Household Saving Behaviour in New Zealand: Why do Cohorts Behave Differently?

Grant M Scobie and John K Gibson

NEW ZEALAND TREASURY

WORKING PAPER 03/32

DECEMBER 2003

THE TREASURY
Kaitohutohu Kaupapa Rawa

NZ TREASURY WORKING PAPER 03/32

Household Saving Behaviour in New Zealand: Why do Cohorts Behave Differently?

MONTH/YEAR

December 2003

AUTHORS

Grant M Scobie
New Zealand Treasury
PO Box 3724
Wellington
New Zealand

Email grant.scobie@treasury.govt.nz
Telephone 64 4 471 5005
Fax 64 4 473-0982

John Gibson
University of Waikato
Private Bag
Hamilton
New Zealand

Email jkgibson@waikato.ac.nz
Telephone 64 7 838-4466
Fax 64 7 838-4331

ACKNOWLEDGEMENTS

The authors thank Ivan Tuckwell for his help in preparing the data files and participants in seminars at the University of Waikato and the Treasury. In particular, Bob Buckle and John Creedy have made helpful suggestions.

NZ TREASURY

New Zealand Treasury
PO Box 3724
Wellington 6000
NEW ZEALAND

Email information@treasury.govt.nz
Telephone 64-4-472 2733
Website www.treasury.govt.nz

DISCLAIMER

Abstract

The aim of this paper is to add to the understanding of saving decisions by households. The saving behaviour of households is found to differ depending on the birth cohort of the household head. This paper seeks to explain why this pattern might exist. It is based on an analysis of synthetic cohorts derived from unit record data taken from the Household Economic Survey (HES) for the March years 1984 to 1998. The need to use synthetic cohorts arises as the HES is not a longitudinal panel survey, but rather a time series of independent cross-sectional samples. We use a range of regression models to separate out the effect of age, birth-year cohort and year on saving rates. The typical saving rates for the cohorts born between 1920 and 1939 are found to be significantly lower relative to the younger and older cohorts studied. This pattern of cohort effects is robust to the inclusion of conditioning variables; to the trimming from the sample of households with either negative or very large ratios of savings to consumption, and to different definitions of saving. Some exploratory 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.

JEL CLASSIFICATION E21 Consumption; Saving
J26 Retirement

KEYWORDS Household saving rates; cohort effects; New Zealand; economic and social policies

Table of Contents

List of Tables

List of Figures

Household Saving Behaviour in New Zealand: Why do Cohorts Behave Differently?

1 Introduction

This paper has two major objectives. The first is to present estimates of the household saving patterns for different aged cohorts. The second is to offer an explanation of why saving behaviour might be different for different aged cohorts.

It is extremely difficult to glean the implications of saving for say retirement income from aggregate data on the household sector. Aside from difficulties of measurement, a low overall level of saving in ageing population could be consistent with high saving by those in their working years offset by dissaving among an expanding older population of retirees. In short, a better understanding of saving by households requires an analysis of micro data based on individual household records.

This study uses individual records from the Household Expenditure Survey (HES) for a 15-year period to construct synthetic cohorts (Section 2). Regression models estimate the average saving rates for each five-year birth cohort, after allowing for age and year effects together with a set of conditioning variables (Section 3). Saving rates are found to differ markedly across cohorts. Section 4 presents some tentative findings, which suggest that household saving behaviour may well be influenced by economic conditions and social policies. It is argued that the cohort patterns of saving may reflect different conditions which faced different cohorts as they moved through their working ages, especially those existing during their peak saving years. Conclusions follow in Section 5.

The study finds that different cohorts do display different saving patterns. Those born from 1920 to 1939 are found to have significantly lower saving rates than older or younger cohorts. The paper finds that these differences are consistent with the fact that each cohort faced a different set of economic and social policies. The environment that prevailed especially during peak earning and saving periods was different for the different cohorts. In particular a more “favourable” environment that prevailed in the period 1950-1980 seems to explain why certain cohorts had lower saving rates. The implication is that extent of public provision of social welfare and retirement benefits together with conditions in labour markets, do influence the rate at which households will save.

2 Construction of the data and the cohort approach

In this section we first describe the data, and how we have constructed measures of saving and synthetic cohorts. We then present a first glimpse of household saving patterns based on these cohorts.

2.1 The Data

We have used the income and expenditure data from the HES to estimate saving as a residual. We do this fully cognizant of the limitations of the HES.[1] Our defence 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.

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 principal, and contributions to savings. In other words we attempt to construct a measure of expenditure, which would result in an "economic” view of saving.[2] 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.

The data cover the years 1983-48 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)[3]. 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.

[1] "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". Background Notes to the Household Economic Survey, Statistics New Zealand (1998), p.17.

[2] For details of the adjustments see Gibson and Scobie (2001).

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

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 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.[4]

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 zero[5].

Ideally, a complete accounting for income and saving requires both flow measures and a household balance sheet to track stock changes. The HES provides no measures of the stock of household wealth. As a consequence, any contributions made by an employer to a private pension fund are not recorded as saving; and any withdrawals from a pension fund are counted as income in the year received rather than "dissaving". Fortunately, pension schemes (outside the public superannuation scheme) play a relatively minor role in New Zealand[6].

Should the unit of analysis 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. It must however be recognised that in multi-generational households the saving of working age members could be offset by the dissaving of younger and elderly members of the same household. The net saving rate in such a household could then differ quite significantly from the saving rates of individual members.

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).

---

[4] Another reason for eliminating the oldest households is that pension income may not be distinguished from other income. Failure to recognise the running down of the underlying pension assets will lead overstating saving by the elderly. See Deaton and Paxson (2000), who report a more hump shaped pattern of saving with respect to age when flows into and from pension funds are included.

[5] Denote saving by S, consumption expenditure by X and disposable income by YD. Then:
S/YD =(S/X) .(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.

[6] In 1997-98, 6 percent of all households received income from a private pension, and this accounted for just over 1 percent of their gross income. A little under one half of the household receiving a pension had a head aged 65-74. Among this group, pensions made up 8.5 percent of gross income. 76 percent of these households had no pension income.

Table 1 reports the number of observations in each year together with the saving rates at the mean, median, 25th and 75th percentiles. The final column reports the ratio of the averages of saving to consumption, as distinct from the average (or quantile) of the ratios. The mean is clearly influenced by extreme values, so frequently it will be helpful to focus on the median saving rate.

Table 1 – Sample size and saving rates by survey year

	Total Sample	Savings Rate (S/X)				$\frac{\overline{S}}{\overline{X}}$
		Mean	25th Percentile	Median	75th Percentile	
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

It should be stressed that there is 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 (e.g. 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. Nearly 32 percent of the sample report negative saving. Inevitably some of this may be due to under-reporting of income. Evidence from the USA suggests that when the differential under-reporting of both income and consumption is allowed for, up to one third of the apparent fall in the saving rate between 1972-73 and 1983-84 may be due to misreporting (Bosworth, Burtless and Sabelhaus (1991)).

2.2 The Use of Synthetic Cohorts

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.[7] But the availability of a time-series of cross-sectional Household Economic Surveys allows us to construct *synthetic panels* following methods described by Shorrocks (1975) and Deaton (1985)

[7] 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 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).

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 individuals.[8]

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 Surveys. Provided the population is not much affected by migration, 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 group 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, 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 five-year birth intervals.

[8] Verbecek and Nijman (1992) note that 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.

Table 2 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.[9] 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 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.

Table 2 – Cohort definitions, cell sizes and saving rates

Year of Birth	Ages Observed	Average Cell Size	Total Sample	Savings Rate (S/X) Mean	25th Percentile	Median	75th Percentile	$\frac{\bar{S}}{\bar{X}}$
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

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.[10] 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.[11]

Table 3 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 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.[12]

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

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

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

[12] We undertake such an analysis in Section 4.

Table 3 — 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

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 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 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 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 1).

Figure 1 – Household Saving Rates by Five-Year Birth Cohort

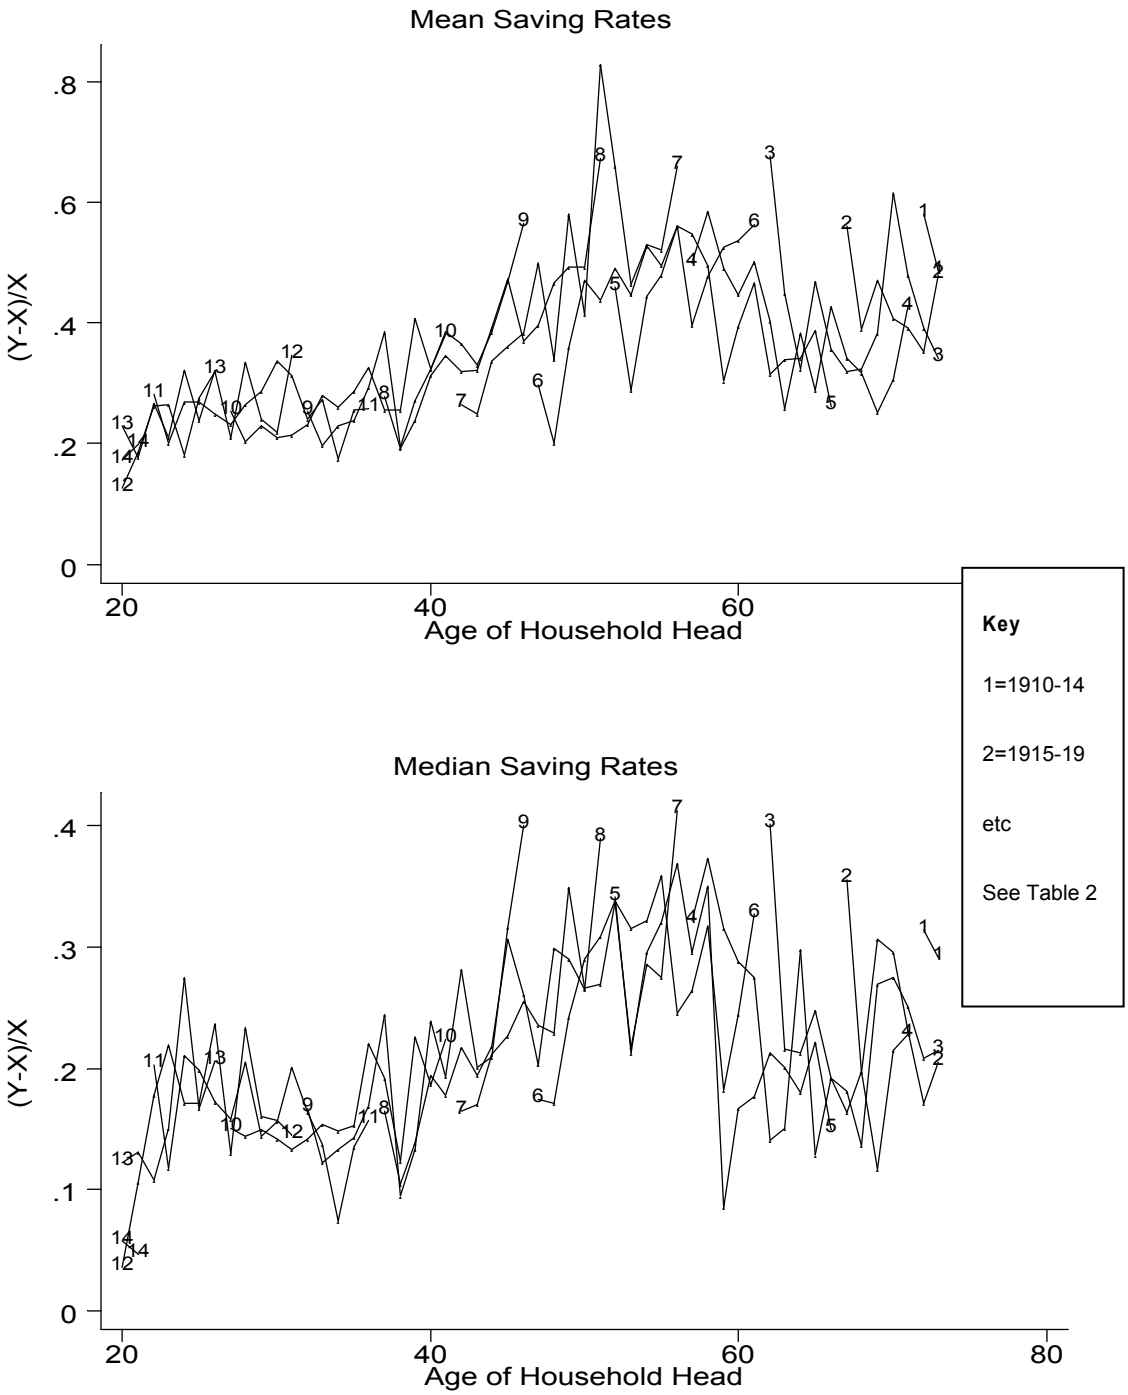

3 The Model and Results

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. The results shown in Figure 1 are consistent with the usual pattern of saving with respect to age. However, saving behaviour may also evolve over time and vary across birth-year cohorts as economies grow and as certain fluctuations affect individuals contemporaneously. In this section we first present the basic model (Section 3.1), which when estimated with the aid of some identifying structure, makes it possible to distinguish the separate effects of age, cohort and time when using data observed over different age intervals for different birth cohorts. The basic results are in Section 3.2 while additional variations follow in Sections 3.3 and 3.4.

3.1 The Basic Model

The results in the previous section show that there is 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, more structure may be needed. One common technique in the literature is to replace individual age effects with a fifth-order polynomial in the age of the household head. This will tend to smooth out much of the noise while still being sufficiently flexible to capture the shape of the underlying age profile. This approach can also be extended to the cohort effects, replacing them with a fifth-order polynomial in year-of-birth.[13] 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.[14]

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 1. The basic model is:

$$s_t^{ch} = f\left(a^h\right) + D^c\gamma + \delta\left(d_t\right) + u_t^{ch} \qquad (1)$$

where:

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

f = a function representing the fifth-order polynomial in age (a);

γ = the coefficient vector capturing the effects of the cohort intercept dummies;

δ = a function representing the time effects, d_t; and

u_t^{ch} = the residual term.

[13] See, for example, Jappelli (1999).
[14] See, for example, Attanasio (1998).

The separate effects of *a*, c, and *a*+c which equals t, cannot be identified in equation (3.1) and any trends in the data can be arbitrarily attributed to year effects, or a combination of age and cohort effects (Deaton, 1997).

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.[15] 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).

3.2 The Initial Results

The basic results of estimating equation (1) are presented in this section. In keeping with our previous search for robust patterns in the data, this equation is estimated for mean saving rates and for three quantiles: the 25th, 50th, and 75th percentiles of the distribution of saving-to-expenditure ratios.

Table 4 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.[16]. In all cases we reject the hypotheses that either the year effects, the age effects, or the cohort effects are jointly zero.[17]

[15] We have adopted this approach following Deaton (1997, p.126). The reparameterisation that is implied is most clearly seen by writing the model in its most general form as:

$$s/x = D^{a}\alpha + D^{b}\gamma + D^{a+b}\delta + u$$

where b is used to denote individual birth-year cohorts, and then stating the restrictions as:

$$\sum_{a+b=t}^{T} D^{a+b} = 0$$
$$\sum_{a+b=t}^{T} (a+b)\cdot D^{a+b} = 0$$

For a full discussion of this see Gibson and Scobie (2001).

[16] To reduce the volume of results, we have not reported the additional 14 coefficients for the survey year effects.

[17] It will be apparent to the reader 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 purpose here is to identify age and particularly cohort effects, rather than to estimate a household saving function, per se.

Table 4 – Cohort effects in individual saving rates, controlling for age and year effects

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

The cohort effects in the mean saving rate are reported in the first column of Table 4 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 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 4 show that the saving rate for the 1930-34 birth cohort is 28 percent lower than for the 1910-14 cohort. While the cohort effect appears to be largely restricted to the upper end of the distribution, it must be stressed that these are the households who contribute the bulk of aggregate saving.[18] It is therefore, important to describe and understand these cohort effects if one is to make any inferences about the future path of aggregate savings.

3.3 Adding Conditioning Variables[19]

To check whether within-cell heterogeneity can explain the marked pattern of cohort effects, 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 4 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 test if the shifts in lifecycle saving profiles can be explained by these demographic, education and family structure effects, rather than by pure cohort effects.

Some of the conditioning variables, such as employment status, are likely to change over the lifecycle, whereas others, such as ethnicity and gender obviously remain fixed. 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[20]. Therefore, the specification of the regression models is:

$$s_t^{ch} = f\left(a^h, c^h\right) + \beta' w_t^{ch} + u_t^{ch} \qquad (2)$$

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} = a vector of conditioning variables; and

u_t^{ch} = the residual term.

[18] Over the fifteen-year sample period, total household savings were estimated at $93bn of which $87bn was accounted for by the households whose disposable incomes fell in the top three deciles on the income distribution.

[19] This section presents only abbreviated results of the findings of the robustness of the basic model as different sets of conditioning variables were added. A complete set is available from the authors on request.

[20] See Attanasio (1998) for an example of interacting education with cohort dummies.

When the gender and ethnicity[21] of the household head are included, 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 did not change the basic pattern of cohort intercepts. As in the previous results (Table 4), households whose head was born ca. 1925-1939 show lower than average saving rates, while the rise in savings rates for the more recently born cohorts is even more apparent than when the demographic controls were absent.

The next model included 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,[22] so it 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 was reported in Table 4, 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.[23]

Figure 2 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 2 is its similarity to results reported by Attanasio (1998, Figure 9). For the U.S., Attanasio finds that 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 strikingly similar.

[21] Both these are fixed over the lifecycle.

[22] 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.

[23] One hypothesis, untested in this study, is that because of lower accumulated wealth and more erratic employment history, Maori and Pacific Island households do not enjoy the same access to credit, inducing a higher level of savings, other factors constant.

Figure 2 – Cohort Effects with Different Sets of Controls

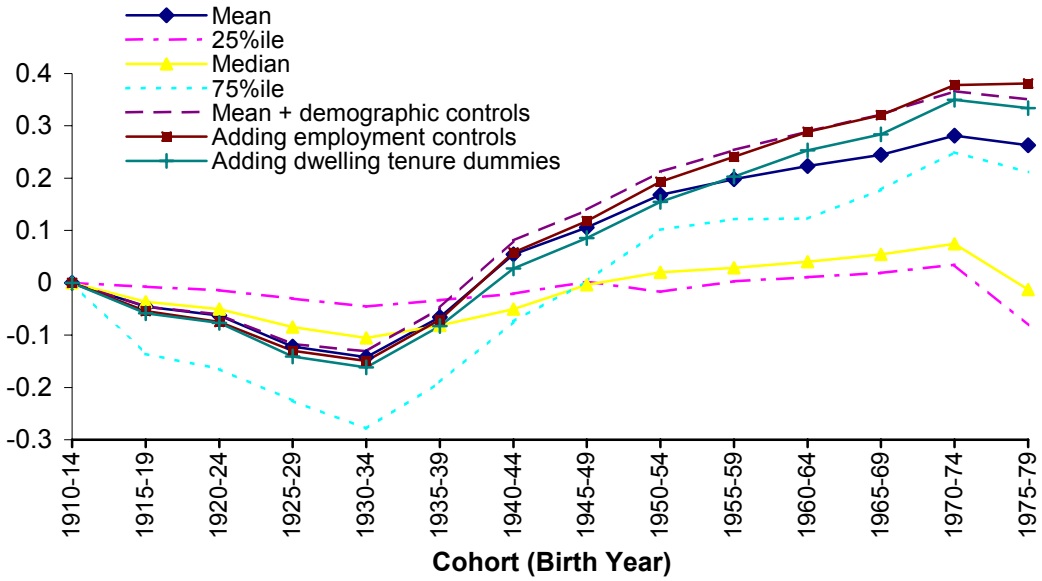

When the basic model of five-year birth-cohorts, is augmented with demographic, employment, family type and tenure status variables, we obtain a reasonably robust description of the underlying data. We therefore use the predictions from this model in to illustrate the shape of the age profile in median household saving rates and to show how that profile has shifted up and down across birth cohorts (Figure 3).

Figure 3 – Smoothed mean savings rate by Cohort

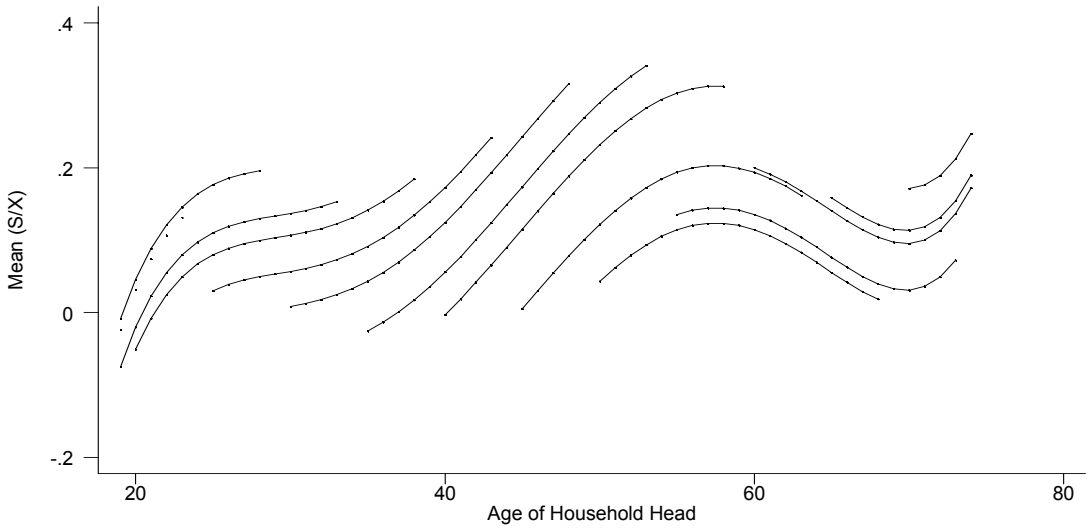

The typical age profile for the average saving rate is somewhat hump shaped with a peak around age 57 but does not become negative at older ages. While the hump shape is consistent with the lifecycle hypothesis, the apparent increase in the saving rate beyond

age 70 is not.[24] 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).

In Figure 2 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. To investigate whether the pattern of cohort effects is just due to some of the extreme values of saving rates that are present in the data various "trimmed samples" that removed extreme values of saving rates were used. In all cases, the models include the demographic, employment, family type and tenure status variables. The results again showed 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, although 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 4. 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.

Figure 4 – Smoothed median savings rate by Cohort: quantile regression

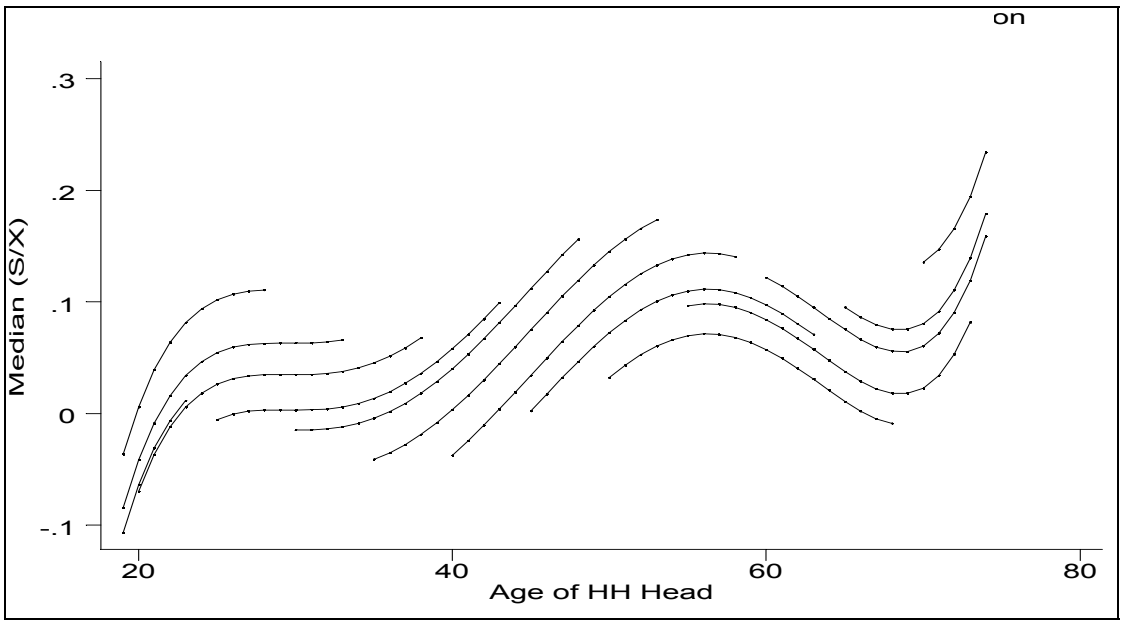

Figure 5 plots the cohort intercepts for each of the models estimated with reduced samples and for the model of mean saving rates estimated on the full sample. Although the magnitude of the cohort effects vary as the sample or estimation method is altered, the relative ranking of cohorts does not change. In all cases, saving rates fall from earlier to later born cohorts between the 1910-14 and 1930-34 cohorts and then the pattern

[24] 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. See Gibson and Scobie (2001) where this same feature was evident in the unrestricted estimates of the mean saving rate.

reverses with later born cohorts exhibiting higher saving rates. These patterns seem sufficiently robust to warrant further investigation.

Figure 5 – Cohort effects with different estimation samples

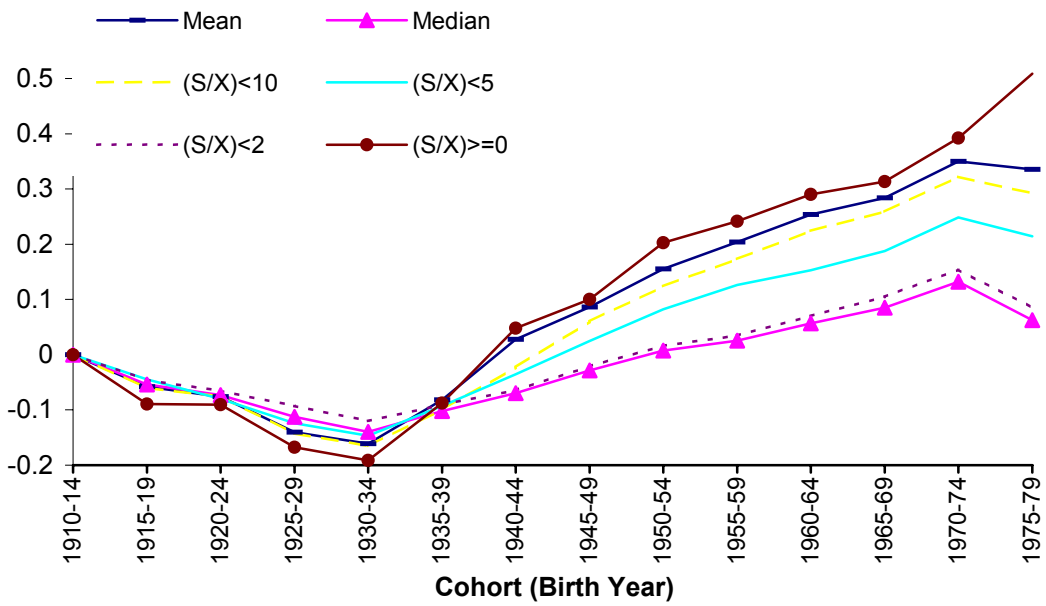

3.4 Alternative Definitions of Saving

As noted in the introduction, consumption expenditures were adjusted to remove those items reasonably viewed as "investment" so as to give a truer picture of the underlying saving rate. 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.

The cohort effects estimated under these different definitions of consumption and saving are plotted in Figure 6. 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 6 is consistent with our other sensitivity checks, including trimming the sample (Figure 5) and controlling for within-cell heterogeneity (Figure 2). 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. The remaining task is to explain this cohort pattern of saving rates.

4 Exploring the Cohort Patterns

Section 3 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 2. They display a distinct V-shaped pattern. In short, saving rates appeared to differ significantly for different cohorts.

Figure 6 – Cohort effects with different definitions of consumption and saving

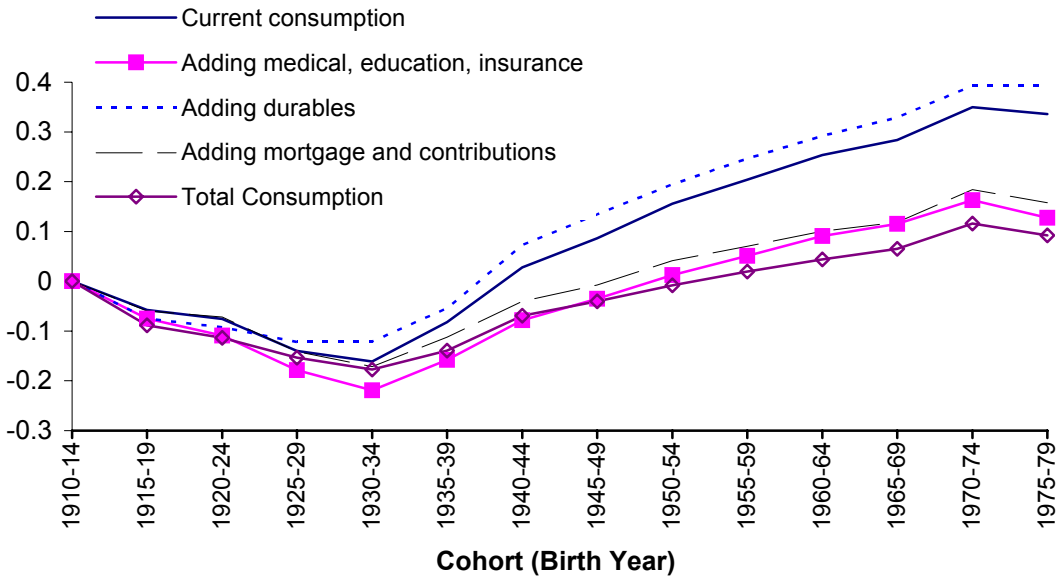

The cohort coefficients are estimated with reference to the 1910-14 birth cohort. Those coefficients for cohorts 3, 4 and 5 are all significantly lower than the reference cohort. In contrast, the coefficients for cohorts 9 through 14 are all significantly greater. In other words, those born from 1920 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 lead 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 had 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 (Claus and Scobie, 2002).

Attanasio (1998) finds a 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 for 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).

Attanasio (1998) 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 the 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 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 go[25]. 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 may also influence the saving rate.

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

In short there are many 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[26].

Thomson (1991) 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; and 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 to 100 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 birth[27], 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 saving[28]. 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

[26] We refer to these effects as cohort effects for convenience. In fact it possible that the effect is really a "time" effect, so that in the absence of certain changes both younger and older cohorts would have behaved similarily. We are grateful to John Creedy for pointing out this difference.

[27] 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.

[28] 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".

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 5 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 3 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.

Table 5 – 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	!920 -1924	Negative	1940-1984	1965-1979
4	1925-1929	Negative	1945-1989	1970-1984
5	1930-1934	Negative	1950-1994	1975-1990
9	!950-1954	Positive	1970-2014	1995-2010
10	1955-1959	Positive	1975-2019	2000-2014

Notes: Refers to lifetime saving rate relative to the 1910-14 reference group (by definition zero). See pattern of cohort dummies in the regressions presented in Section 3. Negative and positive refer to the cohorts that were typically significantly lower or higher in their lifetime saving rates.

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.

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 5).

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[29].

Prior to the major reforms of the late 1980's, the New Zealand labour markets were characterised by central wage fixing, limited flexibility and a high degree of unionisation (see Figure 7). Strong national unions were able to bargain particularly with state sector

[29] Unless otherwise noted, all the values of the economic and social policy variables are taken from various editions of the New Zealand Official Yearbook.

employers (public works, power generation and distribution, the Post Office, forestry and railways) and gain job security for their members.[30] It would seem plausible that the job security (at least perceived) which the unions were able to achieve, might have reduced the incentive for precautionary saving by their members.

Figure 7- Union density and membership in New Zealand: 1936-1999

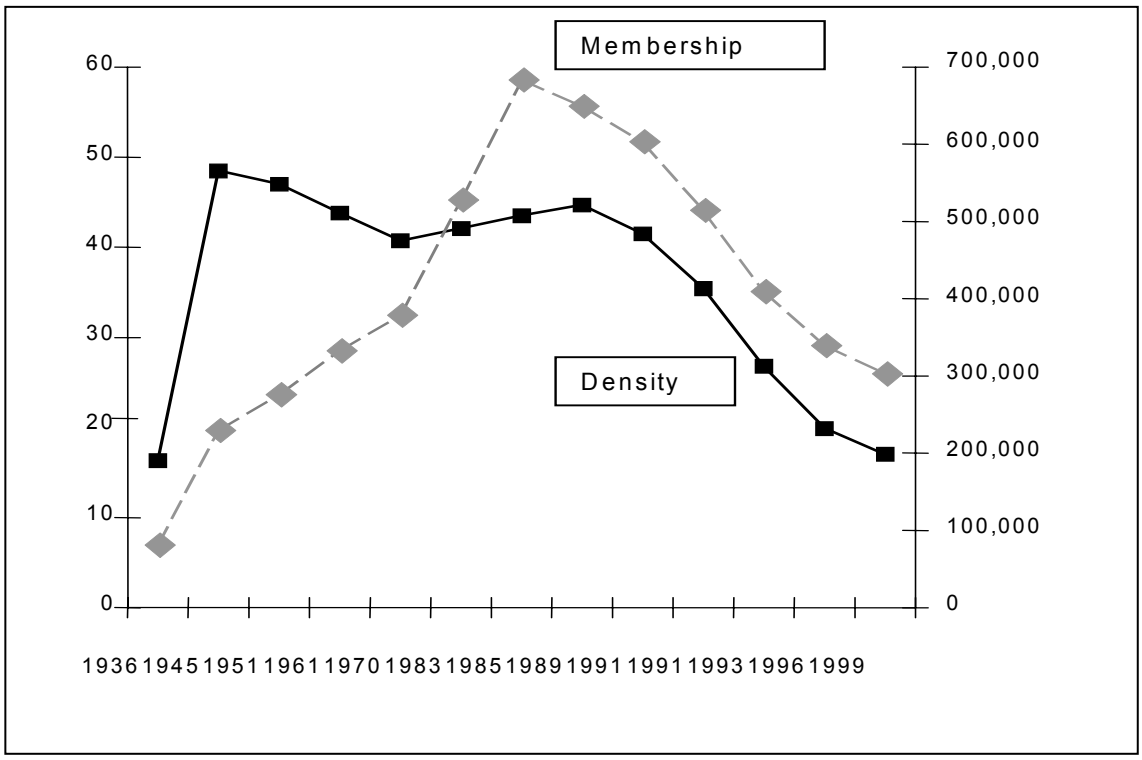

(Membership: RHS; Density (%) LHS)

A further important labour market indicator is the rate of participation.[31] The reference group had labour force participation rates of 94 and 35 percent for male and females 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.

[30] This resulted in over-manning and a major down-sizing after the subsequent liberalisation. We recognise that we cannot rule out the possibility that the direction of causality may run from cohorts to union membership; ie, the patterns were generated by cohorts with different propensities, values and attitudes moving through the time periods.

[31] See Rankin (1990).

Those facing expectations of higher future income growth rates might be expected to have lower rates of saving.[32] 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) cohort.

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 a universal system of retirement income support which is neither means nor asset tested.[33] 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 effect[34]. 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 long-term series for health and education benefits, making any inferences about changes in these policies too speculative at this stage. As an increasing share of the costs of tertiary education has been shifted from the state to individuals over the 1990s, it is to be expected that younger cohorts would have more incentive to make provision for their children's educational costs.[35] This tendency would be reinforced by the marked increase in tertiary participation rates.

Arguably one of the most significant elements of public policy that influence saving behaviour is that relating to the provision of public pensions. Between 1970 and 1979, the payment to a married couple under the New Zealand Superannuation Scheme rose by over 40 percent in real terms (see Figure 8). This rise corresponded to the peak saving years of the very cohorts that display lower lifetime savings. 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.

Another way to assess the real value of the universal superannuation is to compute the pension-to-wage ratio. This 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)[36]. In addition, between

---

[32] The effect of income growth on the saving of an individual is to be distinguished from the effect of economic growth on aggregate saving. Under the simple version of the life-cycle model in which workers save and the retired consume previously accumulated assets, an increase in the growth rate which increases the lifetime resources of the young relative to the elderly will unambiguously increase the aggregate saving rate. See Deaton and Paxson (2000).

[33] 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.

[34] 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).

[35] To the extent that education is seen as an investment, one might expect "saving for the children's college education", a concept well entrenched in the USA become more prevalent in New Zealand.

[36] See Preston (1999).

1985-86 and 1997-98 an income related surcharge on superannuation was in place, affecting up to one third of all recipients. Not only did this reduce the value of the pension to current retirees but it would have created uncertainty about the level of future payments and hence encouraged later cohorts to place less reliance on its eventual receipt.

National survey results from September 1999 indicate that 75 percent of those questioned agreed with the statement that they would have to rely entirely on their on efforts to ensure a comfortable retirement. Almost 90 percent felt that there would be significant changes to the present arrangements, and 88 percent felt that any future payments would be less than today's levels. In fact 52 percent felt there would be no superannuation within 20 to 30 years or that it would be only available to the indigent on a social safety net basis UMR, 1999).

New Zealand has created a public superannuation fund to meet part of the future government liabilities under the universal pension scheme. Taxes are now higher than are needed to meet current obligations, and will be commensurately lower in the future when drawings from the fund meet a part of the costs of the public pension; in fact the scheme is best viewed as a tax-smoothing system[37]. If the costs of higher taxes today are offset in present value by higher expected pensions in the future (or greater certainty that the present levels will be maintained), then lifetime wealth of an individual would be unaltered implying no change in their consumption today. As a consequence while national saving may be unchanged by the introduction of such a scheme, personal saving rates could well decline reflecting the effect of tax and transfer programmes rather than any underlying change in household consumption and saving behaviour. Gokhale, Kotlikoff and Sabelhaus (1996) conclude that the social security scheme which transferred resources from the current young and future generations to the current older ones is a principal factor in explaining the post-war decline in US saving rates.

Figure 8 – Real Payments under New Zealand Superannuation: 1970-2001

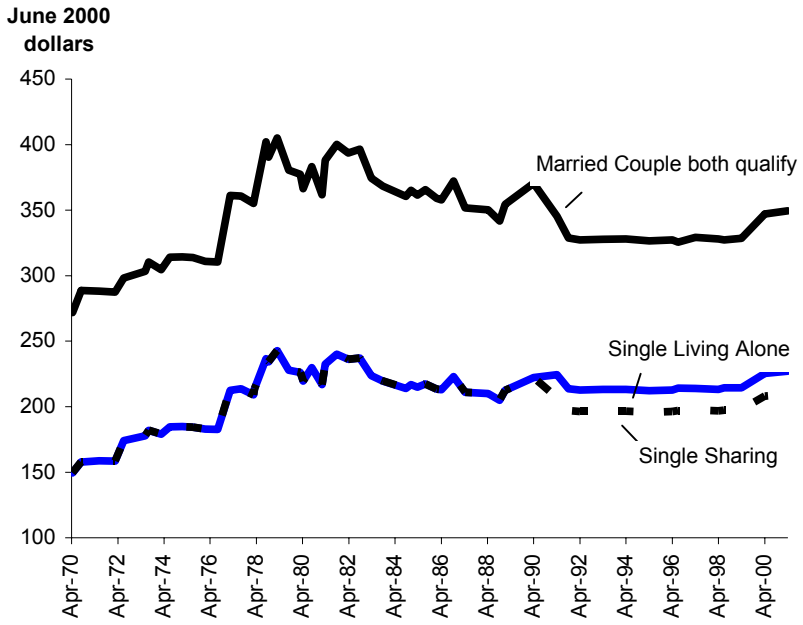

Source: NZ Treasury

[37] See http://www.nzsuperfund.co.nz/.

It has been argued that the universal New Zealand superannuation scheme has had, and continues to have, a significant influence on the saving behaviour of households. In the 1970's, when the scheme was perceived as both more generous and certain, saving rates of the cohorts directly affected were markedly lower. However it is recognised that there is still a debate about this response.

Arguably a pivotal study that examined the effect of public pensions on private saving was Feldstein (1974), and subsequently Feldstein (1996). Based on a life-cycle consumption function including a term for social security wealth, Feldstein argued that the presence of the US social security scheme significantly lowered the level of personal saving in the USA. However, Meguire (2001) challenges these findings and demonstrates that relatively modest changes to the sample period and the definition of the wealth variables can reverse the conclusion.

Accepting for the moment that the USA evidence does support the hypothesis that a public pension scheme reduces the level of personal saving, would we expect this to carry over to New Zealand? The US social security system is funded by a payroll tax and the benefits are earnings related (although not strictly proportional). It is arguable that while the system has in effect operated as a pay-as-you-go scheme (as distinct from a funded scheme), contributors may have perceived the scheme more as an individual retirement account. In that case, the public saving may have been seen as a closer substitute for personal saving than the defined benefit system funded from general revenue that operates in New Zealand.

In short, it appears the evidence is mixed. We would argue however that unit record data of the type used in this study are potentially more revealing of the underlying household responses than aggregate time series data of the type used by Feldstein. It is quite possible that households in the upper two or three deciles of lifetime income, show no response at all to changes in social security wealth. They discount the public pension as either a small share of their retirement income or subject to political uncertainty about whether they would eventually be beneficiaries in any event. When it is recalled that these households provide the majority of the total household saving (over 80 percent in the case of New Zealand) then it should not be surprising if aggregate data failed to reveal any effect of the superannuation scheme on personal saving. In contrast, for the 50 percent of the households in the lower income deciles, there may be some response toward reduced private saving[38], but the aggregate effect of this is likely to be quite small.

In addition to the direct effect of the public pension scheme on the savings of households, there is evidence that the US social security scheme has led to lower retirement ages and labour force participation among older workers (Venti and Wise (1996) and Lumsdaine and Wise (1990)). We would expect that the New Zealand superannuation system especially during the 1970s and 1980s when it became more generous, has played a similar role in contributing to some of the observed decline in the labour force participation rates of older male workers in New Zealand.

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 loan 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

[38] Scobie and Gibson (2003) find that New Zealand superannuation represents the majority of the retirement wealth of the four lowest wealth deciles and that if people wish to achieve consumption smoothing over their life time, the presence of NZS implies these people have no incentive for any other form of retirement saving.

which the number of state houses being built was much greater than in later years,[39] thus increasing the probability that a family would get access to a state house on concessional terms.

Added to this, the Family Benefits (Home Ownership) Act 1958 provided the option to capitalise a universal child benefit and apply that to equity in a first home. This would have reduced the incentive to save for an initial deposit (although lifetime income would not have altered). However the fact that a certain cash grant was replacing a future stream of benefits subject to political risk would suggest that the certainty equivalent of the capitalisation scheme would have reduced the incentive to save. Those born from 1925 to 1934 would have been in their household formation years when this policy was in place. They correspond to Cohorts 4 and 5 (Table 4) with saving rates significantly below those for the reference group. Finally, low savers who were paying off mortgages in the 1970s and early 1980s faced 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.

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 each 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.[40] 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.

5 Conclusions

In this paper we have focussed on the lifetime saving patterns of different cohorts. We have provided robust evidence that people born at different times, all other things being held constant, demonstrate different lifetime saving rates. 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. We examined the effect of different definitions and find that the essential cohort differences in saving rates remain. Furthermore, the cohort patterns remain when we included a wide range of conditioning variables to account for within-cell heterogeneity.

There are two reasons why the finding of cohort differences 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

[39] For example, in 1949-50, over 20 percent of all new dwellings completed were state houses.
[40] A complete accounting would require tracking of asset holdings and changes in asset values.

how to reconcile the divergent saving rates from the national accounts and the HES.[41] 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. It may hold out some glimmer that saving rates will not continue to decline as sharply in the future.

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 social and economic policies on household saving behaviour. This should contribute to better being able to predict the effect of policy changes in the future.

It is then logical to ask why birth cohort appears as such a significant determinant of saving rates. We have posited the hypothesis 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 birth-year alone is an important factor in explaining the distribution of the benefits of the welfare state over the past 50 years.

We argue that individuals would tend to have a lower lifetime saving rate if for example, their working life corresponded to a period of low unemployment, greater job security, real earnings growth, dual income households, generous welfare benefits and they held expectations of an assured state pension. By comparing the lifetime environments of different cohorts one could potentially test this hypothesis.

We have made some modest progress with this 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. There is the possibility that, in fact, the policies themselves are "endogenous"; they reflect the outcomes of the political process which in turn is function of other "truly" exogenous circumstances. The generation whose working lives commenced during the Great Depression and then faced WW II were already 25-40 years old by 1950. Typically they had delayed marriage and family formation and had little net wealth. Arguably they saw state transfers as one way that would both compensate them and substitute for the loss of "saving time" -and voted for such policies accordingly.

Because of the nature of the HES we have relied on synthetic cohorts. This method has its limitations. We do not observe the behaviour of say 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.

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.[42] Unfortunately, because of the divergent trends in the two sources 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

[41] For a comparison of the saving trends in the HES with the national accounts estimates, see Claus and Scobie (2002).
[42] See for example Bosworth, Burtless and Sabelhaus (1991).

aggregate trends. But until we have a reconciliation and can explain the divergent series, then this task remains in the category of unfinished business.

While the results of the cohort saving behaviour (Section 3) seem both significant and robust, our attempts to provide an explanation (Section 4) are partial and tentative. This is a complex area; the saving rates we observe are the resolution of a set of forces encompassing social and cultural norms shaped by the experience of earlier generations, economic conditions over the working life, expectations of future incomes, health status and life expectancy, and myriad state interventions. Arguably, the provision of higher state benefits, or more certainty would be expected to dampen the incentive for private saving.[43] Our preliminary examination of some snippets of evidence is at least consistent with that argument. Further testing of this relationship awaits the development of long-term data series for at least the last one hundred years, together with richer models about how cultural norms and values together with expectations, blend to shape consumption and saving decisions.

[43] Cross-sectionally countries with pay-as-you-go pension schemes funded from general or payroll taxes tend to have lower household saving rates. See Samwick (2000).

References

Attanasio, O., A Cohort Analysis of Saving Behaviour by U.S. Households. *Journal of Human Resources* 33(3): 575-609, 1998.

Attanasio, O. and J. Banks., Trends in Household Saving: A Tale of Two Countries. *Working Paper* No 98/15, The Institute of Fiscal Studies, London, 1998.

Bosworth, B., G. Burtless and J. Sabelhaus, The decline in saving: evidence from household surveys. *Brookings Papers on Economic Activity* 1: 183-256, 1991.

Caballero, R., Earnings Uncertainty and Aggregate Wealth Accumulation. *American Economic Review* 76(3): 676-691, 1991.

Claus, I. and G. Scobie, Saving in New Zealand: Measurement and Trends, *Working Paper 02/02*, New Zealand Treasury, 2002. http://www.treasury.govt.nz/workingpapers/2002/02-2.asp

Deaton, A., Panel Data from Time Series of Cross Sections. *Journal of Econometrics* 30(1): 109-126, 1985.

Deaton, A., *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*, Johns Hopkins University Press: Baltimore, 1997.

Deaton, A. and C. Paxson, Growth and Saving Among Individuals and Households. *Review of Economics and Statistics* 82(2): 212-225, 2000.

Feldstein, M., Social Security, Induced Retirement and Aggregate Capital Accumulation. *Journal of Political Economy* 82:905-926, 1974.

Feldstein, M., Social Security and Savings: New Time Series Evidence. *National Tax Journal* 49:151-163, 1996.

Gibson, J. and G. Scobie, Household Saving Behaviour in New Zealand: A Cohort Analysis. *Working Paper 01/18*, New Zealand Treasury, http://www.treasury.govt.nz/workingpapers/2001/01-18.asp 2001.

Gokhale, J., L. Kotlikoff, and J. Sabelhaus, Understanding the Post-war Decline in U.S. Saving: A Cohort Analysis. *Working Paper* 5571, National Bureau of Economic Research, 1996.

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, 1999.

Lumsdaine, R.L. and D.A.Wise, Aging and Labor Force Participation: A Review of Trends and Explanations. *Working Paper* No. 3420, National Bureau of Economic Research, 1990.

Meguire, P., Social Security and Personal Saving: 1971 and Beyond. Paper presented to the Annual Conference of the New Zealand Association of Economists, Christchurch, New Zealand, June 27-29, 2001.

Preston, D., *Retirement Income in New Zealand: the historical context*. Wellington: Office of the Retirement Commissioner, 1999.

Rankin, K., New Zealand's Labour Supply in a Long Term Perspective. Proceedings of the Sesquicentennial Conference of the New Zealand Association of Economists, Auckland August 20-22 : pp. 519-539, 1990.

Shorrocks, A., The Age-wealth Relationship: A Cross-section and Cohort Analysis. *Review of Economics and Statistics* 57(1): 155-163, 1975.

Thomson, D., *Selfish Generations? The Ageing of New Zealand's Welfare State* (Wellington: Bridget Williams Books), 1991.

UMR Insight Ltd, New Zealand Retirement Income- A Tracking Study. Prepared for The Super 2000Taskforce. Wellington, 1999.

Venti, S. and D. Wise, The Wealth of Cohorts: Retirement Saving and the Changing Assets of Older Americans. Working Paper 5609. National Bureau of Economic Research, 1996.

Verbecek, M. and T. Nijman, Can Cohort Data be Treated as Genuine Panel Data? *Empirical Economics* 17:9-23, 1992.