

Intergenerational Developments in Household Saving Behaviour

Mark Vink

New Zealand Treasury Working Paper 14/23

November 2014



THE TREASURY

Kaitohutohu Kaupapa Rawa

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 14/23**

Intergenerational Developments in Household Saving Behaviour

MONTH/YEAR

November 2014

AUTHOR

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)

978-0-478-43607-5

URL

Treasury website at November 2014:
<http://www.treasury.govt.nz/publications/research-policy/wp/2014/14-23>
Persistent URL: <http://purl.oclc.org/nzt/1704>

ACKNOWLEDGEMENTS

I am grateful to Hamish Low, Andrew Coleman, John Creedy, Grant Scobie, John Gibson, Oscar Parkyn, and Jeff Cope for their insightful comments and suggestions. I would also like to thank Christopher Ball and Talosaga Talosaga for their work in preparing the data. This paper extends research I originally undertook as part of the requirements for the degree of Master of Philosophy in Economics, University of Cambridge, 2013.

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 saving behaviour of different generations of households in New Zealand over the period 1984 to 2010 using data from the Household Economic Survey. The paper employs a life-cycle framework to estimate regression models that identify the influence of age and birth year on household saving rates. Consistent with the life-cycle hypothesis, the results show that household saving rates exhibit a hump shape over the life cycle. The results also indicate significant differences in the average saving rates of households from different birth cohorts. From the baby boomers onward, the average saving rates of each generation exceed those of the generation preceding it. Although there are differences between the Household Economic Survey and national accounts saving measures, which present a caveat to the analysis, the paper's findings provide some insight into demographic influences on national household saving trends. The results suggest that the movement of the baby boomers into their higher saving years has contributed positively to aggregate saving rates, but that future effects of population ageing are likely to be negative. On the other hand, it is possible that the lift in saving rates over recent generations will provide an ongoing positive contribution to aggregate saving rates throughout the projection period ending 2030.

JEL CLASSIFICATION

D14: Household saving
D91: Life cycle models and saving
E21: Consumption; Saving

KEYWORDS

Household saving, life cycle, age, cohorts

Executive Summary

Elevating household saving rates has been an ongoing focus for policymakers in New Zealand. In particular, concern is often expressed about the rates of saving among younger generations. However, owing to the paucity of readily available household-level information, little is actually known about New Zealand households' saving behaviour beneath the aggregate level.

This paper uses household-level data from the Household Economic Survey (HES) over the period 1984 to 2010 to examine the saving rates of different generations of households in New Zealand. The paper employs the life-cycle model as a framework to estimate regression models that identify differences in average household income, consumption expenditure and saving according to birth year (or "cohort") and age. The basic life-cycle model predicts that saving will exhibit a hump shape over the life cycle, with individuals saving during working age and dissaving during retirement.

The paper's results can be summarised in two main findings. First, household saving over the life cycle does exhibit a hump shape, with a peak in saving rates when household heads are aged in their mid-to-late fifties. This life-cycle pattern of saving is associated with a more pronounced hump-shaped profile in disposable income than in consumption expenditure. Second, there is a trend increase in the average saving rates of cohorts born between the early 1930s and those born in the 1980s. As a result, from the baby boomers onward, the average saving rates of each generation exceed those of the generation preceding it. This trend increase in saving rates reflects ongoing rises in disposable income accompanied by more moderate-to-flat growth in consumption expenditure across cohorts. One potential explanation for the rise in saving rates, which is supported by the findings of other research, is that it reflects household responses to an "unfavourable" evolution in the general economic and policy environments faced by successive cohorts.

These two main findings are robust to various sensitivity tests and are consistent with the results of a previous similar study for New Zealand. However, while no better household-level saving data exists, the potential influence of measurement error in HES presents an important caveat to the analysis. This error is evident in the divergence in trends between the aggregate saving rate based on HES data and the corresponding national accounts measure, and it could bias the estimated effects of age and cohort on saving rates.

Although the potential influence of measurement error cautions against making overly strong inferences, the estimates of age and cohort effects may provide some insight into demographic influences on national household saving trends. In particular, an increase in the proportion of the population in high-saving age groups contributed approximately one quarter of the overall trend increase in the aggregate HES saving rate between 1984 and 2010. However, population projections suggest future positive contributions from this source are unlikely, with the ongoing ageing of baby boomers in retirement expected to weigh upon aggregate saving from the 2020s onward. The remaining three quarters of the aggregate trend increase between 1984 and 2010 is attributable to the rise in average saving rates of successive cohorts born since 1930. If the differences in average saving rates between cohorts persist into the future, these cohort effects are likely to continue to make a positive contribution to aggregate saving rates throughout the projection period ending 2030.

Table of Contents

Abstract	i
Executive Summary	ii
1 Introduction	1
2 Data	2
2.1 The Household Economic Survey	2
2.2 The sample and cohort construction	2
2.3 Defining saving and the saving rate	4
2.4 Comparing Household Economic Survey and national accounts saving rates	5
3 Method: estimating cohort and age effects within a life-cycle framework	6
4 Results	8
4.1 Saving rate age profiles by cohort	8
4.2 Age effects	9
4.3 Cohort effects	11
4.4 Time effects	14
5 Extensions and sensitivity analyses	14
5.1 Alternative measures of saving	14
5.2 Household-level analysis	16
5.3 Allowing variation in age, time and cohort effects	18
6 Implications for aggregate saving	21
7 Conclusion	23
References	24
Appendix	26

List of Tables

Appendix Table 1: HES expenditure categories used in preferred saving measure calculation	26
Appendix Table 2: Expenditure items classified as “investment” expenditure	27

List of Figures

Figure 1: Annual aggregate household saving rates, 1984-2010 (March years)	6
Figure 2: 10-year-birth cohort mean HES saving rates by age of household head.....	9
Figure 3: Estimated age effects on household disposable income and consumption expenditure	10
Figure 4: Estimated age effects on household saving rates	10
Figure 5: Estimated cohort effects on household disposable income and consumption expenditure	11
Figure 6: Estimated cohort effects on household saving rates	12
Figure 7: Estimated time effects on household saving rates	14
Figure 8: Estimated cohort effects on household saving rates with different saving definitions.....	15
Figure 9: Estimated age effects on household saving with different saving definitions.....	16
Figure 10: Estimated cohort effects on household saving rates for cohort means and median households	17
Figure 11: Estimated household saving rates by age for cohort means and median households	17
Figure 12: Estimated effects of household characteristics on the median household saving rate	18
Figure 13: Estimated cohort effects on household saving rates	19
Figure 14: Estimated age effects on household saving rates	19
Figure 15: HES survey frequency distribution by age.....	21
Figure 16: Contributions to predicted aggregate saving rate, 1984 to 2030.....	22

Intergenerational Developments in Household Saving Behaviour

1 Introduction

Elevating household saving rates has been an ongoing focus for policymakers in New Zealand. In particular, concern is often expressed about the rates of saving among younger generations (eg, Savings Working Group, 2011). However, owing to the paucity of readily available household-level information, little is actually known about New Zealand households' saving behaviour beneath the aggregate level.

This paper uses household-level data from the Household Economic Survey (HES) over the period 1984 to 2010 to examine the saving rates of different generations of households in New Zealand. The paper employs the life-cycle model as a framework to estimate regression models that identify differences in average household income, consumption expenditure and saving according to birth year (or "cohort") and age. Knowing these differences is useful for understanding the aggregate saving implications of population ageing and of changes in saving behaviour between older and younger generations. It can also be helpful for assessing the potential effects of various policy interventions designed to raise household saving rates.

The paper extends analysis of HES data for the period 1984 to 1998 by Gibson and Scobie (2001 and 2001a) and Scobie and Gibson (2003). This suite of papers represents the only previously published empirical work that examines household saving in New Zealand using household-level saving data.¹

The rest of the paper is organised as follows. Section 2 introduces the dataset and the method for constructing cohorts. Section 3 outlines the approach to estimation, based on a simple life-cycle framework. Section 4 describes the results, followed by sensitivity analysis and extensions in Section 5. Section 6 considers the implications for aggregate saving rates of the foregoing analysis and Section 7 concludes.

¹ Two recent papers, Scobie and Henderson (2009) and Le, Gibson, and Stillman (2012), derived household saving rates over 2004-2006 using household-level net wealth data from the Survey of Family, Income and Employment. Given the short sample period, neither study was able to identify cohort effects on saving.

2 Data

This section introduces the HES dataset in Subsection 2.1, with particular attention to its limitations. Subsection 2.2 briefly describes the HES sample and the sample restrictions, and the approach to constructing cohorts along with its advantages and disadvantages. Subsection 2.3 provides the definitions used for saving and the saving rate, and finally Subsection 2.4 compares the preferred HES aggregate saving rate measure with the national accounts measure of the national household saving rate.

2.1 The Household Economic Survey

Household saving rates are calculated in this paper as the difference between household income and expenditure as recorded in HES. Each HES survey provides a rich set of income, expenditure and demographic data for an independent and representative sample of New Zealand resident households. The primary purpose of HES is to provide information for the calculation of inflation measures and for some components of the National Accounts. The survey is analogous to the Living Costs and Food Survey in the United Kingdom, the Consumer Expenditure Survey in the United States, and the Household Expenditure Survey in Australia. Although these surveys are typically the best source of household-level saving data and are frequently used in the international literature, they have some drawbacks. Indeed, while there are no better household-level saving data available for New Zealand, Statistics New Zealand warns that HES is not designed to measure saving and should be used for this purpose with caution (Bascand, Cope, and Ramsay, 2006).

The main drawback of HES is that it underestimates actual household income and expenditure. The problem stems from under coverage both of the population (for example those in old-age care institutions are not captured) and of the actual total income and expenditure of those households that are surveyed. It is also likely that coverage varies across different types of households, potentially introducing bias into the age and cohort effects identified in this paper. Bascand, Cope and Ramsay (2006) provide detail on the coverage and other differences between HES and the national accounts Household Income and Outlay Account (HIOA) saving measure. Fesseau, Wolff, and Mattonetti (2013) show that these differences, between “micro” and “macro” income and consumption measures, are common across many countries. The HIOA saving measure is also not free of measurement problems, which has been reflected in substantial historical data revisions over recent years as outlined by Gorman, Scobie, and Paek (2013).

2.2 The sample and cohort construction

This paper uses the 15 annual March-year HES surveys from 1983/84 to 1997/98, the same data set used by Gibson and Scobie (2001)², as well as the four June-year HES surveys conducted in 2000/01, 2003/04, 2006/07 and 2009/10 (HES has been triennial since 2001). The individual HES surveys are referred to henceforth by the year in which they ended. The 19 HES surveys provide a total number of 61,985 household observations, with sample sizes for individual surveys mostly ranging between 2,000 and 4,000 observations.

² The dataset includes some improvements to tax calculations which serve to slightly increase household disposable income for the periods 1984-1987 and 1994-1998 compared with the data used by Gibson and Scobie (2001).

The household, rather than the individual, is used as the unit of analysis, reflecting a view that many consumption and saving decisions are made on a household basis.³ Characteristics of the household “head” define each household’s age, gender, ethnicity and work status. This definition may not provide a good guide to the effects of these characteristics on the saving behaviour of a household if its members are highly heterogeneous. The household head is defined, in the first instance, as the household member with the highest gross income. If household members have equal highest gross incomes then the head is defined as the older member. If household members have both equal highest income and equal age (in years), which is sometimes the case particularly for retired couples, the head is assigned to the member with the lowest Statistics New Zealand HES person number. Alternative definitions of household head, such as the Statistics New Zealand’s “reference person” (as in Gibson and Scobie, 2001), or the oldest household member, do not materially alter the results.

Ideally, panel data, which comprise observations from the same sample 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 using the HES data by dividing the sample into cohorts determined by the birth year of the household head, following a method described by Deaton (1985). The average behaviour of these cohorts can be tracked over time and should be consistent with estimates from a sample of genuine panel data on individuals, provided the membership of the population (and each cohort within it) is fixed.

Gibson and Scobie (2001) describe the advantages and disadvantages of using synthetic panels compared with genuine panel data. In summary, the advantages include less severe problems of sample attrition because fresh samples are used in each survey, and less measurement error and outlier bias because cohort averages should reduce the impact of idiosyncratic variability, which is a feature of data on individuals.

Using synthetic panel data from HES has three main disadvantages. The first problem relates to 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 may also occur when there are common age-related changes to income, such as retirement, which affect the relative earnings of household members and therefore the definition of the household head.

Second, the assumption of fixed cohort membership is difficult to maintain. An obvious violation of this assumption stems from New Zealand’s high rates of inward and outward migration. Another particular concern for this study arises if wealth and longevity are positively related. If this is the case, sample cohort averages will reflect the fact that the population is becoming progressively richer as poor individuals die younger, or are absorbed into younger households. In an attempt to address this “wealth-mortality” bias, this paper follows Gibson and Scobie (2001) and eliminates from the sample all households in which the household head was older than 74 when surveyed (a total of 4,894 observations)⁴. Households with heads who are younger than 19 are also eliminated (412 observations) as well as households with negative disposable income (256 observations). These sample restrictions leave a total sample of 56,423 with birth years ranging from 1910 to 1991.

³ Gibson and Scobie (2001) undertake sensitivity analyses using the individual as the unit of analysis, but note that “extreme” assumptions are required to allocate expenditure (which is only recorded at the total household level) across individuals (p20). They find the age and cohort patterns in saving are similar to those found using the household as the unit of analysis.

⁴ $Birth\ year = survey\ year - age\ at\ time\ of\ survey$, where survey year is the year in which the survey ended.

The third problem concerns the small sample sizes in HES. This means some survey year-cohort averages represent very small sample sizes, affecting the precision of estimates. Interacting the 19 HES survey years with individual birth years gives a total of 1,064 year-cohort “cells”. The smallest number of households in an individual year-cohort cell is 6 (for 19-year old household heads surveyed in 2007) and the median is 50. Summing across all survey years, there is a wide range in the total number of observations in each cohort. Cohorts born between 1946 and 1966 each contain more than 1,000 households, compared with less than 100 households in each of the youngest (birth years 1983 to 1991) and oldest (1910 to 1912) cohorts. By age (again summing across all survey years), the range in the number of observations is narrower, although there are still substantially fewer observations at either extreme of the age spectrum compared with mid-life ages.

2.3 Defining saving and the saving rate

The preferred measure of household saving in this paper has been chosen to correspond closely to the HIOA definition of household saving. Specifically, saving is defined as:

$$S = Y - C \quad (1)$$

where:

- $Y =$ HES “total income”
 - net tax and transfer payments⁵
- $C =$ 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 income and expenditure, plus mortgage principal and life and health insurance payments. More detail on this measure, including the New Zealand Household Expenditure codes for the expenditure categories listed, is contained in the Appendix. Alternative definitions of saving are incorporated as sensitivity analysis in Subsection 5.1.

⁵ HES does not record income tax data. The paper uses HES data adjusted by the New Zealand Treasury’s non-behavioural tax-benefit micro simulation model, Taxwell for the 2007 and 2010 surveys, and by a predecessor of Taxwell for the earlier surveys. These models estimate net tax payment and disposable income measures.

Saving rates are defined in the usual way as saving divided by disposable income. For year-cohort cells, saving rates are calculated as the ratio of each cohort's mean (total) saving to mean (total) disposable income⁶, ie,

$$sr_{bt} = \frac{\sum_h S_{hbt}}{\sum_h Y_{hbt}} \quad (2)$$

where:

S_{hbt} = saving of each household h , belonging to cohort b , in survey year t ;

Y_{hbt} = disposable income of each household h , belonging to cohort b , in survey year t .

The aggregation properties of this “ratio-of-averages” measure (as opposed to a “average-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. In addition, the ratio-of-averages measure reduces the effects of outliers and measurement error. These effects are particularly relevant when using disposable income as the denominator because HES records some households as having very low (or zero) disposable income, which leads to extremely high (or non-defined) negative saving rates for those households and an associated bias in average-of-ratio measures.

A disadvantage of the ratio-of-averages measure is its limited use for understanding behaviour at the household level. Ratios of means are more influenced by higher income households than lower income households, and therefore may be uninformative about households at the median or lower parts of the saving distribution. Ratios of other quantiles, such as the median, can be difficult to interpret because the median household by income is not necessarily the same median household by consumption expenditure, so the median saving rate may be derived from two different households. Therefore, to get a better sense of the behaviour of the “typical” household, household-level data, rather than cohort-year cells, are used as the unit of analysis in the extensions discussed in Subsection 5.2.

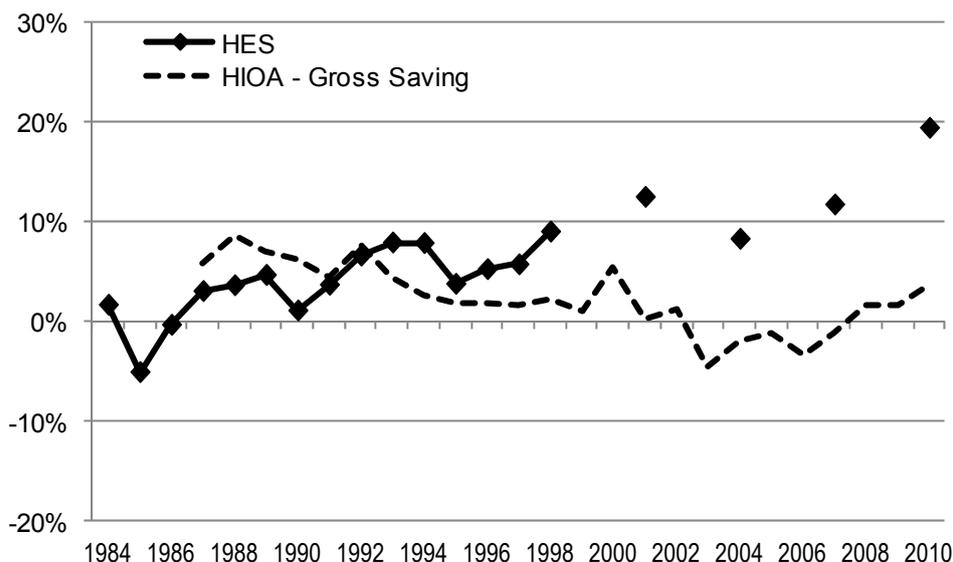
2.4 Comparing Household Economic Survey and national accounts saving rates

Figure 1 shows the aggregate HES saving rate by year alongside the HIOA gross (ie excluding consumption of fixed capital) saving series. The HES saving rate shows a relatively steady upward trend over the sample period. The HIOA saving rate is generally lower than the HES measure and, at least up until 2001, appears to show a downward trend. This divergence in HIOA and HES trends mirrors the divergence in trends seen between analogous measures in the UK and US (Barrett, Crossley, and Milligan, 2010). Although the move to a triennial HES complicates comparison in more recent years, HIOA and HES measures appear to show a similar pattern from 2001, with a dip in the early 2000s and subsequent recovery. Nevertheless, the historic differences between movements in the HES and HIOA measures indicate caution is needed when drawing conclusions about changes in the latter from changes in the former.⁷

⁶ Population weights have not been applied in these calculations or any analysis in this paper because of the inconsistency created by the change in Statistics New Zealand's weighting method from 2001. However, applying Statistics New Zealand's (inconsistent) weights does not lead to materially different results from the unweighted analysis.

⁷ The sample selection, outlined in Subsection 2.2, has little effect on the aggregate HES measure.

Figure 1: Annual aggregate household saving rates, 1984-2010 (March years⁸)



3 Method: estimating cohort and age effects within a life-cycle framework

The basic life-cycle model assumes that individuals smooth consumption over their lifetimes, saving in one period to consume in another. Because income is typically “hump shaped” over the life cycle, the life-cycle model predicts that saving will typically exhibit a corresponding hump shape, with individuals saving during working age and dissaving during retirement. Saving behaviour will therefore differ between different individuals at different points in their life cycles. It may also vary over time, and across cohorts, because of the effects of public policy changes, economic growth, and/or fluctuations in the economic cycle that impact different cohorts (at different ages) contemporaneously. This paper generally follows Gibson and Scobie (2001) by estimating a life-cycle model that allows the separate identification of cohort, age and time effects on saving behaviour, as developed by Deaton (1997).

In Deaton’s model, an individual’s consumption expenditure level is proportional to lifetime wealth (W), known with certainty at birth, with a factor of proportionality (determined by preferences) that depends on age. Interest rates are ignored.⁹ Therefore, for an individual i born in year b and observed in year t (age = $t - b$), with preferences represented by the function g , consumption expenditure can be expressed as:

$$C_{ibt} = g_i(t - b)W_{ib} \quad (3)$$

This model can be adapted to households by assuming lifetime wealth is known at the time of household formation, with the function g representing household preferences. Taking logs of Equation 3 and then averaging over all households with a household head in the cohort born in year b and observed at t gives:

$$\overline{\ln C_{bt}} = \overline{\ln g}(t - b) + \overline{\ln W}_b \quad (4)$$

⁸ HES saving rates are in June years from 2001.

⁹ The influence of changes in the interest rate on saving rates is theoretically ambiguous, depending on the relative strength of income and substitution effects, and most empirical work finds small or nonexistent effects.

so that average consumption expenditure is decomposed into the sum of two components, one that depends only on age and one that depends only on cohort. Equation 4 can be estimated by regressing the mean of the logarithms of consumption expenditure for each cohort in each survey year on a set of age and cohort dummy variables. The coefficients on the age dummies recover preferences about intertemporal choice and the coefficients on the cohort dummies capture the lifetime wealth levels of each cohort. There is no need to assume a functional form for preferences or to measure lifetime wealth levels.

This model is consistent with some level of uncertainty, provided members of each cohort estimate their future earnings correctly on average. It is also possible to incorporate the effects of “one-off” macroeconomic shocks that surprise households by adding a function representing time effects. Because of the identification problem caused by the linear relationship between age, cohort and time, it is not possible capture time effects by simply adding a set of survey-year dummies. To overcome this problem and allow for the separate identification of age, cohort and time effects, a normalisation is used to make the time effects sum to zero and orthogonal to a time trend.¹⁰ This approach effectively attributes any time trends in the data to a combination of age and cohort effects, not to time effects, which capture cyclical fluctuations that average to zero over the long run. The normalization is restrictive, but can be justified on several grounds, as discussed by Attanasio (1998).

Introducing the normalized time effects together with the age and cohort dummies, and adding variables for the mean number of children n^c (defined as individuals aged 15 or younger) and adults n^a in each year-cohort cell to allow their different consumption requirements, Equation 4 can be estimated as:

$$\ln \overline{C}_{bt} = D^a \alpha_C + D^b \beta_C + \gamma_C(d_t) + \delta_{C1}n^a + \delta_{C2}n^c + u_c \quad (5)$$

where D^a is a matrix of age dummies, D^b is a matrix of cohort dummies, the coefficients α_C and β_C are the age and cohort effects on consumption expenditure, γ_C represents time effects, d_t the δ_C s control for the effect of demographic composition, and u_c is the error relating to the sample estimate of log mean consumption expenditure for households born in year b and observed in year t . The estimation uses the logarithm of the mean rather than the mean of the logarithms (shown in Equation 4) to better account for the measurement problems discussion in Subsection 2.1.

Household income can be treated in the same way as consumption expenditure. The underlying relationship corresponding to Equation 5 is that income can be expressed as proportional to lifetime resources, where the factor of proportionality depends on age.¹¹ The corresponding estimated equation for income is thus:

$$\ln \overline{Y}_{bt} = D^a \alpha_Y + D^b \beta_Y + \gamma_Y(d_t) + \delta_{Y1}n^a + \delta_{Y2}n^c + u_c \quad (6)$$

The difference between the logarithm of income and the logarithm of consumption expenditure is a monotone increasing function of the saving-to-income ratio. When the saving ratio is low, this difference is approximately equal to the saving ratio, and so

¹⁰ Specifically, d_t^* is defined as follows $d_t^* = d_t - [(t-1)d_2 - (t-1)d_1]$, where d_t is the usual time dummy. The coefficients of d_t^* give the third to final year coefficients, and the first and second coefficients can be derived from the constraint that all time effects sum to zero and are orthogonal to a time trend (Deaton 1997).

¹¹ This relationship assumes that income has an unchanging age profile. Therefore economic growth only affects the position of each cohort's age profile and not the age profiles themselves.

together, the income and consumption expenditure cohort-age-time decompositions provide a decomposition of the saving rate as:

$$sr_{bt} \approx \ln \overline{Y_{bt}} - \ln \overline{C_{bt}} = D^a(\alpha_Y - \alpha_C) + D^b(\beta_Y - \beta_C) + (\gamma_Y - \gamma_C)(d_t) + u_c \quad (7)$$

where the terms in brackets represent the respective effects on the saving rate of age, cohort, and time. Demographic effects are ignored for reasons outlined in the next section. Rather than rely on the approximation of $(\ln \overline{Y_{bt}} - \ln \overline{C_{bt}})$ as the dependent variable, Equation 7 can be estimated directly using sr_{bt} as the dependent variable and this is main the specification adopted for analysing saving behaviour in this paper.

4 Results

This section reports regression estimates for the effects of age, cohort and time on saving rates corresponding to the equations for consumption expenditure (5), income (6) and saving (7).¹² Although saving rates are the principal focus, the separate consumption and income analysis provides a useful breakdown of the saving patterns. A constant is included in each equation, and variables for the mean number of children and the mean number of adults are also included in the consumption expenditure and income equations, but excluded from the saving equation. Excluding these demographic variables (which enter in the extensions covered in Section 5) from the main saving equation makes little difference to the pattern of results (implying their effects on income and consumption offset one another), but greatly simplifies consideration of aggregate implications covered in Section 6.

Each equation is estimated on 1,064 year-cohort cells using weighted least squares, with the weights equal to the number of household observations in each year-cohort cell. The weighting method provides an efficient way of estimating parameters when using cell averages (as opposed to household-level data), by accounting for the greater variance in cells with fewer observations.

To provide some context to the regression estimates, Subsection 4.1 first shows the raw saving rate data by household head age divided into 10-year-birth cohort groups. Subsections 4.2 and 4.3 help to clarify the observed patterns in the raw data by reporting regression estimates, in the form of figures, for age and cohort effects. Finally, Subsection 4.4 reports estimated time effects for the saving Equation 7.

4.1 Saving rate age profiles by cohort

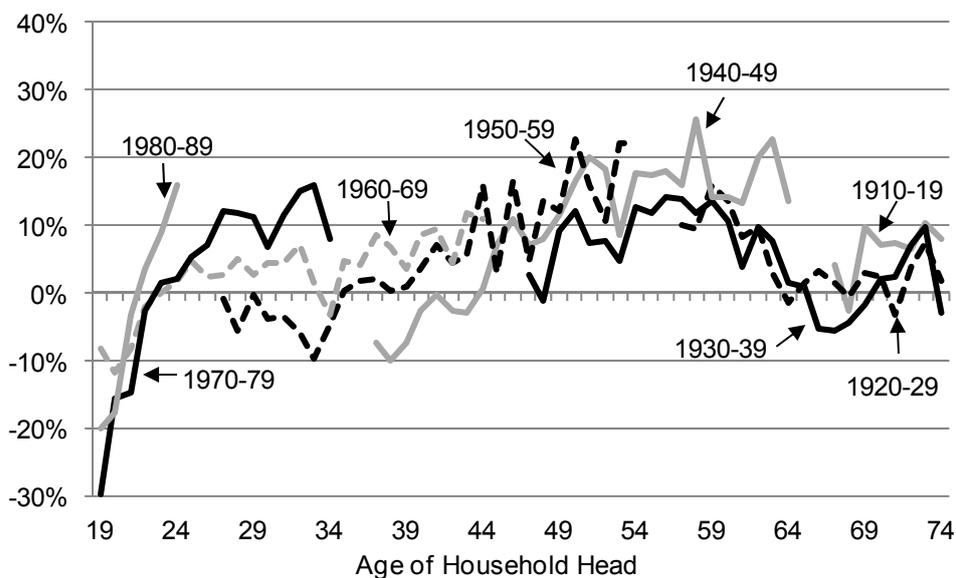
Figure 2 provides a first glimpse of the age and cohort patterns in household saving. Each point in the figure shows the mean cohort saving rate, across surveys, at a particular household head age.¹³ For example, the saving rate at age 35 for the 1950-59 cohort is the mean of the rate for 35-year-old-headed households across the 10 surveys from 1985 to 1994. Although the use of 10-year-birth cohorts smoothes the picture considerably, substantial underlying volatility in the data is still evident. Two interesting patterns are nonetheless observable. First, there is a tendency for cohorts born after the 1930s to have

¹² Real income and consumption measures are used in the estimation. Nominal data are deflated by the Reserve Bank of New Zealand's "Consumers Price Index (excluding interest rates)" measure.

¹³ For each cohort, the figure excludes data points for ages that are observed in two or less survey years. This censoring reduces the noise at the ends of the cohort "lines" without affecting the overall patterns.

successively higher saving rates than previous cohorts. It is less clear whether there is a difference in behaviour between the older three cohorts. Second, there appears to be a hump-shaped profile in saving rates over the life cycle, as predicted by the life-cycle model. Saving rates tend to be lower (or negative) at the younger household head ages, before rising to peak when household heads are in their fifties. Saving rates decline at ages thereafter, but they do not become consistently negative at older household head ages and in fact appear to increase from ages in the mid-to-late sixties.

Figure 2: 10-year-birth cohort mean HES saving rates by age of household head



A rise in saving rates during old age is inconsistent with the predictions of the life-cycle model, and it perhaps reflects some remaining wealth-mortality bias (or household dissolution bias) in the data, despite the truncation of the sample to household heads younger than 74 years old. On the other hand, non-negative saving rates in old age are not surprising here because of the treatment of pension income as income rather than as the drawdown of savings. Jappelli and Modigliani (1998) argue that public pension payments should be treated as dissaving, rather than as transfer income, and that tax and other payments that contribute to these pensions should be treated as saving, rather than as reductions in disposable income. Unless these adjustments are made, they claim that true household saving will be understated during the pre-retirement period and overstated during retirement. Coleman (2006) makes these adjustments for New Zealand and finds a marked impact on measured saving rates in the expected directions. Such tax and pension payment adjustments are not made in this paper.

4.2 Age effects

Figures 3 and 4 show regression estimates for the mean effects of age (an “age profile”) on household disposable income, consumption expenditure and saving rates, compared with the reference household head age of 19 years old.¹⁴ It is worth emphasising that because no one cohort is observed across all household head ages, the single age profile for all cohorts is estimated using observations at different ages from different cohorts, controlling for cohort and time effects.

¹⁴ Each point in Figures 3 and 4 corresponds to an estimated coefficient from a regression of Equations 5 and 6, or 7.

The age profiles shown in Figure 3 show the estimated mean percentage difference that age makes to disposable income and consumption expenditure.¹⁵ The estimates include controls for the number of adults and children in the household. Income exhibits a clear hump shape, rising through younger ages to peak when household heads are in their 50s, before falling at older ages. Consumption expenditure also exhibits a hump shape, but it is less pronounced. The variation in consumption expenditure over the life cycle, which is consistent with findings for other countries, appears to contradict the predictions of the basic life-cycle hypothesis. However, it is not inconsistent with more sophisticated life-cycle models that incorporate, for example, precautionary motives – as discussed by (Browning and Crossley, 2001).

Figure 3: Estimated age effects on household disposable income and consumption expenditure

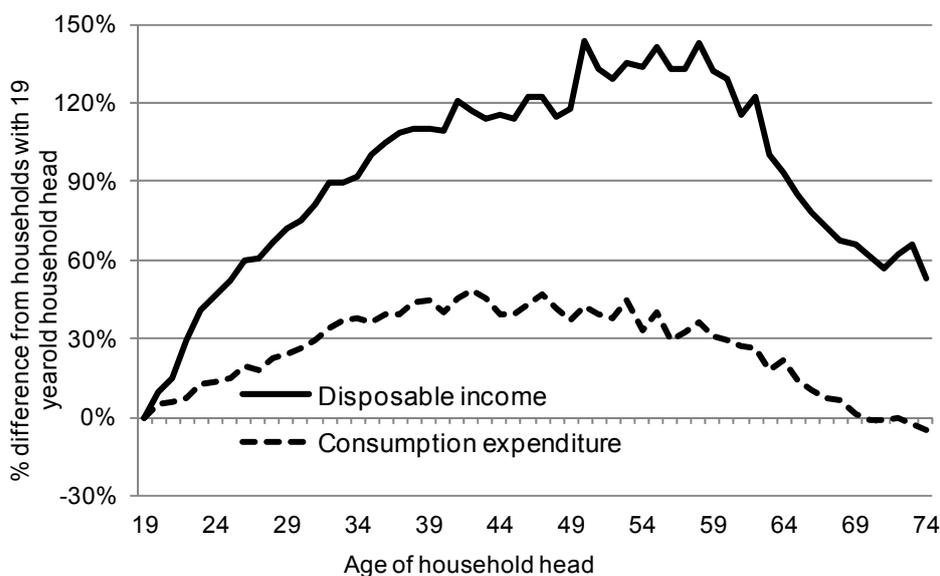
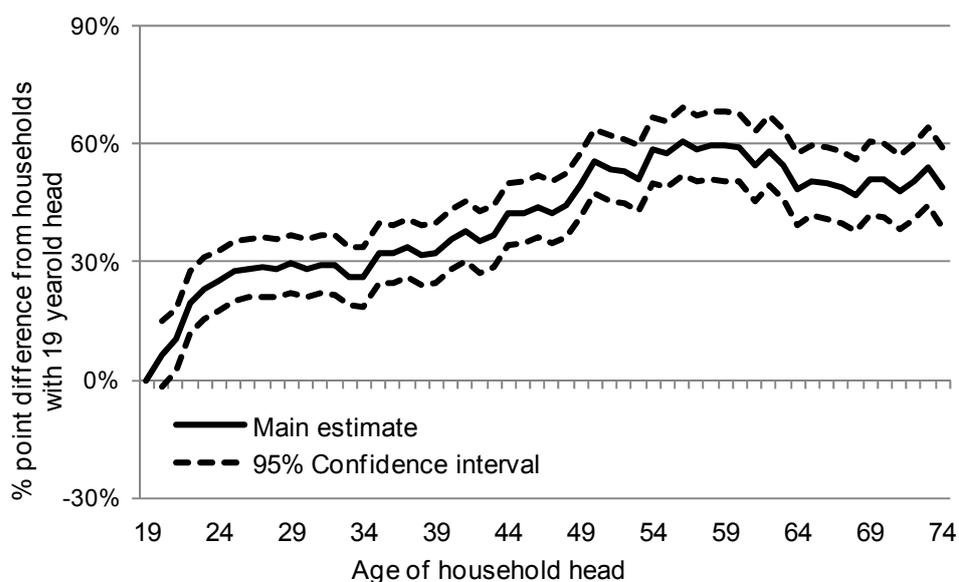


Figure 4: Estimated age effects on household saving rates



Having controlled for cohort and time effects, the estimated age profile in Figure 4, which includes 95 percent confidence intervals, exposes the profile in household saving rates more clearly than Figure 2. Household saving increases sharply when household heads

¹⁵ Estimates are converted from log values into percentage terms by taking the exponent of estimated coefficients and subtracting 1.

are in their early twenties, before rising more gradually to peak when they are in their mid-to-late fifties. Although they decline at ages thereafter, saving rates for households with older household heads remain generally above those of households with heads in their thirties and forties. The late-age rise in saving rates apparent in Figure 2 is also visible in Figure 4.

4.3 Cohort effects

4.3.1 Cohort effect estimates

Estimated cohort effects for the consumption expenditure, income and saving equations are shown in Figures 5 and 6. The reference cohort is set to the 1930 birth cohort. The cohort effects in Figure 5 represent the estimated percentage difference that birth year makes to the mean consumption expenditure or income of a cohort, compared with the 1930 cohort, at any given age between 19 and 74. For example, the mean disposable income of a household headed by someone born in 1960 is estimated to be around 50 per cent higher, at any age, compared with a household headed by someone born in the 1930. Figure 5 shows a trend increase in both mean disposable income and consumption expenditure by cohort between birth years in the 1910s and 1950s, with consumption expenditure rising more rapidly than income across the earlier-born cohorts and less rapidly across the later-born cohorts. For cohorts born after the 1950s, consumption expenditure appears to broadly plateau, while disposable income continues generally to increase by cohort up until birth years in the 1980s. There then appears to be some decline in both disposable income and consumption expenditure through cohorts born during the 1980s. However, there is clearly more year-on-year volatility (and lower statistical significance) in the estimates for these youngest cohorts, reflecting a smaller number of household observations in the sample. More caution is therefore required in interpreting the estimates for these youngest cohorts.

Figure 5: Estimated cohort effects on household disposable income and consumption expenditure

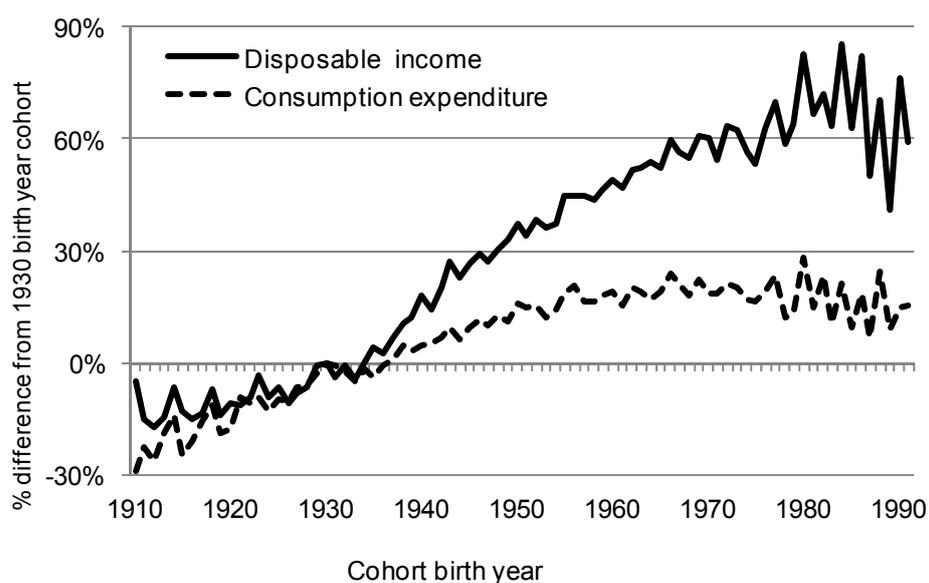


Figure 6: Estimated cohort effects on household saving rates

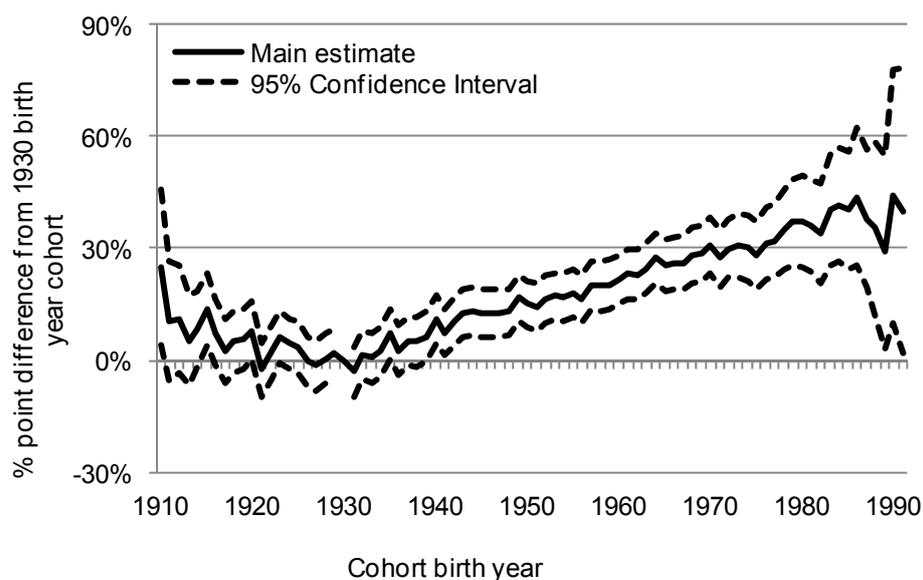


Figure 6 shows the estimated cohort effects on saving rates, along with 95 per cent confidence intervals. These effects indicate the estimated percentage *point* difference that birth year makes to the mean saving rate of a cohort. Consistent with the cohort trends in income and consumption expenditure, there is a general decline in the mean saving rate by cohort between those with birth years from 1910 to around 1930. For cohorts born subsequently there appears to be a near-linear-trend increase in mean saving rates, through to those born in the mid-1980s. The rise in estimated mean saving rates by cohorts over this period is substantial. Saving rates of households with a “baby boomer” head (roughly those born between the mid-1940s and the mid-1960s) are approximately 15 percentage points higher than households with heads in the “silent generation” (born in the 20 years previous). Saving rates of households with heads in Generation X (born between 1965 and 1980) are approximately 12 percentage points higher than those of households with baby-boomer heads. Saving rates of households with heads in Generation Y (born between 1981 and 1991) are approximately 8 percentage points higher again.

The widening confidence intervals towards the right (and left) of Figure 6 clearly illustrate the loss of statistical significance that occurs at both ends of the birth-year spectrum, as a result of smaller samples. In addition, because these cohorts are observed at relatively few ages and in few surveys (compared with the middle cohorts), there is a need to be more cautious in interpreting the estimates as being representative of their lifetime behaviour.

4.3.2 Cohort effects interpretation

According to the basic life-cycle model, cohort effects should be zero, because any difference in lifetime income across cohorts is matched by differences in lifetime consumption. An exception, in which cohort effects would be expected to be positive, is possible if intergenerational bequests, as a proportion of income, increase with the level of income (ie, bequests are a luxury good). Empirically, estimated cohort effects might also be expected to be nonzero if the retirement period is not fully captured in the data. In this case, a positive estimated cohort effect could represent greater accumulation of wealth by that cohort during working life to support consumption during the retirement period. This

might occur, for example, if younger cohorts have lower expectations of the level of publicly provided services provided in old age, or because of increases in life expectancy.

Empirical studies using comparable frameworks to this paper have interpreted variation in estimated cohort effects in different ways. Some, such as Attanasio (1998), Scobie and Gibson (2003) and Chamon and Prasad (2010) assume that they indicate true differences in saving behaviour between cohorts. Others consider nonzero cohort effects as anomalous features of the data (Deaton, 1997) or as representative of measurement error (Dynan, Edelberg, and Palumbo, 2009). In the context of this study, it would seem reasonable to assume that the pattern of cohort effects likely reflects elements of both measurement error and true effects.

The existence of measurement error in the cohort effects is suggested by the difference in trends between the aggregate HES and HIOA saving rate measures shown in Figure 2 together with the fact that the estimated cohort effects account for much of the trend increase in the HES measure as discussed in Section 6. An example of how measurement error may bias the estimated cohort effects relates to the survey coverage issues discussed in Subsection 2.1. If a category of expenditure (or income) is underestimated by HES, and the proportion of household expenditure on this category increases systematically with birth-year cohort (adjusting for age), then this would bias estimated cohort effects upward compared with their true value. Further work being undertaken by Statistics New Zealand may shed more light into how measurement error in HES may affect the findings of this study.

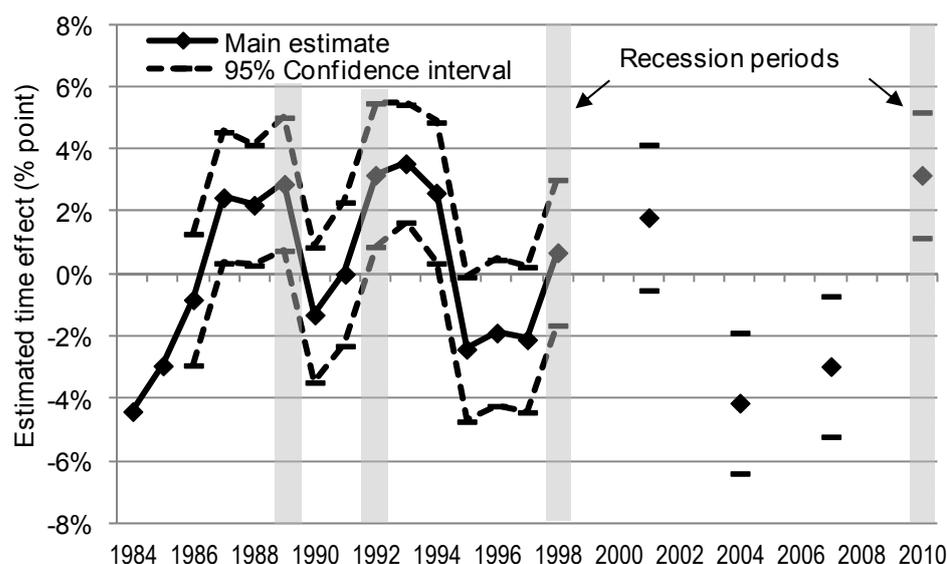
On the other hand, as argued by Scobie and Gibson (2003), the identified differences between cohort saving rates are consistent with the evolution of policy and economic conditions during the last century, and may therefore reflect true cohort effects. In particular, the period from 1950 to 1980 was marked by the prevalence of relatively “favourable conditions” for New Zealand households, with high levels of public sector welfare provision and low unemployment. This benign period would explain why cohorts who were in their peak-earning ages at the time are found to have lower lifetime saving rates than older or younger cohorts. In support of this argument, Talosaga and Vink (2014) provide strong empirical evidence showing the lift in the eligibility age for New Zealand’s public pension, New Zealand Superannuation, led to higher saving rates among affected (younger) cohorts.

Increases in life expectancy provide another plausible explanation for rising cohort effects over recent generations. For example, cohorts born in 1991 (the youngest in the sample) are projected to have a life expectancy at age 50 years old that is more than 10 years longer than cohorts born in 1930 are estimated to have had at the same age (Statistics New Zealand, 2014). The effect of these increases in life expectancies on saving rates will depend on the extent to which households make corresponding adjustments to their expected retirement age. The effects will be lower if longer expected lifetimes are matched with longer expected working lives.

4.4 Time effects

Figure 7 shows estimated time effects for saving Equation 7, along with 95 per cent confidence intervals and an indication of the timing of New Zealand's economic recessions.¹⁶ As noted in Section 3, these time effects sum to zero, are orthogonal to a time trend, and can be interpreted as representing macroeconomic shocks. Consistent with the literature, the estimates suggest recessions are associated with higher saving rates, while booms tend to correlate with lower saving rates.¹⁷ Although the estimated time effects are statistically significant, their exclusion from Equations 5, 6 and 7 does not materially affect the pattern of age and cohort effects discussed above.

Figure 7: Estimated time effects on household saving rates



5 Extensions and sensitivity analyses

This section provides estimates for alternative specifications to those used to generate the main results presented in the previous section. These alternative specifications both provide a check of the main results as well as additional insights into household saving behaviour. Subsection 5.1 considers the analysis using two definitions of saving that have been used in other studies of saving in New Zealand.¹⁸ Subsection 5.2 changes the unit of analysis to the household and examines the effects of different household characteristics on saving behaviour. Finally, Subsection 5.3 considers several examples which relax the assumption that age, cohort and time effects are constant.

5.1 Alternative measures of saving

As discussed in Subsection 2.3, the preferred measure of saving in this paper corresponds to the HIOA definition of household saving. However, other definitions may be preferable from an economic perspective. Gibson and Scobie (2001) use a definition of saving that classifies expenditure on items that provide consumption benefits over more

¹⁶ The timing of these recessions is highly approximate because each HES survey is conducted over a one-year period, with income data and some expenditure components recorded for the year preceding the interview date. This means the date pertaining to some data may vary by up to 24 months between households in the same survey.

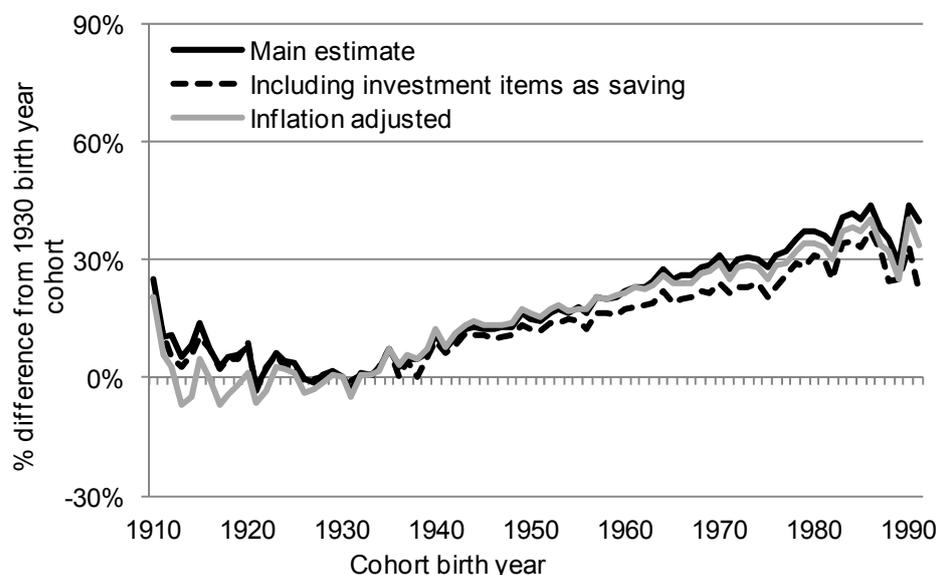
¹⁷ A recent example of the literature discussing household saving and economic recessions is Alan, Crossley and Low (2012).

¹⁸ More detail on the construction of these saving measures is included in the Appendix.

than year; such as consumer durables, health and education; as saving rather than consumption. The argument for this definition is based on the element of durability of these expenditures which means they may be better considered as “investment items” (and therefore a form of saving) rather than consumption.¹⁹ Gorman, Scobie and Paek (2013) make an adjustment for these investment items and show substantial effects on measured household saving at the aggregate level in New Zealand. They also calculate another household saving measure, which includes an adjustment to incorporate the fact that the inflation component of nominal interest charged on outstanding financial liabilities is an implicit capital repayment (not an income payment) to the lender. If the inflation component of interest payments is considered capital, unadjusted household saving rates overstate (understate) the “true” saving of lenders (borrowers), especially when inflation is high.

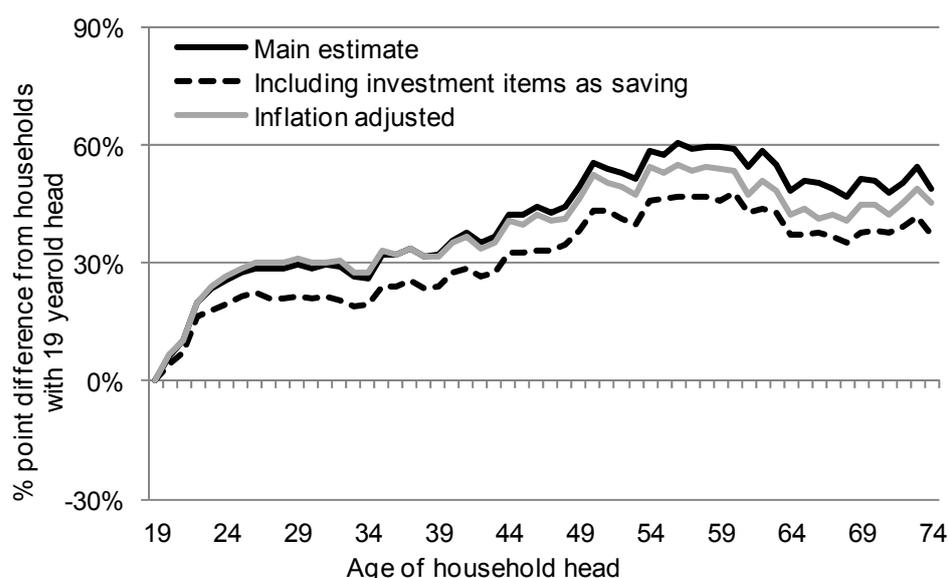
Figures 8 and 9 show the results for age and cohort effects of estimating Equation 7 using these two alternative saving measures. The overall pattern in both age and cohort effects is similar for each of the measures. Of the two alternatives, classifying investment expenditure items as saving leads to larger differences from the HIOA measure, with reduced cohort effects (Figure 8) and a less pronounced age profile (Figure 9). These differences reflect the fact that consumer durable spending as a ratio of income is highest for younger age groups and that this ratio has declined substantially over the sample period as the price of consumer durables to nondurables has fallen.

Figure 8: Estimated cohort effects on household saving rates with different saving definitions



¹⁹ The ideal, but unworkable, approach here would be to exclude changes in the stock of durable expenditures from consumption expenditure and to add to consumption expenditure the value of services obtained from the stock.

Figure 9: Estimated age effects on household saving with different saving definitions



5.2 Household-level analysis

As discussed in Subsection 2.3, the ratio-of-averages saving rate measure used in the paper up until this point provides limited insight into the saving behaviour of typical households. In this section, Equation 7 is re-estimated using household-level saving rates as the dependent variable. Quantile regression is used to reduce the effects of outliers and measurement error, which would lead to substantial bias in least squares regression because of the presence of households with near-zero incomes.²⁰ Using household-level data also allows the model to be conditioned for household characteristics. These characteristics are captured by including dummy variables representing differences in gender, ethnicity, employment status, dwelling tenure, family structure and education.²¹ The influence of these conditioning variables on saving rates is interesting in its own right. Because of the potential for the composition of the sample to change over time, the conditioning variables also provide a useful check on the robustness of the estimates.

Figures 10 and 11 compare the estimated age and cohort effects reported in Subsections 4.2 and 4.3 with those estimated using median regression, with and without conditioning variables. Overall, the results are similar for the different specifications, lending support to the robustness of the main results. The size of the age and cohort effects is somewhat lower at the median than for the main estimates, suggesting that the influence of age and cohort on saving is most marked at the upper end of the income distribution. The results from quantile regressions at the 25th and 75th percentiles (not shown) corroborate this suggestion, with respectively lower and higher degrees of variation in estimated age and cohort effects than at the median.

²⁰ The 79 households with zero recorded income are excluded from the sample, to leave a total sample size of 56,344.

²¹ Questions relating to educational qualifications were not available for some households in the earlier survey years. The conditioned equations are therefore estimated on a reduced sample. Estimates for cohort, age and the other conditioning variables are not materially affected by the reduced sample.

Figure 10: Estimated cohort effects on household saving rates for cohort means and median households

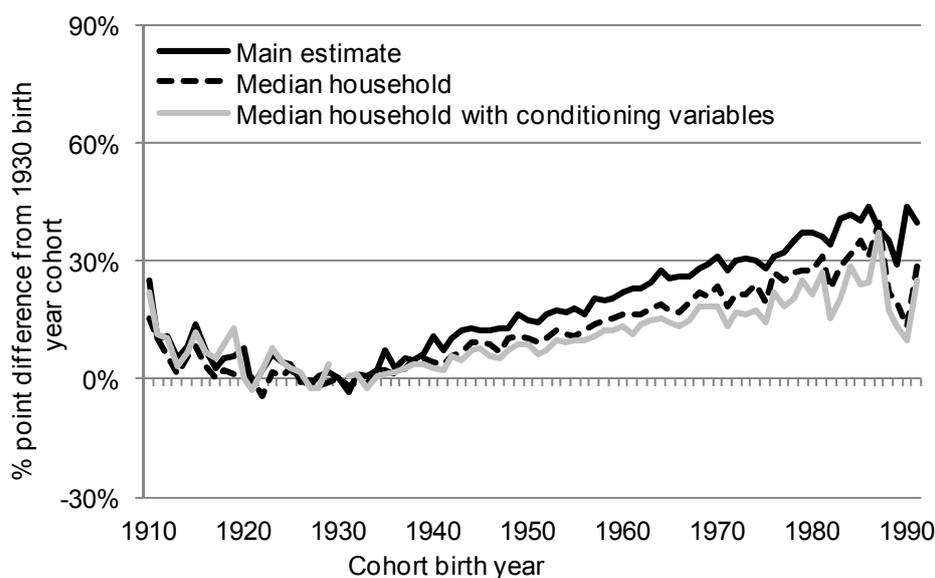
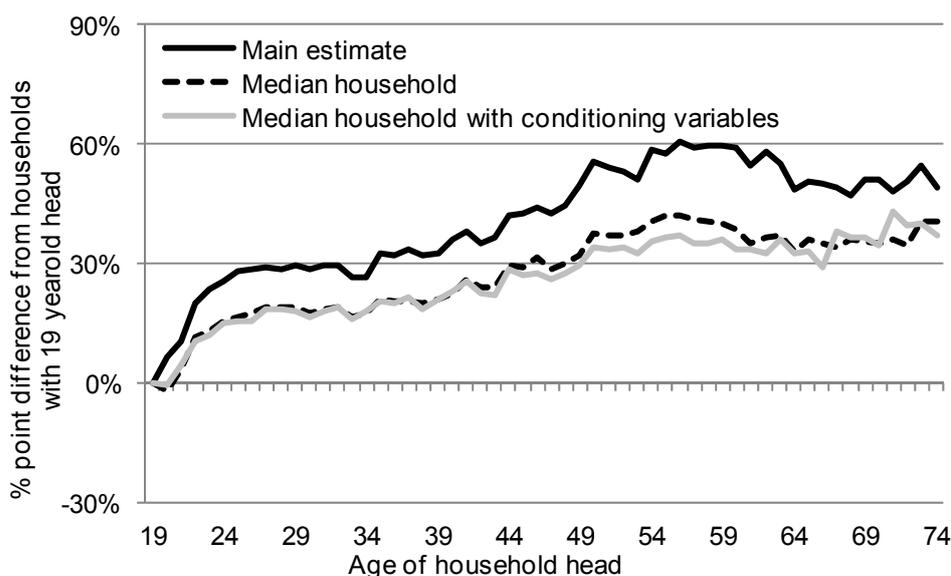


Figure 11: Estimated household saving rates by age for cohort means and median households

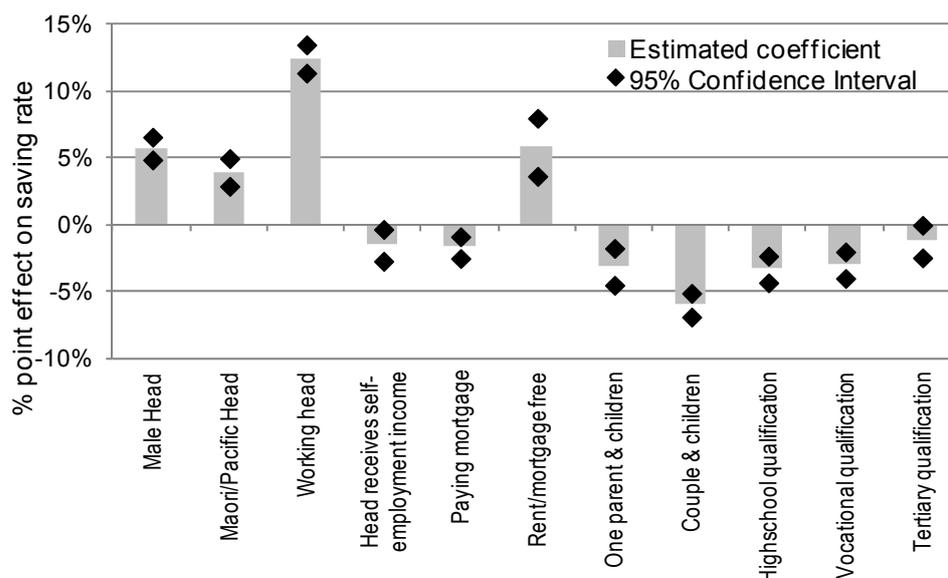


The estimated coefficients of the conditioning dummy variables are shown in Figure 12 and discussed briefly below.²² Gender has relatively large and statistically significant saving effect, with male-headed households saving nearly six percentage points more the female-headed households. Māori or Pacific ethnicity has a smaller positive but significant effect, with Māori or Pacific-headed households' saving rates four percentage points higher than those of households with heads of other ethnicities. Owning a house with a mortgage has little effect on saving rates compared with renting, but owning a house freehold is associated with saving rates six percentage points higher than for renting households. Having children has a negative effect on saving rates, lowering saving rates by two percentage points for sole parents and six percentage points for couples. The effects of basic educational qualifications, while small, are surprisingly negative, at two percentage points each for high school and vocational qualifications, while tertiary

²² It is beyond the scope of this paper to explore these results in detail - this could be a fruitful avenue for future work.

education has no statistically significant effect. This apparent anomaly is partially explained by the correlation of these variables with employment. Household head employment has the largest effect of the conditioning variables at 12 percentage points (compared with not working). Reflecting this and its correlation with other factors, excluding the employment dummy has a significant effect on several of the other coefficients. Most notably it reduces the size and significance of the negative saving effects of educational qualifications, raises the negative effect of sole parenthood, raises the positive effect of male gender and reduces the effect of Maori or Pacific ethnicity.

Figure 12: Estimated effects of household characteristics on the median household saving rate²³



5.3 Allowing variation in age, time and cohort effects

The empirical model used in this paper assumes that age, time and cohort effects are constant and independent from one another. This means, for example, the shape of the estimated age profile of saving does not vary across time, in response to policy or other environmental changes, or by cohort. Rather, average differences between the saving behaviour of cohorts are captured by a level shift (the cohort effect) in the age profile, which is constant across all ages. There are limitations in the extent to which these assumptions can be relaxed because no cohorts are observed across all ages and time periods. However, the following subsections provide some insight into how age and cohort effects may have evolved over time (5.3.1), how age effects may differ by cohort (5.3.2) and how cohort effects may differ by age (5.3.3).

5.3.1 Changes in age and cohort effects over time

This subsection compares estimates for Equation 7 using a sample comprising the first nine surveys (1984 to 1993) with estimates using a sample comprising the remaining eight surveys (1994 to 2010). Time effects have been excluded from the equations because the shorter sample periods make it difficult to separate trends from fluctuations. This exclusion, together with generally smaller sample sizes, reduces the precision of the

²³ The excluded dummies are: “female”, “non-Māori/Pacific”, “not working”, “rented dwelling”, “single adult family”, “couple without children”, “other family types”, and “no secondary-school qualification”.

estimates. Nonetheless, as shown by Figures 13 and 14, the estimated cohort and age effects show similar overall patterns, with some differences, for the two samples. For cohort effects, there is sharper rise for middle cohorts in the earlier sample, perhaps reflecting precautionary-type saving among these cohorts who were of prime-working age during the turbulent economic years of the late 1980s and early 1990s. For age effects, the profile in the later sample is generally flatter, apart from a steep increase between household head ages 19 to 25 years old. In addition, in the later sample the old-age decline in saving occurs at an age around five years older than in the earlier sample. This delayed decline may reflect the influence of the increase in the eligibility age for New Zealand Superannuation, from 60 to 65 years old between 1992 and 2001.

Figure 13: Estimated cohort effects on household saving rates

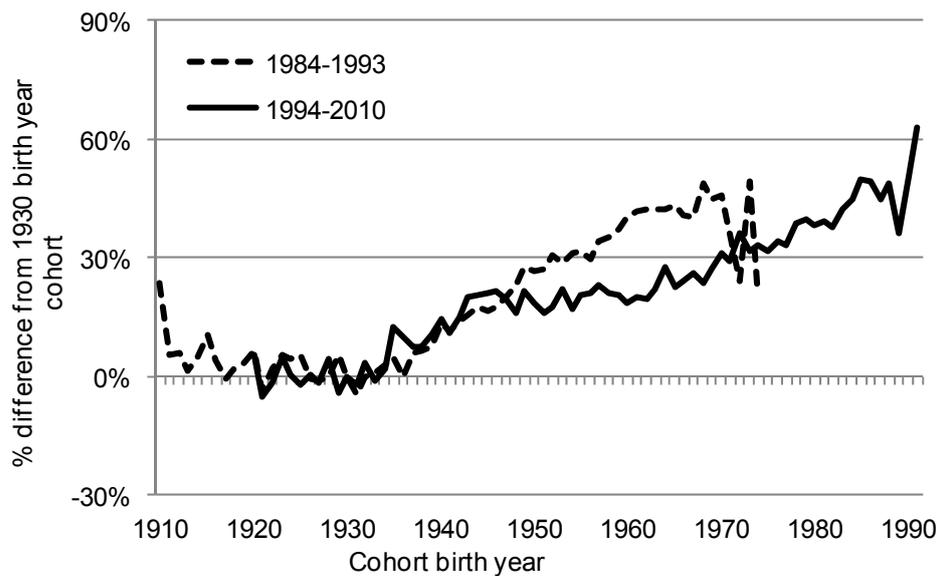
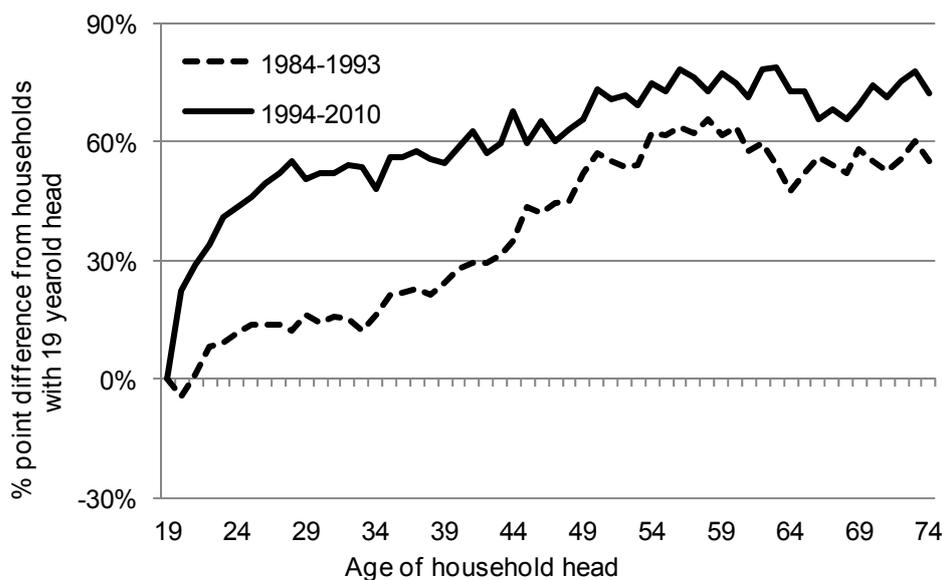


Figure 14: Estimated age effects on household saving rates



5.3.2 Changes in age effects by cohort – accounting for the change in the age of pension eligibility

The increase in the eligibility age for New Zealand Superannuation only affected cohorts born after 1932. If, as suggested by Figure 14, this increase affected the saving behaviour of these cohorts, it may be distorting the estimated age profile. A potential way to address this distortion is to replace the age dummy variables in Equation 7 with dummy variables based on “years-until-expected-retirement”, calculated as *Expected New Zealand Superannuation eligibility age_t – age_t*.²⁴ The expected eligibility ages for New Zealand Superannuation for each birth cohort at time *t* are assumed to be those of government policy at time *t*, and these are set out in more detail by Talosaga and Vink (2014).²⁵ The pattern of estimated cohort and age effects for this alternative specification are very similar to those of the preferred specification presented in Figures 4 and 6. As expected, the lifetime profile has a more pronounced hump shape with a peak just prior to the expected retirement age and a more consistent decline over post-retirement ages.

5.3.3 Changes in cohort effects pre- and post-retirement age

The assumption of constant cohort effects across the life cycle, which is required for the identification of a saving age profile, is a strong one. It implies that differences in saving between cohorts cannot change across time, because of policy changes for example, and that differences in saving between cohorts are never spent within the sample age range. As noted in Subsection 4.3.2, the second implication is only tenable, if it is assumed that higher saving cohorts have higher dissaving at ages above the sample range, or higher bequests. In effect, this assumption involves the reversal of estimated positive and negative cohort effects in later life.

The retirement age is a natural point at which a reversal in cohort effects might be expected. One way to roughly test whether this reversal occurs is to split the sample into those households with heads who are eligible to receive New Zealand Superannuation (as a proxy for retirement) and those who are not, and to rerun the regressions on the two subsamples. The estimated cohort effects for the “pre-retirement subsample” closely match those estimated for the full sample. However, the estimates for the “retirement subsample” are substantially lower, with negative cohort effect point estimates for more than half of the cohorts in the subsample. The majority of negative point estimates reverses the pattern of nearly all positive estimates for cohorts, shown in Figure 6. Although a crude test, this result would appear to be consistent with the life-cycle model’s prediction that cohorts with higher saving rates during working ages have lower rates of saving (or higher rates of dissaving) during retirement.

²⁴ The author thanks Andrew Coleman for this suggestion.

²⁵ In addition to the changes announced in the 1991 Budget, the expected eligibility age variable also accounts for the more gradual increase in the New Zealand Superannuation eligibility age announced in 1989, which involved the eligibility age increasing from 60 to 65 years old between 2006 and 2026. Obviously there may be differences between individuals’ expectations in relation to the future eligibility age and actual government policy at the time.

6 Implications for aggregate saving

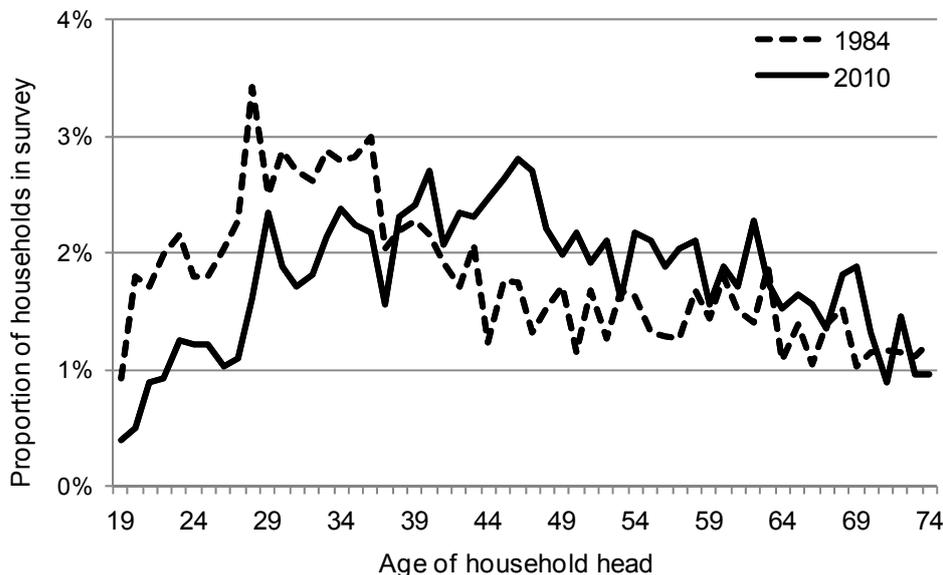
A useful feature of the life-cycle model outlined in Section 3 is its provision of a simple accounting framework for describing changes in the aggregate saving rate according to changes in population structure and income growth. Using this framework, the trend aggregate saving rate can be predicted as the weighted sum of age effects (or “life-cycle effects”) and cohort effects, where the weights are determined by each cohort’s share of aggregate disposable income. Specifically, the trend aggregate saving rate can be estimated as:

$$\bar{s}r_t = \underbrace{\frac{\sum_b Y_{bt}(\hat{\alpha}_s)}{\sum_b Y_{bt}}}_{\text{Life-cycle effects contribution}} + \underbrace{\frac{\sum_b Y_{bt}(\hat{\beta}_s)}{\sum_b Y_{bt}}}_{\text{Cohort effect contribution}} + \hat{\tau} \quad (8)$$

where Y_{bt} is the aggregate disposable income of each birth cohort, $\hat{\alpha}_s$ and $\hat{\beta}_s$ are the respective estimated age and cohort effects on saving as reported in Section 4, and $\hat{\tau}$ is the estimated constant. Growth in the aggregate income of successive cohorts, owing to population and/or economic growth, increases the weighting of younger cohorts in the aggregate, and thereby the relative size of younger cohorts’ contribution to the aggregate saving rate. To illustrate, in a “stripped-down” model where saving occurs pre-retirement and accumulated wealth is spent in retirement (with no cohort effects), economic or population growth leads to an increase in the aggregate saving ratio as the total saving of the young exceeds the dissaving of the elderly.

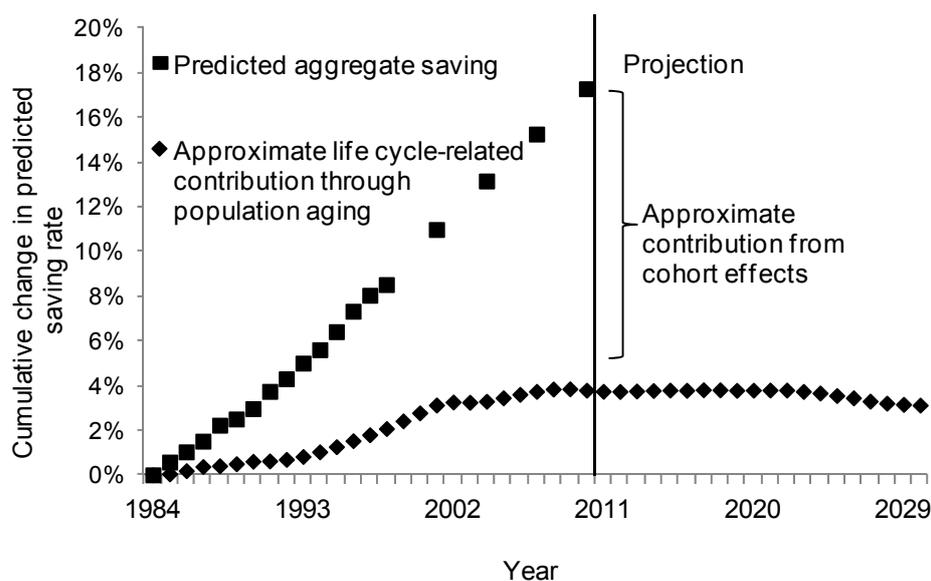
This framework can be used to approximate how changes in New Zealand’s population structure, and the ageing of the baby-boomers in particular, has affected and might affect the aggregate saving rate through life-cycle effects. Clearly, the precision of such approximations is limited to the extent that some population groups, especially the elderly, are not captured in the sample. Figure 15 shows the distribution of households in the HES survey by age of household head in 1984 and 2010. The figure shows the shift between surveys of the baby boomer bulge, from the low-saving young age groups toward the high-saving middle-aged age groups.

Figure 15: HES survey frequency distribution by age



The contribution to the aggregate saving rate of changes in the population structure through the life-cycle channel is approximated following the approach of Dynan, Edelberg, and Palumbo (2009). The income-weighted average of the predicted saving rate of each household head age group in 1984 is multiplied by the change in the share of the population represented by that group in subsequent years. Figure 16 shows this contribution, with projections to 2030 based on Statistics New Zealand population projections (Statistics New Zealand, 2012), alongside the predicted trend in the aggregate saving rate over the sample period.²⁶ The figure shows changes in the population structure through the life-cycle channel contributed approximately one quarter of the trend increase in aggregate saving between 1984 and 2010. The greatest increase in contribution occurred over the late 1990s as the baby boomers entered their high-saving fifties.

Figure 16: Contributions to predicted aggregate saving rate, 1984 to 2030



The total contribution to the aggregate saving ratio through the life-cycle channel combines the contribution from changes in population structure with the contribution from changes in income across cohorts. The latter contribution is almost zero because the low average rate of average income growth across birth year cohorts makes little difference to the distribution of the cohort income weights over time.²⁷ As a result, the difference between the aggregate trend and the life-cycle contribution shown in Figure 16 can be almost entirely attributed to cohort effects.

Statistics New Zealand's population projections to 2030 suggest changes in the population structure are unlikely to provide an additional future boost to aggregate saving through the life-cycle channel, assuming the estimated age profile of saving is unchanged. In fact they are likely to put downward pressure on the aggregate rate from the 2020s as the baby boomers become increasingly represented among the lower-saving elderly. This downward pressure may be greater than shown in Figure 16 to the extent that the oldest age groups are excluded from the sample. On the other hand, the future contribution from cohort effects is likely to be positive throughout the projection period assuming the identified pattern of cohort effects persist and future cohorts (not captured in the sample) save at comparable rates to the youngest cohorts in the sample. This positive contribution

²⁶ The life-cycle-related calculations in Figure 16 are made using Statistics New Zealand's population-by-age estimates, not the HES survey data shown in Figure 15, because the former includes future projections. There is little difference in the results of aggregate calculation using the two alternative data sets.

²⁷ The estimated cohort effects indicate a compound average growth rate of approximately ¾ per cent per birth year between 1910 and 1990.

will be underpinned by the fact that both the cohort effects and projected population sizes of cohorts reaching retirement generally increase through to 2030.²⁸

7 Conclusion

This paper used household-level data from HES between 1984 and 2010 to characterise the life-cycle saving behaviour of different generations of households. The results can be summarised in two main findings. First, household saving over the life cycle exhibits a hump shape, as predicted by the basic life-cycle model, with a peak when household heads are aged in their mid-to-late fifties. The life-cycle pattern of saving is associated with a more pronounced hump-shaped age profile of disposable income than of consumption expenditure. Second, there are significant differences between the estimated average saving rates of different cohorts over the sample age range of 19 to 74 years old. In particular, after accounting for age and one-off time effects, there is a near-linear trend increase in the average saving rates of cohorts born between the early 1930s and those born in the 1980s. As a result, from the baby boomers onward, the saving rates of each generation exceed those of the generation preceding it. This trend increase in saving rates reflects ongoing rises in disposable income accompanied by more moderate-to-flat growth in consumption expenditure across cohorts. One plausible explanation for the rise in saving rates, which is supported by the findings of other research, is that it reflects responses to an “unfavourable” evolution in the general economic and policy environments faced by successive cohorts.

These two findings are robust to various sensitivity tests including the use of alternative measures of saving; the introduction of conditioning variables to account for differences in household characteristics; and relaxing the assumption of constant age, cohort and time effects. The findings are consistent with the results of the only previous similar work in New Zealand (Gibson and Scobie, 2001 and 2001a, and Scobie and Gibson, 2003), which used the same HES data set for a shorter time period, 1984 to 1998. However, while no better household-level saving data exists, the potential influence of measurement error in HES presents an important caveat to the analysis. This error is evident in the divergence in trends between the aggregate saving rate based on HES data and the corresponding national accounts measure, and it could bias the estimated effects of age and cohort on saving rates. Ongoing work by Statistics New Zealand may provide more information about the nature of these potential biases.

Although the potential influence of measurement error cautions against making overly strong inferences, the estimates of age and cohort effects may provide some insight into the underlying influence of demography on national household saving trends. In particular, an increase in the proportion of the population in high-saving age groups contributed approximately one quarter of the overall trend increase in the aggregate HES saving rate between 1984 and 2010. However, population projections suggest that future positive contributions from this source are unlikely with the ongoing ageing of baby boomers in retirement expected to weigh upon the aggregate saving rate from the 2020s onward. The remaining three quarters of the aggregate trend increase between 1984 and 2010 is attributable to the rise in average saving rates of successive cohorts born since 1930. If the differences in average rates between cohorts persist into the future, cohort effects are likely to continue to make a positive contribution to aggregate saving rates throughout the projection period ending 2030.

²⁸ As discussed in Section 5.3.3, it seems likely that differences between the saving behaviour of cohorts would reverse at some point in the retirement period, but the data limitations prevent a precise estimation of how this occurs. Nevertheless, it seems reasonable to assume that the “net contribution” from cohort effects will be positive for as long as both the population size and average saving rate of cohorts reaching retirement exceed those of the (older) cohort before them.

References

- Alan, Sule, Thomas Crossley, and Hamish Low (2012) "Saving on a rainy day, borrowing for a rainy day." *Institute for Fiscal Studies Working Paper*, 12:11.
- Attanasio, Orazio P (1998) "Cohort Analysis of Saving Behavior by US Households." *Journal of Human Resources*, pp. 575-609.
- Barrett, Garry, Thomas F Crossley, and Kevin Milligan (2010) "An international comparison of savings rates from microdata and national accounts." Presentation given at the National Bureau of Economic Research, Summer Institute 2010's Economics of Household Saving Workshop on Saturday 24 July, Cambridge Massachusetts.
- Bascand, Geoff, Jeff Cope, and Diane Ramsay (2006) "Selected issues in the measurement of New Zealand's saving(s)." Paper Presented at Reserve Bank of New Zealand Workshop on Saving 14 November 2006.
- Browning, Martin and Thomas F Crossley (2001) "The life-cycle model of consumption and saving." *Journal of Economic Perspectives*, pp. 3-22.
- Chamon, Marcos D and Eswar S Prasad (2010) "Why are saving rates of urban households in China rising?" *American Economic Journal: Macroeconomics*, pp. 93-130.
- Coleman, Andrew (2006) "The life-cycle model, savings and growth." Paper prepared for Reserve Bank workshop entitled *Housing, Savings, and the Household Balance Sheet*, Wellington, 14 November 2006.
- Deaton, Angus (1985) "Panel data from time series of cross-sections." *Journal of Econometrics*, 30:1, pp. 109-26.
- Deaton, Angus (1997) *The analysis of household surveys: a microeconomic approach to development policy*. Johns Hopkins University Press.
- Dynan, Karen E, Wendy Edelberg, and Michael G Palumbo (2009) "The effects of population aging on the relationship among aggregate consumption, saving, and income." *The American Economic Review*, 99:2, pp. 380-86.
- Fesseau, Maryse, Florence Wolff, and Maria Liviana Mattonetti (2013) "A Cross-country Comparison of Household Income, Consumption and Wealth between Micro Sources and National Accounts Aggregates." OECD Publishing.
- Gibson, John and Grant Scobie (2001) "Household Saving Behaviour in New Zealand: A Cohort Analysis." New Zealand Treasury Working Paper 01/18.
- Gorman, Emma, Grant M Scobie, and Yongjoon Paek (2013) "Measuring Saving Rates in New Zealand: An Update." New Zealand Treasury Working Paper 13/04.
- Jappelli, Tullio and Franco Modigliani (1998) "The age-saving profile and the life-cycle hypothesis." Centro Studi in Economia e Finanza Working Paper.

- Le, Trinh, John Gibson, and Steven Stillman (2012) "Wealth and saving in New Zealand: evidence from the longitudinal survey of family, income and employment." *New Zealand Economic Papers*, 46:2, pp. 93-118.
- Savings Working Group (2011) "Saving New Zealand: Reducing Vulnerabilities and Barriers to Growth and Prosperity." Savings Working Group Final Report to the Minister of Finance, January 2011.
- Scobie, Grant and John Gibson (2003) "Household saving behaviour in New Zealand: why do cohorts behave differently?" New Zealand Treasury Working Paper 03/32.
- Scobie, Grant and Katherine Henderson (2009) "Saving rates of New Zealanders: A net wealth approach." New Zealand Treasury Working Paper 09/04.
- Statistics New Zealand (2007) "Changes to the Household Economic Survey." Statistics New Zealand.
- Statistics New Zealand (2012) "National Population Projections: 2011(base)-2061." Statistics New Zealand.
- Statistics New Zealand (2014) "Complete cohort life tables 1876-2012." Statistics New Zealand.
- Talosaga, Talosaga and Mark Vink (2014) "The effect of public pension eligibility age on household saving: evidence from a New Zealand natural experiment." New Zealand Treasury Working Paper 14/21.

Appendix

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

Preferred Saving Measure

Table 1 provides the expenditure categories and codes relating to the saving measure outlined in Subsection 2.3. HES was redeveloped between 2003 and 2006, resulting in a break in the expenditure time series and a change to the classification of expenditures (Statistics New Zealand, 2007). Although the old and new classifications of expenditure (shown in Table 1) have been matched as closely as possible, some inconsistency between the two periods may remain. HES income measures were relatively unaffected by the redevelopment.

Appendix Table 1: HES expenditure categories used in preferred saving measure calculation

Category	Item reference number codes	
	1984- 2004 surveys	New Zealand Household Expenditure Classification 2007 & 2010 surveys
HES total expenditure	0001-7269	1-14
HES contribution to savings	6900-6909	13.2
HES mortgage principal payments	1210-1217	04.2.01.2
HES life and health insurance payments	6903	11.4.01, 11.4.04
HES purchases of property*	1100-1109, 1538	04.2.01.1
HES sale of property*	1110-1119	14.2.01

*Sale and purchases of property data were not collected in HES from 2006. For earlier surveys these data are excluded from the calculation of saving, as described in Subsection 2.3.

For the 2007 and 2010 HES surveys, the New Zealand Treasury's non-behavioural tax-benefit micro simulation model, *Taxwell*, is used to estimate net tax payment and disposable income measures based on the gross income and demographic information provided by HES. In addition to estimating tax payments, which are not collected in HES, *Taxwell* makes adjustments to some income and benefit data in HES to account for known measurement inaccuracies. A predecessor of *Taxwell* was used to make the same calculations for the earlier surveys.

Alternative saving measure – inflation adjustment

The inflation adjusted measure of saving discussed in Subsection 5.1 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 consumption. 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 5.1.

Appendix Table 2: Expenditure items classified as “investment” expenditure

Category	Item reference number codes	New Zealand Household Expenditure Classification
	1984- 2004 surveys	2007 & 2010 surveys
Health	5200 – 5299, 6000-6099	06
Education	6200-6299, 6702-6703	04.1.01.2.0.01 04.1.01.2.0.03 07.2.05.0.0.12 09.4.01.0.1 10
Durable goods	2100-2179, 2200-2339, 2400-2419, 2500-2519, 4200-4229	05.1.01, 05.1.02, 07.1, 09.1.01 11.5.01.0.3.03 13.1.02.0.0.01 13.1.03.0.0.01
Building-permit fees	1300	04.4.03.2.0.01
Office Equipment	5650-5669,	09.1.02, 05.1.01.0.0.02 05.1.01.0.0.08 05.1.01.0.0.97 09.5.04.0.0.21 - 09.5.04.0.0.26
Sales of durable and capital goods	7000-7269	14
Other capital goods	1100-1199, ,5506, 5507	04.2.01.1 09.3.04.1.0.06 09.3.04.1.0.07 11.5.01, 11.6.03 11.6.04.0.0.00 11.6.04.0.0.09 11.6.04.0.0.11