932	60	35	40	40	65	50	35	35	40	65	30	30	30	30	30	615
1933	40	45	50	30	60	45	45	30	35	55	35	35	25	35	25	590
1934	45	35	35	40	55	40	40	35	40	50	30	35	30	30	35	575
1935	55	40	30	30	50	45	45	35	30	50	35	35	40	35	30	585
1936	45	50	55	30	55	35	40	35	35	50	35	40	35	35	35	610
1937	60	60	50	45	60	40	50	30	35	55	30	35	30	30	30	640
1938	55	60	45	50	60	40	35	30	30	40	30	40	30	35	50	630
1939	50	50	55	45	70	40	50	40	40	55	50	35	45	40	25	690
1940	55	65	65	60	45	60	45	40	35	65	40	45	45	50	35	750
1941	60	65	50	65	65	45	40	45	40	65	40	40	40	45	45	750
1942	65	70	50	50	75	55	55	50	45	80	60	45	50	30	35	815
1943	60	55	45	55	80	65	55	40	45	80	40	40	50	45	45	800
1944	75	65	70	60	85	40	55	50	55	60	35	40	35	40	45	810
1945	70	65	65	75	60	55	40	40	50	70	50	40	50	35	50	815
1946	80	80	60	70	80	65	70	55	55	80	60	45	50	35	65	950
Total	1060	1060	965	940	1210	905	860	755	780	1130	755	710	730	700	705	13265

Note: Actual sample sizes rounded to the nearest multiple of five.

Appendix Table 2 – Average income effects: robustness checks over the pre-treatment period (1984 to 1990)

Specification	(m) Eq 4: difference in trends	Placebo treatment (1992-2001)			Placebo treatment (1988-1997)		
Age	0.201*** (0.000)	0.152*** (0.004)	0.156*** (0.005)	0.179*** (0.002)	0.147*** (0.003)	0.152*** (0.005)	0.174*** (0.003)
Age squared	-0.002*** (0.000)	-0.002*** (0.004)	-0.002*** (0.006)	-0.002*** (0.001)	-0.002*** (0.003)	-0.002*** (0.006)	-0.002*** (0.001)
Time trend (γ_1)	0.005 (0.169)						
Treatment x Time trend (γ_2)	-0.000 (0.119)						
Treated group (β_1)		-0.006 (0.505)	-0.005 (0.557)		-0.001 (0.554)	-0.001 (0.630)	
After (β_2)		-0.049 (0.274)		-0.026 (0.571)	-0.063 (0.181)		-0.036 (0.468)
Difference-in-difference estimator of placebo effect (β_3)		0.017 (0.116)	0.016 (0.152)	0.012 (0.352)	0.004* (0.078)	0.004 (0.109)	0.003 (0.304)
Year effects			Y			Y	
Cohort effects				Y			Y
R squared	0.343	0.423	0.435	0.461	0.424	0.437	0.462
N	133	133	133	133	133	133	133

Appendix Table 3 – Average expenditure effects: robustness checks over the pre-treatment period (1984 to 1990)

Specification	(m) Eq 4: difference in trends	Placebo treatment (1992-2001)			Placebo treatment (1988-1997)		
Age	0.162*** (0.000)	0.130*** (0.005)	0.135*** (0.006)	0.128* (0.080)	0.123** (0.011)	0.127** (0.013)	0.126* (0.084)
Age squared	-0.002*** (0.000)	-0.001*** (0.005)	-0.002*** (0.005)	-0.001** (0.037)	-0.001** (0.010)	-0.001** (0.013)	-0.001** (0.039)
Time trend (γ_1)	-0.008 (0.200)						
Treatment x Time trend (γ_2)	-0.000* (0.075)						
Treated group (β_1)		-0.003 (0.789)	-0.003 (0.782)		-0.000 (0.969)	0.000 (0.976)	
After (β_2)		-0.071 (0.154)		-0.059 (0.326)	-0.083 (0.139)		-0.063 (0.318)
Difference-in-difference estimator of placebo effect (β_3)		0.007 (0.547)	0.007 (0.603)	0.007 (0.678)	0.002 (0.473)	0.002 (0.518)	0.002 (0.664)
Year effects			Y			Y	
Cohort effects				Y			Y
R squared	0.616	0.614	0.655	0.657	0.614	0.656	0.657
N	133	133	133	133	133	133	133

Appendix Table 4 – Labour supply effects: robustness checks over the pre-treatment period (1984 to 1990)

Specification	(m) Eq 4: difference in trends	Placebo treatment (1992-2001)			Placebo treatment (1988-1997)		
Time trend (γ_1)	-1.625* (0.073)						
Treatment x Time trend (γ_2)	-0.000 (0.998)						
Treated group (β_1)		-1.543** (0.042)	-0.414 (0.600)		-0.409** (0.025)	-0.108 (0.579)	
After (β_2)		-3.971 (0.194)		0.424 (0.917)	-3.006 (0.405)		0.156 (0.971)
Difference-in- difference estimator of placebo effect (β_3)		0.217 (0.822)	0.748 (0.480)	0.701 (0.630)	0.013 (0.953)	0.159 (0.520)	0.151 (0.623)
Year effects				Y		Y	
Cohort effects				Y			Y
R squared	0.817	0.863	0.857	0.819	0.863	0.857	0.817
N	133	133	133	133	133	133	133