

TREASURY WORKING PAPER

01/18

Household Saving Behaviour in New Zealand: A Cohort Analysis

John K. Gibson and Grant M. Scobie

ABSTRACT

This paper seeks to improve our understanding of household saving behaviour. It is based on an analysis of unit record data from March years 1984 to 1998 taken from the Household Economic Survey (HES). There are limitations of the data set but it provides the only available estimates of income and expenditure, from which saving is estimated as a residual. The HES is a series of cross-sectional surveys rather than a true panel, so we construct synthetic cohorts rather than tracking individual households. We use a range of regression models to separate out the effect of age, birth-year cohort and year on saving rates. The typical age profile for savings is hump-shaped, peaks around age 57 and does not become negative at older ages. Such a profile appears to have shifted down for the cohorts born between 1920 and 1939 relative to the younger and older cohorts studied. This pattern of cohort effects is robust to the inclusion of conditioning variables and to the trimming from the sample of households with either negative or very large ratios of savings to consumption. Preliminary investigation supports the hypothesis that changes in the economic and policy environment help explain the different saving behaviour of different birth cohorts. Tentative results suggest that more favourable environments are associated with lower rates of lifetime saving, although more research is needed to confirm this finding.

JEL Classification: D19: Household Behaviour; E21: Consumption, Saving; J26: Retirement.

Key Words: Household saving, lifecycle, age, cohorts

John K Gibson, Department of Economics, University of Waikato, Private Bag 3105, Hamilton, New Zealand. Tel [64 7] 838 4466 extn 8160 ; <mailto:jkgibson@waikato.ac.nz>
Grant M Scobie, PCD, The Treasury, PO Box 3724, Wellington, New Zealand.
Tel [64-4] 471-5005; <mailto:grant.scobie@treasury.govt.nz>

Disclaimer: The views expressed in this Working Paper are those of the author(s) and do not necessarily reflect the views of the Treasury. The paper is presented not as policy, but with a view to inform and stimulate wider debate.

Contents

1	INTRODUCTION AND SUMMARY	4
1.1	Key Findings	5
2	HOUSEHOLD SAVING PATTERNS: A DESCRIPTION	7
3	THE COHORT APPROACH AND CONSTRUCTION OF THE DATA	10
4	UNRESTRICTED ESTIMATES OF AGE AND COHORT EFFECTS	13
4.1	Introduction	13
4.2	Method	14
4.3	Results	16
4.4	Results for Population Sub-groups.....	19
4.5	Sensitivity Analyses Using “Individual” Saving Rates	20
5	INTRODUCING THE YEAR EFFECTS	22
5.1	Methods	22
5.2	Results	23
5.3	Results for Population Sub-Groups	24
6	THE MODEL USING INDIVIDUAL HOUSEHOLD OBSERVATIONS	25
6.1	Results for the Basic Model.....	25
6.2	Adding Conditioning Variables	26
6.3	Sensitivity Analyses	30
7	USING DIFFERENT DEFINITIONS OF SAVING	31
8	EXPLORING THE COHORT PATTERNS	32
	REFERENCES	39

List of Tables - [click to see Tables](#)

Table 2.1:	Sample Size and Savings Rates by Survey Year
Table 2.2:	Frequency Distribution of Savings Rates
Table 2.3:	Savings by Household Disposable Income Decile
Table 2.4:	Household income, expenditure and saving by income and expenditure deciles
Table 2.5:	Sample Sizes, Median Ages and Savings Rates by Population Sub-groups
Table 3.1:	Cohort Definitions, Cell Size and Savings Rates
Table 3.2:	Mean and Median Saving Rates, Averages Over Overlapping Ages
Table 6.1:	Cohort Effects in Individual Savings Rates, Controlling for Age and Time Effects
Table 6.2:	Cohort Effects in Mean Savings Rate With Conditioning Variables
Table 6.3:	Cohort Effects in Mean Savings Rate With Household-level Conditioning Variables
Table 6.4:	Sensitivity Analysis for Cohort Effects With Conditioning Variables
Table 7.1:	Cohort Effects in Mean Saving Rates With Different Definitions of Consumption and Saving
Table 8.1	Lifetime Patterns for Different Cohorts
Appendix	Table A: Wald tests of joint statistical significance of age and cohort effects in the unrestricted regressions

List of Figures - [click to see Figures](#)

- Figure 1: Household Saving Rate (HHIO vs. HES Flow measure)
- Figure 2.1 Lorenz Curve for Savings
- Figure 2.2: Comparison of Cumulative Densities for Savings Rates
- Figure 3.1: Household Saving Rates by Five-Year Birth Cohort
- Figure 4.1: Age Effects in (log) Income and Consumption and in the Derived Saving Rate $[\ln(y)-\ln(c)]$
- Figure 4.2: Cohort Effects in (log) Income and Consumption and in the Derived Saving Rate $[\ln(y)-\ln(c)]$
- Figure 4.3: Age and Cohort Effects in Income, Consumption and Saving Rates $[\ln(y)-\ln(c)]$ for Male and Female-headed Households
- Figure 4.4: Age and Cohort Effects in Income, Consumption and Saving Rates $[\ln(y)-\ln(c)]$ by Ethnicity of Household Head
- Figure 4.5: Estimates of Age-Specific Income and Consumption (B_{at} 's) by Survey Year
- Figure 4.6: Age and Cohort Effects in Saving Rates $[\ln(y) - \ln(c)]$ for Individuals and Households
- Figure 5.1: Age, Cohort and Time Decomposition of Mean Savings Rate (Age \times Year Cells)
- Figure 5.2: Relationship Between Survey Year Effects in Mean Saving Rates and Various Macroeconomic Variables
- Figure 5.3: Age, Cohort and Time Decomposition of Median Savings Rate
- Figure 5.4: Age, Cohort and Time Decomposition at 25th and 75th Percentile Savings Rate
- Figure 5.5: Age, Cohort and Time Decomposition of Mean Savings Rate: Male versus Female-headed households
- Figure 5.6: Age, Cohort and Time Decomposition of Mean Savings Rate: Maori and Pacific Islander-headed households Compared with Other Households
- Figure 6.1: Cohort Effects With Different Sets of Controls
- Figure 6.2: Smoothed Savings Rate by Cohort
- Figure 6.3: Smoothed Savings Rate by Cohort by Quantile Regression
- Figure 6.4: Cohort Effects With Different Estimation Samples (Demographics, Employment, Tenure and Family Type Controls)
- Figure 7.1: Cohort Effects With Different Definitions of Saving (Demographics, Employment, Tenure and Family Type Controls)

1 INTRODUCTION AND SUMMARY*

In many countries, concern has been raised by apparent falls in household saving rates. Low levels of household saving are linked to falling national saving and growing current account deficits on the one hand, and concern for the adequacy of retirement income on the other. There is a tendency to argue that savings are “too low” and for governments to attempt to intervene with policies designed to increase saving. These include tax breaks to individuals and education campaigns designed to raise awareness of the importance of saving.¹

Typically saving is measured from the national accounts as the difference between income and consumption in the sense of a flow. In New Zealand these data show an apparent decline in household saving. We have some stock measure of implied saving found by taking the changes in aggregate household wealth². But until the results of the Household Saving Survey currently underway become available, we have no micro-level data on household wealth.

There are well-recognized problems with the definition and measurement of saving from aggregate data. But more crucial is the fact that aggregate saving could be observed to be quite low (or even negative) and this would not necessarily imply a “problem” for household retirement incomes. In an ageing population households in their working lives could be saving while those in their retirement years were drawing down their saving and on balance there may be little net saving observed, despite quite high levels of saving among the saving households.

This is not to suggest that changing demographics alone could explain the fall in the household saving rate in New Zealand. In fact, if anything, household saving rates should be higher at present as much of the baby boom bulge is passing through its peak saving years of the lifecycle. Rather it is to highlight the fact that aggregate data can conceal what is happening at the individual household level. In short, if we really want to understand what is happening to saving we have to examine the dynamics of household saving patterns. This is the primary objective of this paper.

Individual data allows us to focus on particular economic, social and demographic groups and track their saving behaviour. This study uses individual records from the Household Expenditure Survey (HES) for a 15-year period. This is the only micro data set containing detailed information on income and expenditure. The survey does not provide direct estimates of saving. These have to be derived from the income and expenditure data, and the reader is alerted to the limitations that accompany this use of the HES data.

* The authors thank Ivan Tuckwell for his help in preparing the data files, Dave Maré and Roger Hurnard for discussions on methodology, Dean Hyslop for suggestions to improve the paper, and participants in seminars at the University of Waikato and in the Treasury for useful comments.

¹ For a recent synthesis of key issues in household saving, see Attanasio and Banks (2001).

² See Claus and Scobie (2001).

1.1 Key Findings

In summary, acknowledging the limitations of the data, we have established the following results:

- A very large share of total household saving comes from a small number of high-income households.
- Some 30 percent of all households used in the analysis, report negative saving.
- Saving rates are typically lower in households with a female head, in part because female-headed households are disproportionately represented among the retirees and sole parents.
- There is considerable variation in saving rates across households even when other characteristics are common.
- Median rates of saving (or other quantile estimates) are typically more informative than average rates which can be influenced by outliers.
- Ethnicity does not account for differences in saving rates across households.
- Saving rates appear to have risen modestly over the period between 1984 and 1998.
- The distribution of household saving is slightly more unequal in 1998 than in 1984.
- By imposing some identifying restrictions on the data we are able to separate out the effect of age, cohort and year on saving rates. Consumption largely tracks income over the lifecycle but tends to peak earlier. This results in typical hump-shaped pattern of saving rates over the lifecycle.
- Saving rates rise throughout working life, peaking when household heads are between ages 55 and 60 and show a marked tendency to decline then until age 74 (which is the oldest age considered in the sample).
- Households with heads born between 1920 and 1934 had significantly lower lifetime saving rates than those born before or after this period. In particular those born after 1950 have significantly higher lifetime saving rates. Using different definitions of saving results in basically the same patterns.
- These patterns are robust when including controls for a range of household characteristics.
- This pattern of saving rates for different cohorts may reflect the different economic and social environment prevailing at critical points in the lifecycle.
- We assemble some preliminary evidence that the tax and benefit system has treated different cohorts rather differently, and suggest that this helps explain why some cohorts have lower saving rates than others. Their behaviour is seen as a response to the incentives or disincentives to save created by public policy interventions.

Ideally, we would like to be able to use the insights about saving behaviour at the household level to help explain the aggregate saving trends observed in the national income data. Unfortunately, because the trends in the HES show saving rates rising, in contrast to the drop observed in the aggregate data (Figure 1), we clearly cannot use the former to explain the latter. Eventually, we are confident that the insights into individual saving behaviour from the micro data will be useful in explaining aggregate trends. But until we have a reconciliation and can explain the divergent series, then this task remains in the category of unfinished business.³

³ A study aimed at that reconciliation is currently being undertaken by Statistics New Zealand.

The paper is structured as follows. Section 2 explores the basic patterns of household saving revealed in the HES data. The tabulations, while giving a useful first glimpse, do not necessarily reflect the underlying behaviour. For example, the median saving rate for households headed by females might be lower than the corresponding rate for male-headed households. However, as this comparison does not hold constant family type, dwelling tenure, ethnicity, or age (stage of the lifecycle), then arguably these factors could be confounding the comparison. Regression analyses presented in Section 6 are a more satisfactory way to explore the effect on saving of particular characteristics while holding constant other influences.

Section 3 introduces the concept of synthetic cohorts, which form the basis of our analysis. Their construction is necessitated by the fact that we do not have true longitudinal panel data which tracks specific individuals or households through time. Initial estimates of mean and median saving rates by cohort are presented in this section.

Section 4 moves on to estimate the age and cohort effects based on cell means or quantiles. No further structure is imposed at this stage. The section concludes with estimates based on individual income, consumption and saving, where these estimates for individuals are themselves derived from a simple model of the household.

Section 5 introduces controls for year effects. Macroeconomic shocks, which affect, say employment or earnings, may alter saving rates in that year. In order to allow for all three effects (age, cohort and year) we must impose some additional identifying structure as clearly these variables are linearly dependent. The structure we adopt follows Paxson (1996).

In Section 6 we attempt to remove some of the noise in the data by smoothing. This is done through replacing the individual age effects with a fifth-order polynomial in age, and the one-year birth cohort effects with 5-year birth intervals. We introduce a set of conditioning variables (family type, ethnicity, dwelling tenure, labour market status and gender) in order to more clearly separate out the effects of age and cohort. Cohort is found to be an important variable in explaining differences in saving rates across households. In other words, when a person was born explains part of the differences in lifetime saving rates, all other factors being held constant.

Throughout we have used a definition of saving that removes from current consumption items that provide a flow of services over an extended period; ie, we have sought to use an “economic” approach to defining savings. Such items as education expenses or motor vehicles are removed from current consumption. In Section 7 we examine the effect of different definitions on the results of the regression analyses. We find that the essential cohort differences in saving rates hold up under different definitions.

It is then logical to ask why birth cohort appears as such a significant determinant of saving rates. We posit the hypothesis in Section 8 that an individual’s saving rate is in part a reflection of the economic and policy environment prevailing over their working life, and in particular applying during their peak saving years between ages 50 and 60. Certainly Thompson (1991) has argued that the benefits of the welfare state over the past 50 years have been distributed more on the basis of birth year than any notion of need, justice or desert.

We argue that an individual would tend to have a lower lifetime saving rate if for example, their working life corresponded to a period of low unemployment, real earnings growth, generous welfare benefits and they held expectations of an assured

state pension. By comparing the lifetime environments of different cohorts we could test this hypothesis. We make some progress with this potentially complex task, by looking at some indicators of the environment facing different cohorts and finding at least a tentative association with the pattern of lifetime saving displayed by those born between 1920 and 1934 and those born before (1910-14) and after 1950. We conclude, perhaps unsurprisingly, that social welfare policies do seem to matter to the amount people are prepared to save. Further testing of this relationship awaits the development of long-term data series for at least the last one hundred years.

2 HOUSEHOLD SAVING PATTERNS: A DESCRIPTION

In this section we first describe the data source and the way we have used it to derive measures of saving. We then present some tabulations which provide a first glimpse of household saving patterns.

Before setting out on this course we want to be very explicit in acknowledging the limitations of the data. These are summarised by Statistics New Zealand in the following quote:

"For several reasons, care is required in making comparisons of expenditure with income from the Household Economic Survey, as the method of surveying income and expenditure does not provide for consistency at an individual respondent level...."

Consequently, comparisons of total expenditure against total income are not valid at the household level. It follows that any comparisons of average expenditure statistics against average income statistics for groups of households, to estimate savings, for example, could lead to spurious results".⁴

We have used the income and expenditure data from the HES to estimate saving. Our defence for this seeming disregard of the warning rests largely on the fact that there is no other source of micro-data for examining household saving behaviour. Moreover, analysts in other countries have used similar data sources, particularly the Family Expenditure Survey in the United Kingdom (Attanasio and Banks, 1998) and the Consumer Expenditure Survey in the United States (Attanasio, 1998). We have tried to eliminate some outliers, and we have a large sample, which might arguably compensate for the underlying deficiencies. But we accept that our results are only as good as the survey data from which they are derived. To the extent that the survey is not providing a proper picture of household saving, then it must be acknowledged that our results may not be meaningful.

To construct a measure of household saving we subtract Current Consumption

⁴ Household Economic Survey Background Notes 1996/97, p. 17.

Expenditure (CCEx) from Disposable Income (YD)⁵. To estimate CCEx, we take Total HH Expenditure (THHEx) and subtract from it the following categories with a view to developing an "economic" view of saving - ie expenditures that provide consumption benefits over more than one year (eg. consumer durables, educational fees) are excluded from CCEx. The following items are subtracted from total expenditures to derive our measure of consumption expenditure:

<i>Code</i>	<i>Description</i>
11	Net Capital Outlay and Related Expenses
121	Mortgage- Principal Repayments
1300	Building Permit fees
210 -217	Home Appliances
220-222	Large Household Equipment
230-233	Furniture
25	Floor Coverings
42	Road Vehicles
5506	Purchase of Racehorses
5507	Purchase of Livestock
565-566	Office Equipment
58	Recreational Vehicles
62	Educational and Tuitional Services
6702-6703	Kindergarten Playcentre and Pre-School fees
690	Contributions to Savings
71	Sales
72	Trade-ins

The data cover the years 1983-84 to 1997-98. We refer to these years by the latter year; ie, 1984 and 1998. A total sample of 50,624 households was available over these 15 years. We have removed all observations where household disposable income was reported as negative (some 330 households)⁶. Further we have truncated the sample to remove all observations where the age of the household head was reported as less than 19 or greater than 74 at the time of the survey. This left us with a sample of 46,269 households.

Our justification for restricting the sample to the age range 19-74 is that those less than 19 were not considered important for studying lifetime saving patterns, while amongst the elderly, the HES does not cover institutions so those elderly living in rest homes are not included. This means that in the upper age groups we have an incomplete sample based only on those living on their own or as part of another household and this group may not be representative of the full population of the elderly.

Rather than taking the ratio of saving to disposable income, we have chosen to follow Attanasio (1998) and calculate saving rates by the ratio of saving to consumption. This has the advantage of being defined even when reported disposable income is zero⁷.

⁵ Disposable income is defined as the sum of income from all sources including government benefits, less taxes.

⁶ We recognise that these households might include some self-employed unincorporated businesses, whose business expenditures result in negative reported incomes.

⁷ Denote saving by S, consumption expenditure by X and disposable income by YD. Then: $S/YD = (S/X) \cdot (X/YD)$; ie, the ratio of saving to income is a monotonic transform of the ratio of saving to consumption and the saving ratios reported in this study can be converted by multiplying by the average propensity to consume.

A decision was needed as to the unit of analysis. Should it be the individual or the household? There is no clear answer to this; both have advantages and drawbacks. The HES reports income for each individual in the household, and expenditure on a household-wide basis. This means that to compute saving, one needs to either:

- a. allocate expenditure to individuals and then subtract from reported incomes to find individual saving levels; or
- b. combine the incomes of individuals to a total household income and subtract reported expenditure.

We have chosen the second option, believing that many saving decisions are taken on a household basis, and considering that allocating expenditure to individuals would have created some spurious saving estimates, especially for those household members who are not participating in the labour force.

We have defined the age, gender, labour market status and ethnicity of the household based on the reported characteristics of the head. However we also report results based on household shares (eg., the share who are working, who are male, etc).

Table 2.1 reports the number of observations in each year together with the saving rates at the mean, median, 25th and 75th percentiles. The mean is clearly influenced by extreme values, so frequently it will be helpful to focus on the median saving rate. The final column reports the ratio of the averages of saving to consumption, as distinct from the average (or quantile) of the ratios.

The results in Table 2.2 show the enormous underlying variability in the data. Household savings vary between -\$1.56 million and +\$0.78m. Among any one group with the same rate of saving (eg 0.20 to 0.29) consumption expenditures vary from \$2,950 to \$125,474, and their absolute level of saving varies from \$879 to \$32,241.

How do saving rates vary with household income? Predictably, median savings rates rise with income. Indeed, over one-half of total household savings over our 15-year period are generated by those in the richest decile of households. Absolute savings are negative in the three lowest deciles. Of course not all households in these three deciles have negative saving rates, as shown by the values of the 75th percentile (see Table 2.3).

What has happened to the saving rates over the period of the survey? The observations in each income decile were regressed on time. The lower three deciles show a negative trend and the upper income deciles show an increasingly positive trend. In other words, with some increased inequality in the distribution of income, there appears to have been a parallel increase in the inequality of saving. The Lorenz curves in Figure 2.1 plotted for the start and end years of the sample period also show this increased inequality.⁸

The cumulative distributions of the saving rates for the same two years are plotted in Figure 2.2, showing that 30 percent of household have negative estimated saving rates, a further 30 percent have rates between 0.0 and 0.5, while the remaining 40 percent have rates over 0.5; ie, they report income at least 150 percent higher than their consumption spending. The increasing inequality in saving rates, amongst both

⁸ Note that these Lorenz curves are only plotted for households with positive saving rates.

the negative and positive savers, is also apparent from the “straightening” of the cumulative distribution in the latter survey years because this implies that the underlying density function is getting more spread out.

Is saving confined to richer households? It depends on how one defines “rich”. When ranking according to household disposable income, the poorest three deciles show dissaving on average, while the level of saving for the richest decile of households is over five-times the average (see Table 2.4).

To check if this pattern is robust, savings by expenditure decile can be examined because, on average, households that are well-off in terms of income are also well-off in terms of expenditures (the correlation across the deciles in Table 2.4 is 0.99). But when ranking according to expenditures, the level of savings is roughly constant across the poorest nine deciles (implying a falling savings *rate*), while the richest decile in terms of household expenditures shows negative saving on average.

One plausible explanation for these patterns is measurement error. If income and expenditure are independently measured, at least to some extent, households who overstate their incomes will also, on average, overstate their savings, while households who overstate their consumption expenditures will correspondingly understate their saving (Deaton, 1997 p.31). The top deciles of income may contain a large fraction of households with overstated incomes and will thus show the highest savings rates, with the opposite effect for the top deciles of consumption whose overstated expenditures give the appearance of negative saving⁹.

We next present saving rates for various population sub-groups (Table 2.5). Female-headed households have lower saving rates, in part because they are over-represented among the oldest category and as sole parents. Households not in the labour force have saving rates a third less than those in the labour force, but there are no apparent differences by ethnicity. The age pattern of saving rates conforms to the expected hump-shaped lifecycle pattern. Sole parent families have a markedly lower saving rate than all other household types.

However, these tabulations provide only limited information as they do not attempt to hold other factors constant. That is the task of the regression models in Section 6.

3 THE COHORT APPROACH AND CONSTRUCTION OF THE DATA

To study the lifecycle profiles of saving we would ideally have panel data, where the same people are tracked over time. However, the available panel surveys in New Zealand are restricted to cohorts of young people who were born in the 1970s, and so are unsuitable for studying lifecycle phenomena.¹⁰ But the availability of a time-series

⁹ An alternative explanation is that instead of measurement error, the data may be revealing the effects of “transitory income”. A transitory rise in income, not expected to be sustained in a permanent manner, would not necessarily induce a corresponding rise in consumption. Likewise a transitory fall in income, would not necessarily be accompanied by a cut in consumption. This raises the possibility that in terms of permanent income saving among the lower actual income deciles is understated and amongst the higher income levels is overstated.

¹⁰ The “snapshot” offered by a single cross-section is also unsuitable for observing life-cycle patterns because although a variety of ages are observed in a cross-section, they also represent different birth

of cross-sectional Household Economic Surveys allows us to construct *synthetic panels* following methods described by Deaton (1985).

The key idea with synthetic panels is to divide the sample into groups whose membership is assumed to be fixed over time. The average behaviour of these groups is then tracked over time and as long as the sample is continually representative of the population that has fixed composition, estimates from these synthetic panel data should be consistent with estimates from genuine panel data on individuals¹¹.

In the context of saving behaviour, the synthetic panel method requires that we form various cohorts defined by date of birth and then follow them across the successive Household Economic Surveys¹². Provided the population is not much affected by immigration and emigration, and provided that a particular cohort is not so old that its members are dying in significant numbers, each successive survey lets us track movements in the average behaviour of each cohort over time (Deaton, 1997). For example, we can potentially look at the average saving rate of people who are 30-years old in the 1985 survey and connect that to the average saving rate of those who are 31-years old in the 1986 survey because both averages refer to the cohort born in 1955. Not only may these averages have many of the properties of panel data, they may also avoid some of the problems. In particular:

- Cohort data are constructed from fresh samples each year, so problems of sample attrition should be less severe, and
- There may be less bias due to measurement error because we are typically working with a cohort average (or some other quantile), which should reduce the impact of idiosyncratic variability that is a feature of data on individuals.

However, there are at least three practical problems with the use of synthetic panels for studying saving behaviour in the Household Economic Survey. The first is that we do not have data on individual consumption (and hence saving) so we can only follow households, whose cohort is defined by the date of birth of the household head. Hence, we face problems of household dissolution and reformation, where, for example, older people go to live with their children, so that previously “old” households become “young” households in subsequent years. There is no practical way to deal with this problem, given the nature of the data at hand, but we do attempt some sensitivity analyses based on “individual” measures of saving.

The second problem is that the assumption that the membership of the cohort is fixed may sometimes be hard to maintain. For example, if mortality and wealth are negatively related, cohort averages will reflect the fact that the population from which the samples are drawn becomes progressively richer as the poorer individuals die younger (Attanasio and Banks, 1998). This second problem is related to the first,

cohorts. If there are strong cohort effects, a cross-section age profile may be very different from the age profile of any individual, as noted by Shorrocks (1975).

¹¹ We do not observe the behaviour of the currently middle aged when they were young, nor the currently old when they were middle aged. As Attanasio (1998) points out, when extrapolating the lifecycle profile of each cohort for the ages at which it is not observed, it is necessary to use the information on the behaviour of other cohorts to impose some structure on the data.

¹² Verbeek and Nijman (1992) note that by treating averages of cohorts as if they were from genuine panel data may result in inconsistent estimates if the unobservable individual fixed effects are correlated with the explanatory variables. However provided that the true means in each cohort exhibit sufficient time variation and the cohort sizes are sufficiently large (they suggest 100 to 200) then the bias arising from ignoring this errors-in-variables problem is likely to be quite small.

because rather than dying, the poorer elderly also may be absorbed into younger households. We attempt to deal with this problem by restricting the maximum age in our sample to 74 years, although the possibility that wealth-related mortality has begun earlier than this age cannot be discounted.

The third problem is that the overall sample size of the HES (approximately 3500 households per year) means that many of the cell averages would represent rather small samples if they are formed from the interaction of each birth-year with each survey year. These small cell sizes may impair the precision of any estimates formed using the synthetic panel techniques. We respond to this sample size problem by using two types of cohorts: single birth years, for each year between 1910 and 1979, and five-year birth intervals. The narrower, one-year, cohort definition will reduce the amount of heterogeneity within each cell but at the cost of increased imprecision. By using both types of cohorts, we may gain some sense of the robustness of our results.

Table 3.1 contains details of the five-year birth-interval cohorts, including the birth years, the ages observed and the average cell size. Some of the earliest and latest born cohorts are tracked across fewer of the survey years because otherwise the age of these household heads would fall outside the range 19-74 years during the 1984-98 period.¹³ The table also contains estimates of the average saving-to-consumption ratio for each cohort, the same ratios calculated at the median, 25th and 75th percentiles and the ratio of average savings to average consumption.¹⁴ Comparing these averages across cohorts, saving rates appear to fall, then rise and then fall again when moving from the earliest born to the latest born cohorts. However, it is clear that this pattern combines both age and birth cohort effects because the ages over which household heads are observed also vary when moving from one cohort to another.

One way to hold age constant so that any cohort effect can be observed is to focus on the ages where adjacent cohorts overlap. To do this, each cohort's 'age' is based on the median year of birth within the five-year birth interval.¹⁵ Because the cohorts are defined by a five-year interval and we have 15 years of data, each cohort potentially overlaps at ten ages with the next one. For example, Cohort 6 (born in 1935-39) is observed between (median) ages 47 and 61, while Cohort 7 (born in 1940-44) is observed between (median) ages 42 and 56, giving overlapping ages between 47 and 56.¹⁶

Table 3.2 contains estimates of the mean and median saving rate for each pair of adjacent cohorts, averaged over the ages in which the two cohorts overlap. For both the mean and the median, the first four rows of the table, corresponding to households whose heads are born between 1910 and 1934, show a negative cohort effect. Each later born cohort has a lower average saving rate than the earlier born cohort had at the same age. This pattern is reversed when moving from Cohort 5 (household heads born in 1930-34) through to Cohort 11 (born in 1960-64) as each later born cohort has a higher average saving rate than did the earlier born cohort at the same age. This

¹³ Specifically, Cohorts 1-3 and 12-14 with birth years 1910-24 and 1965-79.

¹⁴ Attanasio (1998) points out that it may be easier to work with the ratio of averages rather than the average of ratios when constructing counterfactual experiments that link the saving rates of cohorts to the aggregate saving rate.

¹⁵ For example, for Cohort 1, where household heads are born between 1910-14, we treat the year of birth as 1912.

¹⁶ However, the earliest and latest born cohorts are observed for fewer years and hence have fewer years of overlap.

preliminary view of the raw data suggests that there may well be an important cohort pattern on saving among New Zealand households. However, more formal methods are needed to see if these cohort effects persist and are statistically significant once a greater lifecycle age structure is imposed on the data, and allowance is made for other conditioning variables.

In addition to comparing the average saving rates across cohorts, we can also track the saving rate for each cohort across successive survey years. To do this, each cohort's 'age' is once again based on the median year of birth within the five-year birth interval. Figure 3.1 plots these saving rates against age for each cohort, with the mean saving rate in the top panel and the median saving rate in the bottom panel. To give an example of how these cohorts are tracked, for Cohort 6 who had a median age of 47 in 1984, the 1984 survey was used to calculate the average saving rate for all households whose head was born in 1935-39 and the result is plotted as the first point on the line marked "6" (with a median saving rate of 0.17). The rest of the line comes from the other surveys, tracking those households whose head was born in 1935-39 until they are last observed at (median) age 61 in the 1998 survey (with a median saving rate of 0.33).

The immediate impression from both the mean and median saving rates in Figure 3.1 is the substantial amount of noise in the estimated average saving rates. Because each point is a summary statistic for cells that themselves hold an average of 250 households, the great variability in saving behaviour across households is apparent. But even with the noise, there is somewhat of a "hump" shape in these graphs, with average saving rates being highest from the mid-40's until household heads reach their 60's. Any decline in saving rates after the peak saving years is more apparent at the median than the mean. The cohort effects can also be seen from the variation in saving rates for different cohorts at the same age (i.e., by taking a vertical section anywhere through Figure 3.1).

4 UNRESTRICTED ESTIMATES OF AGE AND COHORT EFFECTS

4.1 Introduction

According to the lifecycle model, a person saves at one stage of his or her life to consume in another period. Therefore, saving behaviour should differ for different individuals at different stages of their lifecycles and may also evolve over time and across (birth-year) cohorts as economies grow and as certain fluctuations affect individuals contemporaneously. However, as noted in Section 1, without the aid of some identifying structure, it is impossible to distinguish the separate effects of age, cohort and time when using data observed over different age intervals for different birth cohorts. Of course, a possible limitation that arises from this need for an identifying structure, is that the resulting estimates simply reflect the assumptions chosen.

In this section we use a relatively unrestricted approach to estimating age and cohort effects, based on methods outlined by Paxson (1996) and Deaton (1997). Because there is no attempt to estimate time effects (i.e., year of observation), we do not need to impose a full identifying structure. Hence, the results in this section can serve as a cross-check for those derived from the more structured approaches used in the following sections.

Three other features of the approach used in this section are notable. First, it is based on cell means (or quantiles) rather than individual observations,¹⁷ so this averaging may reduce the impact of outliers and measurement error. Second, the approach allows us to examine age and cohort effects in saving rates jointly with the age and cohort effects in household income and consumption. This is useful because in our data saving is derived as a residual, so it may be informative to view patterns in the original data series, particularly given interest in shifting patterns of income and consumption in New Zealand (Borland, 2000). Third, we focus on households rather than on the individuals who are the decision-making units in theoretical treatments of saving behaviour. We also use the age and birth year of the household head to determine what cohort the household belongs to and at what point in its lifecycle it has been observed. To the best of our knowledge, this use of households rather than individuals is common to all but two studies of age and cohort effects in saving behaviour.¹⁸ Our approach can be justified by the fact that many consumption and saving choices are taken at household level (Attanasio, 1998) and it avoids the difficulty of imputing individual-level measures from household-level data. On the other hand, households may contain members who are at different stages in their lifecycles, so the age of the head may not be a good guide to the saving plans and behaviour of all members. Therefore, we do report some sensitivity analyses that are based on an individual version of an empirical lifecycle model, following the lead of Deaton and Paxson (2000).

4.2 Method

In the standard lifecycle model without uncertainty, an individual's consumption level is proportional to lifetime wealth – comprised of assets at birth and the discounted present value of expected future earnings – with a factor of proportionality that depends on age and the real interest rate. Ignoring the effect of interest rates, for individual i who is observed in year t and was born in year b ,

$$c_{ibt} = g_i(t - b)W_{ib} \quad (4.1)$$

where c is consumption and W is wealth, and noting that age, $a = t - b$. Adapting the model to households, lifetime resources are assumed to be set at the time of household formation and the task is to allocate consumption over time, according to the household's preferences as represented by the function g . If we take logarithms of equation (4.1) and then average over all households whose head is in the cohort born at time b and observed at t , we obtain

¹⁷ These cells are formed from the interaction of survey year and year of birth (of the household head).

¹⁸ Gokhale, Kotlikoff and Sabelhaus (1996) use cross-sectional profiles and population data to distribute aggregate data into cells defined by individual's age and sex. This required them to allocate total income earned by married couples evenly between the husband and the wife and to allocate household consumption amongst adults (after equalising children's consumption and reallocating it to the co-resident parents) using a combination of commodity-specific sharing rules and simple equal-per-adult allocations. The burgeoning literature on the lack of pooling of household resources (see, for example, Haddad, 1999) makes these assumptions appear questionable. Deaton and Paxson (2000) retrieve 'individual' consumption (income) by regressing household consumption (income) levels on the numbers of people in the household in each age from 0 to 99 and then apply the age and cohort decomposition described by equation (4.5). However, this approach assumes that people living together consume, earn and save in the same way that they would if living apart, which ignores issues such as economies of scale in household consumption.

$$\overline{\ln c_{bt}} = \overline{\ln g}(t - b) + \overline{\ln W}_b \quad (4.2)$$

so that average consumption is the sum of two components, one of which depends only on age and one of which depends only on birth-year cohort.

To estimate equation (4.2), the sample average of the logarithm of consumption for each cohort in each survey year can be regressed on a set of age and cohort dummy variables. However, Deaton (1997, p.346) suggests also including in the regression the average number of adults and children per household in each cohortxyear cell to make allowance for the differential consumption requirements of adults and children. These family composition effects can be thought of as an argument of the preference function g in equation (4.1). We follow this approach, but distinguish the number of prime age adults, n^a from the number of elderly adults, n^e (with the number of children denoted n^c).¹⁹ Therefore, the equation to be estimated is:

$$\overline{\ln c_{bt}} = D^a \mathbf{a}_c + D^b \mathbf{g}_c + \mathbf{d}_{c1} n^c + \mathbf{d}_{c2} n^a + \mathbf{d}_{c3} n^e + u_c \quad (4.3)$$

where D^a is a matrix of age dummies, D^b is a matrix of cohort (year of birth) dummies, the coefficients \mathbf{a}_c and \mathbf{g}_c are the age and cohort effects in consumption, the \mathbf{d} 's control for the effect of demographic composition and u_c is the error surrounding the sample estimate of average log consumption for households with heads born in year b and observed in year t . In addition to decomposing the mean of $\ln c_{bt}$ into age and cohort effects, the approach in equation (4.3) can be used for quantiles, and in our empirical application we apply the method to the cell median and the 25th and the 75th percentile of (log) household consumption.²⁰

In the same way that a time series of cross-sectional household surveys allows us to track the average consumption for different birth cohorts over time, so too does it allow us to track the average income of those cohorts over time. The underlying relationship, corresponding to equation (4.1), is that income at any age can be expressed proportional to lifetime resources, where the factor of proportionality depends on age (Deaton and Paxson, 2000). Hence, we can empirically estimate a counterpart to equation (4.3) but this time for log household incomes, where again they are attributed to cohorts based on the year of birth of the head:

$$\overline{\ln y_{bt}} = D^a \mathbf{a}_y + D^b \mathbf{g}_y + \mathbf{d}_{y1} n^c + \mathbf{d}_{y2} n^a + \mathbf{d}_{y3} n^e + u_y . \quad (4.4)$$

The difference between the logarithm of income and the logarithm of consumption is a monotone increasing function of both the saving-to-income and the saving-to-consumption ratios and obeys the inequality:²¹

$$s/y \leq \left[\overline{\ln y} - \overline{\ln c} \right] \leq s/x .$$

¹⁹ The transition from child to adult is assumed to occur at age 15, and from adult to elderly at age 65.

²⁰ All cell averages and quantiles formed for those households with a head born in b and observed in t are weighted by the household sampling weights recorded in the HES datafiles.

²¹ As the saving rate rises, the difference between log income and log consumption gives a closer approximation to the savings-to-income rate than to the savings-to-consumption ratio.

When the savings ratio is low, the difference between log income and log consumption is approximately equal to (either) savings ratio, so the cohort-age decomposition of log consumption and log income automatically gives a cohort-age decomposition of the savings ratio (Deaton, 1997). Subtracting equation (4.3) from (4.4) and ignoring the demographic variables gives:

$$s/x \approx \overline{\ln y} - \overline{\ln c} = D^a (\mathbf{a}_y - \mathbf{a}_c) + D^b (\mathbf{g}_y - \mathbf{g}_c) + u_y - u_c \quad (4.5)$$

where $(\mathbf{a}_y - \mathbf{a}_c)$ is the estimated age effect in the saving rate and $(\mathbf{g}_y - \mathbf{g}_c)$ is the estimated cohort effect.

This age-cohort decomposition of log consumption and log income is based on samples with 840 observations, where those observations are themselves cell averages (or quantiles). These cells are formed by the interaction of 15 survey years (1984-1998) and birth-year cohorts that ranged from 1910 to 1979. Note however that the cohorts at either end of the range do not appear in all 15 surveys because the sample is restricted to households whose head was between the ages of 19-74 at the time of the survey. The smallest number of households used to construct a cohort×year cell is 11 (for 19-year-olds in 1987), while the median cell size is 53.²² The March quarter CPI was used to convert all monetary values to December 1993 prices, prior to forming the cell averages of household incomes and consumption.

4.3 Results

Figure 4.1 shows the estimated age effects in consumption and income for the mean, median, and 25th and 75th percentiles of the distribution of households within cells. The average of current consumption expenditure appears to peak by the time household heads have reached age 40, while the age profile of average household income is everywhere higher and appears to peak somewhat later, although there is some variability in the graph.²³ Consequently, the saving rate (shown by the graph with a bold line and using the scale on the right-hand side), peaks in the decade after the household head reaches age 50 and then declines somewhat in the 60's but still remains well above zero. This apparent lack of dissaving amongst households headed by the elderly has been widely remarked upon in the literature, with most explanations focussing on bequest motives or differential mortality (Johnson and Stears, 1998; Attanasio and Hoynes, 2000), whereby the poor – who are likely to have low or negative saving rates – die earlier than the rich and so cease to be a downward drag on the average. When using household data, another form of “mortality” is also relevant; those households who have consumed all of their resources may be absorbed into other households (e.g., elderly parents moving in with children).

The age effects in median consumption, income and saving rates appear to be more favourable to the lifecycle hypothesis, with dissaving apparent at both young and old age. The ‘hump’ in savings for the decade beginning about age 50 is also more apparent, with saving rates rising by about ten percentage points, which is equivalent

²² Paxson (1996) carries out a similar analysis where the smallest cell size is only five households, although her median cell sizes are approximately 200.

²³ Each point on the income and consumption graphs corresponds to a coefficient from a regression. We have not applied any smoothing to these points and we do not show the standard error that surrounds each point. However, the age and cohort effects are both jointly statistically significant (see Appendix Table A).

to one-third of the amplitude shown over the lifecycle. It appears that as household heads advance past age 60 median household incomes fall more rapidly than do mean incomes and it is this, rather than differences in the age profiles of consumption at the mean and the median, which causes the differences in the age-saving profiles.

The bottom two panels of Figure 4.1 show the age profiles at the lower and upper end of the distribution. At any age, the income of household at the 25th percentile of the within-cell distribution is less than the current consumption expenditure of the household at the 25th percentile of the expenditure distribution (i.e., there is apparent dissaving), while the reverse pattern holds at the 75th percentile. This pattern may simply reflect the greater dispersion of household incomes compared to consumption. Regardless of the cause, it is apparent that age-related attrition amongst households at the lower end of the income and consumption distributions will tend to raise the mean saving rate. The 'hump' in savings is less apparent amongst the richer households, with a near-linear trend in the savings rate from age 25 to age 55. The drop in saving rates after age 60 is also less marked for these rich households.

The cohort effects are shown in Figure 4.2, where cohorts are relabelled according to their age in the first year of the survey (1984), so older cohorts are further to the right. There is considerable year-to-year variability in these cohort graphs, but the cohort effect in mean, median and 25th percentile consumption appears to be rising slightly with age in 1984. In other words, at any given age, the real consumption level of households whose head comes from an earlier birth year is somewhat higher than the consumption level of the more recent cohorts. In contrast, there is no overall trend in the cohort effect for the consumption of rich households, with both younger and older cohorts appearing to have higher consumption than the cohorts whose household head was middle-aged in 1984. These cohort effects in real consumption may be consistent with the reports (see, for example, Stephens *et al.*, 2000) that living standards for many households in the middle and bottom of the distribution may have been stagnant or even declining since the period of economic reforms beginning around 1984.²⁴

The cohort effects in mean and median household income also rise slightly with age at time of the first survey, so, for example, a household headed by someone in their 20's in 1984 has real income that appears about two percent lower than a household in its 20's had 40 years before.²⁵ But much more apparent than this slight trend is the 'dip' in average household incomes for those households headed by someone born between ca. 1930-1945. For example, a household headed by a 49-year old in 1984 has income that is about 12 percent lower than the income of a 49-year old household 11 years later, in 1995.²⁶ One possible explanation for this 'dip' in incomes is that the household heads aged ca. 40-55 were especially vulnerable to the effects of New Zealand's economic restructuring from 1984 onwards. Just as this group would expect to be in their peak earning years, they instead faced a decade of unprecedentedly high unemployment. By the time recovery came, many of these people may have been too

²⁴ This claim also relies upon the CPI being the correct measure for putting the consumption standard of living observed in different survey years onto a consistent basis. To the extent that there is unmeasured quality change, living standards of younger cohorts may in fact exceed those of older cohorts, but this is not a topic that we consider further. Moreover, there is a significant downward trend in household size across the survey years, so the trend in living standards for individuals may not be the same as the trend in living standards for households.

²⁵ A household whose head was in their 20's, 40 years earlier would be in its 60's at the time of the first survey.

²⁶ This is the gap from the income graph in the first panel of Figure 4.2 between ca. 9.3 on the log scale for a 38-year old in 1984 and 9.2 on the log scale for a 49-year old.

old to be re-employed, so they may have suffered a permanent fall in their lifetime incomes. But if this is the explanation for the cohort effects in average income, it is puzzling that the dip for those born ca. 1930-1945 is much more apparent for rich households than the poor (see the bottom panels of Figure 4.2). Most accounts of New Zealand's economic restructuring suggest that it has been the poor who have borne the brunt, so one would expect to see the fall in incomes for the 1930-1945 cohorts to be more apparent at the 25th percentile and less apparent at the 75th percentile²⁷.

One reason that the dip in household incomes for cohorts born ca. 1930-45 is so noticeable at the 75th percentile of the income distribution is that for these rich households, later-born cohorts have higher real incomes than other households. This pattern of the latter-born having greater lifetime resources is the 'usual' one in the presence of economic growth (see, for example, results for four countries reported by Paxson (1996)). In contrast, the cohort effect in income at the 25th percentile is the opposite of what economic growth would be expected to bring; age-income profiles appear to have declined sharply when moving from older to younger cohorts, with this downwards shift occurring at a faster rate than for the age-consumption profiles.

The combination of little change in average consumption across cohorts and a dip in the average incomes of certain birth-year cohorts produces a noticeable cohort effect in the saving rate²⁸. Both younger and older cohorts have higher average household saving rates than do those households whose head was between 40 and 60 years old in 1984. The fall in the saving rates for households with heads in this age range is about 10 percentage points at the mean and median and about 15 percentage points at the 75th percentile. The pattern of cohort effects at the 75th percentile is almost 'V' shaped, so it is possible that downward trends in aggregate saving rates might be temporary; older cohorts with higher saving rates have been replaced by middle-aged cohorts with low saving rates, but those middle-aged will, in turn, be replaced by younger cohorts with high saving rates.

In summary, this relatively unrestricted approach to decomposing income, consumption and saving rates into age and cohort effects has suggested the following: Consumption largely tracks income over the lifecycle but appears to peak earlier. Consequently there is somewhat of a 'hump' in saving rates for households whose head is between 50-60 years old. This 'hump' pattern (which, along with dissaving by the elderly is consistent with predictions of the lifecycle model) is especially apparent when using cell medians. There is a slight linear trend in consumption and income across cohorts, with the later born appearing to have lower lifetime wealth but the more noticeable effect is the decline in incomes for those cohorts aged ca. 40-55 in 1984. This produces a corresponding fall in saving rates for these cohorts.

²⁷ Future work could examine the effect that industry affiliation has on the cohort patterns found here. The reform of the NZ economy was uneven across industries in both its intensity and its sequencing, so we cannot rule out that possibility that the stronger cohort patterns found at the 75th percentile just reflect the attachment to certain industries by that group of households. Note, however, that the test of this hypothesis would require data on industry affiliation at the time of the reforms, whereas we can only obtain industry affiliation at the time of the survey".

²⁸ It is possible that the discrepancy between the timing of income and expenditures in the survey affects some types of households more than others, and that these households tend to be richer. One way to test whether our cohort patterns reflect this timing, and possibly seasonal influences, might be to control for the month that the household was observed (although the sample is staggered evenly throughout the year). Again, we leave that for future work.

The results reported below will test whether this cohort effect persists when more structured approaches are used to estimate year effects along with age and cohort effects. But before that, some disaggregated results for population sub-groups will be presented, because some explanations for falling aggregate saving rates revolve around changes in the population composition. We also present some sensitivity analyses to see whether the age and cohort patterns reported here are also present when the unit of analysis shifts from households to individuals.

4.4 Results for Population Sub-groups

The age and cohort effects in (log) income and consumption and in the derived saving rate for male-headed and female-headed households are plotted in Figure 4.3. These graphs are based on the decomposition of cell medians and so can be compared with the full-sample estimates in the upper right-hand panel of Figures 4.1 and 4.2. Female-headed households have lower saving rates over their lifecycle and they exhibit a much less pronounced rise than do male saving rates in the 40-60 years-old range. The age-profile of consumption appears to more closely track the age-profile of income for female-headed households whereas male-headed households appear to be able to detach their consumption from their income, at least between their early 40's and late 50's. The other interesting feature of the age profiles is the rise in the saving rate when a male head advances through his 60's and 70's; this rise is less apparent in the aggregate results. One reason may be the greater life expectancy of females, which causes a rising proportion of female-headed households amongst the elderly;²⁹ because these households have a lower saving rate this causes the overall saving rate for the elderly to fall.

The cohort graphs in the bottom panels of Figure 4.3 show that the lower saving rate for those households headed by someone who was ca. 45-55 in 1984 is a feature for both male- and female-headed households. There is also a strong cohort effect in income for households headed by women, but not for households headed by a man. For example, a household headed by a women in her 30's in 1984 has real income that is about 25 percent higher than a household headed by a women in her 30's had 20 years before (i.e., a household headed by a women in her 50's in 1984).

Figure 4.4 plots the age and cohort effects in income, consumption and saving rates according to the ethnicity of the household head, using a simple split of the sample into Maori or Pacific Islander and all other ethnic groups (predominantly European/Pakeha). For the Maori and Pacific Islander-headed households, the age profile of consumption appears to track the age profile of income so saving rates do not appear to exhibit hump-shaped patterns over the lifecycle. It is however difficult to observe the underlying patterns because of the high year-on-year variability and this problem of noise in the data is even more apparent in the graphs of the cohort effects. Hence, a more structured approach may be needed to studying the effect of ethnicity on saving rates (e.g., restricting the effects of ethnicity to the intercepts rather than the slopes of age-saving profiles).

The age and cohort decompositions can also be carried out with other population sub-groups, such as employed household heads compared with those who are unemployed or out of the labour force. However, these other variables that can be used to disaggregate the sample are likely to vary over the life cycle, which raises doubts about

²⁹ In the 55-64 year-old group, 65 percent of household heads are male, but in the 65-74 year-old group only 56 percent of household heads are male.

the appropriateness of the sampling assumption that is made when comparing people of different ages observed in different survey years. Possibly for this reason, statistical tests of whether age and cohort effects are (jointly) zero are often not rejected when the decomposition is applied to these sub-samples (Appendix Table A).

4.5 Sensitivity Analyses Using “Individual” Saving Rates

The available data force us to use households as the unit of analysis rather than the individuals who are the decision-makers in theoretical treatments of saving behaviour. The age of the household head is also used to determine the stage in the lifecycle for the household, even though there may be active savers within the household who are at different lifecycle stages. If there is some relationship between savings (or wealth) and being a household head, then working with the head’s age rather than the individual’s age may bias the estimated age profile of saving rates. For example, if low savers amongst the elderly move in with their children once they have exhausted their financial resources, while high savers continue as household heads, the age profile of saving will not decline with age in a way that reflects the average behaviour of the elderly (Deaton and Paxson, 2000).

There are enormous difficulties in retrieving estimates of individual income, consumption and savings from the household data that we have available. However, treating the results just as a sensitivity analysis, we can at least follow the example of Deaton and Paxson (2000) in estimating an individual version of an empirical lifecycle model. The key assumption of this approach is that households are simply a veil for the individuals within them, with no effect on anyone’s income, consumption, or savings. Household consumption is simply the sum of individual consumption, so the following regression can be estimated for a single cross-section from year t .

$$c_{ht} = \sum_{a=0}^N n_{ah} \mathbf{b}_{at} + v_{ht}, \quad t = 1, \dots, T \quad (4.6)$$

where c_{ht} is the observed consumption of household h in time t , and n_{ah} is the number of people aged a in household h at time t . More specifically, the β ’s are recovered from a cross-sectional regression of household consumption levels on the numbers of people in the household in each age from 0 to 99, and serve as estimates of “individual” consumption. A similar approach can be used to measure “individual” income. Although these are extreme assumptions, ignoring economies of scale and the division of labour between market and household production, they at least allow the ‘sharing rules’ for income and consumption to be derived from the data rather than from assumptions (compare with Gokhale *et. al.*, 1996).

Figure 4.5 plots the β ’s for each of the 15 survey years. In addition to the negative values below age 20 and above age 75, it is also apparent that this regression procedure produces substantial noise for the older ages, where the number of observations are small (or missing).³⁰

³⁰ With 15 survey years and 20 ages above 79, we should have 300 coefficients for the 80-99 year-olds in each regression but we only have 146 due to some survey years having no one of each particular age in the 80-99 range.

The individuals within households are assumed to each follow their own lifecycle trajectory, in which both consumption and income at each age is the product of an age effect and a lifecycle wealth effect. Hence, we can write a counterpart to equation (4.1):

$$\mathbf{b}_{at} = f(a)W_{t-a} \quad (4.7)$$

where $t-a$ gives the year of birth. These measures of individual consumption (and income) can be decomposed into age effects and cohort (wealth) effects. As Deaton and Paxson (2000, p.220) point out, provided that the β 's are positive, this decomposition can be performed by pooling the β 's estimated from all of the survey years, taking logarithms and regressing on a set of age and cohort dummies. This corresponds to using equations (4.3) and (4.4), which is the first and only stage when working with household-level data, because the β 's serve as an estimate of the average consumption (or income) for people of age a , observed in year t .

The major practical problem, however, is that the regression in equation (4.6) does not always result in positive values for the β 's, especially for those at very young and very old ages. Consequently, the decomposition using equations (4.3) and (4.4) is not possible because the logarithm of negative values is undefined.

Deaton and Paxson use an iterative matrix procedure to deal with the problem of negative estimates of individual income and consumption. They start with β_{at} estimates, similar to those reported in Figure 4.5, but with the income coefficients constrained to be zero for ages 80 and above. The negative values are set equal to one (so that logs can be taken) and the equation (4.4) decomposition is carried out to get initial estimates of age and cohort effects. These are then used in several iterations to produce a full set of cohorts effects for all ages observed by the survey. In principle, the vectors of cohort and age effects can then be estimated by minimizing the sum over all cohorts and ages of the squared residuals between the β 's and the product of the age and cohort effects. This is equivalent to extracting an eigenvector from the squared matrix of age and cohort effects. However, in their empirical application they are forced to impose the restriction that the cohort effects in income and consumption are identical, so no cohort effects in saving rates can be measured.

We use a simpler approach to ensuring non-negativity of the estimated β 's for individual consumption and income. This approach also has the advantage of allowing us to obtain cohort effects in individual saving rates. By using non-linear least squares to estimate equation (4.6), but with the β 's squared, we implicitly find the set of age-specific individual consumptions that minimise the residual sum of squares subject to the constraint that they be non-negative. For example, if we write a fragment of equation (4.6):

$$c_{ht} = \dots + \mathbf{b}_{19ht}n_{19ht} + \mathbf{b}_{20ht}n_{20ht} + \dots + v_{ht}$$

what we estimate is:

$$c_{ht} = \dots + (\mathbf{b}_{19ht} \cdot \mathbf{b}_{19ht})n_{19ht} + (\mathbf{b}_{20ht} \cdot \mathbf{b}_{20ht})n_{20ht} + \dots + v_{ht}$$

so that regardless of whether the estimates produced by the non-linear least squares algorithm, say, $\hat{\mathbf{b}}_{20ht}$, are positive or negative, they must be squared (and hence,

become positive) in order to restore the underlying linear regression in equation (4.6).³¹ We also impose some further structure, by requiring the β 's to be equal within the following age groups: 0-4, 5-9, 10-14, 15-19 and 80-99 years. This helps with the convergence of the non-linear least squares algorithm and should not affect our results, because we focus on the saving behaviour of individuals between ages 20-75.

Figure 4.6 plots the derived saving rates (more precisely, the difference between the logarithms of income and consumption) for individuals and households.³² The age profiles for the individual saving rates are considerably more hump-shaped than are the profiles using the household data and the age of the household head. This follows from the fact that both approaches should have the same mean saving rate (under the assumption that households have no effect on anyone's income or consumption) but the individual approach allows greater dissaving amongst the young and the old. Deaton and Paxson (2000) also find that the individual approach gives age profiles that are more in accord with the lifecycle hypothesis. Although the individual approach shows a greater rise in the saving rate, from approximately age 40 onwards, once the age effects are rescaled (the bottom panels of Figure 4.6) the age patterns are similar except in the retirement phase. The household-level saving rate rises as the household head's age rises from the early 60s to the mid 70's while the individual saving rate declines. This discrepancy may reflect the tendency of the less wealthy elderly to be absorbed into younger households (e.g., their children) where the economically active adults are building up their savings.

The cohort effects in the individual and household saving rates follow the same pattern as the age effects, in the sense that using the individual approach uncovers much greater variability. But once the cohort effects are rescaled, in the bottom panel of Figure 4.6, a similar group emerges as having the lowest saving rates. This group is comprised of those who were aged between ca. 45 and 60 years in 1984 (or equivalent, were born in ca. 1924-39). Hence, using the available methods for deriving individual saving rates, shows that the pattern of cohort effects discovered in the household data also is reproduced in the individual data.

5 INTRODUCING THE YEAR EFFECTS

5.1 Methods

In this section we introduce controls for time (i.e., survey year) that reflect the fact that saving behaviour may change over time as various macroeconomic fluctuations affect individuals contemporaneously. A straight-forward extension to the model would be to simply add a set of survey-year fixed effects, D^{a+b} ($a+b=t, t+1, \dots, T$, where t and T are the first and last available cross-sections):

$$s/x = D^a \mathbf{a} + D^b \mathbf{g} + D^{a+b} \mathbf{d} + u \quad (5.1)$$

where D^a and D^b are the age and birth year dummies as before. However, the superscripts highlight the fact that the year in which each household is observed equals the age of the household head, a plus their year of birth, b . The separate effect of a , b , and $a+b$ therefore cannot be identified in equation (5.1) and any trends in the data can

³¹ This approach to imposing inequality restrictions is noted by White (1993, p.107).

³² The household estimates are obtained directly from Figure 4.1.

be arbitrarily attributed to year effects, or a combination of age and cohort effects (Deaton and Paxson, 1994).

On the other hand, if the year effects are dropped from the model it rules out any uncertainty, such as due to macroeconomic shocks that surprise all members of a cohort. A less extreme assumption than dropping the year effects is to include them but in a normalised form so that they sum to zero and are orthogonal to a time trend:

$$\begin{aligned} \sum_{a+b=t}^T D^{a+b} &= 0 \\ \sum_{a+b=t}^T (a+b) \cdot D^{a+b} &= 0 \end{aligned} \tag{5.2}$$

This is equivalent to assuming that all trends in the data can be interpreted as a combination of age and cohort effects and are therefore, by definition, predictable. The time effects then reflect additive macroeconomic shocks or the residual influence of non-systematic measurement error (Jappelli, 1999).

We introduce these constrained year effects into an age-cohort decomposition of household saving rates that is based on the sample of 840 cell averages that were used in Section 4. Unlike the previous section, we examine the ratio of household saving to consumption directly rather than studying the age, cohort and time patterns in income and consumption. However, we continue our previous modelling methods, of dating the household by the age and birth year of the head and studying the patterns at cell means, medians and 25th and 75th percentiles of the saving rate.

5.2 Results

Figure 5.1 plots the age, cohort, and survey-year effects in mean saving rates. The panel at the bottom right-hand side of the graph shows that all three effects are statistically significant at the $p < 0.01$ level. Comparing Figure 5.1 with the saving rate graph from Figure 4.1 shows that the addition of the time effects does not alter the basic shape of the age-profile of saving rates but the cohort effect becomes more pronounced. Specifically, the fall in saving rates for those households whose head was aged ca. 40-60 in 1984 is more apparent once the time effects are included. These time effects are negative in 1984, 1990, and 1994-97 (noting that they are constrained to sum to zero across all years), which are not necessarily the years with the lowest average saving rates. In fact, there is only a correlation of 0.22 between mean saving rates and the time effects in saving rates that remain after age and cohort effects are removed.³³

What macroeconomic variables might explain these time effects, noting that they are meant to be capturing additive macroeconomic shocks? Figure 5.2 plots the relationship between the time effects and the real growth rate, the government saving rate, the unemployment rate, and the real interest rate. The trend lines are shown both for the full sample of 15 years and for the sample that excludes 1984, which looks to be an outlier on several of the graphs. The only two relationships that are statistically significant – both are negative – are with the real growth rate and with the ratio of

³³ This correlation rises to 0.44 if 1984 is excluded, but this is still statistically insignificant ($p < 0.12$).

One reason for excluding 1984 may be that incomes, and hence saving, was artificially suppressed by the wage and price freeze.

government saving to GDP. Hence, unpredictable shocks that either raise the real growth rate or raise the rate of public sector saving appear to reduce the household saving rate, holding constant any other macroeconomic influences working via the age or cohort effects.

These patterns may conflict with lifecycle explanations, because growth should raise mean saving rates by shifting resources toward higher savers, but this a long run effect and so may differ from the one plotted in Figure 5.1.³⁴ There is no significant relationship between the time effect in mean saving rates and either the unemployment rate,³⁵ or the real interest rate (nor with the nominal interest rate nor the inflation rate, neither of which relationship is plotted).

The age, cohort and time effects in the median saving rate are similar to the patterns in mean saving rates (Figure 5.3). However, there is less variation in saving rates over the lifecycle or across cohorts at the median compared with at the mean.

In Figure 5.4 it is apparent that there is even less variation in the age and cohort profiles of saving at the 25th percentile, and the hypothesis that the cohort effects are (jointly) zero cannot be rejected ($p < 0.24$). In contrast, the cohort effects are sharply defined at the 75th percentile, with those households headed by someone who was aged ca. 45-60 in 1984 having saving rates that were up to 70 percent lower than the saving rates that similarly aged households from the reference group had in earlier years (Figure 5.4). This contrast in the cohort patterns, which is consistent with the evidence in Figure 4.2, emphasises the diversity amongst households in the patterns of saving rates. But despite the dissimilarity of the cohort effects at the 25th and 75th percentiles, the time effects are largely the same, being negative in 1984-85, 1990, 1994-97.

5.3 Results for Population Sub-Groups

The age, cohort and time effects in saving rates for male-headed and female-headed households are plotted in Figure 5.5. The cohort effect is somewhat less pronounced for the female-headed households, but it is still statistically significant. There are also some differences in the time effects, which are negative for female-headed households in 1997.

Figure 5.6 plots the same decomposition for households headed by Maori or Pacific Islanders, compared with households headed by other ethnic groups. Although saving rates drift upwards as Maori and Pacific Island household heads age, they show no rapid rise from age 40 onwards, in contrast to the somewhat hump-shaped pattern for households headed by other ethnic groups. The cohort pattern for the ethnic minority groups is also less noticeable, with saving rates just trending downwards with the birth year of the household head, but it remains statistically significant. There is also a different pattern of time effects across the two groups of households; the year effects are not statistically significant for the Maori and Pacific Islands households, so it may be true that the saving behaviour of this group of households has been affected rather

³⁴ There is also a negative correlation (-0.12) between the real growth rate and the mean saving rate (without control for age or cohort effects), although it is not statistically significant.

³⁵ Attanasio and Banks (1998) find that saving rates and unemployment rates are positively related in the U.K. and suggest that it may reflect the impact of economic uncertainty on precautionary saving by households. For the mean saving rate (without control for age or cohort effects) the relationship with unemployment rates is positive (0.43) but statistically insignificant ($p < 0.12$).

differently by the economic reforms since 1984, although the saw-toothed pattern of year effects in Figure 5.6 is hard to interpret.

6 THE MODEL USING INDIVIDUAL HOUSEHOLD OBSERVATIONS

The results in the previous section show that even when further structure is added to the model, in the form of the (constrained) year effects, there is still considerable noise in both the estimated age profiles of saving rates and the birth-year cohort effects. Hence, to see the underlying patterns more clearly, even more structure may be needed. One common technique in the literature is to replace the age dummies with a fifth-order polynomial in the age of the household head; this appears to smooth out much of the noise while still being flexible enough to illustrate the shape of the underlying age profile. This approach can also be extended to the cohort dummies, replacing them with a fifth-order polynomial in year-of-birth.³⁶ An alternative approach to smoothing out the noise in the cohort effects is to use broader cohorts, such as five-year intervals of birth-year.³⁷

In this section we follow the approach of Attanasio (1998) and model the age profile with a fifth-order polynomial and the cohort effects with a set of five-year intervals of birth-year, which were defined in Table 3.1. To distinguish from the previous approach, we replace the index b for birth-year with c to indicate the five-year birth cohort. A further change is that in contrast to the methods used in the previous two sections, the unit of analysis is shifted to individual households, so that we can (eventually) study the effects of conditioning variables that may vary within the age/birth-year cells that we previously modelled. Hence, our basic model in this section is:

$$s_t^{ch} = f(a^h) + D^c \mathbf{g} + \mathbf{d}(d_t) + u_t^{ch} \quad (6.1)$$

where s_t^{ch} is the saving rate for household h , observed in year t and belonging to (five-year) birth-cohort c , the function f represents the fifth-order polynomial in the age of the household head (a^h), the coefficient vector \mathbf{g} captures the effects of the cohort intercept dummies, the function \mathbf{d} represents the time effects, d_t and u_t^{ch} is the residual. We continue the assumption made in the previous section, that all the (linear) trends observed in the data can be attributed to age and cohort effects, so that time effects are orthogonal to a linear trend and average to zero (see equation 5.2).³⁸ In keeping with our previous search for robust patterns in the data, equation 6.1 is estimated on mean saving rates and on three quantiles: the 25th, 50th, and 75th percentiles of the distribution of saving-to-expenditure ratios.

6.1 Results for the Basic Model

Table 6.1 reports the estimates from four regressions on the individual household saving rates. The separate intercepts for each five-year birth cohort are reported, along with the coefficients on the fifth-order polynomial in the age of the household head. The results for the survey year effects are not reported, but follow a similar pattern to that

³⁶ See for example, Jappelli, 1999.

³⁷ See for example, Attanasio, 1998.

³⁸ This follows from Deaton (1997), p.126.

displayed in Figure 5.1. In all cases we reject the hypotheses that either the year effects, the age effects, or the cohort effects are jointly zero.³⁹

The cohort effects in the mean saving rate are reported in the first column of Table 6.1 and follow a somewhat 'V' shaped pattern. Relative to the reference group, which is households headed by someone born in 1910-14, saving rates fall across later born cohorts, reaching their lowest point for households headed by someone from the 1930-34 birth cohort, where the saving rate is 14 percentage points below the reference group. Thereafter, the mean saving rate increases monotonically across the more recent cohorts, until it peaks amongst those households headed by someone from the 1970-74 birth cohort, where it is 28 percentage points above the reference group. This pattern is quite similar to that observed in Sections 4 and 5, when using the cell averages and single birth year cohorts. In those sections it was observed that younger and older cohorts had higher saving rates than did households whose head was between 40 and 60 years old in 1984. The current results allow a more precise dating of the dip in saving rates, which seems to especially affect those households headed by someone born in the 1925-34 period. This result carries the possible implication that downward trends in aggregate saving rates might be temporary, as middle-aged cohorts with low saving rates will eventually be replaced by younger cohorts with higher saving rates.

However, there are fewer grounds for such optimism when considering median saving rates, which exhibit the same 'dip' for cohorts born ca. 1925-39 but do not show any statistically significant rise in saving rates across the more recently born cohorts. It appears that the results for mean saving rates are being caused mainly by the behaviour of households in the upper end of the distribution; at the 25th percentile there are no significant cohort effects, whereas at the 75th percentile the 'V' shape is accentuated. Amongst these households with high saving rates, the results in the final column of Table 6.1 show that the saving rate for the 1930-34 birth cohort is 28 percent lower than for the 1910-14 cohort. This concentration of the cohort effects at the 75th percentile is also consistent with the results found from cell averages in Section 4. Although the cohort effect appears to be largely restricted to the upper end of the distribution, these are the households who contribute the bulk of aggregate saving, so it is important to adequately describe and understand these cohort effects if one is to make accurate predictions about the future path of aggregate savings.

6.2 Adding Conditioning Variables

To check whether within-cell heterogeneity can explain the cohort effects that we have previously found, the regression models are augmented with various conditioning variables, controlling for demographics, education, employment, family structure and dwelling tenure. For example, one possible cause of the cohort effects in Table 6.1 is that there are differences in family structure across birth year cohorts due, say, to the impact of changing social conditions and welfare policies on the prevalence of sole parenthood. By checking to see if the pattern of cohort effects changes when these conditioning variables are introduced, we may see if the shifts in lifecycle saving profiles can be explained by these demographic, education and family structure effects.

³⁹ It will be apparent from the low value of the R^2 statistics, that saving rates at the level of individual households are explained by much more than just age, cohort and year. However our objective is to examine age and cohort effects rather than explain the level of household saving.

Some of the conditioning variables, such as employment status, are likely to change over the lifecycle, whereas others, such as ethnicity and gender, are constant over the lifecycle. In both cases, the conditioning variables are allowed to shift the intercept of the estimated age profile of saving but because of the small sample sizes we do not consider interaction effects where the shape of the age profile can differ between, say, education groups. Therefore, the specification of the regression models is:

$$s_t^{ch} = f(a^h, c^h) + \mathbf{b}'w_t^{ch} + u_t^{ch} \quad (6.2)$$

where:

s_t^{ch} = the saving rate for household h , observed in year t and belonging to (five-year) birth-cohort c ;

w_t^{ch} = the conditioning variables; and

u_t^{ch} = the residual.

In addition to being a check on the robustness of the cohort effects found previously, the coefficients on the conditioning variables in the \mathbf{b} vector may also be of intrinsic interest. These coefficients can tell us, for example, whether there are significant differences in saving rates between male and female headed households of the same age and from the same birth cohort. We also check to see whether the results change when based on the characteristics of the household head compared with when they are based on the proportion of household members having the characteristic (e.g., we use, alternatively, as a conditioning variable, a dummy variable for whether the household head is of Maori or Pacific Island ethnicity, and a variable that records the proportion of household members who are of Maori or Pacific Island ethnicity).

Table 6.2 contains the results of several models that add increasingly broader sets of conditioning variables. All of the specifications include, along with the variables reported in the table, a fifth order polynomial in age and a set of survey year variables constrained to sum to zero and to be orthogonal to a linear trend. The model in the first column of Table 6.2 includes the gender and ethnicity of the household head, which can both be considered fixed over the lifecycle. There is a strong effect of gender, with male-headed households having saving rates approximately nine percentage points higher, but there is no apparent effect of ethnicity on savings. The addition of these demographic controls does not change the basic pattern of cohort intercepts, with households whose head was born ca. 1925-1939 having lower than average saving rates, while the rise in saving rates for the more recently born cohorts is more apparent than when the demographic controls were absent.

In column (ii) of Table 6.2, the model includes variables for whether the household head is either employed or unemployed, and another variable for whether the head receives self-employment income. In comparison with the reference category, which is households whose head is out of the labour force, average saving rates are seven percent lower if the head is unemployed and 17 percent higher if the head is working. The saving rate appears about 13 percent higher when the household head receives self-employment income. The difference in income levels between the self-employed and other households may be too small to explain this large jump in saving rates,⁴⁰ so it

⁴⁰ Using all 15 surveys from 1984-98, the average disposable income of households headed by someone receiving self-employment income was \$41,900 (in December 1993 prices), while the average for other households where the head is employed is \$39,600.

may be evidence for theoretical arguments that uninsurable income risk, which is likely to be greater for the self-employed, raises the level of wealth accumulation (Caballero, 1991).

The addition of these three employment variables reinforces the basic cohort pattern in saving rates that is reported in Table 6.1, with higher saving rates amongst the later born cohorts and lower saving rates amongst the households whose head was born ca. 1920-1939. Adding the employment variables also affects the results for the other demographic controls, halving the coefficient on gender and producing a significantly positive coefficient on ethnicity. Hence, the lower saving rates of female-headed households are partly because of their lower employment rates, while households headed by Maori and Pacific Islanders would have higher than average saving rates if their household heads had employment probabilities that were the same as the rest of the population.

Adding variables that measure family type and dwelling tenure to the model does not cause much change in the pattern of the cohort effects (column (iii), Table 6.2). However, the family type and tenure variables, which all have strong life cycle patterns,⁴¹ turn out to be significant determinants of saving rates. Sole parents and couples with children have saving rates that are at least ten percentage points lower than for the reference group comprised of single adults, couples without children and mixed and non-family households. Amongst the families with children, the saving rate of those households where there is just a sole-parent is four percentage points lower than the saving rate of couples with children ($p < 0.01$). Hence, the rise over time in the proportion of sole parent households might be expected to depress the aggregate saving rate.

In comparison with the reference group who live in rented dwellings, saving rates are higher for households paying a mortgage on their own dwelling and are especially higher if the mortgage has been repaid on the dwelling (and if it is rent-free, which is the tenure type for about three percent of households). The coefficients of the model in column (iii) imply a 25 percent higher saving rate for a household that has repaid its mortgage, compared with a household headed by someone of the same age and birth-cohort that is still paying a mortgage. This result underscores the extent to which New Zealanders view the acquisition of residential property as a form of retirement savings. The fact that the cohort effects do not change greatly once the home ownership variables are added also seems to count against any hypothesis that changes in aggregate saving rates have resulted from the greater capital gains made by some cohorts due to their ownership of real estate during periods of rapid inflation.

The literature on savings often mentions education as a further source of variation in saving rates for households headed by persons of the same age and birth-cohort (Attanasio, 1998; Jappelli, 1999). Two problems with studying the effect of education in the HES are that questions on educational qualifications were not asked in one survey year (1987) and in other years were treated as "not applicable" for three-quarters of the household heads aged 65 years and above ($n=4546$). Dropping observations where the qualifications variables were either not applicable, not specified, or missing reduces the sample by one-fifth ($n=37922$) and especially excludes households headed by older persons.

⁴¹ For example, using a fifth-order polynomial to control for age of the household head, a probit model shows that sole-parenthood is significantly more likely amongst household heads born since 1970.

This change in the estimation sample means that two regressions are required to see what impact the education variables have. The first model (in column (iv)) controls for the change in sample, while the second (in column (v)) adds variables for whether the household head has no qualifications, school qualifications or vocation qualifications, with those with university qualifications as the excluded category. The comparison of columns (v) and (iv) shows little change in the estimated cohort effects with the addition of the qualifications variables, which anyway do not attract especially large or statistically significant coefficients. Therefore, we ignore any further role for education in our analyses.

Figure 6.1 plots the cohort intercepts estimated at the mean, median, 25th and 75th percentiles of the distribution of saving rates, along with the intercepts from the models that include conditioning variables. It is evident that the introduction of controls for within-cell heterogeneity does not greatly modify the relative magnitude of the cohort dummies, tending to cause variation only for the most recently born cohorts. There is rather more variation in the patterns of cohort effects estimated at different points in the distribution, so we return to that point below, in considering whether the results are robust to increasingly severe trimming of outliers from the estimation sample. The other notable feature of Figure 6.1 is its similarity to results reported by Attanasio (1998, Figure 9); in the U.S. the household saving rate falls for the first four five-year birth-cohorts from 1910-14 to 1925-29, and then rises for each of the younger cohorts. With the exception of the later dating of the turning point in New Zealand (the 1930-34 cohort) and the inclusion of cohorts born post-1959, the patterns in the two countries are quite similar.

One concern with the controls for within-cell heterogeneity that have been introduced in Table 6.2 may be that household saving rates are allowed to vary within-cells only according to the characteristics of the household head. To see whether this restrictive way of controlling for heterogeneity matters, we replace these household head variables with variables that measure the prevalence of a characteristic at the household level. For example, instead of using a dummy variable for whether the household head is of Maori or Pacific Island ethnicity we use the proportion of household members who are of Maori or Pacific Island ethnicity. These alternative estimates, reported in the first three columns of Table 6.3, do not change the patterns amongst the estimated cohort effects. The patterns amongst the conditioning variables are also quite similar to those that were found when the characteristics of the household head were used.

A further check on the robustness of the cohort effects to different conditioning variables was made by adding a set of income source variables to the model (column (iv), Table 6.3). Unsurprisingly, saving rates were lower, the greater the proportion of household income from New Zealand Superannuation, private pensions and welfare benefits, while even reliance on wage and self-employment income is associated with lower saving rates than for the reference category of "other" income. But the most important result is that these income proportions do not alter the pattern of cohort effects found previously; the lowest saving rates are still found amongst those born in the 1930-34 period.

Given the results in Tables 6.2 and 6.3, the basic model of five-year birth-cohorts, augmented with demographic, employment, family type and tenure status variables, appears to give a reasonably robust description of the underlying data. We therefore use the predictions from the model in column (iii) of Table 6.2 to illustrate the shape of the age profile in mean household saving rates and to show how that profile has shifted

up and down across birth cohorts (Figure 6.2). The typical age profile for the average saving rate is somewhat hump shaped with a peak around age 57 but does not become negative at older ages. While the hump shape is consistent with the lifecycle hypothesis, the apparent increase in the saving rate beyond age 70 is not.⁴² The downward shift in the saving profile for earlier born cohorts, up until the fifth oldest one, is also apparent (earlier born cohorts are shown by the start and end points for their graph occurring at an older age, which is the age at the time of the first survey in 1984).

6.3 Sensitivity Analyses

In Figure 6.1 it is apparent that the pattern of cohort effects is more apparent at the mean and 75th percentile of the distribution of saving rates than it is at median and 25th percentile. The quantile regression at the median is based on least absolute deviations of the residuals, rather than least squares, and so is less sensitive to the presence of outliers. It is therefore worth investigating whether the pattern of cohort effects that has been found is just due to some of the extreme values of saving rates that were documented in Table 2.2.

Table 6.4 contains the results of the model estimated using either robust methods (quantile regression on median saving rates) or various trimmed samples that remove extreme values of saving rates. In all cases, the models include the demographic, employment, family type and tenure status variables. The results in column (i) show that the median saving rate falls from the earliest born cohorts until those born in 1930-34 and then rises across the cohorts born in later years. Hence, the pattern is the same as for the mean saving rates but the rise in saving rates for the most recently born cohorts is not as marked.

The predictions from this quantile regression model give the smoothed median saving rates in Figure 6.3. The pattern is similar to that for the mean saving rate, except that median saving rates are everywhere lower so that there is negative saving at the start of the lifecycle and around age 65, and the downward shifts in the profile when moving from later to earlier birth cohorts are smaller.

Removing the 24 households in the sample with $s/x \geq 10$ shifts all of the cohort intercepts down but does not alter the relative position of any of them, so that the same basic cohort pattern is apparent (Table 6.4, column (ii)). Removing households with these extremely high saving ratios also reverses the coefficient on one of the conditioning variables – the receipt of self-employment income. It appears that self-employment is more prevalent amongst households that report these extremely high saving rates – almost 40 percent of the households with $s/x \geq 10$ are headed by someone with self-employment income, as opposed to a self-employment rate of only 11 percent across the whole sample.

Further trimming of the sample, by removing those households with $s/x \geq 5$ causes a slight ‘flattening’ of the cohort pattern, with the fall in saving rates from the 1910-14 cohort to the 1930-34 cohort being about two percentage points smaller and the rise in saving rates through to the 1975-79 cohort being about 10 percentage points smaller (column (iii)). Taking a deeper cut of the sample, by removing households in the top

⁴² The use of a fifth-order polynomial does mean that the smoothed saving rate will exhibit four turning points but the predicted rise in the saving rate beyond age 70 does not appear to be a result of over-fitting the data. This same feature was evident in the unrestricted estimates of the mean saving rate in Section 4.

three percentiles, with saving rates of $s/x \geq 2$ causes a further flattening of the pattern, with reductions of about 10 percentage points in the cohort intercepts for the later born cohorts (column (iv)). The pattern of cohort effects in mean saving rates amongst households where $s/x \geq 2$ closely approximates the pattern in median saving rates estimated over the full sample. On the other hand, if the sample is trimmed by removing households with zero or negative saving rates, the basic 'V' shaped is accentuated, with saving rates amongst the 1930-34 cohort being 19 percent below those for the reference group and saving rates for the 1975-79 cohort being 50 percent above those for the reference group.

Figure 6.4 plots the cohort intercepts for each of the models reported in Table 6.4 and for the model of mean saving rates estimated on the full sample (Table 6.2, column (iii)). Although the magnitude of the cohort effects vary as the sample or estimation method is altered, the relative ranking of cohorts does not change. In all cases, saving rates fall from earlier to later born cohorts between the 1910-14 and 1930-34 cohorts and then the pattern reverses with later born cohorts exhibiting higher saving rates. These patterns seem sufficiently robust to warrant further investigation.

7 USING DIFFERENT DEFINITIONS OF SAVING

The definition of consumption that we have used when deriving the saving rate excludes items that are more properly considered as forms of investment and hence are a type of saving. In particular, to obtain the estimate of "current" consumption expenditure we removed from HES total expenditure, expenses on education, life and health insurance, purchases of durable goods, medical expenses, repayments of mortgage principle, and contributions to savings. Consequently, our consumption and saving variables differ from those that may have been available for previous studies and from the definitions used for national aggregates. In this section, we assess whether the cohort effects that we have found previously are sensitive to the re-inclusion of some of these items in the household consumption variable. Rather than repeating the entire analysis from Section 6, we concentrate only on the regression model for mean saving rates with the demographic, employment, family type and tenure status control variables.

Table 7.1 contains the estimates for four regressions on alternative definitions of the saving rate, each based on the addition of some categories of expenditure to the current consumption variable. The table also includes the estimated mean saving rate under each definition. For the purposes of comparison, the results using our preferred measure of saving from column (iii) of Table 6.2 are also reported.

If spending on education, health and insurance is considered as a form of consumption rather than investment, the mean saving rate falls by about six percentage points. The "V" shaped pattern of cohort effects is also flattened somewhat, with reductions of about 20 percent in the cohort intercepts for the later born cohorts. Treating the repayment of mortgage principle and contributions to household savings as a form of consumption has a similar effect on the mean saving rate and on the pattern of cohort effects. Neither of these amendments to the consumption and saving variables alter the finding that households headed by someone from the 1930-34 birth cohort have the lowest saving rates, conditional on age, survey year and the other control variables.

If spending on durable goods and vehicles is considered as a form of consumption rather than investment, the average saving rate is halved and the “V” shaped pattern of cohort effects is slightly accentuated, with the latter born cohorts having saving rates that are about five percent higher than under our preferred definition of saving. When all of these expenditure categories are included, and the saving rate is derived from total consumption rather than from current consumption, the cohort intercepts fall for all cohorts and especially for the more recent ones but the basic “V” shape is still apparent.

The cohort effects estimated under these different definitions of consumption and saving are plotted in Figure 7.1. These graphs illustrate the robustness of the relative cohort effects and especially the location of lower saving rates amongst those born ca. 1925-1939. The pattern of cohort effects in Figure 7.1 is consistent with our other sensitivity checks, including trimming the sample (Figure 6.4) and controlling for within-cell heterogeneity (Figure 6.1). Hence, we are confident that this lower saving rate for those born ca. 1925-1939 is a genuine feature of the saving behaviour of New Zealand households rather than just some artefact of the data or of our econometric procedures. Hence, the next task is to explain this cohort pattern of saving rates.

8 EXPLORING THE COHORT PATTERNS

Section 6 reported on the pattern of coefficients for the birth year cohorts. A range of these estimates for different quantiles and with different sets of controls was reported in Figure 6.1. They display a distinct V-shaped pattern. In short, saving rates appeared to differ significantly for different cohorts.

As noted, these cohort coefficients are estimated with reference to the 1910-14 birth cohort. In Table 6.2, column (iii) the coefficients for cohorts 3, 4 and 5 are all significantly lower than the reference cohort. In contrast, the coefficients for cohorts 9 through 14 are all significantly greater. In other words, those born from 1920 to the mid-1930s have demonstrably lower lifetime saving rates, while those born after 1950 have significantly higher rates of saving.

At first glance this result may seem surprising. Anecdotal evidence might have suggested that those born in the early inter-war period would have been conditioned by wars and the Great Depression, which could have led to higher saving rates, at the least that part of saving driven by a precautionary motive. In contrast, the post WWII cohorts facing greater economic growth and security, together with liberalised financial markets after 1986, might have been expected to have displayed greater profligacy, and have lower, not higher, rates of saving.

While these effects may be responsible for some influence on the estimated coefficients, clearly some other forces have operated to override them and produce a ‘V’ rather than an ‘inverted-V’ shaped pattern of saving by birth cohorts.

It is important to explore further the cohort pattern of saving. As cohorts with different patterns of saving move through their life cycle, they may influence the aggregate pattern of saving. The cohorts with significantly lower saving rates were aged between

45 and 60 in 1980, and entering their peak saving years. This is precisely the time that aggregate household saving was observed to start declining (Choy, 2000)⁴³.

Attanasio (1998) finds a remarkably similar 'V-shaped' pattern of cohort saving behaviour. He notes that the lower saving rates of the group aged in their 40s and 50s in the 1980s "are those mainly responsible for the decline in aggregate saving...because those cohorts were in the part of their lifecycle when saving are highest" (p.604). While he adds that even in the USA where the data are much more consistent, it is not possible to precisely match the aggregate and micro level data. Nevertheless, the estimated cohort effects "explain a substantial part of the decline in the aggregate saving rate".

Attanasio continues, noting that:

"The main deficiency of the analysis is its failure to explain why those particular cohorts did not save 'enough'. A plausible hypothesis, that I have not tested explicitly, is that the negative cohort effects for the middle cohorts are linked to increases in social security entitlements that the same cohorts have enjoyed" (p.604).

The cohort patterns found in the present study, as well as mirroring those found for the USA by Attanasio, are quite robust. As shown in Figure 6.1, with the exception of the 25th percentile, all the estimates share the inverted-V pattern. The pattern is particularly marked for the 75th percentile. This is a potentially important finding in terms of understanding the decline in aggregate savings. The 75th percentile represents higher income households, which account for a significant part of the total absolute savings.

Attanasio posited that more generous public pensions might have explained the different lifetime saving patterns of different cohorts. If over an individual's lifetime there is increasingly generous provision by the state for public pensions, it seems entirely plausible that this would shape expectations about the level of publicly subsidised pension that one might receive. Those expectations could then in turn influence an individual's decisions about the optimal allocation of consumption over their lifetime. Knowing (or at least predicting) that there would be a generous state pension to underpin consumption levels after retirement may lead to higher consumption prior to retirement from a given level of lifetime wealth. The consequence would be that the lifetime saving rate would be lower than in the absence of the public pension scheme, or with a less generous scheme.

If this were the case one would expect to see lower saving rates among those cohorts whose expectations were for the receipt of a more generous pension and higher rates among those who expected to receive lower real pension payments.

The impact of the social and economic environment over the lifetime of an individual (household) could be called the direct effect. There is also the indirect effect transmitted through family and those close to a person. These are the values and norms that are transmitted to them by their parents and others, and which in turn were formed by the conditions prevailing in an earlier time and shaping the values of their

⁴³ It should be recalled however, that because of conflicting trends in the estimates of saving rates between the HES and the SNA data, some caution is needed in using results from the HES (with an overall rising rate of saving) to "explain" the falls observed in the saving rates estimated from the national accounts.

parents. In that sense, the behaviour patterns observed at any one time are a function of the prevailing climate, the expectations of the future climate, together with the effects of all previous environments. Those in the most immediate past could be expected to have the greatest effect, with the impact tailing off the further back we go⁴⁴. We have not tried to explicitly allow for these indirect influences of past conditions, which might shape the behaviour of a particular cohort.

To proceed further with this hypothesis, it is necessary to posit some mechanism of how and when expectations are formed. Clearly this is a complex process, and one that would reflect the person's perception of their economic and social environment. In addition, the experiences of their parents and that of their childhood could well condition their perceptions and the need for saving. The environment prevailing during their working lives will affect their labour market experience and earnings (the ability to save) while the provision of social services (health, education, housing, pensions) and welfare (sickness, disability, unemployment, family and single mother benefits) will influence the need to save.

Other factors including capital gains on housing, real interest rates, rates of income growth, access to credit and life expectancy could all be expected to impinge on the decision to allocate income to current or future consumption. In addition, the desire to make bequests will also influence the saving rate.

In short there are a wide range of possible indicators that might affect the decisions of individuals with regard to their saving rate. Some of these will operate throughout their working lives. Given that the peak saving years are typically between ages 45 and 60, an individual's perceptions and expectations during this period, would arguably have a significant influence on their saving behaviour.

Essentially there are two steps in the argument: the first, that different cohorts operated in different environments; and the second that these different environments shaped the responses of different cohorts, particularly in the present case, with respect to their household saving behaviour.

Thomson (1996) documents changes that would support the first of these steps. He argues that over "the last 50 years welfare states have been very uneven in the benefits they provide for successive generations, that is for people born in different decades" (p.1). "The prizes and penalties of living in a welfare state are distributed more on the basis of birth date than of need, justice or desert. In New Zealand the big winners in this have been the 'welfare generation' – those born between about 1920 and 1945" (1991, p.1).

Testing these hypotheses is clearly a challenge on at least four counts. First we have little theoretical guidance as to how expectations are formed; in particular what relative weight should be assigned to each of the three critical periods:

- The experience of the previous generations particularly parents and grandparents which through norms and values could be expected to shape the saving behaviour of a particular cohort;
- The conditions prevailing during their working life and in particular applying during the peak saving years

⁴⁴ Counter examples can be found where some traditions are handed on virtually unchanged through many generations.

- Their expectations throughout their working lives about the level and eligibility for state support of retirement income.

Second, we need long time series, arguably covering the last 70 years to provide a quantitative assessment of the different policy environments enjoyed by different cohorts; third, we only have 14 observations of the “dependent” variable, ie, the cohort dummies that relate lifetime saving behaviour to year of birth⁴⁵, meaning the scope for any statistical tests is limited; and fourth, as the working lives and saving periods of the later cohorts extend into the future, some forecasts of future conditions will be involved in order to compare their behaviour with that of individuals from the older cohorts who are either retired or dead.

In what follows we “test” the hypotheses in a very elementary way. Basically we look at a series of indicators for which we can obtain at least partial data. Often we focus only on selected years or periods that might be “typical or representative”, or occur at that time of peak saving⁴⁶. In effect we are conducting in a loose manner a non-parametric sign test as an initial step. Are changes in the indicators broadly consistent with the hypothesis that the savings patterns of different birth cohorts could have been influenced by the economic and social policy environment prevailing at key points in the lifetimes? The objective is to make a preliminary foray to establish whether the patterns of some key variables that arguably affect the saving decisions of individuals are consistent with the hypothesis that the cohort differences reflect the external environment. In particular, we examine both labour market indicators (the ability to save) and some measures of public pensions and welfare (affecting the incentive to save).

Table 8.1 summarises the working life (assumed to be 40 years from age 20) and the peak earning years (assumed to be from ages 45 to 60) for each of the key birth year cohorts selected for this section. The first cohort (1910-14) is the reference cohort in the sense that the regression coefficients for cohort presented in Section 6 are referenced to this cohort, which has a value of zero. The lifetime saving rate of the adjoining cohort (1915-19) is not significantly different. The next three cohorts covering birth years from 1920 to 1934 are the group that typically have demonstrated significantly lower lifetime saving rates, while the last two (covering 1950 to 1959) are representative of those showing a significantly greater level of lifetime saving rate.

In what follows we examine some selected aspects of the economic and social environment facing the different cohorts both over their working lives as a whole, and in particular during their peak saving years. The question posed is the following: do those indicators vary in a way that is consistent with the observed cohort patterns in saving rates? We would expect to find that the proxies chosen for the economic and social environment adopted values less favourable for household saving rates during the critical years of the low saving cohorts, while the same indicator should be more favourable in years corresponding to the high saving cohorts.

⁴⁵ It is true that we could estimate the saving rate models with many more cohort dummies; in fact, potentially one for each birth year of all the individuals (or household heads) in the sample, as we have done in Sections 4 and 5. This would span some 80 years. However these estimates would tend to be noisy; making it difficult to estimate relationships with the policy variables which show much less year-to-year variance.

⁴⁶ Clearly this approach could be enhanced with more continuous time-series data from the 1930s to the present, but that is not an insignificant task, and one we assign to the category of “future research directions”.

Because of the magnitude of the task of assembling annual data on a wide range of variables in a consistent manner for 70 years, we have chosen to use selected years to illustrate the results. We focus on three cohorts: 1 (born 1910-14), 4 (born 1925-29) and 9 (born 1950-54), and will refer to these as the reference, early and late cohorts. Typically we will look at the values of the indicator variable prevailing during their peak saving years (given in Table 8.1).

We start with some key indicators relating to the labour market. The extent of unemployment is a critical factor affecting the expected flow of earnings. Those cohorts facing a lower probability of unemployment would be expected to have less incentive for precautionary saving. The unemployment rates (based on the average of the census years) facing the reference and early cohorts were 1.2 and 3.3 percent respectively, while based largely on projections the late cohort could face an average of 6 percent.

A further important labour market indicator is the rate of participation. The reference group had labour force participation rates of 94 and 35 percent for male and female workers. In contrast the early cohort faced 89 and 46 percent. The significant increase for women would be consistent with the lower saving rates of this group. The late cohort could face rates of 85 and 70 percent, suggesting a possible drop in the saving rates in future.

Those facing expectations of higher future income growth rates might be expected to have lower rates of saving. It is certainly true that real income growth rates (both GDP and household disposable incomes) were higher for the early cohort and lower for the later (higher saving) cohort⁴⁷.

Increasingly in New Zealand over the last century the state has assumed the role of saving for the household through the provision of subsidised (or free) education and health services, social insurance (sickness and unemployment benefits), family support (family allowances, capitalisation of family benefits, support to solo mothers) and generous levels of retirement income support.⁴⁸ It is to be expected that the incentive for voluntary saving by households would be reduced in the presence of these programmes; further, the more generous the programmes the greater the disincentive effect⁴⁹. As a result, we would expect that the lower savers faced more and the higher savers faced less generous state subsidies in health, education and welfare.

Much work remains to develop consistent series for health and education benefits. In the case of pensions, the pension:wage ratio was low to medium for the reference group, was markedly higher for the early cohorts (typically over 80 percent for a decade from the late 1970s), and lower for the late cohorts (expected to average 65 percent)⁵⁰.

⁴⁷ We note however that this contrasts with the view that growth raises saving by redistributing resources away from older dis-saving cohorts towards those who are younger and more active savers.

⁴⁸ It could be argued that in the absence of these state funded programmes private charities provided much of the social insurance, and by so doing had a similar effect on dampening the incentive to save as does state provision. This is undoubtedly true to some extent, but we would expect that most people would see the state system as more certain than relying on private charity, eligibility for which might have depended on certain behavioural patterns and religious proclivities seen as desirable by the providers.

⁴⁹ As an illustration, between 1900 and 1935 the total payments for civil pensions and family allowances rose from \$314,000 to \$4,109,658, which per head of European population corresponded to \$0.40 to \$2.75. By 1941, all pensions and social security had reached over \$15 per head (NZOYB, 1939 (p.518) and 1942 (p.506).

⁵⁰ See Preston (1999).

During the 1970s when the early cohort was at its peak saving period the real weekly pension for a married person rose from \$135 to over \$200 (in constant June 2000 terms). Over the next decade this fell and reached a low of \$163 in September 1996. By June 2000 it had recovered to \$174. In other words the early cohort faced the prospects of very substantial real increases in the state pension, a position that was not subsequently sustained. This is consistent with the early cohort making less retirement provision than the later cohort who face lower real pensions and greater uncertainty about their viability.

For many households, saving for housing is an important element of voluntary saving. If there was a significant probability of being allocated a state house at a subsidised rate with generous conditions for purchase it is likely that the incentive for voluntary saving would be commensurately diminished. In fact, lower savers faced a housing market in which the number of state houses being built was much greater than in later years⁵¹. Added to this, low savers who were paying off mortgages in the 1970s and early 1980s face negative real interest rates, further reducing the cost of housing and permitting consumption levels to be higher than they would have been had housing costs taken a greater share of disposable incomes.

In the past family sizes were larger. Having more children increases the probability that the parents will receive support in retirement for their children. Reduced family size in later years would be consistent with later cohorts being higher savers than the older cohorts with larger families.

As an increasing share of the costs of tertiary education has been shifted from the state to individuals, it is to be expected that younger cohorts would have more incentive to make provision for their children's educational costs. This tendency would be reinforced by the marked increase in tertiary participation rates.

Up to this point our consideration of the different saving behaviour of older and younger cohorts has focussed largely on the state provided benefits that they could have expected. Of course to meet these costs taxes had to be paid, so a full intergenerational accounting requires us to incorporate not only market earnings and state benefits, but taxation payments as well.⁵² This has not been attempted here. Thomson (1991) compares a prototypical early family (born 1930) with a late family (born 1955) and traces their lifetime earnings, taxes and benefits. He finds that the benefit:tax ratio for the so-called early family was 2.3 while the corresponding ratio for the late family was 0.6 to 0.8 (p.172). This result is strikingly consistent with the pattern of lifetime saving behaviour displayed by different birth cohorts in the present study.

In this section we have focussed on the lifetime saving patterns of different cohorts. People born at different times, all other things being held constant, demonstrate different lifetime saving rates. There are two reasons why this result is important. In the first place, it might help to explain the changes in aggregate saving behaviour. Unfortunately, in the case of New Zealand, we cannot carry this too far until we have a better understanding of how to reconcile the saving rates from the SNA and the HES. Our results here suggest younger cohorts have higher saving rates than their parents, but we are not yet able to link that to the observed decline in aggregate household saving rates.

⁵¹ For example, in 1949-50, over 20 percent of all new dwellings completed were state houses.

⁵² A complete accounting would require tracking of asset holdings and changes in asset values.

In the second place, if the different saving behaviour of different cohorts is due in part to the economic and social climate prevailing during their lifetimes, and in particular if different policies were operative for different cohorts, then we can at least start to better understand the impact of policies on household saving behaviour. This should contribute to better being able to predict the effect of policy changes in the future.

We stress that while the results of the cohort saving behaviour (Section 6) seem robust and significant, our attempt to provide an explanation is partial and tentative. This is a complex area and the saving rates we observe are the resolution of a set of forces encompassing social and cultural norms, economic conditions and a myriad of state interventions. Arguably, the provision of higher state benefits, or more certainty would be expected to dampen the incentive for private saving. Our preliminary examination of some snippets of evidence is at least consistent with that argument.

REFERENCES

- Attanasio, O. 1998. A cohort analysis of saving behavior by U.S. households. *Journal of Human Resources* 33(3): 575-609.
- Attanasio, O. and Banks, J. 1998. Trends in household saving: a tale of two countries. *Working Paper* No 98/15, The Institute of Fiscal Studies, London.
- Attanasio, O. and Banks, J. 2001. The assessment: household saving – issues in theory and policy. *Oxford Review of Economic Policy* 17(1):1-19.
- Attanasio, O., and Hoynes, H. 2000. Differential mortality and wealth accumulation. *Journal of Human Resources* 35(1): 1-29.
- Borland, J. 2000. Economic explanations of earnings distribution trends in the international literature and application to New Zealand. Treasury Working Paper 00/16, The Treasury, Wellington.
- Caballero, R., 1991. Earnings uncertainty and aggregate wealth accumulation. *American Economic Review* 76(3): 676-691.
- Claus, I. and Scobie, G.M. 2001. Household net worth: an international comparison. The Treasury, Wellington, Working Paper 01/19, September.
- Choy, W-K. 2000. Determinants of New Zealand National and Household Saving rates: A Cointegration Approach. Paper presented to the Conference of the New Zealand Association of Economists, Wellington, July.
- Deaton, A. 1985. Panel data from time series of cross sections. *Journal of Econometrics* 30(1): 109-126.
- Deaton, A. 1997. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy* Johns Hopkins University Press: Baltimore.
- Deaton, A., and Paxson, C. 2000. Growth and saving among individuals and households. *Review of Economics and Statistics* 82(2): 212-225.
- Gokhale, J., Kotlikoff, L., and Sabelhaus, J. 1996. Understanding the postwar decline in U.S. saving: a cohort analysis. Working Paper 5571, National Bureau of Economic Research.
- Haddad, L. 1999. The income earned by women: impacts on welfare outcomes. *Agricultural Economics* 20(1): 135-141.
- Jappelli, T. The age-wealth profile and the life-cycle hypothesis: a cohort analysis with a time series of cross-sections of Italian households. *Review of Income and Wealth* 45(1): 57-75.
- Johnson, P., and Stears, G. 1998. Why are older pensioners poorer? *Oxford Bulletin of Economics and Statistics* 60(3): 271-290.

- Paxson, C. 1996. Saving and growth: evidence from micro data. *European Economic Review* 40(2): 255-288.
- Preston, D. 1999 Retirement Income in New Zealand: the historical context. Wellington: Office of the Retirement Commissioner.
- Shorrocks, A. 1975. The age-wealth relationship: a cross-section and cohort analysis. *Review of Economics and Statistics* 57(1): 155-163.
- Stephens, R., Frater, P. and Waldegrave, C. 2000. Below the line: an analysis of income poverty in New Zealand, 1984 - 1998", GSBGM Discussion Paper 2/00, Victoria University of Wellington.
- Thomson, D. 1991 *Selfish Generations? The Ageing of New Zealand's Welfare State* (Wellington: Bridget Williams Books).
- Thomson, D. 1996 *Selfish Generations? How Welfare States Grow Old* (Wellington: Bridget Williams Books).
- Verbecek, M. and Th. Nijman 1992. Can Cohort Data be Treated as Genuine Panel Data? *Empirical Economics* 17:9-23.
- White, K. 1993. *Shazam User's Reference Manual Volume 7.0* McGraw-Hill, New York.

Table 2.1: Sample Size and Savings Rates by Survey Year

	Total Sample	Mean	Savings Rate (S/X)		75 th Percentile	$\frac{\bar{S}}{\bar{X}}$
			25 th Percentile	Median		
1984	3331	0.376	-0.058	0.227	0.613	0.189
1985	3295	0.287	-0.106	0.168	0.498	0.139
1986	3174	0.318	-0.078	0.201	0.551	0.177
1987	3210	0.341	-0.084	0.209	0.581	0.204
1988	4021	0.347	-0.043	0.212	0.563	0.200
1989	3142	0.358	-0.080	0.207	0.601	0.200
1990	3047	0.313	-0.110	0.188	0.560	0.166
1991	2674	0.340	-0.099	0.193	0.575	0.227
1992	2712	0.380	-0.062	0.217	0.609	0.251
1993	4244	0.415	-0.057	0.222	0.621	0.270
1994	2839	0.338	-0.096	0.197	0.546	0.235
1995	2695	0.336	-0.110	0.191	0.607	0.234
1996	2621	0.355	-0.111	0.186	0.572	0.246
1997	2642	0.359	-0.091	0.193	0.578	0.259
1998	2622	0.428	-0.081	0.238	0.676	0.320
Total	46269	0.353	-0.086	0.202	0.584	0.222

Note: Year is the year the survey ended, so 1984 is the survey year from April 1983 – March 1984.

Table 2.2: Frequency Distribution of Savings Rates

S/X	Sample Size	Mean Savings	Minimum Savings	Maximum Savings	Mean Expenditure	Minimum Expenditure	Maximum Expenditure
< 0.0	14755	-10238	-1561727	-1	33864	947	1566493
< 0.1	4149	1419	0	10562	28268	2311	127874
< 0.2	4151	4164	600	21859	27991	5258	130329
< 0.3	3724	6733	879	32241	27038	2950	125474
< 0.4	3294	9270	1217	69819	26546	3606	227809
< 0.5	2759	11505	1814	98661	25726	3882	237113
< 0.6	2341	13768	130	91261	25142	250	167609
< 0.7	1923	15740	3110	80006	24288	4694	127835
< 0.8	1557	17239	3190	94659	23092	4373	120173
< 0.9	1178	19068	3477	193140	22450	4173	220203
< 1.0	1013	20721	1496	103943	21851	1526	112267
< 2.0	4036	27465	3430	236500	20491	2870	129614
< 5.0	1232	50013	5809	668172	18239	1869	151723
< 10.0	133	103283	4835	780217	15809	945	122334
≥ 10	24	110485	14346	444324	8192	593	30593
Total	46269	6188	-1561727	780217	27862	250	1566493

Table 2.3: Savings by Household Disposable Income Decile

	Mean Income	Total Savings (\$m)	Trend in mean S/X	Savings Rate (S/X)				$\frac{\bar{S}}{\bar{X}}$
				Mean	25 th Percentile	Median	75 th Percentile	
1	8540	-8720	-0.005	-0.115	-0.539	-0.166	0.184	-0.404
2	14188	-2870	-0.005	0.088	-0.247	-0.004	0.292	-0.118
3	17698	-1150	-0.001	0.153	-0.171	0.056	0.339	-0.041
4	21807	96	-0.006	0.190	-0.150	0.091	0.397	0.003
5	26428	2753	0.002	0.275	-0.091	0.146	0.479	0.074
6	31400	5885	0.003	0.343	-0.031	0.224	0.547	0.142
7	36933	9547	0.002	0.424	0.035	0.287	0.629	0.208
8	43686	13790	0.005	0.493	0.083	0.348	0.692	0.266
9	53357	21280	0.007	0.614	0.164	0.448	0.851	0.361
10	86465	52390	0.036	1.069	0.294	0.669	1.299	0.675
ALL	34049	93001	0.004	0.353	-0.086	0.202	0.584	0.222

Note: Decile 1 is poorest and decile 10 is richest. All monetary figures are in December 1993 values, with the March Qtr CPI used as the deflator. The time trend in mean S/X comes from a linear regression on a time index of the savings rate for each decile in each year.

Table 2.4: Household income, expenditure and saving by income and expenditure deciles

Decile	By income decile			By expenditure decile		
	Income	Expenditure	Saving	Income	Expenditure	Saving
1	8540	14341	-5801	14792	8007	6785
2	14188	16094	-1906	19123	12584	6539
3	17698	18461	-764	22330	15933	6396
4	21807	21743	64	25873	19156	6716
5	26428	24597	1831	29891	22400	7491
6	31400	27485	3915	33842	25882	7960
7	36933	30582	6351	37448	29922	7526
8	43686	34509	9177	42491	34987	7504
9	53357	39202	14155	48761	42466	6295
10	86465	51607	34858	65951	67288	-1338
All	34049	27862	6187	34049	27862	6187

Note:

Figures are weighted averages over all households in each decile. The deciles are formed from the combined sample of households from the 1983-84 to 1997-98 Household Economic Surveys, with all values put in December 1993 prices (using the March quarter CPI for each year).

Expenditure is "Current Consumption Expenditure" which excludes expenses on medical, education, durables, mortgage principal, other capital expenditures, contributions to savings and life and health insurance payments. Income is "Household Disposable Income".

Table 2.5: Sample Sizes, Median Ages and Savings Rates by Population Sub-groups

	Total Sample	Median Age	Savings Rate (S/X)			$\frac{\bar{S}}{\bar{X}}$	
			Mean	25 th Percentile	Median		75 th Percentile
<i>Characteristics of the household head</i>							
Maori/Pacific	4680	37	0.325	-0.070	0.196	0.565	0.215
Other	41589	43	0.356	-0.087	0.203	0.586	0.223
Male	28531	43	0.389	-0.065	0.231	0.616	0.240
Female	17738	42	0.298	-0.111	0.160	0.527	0.189
Working	29716	39	0.395	-0.050	0.244	0.629	0.251
Job seeker	940	37	0.155	-0.176	0.059	0.375	0.089
Not in Lab. Force	15613	60	0.285	-0.133	0.133	0.505	0.140
<i>Age of household head</i>							
19-24	3260	22	0.217	-0.125	0.147	0.448	0.130
25-34	10910	30	0.261	-0.102	0.159	0.491	0.157
35-44	10898	39	0.302	-0.106	0.173	0.532	0.181
45-54	7934	49	0.472	-0.035	0.281	0.703	0.312
55-64	6959	60	0.465	-0.073	0.257	0.723	0.325
65-74	6308	69	0.396	-0.065	0.224	0.632	0.236
<i>Household type</i>							
Single adult	7831	57	0.395	-0.127	0.181	0.597	0.200
Couple, no children	11428	56	0.457	-0.026	0.288	0.717	0.320
Sole parent	4213	38	0.154	-0.181	0.058	0.347	0.071
Couple + children	17197	39	0.315	-0.086	0.191	0.535	0.198
Mixed family	3416	42	0.394	-0.031	0.257	0.648	0.248
Non-family	2184	26	0.265	-0.093	0.199	0.530	0.171
<i>Dwelling Tenure</i>							
Renting	11820	33	0.203	-0.134	0.121	0.432	0.126
Rent-free	1502	38	0.726	-0.041	0.341	1.001	0.383
Own, mortgaged	17070	38	0.271	-0.068	0.191	0.511	0.169
Own mortgage-free	15877	58	0.506	-0.066	0.279	0.775	0.346
Total	46269	43	0.353	-0.086	0.202	0.584	0.222

Table 3.1: Cohort Definitions, Cell Size and Savings Rates

Year of Birth	Ages Observed	Average Cell Size	Total Sample	Savings Rate (S/X)			\bar{S} \bar{X}	
				Mean	25 th Percentile	Median		75 th Percentile
1910-14	70-74	116	581	0.497	0.002	0.317	0.802	0.311
1915-19	65-74	174	1,743	0.431	-0.037	0.241	0.674	0.263
1920-24	60-74	201	3,009	0.426	-0.069	0.235	0.657	0.267
1925-29	55-73	235	3,518	0.396	-0.083	0.216	0.647	0.253
1930-34	50-68	224	3,361	0.405	-0.076	0.242	0.631	0.266
1935-39	45-63	222	3,324	0.449	-0.051	0.265	0.711	0.277
1940-44	40-58	270	4,047	0.474	-0.053	0.251	0.701	0.304
1945-49	35-53	330	4,955	0.383	-0.070	0.221	0.600	0.257
1950-54	30-48	375	5,624	0.322	-0.112	0.187	0.559	0.196
1955-59	25-43	382	5,726	0.277	-0.103	0.162	0.507	0.164
1960-64	20-38	362	5,437	0.261	-0.103	0.158	0.470	0.164
1965-69	19-33	206	3,087	0.254	-0.097	0.167	0.495	0.173
1970-74	19-28	155	1,553	0.260	-0.099	0.168	0.521	0.194
1975-79	19-23	61	304	0.156	-0.236	0.030	0.392	0.033
All cohorts	19-74	257	46269	0.353	-0.086	0.202	0.584	0.222

Table 3.2: Mean and Median Saving Rates, Averages Over Overlapping Ages

Cohorts	Ages of Overlap	Average of Means		Average of Medians	
1, 2	72-76	0.463,	0.412	0.293,	0.249
2, 3	67-76	0.428,	0.396	0.263,	0.222
3, 4	62-71	0.440,	0.340	0.238,	0.183
4, 5	57-66	0.428,	0.385	0.248,	0.204
5, 6	52-61	0.444,	0.502	0.261,	0.291
6, 7	47-56	0.428,	0.550	0.276,	0.288
7, 8	42-51	0.426,	0.440	0.238,	0.264
8, 9	37-46	0.324,	0.368	0.197,	0.231
9, 10	32-41	0.277,	0.301	0.164,	0.177
10, 11	27-36	0.250,	0.264	0.154,	0.153
11, 12	22-31	0.270,	0.266	0.176,	0.180
12, 13	17-26	0.189,	0.281	0.125,	0.171
13, 14	17-21	0.291,	0.127	0.162,	0.035

Note:

Cohort 1 is those households whose head is born 1910-14, while cohort 14 is those whose head is born 1975-79.

Table 6.1: Cohort Effects in Individual Savings Rates, Controlling for Age and Time Effects

	Mean	25 th Percentile	Median	75 th Percentile
Cohort 2 (b. 1915-19)	-0.045 (1.06)	-0.007 (0.25)	-0.036 (1.14)	-0.136 (3.05)**
Cohort 3 (b. 1920-24)	-0.063 (1.41)	-0.014 (0.47)	-0.051 (1.51)	-0.165 (3.42)**
Cohort 4 (b. 1925-29)	-0.122 (2.55)*	-0.030 (0.90)	-0.084 (2.30)*	-0.225 (4.31)**
Cohort 5 (b. 1930-34)	-0.142 (2.65)**	-0.046 (1.24)	-0.106 (2.62)**	-0.279 (4.84)**
Cohort 6 (b. 1935-39)	-0.066 (1.05)	-0.033 (0.83)	-0.081 (1.85)+	-0.190 (3.03)**
Cohort 7 (b. 1940-44)	0.054 (0.72)	-0.021 (0.48)	-0.051 (1.07)	-0.074 (1.10)
Cohort 8 (b. 1945-49)	0.106 (1.42)	0.002 (0.05)	-0.004 (0.08)	0.003 (0.04)
Cohort 9 (b. 1950-54)	0.168 (2.08)*	-0.017 (0.35)	0.020 (0.37)	0.102 (1.31)
Cohort 10 (b. 1955-59)	0.198 (2.32)*	0.003 (0.05)	0.029 (0.50)	0.121 (1.48)
Cohort 11 (b. 1960-64)	0.223 (2.48)*	0.011 (0.20)	0.040 (0.66)	0.122 (1.41)
Cohort 12 (b. 1965-69)	0.245 (2.61)**	0.019 (0.32)	0.054 (0.84)	0.178 (1.93)+
Cohort 13 (b. 1970-74)	0.281 (2.89)**	0.035 (0.55)	0.074 (1.08)	0.250 (2.56)*
Cohort 14 (b. 1975-79)	0.263 (2.39)*	-0.081 (1.12)	-0.013 (0.16)	0.211 (1.87)+
Age	1.332 (8.61)**	0.759 (7.32)**	0.998 (8.78)**	1.644 (10.04)**
Age ²	-0.067 (8.90)**	-0.038 (7.64)**	-0.051 (9.29)**	-0.083 (10.52)**
Age ³	0.002 (9.18)**	0.001 (7.95)**	0.001 (9.76)**	0.002 (10.93)**
Age ⁴	0.000 (9.35)**	0.000 (8.21)**	0.000 (10.09)**	0.000 (11.16)**
Age ⁵	0.000 (9.40)**	0.000 (8.41)**	0.000 (10.29)**	0.000 (11.21)**
Constant	-10.098 (8.18)**	-5.910 (7.10)**	-7.390 (8.11)**	-12.190 (9.26)**
R^2	0.0135	0.0042	0.0064	0.0144
Cohort effects = 0	$P < 0.000$	$P < 0.003$	$P < 0.000$	$P < 0.000$
Age effects = 0	$P < 0.000$	$P < 0.000$	$P < 0.000$	$P < 0.000$
Year effects = 0	$P < 0.001$	$P < 0.000$	$P < 0.005$	$P < 0.000$

Note:

Coefficients weighted by population sampling weights. Absolute value of robust t-statistics in parentheses; + significant at 10% level * significant at 5% level; ** significant at 1% level. The sample has $N=46269$ observations. Each regression also includes 13 time dummies, whose coefficients are constrained to sum up to zero and to be orthogonal to a linear trend.

Table 6.2: Cohort Effects in Mean Savings Rate With Conditioning Variables

	(i)	(ii)	(iii)	(iv)	(v)
Cohort 2	-0.046 (1.07)	-0.054 (1.28)	-0.058 (1.36)	-0.117 (1.26)	-0.118 (1.28)
Cohort 3	-0.061 (1.36)	-0.075 (1.69)+	-0.076 (1.70)+	-0.291 (2.36)*	-0.292 (2.37)*
Cohort 4	-0.117 (2.44)*	-0.130 (2.73)**	-0.140 (2.93)**	-0.433 (3.36)**	-0.433 (3.37)**
Cohort 5	-0.131 (2.44)*	-0.149 (2.80)**	-0.161 (3.02)**	-0.479 (3.52)**	-0.477 (3.50)**
Cohort 6	-0.048 (0.76)	-0.070 (1.12)	-0.082 (1.31)	-0.452 (3.07)**	-0.448 (3.04)**
Cohort 7	0.080 (1.06)	0.058 (0.77)	0.028 (0.37)	-0.389 (2.38)*	-0.383 (2.34)*
Cohort 8	0.139 (1.86)+	0.118 (1.58)	0.086 (1.16)	-0.376 (2.19)*	-0.369 (2.14)*
Cohort 9	0.212 (2.60)**	0.194 (2.37)*	0.155 (1.91)+	-0.342 (1.83)+	-0.336 (1.79)+
Cohort 10	0.254 (2.93)**	0.242 (2.78)**	0.204 (2.36)*	-0.338 (1.67)+	-0.330 (1.63)
Cohort 11	0.290 (3.16)**	0.289 (3.14)**	0.254 (2.78)**	-0.339 (1.55)	-0.331 (1.50)
Cohort 12	0.322 (3.38)**	0.321 (3.35)**	0.284 (2.99)**	-0.353 (1.51)	-0.343 (1.46)
Cohort 13	0.366 (3.70)**	0.378 (3.79)**	0.350 (3.54)**	-0.334 (1.33)	-0.323 (1.28)
Cohort 14	0.351 (3.13)**	0.382 (3.39)**	0.336 (2.99)**	-0.392 (1.43)	-0.376 (1.37)
Male head	0.093 (9.31)**	0.042 (3.75)**	0.037 (3.17)**	0.049 (3.72)**	0.051 (3.72)**
Maori/Pacific head	0.005 (0.38)	0.049 (3.95)**	0.095 (7.69)**	0.085 (6.58)**	0.080 (6.16)**
Working head		0.172 (10.62)**	0.159 (9.50)**	0.140 (7.41)**	0.143 (7.37)**
Unemployed head		-0.069 (2.49)*	-0.060 (2.23)*	-0.079 (2.78)**	-0.079 (2.79)**
Head has self-employment		0.126 (5.96)**	0.075 (3.59)**	0.059 (2.66)**	0.059 (2.64)**
Income			0.043 (4.54)**	0.056 (5.58)**	0.060 (5.96)**
Paying mortgage			0.297 (19.56)**	0.334 (18.91)**	0.336 (19.22)**
Rent/mortgage-free			-0.147 (9.76)**	-0.163 (10.20)**	-0.163 (10.26)**
Sole parent w/ children			-0.103 (8.70)**	-0.115 (8.97)**	-0.115 (8.93)**
Couple with children					-0.022 (1.24)
No school qualifications					-0.045 (2.24)*
School quals only					-0.075 (4.52)**
Vocational qualification ^a					-7.519 (4.83)**
Constant	-9.849 (8.01)**	-8.741 (6.94)**	-7.852 (6.36)**	-7.607 (4.97)**	
Sample size	46269	46269	46269	37922	37922
R-squared	0.016	0.024	0.040	0.044	0.044

Note:

Coefficients weighted by population sampling weights. Absolute value of robust t-statistics in parentheses; + significant at 10% level * significant at 5% level; ** significant at 1% level. Each regression includes a fifth-order polynomial in age of the household head and 13 time dummies, with coefficients constrained to sum up to zero and to be orthogonal to a linear trend. Excluded categories are female head, non-Maori and non-Pacific head, head not in the labour force, dwelling is rented, family type is single adult, couple without children, mixed and non-family, head has a university qualifications. For each regression, tests of the hypothesis that either the age, cohort or time effects are jointly zero are rejected at the $p < 0.001$ level.

Columns (iv) and (v) exclude observations from 1987 (data on qualifications unavailable), exclude most households where the head's age exceeded 65, because the qualifications question was not usually asked of this group, and excludes those for whom qualifications were not specified.

^aIncludes trade and "other" qualifications.

Table 6.3: Cohort Effects in Mean Savings Rate With Household-level Conditioning Variables

	(i)	(ii)	(iii)	(iv)
Cohort 2	-0.046 (1.09)	-0.053 (1.26)	-0.057 (1.34)	-0.052 (1.26)
Cohort 3	-0.061 (1.38)	-0.072 (1.63)	-0.075 (1.68)+	-0.046 (1.05)
Cohort 4	-0.124 (2.60)**	-0.134 (2.83)**	-0.146 (3.05)**	-0.121 (2.57)*
Cohort 5	-0.143 (2.68)**	-0.156 (2.95)**	-0.169 (3.18)**	-0.145 (2.78)**
Cohort 6	-0.066 (1.05)	-0.091 (1.47)	-0.101 (1.62)	-0.090 (1.50)
Cohort 7	0.054 (0.73)	0.023 (0.30)	0.004 (0.06)	0.036 (0.48)
Cohort 8	0.103 (1.39)	0.071 (0.96)	0.054 (0.73)	0.092 (1.26)
Cohort 9	0.169 (2.10)*	0.137 (1.69)+	0.118 (1.46)	0.150 (1.87)+
Cohort 10	0.200 (2.34)*	0.171 (1.99)*	0.155 (1.81)+	0.185 (2.18)*
Cohort 11	0.223 (2.48)*	0.195 (2.15)*	0.186 (2.07)*	0.221 (2.47)*
Cohort 12	0.244 (2.61)**	0.212 (2.24)*	0.206 (2.20)*	0.245 (2.62)**
Cohort 13	0.285 (2.94)**	0.274 (2.79)**	0.275 (2.82)**	0.319 (3.28)**
Cohort 14	0.268 (2.43)*	0.289 (2.60)**	0.271 (2.45)*	0.287 (2.59)**
Share of HH who are MALE	0.214 (9.09)**	0.164 (6.83)**	0.151 (6.26)**	0.137 (5.83)**
Share of HH who are Maori/Pacific	-0.004 (0.34)	0.083 (5.88)**	0.132 (9.35)**	0.160 (11.70)**
Share of HH who are working		0.304 (15.59)**	0.295 (12.64)**	0.198 (6.68)**
Share of HH who are unemployed		-0.144 (3.27)**	-0.123 (2.74)**	-0.140 (3.17)**
Share of HH with self-employment income		0.253 (6.01)**	0.164 (3.91)**	-0.007 (0.14)
Mortgage being paid on dwelling			0.037 (3.98)**	0.015 (1.54)
Dwelling is rent/mortgage-free			0.293 (19.48)**	0.244 (17.22)**
Sole parent with children			-0.086 (5.10)**	-0.074 (4.48)**
Couple with children			-0.022 (1.52)	-0.074 (4.71)**
<i>Share of household income from:</i>				
Social Welfare Benefits				-0.489 (9.09)**
Private Pensions				-0.550 (3.12)**
New Zealand Superannuation				-0.793 (12.53)**
Self-Employment				-0.083 (1.88)+
Wages and Salary				-0.293 (6.44)**
Constant	-10.029 (8.16)**	-6.275 (4.80)**	-6.064 (4.70)**	-3.822 (2.85)**
Sample size	46269	46269	46269	46227
R-squared	0.017	0.033	0.047	0.066

Note:

Coefficients weighted by population sampling weights. Absolute value of robust t-statistics in parentheses; + significant at 10% level * significant at 5% level; ** significant at 1% level. Each regression includes a fifth-order polynomial in age of the household head and 13 time dummies, with coefficients constrained to sum up to zero and to be orthogonal to a linear trend. For each regression, tests of the hypothesis that either the age, cohort or time effects are jointly zero are rejected at the $p < 0.001$ level.

Results in column (iv) exclude 42 observations with zero household disposable income and the excluded income category is “other”, which includes various forms of non-labour income.

Table 6.4: Sensitivity Analysis for Cohort Effects With Conditioning Variables

	Median	Mean (S/X) on trimmed samples			
	(S/X)	(S/X)<10	(S/X)<5	(S/X)<2	(S/X)≥0
Cohort 2	-0.054 (2.09)*	-0.060 (1.48)	-0.045 (1.26)	-0.046 (1.59)	-0.089 (1.83)+
Cohort 3	-0.073 (2.68)**	-0.076 (1.72)+	-0.080 (2.14)*	-0.066 (2.19)*	-0.091 (1.76)+
Cohort 4	-0.112 (3.74)**	-0.142 (3.06)**	-0.124 (3.05)**	-0.093 (2.85)**	-0.167 (2.98)**
Cohort 5	-0.139 (4.21)**	-0.165 (3.23)**	-0.146 (3.26)**	-0.120 (3.37)**	-0.191 (3.01)**
Cohort 6	-0.102 (2.84)**	-0.099 (1.75)+	-0.093 (1.86)+	-0.091 (2.32)*	-0.087 (1.15)
Cohort 7	-0.070 (1.79)+	-0.023 (0.36)	-0.036 (0.65)	-0.065 (1.53)	0.048 (0.51)
Cohort 8	-0.029 (0.69)	0.060 (0.89)	0.025 (0.43)	-0.021 (0.46)	0.100 (1.10)
Cohort 9	0.008 (0.17)	0.124 (1.72)+	0.082 (1.32)	0.016 (0.34)	0.203 (2.02)*
Cohort 10	0.026 (0.54)	0.173 (2.28)*	0.126 (1.91)+	0.035 (0.68)	0.242 (2.27)*
Cohort 11	0.057 (1.13)	0.224 (2.78)**	0.153 (2.20)*	0.070 (1.31)	0.290 (2.57)*
Cohort 12	0.085 (1.60)	0.259 (3.09)**	0.187 (2.57)*	0.105 (1.85)+	0.314 (2.68)**
Cohort 13	0.132 (2.32)*	0.322 (3.69)**	0.249 (3.25)**	0.154 (2.58)**	0.393 (3.21)**
Cohort 14	0.063 (0.96)	0.292 (2.96)**	0.214 (2.42)*	0.085 (1.24)	0.509 (3.59)**
Male head	0.022 (3.78)**	0.031 (3.39)**	0.028 (3.43)**	0.018 (2.87)**	0.036 (2.39)*
Maori/Pacific head	0.075 (8.71)**	0.098 (8.33)**	0.096 (9.04)**	0.080 (9.34)**	0.056 (3.73)**
Working head	0.155 (21.62)**	0.169 (15.38)**	0.168 (17.36)**	0.146 (19.14)**	0.048 (1.94)+
Unemployed head	-0.018 (0.93)	-0.041 (1.84)+	-0.026 (1.18)	-0.020 (1.13)	-0.125 (3.33)**
Head has self-employment	-0.059 (6.73)**	0.059 (3.21)**	0.017 (1.08)	-0.085 (7.58)**	0.312 (11.40)**
Income	0.051 (7.07)**	0.043 (4.73)**	0.047 (5.58)**	0.047 (6.69)**	-0.001 (0.08)
Paying mortgage	0.149 (19.07)**	0.277 (20.78)**	0.241 (21.91)**	0.134 (15.53)**	0.344 (18.15)**
Rent/mortgage-free	0.149 (19.07)**	0.277 (20.78)**	0.241 (21.91)**	0.134 (15.53)**	0.344 (18.15)**
Sole parent w/ children	-0.111 (10.97)**	-0.142 (10.55)**	-0.137 (11.54)**	-0.113 (11.52)**	-0.136 (6.77)**
Couple with children	-0.056 (8.94)**	-0.092 (8.83)**	-0.085 (9.59)**	-0.052 (7.54)**	-0.125 (8.56)**
Constant	-6.155 (8.21)**	-7.003 (6.57)**	-5.716 (5.96)**	-5.152 (6.87)**	-6.690 (4.22)**
Sample size	46269	46245	46112	44880	31514
R-squared	0.022	0.054	0.053	0.042	0.054

Note:

Coefficients weighted by population sampling weights. Absolute value of robust t-statistics in parentheses; + significant at 10% level * significant at 5% level; ** significant at 1% level. Each regression includes a fifth-order polynomial in age of the household head and 13 time dummies, with coefficients constrained to sum up to zero and to be orthogonal to a linear trend. Excluded categories are female head, non-Maori and non-Pacific head, head not in the labour force, dwelling is rented, family type is single adult, couple without children, mixed and non-family. For each regression, tests of the hypothesis that either the age, cohort or time effects are jointly zero are rejected at the $p < 0.001$ level.

Table 7.1: Cohort Effects in Mean Saving Rates With Different Definitions of Consumption and Saving

	Components added to current consumption				
	Using current consumption expenditure ^a	Education, health and insurance ^b	Durables ^c	Mortgage repayments and savings contributions ^d	Using total consumption ^e
Mean S/X	0.353	0.293	0.177	0.263	0.065
Cohort 2	-0.058 (1.36)	-0.075 (1.83)+	-0.075 (1.84)+	-0.060 (1.46)	-0.088 (2.30)*
Cohort 3	-0.076 (1.70)+	-0.109 (2.55)*	-0.092 (2.16)*	-0.072 (1.68)+	-0.114 (2.84)**
Cohort 4	-0.140 (2.93)**	-0.178 (3.86)**	-0.122 (2.65)**	-0.141 (3.07)**	-0.153 (3.55)**
Cohort 5	-0.161 (3.02)**	-0.219 (4.28)**	-0.121 (2.40)*	-0.172 (3.37)**	-0.177 (3.78)**
Cohort 6	-0.082 (1.31)	-0.158 (2.62)**	-0.054 (0.98)	-0.112 (1.88)+	-0.139 (2.71)**
Cohort 7	0.028 (0.37)	-0.078 (1.08)	0.072 (1.03)	-0.040 (0.56)	-0.069 (1.06)
Cohort 8	0.086 (1.16)	-0.035 (0.49)	0.134 (1.97)*	-0.008 (0.11)	-0.040 (0.64)
Cohort 9	0.155 (1.91)+	0.012 (0.16)	0.195 (2.64)**	0.041 (0.54)	-0.008 (0.12)
Cohort 10	0.204 (2.36)*	0.051 (0.62)	0.247 (3.16)**	0.070 (0.86)	0.019 (0.27)
Cohort 11	0.254 (2.78)**	0.091 (1.04)	0.292 (3.57)**	0.100 (1.16)	0.044 (0.58)
Cohort 12	0.284 (2.99)**	0.115 (1.26)	0.329 (3.86)**	0.118 (1.32)	0.065 (0.84)
Cohort 13	0.350 (3.54)**	0.163 (1.71)+	0.393 (4.43)**	0.184 (1.97)*	0.116 (1.43)
Cohort 14	0.336 (2.99)**	0.128 (1.18)	0.393 (3.94)**	0.157 (1.49)	0.092 (1.00)
Male head	0.037 (3.17)**	0.037 (3.38)**	0.021 (1.98)*	0.031 (2.81)**	0.019 (1.93)+
Maori/Pacific head	0.095 (7.69)**	0.112 (9.29)**	0.103 (9.28)**	0.101 (8.78)**	0.118 (11.57)**
Working head	0.159 (9.50)**	0.149 (9.24)**	0.117 (7.44)**	0.116 (7.17)**	0.080 (5.42)**
Unemployed head	-0.060 (2.22)*	-0.044 (1.71)+	-0.031 (1.25)	-0.044 (1.71)+	-0.012 (0.52)
Head has self-employmt Income	0.075 (3.59)**	0.068 (3.47)**	0.063 (3.64)**	0.067 (3.47)**	0.052 (3.46)**
Paying mortgage	0.043 (4.54)**	0.038 (4.20)**	0.026 (3.25)**	-0.058 (6.79)**	-0.052 (7.41)**
Rent/mortgage-free	0.297 (19.56)**	0.261 (18.17)**	0.209 (16.10)**	0.257 (18.03)**	0.158 (13.42)**
Sole parent w/ children	-0.147 (9.76)**	-0.141 (9.76)**	-0.113 (8.58)**	-0.112 (7.88)**	-0.083 (6.93)**
Couple with children	-0.103 (8.70)**	-0.113 (10.06)**	-0.089 (8.48)**	-0.086 (7.93)**	-0.081 (8.81)**
Constant	-7.852 (6.36)**	-7.462 (6.29)**	-6.982 (6.16)**	-5.802 (4.97)**	-5.234 (5.01)**
R-squared	0.040	0.024	0.040	0.044	0.044

Note:

Coefficients weighted by population sampling weights. Absolute value of robust t-statistics in parentheses; + significant at 10% level * significant at 5% level; ** significant at 1% level. Each regression includes a fifth-order polynomial in age of the household head and 13 time dummies, with coefficients constrained to sum up to zero and to be orthogonal to a linear trend. Excluded categories are female head, non-Maori and non-Pacific head, head not in the labour force, dwelling is rented, family type is single adult, couple without children, mixed and non-family, head has a university qualifications. $N=46269$.

^a The estimates in this column are reproduced from column (iii) of Table 6.2.

^b Insurance includes both health and life insurance.

^c Durables includes motor vehicles.

^d The mortgage repayments are only for the principle.

^e Total consumption is current consumption expenditures plus medical, health and insurance spending, plus durables expenditures, plus mortgage principle repayments and contributions to savings.

Table 8.1 Lifetime Patterns for Different Cohorts

Cohort Number	Birth years	Saving Rate (a)	Working life	Peak Saving Years
1	1910-1914	0 (b)	1930-1974	1955-1969
2	1915-1919	0 (b)	1935-1980	1960-1974
3	1920 -1924	Negative	1940-1984	1965-1979
4	1925-1929	Negative	1945-1989	1970-1984
5	1930-1934	Negative	1950-1994	1975-1990
9	1950-1954	Positive	1970-2014	1995-2010
10	1955-1959	Positive	1975-2019	2000-2014

Notes:

- a. Refers to lifetime saving rate relative to the 1910-14 reference group..
- b. 1910-14 is the reference group (by definition zero). See pattern of cohort dummies in the regressions presented in Section 6. Negative and positive refer to the cohorts that were typically significantly lower or higher in their lifetime saving rates.

Appendix Table A: Wald tests of joint statistical significance of age and cohort effects in the unrestricted 'Deaton Chapter 6' regressions

	ln (disposable income)		ln (current consumption)	
	Age	Cohort	Age	Cohort
Mean	$\chi^2_{(55)}=347.55$ $p<0.000$	$\chi^2_{(69)}=124.49$ $p<0.000$	$\chi^2_{(55)}=229.33$ $p<0.000$	$\chi^2_{(69)}=139.64$ $p<0.000$
25 th percentile	$\chi^2_{(55)}=207.94$ $p<0.000$	$\chi^2_{(69)}=155.45$ $p<0.000$	$\chi^2_{(55)}=166.96$ $p<0.000$	$\chi^2_{(69)}=118.46$ $p<0.000$
50 th percentile	$\chi^2_{(55)}=400.08$ $p<0.000$	$\chi^2_{(69)}=152.86$ $p<0.000$	$\chi^2_{(55)}=209.08$ $p<0.000$	$\chi^2_{(69)}=112.93$ $p<0.001$
75 th percentile	$\chi^2_{(55)}=668.65$ $p<0.000$	$\chi^2_{(69)}=181.68$ $p<0.000$	$\chi^2_{(55)}=191.04$ $p<0.000$	$\chi^2_{(69)}=111.68$ $p<0.001$
<i>Population sub-groups (all based on cell medians)</i>				
Household head is Maori or Pacific Islander	$\chi^2_{(55)}=73.08$ $p<0.052$	$\chi^2_{(68)}=93.63$ $p<0.022$	$\chi^2_{(55)}=74.82$ $p<0.040$	$\chi^2_{(68)}=110.44$ $p<0.000$
Household head not Maori or Pacific Islander	$\chi^2_{(55)}=469.14$ $p<0.000$	$\chi^2_{(69)}=166.91$ $p<0.000$	$\chi^2_{(55)}=200.18$ $p<0.000$	$\chi^2_{(69)}=87.65$ $p<0.064$
Household head is female	$\chi^2_{(55)}=132.13$ $p<0.000$	$\chi^2_{(69)}=106.05$ $p<0.003$	$\chi^2_{(55)}=117.07$ $p<0.000$	$\chi^2_{(69)}=92.69$ $p<0.031$
Household head is male	$\chi^2_{(55)}=391.96$ $p<0.000$	$\chi^2_{(69)}=138.63$ $p<0.000$	$\chi^2_{(55)}=201.90$ $p<0.000$	$\chi^2_{(69)}=111.38$ $p<0.001$
Household head is working (full/part-time)	$\chi^2_{(55)}=126.79$ $p<0.000$	$\chi^2_{(69)}=87.23$ $p<0.069$	$\chi^2_{(55)}=47.93$ $p<0.740$	$\chi^2_{(69)}=85.14$ $p<0.091$
Household head is not working	$\chi^2_{(55)}=80.88$ $p<0.014$	$\chi^2_{(69)}=91.09$ $p<0.039$	$\chi^2_{(55)}=66.04$ $p<0.147$	$\chi^2_{(69)}=67.05$ $p<0.545$
Household head has no school qualifications	$\chi^2_{(55)}=133.84$ $p<0.000$	$\chi^2_{(69)}=145.71$ $p<0.000$	$\chi^2_{(55)}=187.20$ $p<0.000$	$\chi^2_{(69)}=178.83$ $p<0.001$
Household head has at least School Certificate	$\chi^2_{(55)}=447.66$ $p<0.000$	$\chi^2_{(69)}=189.56$ $p<0.000$	$\chi^2_{(55)}=159.51$ $p<0.000$	$\chi^2_{(69)}=133.69$ $p<0.001$
<i>Household type^a</i>				
Single adult	$\chi^2_{(55)}=207.69$ $p<0.000$	$\chi^2_{(69)}=131.35$ $p<0.000$	$\chi^2_{(55)}=126.80$ $p<0.000$	$\chi^2_{(69)}=107.96$ $p<0.002$
Couple without children	$\chi^2_{(55)}=322.50$ $p<0.000$	$\chi^2_{(69)}=69.18$ $p<0.018$	$\chi^2_{(55)}=197.11$ $p<0.000$	$\chi^2_{(69)}=62.82$ $p<0.687$
Sole parent	$\chi^2_{(55)}=41.27$ $p<0.916$	$\chi^2_{(69)}=94.27$ $p<0.024$	$\chi^2_{(55)}=81.67$ $p<0.012$	$\chi^2_{(69)}=81.89$ $p<0.138$
Couple with children	$\chi^2_{(55)}=253.05$ $p<0.000$	$\chi^2_{(69)}=127.21$ $p<0.000$	$\chi^2_{(55)}=227.03$ $p<0.000$	$\chi^2_{(69)}=79.50$ $p<0.161$
<i>Dwelling tenure^a</i>				
Owned	$\chi^2_{(55)}=422.72$ $p<0.000$	$\chi^2_{(69)}=194.61$ $p<0.000$	$\chi^2_{(55)}=267.36$ $p<0.000$	$\chi^2_{(69)}=104.73$ $p<0.004$
Rented	$\chi^2_{(55)}=199.30$ $p<0.000$	$\chi^2_{(69)}=113.26$ $p<0.001$	$\chi^2_{(55)}=115.21$ $p<0.000$	$\chi^2_{(69)}=67.93$ $p<0.514$
Owned, with mortgage	$\chi^2_{(55)}=194.58$ $p<0.000$	$\chi^2_{(69)}=99.36$ $p<0.008$	$\chi^2_{(55)}=135.79$ $p<0.000$	$\chi^2_{(69)}=129.32$ $p<0.004$
Owned, mortgage-free	$\chi^2_{(55)}=237.05$ $p<0.000$	$\chi^2_{(69)}=102.55$ $p<0.006$	$\chi^2_{(55)}=168.54$ $p<0.000$	$\chi^2_{(69)}=127.55$ $p<0.000$


Note: Results are Wald tests of the hypothesis that coefficients on either age or cohort dummy variables are jointly zero in a seemingly-unrelated regression model where $\ln Y$ and $\ln X$ are each regressed on age and cohort dummies and on cell averages of the number of children, adults and old people per household. The sample is 840 cell averages (medians or means as noted), where cells are based on the interaction of age (19-74 inclusive) and survey year (1984-98). The regressions on household type sub-groups exclude the variables counting number of children, adults and elderly, to avoid singularity problems.

The population sub-groups formed according to school qualifications of the household head exclude data from 1987 and exclude observations where the qualification is not applicable (applied to most household heads older than 64) or not specified.

^aThe mixed family and non-family household types are excluded.

^bThe rent-free category is excluded

**Figure 1: Household Saving Rate
(Household Income and Outlay Accounts vs. HES Flow measure)**



**Figure 2.1 Lorenz Curve for Savings
(Positive Savers Only)**




Figure 2.2: Comparison of Cumulative Densities for Savings Rates


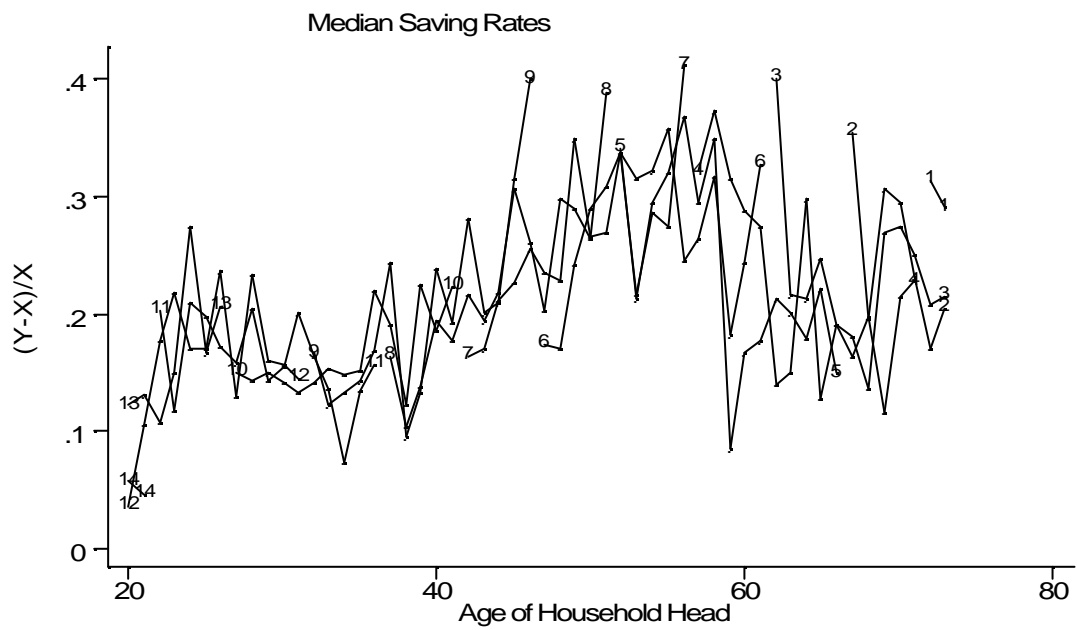



Figure 3.1: Household Saving Rates by Five-Year Birth Cohort



Key		
1=1910-14	5=1930-34	10=1955-59
2=1915-19	6=1935-39	11=1960-64
3=1920-24	7=1940-44	12=1965-69
4=1925-29	8=1945-49	13=1970-74
	9=1950-54	14=1975-79

Figure 4.1: Age Effects in (log) Income and Consumption and in the Derived Saving Rate [$\ln(y)-\ln(c)$]

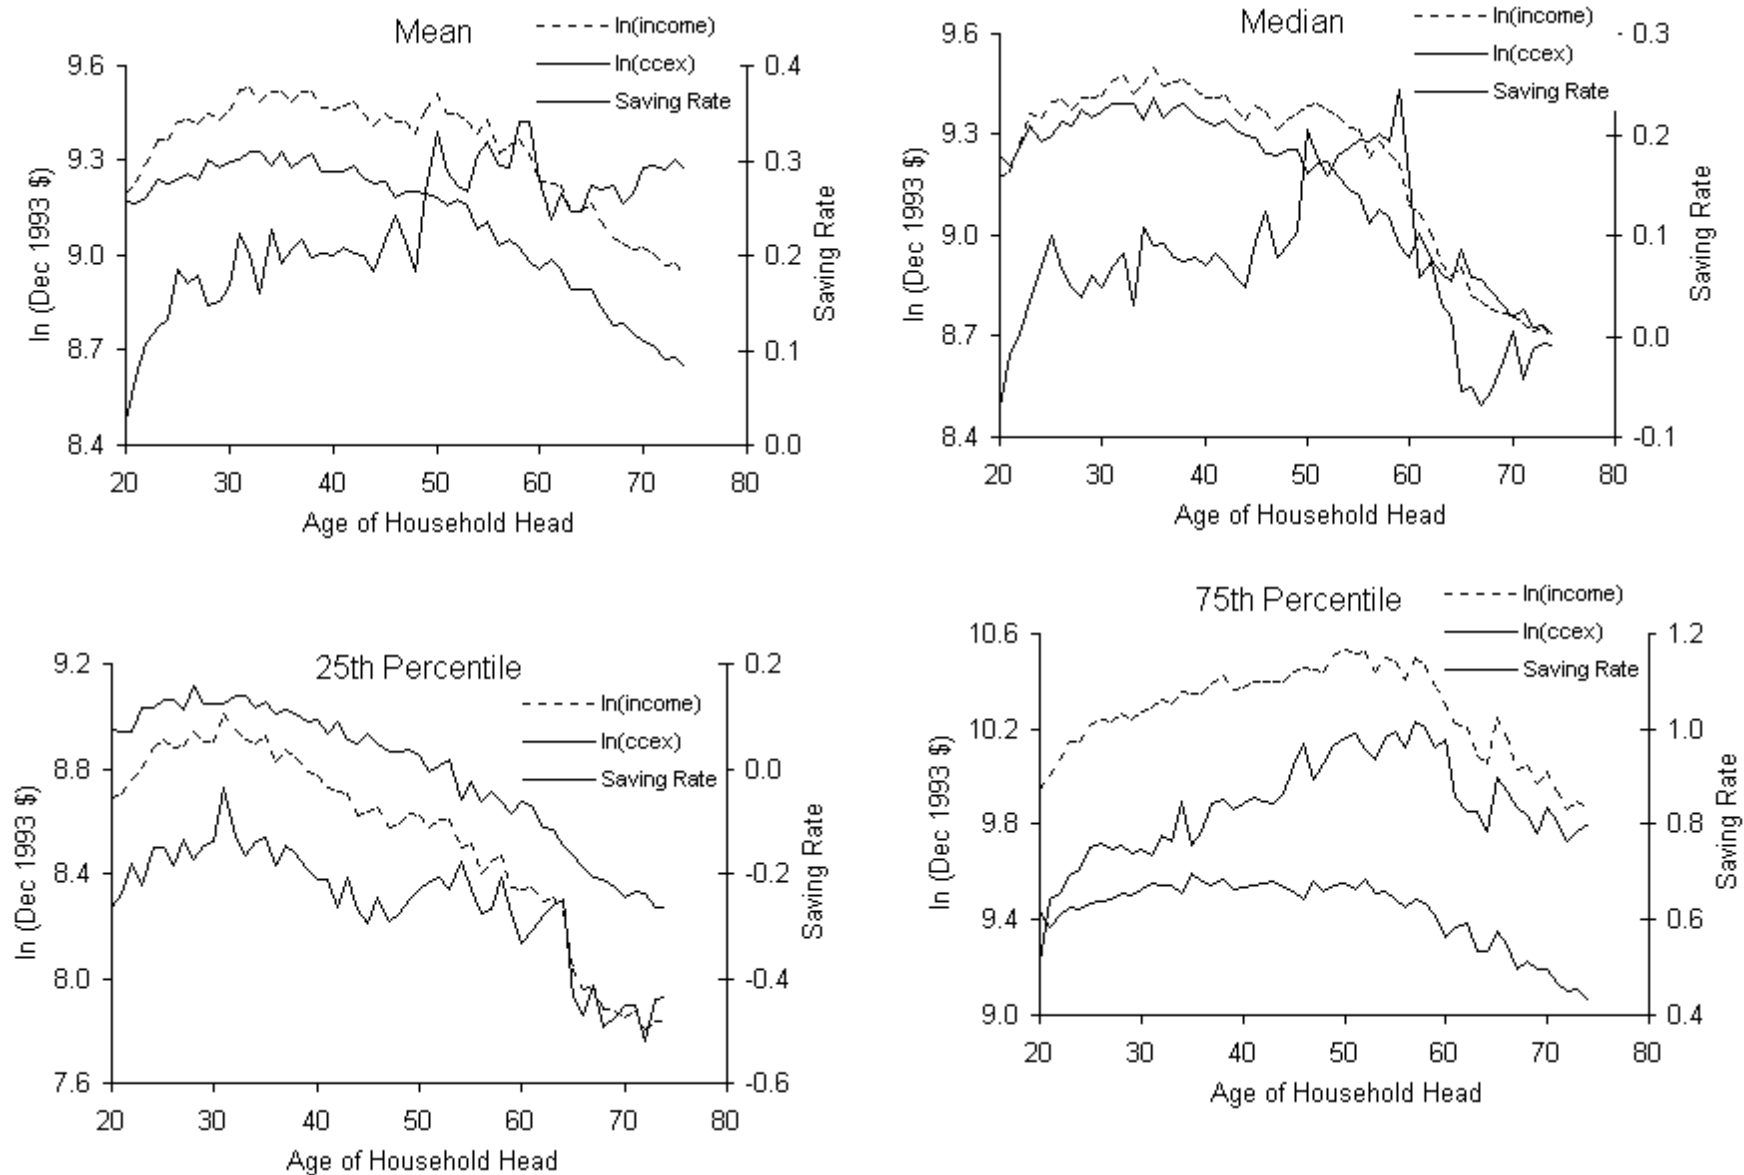


Figure 4.2: Cohort Effects in (log) Income and Consumption and in the Derived Saving Rate [$\ln(y)-\ln(c)$]






Figure 4.3: Age and Cohort Effects in Income, Consumption and Saving Rates $[\ln(y)-\ln(c)]$ for Male and Female-headed Households




Figure 4.4: Age and Cohort Effects in Income, Consumption and Saving Rates $[\ln(y)-\ln(c)]$ by Ethnicity of Household Head

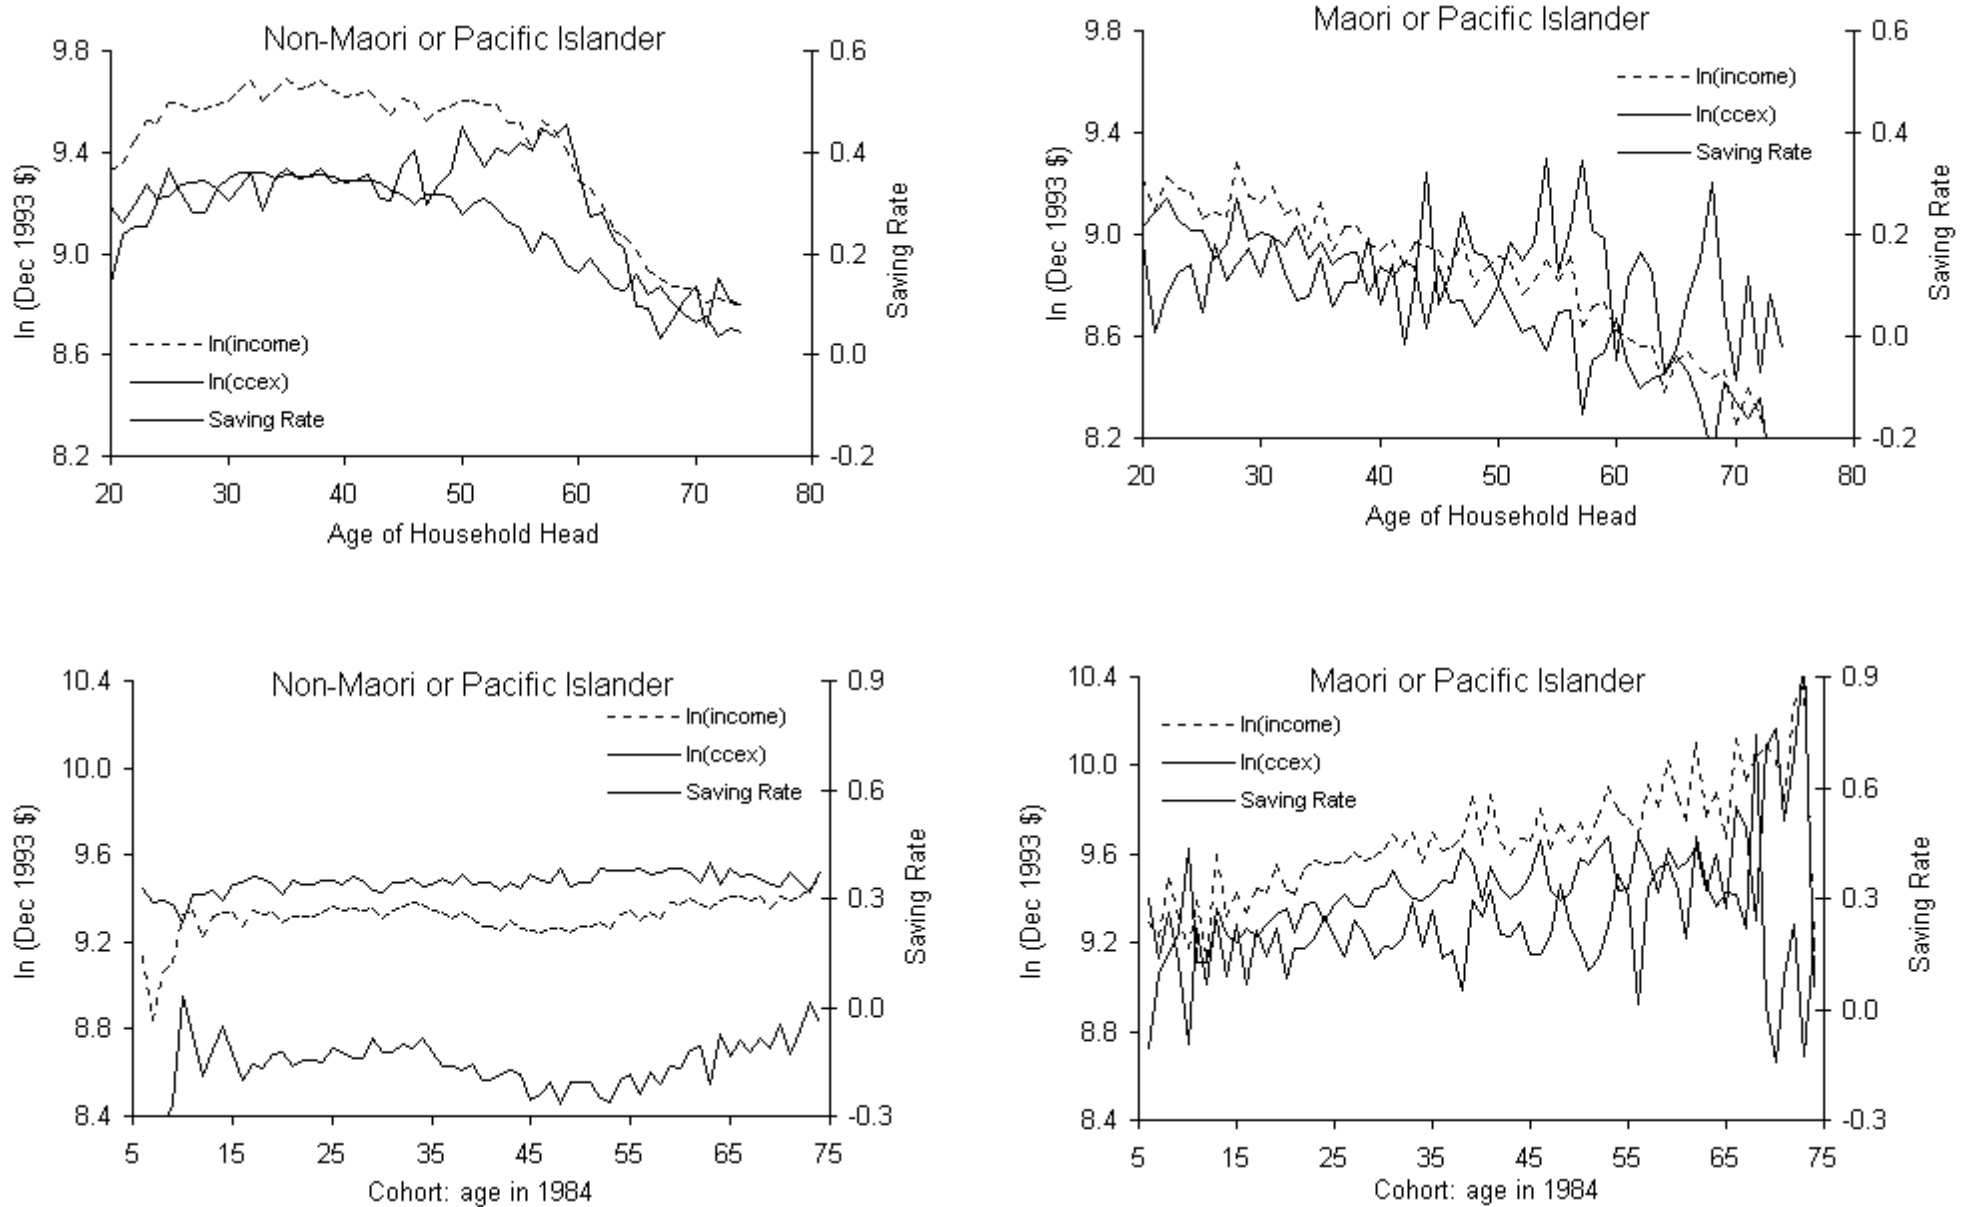


Figure 4.5: Estimates of Age-Specific Income and Consumption (B_{at} 's) by Survey Year




Figure 4.6: Age and Cohort Effects in Saving Rates [$\ln(y) - \ln(c)$] for Individuals and Households




Figure 5.1: Age, Cohort and Time Decomposition of Mean Savings Rate (Age×Year Cells)




Figure 5.2: Relationship Between Survey Year Effects in Mean Saving Rates and Various Macroeconomic Variables




Figure 5.3: Age, Cohort and Time Decomposition of Median Savings Rate





Figure 5.4: Age, Cohort and Time Decomposition at 25th and 75th Percentile Savings Rate

25th Percentile



75th Percentile




Figure 5.5: Age, Cohort and Time Decomposition of Mean Savings Rate: Male versus Female-headed households


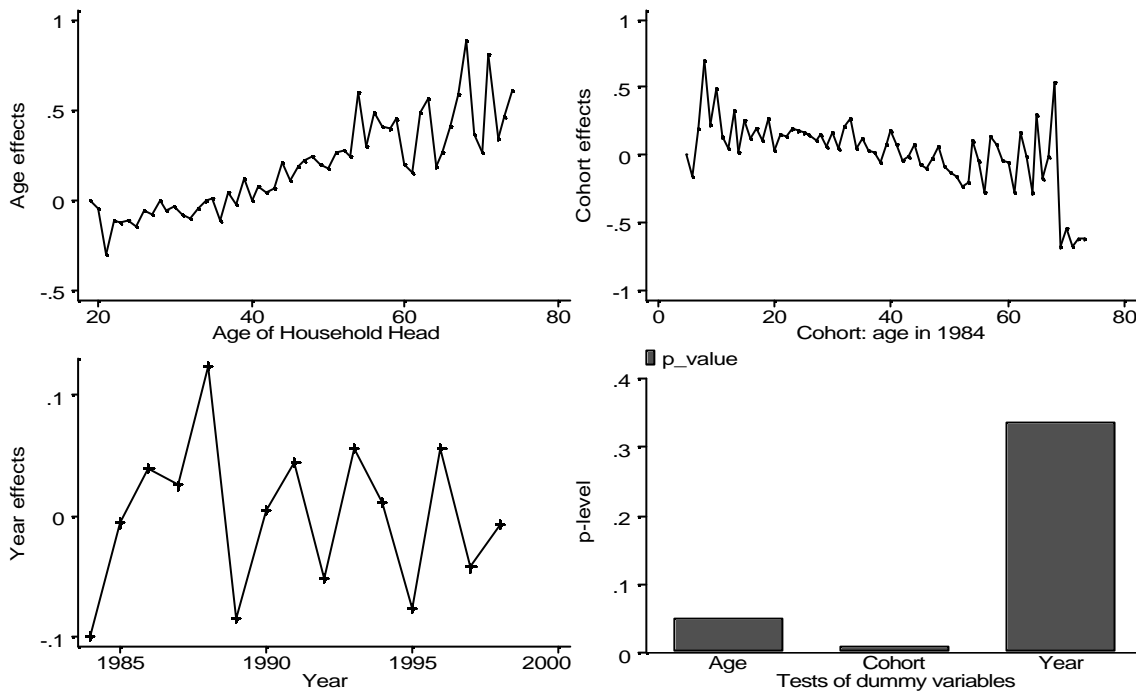


Figure 5.6: Age, Cohort and Time Decomposition of Mean Savings Rate: Maori and Pacific Islander-headed households Compared with Other Households

Maori and Pacific Islander-Headed Households



Households Headed by Other Ethnic Groups




Figure 6.1: Cohort Effects With Different Sets of Controls

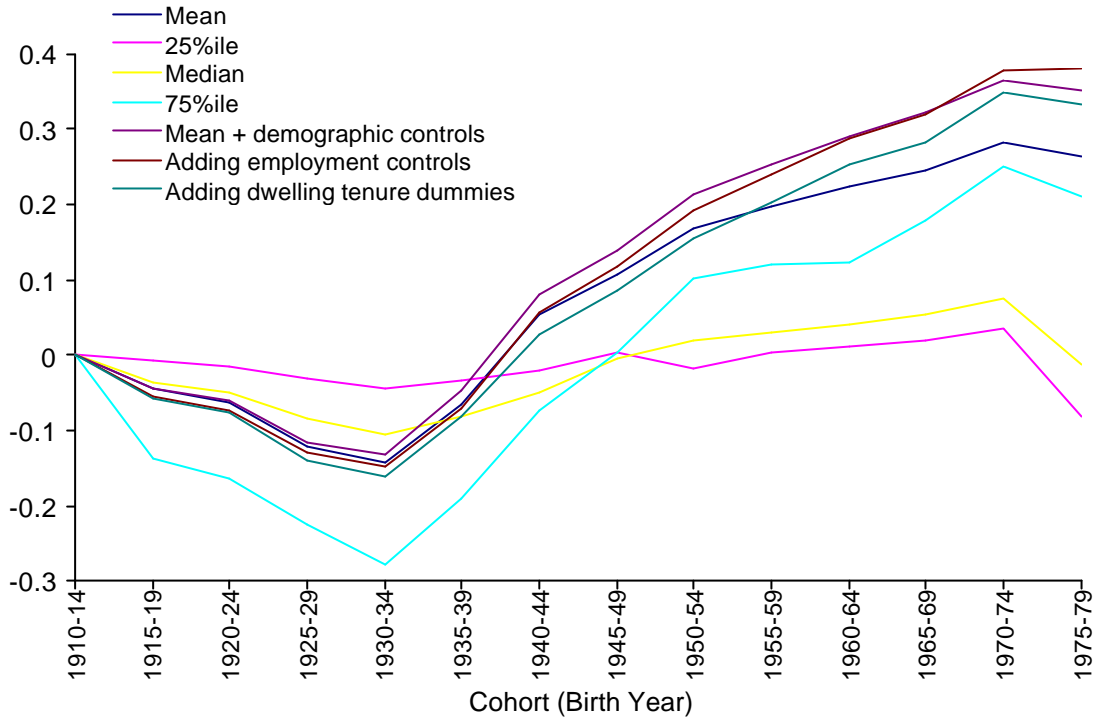


Figure 6.2: Smoothed Savings Rate by Cohor




Figure 6.3: Smoothed Savings Rate by Cohort: Quantile Regression




Figure 6.4: Cohort Effects With Different Estimation Samples
(Demographics, Employment, Tenure and Family Type Controls)

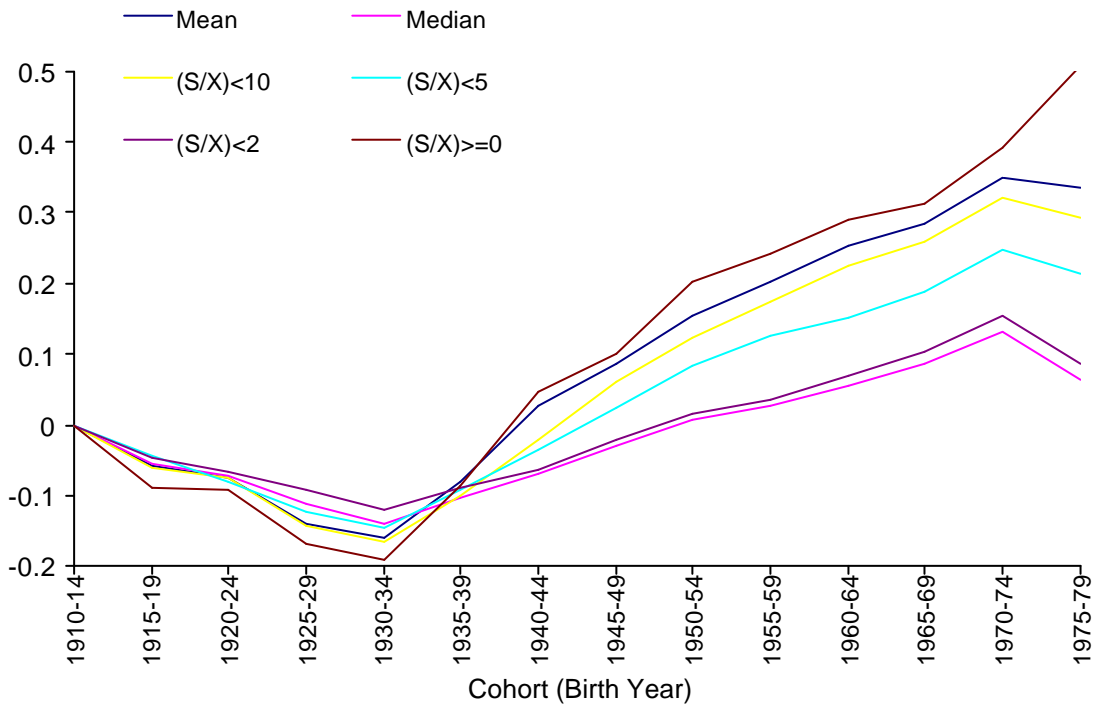


Figure 7.1: Cohort Effects With Different Definitions of Saving
 (Demographics, Employment, Tenure and Family Type Controls)

